Intermediate Quantitative Economics with Python

Thomas J. Sargent & John Stachurski

Nov 08, 2026

CONTENTS

I Tools and Techniques **7**

1 Modeling COVID-19 **9**
- 1.1 Overview . . . 9
- 1.2 The SEIR model . . . 10
- 1.3 Implementation . . . 11
- 1.4 Experiments . . . 12
- 1.5 Ending lockdown . . . 17

2 Linear Algebra **21**
- 2.1 Overview . . . 21
- 2.2 Vectors . . . 22
- 2.3 Matrices . . . 29
- 2.4 Solving systems of equations . . . 31
- 2.5 Eigenvalues and eigenvectors . . . 36
- 2.6 Further topics . . . 39
- 2.7 Exercises . . . 41

3 QR Decomposition **45**
- 3.1 Overview . . . 45
- 3.2 Matrix factorization . . . 45
- 3.3 Gram-Schmidt process . . . 45
- 3.4 Some code . . . 47
- 3.5 Example . . . 48
- 3.6 Using QR decomposition to compute eigenvalues . . . 50
- 3.7 QR and PCA . . . 51

4 Circulant Matrices **55**
- 4.1 Overview . . . 55
- 4.2 Constructing a Circulant Matrix . . . 55
- 4.3 Connection to Permutation Matrix . . . 57
- 4.4 Examples with Python . . . 58
- 4.5 Associated Permutation Matrix . . . 62
- 4.6 Discrete Fourier Transform . . . 64

5 Singular Value Decomposition (SVD) **71**
- 5.1 Overview . . . 71
- 5.2 The setting . . . 71
- 5.3 Singular value decomposition . . . 72
- 5.4 Four fundamental subspaces . . . 73
- 5.5 Eckart-Young theorem . . . 76

5.6 Full and reduced SVD's . . . 77
5.7 Polar decomposition . . . 80
5.8 Application: principal components analysis (PCA) . . . 81
5.9 Relationship of PCA to SVD . . . 83
5.10 PCA with eigenvalues and eigenvectors . . . 83
5.11 Connections . . . 85
5.12 Exercises . . . 90

6 VARs and DMDs **97**
6.1 Overview . . . 97
6.2 First-order vector autoregressions . . . 98
6.3 A reduced-rank VAR . . . 100
6.4 The DMD algorithm . . . 102
6.5 Right eigenvectors: DMD modes . . . 103
6.6 Left eigenvectors and a modal representation . . . 104
6.7 Long-run responses . . . 106
6.8 A simulated example . . . 106
6.9 Source for some Python code . . . 111
6.10 Exercises . . . 111

7 Using Newton's Method to Solve Economic Models **115**
7.1 Overview . . . 115
7.2 Fixed point computation using Newton's method . . . 116
7.3 Root-Finding in one dimension . . . 123
7.4 Multivariate Newton's method . . . 124
7.5 Exercises . . . 133

II Elementary Statistics **139**

8 Elementary Probability with Matrices **141**
8.1 Sketch of basic concepts . . . 142
8.2 What does probability mean? . . . 142
8.3 Representing probability distributions . . . 144
8.4 Univariate probability distributions . . . 144
8.5 Bivariate probability distributions . . . 146
8.6 Marginal probability distributions . . . 146
8.7 Conditional probability distributions . . . 147
8.8 Transition probability matrix . . . 147
8.9 Application: forecasting a time series . . . 148
8.10 Statistical independence . . . 148
8.11 Means and variances . . . 149
8.12 Matrix representations of some bivariate distributions . . . 149
8.13 A continuous bivariate random vector . . . 156
8.14 Sum of two independently distributed random variables . . . 165
8.15 Coupling . . . 166
8.16 Copula functions . . . 167
8.17 Exercises . . . 174

9 Some Probability Distributions **181**
9.1 Some Discrete Probability Distributions . . . 181
9.2 Geometric distribution . . . 181
9.3 Pascal (negative binomial) distribution . . . 182
9.4 Newcomb–Benford distribution . . . 183
9.5 Univariate Gaussian distribution . . . 185

9.6 Uniform Distribution . . . 185
9.7 A Mixed Discrete-Continuous Distribution . . . 186
9.8 Drawing a Random Number from a Particular Distribution . . . 187

10 LLN and CLT **195**
10.1 Overview . . . 195
10.2 Relationships . . . 196
10.3 LLN . . . 196
10.4 CLT . . . 200
10.5 Exercises . . . 206

11 Two Meanings of Probability **213**
11.1 Overview . . . 213
11.2 Frequentist Interpretation . . . 214
11.3 Bayesian Interpretation . . . 220
11.4 Comparing the two interpretations . . . 228
11.5 Role of a Conjugate Prior . . . 229

12 Multivariate Hypergeometric Distribution **231**
12.1 Overview . . . 231
12.2 The administrator's problem . . . 231
12.3 Usage . . . 235

13 Multivariate Normal Distribution **241**
13.1 Overview . . . 241
13.2 The multivariate normal distribution . . . 242
13.3 Bivariate example . . . 245
13.4 Trivariate example . . . 249
13.5 One dimensional intelligence (IQ) . . . 250
13.6 Information as surprise . . . 254
13.7 Cholesky factor magic . . . 256
13.8 Math and verbal intelligence . . . 256
13.9 Univariate time series analysis . . . 259
13.10 Stochastic difference equation . . . 264
13.11 Application to stock price model . . . 265
13.12 Filtering foundations . . . 267
13.13 Classic factor analysis model . . . 271
13.14 PCA and factor analysis . . . 274
13.15 Exercises . . . 279

14 Fault Tree Uncertainties **289**
14.1 Overview . . . 289
14.2 The lognormal distribution . . . 290
14.3 The convolution theorem . . . 291
14.4 Approximating continuous distributions . . . 292
14.5 Discretizing the lognormal distribution . . . 296
14.6 Convolving probability mass functions . . . 298
14.7 Fault tree analysis . . . 301
14.8 Failure rates unknown . . . 302
14.9 Application: waste hoist failure rate . . . 304
14.10 Exercises . . . 309

15 Introduction to Artificial Neural Networks **317**
15.1 Overview . . . 318
15.2 A Deep (but not Wide) Artificial Neural Network . . . 318

15.3 Calibrating Parameters . . . 319
15.4 Back Propagation and the Chain Rule . . . 320
15.5 Training Set . . . 321
15.6 Example 1 . . . 324
15.7 How Deep? . . . 325
15.8 Example 2 . . . 326

16 Randomized Response Surveys **329**
16.1 Overview . . . 329
16.2 Warner's Strategy . . . 329
16.3 Comparing Two Survey Designs . . . 331
16.4 Concluding Remarks . . . 337

17 Expected Utilities of Random Responses **339**
17.1 Overview . . . 339
17.2 Privacy measures . . . 339
17.3 Zoo of concepts . . . 340
17.4 Respondent's expected utility . . . 341
17.5 Utilitarian view of survey design . . . 345
17.6 Criticisms of proposed privacy measures . . . 348
17.7 Concluding remarks . . . 352

III Bayes Law 355

18 Non-Conjugate Priors **357**
18.1 Overview . . . 357
18.2 The coin-flipping model . . . 358
18.3 MCMC reproduces the conjugate posterior . . . 360
18.4 Non-conjugate priors . . . 362
18.5 Variational inference . . . 367
18.6 Where to next . . . 371

19 Posterior Distributions for AR(1) Parameters **373**
19.1 Overview . . . 373
19.2 Estimation . . . 375
19.3 Conclusion . . . 385

20 Forecasting an AR(1) Process **387**
20.1 A Univariate First-Order Autoregressive Process . . . 388
20.2 Implementation . . . 389
20.3 Predictive Distributions of Path Properties . . . 394
20.4 A Wecker-Like Algorithm . . . 395
20.5 Using Simulations to Approximate a Posterior Distribution . . . 395
20.6 Calculating Sample Path Statistics . . . 396
20.7 Original Wecker Method . . . 398
20.8 Extended Wecker Method . . . 399
20.9 Comparison . . . 401
20.10 Conclusion . . . 403

IV Statistics and Information 405

21 Statistical Divergence Measures **407**
21.1 Overview . . . 407

21.2 Primer on entropy, cross-entropy, KL divergence . . . 408
21.3 Two Beta distributions: running example . . . 409
21.4 Kullback–Leibler divergence . . . 410
21.5 Jensen-Shannon divergence . . . 412
21.6 Chernoff entropy . . . 414
21.7 Comparing divergence measures . . . 415
21.8 KL divergence and maximum-likelihood estimation . . . 418
21.9 Related lectures . . . 419
21.10 Exercises . . . 420

22 Likelihood Ratio Processes **427**
22.1 Overview . . . 427
22.2 Likelihood Ratio Process . . . 428
22.3 Nature permanently draws from density g . . . 430
22.4 Peculiar property . . . 433
22.5 Nature permanently draws from density f . . . 433
22.6 Likelihood ratio test . . . 436
22.7 Hypothesis testing and classification . . . 445
22.8 Markov chains . . . 458
22.9 Related lectures . . . 461
22.10 Exercises . . . 461

23 Heterogeneous Beliefs and Financial Markets **469**
23.1 Overview . . . 469
23.2 Review: likelihood ratio processes . . . 470
23.3 Blume and Easley's setting . . . 471
23.4 Nature and agents' beliefs . . . 472
23.5 A socialist risk-sharing arrangement . . . 472
23.6 Social planner's allocation problem . . . 473
23.7 If you're so smart, ... 475
23.8 Competitive equilibrium prices . . . 476
23.9 Simulations . . . 480
23.10 Related lectures . . . 487
23.11 Exercises . . . 487

24 Likelihood Processes For VAR Models **515**
24.1 Overview . . . 515
24.2 VAR model setup . . . 516
24.3 Likelihood ratio process . . . 519
24.4 Example 1: two AR(1) processes . . . 520
24.5 Example 2: bivariate VAR models . . . 521
24.6 KL divergence rate . . . 526
24.7 Application: Samuelson multiplier-accelerator . . . 528
24.8 Exercises . . . 535

25 Mean of a Likelihood Ratio Process **543**
25.1 Overview . . . 543
25.2 Mathematical expectation of likelihood ratio . . . 544
25.3 Importance sampling . . . 547
25.4 Approximating a cumulative likelihood ratio . . . 548
25.5 Distribution of sample mean . . . 551
25.6 Choosing a sampling distribution . . . 554
25.7 Exercises . . . 558

26 A Problem that Stumped Milton Friedman **565**

26.1 Overview . . . 565
26.2 Source of the Problem . . . 566
26.3 Neyman-Pearson formulation . . . 567
26.4 Wald's sequential formulation . . . 569
26.5 Links between A, B and α, β . . . 571
26.6 Simulations . . . 573
26.7 Related lectures . . . 584
26.8 Exercises . . . 584

27 A Bayesian Formulation of Friedman and Wald's Problem **599**
27.1 Overview . . . 599
27.2 A Dynamic Programming Approach . . . 601
27.3 Implementation . . . 606
27.4 Analysis . . . 609
27.5 Exercises . . . 616

28 Exchangeability and Bayesian Updating **621**
28.1 Overview . . . 621
28.2 Independently and identically distributed . . . 622
28.3 A setting in which past observations are informative . . . 623
28.4 Relationship between IID and exchangeable . . . 624
28.5 Exchangeability . . . 625
28.6 Bayes' Law . . . 626
28.7 More details about Bayesian updating . . . 627
28.8 Appendix . . . 635
28.9 Sequels . . . 643
28.10 Exercises . . . 644

29 Likelihood Ratio Processes and Bayesian Learning **653**
29.1 Overview . . . 653
29.2 The setting . . . 654
29.3 Likelihood ratio processes and Bayes' law . . . 655
29.4 Another timing protocol . . . 661
29.5 Behavior of posterior probability $\{\pi_t\}$ under subjective probability distribution . . . 663
29.6 Initial prior is verified by paths drawn from subjective conditional densities . . . 671
29.7 Drilling down a little bit . . . 672
29.8 Related lectures . . . 674
29.9 Exercises . . . 674

30 Blackwell's Theorem on Comparing Experiments **681**
30.1 Overview . . . 681
30.2 Experiments and stochastic transformations . . . 683
30.3 Three equivalent criteria . . . 685
30.4 The main theorem . . . 687
30.5 Kihlstrom's Bayesian interpretation . . . 688
30.6 Simulating the Blackwell order with many states . . . 695
30.7 The DeGroot uncertainty function . . . 697
30.8 Application 1: demand for information about product quality . . . 700
30.9 Application 2: sequential experimental design . . . 702
30.10 The Data Processing Inequality and Coarse-Graining . . . 705
30.11 Relation to Bayesian likelihood-ratio learning . . . 708
30.12 Summary . . . 709
30.13 Exercises . . . 710

31 Information and Market Equilibrium **715**

31.1 Overview . . . 715
31.2 Setup . . . 717
31.3 Price revelation . . . 718
31.4 Invertibility and the elasticity of substitution . . . 720
31.5 Bayesian price expectations in a dynamic economy . . . 726
31.6 Simulating Bayesian learning from prices . . . 729
31.7 Exercises . . . 736

32 Incorrect Models **751**
32.1 Overview . . . 751
32.2 Sampling from compound lottery H . . . 754
32.3 Type 1 agent . . . 756
32.4 What a type 1 agent learns when mixture H generates data . . . 756
32.5 Kullback-Leibler divergence governs limit of π_t . . . 760
32.6 Type 2 agent . . . 765
32.7 Concluding remarks . . . 766
32.8 Exercises . . . 768

33 Bayesian versus Frequentist Decision Rules **779**
33.1 Overview . . . 779
33.2 Setup . . . 780
33.3 Frequentist Decision Rule . . . 783
33.4 Bayesian Decision Rule . . . 790
33.5 Was the Navy Captain's Hunch Correct? . . . 798
33.6 More Details . . . 800
33.7 Distribution of Bayesian Decision Rule's Time to Decide . . . 800
33.8 Probability of Making Correct Decision . . . 805
33.9 Distribution of Likelihood Ratios at Neyman-Pearson's t . . . 808
33.10 Exercises . . . 810

34 Merging of Opinions: The Blackwell–Dubins Theorem **817**
34.1 Overview . . . 817
34.2 Probability measures on sequence spaces . . . 818
34.3 The likelihood-ratio martingale . . . 820
34.4 The Blackwell–Dubins theorem . . . 822
34.5 The Beta–Bernoulli model . . . 823
34.6 Failure of merging: mutual singularity . . . 834
34.7 Kakutani's theorem: when does merging hold? . . . 836
34.8 Extension to continuous time . . . 840
34.9 Applications . . . 841
34.10 The rate of merging . . . 843
34.11 Summary and extensions . . . 845
34.12 Exercises . . . 846

35 Survival and Long-Run Dynamics under Recursive Preferences **853**
35.1 Overview . . . 853
35.2 Environment . . . 854
35.3 Planner's problem . . . 856
35.4 Survival conditions . . . 859
35.5 Wealth dynamics decomposition . . . 859
35.6 Survival regions . . . 867
35.7 Three survival channels . . . 871
35.8 Varying the IES . . . 873
35.9 Asymptotic results . . . 874
35.10 The separable case . . . 877

35.11 Asset pricing implications . . . 879
35.12 Optimistic and pessimistic distortions . . . 880
35.13 Long-run consumption distribution . . . 882
35.14 Summary . . . 884

36 The Design and Price of Information **885**
36.1 Overview . . . 885
36.2 The decision problem . . . 886
36.3 Experiments and their value . . . 887
36.4 Blackwell's order is only partial . . . 889
36.5 The seller's problem . . . 891
36.6 Two types . . . 891
36.7 A continuum of types . . . 894
36.8 Concluding remarks . . . 901
36.9 Exercises . . . 902

37 Market Diffusion with Two-Sided Learning **913**
37.1 Overview . . . 913
37.2 The market . . . 914
37.3 Two-sided learning . . . 915
37.4 Efficient experimentation . . . 917
37.5 Equilibrium . . . 919
37.6 Too much experimentation, then too little . . . 922
37.7 Diffusion over time . . . 923
37.8 Martingale properties . . . 926
37.9 Concluding remarks . . . 927
37.10 Exercises . . . 927

38 Which Moments to Match? **939**
38.1 Overview . . . 939
38.2 The estimation problem . . . 941
38.3 The auxiliary model and its score . . . 941
38.4 The estimator . . . 942
38.5 Why the score is the right thing to match . . . 943
38.6 Efficiency . . . 944
38.7 Comparison with matching the auxiliary parameters . . . 945
38.8 Choosing the score generator . . . 945
38.9 A first example: an AR(1) inside an AR(2) . . . 946
38.10 A second example: stochastic volatility . . . 950
38.11 Testing the structural model . . . 954
38.12 Applications . . . 955
38.13 Recovering the latent state . . . 956
38.14 Related lectures . . . 956
38.15 Exercises . . . 956

V Linear Programming **967**

39 Optimal Transport **969**
39.1 Overview . . . 969
39.2 The Optimal Transport Problem . . . 970
39.3 The Linear Programming Approach . . . 971
39.4 The Dual Problem . . . 978
39.5 The Python Optimal Transport Package . . . 980

40 Von Neumann Growth Model (and a Generalization) 985
40.1 Notation . . . 990
40.2 Model Ingredients and Assumptions . . . 991
40.3 Dynamic Interpretation . . . 993
40.4 Duality . . . 994
40.5 Interpretation as Two-player Zero-sum Game . . . 995

VI Introduction to Dynamics 1003

41 Finite Markov Chains 1005
41.1 Overview . . . 1005
41.2 Definitions . . . 1006
41.3 Simulation . . . 1008
41.4 Marginal Distributions . . . 1010
41.5 Irreducibility and Aperiodicity . . . 1012
41.6 Stationary Distributions . . . 1015
41.7 Ergodicity . . . 1019
41.8 Computing Expectations . . . 1020
41.9 Exercises . . . 1021

42 Inventory Dynamics 1029
42.1 Overview . . . 1029
42.2 Sample paths . . . 1030
42.3 Marginal distributions . . . 1033
42.4 Exercises . . . 1036

43 Linear State Space Models 1041
43.1 Overview . . . 1041
43.2 The Linear State Space Model . . . 1042
43.3 Distributions and Moments . . . 1048
43.4 Stationarity and Ergodicity . . . 1055
43.5 Noisy Observations . . . 1060
43.6 Prediction . . . 1061
43.7 Code . . . 1062
43.8 Exercises . . . 1062

44 Samuelson Multiplier-Accelerator 1065
44.1 Overview . . . 1065
44.2 Details . . . 1067
44.3 Implementation . . . 1070
44.4 Stochastic shocks . . . 1079
44.5 Government spending . . . 1081
44.6 Wrapping everything into a class . . . 1083
44.7 Using the LinearStateSpace class . . . 1088
44.8 Pure multiplier model . . . 1095
44.9 Summary . . . 1098

45 The Acceleration Principle and the Nature of Business Cycles 1099
45.1 Overview . . . 1099
45.2 Empirical foundation for the acceleration principle . . . 1101
45.3 Acceleration enables oscillations . . . 1102
45.4 A linear system with shocks . . . 1105
45.5 From autocovariances to spectra . . . 1106
45.6 Spectral peaks in the Hansen-Samuelson model . . . 1107

45.7 Real roots can produce peaks in general models . . . 1111
45.8 A calibrated model in the frequency domain . . . 1114
45.9 Summary . . . 1124
45.10 Exercises . . . 1125

46 Kesten Processes and Firm Dynamics 1133
46.1 Overview . . . 1133
46.2 Kesten processes . . . 1134
46.3 Heavy tails . . . 1137
46.4 Application: firm dynamics . . . 1139
46.5 Exercises . . . 1140

47 Wealth Distribution Dynamics 1147
47.1 Overview . . . 1147
47.2 Lorenz Curves and the Gini Coefficient . . . 1148
47.3 A Model of Wealth Dynamics . . . 1151
47.4 Implementation . . . 1152
47.5 Applications . . . 1155
47.6 Exercises . . . 1159

48 A First Look at the Kalman Filter 1163
48.1 Overview . . . 1163
48.2 The basic idea . . . 1164
48.3 Convergence . . . 1172
48.4 Implementation . . . 1172
48.5 Exercises . . . 1173

49 Another Look at the Kalman Filter 1181
49.1 A worker's output . . . 1182
49.2 A firm's wage-setting policy . . . 1183
49.3 A state-space representation . . . 1183
49.4 An innovations representation . . . 1184
49.5 Some computational experiments . . . 1187
49.6 Future extensions . . . 1197

50 The Kalman Filter and Vector Autoregressions 1199
50.1 Overview . . . 1199
50.2 The state space system . . . 1200
50.3 The Kalman filter . . . 1200
50.4 The Gram-Schmidt process . . . 1203
50.5 Hidden Markov model . . . 1203
50.6 Estimation . . . 1203
50.7 Vector autoregressions and the Kalman filter . . . 1204
50.8 Spectral factorization identity . . . 1205
50.9 Python implementation . . . 1207
50.10 Where this leads . . . 1212
50.11 Summary . . . 1213
50.12 Exercises . . . 1213

51 Vector Autoregressions for Subsets of Variables 1219
51.1 Overview . . . 1219
51.2 The large system . . . 1220
51.3 The small system . . . 1221
51.4 Four representations . . . 1221
51.5 Two cases where nothing is lost . . . 1224

51.6 Code . . . 1225
51.7 Example 1: a bivariate VAR(2) . . . 1228
51.8 Example 2: an omitted interest rate . . . 1235
51.9 Summary . . . 1237
51.10 Exercises . . . 1237

52 Organization Capital **1241**
52.1 Overview . . . 1241
52.2 The basic idea . . . 1242
52.3 Personnel information as organization capital . . . 1242
52.4 Industry equilibrium . . . 1249
52.5 Bayesian screening simulation . . . 1251
52.6 Team information . . . 1253
52.7 Firm-specific human capital . . . 1255
52.8 Costs of transferring organization capital . . . 1256
52.9 Summary and implications . . . 1258

53 Two Models of Measurements and the Investment Accelerator **1261**
53.1 Overview . . . 1261
53.2 The economic model . . . 1263
53.3 Measurement errors . . . 1267
53.4 A classical model of measurements initially collected by an agency . . . 1269
53.5 A model of optimal estimates reported by an agency . . . 1276
53.6 Simulation . . . 1281
53.7 Summary . . . 1287
53.8 Exercises . . . 1288

VII Search **1293**

54 Job Search I: The McCall Search Model **1295**
54.1 Overview . . . 1296
54.2 The McCall Model . . . 1296
54.3 Computing the Optimal Policy: Take 1 . . . 1299
54.4 Computing an Optimal Policy: Take 2 . . . 1306
54.5 Continuous Offer Distribution . . . 1308
54.6 Volatility . . . 1310
54.7 Exercises . . . 1315

55 Job Search II: Search and Separation **1319**
55.1 Overview . . . 1320
55.2 The model . . . 1320
55.3 Solving the model . . . 1321
55.4 Code . . . 1322
55.5 A simplifying transformation . . . 1326
55.6 Implementation . . . 1327
55.7 Impact of parameters . . . 1328
55.8 Exercises . . . 1331

56 Job Search III: Search with Separation and Markov Wages **1335**
56.1 Model setup . . . 1336
56.2 Code . . . 1337
56.3 Improving efficiency . . . 1341
56.4 Sensitivity analysis . . . 1343
56.5 Employment simulation . . . 1344

56.6 Ergodic property . . . 1347
56.7 Cross-sectional analysis . . . 1348
56.8 Lower unemployment compensation (c=0.5) . . . 1352
56.9 Exercises . . . 1352

57 Job Search IV: Fitted Value Function Iteration 1355
57.1 Overview . . . 1355
57.2 Model . . . 1356
57.3 Solution method . . . 1356
57.4 Implementation . . . 1359
57.5 Simulation . . . 1364
57.6 Exercises . . . 1371

58 Job Search V: Risk-Sensitive Preferences 1375
58.1 Overview . . . 1375
58.2 Outline . . . 1376
58.3 Introduction to risk-sensitivity . . . 1376
58.4 Back to job search . . . 1383

59 Job Search VI: Persistent and Transitory Wage Shocks 1391
59.1 Overview . . . 1391
59.2 The model . . . 1392
59.3 Implementation . . . 1393
59.4 Unemployment duration . . . 1397
59.5 Exercises . . . 1399

60 Job Search VII: Modeling Career Choice 1401
60.1 Overview . . . 1401
60.2 Model . . . 1402
60.3 Implementation . . . 1404
60.4 Exercises . . . 1409

61 Job Search VIII: On-the-Job Search 1415
61.1 Overview . . . 1415
61.2 Model . . . 1416
61.3 Implementation . . . 1418
61.4 Solving for Policies . . . 1420
61.5 Exercises . . . 1423

62 Job Search IX: Search with Learning 1429
62.1 Overview . . . 1429
62.2 Model . . . 1430
62.3 Take 1: Solution by VFI . . . 1434
62.4 Take 2: A More Efficient Method . . . 1439
62.5 Another Functional Equation . . . 1439
62.6 Solving the RWFE . . . 1439
62.7 Implementation . . . 1440
62.8 Exercises . . . 1441
62.9 Solutions . . . 1441
62.10 Appendix A . . . 1443
62.11 Appendix B . . . 1445
62.12 Examples . . . 1449

VIII Reinforcement Learning 1461

63 Inventory Management via Q-Learning 1463
63.1 Introduction . . . 1463
63.2 The Model . . . 1464
63.3 Solving via Value Function Iteration . . . 1465
63.4 Q-Learning . . . 1469

64 Risk-Sensitive Inventory Management via Q-Learning 1477
64.1 Introduction . . . 1477
64.2 The Model . . . 1478
64.3 Solving via Value Function Iteration . . . 1478
64.4 Interpreting the Outcomes . . . 1483
64.5 Q-Learning . . . 1484
64.6 Conclusion . . . 1491

65 Job Search X: Search with Q-Learning 1493
65.1 Overview . . . 1493
65.2 Review of McCall model . . . 1494
65.3 Implied quality function Q . . . 1497
65.4 From probabilities to samples . . . 1498
65.5 Q-learning . . . 1499
65.6 Employed worker can't quit . . . 1507
65.7 Possible extensions . . . 1508

IX Introduction to Optimal Savings 1509

66 Optimal Savings I: Cake Eating 1511
66.1 Overview . . . 1511
66.2 The model . . . 1512
66.3 The value function . . . 1513
66.4 The optimal policy . . . 1515
66.5 The Euler equation . . . 1516
66.6 Exercises . . . 1519

67 Optimal Savings II: Numerical Cake Eating 1523
67.1 Overview . . . 1523
67.2 Reviewing the Model . . . 1524
67.3 Value Function Iteration . . . 1524
67.4 Exercises . . . 1533

68 Optimal Savings III: Stochastic Returns 1539
68.1 Overview . . . 1539
68.2 The Model . . . 1540
68.3 Computation . . . 1544
68.4 Exercises . . . 1552

69 Optimal Savings IV: Time Iteration 1555
69.1 Overview . . . 1555
69.2 The Euler Equation . . . 1556
69.3 Implementation . . . 1558
69.4 Exercises . . . 1564

70 Optimal Savings V: The Endogenous Grid Method 1567

70.1 Overview . . . 1567
70.2 Key Idea . . . 1568
70.3 Implementation . . . 1569

71 Optimal Savings VI: EGM with JAX 1575
71.1 Overview . . . 1575
71.2 Implementation . . . 1576
71.3 Exercises . . . 1580

X Household Problems 1585

72 The Income Fluctuation Problem I: Discretization and VFI 1587
72.1 Overview . . . 1587
72.2 Set Up . . . 1588
72.3 Code . . . 1589
72.4 Exercises . . . 1593

73 The Income Fluctuation Problem II: Optimistic Policy Iteration 1597
73.1 Overview . . . 1597
73.2 Model and Primitives . . . 1598
73.3 Operators and Policies . . . 1598
73.4 Value Function Iteration . . . 1600
73.5 Optimistic Policy Iteration . . . 1601
73.6 Timing Comparison . . . 1602
73.7 Exercises . . . 1605

74 The Income Fluctuation Problem III: The Endogenous Grid Method 1609
74.1 Overview . . . 1609
74.2 The Household Problem . . . 1610
74.3 Computation . . . 1612
74.4 NumPy Implementation . . . 1613
74.5 JAX Implementation . . . 1616
74.6 Simulation . . . 1624

75 The Income Fluctuation Problem IV: Transient Income Shocks 1627
75.1 Overview . . . 1627
75.2 The Household Problem . . . 1628
75.3 NumPy Implementation . . . 1629
75.4 JAX Implementation . . . 1633
75.5 Simulation . . . 1640
75.6 Wealth Inequality . . . 1642
75.7 Exercises . . . 1645

76 The Income Fluctuation Problem V: Stochastic Returns on Assets 1649
76.1 Overview . . . 1650
76.2 The Model . . . 1650
76.3 Solution Algorithm . . . 1652
76.4 Implementation . . . 1653
76.5 Simulation . . . 1656
76.6 Wealth Inequality . . . 1659
76.7 Exercises . . . 1660

XI LQ Control 1665

77 LQ Control: Foundations 1667
77.1 Overview . . . 1667
77.2 Introduction . . . 1668
77.3 Optimality – finite horizon . . . 1671
77.4 Implementation . . . 1673
77.5 Extensions and comments . . . 1678
77.6 Further applications . . . 1680
77.7 Exercises . . . 1688

78 Lagrangian for LQ Control 1697
78.1 Overview . . . 1697
78.2 Undiscounted LQ DP Problem . . . 1698
78.3 Lagrangian . . . 1699
78.4 State-Costate Dynamics . . . 1700
78.5 Reciprocal Pairs Property . . . 1700
78.6 Schur decomposition . . . 1701
78.7 Application . . . 1702
78.8 Other Applications . . . 1707
78.9 Discounted Problems . . . 1707

79 Eliminating Cross Products 1711
79.1 Overview . . . 1711
79.2 Undiscounted Dynamic Programming Problem . . . 1711
79.3 Kalman Filter . . . 1712
79.4 Duality table . . . 1713

80 The Permanent Income Model 1715
80.1 Overview . . . 1715
80.2 The Savings Problem . . . 1716
80.3 Alternative Representations . . . 1723
80.4 Two Classic Examples . . . 1726
80.5 Further Reading . . . 1729
80.6 Appendix: The Euler Equation . . . 1730

81 Permanent Income II: LQ Techniques 1731
81.1 Overview . . . 1731
81.2 Setup . . . 1732
81.3 The LQ Approach . . . 1734
81.4 Implementation . . . 1735
81.5 Two Example Economies . . . 1738

82 Robust Permanent Income and Pricing 1749
82.1 Overview . . . 1749
82.2 Risk-sensitive recursive preferences . . . 1750
82.3 A preference for robustness . . . 1753
82.4 The permanent income economy . . . 1753
82.5 Observational equivalence . . . 1754
82.6 Estimation . . . 1757
82.7 Asset pricing and the market price of risk . . . 1760
82.8 A risk-sensitive regulator . . . 1762
82.9 Exercises . . . 1765
82.10 Related lectures . . . 1768

83 Certainty Equivalence **1769**
83.1 Overview . . . 1769
83.2 A central problem of empirical economics . . . 1770
83.3 A formal setup . . . 1770
83.4 Estimated rules are not enough . . . 1771
83.5 An optimization problem . . . 1771
83.6 A linear-quadratic DP problem and certainty equivalence . . . 1772
83.7 A trouble with ad hoc expectations . . . 1775
83.8 Exercises . . . 1775
83.9 Concluding remarks . . . 1776

84 Certainty Equivalence and Model Uncertainty **1777**
84.1 Overview . . . 1777
84.2 Recap: ordinary certainty equivalence . . . 1778
84.3 Model uncertainty and robustness . . . 1780
84.4 Value function under robustness . . . 1782
84.5 Risk-sensitive preferences . . . 1783
84.6 Application: permanent income model . . . 1783
84.7 Summary . . . 1788
84.8 Exercises . . . 1788

85 The LQ Permanent Income Model **1793**
85.1 Overview . . . 1793
85.2 The standard LQ permanent income model . . . 1794
85.3 Exercises . . . 1800

86 Consumption Smoothing with Incomplete and Complete Markets **1807**
86.1 Overview . . . 1807
86.2 A brief review . . . 1808
86.3 Spreading consumption cross sections . . . 1809
86.4 A borrowers and lenders economy . . . 1811
86.5 Consumption smoothing with complete markets . . . 1813
86.6 Exercises . . . 1815

87 Robust Consumption Smoothing and Precautionary Savings **1821**
87.1 Overview . . . 1821
87.2 A brief review . . . 1822
87.3 A robust permanent income model . . . 1824
87.4 Observational equivalence . . . 1827
87.5 The scalar worst-case model . . . 1832
87.6 Frequency domain interpretation . . . 1836
87.7 Detection error probabilities . . . 1838
87.8 Robustness of decision rules . . . 1841
87.9 Concluding remarks . . . 1841
87.10 Exercises . . . 1842
87.11 Related lectures . . . 1845

88 A Robust LQ Bewley Model **1847**
88.1 Overview . . . 1847
88.2 Mapping the Bewley economy into the HST framework . . . 1848
88.3 The robustness scalar for this economy . . . 1849
88.4 The Bewley observational-equivalence locus . . . 1849
88.5 Equilibrium with heterogeneous types . . . 1850
88.6 Where the agents genuinely differ . . . 1851
88.7 Computation . . . 1851

88.8 Concluding remarks . . . 1857
88.9 Exercises . . . 1857
88.10 Related lectures . . . 1861

89 Production Smoothing via Inventories 1863
89.1 Overview . . . 1863
89.2 Example 1 . . . 1868
89.3 Inventories Not Useful . . . 1869
89.4 Inventories Useful but are Hardwired to be Zero Always . . . 1869
89.5 Example 2 . . . 1871
89.6 Example 3 . . . 1872
89.7 Example 4 . . . 1873
89.8 Example 5 . . . 1875
89.9 Example 6 . . . 1876
89.10 Exercises . . . 1879

XII Bounded Rationality in Macroeconomics 1885

90 A Peculiar Definition of Bounded Rationality 1887
90.1 Overview . . . 1887
90.2 Rational expectations as a fixed point . . . 1892
90.3 One recursion, many learners . . . 1897
90.4 Money and prices . . . 1901
90.5 Two currencies . . . 1904
90.6 Where this leaves us . . . 1905
90.7 Exercises . . . 1906

91 Adaptive Agents in an Overlapping Generations Monetary Economy 1913
91.1 Overview . . . 1913
91.2 The environment . . . 1914
91.3 Part 1: a stochastic deficit . . . 1915
91.4 Part 2: a constant deficit and two steady states . . . 1921
91.5 A government learning the Phillips curve . . . 1932
91.6 Concluding remarks . . . 1936
91.7 Exercises . . . 1936

92 Learning, Approximation, and Equilibrium Computation 1951
92.1 Overview . . . 1951
92.2 A continuous-state functional equation . . . 1952
92.3 A benchmark solution . . . 1955
92.4 Parameterized expectations . . . 1957
92.5 Learning is equilibrium computation . . . 1960
92.6 Learning without a functional form . . . 1964
92.7 Concluding remarks . . . 1967
92.8 Exercises . . . 1968

93 Exchange Rate Indeterminacy, Learning, and Experiments 1977
93.1 Overview . . . 1977
93.2 The two-currency economy . . . 1978
93.3 Newton–Raphson learning . . . 1980
93.4 Evidence from the laboratory . . . 1987
93.5 A genetic algorithm economy . . . 1988
93.6 Concluding remarks . . . 1994
93.7 Exercises . . . 1995

94 Networks, Genetic Algorithms, and Classifier Systems **2005**
94.1 Overview . . . 2005
94.2 The perceptron: an agent as a discriminant function . . . 2006
94.3 Associative memory: the Hopfield network . . . 2009
94.4 The genetic algorithm . . . 2012
94.5 Classifier systems . . . 2019
94.6 Evolutionary programming . . . 2024
94.7 Concluding remarks . . . 2025
94.8 Exercises . . . 2026

95 Money as a Medium of Exchange among Artificially Intelligent Agents **2037**
95.1 Overview . . . 2037
95.2 The Kiyotaki-Wright environment . . . 2039
95.3 Classifier systems for the Kiyotaki-Wright environment . . . 2042
95.4 Implementation . . . 2046
95.5 Economy A1.1: does a medium of exchange emerge? . . . 2053
95.6 Economy A2.1: when theory predicts speculation . . . 2056
95.7 The bucket brigade's discount factor . . . 2060
95.8 Economy B.1: selecting between two equilibria . . . 2063
95.9 The genetic algorithm . . . 2065
95.10 Economy A1.2: learning from random rules . . . 2067
95.11 Economy C: fiat money . . . 2069
95.12 Economy D: five goods, five types . . . 2072
95.13 Concluding remarks . . . 2075
95.14 Exercises . . . 2076

96 1993 Prospects for Bounded Rationality in Macroeconomics **2083**
96.1 Overview . . . 2083
96.2 The quest, restated . . . 2084
96.3 The debit side: how much we must hard-wire . . . 2084
96.4 Why the econometricians have not returned the compliment . . . 2086
96.5 The credit side: selection, computation, and a returned gift . . . 2092
96.6 Reading the ledger against the quest . . . 2094
96.7 Postscript: what became of the program . . . 2095
96.8 Exercises . . . 2100

XIII Phillips Curve Tradeoffs 2109

97 The Rise and Fall of U.S. Inflation **2111**
97.1 Overview . . . 2111
97.2 Facts . . . 2112
97.3 The Phillips curve in the data . . . 2114
97.4 Two interpretations . . . 2118
97.5 Ignoring the Lucas Critique . . . 2119
97.6 A premature summary: triumph or vindication? . . . 2121
97.7 Data patterns after 1999 . . . 2124
97.8 A note on notation . . . 2130
97.9 Exercises . . . 2133

98 The Credibility Problem **2141**
98.1 Overview . . . 2141
98.2 A one-period economy . . . 2142
98.3 Nash and Ramsey outcomes . . . 2143
98.4 Best response dynamics . . . 2146

98.5 Least squares learning converges to Nash . . . 2148
98.6 More foresight . . . 2151
98.7 Appendix: stochastic approximation . . . 2151
98.8 Exercises . . . 2152

99 Credible Government Policies **2161**
99.1 Overview . . . 2161
99.2 The repeated economy . . . 2163
99.3 Recursive strategies and promised values . . . 2166
99.4 The Abreu–Pearce–Stacchetti method . . . 2167
99.5 The best and the worst equilibrium . . . 2169
99.6 Examples of recursive equilibria . . . 2173
99.7 Multiplicity . . . 2176
99.8 Numerical example . . . 2178
99.9 Interpretations . . . 2179
99.10 Exercises . . . 2181

100 Adaptive Expectations and the Phelps Problem **2191**
100.1 Overview . . . 2191
100.2 Adaptive expectations . . . 2192
100.3 The Phelps problem . . . 2194
100.4 Disinflation paths . . . 2196
100.5 The general Phelps problem . . . 2198
100.6 Testing the natural-rate hypothesis . . . 2200
100.7 Sacrifice ratios and the subversion of the Phelps model . . . 2201
100.8 Exercises . . . 2203

101 Optimal Misspecified Beliefs **2207**
101.1 Overview . . . 2207
101.2 An experiment in Bray's lab . . . 2208
101.3 Optimal misspecification . . . 2210
101.4 Comparing the true and forecasting models . . . 2214
101.5 How the equilibrium depends on ρ . . . 2216
101.6 Lessons . . . 2217
101.7 Exercises . . . 2218

102 Self-Confirming Equilibria **2223**
102.1 Overview . . . 2223
102.2 Objects in the Phelps problem . . . 2224
102.3 The actual Phillips curve . . . 2225
102.4 Self-confirming equilibria . . . 2225
102.5 The special case solved by hand . . . 2226
102.6 Uniqueness and E-stability . . . 2229
102.7 Equilibrium with misspecified beliefs . . . 2232
102.8 Exercises . . . 2240

103 Adaptive Learning and Escape Dynamics **2247**
103.1 Overview . . . 2247
103.2 A primer on recursive algorithms . . . 2249
103.3 The adaptive model . . . 2255
103.4 Mean dynamics near the self-confirming equilibrium . . . 2258
103.5 Least squares learning: a decreasing gain . . . 2260
103.6 Constant gain and escape dynamics . . . 2264
103.7 The escape route and the induction hypothesis . . . 2266
103.8 Relation to equilibria under forecast misspecification . . . 2269

103.9 Role of the discount factor . . . 2269
103.10 Anticipated utility . . . 2271
103.11 Conclusions . . . 2271
103.12 Exercises . . . 2272

104 Escaping Nash Inflation **2279**
104.1 Overview . . . 2279
104.2 The model . . . 2280
104.3 The self-confirming equilibrium . . . 2282
104.4 Adaptation and the mean dynamics . . . 2284
104.5 Escape dynamics as a control problem . . . 2285
104.6 The changed-measure escape path . . . 2287
104.7 The race of four ODEs . . . 2289
104.8 Mean dynamics reinforce the escape . . . 2291
104.9 The experimentation trap . . . 2292
104.10 Simulating escapes . . . 2293
104.11 Escape frequency and model richness . . . 2296
104.12 Escaping volatile inflation . . . 2296
104.13 Exercises . . . 2298

105 Priors, Escapes, and Learning Cycles **2309**
105.1 Overview . . . 2309
105.2 The static model . . . 2310
105.3 The self-confirming equilibrium . . . 2312
105.4 Drifting beliefs and the Kalman filter . . . 2312
105.5 Mean dynamics and E-stability . . . 2314
105.6 Learning cycles . . . 2316
105.7 Escape dynamics and the direction of escape . . . 2320
105.8 Sims's nonconvergence . . . 2320
105.9 Conclusion . . . 2327
105.10 Exercises . . . 2328

106 Econometric Policy Evaluation **2341**
106.1 Overview . . . 2341
106.2 The model . . . 2342
106.3 The likelihood function . . . 2343
106.4 Data . . . 2344
106.5 Beliefs and recommendations . . . 2345
106.6 Estimates . . . 2347
106.7 Recommendations and misses . . . 2349
106.8 Misses favor vindication . . . 2354
106.9 What came later . . . 2358
106.10 Exercises . . . 2359

107 Drifts and Volatilities **2365**
107.1 Overview . . . 2365
107.2 Bad policy or bad luck? . . . 2367
107.3 A VAR with drifting coefficients and stochastic volatility . . . 2368
107.4 The data . . . 2370
107.5 Priors . . . 2372
107.6 Simulating the posterior . . . 2373
107.7 What the data say . . . 2375
107.8 Why classical tests missed the drift . . . 2394
107.9 Bad policy or bad luck? A verdict . . . 2395
107.10 Appendix: the sampler in detail . . . 2396

107.11 Exercises . . . 2400

108 Drifts and Volatilities after 2000 **2409**
108.1 Overview . . . 2409
108.2 The model and the sampler . . . 2410
108.3 The extended sample . . . 2410
108.4 Fitting the extended sample . . . 2413
108.5 Did coefficient drift continue? . . . 2414
108.6 Volatility after the Great Moderation . . . 2415
108.7 Core inflation, the natural rate, and persistence . . . 2419
108.8 What the additional observations change . . . 2423
108.9 Exercises . . . 2423

109 The Lost Conquest: Fed Policy in the 2020s **2429**
109.1 Overview . . . 2429
109.2 The three elements . . . 2431
109.3 The Fed's drifting-coefficients beliefs . . . 2432
109.4 The Fed's Phelps problem . . . 2435
109.5 Why the Fed was slow: a counterfactual . . . 2441
109.6 Real-time data and the delay . . . 2442
109.7 How beliefs shape the Phelps rule . . . 2446
109.8 From policy to measured beliefs . . . 2448
109.9 Was the conquest lost? . . . 2454
109.10 Exercises . . . 2455

XIV Optimal Growth 2465

110 Cass-Koopmans Model **2467**
110.1 Overview . . . 2467
110.2 The Model . . . 2468
110.3 Planning Problem . . . 2469
110.4 Shooting Algorithm . . . 2473
110.5 Setting Initial Capital to Steady State Capital . . . 2477
110.6 A Turnpike Property . . . 2479
110.7 A Limiting Infinite Horizon Economy . . . 2482
110.8 Stable Manifold and Phase Diagram . . . 2483
110.9 Concluding Remarks . . . 2485

111 Cass-Koopmans Competitive Equilibrium **2489**
111.1 Overview . . . 2489
111.2 Review of Cass-Koopmans Model . . . 2490
111.3 Competitive Equilibrium . . . 2491
111.4 Market Structure . . . 2492
111.5 Firm Problem . . . 2492
111.6 Household Problem . . . 2493
111.7 Computing a Competitive Equilibrium . . . 2495
111.8 Yield Curves and Hicks-Arrow Prices . . . 2503

112 Cass-Koopmans Model with Distorting Taxes **2505**
112.1 Overview . . . 2505
112.2 The Economy . . . 2505
112.3 Equilibrium . . . 2507
112.4 No-arbitrage Condition . . . 2507
112.5 Household's First Order Condition . . . 2508

112.6 Computing Equilibria . . . 2509
112.7 Python Code . . . 2509
112.8 Some functional forms . . . 2513
112.9 Computation . . . 2514
112.10 Exogenous growth . . . 2537

113 Two-Country Model with Distorting Taxes **2549**
113.1 Overview . . . 2549
113.2 A Two-Country Cass-Koopmans Model . . . 2549

XV Multiple Agent Models **2563**

114 A Lake Model of Employment and Unemployment **2565**
114.1 Overview . . . 2565
114.2 The model . . . 2566
114.3 Implementation . . . 2568
114.4 Dynamics of an individual worker . . . 2575
114.5 Exercises . . . 2578

115 Lake Model with an Endogenous Job Finding Rate **2585**
115.1 Overview . . . 2585
115.2 Set Up . . . 2586
115.3 Fiscal policy . . . 2591
115.4 Exercises . . . 2595

116 Rational Expectations Equilibrium **2599**
116.1 Overview . . . 2599
116.2 Rational expectations equilibrium . . . 2602
116.3 Computing an equilibrium . . . 2605
116.4 Concluding remarks . . . 2607
116.5 Exercises . . . 2608

117 Investment Under Uncertainty **2615**
117.1 Overview . . . 2615
117.2 The industry . . . 2616
117.3 The firm and the price of installed capital . . . 2617
117.4 Rational expectations equilibrium . . . 2618
117.5 Equilibrium as a planning problem . . . 2619
117.6 Recursive competitive equilibrium . . . 2620
117.7 A computable version . . . 2621
117.8 Long run behavior with serially independent demand . . . 2624
117.9 Long run behavior with serially correlated demand . . . 2626
117.10 Relation to the rational expectations lecture . . . 2631
117.11 Exercises . . . 2632

118 Equilibrium Search and Unemployment **2637**
118.1 Overview . . . 2637
118.2 The economy . . . 2638
118.3 A worker's problem . . . 2639
118.4 Employment, wages, and two thresholds . . . 2640
118.5 A computable version . . . 2642
118.6 Economy-wide equilibrium . . . 2646
118.7 Relation to other lectures . . . 2651
118.8 Exercises . . . 2654

119 Optimal Growth Under Uncertainty and Tobin's q **2667**
119.1 Overview . . . 2667
119.2 The Brock-Mirman planning problem . . . 2668
119.3 Long-run behavior . . . 2669
119.4 Irreversible investment and Tobin's q . . . 2670
119.5 Differentiability of the value function . . . 2671
119.6 Computing the model . . . 2674
119.7 Exercises . . . 2679

120 Stability in Linear Rational Expectations Models **2683**
120.1 Overview . . . 2683
120.2 Linear Difference Equations . . . 2684
120.3 Illustration: Cagan's Model . . . 2686
120.4 Some Python Code . . . 2687
120.5 Alternative Code . . . 2689
120.6 Another Perspective . . . 2691
120.7 Log money Supply Feeds Back on Log Price Level . . . 2694
120.8 Big P, Little p Interpretation . . . 2697
120.9 Fun with SymPy . . . 2699

121 Markov Perfect Equilibrium **2703**
121.1 Overview . . . 2703
121.2 Background . . . 2704
121.3 Linear Markov Perfect Equilibria . . . 2705
121.4 Application . . . 2708
121.5 Exercises . . . 2711

122 Uncertainty Traps **2721**
122.1 Overview . . . 2721
122.2 The Model . . . 2722
122.3 Implementation . . . 2725
122.4 Results . . . 2726
122.5 Exercises . . . 2727

123 The Aiyagari Model **2735**
123.1 Overview . . . 2735
123.2 The Economy . . . 2736
123.3 Implementation . . . 2738
123.4 Exercises . . . 2748

124 The Aiyagari Model with Endogenous Grid Method **2755**
124.1 Overview . . . 2755
124.2 The Economy . . . 2756
124.3 Implementation with EGM . . . 2758
124.4 Computing Aggregate Capital by Simulation . . . 2762
124.5 Computing Equilibrium . . . 2763
124.6 Wealth Distribution . . . 2765
124.7 Summary and Comparison . . . 2767
124.8 Exercises . . . 2768

125 Transitions in an Overlapping Generations Model **2769**
125.1 Overview . . . 2769
125.2 Setting . . . 2770
125.3 Production . . . 2771
125.4 Government . . . 2771

125.5 Activities in factor markets . . . 2771
125.6 Representative firm's problem . . . 2772
125.7 Individuals' problems . . . 2773
125.8 Equilibrium . . . 2774
125.9 Next steps . . . 2774
125.10 Closed form solution . . . 2774
125.11 A computational strategy . . . 2786
125.12 Exercises . . . 2795

126 A Long-Lived, Heterogeneous Agents, Overlapping Generations Model 2799
126.1 Overview . . . 2799
126.2 Environment . . . 2800
126.3 Production . . . 2801
126.4 Government . . . 2801
126.5 Activities in factor markets . . . 2802
126.6 Representative firm's problem . . . 2803
126.7 Households' problems . . . 2803
126.8 Population dynamics . . . 2803
126.9 Equilibrium . . . 2804
126.10 Implementation . . . 2804
126.11 Computing a steady state . . . 2808
126.12 Transition dynamics . . . 2814
126.13 Experiment 1: an immediate tax cut . . . 2819
126.14 Experiment 2: a preannounced tax cut . . . 2825
126.15 Exercises . . . 2830

127 Two Computations to Fund Social Security 2835
127.1 Overview . . . 2836
127.2 Model . . . 2837
127.3 Two experiments . . . 2842
127.4 Computation strategy . . . 2843
127.5 Calibration . . . 2846
127.6 Individual optimality . . . 2850
127.7 Transition path computation . . . 2860
127.8 Experiment 1: compensation through debt . . . 2872
127.9 Experiment 2: government capital accumulation . . . 2886
127.10 Distribution surfaces . . . 2893
127.11 Concluding remarks . . . 2899
127.12 Exercises . . . 2900

128 Social Security Trust Fund Depletion 2905
128.1 Overview . . . 2905
128.2 The institution . . . 2907
128.3 A model we can solve quickly . . . 2908
128.4 Aggregation and the ledger . . . 2918
128.5 The economy before the news . . . 2918
128.6 Solving for the equilibrium payable path . . . 2920
128.7 Claiming and retirement along the transition . . . 2922
128.8 How this compares with the full model . . . 2927
128.9 Exercises . . . 2929

129 Linear Quadratic Mean Field Games 2943
129.1 Overview . . . 2943
129.2 The environment . . . 2945
129.3 Solving the individual problem . . . 2946

129.4 A state-costate system . . . 2947
129.5 The equilibrium Riccati equation . . . 2948
129.6 Computing equilibria . . . 2950
129.7 The scalar case . . . 2952
129.8 An industry equilibrium with capital accumulation . . . 2955
129.9 Aggregate shocks and identification . . . 2963
129.10 Persistence . . . 2966
129.11 Multiproduct price setting with Kimball demand . . . 2968
129.12 The planner . . . 2973
129.13 Summary . . . 2978
129.14 Exercises . . . 2978
129.15 Further reading . . . 2990

130 Linear Quadratic Mean Field Games: Extensions 2991
130.1 Overview . . . 2991
130.2 What we need from the previous lecture . . . 2992
130.3 State-action cross terms . . . 2993
130.4 Forcing processes and pass-through . . . 2998
130.5 Coupling through higher moments . . . 3000
130.6 Exercises . . . 3004
130.7 Further reading . . . 3009

XVI Asset Pricing and Finance 3011

131 Asset Pricing: Finite State Models 3013
131.1 Overview . . . 3013
131.2 Pricing Models . . . 3014
131.3 Prices in the Risk-Neutral Case . . . 3015
131.4 Risk Aversion and Asset Prices . . . 3020
131.5 Exercises . . . 3029

132 Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns 3035
132.1 Overview . . . 3036
132.2 Euler equation . . . 3036
132.3 The Euler equation under lognormality . . . 3037
132.4 The restricted system and its likelihood . . . 3038
132.5 Likelihood implementation . . . 3040
132.6 Preference parameters and likelihood ratio tests . . . 3051
132.7 Predictability and the R-squared restriction . . . 3054
132.8 Return-difference tests . . . 3055
132.9 Empirical MLE estimation . . . 3057
132.10 Residual diagnostics . . . 3061
132.11 Connection to the equity premium puzzle . . . 3063
132.12 Another approach . . . 3064

133 Estimating Euler Equations by Generalized Method of Moments 3065
133.1 Overview . . . 3065
133.2 The economic model . . . 3066
133.3 From conditional to unconditional moments . . . 3067
133.4 GMM criterion and asymptotic theory . . . 3073
133.5 Covariance estimation and the choice of instruments . . . 3074
133.6 Empirical GMM estimation . . . 3081
133.7 Summary . . . 3083

134 Competitive Equilibria with Arrow Securities **3085**
134.1 Overview . . . 3085
134.2 The setting . . . 3087
134.3 Markov asset prices . . . 3088
134.4 Recursive formulation . . . 3090
134.5 A rational expectations equilibrium . . . 3091
134.6 State variable degeneracy . . . 3093
134.7 Computing a competitive equilibrium . . . 3094
134.8 Finite horizon . . . 3097
134.9 Python code . . . 3098
134.10 Examples . . . 3100
134.11 Concluding remarks . . . 3112
134.12 Related lectures . . . 3113
134.13 Exercises . . . 3114

135 Heterogeneous Beliefs and Bubbles **3127**
135.1 Overview . . . 3127
135.2 Structure of the model . . . 3128
135.3 Solving the model . . . 3130
135.4 Exercises . . . 3135

136 Speculative Behavior with Bayesian Learning **3139**
136.1 Overview . . . 3139
136.2 Structure of the model . . . 3140
136.3 Information and beliefs . . . 3141
136.4 Source of heterogeneous priors . . . 3141
136.5 Beta priors . . . 3142
136.6 Market prices with learning . . . 3143
136.7 Two Traders . . . 3144
136.8 Concluding remarks . . . 3151
136.9 Exercise . . . 3151

137 Rational Learning and Rational Expectations **3155**
137.1 Overview . . . 3155
137.2 The economy . . . 3156
137.3 The rational learning equilibrium . . . 3158
137.4 Computing the equilibrium . . . 3159
137.5 Posterior concentration . . . 3162
137.6 Equilibrium trades and prices . . . 3164
137.7 Convergence of posterior assessments . . . 3165
137.8 Sharpening the convergence result . . . 3166
137.9 Obstacles to convergence . . . 3168
137.10 Learning within versus learning about a rational expectations equilibrium . . . 3169
137.11 Summary . . . 3170
137.12 Exercises . . . 3170

138 Least Squares Learning in Self-Referential Models **3175**
138.1 Overview . . . 3175
138.2 The self-referential structure . . . 3176
138.3 Why a differential equation governs the limit . . . 3177
138.4 The governing ODE . . . 3178
138.5 Five examples . . . 3183
138.6 Example 1: Ordinary linear stochastic difference equations . . . 3183
138.7 Example 2: Bray's cobweb model . . . 3184
138.8 Counterexample: Unstable Bray feedback . . . 3187

138.9 Example 3: Bray-Savin supply-shifter model . . . 3189
138.10 Example 4: Hyperinflation / asset prices . . . 3192
138.11 Example 5: Investment under uncertainty . . . 3195
138.12 Connection to rational learning . . . 3198
138.13 Summary . . . 3198
138.14 Exercises . . . 3198

139 Affine Models of Asset Prices **3205**
139.1 Overview . . . 3205
139.2 The model . . . 3206
139.3 Pricing risky assets . . . 3209
139.4 Affine term structure of yields . . . 3211
139.5 Risk premiums . . . 3220
139.6 Risk-neutral probabilities . . . 3225
139.7 Distorted beliefs . . . 3228
139.8 Concluding remarks . . . 3233

140 The Recovery Theorem **3235**
140.1 Overview . . . 3235
140.2 Model setup . . . 3236
140.3 The recovery theorem . . . 3239
140.4 Numerical example . . . 3242
140.5 Natural vs. risk-neutral distributions . . . 3246
140.6 Effect of risk aversion . . . 3248
140.7 Recovering the discount rate . . . 3250
140.8 Tail risk: natural vs. risk-neutral probabilities of catastrophe . . . 3251
140.9 From option prices to transition prices . . . 3252
140.10 Testing efficient markets . . . 3253
140.11 Limitations and extensions . . . 3254
140.12 Exercises . . . 3255

141 Long-term Risk: An Operator Approach **3261**
141.1 Overview . . . 3261
141.2 Multiplicative functionals . . . 3263
141.3 Semigroups . . . 3265
141.4 The generator . . . 3268
141.5 Principal eigenfunctions . . . 3272
141.6 A finite-state Markov chain . . . 3275
141.7 The affine diffusion example . . . 3283
141.8 Long-run risk prices . . . 3292
141.9 Assumptions behind the scenes . . . 3302
141.10 Summary . . . 3303
141.11 Exercises . . . 3304

142 Misspecified Recovery **3313**
142.1 Overview . . . 3313
142.2 Three transition matrices . . . 3314
142.3 Martingale component . . . 3318
142.4 From matrices to the general framework . . . 3320
142.5 When recovery fails . . . 3327
142.6 Additional state vector . . . 3344
142.7 Measuring the martingale component . . . 3345
142.8 Lessons . . . 3346
142.9 Exercises . . . 3346

XVII Data and Empirics 3349

143 Pandas for Panel Data 3351
- 143.1 Overview . . . 3351
- 143.2 Slicing and Reshaping Data . . . 3352
- 143.3 Merging Dataframes and Filling NaNs . . . 3357
- 143.4 Grouping and Summarizing Data . . . 3362
- 143.5 Final Remarks . . . 3368
- 143.6 Exercises . . . 3368

144 Linear Regression in Python 3373
- 144.1 Overview . . . 3373
- 144.2 Simple Linear Regression . . . 3374
- 144.3 Extending the Linear Regression Model . . . 3380
- 144.4 Endogeneity . . . 3382
- 144.5 Summary . . . 3386
- 144.6 Exercises . . . 3386

145 Maximum Likelihood Estimation 3391
- 145.1 Overview . . . 3392
- 145.2 Set up and assumptions . . . 3392
- 145.3 Conditional distributions . . . 3396
- 145.4 Maximum likelihood estimation . . . 3397
- 145.5 MLE with numerical methods . . . 3399
- 145.6 Maximum likelihood estimation with `statsmodels` . . . 3405
- 145.7 Summary . . . 3409
- 145.8 Exercises . . . 3410

146 A Linear Model of Unemployment 3413
- 146.1 Overview . . . 3413
- 146.2 The natural rate versus hysteresis . . . 3414
- 146.3 The data . . . 3414
- 146.4 A linear model of unemployment . . . 3416
- 146.5 A random walk would wander off . . . 3417
- 146.6 Monthly versus annual . . . 3419
- 146.7 What the model misses . . . 3419
- 146.8 Exercises . . . 3423

147 Unemployment Dynamics with Asymmetric Shocks 3427
- 147.1 The data . . . 3428
- 147.2 A model with asymmetric shocks . . . 3428
- 147.3 Bayesian estimation . . . 3431
- 147.4 Comparing models with cross-validation . . . 3433
- 147.5 Conclusion . . . 3437
- 147.6 Exercises . . . 3437

148 Two Illustrations of the Quantity Theory of Money 3439
- 148.1 Overview . . . 3440
- 148.2 Lucas's moving averages . . . 3441
- 148.3 Whiteman's reinterpretation . . . 3441
- 148.4 Data . . . 3442
- 148.5 Scatter plots . . . 3444
- 148.6 Regressions on filtered data . . . 3445
- 148.7 Estimates of $h(0)$ from a VAR . . . 3446
- 148.8 A model for monetary policy analysis . . . 3448

148.9 Three things worth checking . . . 3449
148.10 Solving the model . . . 3450
148.11 The Sargent-Surico model in canonical form . . . 3452
148.12 Bayesian estimation . . . 3456
148.13 How monetary policy moves the low-frequency slopes . . . 3461
148.14 Estimating with Hamiltonian Monte Carlo . . . 3464
148.15 Concluding remarks . . . 3477
148.16 Exercises . . . 3478

XVIII Auctions **3483**

149 First-Price and Second-Price Auctions **3485**
149.1 First-price sealed-bid auction (FPSB) . . . 3485
149.2 Second-price sealed-bid auction (SPSB) . . . 3487
149.3 Characterization of SPSB auction . . . 3487
149.4 Uniform distribution of private values . . . 3487
149.5 Setup . . . 3487
149.6 First price sealed bid auction . . . 3488
149.7 Second price sealed bid auction . . . 3488
149.8 Python code . . . 3488
149.9 Revenue equivalence theorem . . . 3490
149.10 Calculation of bid price in FPSB . . . 3492
149.11 χ^2 Distribution . . . 3494
149.12 Code summary . . . 3497
149.13 Exercises . . . 3501
149.14 Further reading . . . 3506

150 Multiple Good Allocation Mechanisms **3507**
150.1 Overview . . . 3507
150.2 Ascending bids auction for multiple goods . . . 3507
150.3 A benevolent planner . . . 3508
150.4 Comparing the two mechanisms . . . 3508
150.5 Ascending bid auction . . . 3509
150.6 Pseudocode . . . 3509
150.7 An example . . . 3511
150.8 A Python class . . . 3519
150.9 Robustness checks . . . 3527
150.10 A Groves-Clarke mechanism . . . 3539
150.11 An example solved by hand . . . 3540
150.12 Another Python class . . . 3543
150.13 Is the greedy mechanism a VCG mechanism? . . . 3549
150.14 Exercises . . . 3551

XIX Other **3555**

151 Troubleshooting **3557**
151.1 Fixing Your Local Environment . . . 3557
151.2 Reporting an Issue . . . 3558

152 References **3559**

153 Execution Statistics **3561**

Bibliography **3567**

Proof Index **3591**

Index **3597**

This website presents a set of lectures on quantitative economic modeling.

- Tools and Techniques
  - *Modeling COVID-19*
  - *Linear Algebra*
  - *QR Decomposition*
  - *Circulant Matrices*
  - *Singular Value Decomposition (SVD)*
  - *VARs and DMDs*
  - *Using Newton's Method to Solve Economic Models*
- Elementary Statistics
  - *Elementary Probability with Matrices*
  - *Some Probability Distributions*
  - *LLN and CLT*
  - *Two Meanings of Probability*
  - *Multivariate Hypergeometric Distribution*
  - *Multivariate Normal Distribution*
  - *Fault Tree Uncertainties*
  - *Introduction to Artificial Neural Networks*
  - *Randomized Response Surveys*
  - *Expected Utilities of Random Responses*
- Bayes Law
  - *Non-Conjugate Priors*
  - *Posterior Distributions for AR(1) Parameters*
  - *Forecasting an AR(1) Process*
- Statistics and Information
  - *Statistical Divergence Measures*
  - *Likelihood Ratio Processes*
  - *Heterogeneous Beliefs and Financial Markets*
  - *Likelihood Processes For VAR Models*
  - *Mean of a Likelihood Ratio Process*
  - *A Problem that Stumped Milton Friedman*
  - *A Bayesian Formulation of Friedman and Wald's Problem*
  - *Exchangeability and Bayesian Updating*
  - *Likelihood Ratio Processes and Bayesian Learning*
  - *Blackwell's Theorem on Comparing Experiments*
  - *Information and Market Equilibrium*

  - *Incorrect Models*
  - *Bayesian versus Frequentist Decision Rules*
  - *Merging of Opinions: The Blackwell–Dubins Theorem*
  - *Survival and Long-Run Dynamics under Recursive Preferences*
  - *The Design and Price of Information*
  - *Market Diffusion with Two-Sided Learning*
  - *Which Moments to Match?*
- Linear Programming
  - *Optimal Transport*
  - *Von Neumann Growth Model (and a Generalization)*
- Introduction to Dynamics
  - *Finite Markov Chains*
  - *Inventory Dynamics*
  - *Linear State Space Models*
  - *Samuelson Multiplier-Accelerator*
  - *The Acceleration Principle and the Nature of Business Cycles*
  - *Kesten Processes and Firm Dynamics*
  - *Wealth Distribution Dynamics*
  - *A First Look at the Kalman Filter*
  - *Another Look at the Kalman Filter*
  - *The Kalman Filter and Vector Autoregressions*
  - *Vector Autoregressions for Subsets of Variables*
  - *Organization Capital*
  - *Two Models of Measurements and the Investment Accelerator*
- Search
  - *Job Search I: The McCall Search Model*
  - *Job Search II: Search and Separation*
  - *Job Search III: Search with Separation and Markov Wages*
  - *Job Search IV: Fitted Value Function Iteration*
  - *Job Search V: Risk-Sensitive Preferences*
  - *Job Search VI: Persistent and Transitory Wage Shocks*
  - *Job Search VII: Modeling Career Choice*
  - *Job Search VIII: On-the-Job Search*
  - *Job Search IX: Search with Learning*
- Reinforcement Learning
  - *Inventory Management via Q-Learning*

  - *Risk-Sensitive Inventory Management via Q-Learning*
  - *Job Search X: Search with Q-Learning*
- Introduction to Optimal Savings
  - *Optimal Savings I: Cake Eating*
  - *Optimal Savings II: Numerical Cake Eating*
  - *Optimal Savings III: Stochastic Returns*
  - *Optimal Savings IV: Time Iteration*
  - *Optimal Savings V: The Endogenous Grid Method*
  - *Optimal Savings VI: EGM with JAX*
- Household Problems
  - *The Income Fluctuation Problem I: Discretization and VFI*
  - *The Income Fluctuation Problem II: Optimistic Policy Iteration*
  - *The Income Fluctuation Problem III: The Endogenous Grid Method*
  - *The Income Fluctuation Problem IV: Transient Income Shocks*
  - *The Income Fluctuation Problem V: Stochastic Returns on Assets*
- LQ Control
  - *LQ Control: Foundations*
  - *Lagrangian for LQ Control*
  - *Eliminating Cross Products*
  - *The Permanent Income Model*
  - *Permanent Income II: LQ Techniques*
  - *Robust Permanent Income and Pricing*
  - *Certainty Equivalence*
  - *Certainty Equivalence and Model Uncertainty*
  - *The LQ Permanent Income Model*
  - *Consumption Smoothing with Incomplete and Complete Markets*
  - *Robust Consumption Smoothing and Precautionary Savings*
  - *A Robust LQ Bewley Model*
  - *Production Smoothing via Inventories*
- Bounded Rationality in Macroeconomics
  - *A Peculiar Definition of Bounded Rationality*
  - *Adaptive Agents in an Overlapping Generations Monetary Economy*
  - *Learning, Approximation, and Equilibrium Computation*
  - *Exchange Rate Indeterminacy, Learning, and Experiments*
  - *Networks, Genetic Algorithms, and Classifier Systems*
  - *Money as a Medium of Exchange among Artificially Intelligent Agents*

  - *1993 Prospects for Bounded Rationality in Macroeconomics*
- Phillips Curve Tradeoffs
  - *The Rise and Fall of U.S. Inflation*
  - *The Credibility Problem*
  - *Credible Government Policies*
  - *Adaptive Expectations and the Phelps Problem*
  - *Optimal Misspecified Beliefs*
  - *Self-Confirming Equilibria*
  - *Adaptive Learning and Escape Dynamics*
  - *Escaping Nash Inflation*
  - *Priors, Escapes, and Learning Cycles*
  - *Econometric Policy Evaluation*
  - *Drifts and Volatilities*
  - *Drifts and Volatilities after 2000*
  - *The Lost Conquest: Fed Policy in the 2020s*
- Optimal Growth
  - *Cass-Koopmans Model*
  - *Cass-Koopmans Competitive Equilibrium*
  - *Cass-Koopmans Model with Distorting Taxes*
  - *Two-Country Model with Distorting Taxes*
- Multiple Agent Models
  - *A Lake Model of Employment and Unemployment*
  - *Lake Model with an Endogenous Job Finding Rate*
  - *Rational Expectations Equilibrium*
  - *Investment Under Uncertainty*
  - *Equilibrium Search and Unemployment*
  - *Optimal Growth Under Uncertainty and Tobin's q*
  - *Stability in Linear Rational Expectations Models*
  - *Markov Perfect Equilibrium*
  - *Uncertainty Traps*
  - *The Aiyagari Model*
  - *The Aiyagari Model with Endogenous Grid Method*
  - *Transitions in an Overlapping Generations Model*
  - *A Long-Lived, Heterogeneous Agents, Overlapping Generations Model*
  - *Two Computations to Fund Social Security*
  - *Social Security Trust Fund Depletion*

  - *Linear Quadratic Mean Field Games*
  - *Linear Quadratic Mean Field Games: Extensions*
- Asset Pricing and Finance
  - *Asset Pricing: Finite State Models*
  - *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*
  - *Estimating Euler Equations by Generalized Method of Moments*
  - *Competitive Equilibria with Arrow Securities*
  - *Heterogeneous Beliefs and Bubbles*
  - *Speculative Behavior with Bayesian Learning*
  - *Rational Learning and Rational Expectations*
  - *Least Squares Learning in Self-Referential Models*
  - *Affine Models of Asset Prices*
  - *The Recovery Theorem*
  - *Long-term Risk: An Operator Approach*
  - *Misspecified Recovery*
- Data and Empirics
  - *Pandas for Panel Data*
  - *Linear Regression in Python*
  - *Maximum Likelihood Estimation*
  - *A Linear Model of Unemployment*
  - *Unemployment Dynamics with Asymmetric Shocks*
  - *Two Illustrations of the Quantity Theory of Money*
- Auctions
  - *First-Price and Second-Price Auctions*
  - *Multiple Good Allocation Mechanisms*
- Other
  - *Troubleshooting*
  - *References*
  - *Execution Statistics*

Part I

Tools and Techniques

CHAPTER

ONE

MODELING COVID-19

Contents

- *Modeling COVID-19*
  - *Overview*
  - *The SEIR model*
  - *Implementation*
  - *Experiments*
  - *Ending lockdown*

1.1 Overview

This is a Python version of the code for analyzing the COVID-19 pandemic provided by Andrew Atkeson.

See, in particular

- NBER Working Paper No. 26867
- COVID-19 Working papers and code

The purpose of his notes is to introduce economists to quantitative modeling of infectious disease dynamics.

Dynamics are modeled using a standard SEIR (Susceptible-Exposed-Infected-Removed) model of disease spread.

The model dynamics are represented by a system of ordinary differential equations.

The main objective is to study the impact of suppression through social distancing on the spread of the infection.

The focus is on US outcomes but the parameters can be adjusted to study other countries.

We will use the following standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
from numpy import exp
```

We will also use SciPy's numerical routine odeint for solving differential equations.

```
from scipy.integrate import odeint
```

This routine calls into compiled code from the FORTRAN library odepack.

1.2 The SEIR model

In the version of the SEIR model we will analyze there are four states.

All individuals in the population are assumed to be in one of these four states.

The states are: susceptible (S), exposed (E), infected (I), and removed (R).

Assumption 1.2.1

- Those in state R have been infected and either recovered or died.
- Those who have recovered are assumed to have acquired immunity.
- Those in the exposed group are not yet infectious.

1.2.1 Time path

The flow across states follows the path $S \to E \to I \to R$.

When the effective reproduction number exceeds one, infections initially increase but eventually decline as the susceptible population falls.

The interest is primarily in

- the number of infections at a given time (which determines whether or not the health care system is overwhelmed) and
- how long the caseload can be deferred (hopefully until a vaccine arrives)

Using lowercase letters for the fraction of the population in each state, the dynamics are

$$\begin{aligned} \dot{s}(t) &= -\beta(t)\, s(t)\, i(t) \\ \dot{e}(t) &= \beta(t)\, s(t)\, i(t) - \sigma e(t) \\ \dot{i}(t) &= \sigma e(t) - \gamma i(t) \end{aligned} \tag{1.1}$$

In these equations,

- $\beta(t)$ is called the **transmission rate** (the rate at which individuals bump into others and expose them to the virus).
- σ is called the **infection rate** (the rate at which those who are exposed become infected).
- γ is called the **removal rate** (the rate at which infected people recover or die).
- The dot symbol $\dot{y}$ represents the time derivative dy/dt.

We do not need to model the fraction r of the population in state R separately because the states form a partition.

In particular, the "removed" fraction of the population is $r = 1 - s - e - i$.

We will also track $c = i + r$, which is the cumulative caseload (i.e., all those who have or have had the infection).

The system (1.1) can be written in vector form as

$$\dot{x} = F(x, t), \qquad x := (s, e, i) \tag{1.2}$$

for a suitable definition of F (see the code below).

1.2.2 Parameters

Both σ and γ are thought of as fixed, biologically determined parameters.

As in Atkeson's note, we set

- $\sigma = 1/5.2$ to reflect an average incubation period of 5.2 days.
- $\gamma = 1/18$ to match an average illness duration of 18 days.

The transmission rate is modeled as

- $\beta(t) := R(t)\gamma$ where $R(t)$ is the **effective reproduction number** at time t.

(The notation is slightly confusing, since $R(t)$ is different to R, the symbol that represents the removed state.)

1.3 Implementation

First we set the population size to match the US.

```
pop_size = 3.3e8
```

Next we fix parameters as described above.

```
γ = 1 / 18
σ = 1 / 5.2
```

Now we construct a function that represents F in (1.2)

```
def F(x, t, R0=1.6):
    """
    Time derivative of the state vector.

        * x is the state vector (array_like)
        * t is time (scalar)
        * R0 is the effective transmission rate, defaulting to a constant

    """
    s, e, i = x

    # New exposure of susceptibles
    β = R0(t) * γ if callable(R0) else R0 * γ
    ne = β * s * i

    # Time derivatives
    ds = - ne
    de = ne - σ * e
    di = σ * e - γ * i

    return ds, de, di
```

Note that `R0` can be either constant or a given function of time.

The initial conditions are calibrated to a population of 330 million.

The value $i_0 = 10^{-7}$ represents 33 initially infected people, while $e_0 = 4i_0$ represents 132 exposed people.

Setting $s_0 = 1 - i_0 - e_0$ assigns the remaining population to the susceptible state and sets the initial removed fraction to zero.

```
# initial conditions of s, e, i
i_0 = 1e-7
e_0 = 4 * i_0
s_0 = 1 - i_0 - e_0
```

In vector form the initial condition is

```
x_0 = s_0, e_0, i_0
```

We solve for the time path numerically using `odeint`, at a sequence of dates `t_vec`.

```
def solve_path(R0, t_vec, x_init=x_0):
    """
    Solve for i(t) and c(t) via numerical integration,
    given the time path for R0.

    """
    G = lambda x, t: F(x, t, R0)
    s_path, e_path, i_path = odeint(G, x_init, t_vec).transpose()

    c_path = 1 - s_path - e_path       # cumulative cases
    return i_path, c_path
```

1.4 Experiments

Let's run some experiments using this code.

The time period we investigate will be 550 days, or around 18 months:

```
t_length = 550
grid_size = 1000
t_vec = np.linspace(0, t_length, grid_size)
```

1.4.1 Experiment 1: constant R0 case

Let's start with the case where `R0` is constant.

We calculate the time path of infected people under different assumptions for `R0`:

```
R0_vals = np.linspace(1.6, 3.0, 6)
labels = [f'$R_0 = {r:.2f}$' for r in R0_vals]
i_paths, c_paths = [], []

for r in R0_vals:
    i_path, c_path = solve_path(r, t_vec)
    i_paths.append(i_path)
    c_paths.append(c_path)
```

Here's some code to plot the time paths.

```
def plot_paths(paths, labels, ylabel, times=t_vec):

    fig, ax = plt.subplots()
```

(continues on next page)

(continued from previous page)

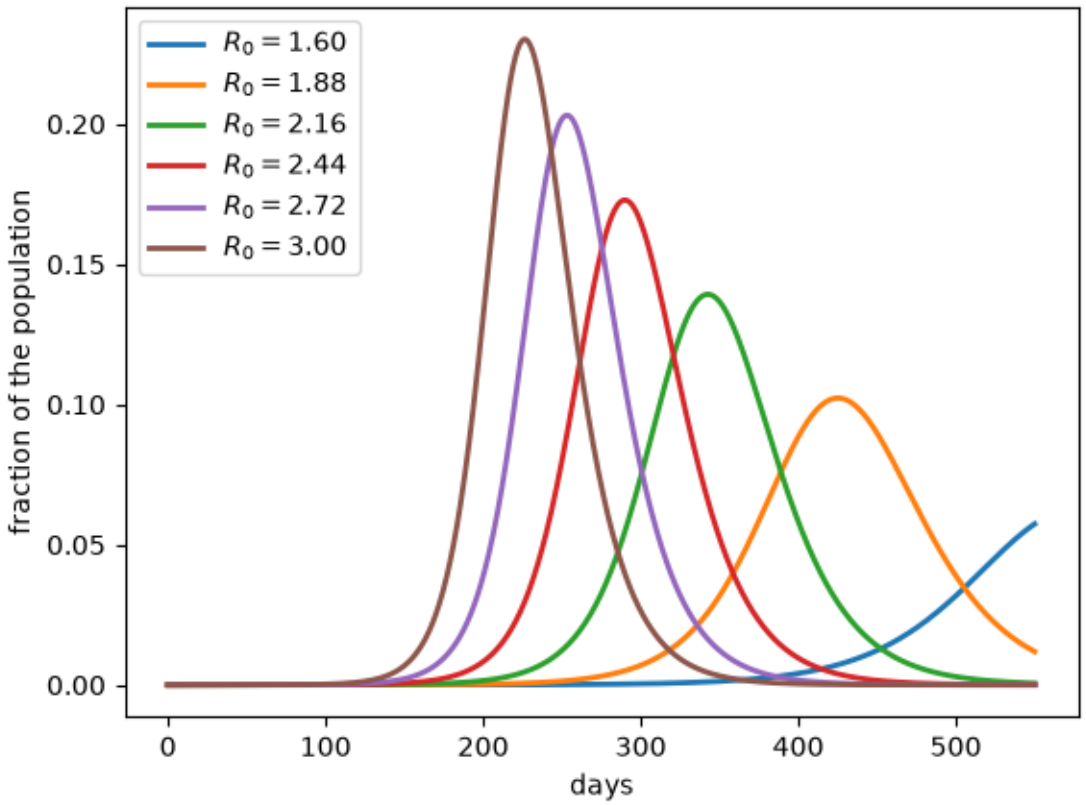
```
    for path, label in zip(paths, labels):
        ax.plot(times, path, lw=2, label=label)

    ax.set_xlabel('days')
    ax.set_ylabel(ylabel)
    ax.legend(loc='upper left')

    plt.show()
```

Let's plot current cases as a fraction of the population.

```
plot_paths(i_paths, labels, ylabel='fraction of the population')
```

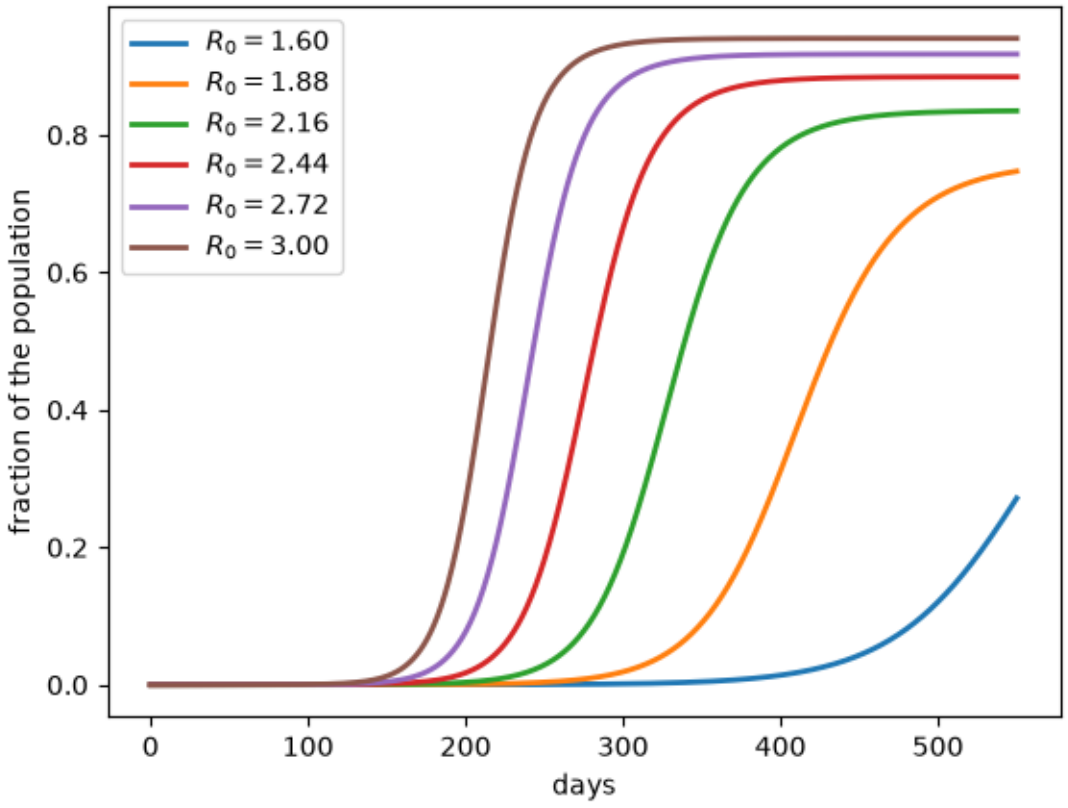


As expected, lower effective transmission rates defer the peak of infections.

They also lead to a lower peak in current cases.

Here are cumulative cases, as a fraction of population:

```
plot_paths(c_paths, labels, ylabel='fraction of the population')
```



1.4.2 Experiment 2: changing mitigation

Let's look at a scenario where mitigation (e.g., social distancing) is successively imposed.

Here's a specification for `R0` as a function of time.

```
def R0_mitigating(t, r0=3, η=1, r_bar=1.6):
    R0 = r0 * exp(- η * t) + (1 - exp(- η * t)) * r_bar
    return R0
```

The idea is that `R0` starts off at 3 and falls to 1.6.

This is due to progressive adoption of stricter mitigation measures.

The parameter `η` controls the rate, or the speed at which restrictions are imposed.

Since t is measured in days, η is measured per day, and its reciprocal, $1/\eta$, is the adjustment period.

The values below correspond to adjustment periods of 5, 10, 20, 50, and 100 days:

```
η_vals = 1/5, 1/10, 1/20, 1/50, 1/100
labels = [fr'$\eta = {η:.2f}$' for η in η_vals]
```

This is what the time path of `R0` looks like at these alternative rates:

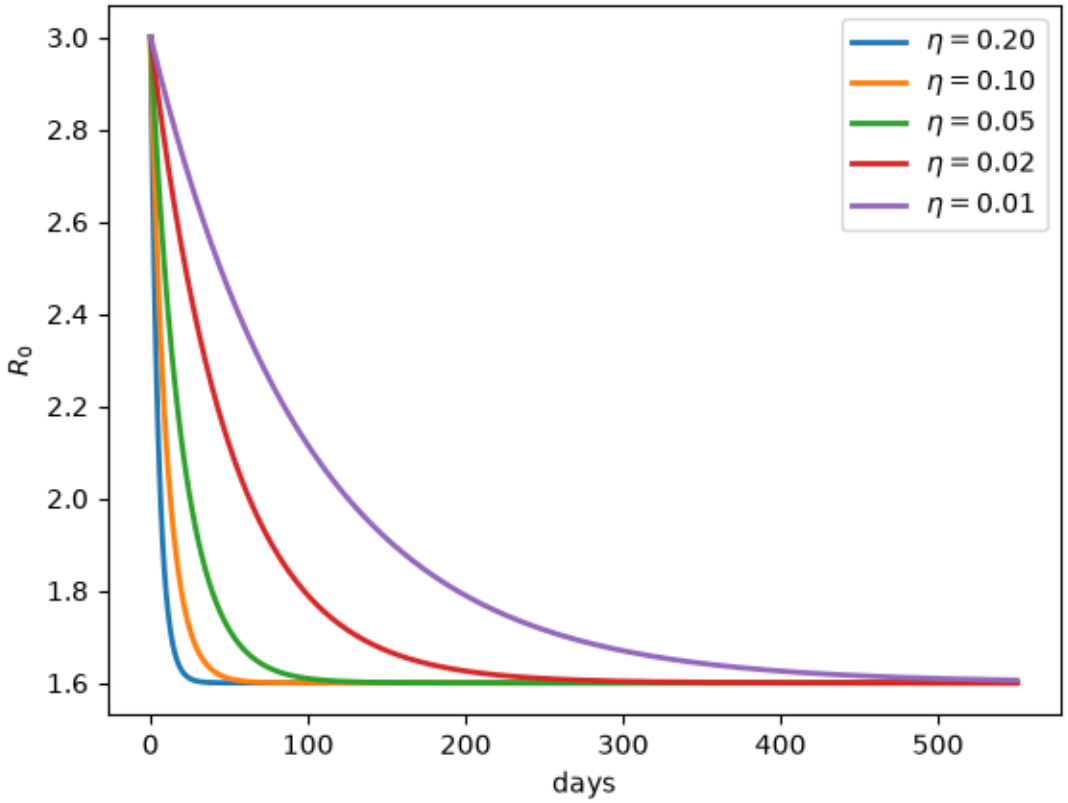
```
fig, ax = plt.subplots()

for η, label in zip(η_vals, labels):
    ax.plot(t_vec, R0_mitigating(t_vec, η=η), lw=2, label=label)
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel('days')
ax.set_ylabel('$R_0$')
ax.legend()
plt.show()
```



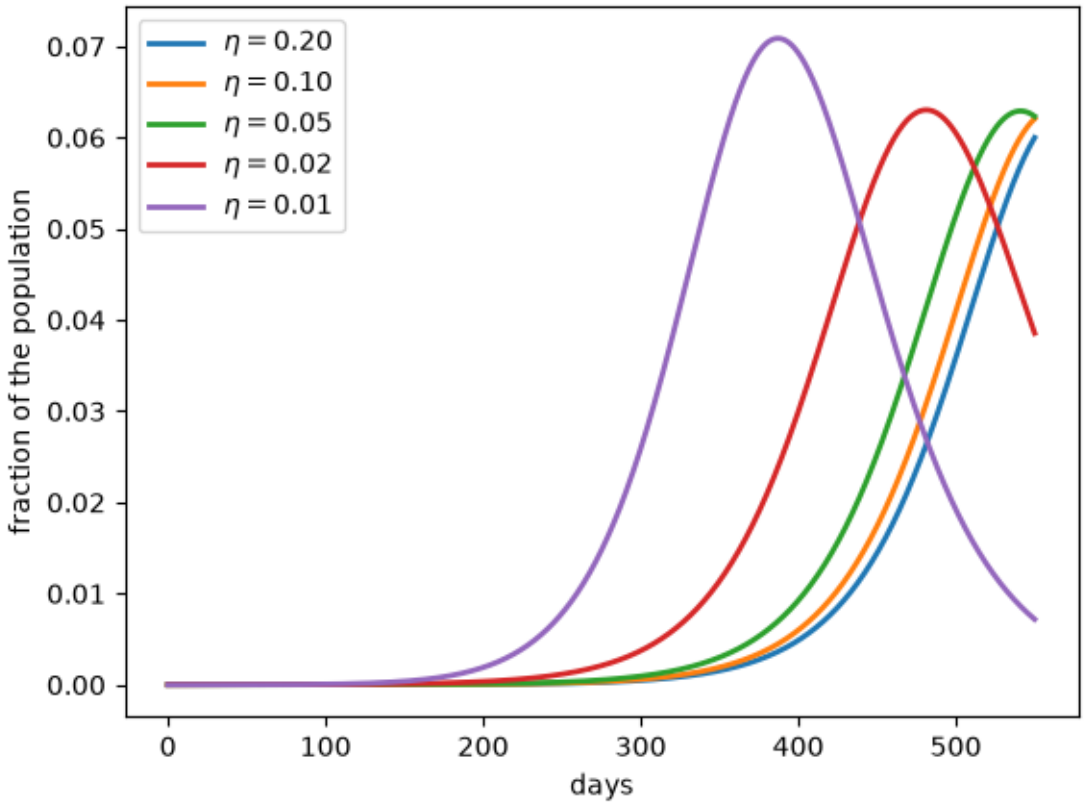
Let's calculate the time path of infected people:

```
i_paths, c_paths = [], []

for η in η_vals:
    R0 = lambda t: R0_mitigating(t, η=η)
    i_path, c_path = solve_path(R0, t_vec)
    i_paths.append(i_path)
    c_paths.append(c_path)
```

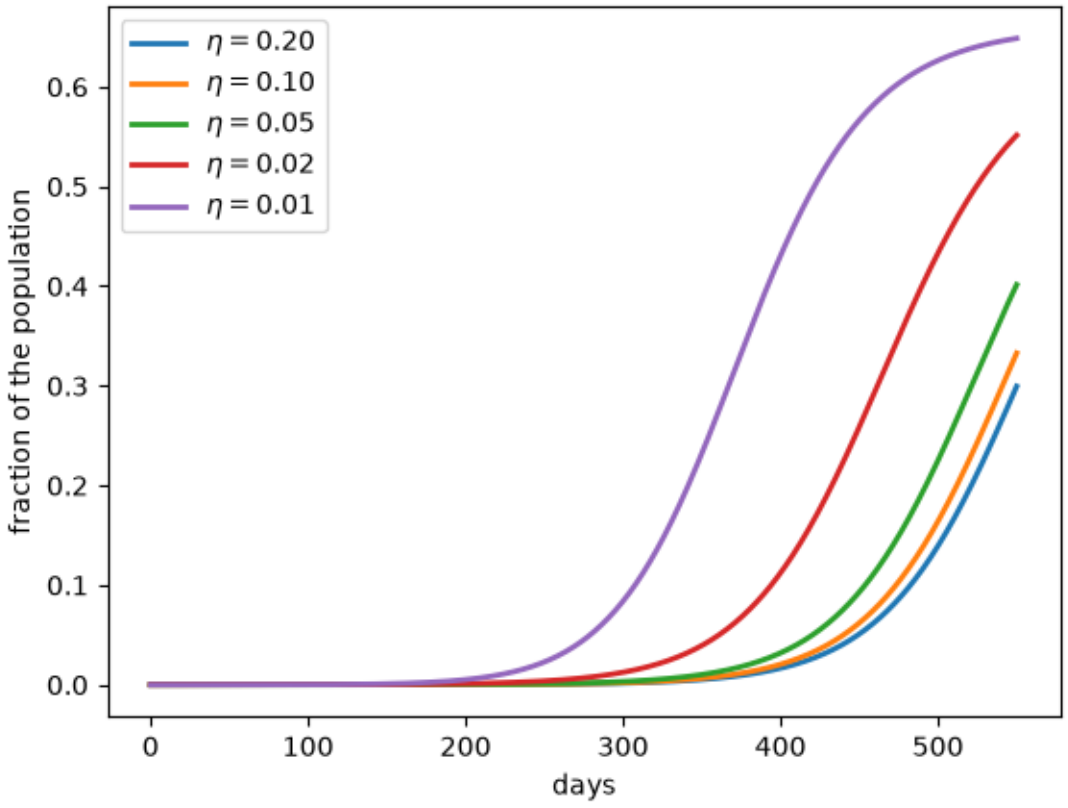
These are current cases under the different scenarios:

```
plot_paths(i_paths, labels, ylabel='fraction of the population')
```



Here are cumulative cases, as a fraction of population:

```
plot_paths(c_paths, labels, ylabel='fraction of the population')
```



Faster implementation of mitigation mainly delays the infection peak, with a smaller effect on its height.

1.5 Ending lockdown

The following replicates additional results by Andrew Atkeson on the timing of lifting lockdown.

Consider these two mitigation scenarios:

1. $R_t = 0.5$ for 30 days and then $R_t = 2$ for the remaining 17 months. This corresponds to lifting lockdown in 30 days.
2. $R_t = 0.5$ for 120 days and then $R_t = 2$ for the remaining 14 months. This corresponds to lifting lockdown in 4 months.

The parameters considered here start the model with 25,000 active infections and 75,000 agents already exposed to the virus and thus soon to be contagious.

```
# initial conditions
i_0 = 25_000 / pop_size
e_0 = 75_000 / pop_size
s_0 = 1 - i_0 - e_0
x_0 = s_0, e_0, i_0
```

Let's calculate the paths:

```
R0_paths = (lambda t: 0.5 if t < 30 else 2,
            lambda t: 0.5 if t < 120 else 2)
```

(continues on next page)

(continued from previous page)

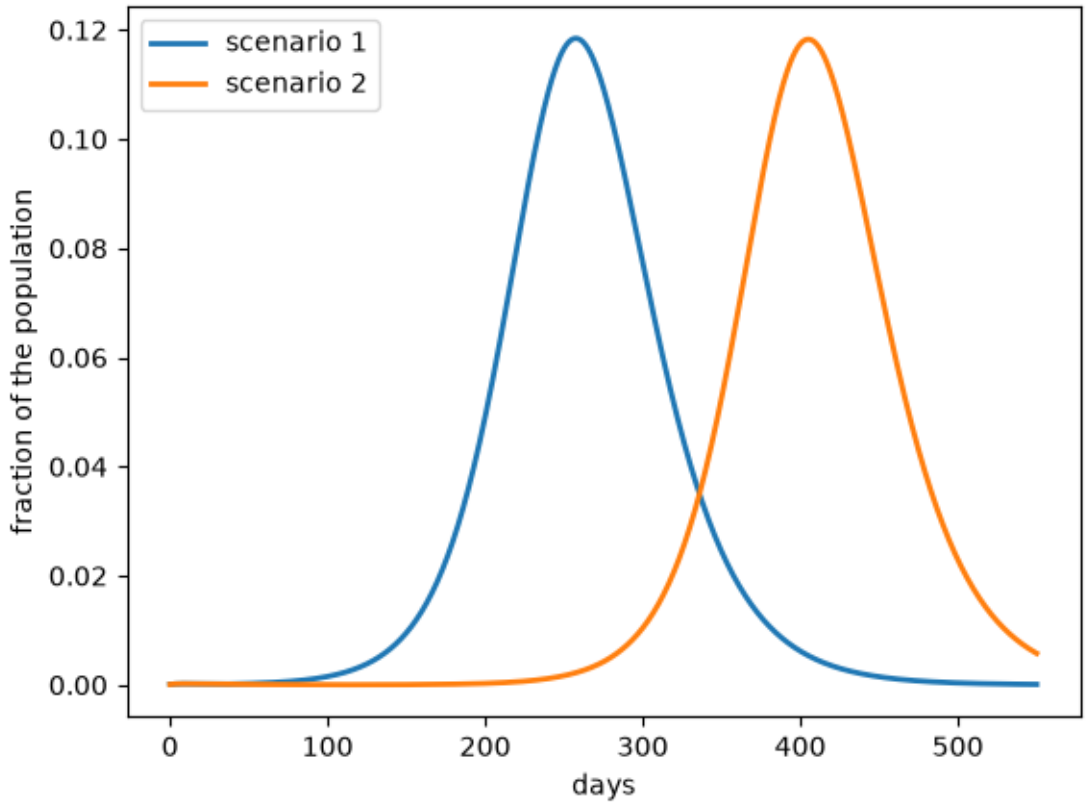
```
labels = [f'scenario {i}' for i in (1, 2)]

i_paths, c_paths = [], []

for R0 in R0_paths:
    i_path, c_path = solve_path(R0, t_vec, x_init=x_0)
    i_paths.append(i_path)
    c_paths.append(c_path)
```

Here is the fraction of the population with an active infection:

```
plot_paths(i_paths, labels, ylabel='fraction of the population')
```



Both scenarios produce approximately the same infection peak, but the longer lockdown delays it.

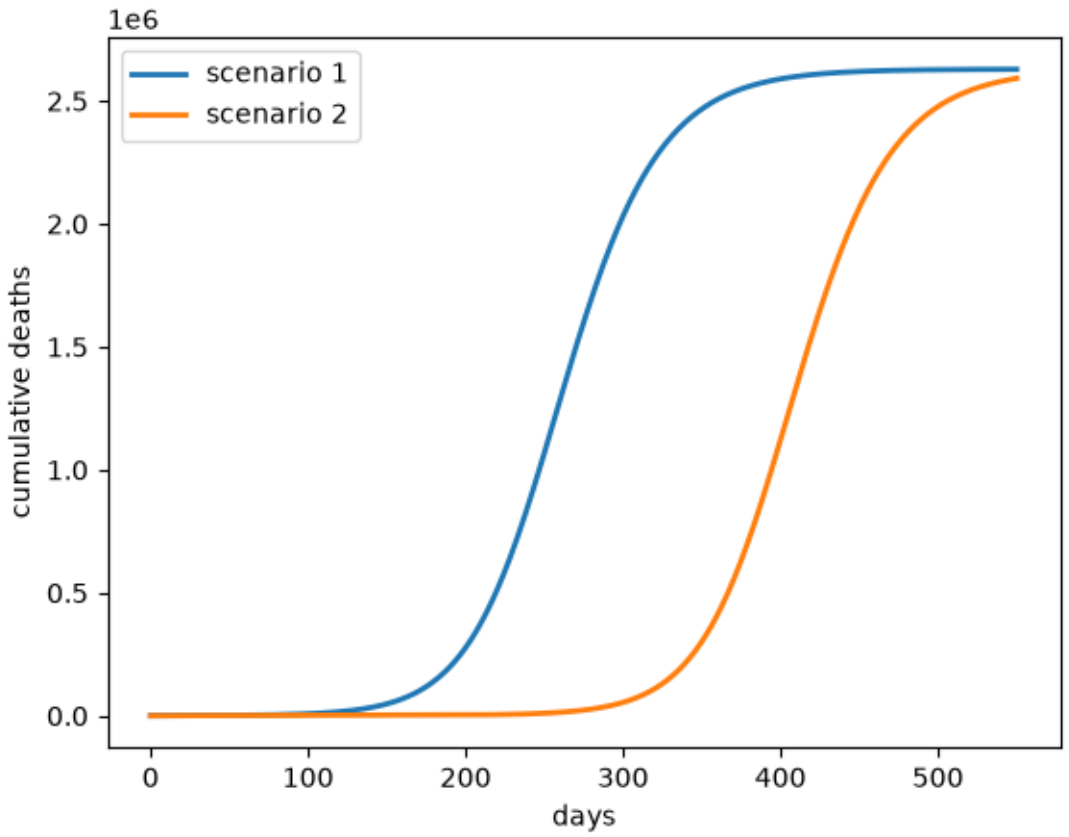
What kind of mortality can we expect under these scenarios?

Suppose that 1% of cases result in death.

```
ν = 0.01
```

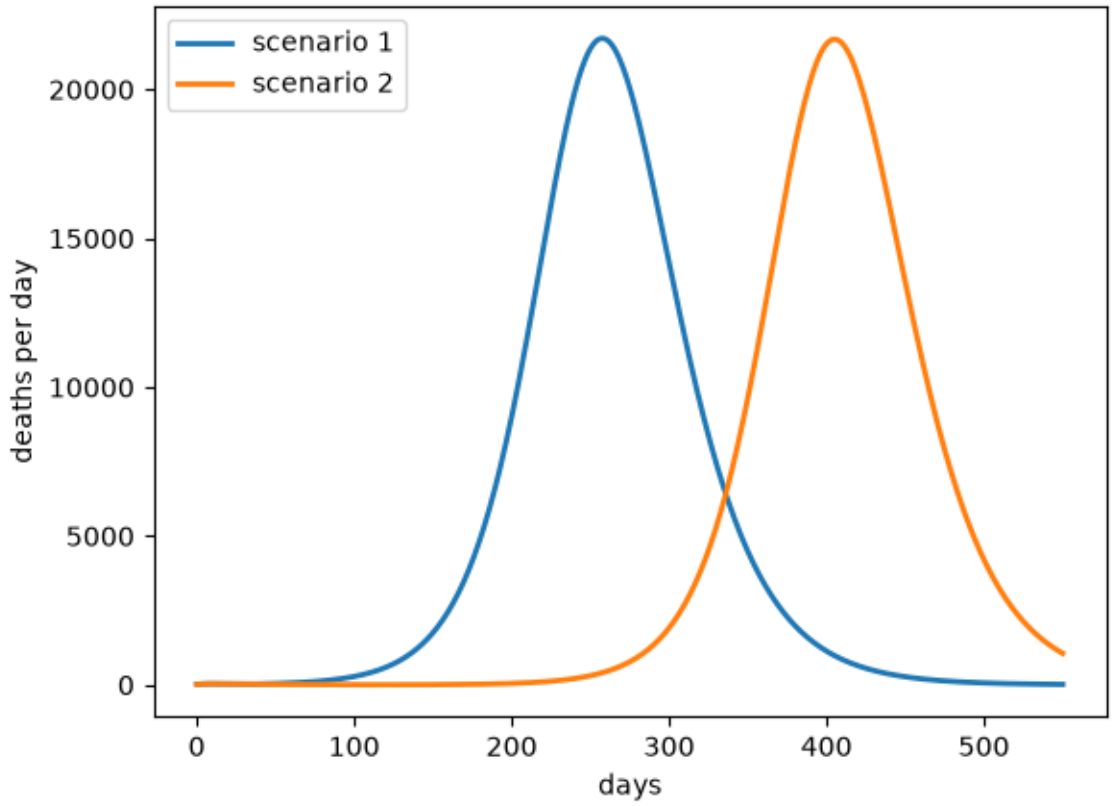
This is the cumulative number of deaths:

```
paths = [(c_path - i_path) * ν * pop_size
         for c_path, i_path in zip(c_paths, i_paths)]
plot_paths(paths, labels, ylabel='cumulative deaths')
```



This is the number of deaths per day:

```
paths = [path * ν * γ * pop_size for path in i_paths]
plot_paths(paths, labels, ylabel='deaths per day')
```



Pushing the peak of the curve further into the future may reduce cumulative deaths if a vaccine is found.

CHAPTER

TWO

LINEAR ALGEBRA

Contents

- *Linear Algebra*
  - *Overview*
  - *Vectors*
  - *Matrices*
  - *Solving systems of equations*
  - *Eigenvalues and eigenvectors*
  - *Further topics*
  - *Exercises*

2.1 Overview

Linear algebra is one of the most useful branches of applied mathematics for economists to invest in.

For example, many applied problems in economics and finance require the solution of a linear system of equations, such as

$$\begin{aligned} y_1 &= ax_1 + bx_2 \\ y_2 &= cx_1 + dx_2 \end{aligned}$$

or, more generally,

$$\begin{aligned} y_1 &= a_{11}x_1 + a_{12}x_2 + \cdots + a_{1k}x_k \\ &\vdots \\ y_n &= a_{n1}x_1 + a_{n2}x_2 + \cdots + a_{nk}x_k \end{aligned} \tag{2.1}$$

The objective here is to solve for the "unknowns" $x_1, \ldots, x_k$ given $a_{11}, \ldots, a_{nk}$ and $y_1, \ldots, y_n$.

When considering such problems, it is essential that we first consider at least some of the following questions

- Does a solution actually exist?
- Are there in fact many solutions, and if so how should we interpret them?
- If no solution exists, is there a best "approximate" solution?

- If a solution exists, how should we compute it?

These are the kinds of topics addressed by linear algebra.

In this lecture we will cover the basics of linear and matrix algebra, treating both theory and computation.

We admit some overlap with NumPy, where operations on NumPy arrays were first explained.

Note that this lecture is more theoretical than most, and contains background material that will be used in applications as we go along.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
from matplotlib import cm
from mpl_toolkits.mplot3d import Axes3D
from scipy.linalg import inv, solve, det, eig
```

2.2 Vectors

A **vector** of length n is just a sequence (or array, or tuple) of n numbers, which we write as $x = (x_1, \ldots, x_n)$ or $x = [x_1, \ldots, x_n]$.

We will write these sequences either horizontally or vertically as we please.

(Later, when we wish to perform certain matrix operations, it will become necessary to distinguish between the two)

The set of all n-vectors is denoted by $\mathbb{R}^n$.

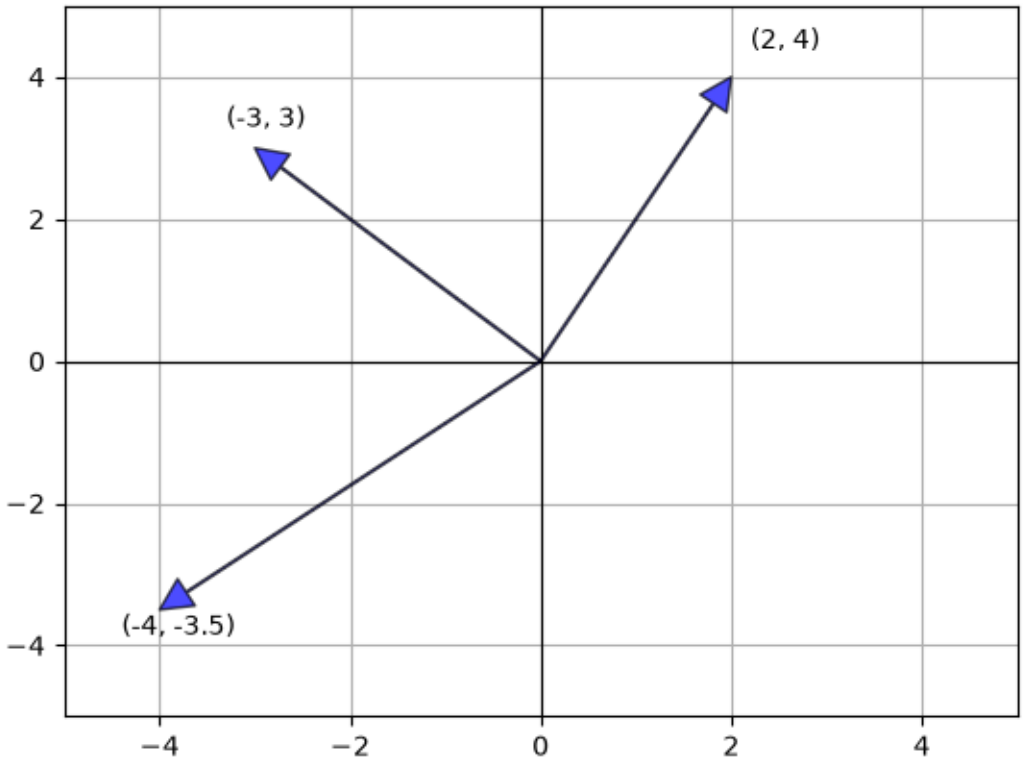
For example, $\mathbb{R}^2$ is the plane, and a vector in $\mathbb{R}^2$ is just a point in the plane.

Traditionally, vectors are represented visually as arrows from the origin to the point.

The following figure represents three vectors in this manner

```
fig, ax = plt.subplots()
# Draw the axes through the origin
ax.axhline(0, color='black', lw=0.8)
ax.axvline(0, color='black', lw=0.8)

ax.set(xlim=(-5, 5), ylim=(-5, 5))
ax.grid()
vecs = ((2, 4), (-3, 3), (-4, -3.5))
for v in vecs:
    ax.annotate('', xy=v, xytext=(0, 0),
                arrowprops=dict(facecolor='blue',
                shrink=0,
                alpha=0.7,
                width=0.5))
    ax.text(1.1 * v[0], 1.1 * v[1], str(v))
plt.show()
```



2.2.1 Vector operations

The two most common operators for vectors are addition and scalar multiplication, which we now describe.

As a matter of definition, when we add two vectors, we add them element-by-element

$$x + y = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} + \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} := \begin{bmatrix} x_1 + y_1 \\ x_2 + y_2 \\ \vdots \\ x_n + y_n \end{bmatrix}$$

Scalar multiplication is an operation that takes a number γ and a vector x and produces

$$\gamma x := \begin{bmatrix} \gamma x_1 \\ \gamma x_2 \\ \vdots \\ \gamma x_n \end{bmatrix}$$

Scalar multiplication is illustrated in the next figure

```
fig, ax = plt.subplots()
# Draw the axes through the origin
ax.axhline(0, color='black', lw=0.8)
ax.axvline(0, color='black', lw=0.8)

ax.set(xlim=(-5, 5), ylim=(-5, 5))
x = (2, 2)
ax.annotate('', xy=x, xytext=(0, 0),
            arrowprops=dict(facecolor='blue',
```

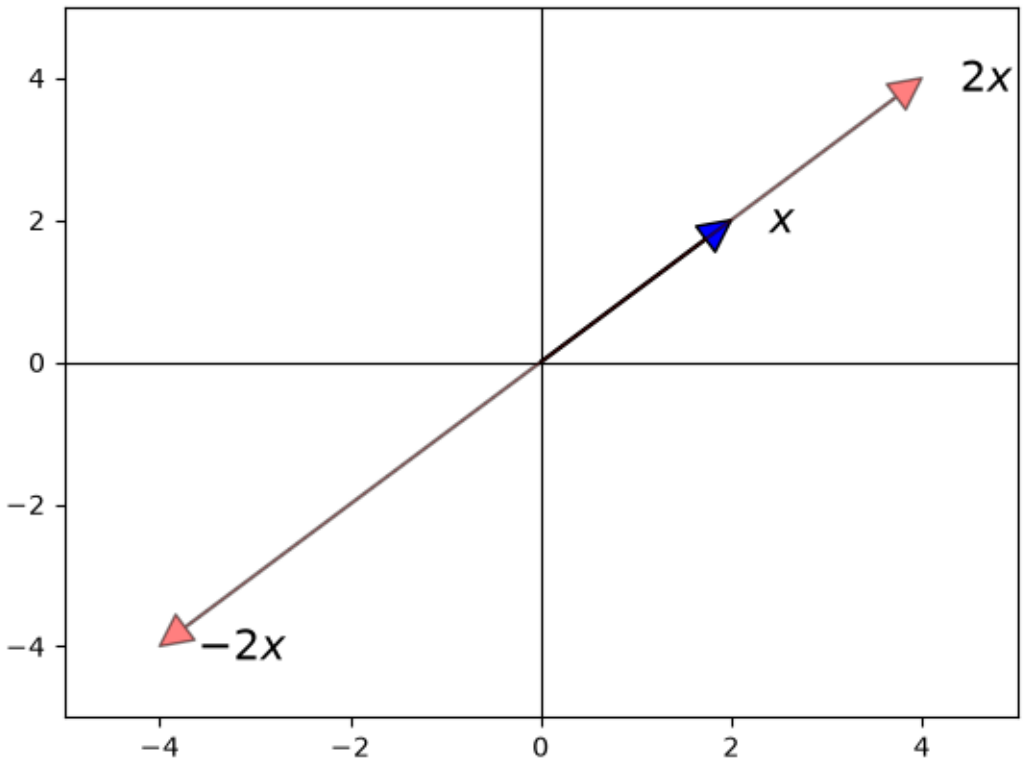
(continues on next page)

(continued from previous page)

```
            shrink=0,
            alpha=1,
            width=0.5))
ax.text(x[0] + 0.4, x[1] - 0.2, '$x$', fontsize='16')


scalars = (-2, 2)
x = np.array(x)

for s in scalars:
    v = s * x
    ax.annotate('', xy=v, xytext=(0, 0),
                arrowprops=dict(facecolor='red',
                shrink=0,
                alpha=0.5,
                width=0.5))
    ax.text(v[0] + 0.4, v[1] - 0.2, f'${s} x$', fontsize='16')
plt.show()
```



In Python, a vector can be represented as a list or tuple, such as `x = (2, 4, 6)`, but is more commonly represented as a NumPy array.

One advantage of NumPy arrays is that scalar multiplication and addition have very natural syntax

```
x = np.ones(3)             # Vector of three ones
y = np.array((2, 4, 6))    # Converts tuple (2, 4, 6) into array
x + y
```

```
array([3., 5., 7.])
```

```
4 * x
```

```
array([4., 4., 4.])
```

2.2.2 Inner product and norm

The **inner product** of vectors $x, y \in \mathbb{R}^n$ is defined as

$$x^\top y := \sum_{i=1}^n x_i y_i$$

Two vectors are called **orthogonal** if their inner product is zero.

The **norm** of a vector x represents its "length" (i.e., its distance from the zero vector) and is defined as

$$\|x\| := \sqrt{x^\top x} := \left(\sum_{i=1}^n x_i^2 \right)^{1/2}$$

The expression $\|x - y\|$ is thought of as the distance between x and y.

Continuing on from the previous example, the inner product and norm can be computed as follows

```
np.sum(x * y)           # Inner product of x and y, method 1
```

```
np.float64(12.0)
```

```
x @ y                   # Inner product of x and y, method 2 (preferred)
```

```
np.float64(12.0)
```

The `@` operator is preferred because it uses optimized BLAS libraries that implement fused multiply-add operations, providing better performance and numerical accuracy compared to the separate multiply and sum operations.

```
np.sqrt(np.sum(x**2))   # Norm of x, take one
```

```
np.float64(1.7320508075688772)
```

```
np.sqrt(x @ x)          # Norm of x, take two (preferred)
```

```
np.float64(1.7320508075688772)
```

```
np.linalg.norm(x)       # Norm of x, take three
```

```
np.float64(1.7320508075688772)
```

2.2.3 Span

Given a set of vectors $A := \{a_1, \ldots, a_k\}$ in $\mathbb{R}^n$, it's natural to think about the new vectors we can create by performing linear operations.

New vectors created in this manner are called **linear combinations** of A.

In particular, $y \in \mathbb{R}^n$ is a linear combination of $A := \{a_1, \ldots, a_k\}$ if

$$y = \beta_1 a_1 + \cdots + \beta_k a_k \text{ for some scalars } \beta_1, \ldots, \beta_k$$

In this context, the values $\beta_1, \ldots, \beta_k$ are called the **coefficients** of the linear combination.

The set of linear combinations of A is called the **span** of A.

The next figure shows the span of $A = \{a_1, a_2\}$ in $\mathbb{R}^3$.

The span is a two-dimensional plane passing through these two points and the origin.

```
ax = plt.figure().add_subplot(projection='3d')

x_min, x_max = -5, 5
y_min, y_max = -5, 5

α, β = 0.2, 0.1

ax.set(xlim=(x_min, x_max), ylim=(x_min, x_max), zlim=(x_min, x_max),
       xticks=(0,), yticks=(0,), zticks=(0,))

gs = 3
z = np.linspace(x_min, x_max, gs)
x = np.zeros(gs)
y = np.zeros(gs)
ax.plot(x, y, z, 'k-', lw=2, alpha=0.5)
ax.plot(z, x, y, 'k-', lw=2, alpha=0.5)
ax.plot(y, z, x, 'k-', lw=2, alpha=0.5)


# Fixed linear function, to generate a plane
def f(x, y):
    return α * x + β * y

# Vector locations, by coordinate
x_coords = np.array((3, 3))
y_coords = np.array((4, -4))
z = f(x_coords, y_coords)
for i in (0, 1):
    ax.text(x_coords[i], y_coords[i], z[i], f'$a_{i+1}$', fontsize=14)

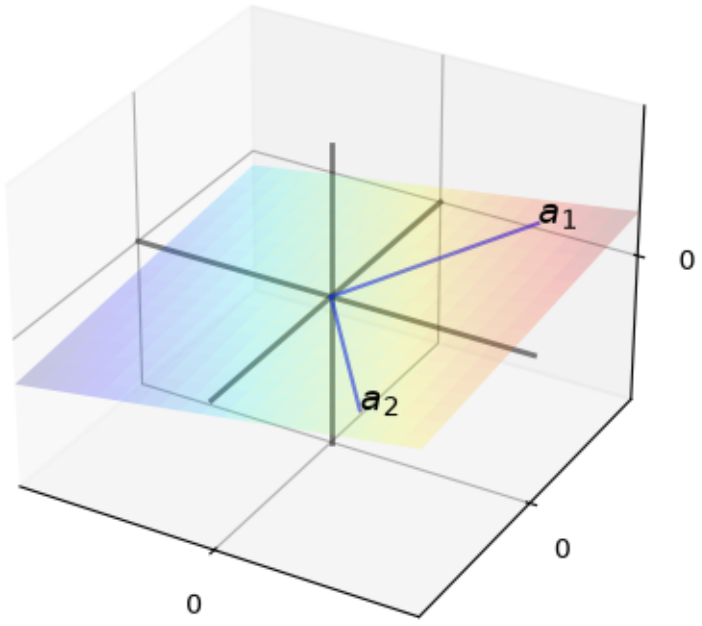
# Lines to vectors
for i in (0, 1):
    x = (0, x_coords[i])
    y = (0, y_coords[i])
    z = (0, f(x_coords[i], y_coords[i]))
    ax.plot(x, y, z, 'b-', lw=1.5, alpha=0.6)


# Draw the plane
grid_size = 20
xr2 = np.linspace(x_min, x_max, grid_size)
```

(continues on next page)

(continued from previous page)

```
yr2 = np.linspace(y_min, y_max, grid_size)
x2, y2 = np.meshgrid(xr2, yr2)
z2 = f(x2, y2)
ax.plot_surface(x2, y2, z2, rstride=1, cstride=1, cmap=cm.jet,
                linewidth=0, antialiased=True, alpha=0.2)
plt.show()
```



Examples

If A contains only one nonzero vector $a_1 \in \mathbb{R}^2$, then its span is just the scalar multiples of a_1, which is the unique line passing through both a_1 and the origin.

If $A = \{e_1, e_2, e_3\}$ consists of the **canonical basis vectors** of $\mathbb{R}^3$, that is

$$e_1 := \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad e_2 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad e_3 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

then the span of A is all of $\mathbb{R}^3$, because, for any $x = (x_1, x_2, x_3) \in \mathbb{R}^3$, we can write

$$x = x_1 e_1 + x_2 e_2 + x_3 e_3$$

Now consider $A_0 = \{e_1, e_2, e_1 + e_2\}$.

If $y = (y_1, y_2, y_3)$ is any linear combination of these vectors, then $y_3 = 0$ (check it).

Hence A_0 fails to span all of $\mathbb{R}^3$.

2.2.4 Linear independence

As we'll see, it's often desirable to find families of vectors with relatively large span, so that many vectors can be described by linear operators on a few vectors.

The condition we need for a set of vectors to have a large span is what's called linear independence.

In particular, a collection of vectors $A := \{a_1, \ldots, a_k\}$ in $\mathbb{R}^n$ is said to be

- **linearly dependent** if some strict subset of A has the same span as A.
- **linearly independent** if it is not linearly dependent.

Put differently, a set of vectors is linearly independent if no vector is redundant to the span and linearly dependent otherwise.

To illustrate the idea, recall *the figure* that showed the span of vectors $\{a_1, a_2\}$ in $\mathbb{R}^3$ as a plane through the origin.

If we take a third vector a_3 and form the set $\{a_1, a_2, a_3\}$, this set will be

- linearly dependent if a_3 lies in the plane
- linearly independent otherwise

As another illustration of the concept, since $\mathbb{R}^n$ can be spanned by n vectors (see the discussion of canonical basis vectors above), any collection of $m > n$ vectors in $\mathbb{R}^n$ must be linearly dependent.

The following statements are equivalent to linear independence of $A := \{a_1, \ldots, a_k\} \subset \mathbb{R}^n$

1. No vector in A can be formed as a linear combination of the other elements.
2. If $\beta_1 a_1 + \cdots + \beta_k a_k = 0$ for scalars $\beta_1, \ldots, \beta_k$, then $\beta_1 = \cdots = \beta_k = 0$.

(The zero in the first expression is the origin of $\mathbb{R}^n$)

2.2.5 Unique representations

Another nice thing about sets of linearly independent vectors is that each element in the span has a unique representation as a linear combination of these vectors.

In other words, if $A := \{a_1, \ldots, a_k\} \subset \mathbb{R}^n$ is linearly independent and

$$y = \beta_1 a_1 + \cdots + \beta_k a_k$$

then no other coefficient sequence $\gamma_1, \ldots, \gamma_k$ will produce the same vector y.

Indeed, if we also have $y = \gamma_1 a_1 + \cdots + \gamma_k a_k$, then

$$(\beta_1 - \gamma_1) a_1 + \cdots + (\beta_k - \gamma_k) a_k = 0$$

Linear independence now implies $\gamma_i = \beta_i$ for all i.

2.3 Matrices

Matrices are a neat way of organizing data for use in linear operations.

An $n \times k$ matrix is a rectangular array A of numbers with n rows and k columns:

$$A = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1k} \\ a_{21} & a_{22} & \cdots & a_{2k} \\ \vdots & \vdots & & \vdots \\ a_{n1} & a_{n2} & \cdots & a_{nk} \end{bmatrix}$$

Often, the numbers in the matrix represent coefficients in a system of linear equations, as discussed at the start of this lecture.

For obvious reasons, the matrix A is also called a vector if either $n = 1$ or $k = 1$.

In the former case, A is called a **row vector**, while in the latter it is called a **column vector**.

If $n = k$, then A is called **square**.

The matrix formed by replacing a_{ij} by a_{ji} for every i and j is called the **transpose** of A and denoted A' or $A^\top$.

If $A = A^\top$, then A is called **symmetric**.

For an $n \times n$ matrix A, the n entries $a_{11}, \ldots, a_{nn}$ form the **principal diagonal**.

A is called **diagonal** if the only nonzero entries are on the principal diagonal.

If, in addition to being diagonal, each element along the principal diagonal is equal to 1, then A is called the **identity matrix** and denoted by I.

2.3.1 Matrix operations

Just as was the case for vectors, a number of algebraic operations are defined for matrices.

Scalar multiplication and addition are immediate generalizations of the vector case:

$$\gamma A = \gamma \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nk} \end{bmatrix} := \begin{bmatrix} \gamma a_{11} & \cdots & \gamma a_{1k} \\ \vdots & \vdots & \vdots \\ \gamma a_{n1} & \cdots & \gamma a_{nk} \end{bmatrix}$$

and

$$A + B = \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nk} \end{bmatrix} + \begin{bmatrix} b_{11} & \cdots & b_{1k} \\ \vdots & \vdots & \vdots \\ b_{n1} & \cdots & b_{nk} \end{bmatrix} := \begin{bmatrix} a_{11} + b_{11} & \cdots & a_{1k} + b_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} + b_{n1} & \cdots & a_{nk} + b_{nk} \end{bmatrix}$$

In the latter case, the matrices must have the same shape in order for the definition to make sense.

We also have a convention for *multiplying* two matrices.

The rule for matrix multiplication generalizes the idea of inner products discussed above and is designed to make multiplication play well with basic linear operations.

If A and B are two matrices, then their product AB is formed by taking as its i, j-th element the inner product of the i-th row of A and the j-th column of B.

There are many tutorials to help you visualize this operation, such as this one, or the discussion on the Wikipedia page.

If A is $n \times k$ and B is $j \times m$, then to multiply A and B we require $k = j$, and the resulting matrix AB is $n \times m$.

As perhaps the most important special case, consider multiplying $n \times k$ matrix A and $k \times 1$ column vector x.

According to the preceding rule, this gives us an $n \times 1$ column vector

$$Ax = \begin{bmatrix} a_{11} & \cdots & a_{1k} \\ \vdots & \vdots & \vdots \\ a_{n1} & \cdots & a_{nk} \end{bmatrix} \begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix} := \begin{bmatrix} a_{11}x_1 + \cdots + a_{1k}x_k \\ \vdots \\ a_{n1}x_1 + \cdots + a_{nk}x_k \end{bmatrix} \tag{2.2}$$

Note

AB and BA are not generally the same thing.

Another important special case is the identity matrix.

You should check that if A is $n \times k$ and I is the $k \times k$ identity matrix, then $AI = A$.

If I is the $n \times n$ identity matrix, then $IA = A$.

2.3.2 Matrices in NumPy

NumPy arrays are also used as matrices, and have fast, efficient functions and methods for all the standard matrix operations[1].

You can create them manually from tuples of tuples (or lists of lists) as follows

```
A = ((1, 2),
     (3, 4))

type(A)
```

```
tuple
```

```
A = np.array(A)

type(A)
```

```
numpy.ndarray
```

```
A.shape
```

```
(2, 2)
```

The `shape` attribute is a tuple giving the number of rows and columns — see here for more discussion.

To get the transpose of `A`, use `A.transpose()` or, more simply, `A.T`.

There are many convenient functions for creating common matrices (matrices of zeros, ones, etc.) — see here.

Since operations are performed elementwise by default, scalar multiplication and addition have very natural syntax

```
A = np.identity(3)
B = np.ones((3, 3))
2 * A
```

[1] Although there is a specialized matrix data type defined in NumPy, it's more standard to work with ordinary NumPy arrays. See this discussion.

```
array([[2., 0., 0.],
       [0., 2., 0.],
       [0., 0., 2.]])
```

```
A + B
```

```
array([[2., 1., 1.],
       [1., 2., 1.],
       [1., 1., 2.]])
```

To multiply matrices we use the `@` symbol.

In particular, `A @ B` is matrix multiplication, whereas `A * B` is element-by-element multiplication.

See here for more discussion.

2.3.3 Matrices as maps

Each $n \times k$ matrix A can be identified with a function $f(x) = Ax$ that maps $x \in \mathbb{R}^k$ into $y = Ax \in \mathbb{R}^n$.

These kinds of functions have a special property: they are **linear**.

A function $f \colon \mathbb{R}^k \to \mathbb{R}^n$ is called **linear** if, for all $x, y \in \mathbb{R}^k$ and all scalars α, β, we have

$$f(\alpha x + \beta y) = \alpha f(x) + \beta f(y)$$

You can check that this holds for the function $f(x) = Ax + b$ when b is the zero vector and fails when b is nonzero.

In fact, it's known that f is linear if and *only if* there exists a matrix A such that $f(x) = Ax$ for all x.

2.4 Solving systems of equations

Recall again the system of equations (2.1).

If we compare (2.1) and (2.2), we see that (2.1) can now be written more conveniently as

$$y = Ax \tag{2.3}$$

The problem we face is to determine a vector $x \in \mathbb{R}^k$ that solves (2.3), taking y and A as given.

This is a special case of a more general problem: Find an x such that $y = f(x)$.

Given an arbitrary function f and a y, is there always an x such that $y = f(x)$?

If so, is it always unique?

The answer to both these questions is negative, as the next figure shows

```
def f(x):
    return 0.6 * np.cos(4 * x) + 1.4


xmin, xmax = -1, 1
x = np.linspace(xmin, xmax, 160)
y = f(x)
ya, yb = np.min(y), np.max(y)
```

(continues on next page)

(continued from previous page)

```
# Extra height keeps the labels clear of the axes in both panels.
fig, axes = plt.subplots(2, 1, figsize=(10, 10))

for ax in axes:
    # Draw the axes through the origin
    ax.axhline(0, color='black', lw=0.8)
    ax.axvline(0, color='black', lw=0.8)

    ax.set(ylim=(-0.6, 3.2), xlim=(xmin, xmax),
           yticks=(), xticks=())

    ax.plot(x, y, 'k-', lw=2, label='$f$')
    ax.fill_between(x, ya, yb, facecolor='blue', alpha=0.05)
    ax.vlines([0], ya, yb, lw=3, color='blue', label='range of $f$')
    ax.text(0.04, -0.3, '$0$', fontsize=16)

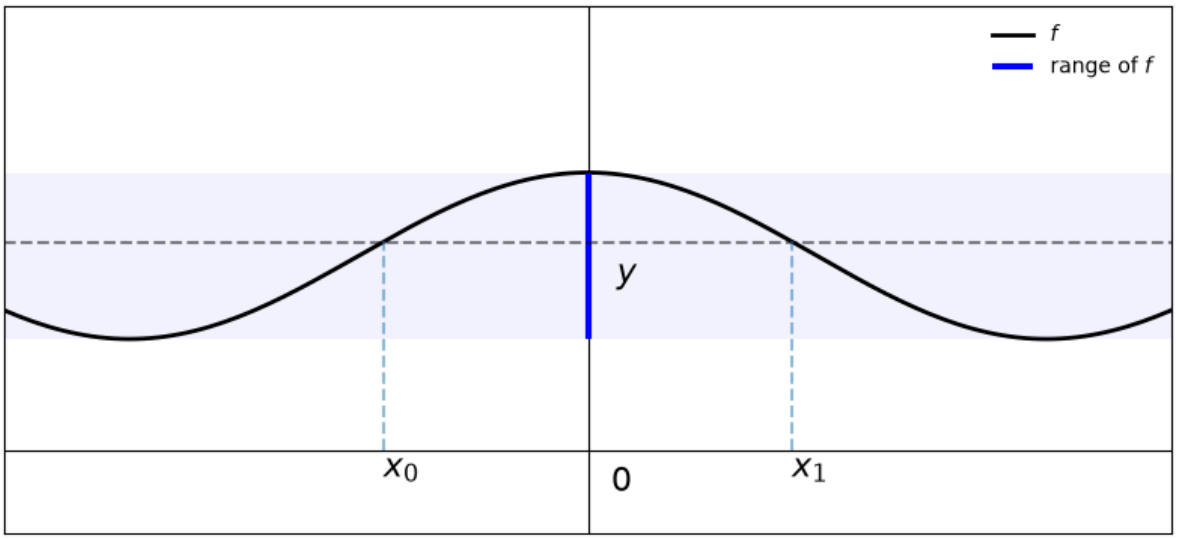
ax = axes[0]

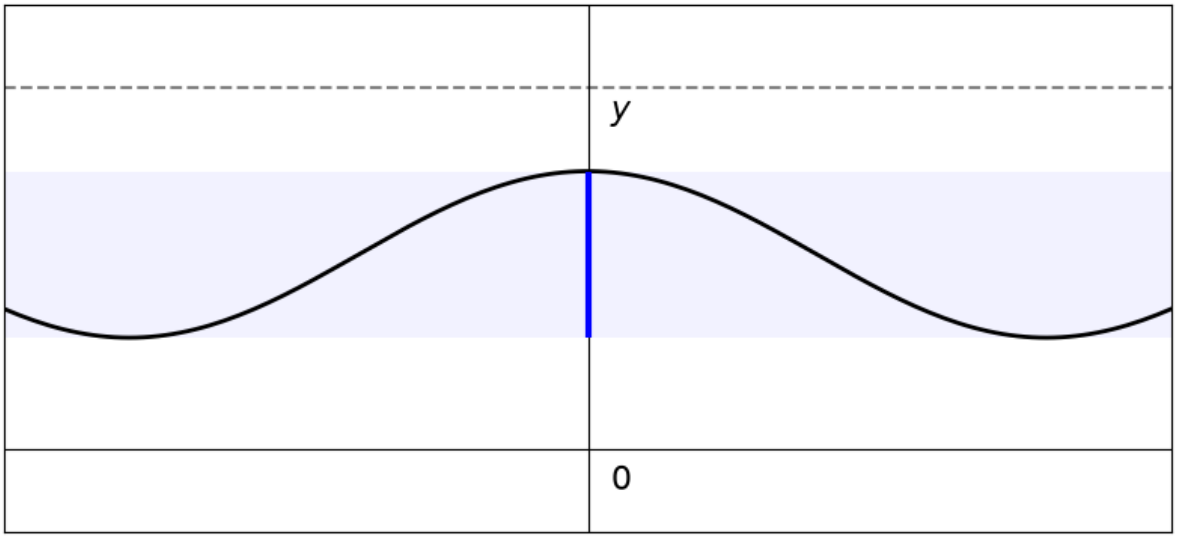
ax.legend(loc='upper right', frameon=False)
ybar = 1.5
ax.plot(x, x * 0 + ybar, 'k--', alpha=0.5)
ax.text(0.05, 0.8 * ybar, '$y$', fontsize=16)
for i, z in enumerate((-0.35, 0.35)):
    ax.vlines(z, 0, f(z), linestyle='--', alpha=0.5)
    ax.text(z, -0.2, f'$x_{i}$', fontsize=16)

ax = axes[1]

ybar = 2.6
ax.plot(x, x * 0 + ybar, 'k--', alpha=0.5)
ax.text(0.04, 0.91 * ybar, '$y$', fontsize=16)

plt.show()
```





In the first plot, there are multiple solutions, as the function is not one-to-one, while in the second there are no solutions, since y lies outside the range of f.

Can we impose conditions on A in (2.3) that rule out these problems?

In this context, the most important thing to recognize about the expression Ax is that it corresponds to a linear combination of the columns of A.

In particular, if $a_1, \ldots, a_k$ are the columns of A, then

$$Ax = x_1 a_1 + \cdots + x_k a_k$$

Hence the range of $f(x) = Ax$ is exactly the span of the columns of A.

We want the range to be large so that it contains arbitrary y.

As you might recall, the condition that we want for the span to be large is *linear independence*.

A happy fact is that linear independence of the columns of A also gives us uniqueness.

Indeed, it follows from our *earlier discussion* that if $\{a_1, \ldots, a_k\}$ are linearly independent and $y = Ax = x_1 a_1 + \cdots + x_k a_k$, then no $z \neq x$ satisfies $y = Az$.

2.4.1 The square matrix case

Let's discuss some more details, starting with the case where A is $n \times n$.

This is the familiar case where the number of unknowns equals the number of equations.

For arbitrary $y \in \mathbb{R}^n$, we hope to find a unique $x \in \mathbb{R}^n$ such that $y = Ax$.

In view of the observations immediately above, if the columns of A are linearly independent, then their span, and hence the range of $f(x) = Ax$, is all of $\mathbb{R}^n$.

Hence there always exists an x such that $y = Ax$.

Moreover, the solution is unique.

In particular, the following are equivalent

1. The columns of A are linearly independent.
2. For any $y \in \mathbb{R}^n$, the equation $y = Ax$ has a unique solution.

The property of having linearly independent columns is sometimes expressed as having **full column rank**.

Inverse matrices

Can we give some sort of expression for the solution?

If y and A are scalar with $A \neq 0$, then the solution is $x = A^{-1} y$.

A similar expression is available in the matrix case.

In particular, if square matrix A has full column rank, then it possesses a multiplicative **inverse matrix** A^{-1}, with the property that $AA^{-1} = A^{-1}A = I$.

As a consequence, if we pre-multiply both sides of $y = Ax$ by A^{-1}, we get $x = A^{-1} y$.

This is the solution that we're looking for.

Determinants

Another quick comment about square matrices is that to every such matrix we assign a unique number called the **determinant** of the matrix — you can find the expression for it here.

If the determinant of A is not zero, then we say that A is **nonsingular**.

Perhaps the most important fact about determinants is that A is nonsingular if and only if A is of full column rank.

This gives us a useful one-number summary of whether or not a square matrix can be inverted.

2.4.2 More rows than columns

This is the $n \times k$ case with $n > k$.

This case is very important in many settings, not least in the setting of linear regression (where n is the number of observations, and k is the number of explanatory variables).

Given arbitrary $y \in \mathbb{R}^n$, we seek an $x \in \mathbb{R}^k$ such that $y = Ax$.

In this setting, the existence of a solution is highly unlikely.

Without much loss of generality, let's go over the intuition focusing on the case where the columns of A are linearly independent.

It follows that the span of the columns of A is a k-dimensional subspace of $\mathbb{R}^n$.

This span is very "unlikely" to contain arbitrary $y \in \mathbb{R}^n$.

To see why, recall the *figure above*, where $k = 2$ and $n = 3$.

Imagine an arbitrarily chosen $y \in \mathbb{R}^3$, located somewhere in that three-dimensional space.

What's the likelihood that y lies in the span of $\{a_1, a_2\}$ (i.e., the two dimensional plane through these points)?

In a sense, it must be very small, since this plane has zero "thickness".

As a result, in the $n > k$ case we usually give up on existence.

However, we can still seek the best approximation, for example, an x that makes the distance $\|y - Ax\|$ as small as possible.

To solve this problem, one can use either calculus or the theory of orthogonal projections.

The solution is known to be $\hat{x} = (A^\top A)^{-1} A^\top y$ — see for example chapter 3 of these notes.

2.4.3 More columns than rows

This is the $n \times k$ case with $n < k$, so there are fewer equations than unknowns.

In this case there are either no solutions or infinitely many — in other words, uniqueness never holds.

For example, consider the case where $k = 3$ and $n = 2$.

Thus, the columns of A consist of 3 vectors in $\mathbb{R}^2$.

This set can never be linearly independent, since it is possible to find two vectors that span $\mathbb{R}^2$.

(For example, use the canonical basis vectors)

It follows that one column is a linear combination of the other two.

For example, let's say that $a_1 = \alpha a_2 + \beta a_3$.

Then if $y = Ax = x_1 a_1 + x_2 a_2 + x_3 a_3$, we can also write

$$y = x_1(\alpha a_2 + \beta a_3) + x_2 a_2 + x_3 a_3 = (x_1\alpha + x_2)a_2 + (x_1\beta + x_3)a_3$$

In other words, uniqueness fails.

2.4.4 Linear equations with SciPy

Here's an illustration of how to solve linear equations with SciPy's `linalg` submodule.

All of these routines are Python front ends to time-tested and highly optimized FORTRAN code

```
A = ((1, 2), (3, 4))
A = np.array(A)
y = np.ones((2, 1))  # Column vector
det(A)  # Check that A is nonsingular, and hence invertible
```

```
np.float64(-2.0)
```

```
A_inv = inv(A)  # Compute the inverse
A_inv
```

```
array([[-2. ,  1. ],
       [ 1.5, -0.5]])
```

```
x = A_inv @ y  # Solution
A @ x          # Should equal y
```

```
array([[1.],
       [1.]])
```

```
solve(A, y)  # Produces the same solution
```

```
array([[-1.],
       [ 1.]])
```

Observe how we can solve for $x = A^{-1}y$ either via `inv(A) @ y`, or using `solve(A, y)`.

To solve $Ax = y$, we generally prefer `solve(A, y)` because it avoids the unnecessary step of computing the inverse of A.

To obtain the least-squares solution $\hat{x} = (A^\top A)^{-1} A^\top y$, use `scipy.linalg.lstsq(A, y)`.

2.5 Eigenvalues and eigenvectors

Let A be an $n \times n$ square matrix.

If λ is scalar and v is a non-zero vector in $\mathbb{R}^n$ such that

$$Av = \lambda v$$

then we say that λ is an **eigenvalue** of A, and v is an **eigenvector**.

Thus, an eigenvector of A is a vector such that when the map $f(x) = Ax$ is applied, v is merely scaled.

The next figure shows two eigenvectors (blue arrows) and their images under A (red arrows).

As expected, the image Av of each v is just a scaled version of the original

```
A = ((1, 2),
     (2, 1))
A = np.array(A)
evals, evecs = eig(A)
evecs = evecs[:, 0], evecs[:, 1]

fig, ax = plt.subplots()
# Draw the axes through the origin
ax.axhline(0, color='black', lw=0.8)
ax.axvline(0, color='black', lw=0.8)
ax.grid(alpha=0.4)

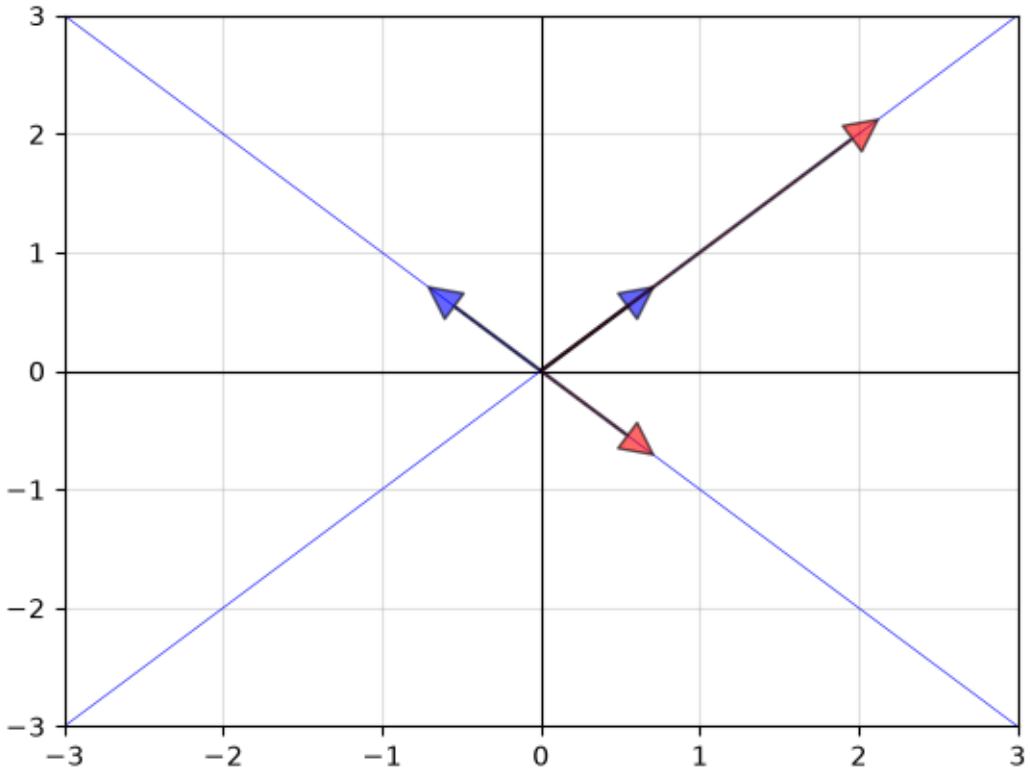
xmin, xmax = -3, 3
ymin, ymax = -3, 3
ax.set(xlim=(xmin, xmax), ylim=(ymin, ymax))

# Plot each eigenvector
for v in evecs:
    ax.annotate('', xy=v, xytext=(0, 0),
                arrowprops=dict(facecolor='blue',
                shrink=0,
                alpha=0.6,
                width=0.5))

# Plot the image of each eigenvector
for v in evecs:
    v = A @ v
    ax.annotate('', xy=v, xytext=(0, 0),
                arrowprops=dict(facecolor='red',
                shrink=0,
                alpha=0.6,
                width=0.5))

# Plot the lines they run through
x = np.linspace(xmin, xmax, 3)
for v in evecs:
    a = v[1] / v[0]
    ax.plot(x, a * x, 'b-', lw=0.4)

plt.show()
```



The eigenvalue equation is equivalent to $(A - \lambda I)v = 0$, and this has a nonzero solution v only when the columns of $A - \lambda I$ are linearly dependent.

This in turn is equivalent to stating that the determinant is zero.

Hence to find all eigenvalues, we can look for λ such that the determinant of $A - \lambda I$ is zero.

This problem can be expressed as one of solving for the roots of a polynomial in λ of degree n.

This in turn implies the existence of n solutions in the complex plane, although some might be repeated.

Some nice facts about the eigenvalues of a square matrix A are as follows

1. The determinant of A equals the product of the eigenvalues.
2. The trace of A (the sum of the elements on the principal diagonal) equals the sum of the eigenvalues.
3. If A is symmetric, then all of its eigenvalues are real.
4. If A is invertible and $\lambda_1, \ldots, \lambda_n$ are its eigenvalues, then the eigenvalues of A^{-1} are $1/\lambda_1, \ldots, 1/\lambda_n$.

A corollary of the first statement is that a matrix is invertible if and only if all its eigenvalues are nonzero.

Using SciPy, we can solve for the eigenvalues and eigenvectors of a matrix as follows

```
A = ((1, 2),
     (2, 1))

A = np.array(A)
evals, evecs = eig(A)
evals
```

```
array([ 3.+0.j, -1.+0.j])
```

```
evecs
```

```
array([[ 0.70710678, -0.70710678],
       [ 0.70710678,  0.70710678]])
```

Note that the *columns* of `evecs` are the eigenvectors.

Since any nonzero scalar multiple of an eigenvector is an eigenvector with the same eigenvalue (check it), the eig routine normalizes the length of each eigenvector to one.

2.5.1 Generalized eigenvalues

It is sometimes useful to consider the **generalized eigenvalue problem**, which, for given matrices A and B, seeks generalized eigenvalues λ and eigenvectors v such that

$$Av = \lambda Bv$$

This can be solved in SciPy via `scipy.linalg.eig(A, B)`.

Of course, if B is square and invertible, then we can treat the generalized eigenvalue problem as an ordinary eigenvalue problem $B^{-1}Av = \lambda v$, but this is not always the case.

2.6 Further topics

We round out our discussion by briefly mentioning several other important topics.

2.6.1 Series expansions

Recall the usual summation formula for a geometric progression, which states that if $|a| < 1$, then $\sum_{k=0}^{\infty} a^k = (1-a)^{-1}$.

A generalization of this idea exists in the matrix setting.

Matrix norms

Let A be a square matrix, and let

$$\|A\| := \max_{\|x\|=1} \|Ax\|$$

The norms on the right-hand side are ordinary vector norms, while the norm on the left-hand side is a **matrix norm** — in this case, the so-called **spectral norm**.

For example, for a square matrix S, the condition $\|S\| < 1$ means that S is **contractive**, in the sense that it pulls all vectors towards the origin[2].

[2] Suppose that $\|S\| < 1$.
Take any nonzero vector x, and let $r := \|x\|$.
We have $\|Sx\| = r\|S(x/r)\| \leq r\|S\| < r = \|x\|$.
Hence every point is pulled towards the origin.

Neumann's theorem

Let A be a square matrix and let $A^k := AA^{k-1}$ with $A^1 := A$.

In other words, A^k is the k-th power of A.

Neumann's theorem states the following: If $\|A^k\| < 1$ for some $k \in \mathbb{N}$, then $I - A$ is invertible, and

$$(I - A)^{-1} = \sum_{k=0}^{\infty} A^k \tag{2.4}$$

Spectral radius

A result known as Gelfand's formula tells us that, for any square matrix A,

$$\rho(A) = \lim_{k \to \infty} \|A^k\|^{1/k}$$

Here $\rho(A)$ is the **spectral radius**, defined as $\max_i |\lambda_i|$, where $\{\lambda_i\}_i$ is the set of eigenvalues of A.

As a consequence of Gelfand's formula, if all eigenvalues are strictly less than one in modulus, there exists a k with $\|A^k\| < 1$.

In which case (2.4) is valid.

2.6.2 Positive definite matrices

Let A be a symmetric $n \times n$ matrix.

We say that A is

1. **positive definite** if $x^\top Ax > 0$ for every $x \in \mathbb{R}^n \setminus \{0\}$
2. **positive semi-definite** or **nonnegative definite** if $x^\top Ax \geq 0$ for every $x \in \mathbb{R}^n$

Analogous definitions exist for negative definite and negative semi-definite matrices.

It is notable that if A is positive definite, then all of its eigenvalues are strictly positive, and hence A is invertible (with positive definite inverse).

2.6.3 Differentiating linear and quadratic forms

The following formulas are useful in many economic contexts.

Let

- z, x and a all be $n \times 1$ vectors
- A be an $n \times n$ matrix
- B be an $m \times n$ matrix and y be an $m \times 1$ vector

Then

1. $\frac{\partial a^\top x}{\partial x} = a$
2. $\frac{\partial Ax}{\partial x} = A^\top$
3. $\frac{\partial x^\top Ax}{\partial x} = (A + A^\top)x$
4. $\frac{\partial y^\top Bz}{\partial y} = Bz$

5. $\frac{\partial y^\top B z}{\partial B} = y z^\top$

Exercise 2.7.1 below asks you to apply these formulas.

2.6.4 Further reading

The documentation of the `scipy.linalg` submodule can be found here.

Chapters 2 and 3 of A Primer in Econometric Theory contain a discussion of linear algebra along the same lines as above, with solved exercises.

If you don't mind a slightly abstract approach, a nice intermediate-level text on linear algebra is [Jänich, 1994].

2.7 Exercises

Exercise 2.7.1

Let x be a given $n \times 1$ vector and consider the problem

$$v(x) = \max_{y,u} \{-y^\top P y - u^\top Q u\}$$

subject to the linear constraint

$$y = Ax + Bu$$

Here

- P is an $n \times n$ matrix and Q is an $m \times m$ matrix
- A is an $n \times n$ matrix and B is an $n \times m$ matrix
- both P and Q are symmetric and positive semidefinite
- $H := Q + B^\top P B$ is positive definite

(What must the dimensions of y and u be to make this a well-posed problem?)

One way to solve the problem is to form the Lagrangian

$$\mathcal{L} = -y^\top P y - u^\top Q u + \lambda^\top [Ax + Bu - y]$$

where λ is an $n \times 1$ vector of Lagrange multipliers.

Try applying the formulas given above for differentiating quadratic and linear forms to obtain the first-order conditions for maximizing $\mathcal{L}$ with respect to y, u and minimizing it with respect to λ.

Show that these conditions imply that

1. $\lambda = -2Py$.
2. The optimizing choice of u satisfies $u = -(Q + B^\top P B)^{-1} B^\top P A x$.
3. The function v satisfies $v(x) = -x^\top \tilde{P} x$ where $\tilde{P} = A^\top P A - A^\top P B (Q + B^\top P B)^{-1} B^\top P A$.

As we will see, in economic contexts Lagrange multipliers often are shadow prices.

Note

If we don't care about the Lagrange multipliers, we can substitute the constraint into the objective function, and then just maximize $-(Ax + Bu)^\top P(Ax + Bu) - u^\top Qu$ with respect to u.
You can verify that this leads to the same maximizer.

Solution

We have an optimization problem:

$$v(x) = \max_{y,u}\{-y^\top Py - u^\top Qu\}$$

s.t.

$$y = Ax + Bu$$

with primitives

- P is a symmetric and positive semidefinite $n \times n$ matrix
- Q is a symmetric and positive semidefinite $m \times m$ matrix
- A is an $n \times n$ matrix
- B is an $n \times m$ matrix

Assume that $H := Q + B^\top PB$ is positive definite.

The associated Lagrangian is:

$$L = -y^\top Py - u^\top Qu + \lambda^\top [Ax + Bu - y]$$

Step 1.

Differentiating Lagrangian equation w.r.t y and setting its derivative equal to zero yields

$$\frac{\partial L}{\partial y} = -(P + P^\top)y - \lambda = -2Py - \lambda = 0\,,$$

since P is symmetric.

Accordingly, the first-order condition for maximizing L w.r.t. y implies

$$\lambda = -2Py$$

Step 2.

Differentiating Lagrangian equation w.r.t. u and setting its derivative equal to zero yields

$$\frac{\partial L}{\partial u} = -(Q + Q^\top)u + B^\top \lambda = -2Qu + B^\top \lambda = 0$$

Substituting $\lambda = -2Py$ gives

$$Qu + B^\top Py = 0$$

Substituting the linear constraint $y = Ax + Bu$ into the above equation gives

$$Qu + B^\top P(Ax + Bu) = 0$$

$$(Q + B^\top PB)u + B^\top PAx = 0$$

which is the first-order condition for maximizing L w.r.t. u.

Thus, the optimal choice of u must satisfy

$$u = -(Q + B^\top PB)^{-1} B^\top PAx \,,$$

which follows from the definition of the first-order conditions for Lagrangian equation.

Step 3.

Rewriting our problem by substituting the constraint into the objective function, we get

$$v(x) = \max_u \{-(Ax + Bu)^\top P(Ax + Bu) - u^\top Qu\}$$

Since we know the optimal choice of u satisfies $u = -(Q + B^\top PB)^{-1} B^\top PAx$, then

$$v(x) = -(Ax + Bu)^\top P(Ax + Bu) - u^\top Qu \quad \text{with} \quad u = -(Q + B^\top PB)^{-1} B^\top PAx$$

To evaluate the function

$$\begin{aligned} v(x) &= -(Ax + Bu)^\top P(Ax + Bu) - u^\top Qu \\ &= -(x^\top A^\top + u^\top B^\top)P(Ax + Bu) - u^\top Qu \\ &= -x^\top A^\top PAx - u^\top B^\top PAx - x^\top A^\top PBu - u^\top B^\top PBu - u^\top Qu \\ &= -x^\top A^\top PAx - 2u^\top B^\top PAx - u^\top (Q + B^\top PB)u \end{aligned}$$

For simplicity, set $S := (Q + B^\top PB)^{-1} B^\top PA$, so that $u = -Sx$.

Regarding the second term $-2u^\top B^\top PAx$,

$$\begin{aligned} -2u^\top B^\top PAx &= +2x^\top S^\top B^\top PAx \\ &= 2x^\top A^\top PB(Q + B^\top PB)^{-1} B^\top PAx \end{aligned}$$

Notice that the term $(Q + B^\top PB)^{-1}$ is symmetric as both P and Q are symmetric.

Regarding the third term $-u^\top (Q + B^\top PB)u$,

$$\begin{aligned} -u^\top (Q + B^\top PB)u &= -x^\top S^\top (Q + B^\top PB)Sx \\ &= -x^\top A^\top PB(Q + B^\top PB)^{-1} B^\top PAx \end{aligned}$$

Hence, the summation of second and third terms is $x^\top A^\top PB(Q + B^\top PB)^{-1} B^\top PAx$.

This implies that

$$\begin{aligned} v(x) &= -x^\top A^\top PAx - 2u^\top B^\top PAx - u^\top (Q + B^\top PB)u \\ &= -x^\top A^\top PAx + x^\top A^\top PB(Q + B^\top PB)^{-1} B^\top PAx \\ &= -x^\top [A^\top PA - A^\top PB(Q + B^\top PB)^{-1} B^\top PA]x \end{aligned}$$

Therefore, the solution to the optimization problem $v(x) = -x^\top \tilde{P} x$ follows the above result by denoting $\tilde{P} := A^\top PA - A^\top PB(Q + B^\top PB)^{-1} B^\top PA$

CHAPTER

THREE

QR DECOMPOSITION

3.1 Overview

This lecture describes the QR decomposition and how it relates to

- Orthogonal projection and least squares
- A Gram-Schmidt process
- Eigenvalues and eigenvectors

We'll write some Python code to help consolidate our understandings.

3.2 Matrix factorization

In this lecture, we consider a real matrix $A \in \mathbb{R}^{n \times m}$ with $m \geq n$.

A QR decomposition (also called a QR factorization) of A takes the form

$$A = QR, \qquad Q \in \mathbb{R}^{n \times n}, \qquad R \in \mathbb{R}^{n \times m}, \qquad Q^\top Q = I_n.$$

The matrix R has zeros below its main diagonal.

Hence, R is upper triangular when $m = n$ and upper trapezoidal when $m > n$.

We'll use a *Gram-Schmidt process* to compute a QR decomposition.

Because doing so is so educational, we'll write our own Python code to do the job.

3.3 Gram-Schmidt process

We'll start with a *square* matrix A.

If a square matrix A is nonsingular, then a QR factorization is unique up to signs.

We'll deal with a rectangular matrix A later.

Actually, our algorithm will work with a rectangular A that is not square.

3.3.1 Gram-Schmidt process for square A

Here we apply a Gram-Schmidt process to the *columns* of matrix A.

In particular, let

$$A = \begin{bmatrix} a_1 & a_2 & \cdots & a_n \end{bmatrix}$$

Let $||\cdot||$ denote the L2 norm.

The Gram-Schmidt algorithm repeatedly combines the following two steps in a particular order

- **normalize** a vector to have unit norm
- **orthogonalize** the next vector

To begin, we set $u_1 = a_1$ and then *normalize*:

$$u_1 = a_1, \quad e_1 = \frac{u_1}{||u_1||}$$

We *orthogonalize* first to compute u_2 and then *normalize* to create e_2:

$$u_2 = a_2 - (a_2 \cdot e_1)e_1, \quad e_2 = \frac{u_2}{||u_2||}$$

We invite the reader to verify that e_1 is orthogonal to e_2 by checking that $e_1 \cdot e_2 = 0$.

The Gram-Schmidt procedure continues iterating.

Thus, for $k = 2, \ldots, n-1$ we construct

$$u_{k+1} = a_{k+1} - (a_{k+1} \cdot e_1)e_1 - \cdots - (a_{k+1} \cdot e_k)e_k, \quad e_{k+1} = \frac{u_{k+1}}{||u_{k+1}||}$$

Here $(a_j \cdot e_i)$ can be interpreted as the linear least squares *regression coefficient* of a_j on e_i

- it is the inner product of a_j and e_i divided by the inner product of e_i where $e_i \cdot e_i = 1$, as *normalization* has assured us.
- this regression coefficient has an interpretation as being a *covariance* divided by a *variance*

It can be verified that

$$A = \begin{bmatrix} a_1 & a_2 & \cdots & a_n \end{bmatrix} = \begin{bmatrix} e_1 & e_2 & \cdots & e_n \end{bmatrix} \begin{bmatrix} a_1 \cdot e_1 & a_2 \cdot e_1 & \cdots & a_n \cdot e_1 \\ 0 & a_2 \cdot e_2 & \cdots & a_n \cdot e_2 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_n \cdot e_n \end{bmatrix}$$

Thus, we have constructed the decomposition

$$A = QR$$

where

$$Q = \begin{bmatrix} e_1 & e_2 & \cdots & e_n \end{bmatrix}$$

and

$$R = \begin{bmatrix} a_1 \cdot e_1 & a_2 \cdot e_1 & \cdots & a_n \cdot e_1 \\ 0 & a_2 \cdot e_2 & \cdots & a_n \cdot e_2 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_n \cdot e_n \end{bmatrix}$$

3.3.2 A not square

Now suppose that A is an $n \times m$ matrix where $m > n$.

Then a QR decomposition is

$$A = \begin{bmatrix} a_1 & a_2 & \cdots & a_m \end{bmatrix} = \begin{bmatrix} e_1 & e_2 & \cdots & e_n \end{bmatrix} \begin{bmatrix} a_1 \cdot e_1 & a_2 \cdot e_1 & \cdots & a_n \cdot e_1 & a_{n+1} \cdot e_1 & \cdots & a_m \cdot e_1 \\ 0 & a_2 \cdot e_2 & \cdots & a_n \cdot e_2 & a_{n+1} \cdot e_2 & \cdots & a_m \cdot e_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & a_n \cdot e_n & a_{n+1} \cdot e_n & \cdots & a_m \cdot e_n \end{bmatrix}$$

Here, R is an $n \times m$ upper trapezoidal matrix.

This implies that

$$\begin{aligned} a_1 &= (a_1 \cdot e_1)e_1 \\ a_2 &= (a_2 \cdot e_1)e_1 + (a_2 \cdot e_2)e_2 \\ &\vdots \quad \vdots \\ a_n &= (a_n \cdot e_1)e_1 + (a_n \cdot e_2)e_2 + \cdots + (a_n \cdot e_n)e_n \\ a_{n+1} &= (a_{n+1} \cdot e_1)e_1 + (a_{n+1} \cdot e_2)e_2 + \cdots + (a_{n+1} \cdot e_n)e_n \\ &\vdots \quad \vdots \\ a_m &= (a_m \cdot e_1)e_1 + (a_m \cdot e_2)e_2 + \cdots + (a_m \cdot e_n)e_n \end{aligned}$$

3.4 Some code

Now let's write some homemade Python code to implement a QR decomposition by deploying the Gram-Schmidt process described above.

```
import numpy as np
from scipy.linalg import qr

rng = np.random.default_rng(1234)
```

The `qr_decomposition` function below illustrates the square case and its extension to wide $n \times m$ matrices with $m \geq n$, assuming that the first n columns are linearly independent.

```
def qr_decomposition(A):
    n, m = A.shape # get the shape of A

    Q = np.empty((n, n)) # initialize matrix Q
    u = np.empty((n, n)) # initialize matrix u

    u[:, 0] = A[:, 0]
    Q[:, 0] = u[:, 0] / np.linalg.norm(u[:, 0])

    for i in range(1, n):

        u[:, i] = A[:, i]
        for j in range(i):
            u[:, i] -= (A[:, i] @ Q[:, j]) * Q[:, j] # get each u vector

        Q[:, i] = u[:, i] / np.linalg.norm(u[:, i]) # compute each e vector
```

(continues on next page)

(continued from previous page)

```
    R = np.zeros((n, m))
    for i in range(n):
        for j in range(i, m):
            R[i, j] = A[:, j] @ Q[:, i]

    return Q, R
```

The preceding code is fine but can benefit from some further housekeeping.

We want to do this because later in this notebook we want to compare results from using our homemade code above with the code for a QR that the Python SciPy package delivers.

There can be sign differences between the Q and R matrices produced by different numerical algorithms.

All of these are valid QR decompositions because of how the sign differences cancel out when we compute QR.

To make the results from our homemade function and the QR module in SciPy comparable, we'll require the diagonal entries of R to be nonnegative.

We do this by adjusting the signs of the columns in Q and the corresponding rows in R appropriately.

To accomplish this we'll define a pair of functions.

```
def diag_sign(A):
    """Compute a diagonal sign matrix, mapping zero signs to one."""

    signs = np.where(np.diag(A) < 0, -1.0, 1.0)
    D = np.diag(signs)

    return D

def adjust_sign(Q, R):
    """
    Adjust the signs of the columns of Q and rows of R
    so that R has a nonnegative diagonal.
    """

    D = diag_sign(R)

    Q[:, :] = Q @ D
    R[:, :] = D @ R

    return Q, R
```

3.5 Example

Now let's do an example.

```
A = np.array([[1.0, 1.0, 0.0], [1.0, 0.0, 1.0], [0.0, 1.0, 1.0]])
# A = np.array([[1.0, 0.5, 0.2], [0.5, 0.5, 1.0], [0.0, 1.0, 1.0]])
# A = np.array([[1.0, 0.5, 0.2], [0.5, 0.5, 1.0]])

A
```

```
array([[1., 1., 0.],
       [1., 0., 1.],
       [0., 1., 1.]])
```

```
Q, R = adjust_sign(*qr_decomposition(A))
```

```
Q
```

```
array([[ 0.70710678,  0.40824829, -0.57735027],
       [ 0.70710678, -0.40824829,  0.57735027],
       [ 0.        ,  0.81649658,  0.57735027]])
```

```
R
```

```
array([[1.41421356, 0.70710678, 0.70710678],
       [0.        , 1.22474487, 0.40824829],
       [0.        , 0.        , 1.15470054]])
```

Let's compare outcomes with what the SciPy package produces

```
Q_scipy, R_scipy = adjust_sign(*qr(A))
```

```
print('Our Q: \n', Q)
print('\n')
print('SciPy Q: \n', Q_scipy)
```

```
Our Q:
 [[ 0.70710678  0.40824829 -0.57735027]
 [ 0.70710678 -0.40824829  0.57735027]
 [ 0.          0.81649658  0.57735027]]


SciPy Q:
 [[ 0.70710678  0.40824829 -0.57735027]
 [ 0.70710678 -0.40824829  0.57735027]
 [ 0.          0.81649658  0.57735027]]
```

```
print('Our R: \n', R)
print('\n')
print('SciPy R: \n', R_scipy)
```

```
Our R:
 [[1.41421356 0.70710678 0.70710678]
 [0.         1.22474487 0.40824829]
 [0.         0.         1.15470054]]


SciPy R:
 [[1.41421356 0.70710678 0.70710678]
 [0.         1.22474487 0.40824829]
 [0.         0.         1.15470054]]
```

The above outcomes give us the good news that our homemade function agrees with what SciPy produces.

Now let's do a QR decomposition for a rectangular matrix A that is $n \times m$ with $m > n$.

```
A = np.array([[1, 3, 4], [2, 0, 9]])
```

```
Q, R = adjust_sign(*qr_decomposition(A))
Q, R
```

```
(array([[ 0.4472136 ,  0.89442719],
        [ 0.89442719, -0.4472136 ]]),
 array([[ 2.23606798,  1.34164079,  9.8386991 ],
        [ 0.        ,  2.68328157, -0.4472136 ]]))
```

```
Q_scipy, R_scipy = adjust_sign(*qr(A))
Q_scipy, R_scipy
```

```
(array([[ 0.4472136 ,  0.89442719],
        [ 0.89442719, -0.4472136 ]]),
 array([[ 2.23606798,  1.34164079,  9.8386991 ],
        [ 0.        ,  2.68328157, -0.4472136 ]]))
```

3.6 Using QR decomposition to compute eigenvalues

Now for a useful fact about the QR algorithm.

The following iterations on the QR decomposition can be used to compute *eigenvalues* of a *square* matrix A.

Here is the algorithm:

1. Set $A_0 = A$ and form $A_0 = Q_0 R_0$
2. Form $A_1 = R_0 Q_0$. Note that A_1 is similar to A_0 (easy to verify) and so has the same eigenvalues.
3. Form $A_1 = Q_1 R_1$ (i.e., form the QR decomposition of A_1).
4. Form $A_2 = R_1 Q_1$ and then $A_2 = Q_2 R_2$.
5. Iterate to convergence.
6. Compute eigenvalues of A and compare them to the diagonal values of the limiting A_n found from this process.

Remark 3.6.1

This algorithm is close to one of the most efficient ways of computing eigenvalues!

Let's write some Python code to try out the algorithm

```
def qr_eigvals(A, tol=1e-12, maxiter=1000):
    "Find the eigenvalues of A using QR decomposition."

    A_old = np.array(A, dtype=float, copy=True)
    A_new = np.array(A, dtype=float, copy=True)

    diff = np.inf
    i = 0
    while (diff > tol) and (i < maxiter):
        A_old[:, :] = A_new
```

(continues on next page)

(continued from previous page)

```
        Q, R = qr_decomposition(A_old)

        A_new[:, :] = R @ Q

        diff = np.abs(A_new - A_old).max()
        i += 1

    eigvals = np.diag(A_new)

    return eigvals
```

Now let's try the code and compare the results with `np.linalg.eigvals`.

Here goes

```
# experiment this with one random A matrix
A = rng.random((3, 3))
```

```
sorted(qr_eigvals(A))
```

```
[np.float64(0.3056734729625916),
 np.float64(0.3842419370857368),
 np.float64(1.5699606602423861)]
```

Compare with NumPy's implementation.

```
sorted(np.linalg.eigvals(A))
```

```
[np.float64(0.3056734729643013),
 np.float64(0.38424193708402865),
 np.float64(1.569960660242386)]
```

3.7 QR and PCA

There are interesting connections between the QR decomposition and principal components analysis (PCA).

Here are some.

1. Let $X^\top$ be a $k \times n$ random matrix where the jth column is a random draw from $N(\mu, \Sigma)$ where μ is $k \times 1$ vector of means and Σ is a $k \times k$ covariance matrix. We want $n \gg k$ – this is an "econometrics example".
2. Form $X^\top = QR$ where Q is $k \times k$ and R is $k \times n$.
3. Form the eigenvalues of $RR^\top$, i.e., we'll compute $RR^\top = \tilde{P}\Lambda\tilde{P}^\top$.
4. Form $X^\top X = Q\tilde{P}\Lambda\tilde{P}^\top Q^\top$ and compare it with the eigen decomposition $X^\top X = P\hat{\Lambda}P^\top$.
5. It will turn out that $\Lambda = \hat{\Lambda}$ and that $P = Q\tilde{P}$.

Let's verify conjecture 5 with some Python code.

Start by simulating a random (n, k) matrix X.

```
k = 5
n = 1000
```

(continues on next page)

(continued from previous page)

```
# generate some random moments
μ = rng.random(size=k)
C = rng.random((k, k))
Σ = C.T @ C
```

```
# Each row of X is a draw from the multivariate normal distribution
X = rng.multivariate_normal(μ, Σ, size=n)
```

```
X.shape
```

```
(1000, 5)
```

Let's apply the QR decomposition to $X^\top$.

```
Q, R = adjust_sign(*qr_decomposition(X.T))
```

Check the shapes of Q and R.

```
Q.shape, R.shape
```

```
((5, 5), (5, 1000))
```

Now we can construct $RR^\top = \tilde{P}\Lambda\tilde{P}^\top$ and form an eigen decomposition.

```
RR = R @ R.T

λ, P_tilde = np.linalg.eigh(RR)
Λ = np.diag(λ)
```

We can also apply the decomposition to $X^\top X = P\hat{\Lambda}P^\top$.

```
XX = X.T @ X

λ_hat, P = np.linalg.eigh(XX)
Λ_hat = np.diag(λ_hat)
```

Compare the eigenvalues that are on the diagonals of Λ and $\hat{\Lambda}$.

```
λ, λ_hat
```

```
(array([  56.69677761,  225.16561049,  667.8151515 , 1413.45780269,
        8882.27383225]),
 array([  56.69677761,  225.16561049,  667.8151515 , 1413.45780269,
        8882.27383225]))
```

Let's compare P and $Q\tilde{P}$.

Again we need to be careful about sign differences between the columns of P and $Q\tilde{P}$.

```
QP_tilde = Q @ P_tilde

np.abs(P @ diag_sign(P) - QP_tilde @ diag_sign(QP_tilde)).max()
```

```
np.float64(5.162537064506978e-15)
```

Let's verify that $X^\top X$ can be decomposed as $Q\tilde{P}\Lambda\tilde{P}^\top Q^\top$.

```
QPΛPQ = Q @ P_tilde @ Λ @ P_tilde.T @ Q.T
```

```
np.abs(QPΛPQ - XX).max()
```

```
np.float64(2.091837814077735e-11)
```

CHAPTER

FOUR

CIRCULANT MATRICES

4.1 Overview

This lecture describes circulant matrices and some of their properties.

Circulant matrices are useful because multiplying by them is closely connected to convolution, and their eigenvectors can be constructed using the Discrete Fourier Transform.

We use circulant matrices to connect several useful concepts, including

- convolution
- Fourier transforms
- permutation matrices

For background on eigenvalues and eigenvectors, see *Linear Algebra*; for another use of Fourier transforms and convolution, see *Fault Tree Uncertainties*.

Circulant matrices are also widely used in machine learning, for example, in image processing.

We begin by importing some Python packages

```
import numpy as np
import matplotlib.pyplot as plt
```

```
np.set_printoptions(precision=3, suppress=True)
```

4.2 Constructing a Circulant Matrix

To construct an $N \times N$ circulant matrix, we need only the first row, say,

$$\begin{bmatrix} c_0 & c_1 & c_2 & c_3 & c_4 & \cdots & c_{N-1} \end{bmatrix}.$$

After setting entries in the first row, the remaining rows of a circulant matrix are determined as follows:

$$C = \begin{bmatrix} c_0 & c_1 & c_2 & c_3 & c_4 & \cdots & c_{N-1} \\ c_{N-1} & c_0 & c_1 & c_2 & c_3 & \cdots & c_{N-2} \\ c_{N-2} & c_{N-1} & c_0 & c_1 & c_2 & \cdots & c_{N-3} \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ c_3 & c_4 & c_5 & c_6 & c_7 & \cdots & c_2 \\ c_2 & c_3 & c_4 & c_5 & c_6 & \cdots & c_1 \\ c_1 & c_2 & c_3 & c_4 & c_5 & \cdots & c_0 \end{bmatrix} \tag{4.1}$$

This pattern can be formalized as follows.

Definition 4.2.1 (Circulant matrix)

An $N \times N$ matrix C is **circulant** if there are numbers $c_0, \dots, c_{N-1}$ such that

$$C_{ij} = c_{(j-i) \bmod N}, \qquad 0 \leq i, j \leq N-1.$$

Equivalently, each row is obtained from the previous row by shifting entries one step to the right.

It is also possible to construct a circulant matrix by creating the transpose of the above matrix, in which case only the first column needs to be specified.

Let's write some Python code to generate a circulant matrix.

```
def construct_circulant(row):

    N = row.size

    C = np.empty((N, N))

    for i in range(N):

        C[i, i:] = row[:N-i]
        C[i, :i] = row[N-i:]

    return C
```

```
# a simple case when N = 3
construct_circulant(np.array([1., 2., 3.]))
```

```
array([[1., 2., 3.],
       [3., 1., 2.],
       [2., 3., 1.]])
```

4.2.1 Some Properties of Circulant Matrices

Here are some useful properties:

Suppose that A and B are both circulant matrices of the same order and constructed using the same cyclic shift convention.

Then it can be verified that

- The transpose of a circulant matrix is a circulant matrix.
- $A + B$ is a circulant matrix
- AB is a circulant matrix
- $AB = BA$

Now consider a circulant matrix with first row

$$c = \begin{bmatrix} c_0 & c_1 & \cdots & c_{N-1} \end{bmatrix}$$

and consider a vector

$$a = \begin{bmatrix} a_0 & a_1 & \cdots & a_{N-1} \end{bmatrix}$$

The **convolution** of vectors c and a is defined as the vector $b = c * a$ with components

$$b_k = \sum_{i=0}^{N-1} c_{k-i} a_i \tag{4.2}$$

Here and below, indices such as $k - i$ are interpreted modulo N.

We use $*$ to denote **convolution** via the calculation described in equation (4.2).

It can be verified that the vector b satisfies

$$b = C^\top a$$

where $C^\top$ is the transpose of the circulant matrix defined in equation (4.1).

4.3 Connection to Permutation Matrix

A good way to construct a circulant matrix is to use a **permutation matrix**.

Before defining a permutation **matrix**, we'll define a **permutation**.

A **permutation** of a set of the set of non-negative integers $\{0, 1, 2, ...\}$ is a one-to-one mapping of the set into itself.

A permutation of a set $\{1, 2, ..., n\}$ rearranges the n integers in the set.

A permutation matrix is obtained by permuting the rows of an $n \times n$ identity matrix according to a permutation of the numbers 1 to n.

Thus, every row and every column contain precisely a single 1 with 0 everywhere else.

Every permutation corresponds to a unique permutation matrix.

For example, the $N \times N$ matrix

$$P = \begin{bmatrix} 0 & 1 & 0 & 0 & \cdots & 0 \\ 0 & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \\ 1 & 0 & 0 & 0 & \cdots & 0 \end{bmatrix} \tag{4.3}$$

serves as a **cyclic shift** operator that, when applied to an $N \times 1$ vector h, shifts entries in rows 2 through N up one row and shifts the entry in row 1 to row N.

Eigenvalues of the cyclic shift permutation matrix P defined in equation (4.3) can be computed by constructing

$$P - \lambda I = \begin{bmatrix} -\lambda & 1 & 0 & 0 & \cdots & 0 \\ 0 & -\lambda & 1 & 0 & \cdots & 0 \\ 0 & 0 & -\lambda & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & 1 \\ 1 & 0 & 0 & 0 & \cdots & -\lambda \end{bmatrix}$$

and solving

$$\det(P - \lambda I) = (-1)^N (\lambda^N - 1) = 0$$

Eigenvalues λ_i can be complex.

Magnitudes $\mid \lambda_i \mid$ of these eigenvalues λ_i all equal 1.

Thus, **singular values** of the permutation matrix P defined in equation (4.3) all equal 1.

It can be verified that permutation matrices are orthogonal matrices:

$$PP^\top = I$$

4.4 Examples with Python

Let's write some Python code to illustrate these ideas.

```
def construct_cyclic_shift_matrix(N):

    P = np.zeros((N, N))

    for i in range(N-1):
        P[i, i+1] = 1
    P[-1, 0] = 1

    return P
```

```
P4 = construct_cyclic_shift_matrix(4)
P4
```

```
array([[0., 1., 0., 0.],
       [0., 0., 1., 0.],
       [0., 0., 0., 1.],
       [1., 0., 0., 0.]])
```

```
# compute the eigenvalues and eigenvectors
λ, Q = np.linalg.eig(P4)
```

```
for i in range(4):
    print(f'λ{i} = {λ[i]:.1f} \nvec{i} = {Q[:, i]}\n')
```

```
λ0 = -1.0+0.0j
vec0 = [-0.5+0.j  0.5+0.j -0.5+0.j  0.5+0.j]

λ1 = 0.0+1.0j
vec1 = [-0. -0.5j  0.5+0.j   0. +0.5j -0.5-0.j ]

λ2 = 0.0-1.0j
vec2 = [-0. +0.5j  0.5-0.j   0. -0.5j -0.5+0.j ]

λ3 = 1.0+0.0j
vec3 = [-0.5+0.j -0.5+0.j -0.5+0.j -0.5+0.j]
```

In graphs below, we shall portray eigenvalues of a shift permutation matrix in the complex plane.

These eigenvalues are uniformly distributed along the unit circle.

They are the n **roots of unity**, meaning they are the n numbers z that solve $z^n = 1$, where z is a complex number.

In particular, the n roots of unity are

$$z = \exp\left(\frac{2\pi jk}{N}\right), \quad k = 0, \ldots, N-1$$

where j denotes the purely imaginary unit number.

```
fig, ax = plt.subplots(2, 2, figsize=(10, 10))

for i, N in enumerate([3, 4, 6, 8]):

    row_i = i // 2
    col_i = i % 2

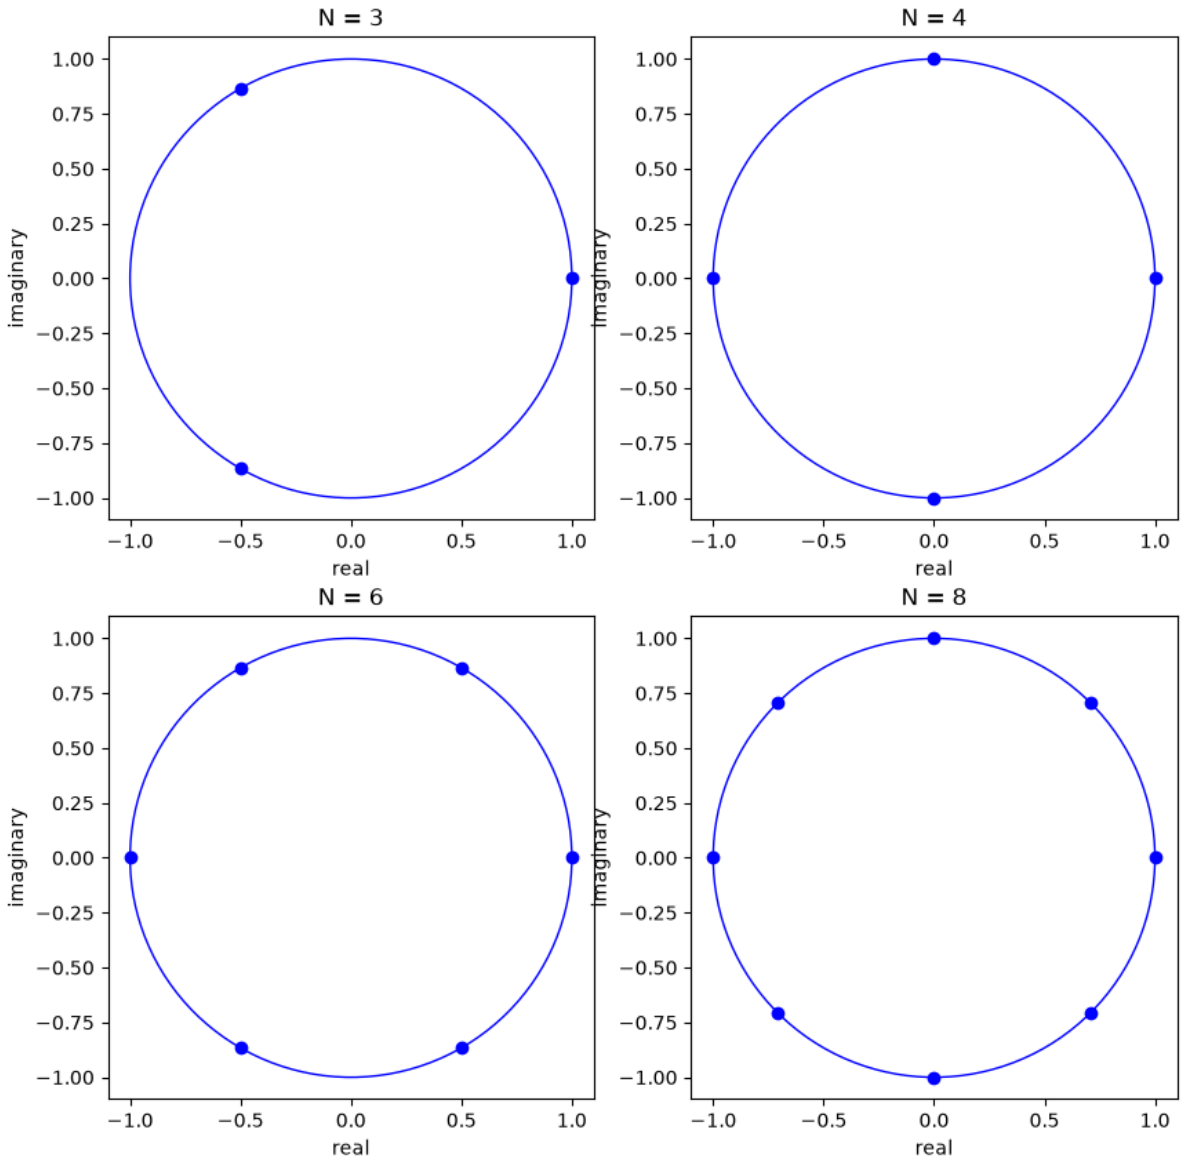
    P = construct_cyclic_shift_matrix(N)
    λ, Q = np.linalg.eig(P)

    circ = plt.Circle((0, 0), radius=1, edgecolor='b', facecolor='None')
    ax[row_i, col_i].add_patch(circ)

    for j in range(N):
        ax[row_i, col_i].scatter(λ[j].real, λ[j].imag, c='b')

    ax[row_i, col_i].set_title(f'N = {N}')
    ax[row_i, col_i].set_xlabel('real')
    ax[row_i, col_i].set_ylabel('imaginary')

plt.show()
```



For a vector of coefficients $\{c_i\}_{i=0}^{n-1}$, eigenvectors of P are also eigenvectors of

$$C = c_0 I + c_1 P + c_2 P^2 + \cdots + c_{N-1} P^{N-1}.$$

Consider an example in which $N = 8$ and let $w = e^{-2\pi j/N}$.

It can be verified that the matrix F_8 of eigenvectors of P_8 is

$$F_8 = \begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ 1 & w & w^2 & \cdots & w^7 \\ 1 & w^2 & w^4 & \cdots & w^{14} \\ 1 & w^3 & w^6 & \cdots & w^{21} \\ 1 & w^4 & w^8 & \cdots & w^{28} \\ 1 & w^5 & w^{10} & \cdots & w^{35} \\ 1 & w^6 & w^{12} & \cdots & w^{42} \\ 1 & w^7 & w^{14} & \cdots & w^{49} \end{bmatrix}$$

The matrix F_8 defines a Discrete Fourier Transform.

To convert it into an orthogonal eigenvector matrix, we can simply normalize it by dividing every entry by $\sqrt{8}$.

- stare at the first column of F_8 above to convince yourself of this fact

The eigenvalues corresponding to each eigenvector are $\{w^j\}_{j=0}^7$ in order.

```
def construct_F(N):

    w = np.e ** (-complex(0, 2*np.pi/N))

    F = np.ones((N, N), dtype=complex)
    for i in range(1, N):
        F[i, 1:] = w ** (i * np.arange(1, N))

    return F, w
```

```
F8, w = construct_F(8)
```

```
w
```

```
(0.7071067811865476-0.7071067811865475j)
```

```
F8
```

```
array([[ 1.   +0.j   ,  1.   +0.j   ,  1.   +0.j   ,  1.   +0.j   ,
         1.   +0.j   ,  1.   +0.j   ,  1.   +0.j   ,  1.   +0.j   ],
       [ 1.   +0.j   ,  0.707-0.707j,  0.   -1.j   , -0.707-0.707j,
        -1.   -0.j   , -0.707+0.707j, -0.   +1.j   ,  0.707+0.707j],
       [ 1.   +0.j   ,  0.   -1.j   , -1.   -0.j   , -0.   +1.j   ,
         1.   +0.j   ,  0.   -1.j   , -1.   -0.j   , -0.   +1.j   ],
       [ 1.   +0.j   , -0.707-0.707j, -0.   +1.j   ,  0.707-0.707j,
        -1.   -0.j   ,  0.707+0.707j,  0.   -1.j   , -0.707+0.707j],
       [ 1.   +0.j   , -1.   -0.j   ,  1.   +0.j   , -1.   -0.j   ,
         1.   +0.j   , -1.   -0.j   ,  1.   +0.j   , -1.   -0.j   ],
       [ 1.   +0.j   , -0.707+0.707j,  0.   -1.j   ,  0.707+0.707j,
        -1.   -0.j   ,  0.707-0.707j, -0.   +1.j   , -0.707-0.707j],
       [ 1.   +0.j   , -0.   +1.j   , -1.   -0.j   ,  0.   -1.j   ,
         1.   +0.j   , -0.   +1.j   , -1.   -0.j   ,  0.   -1.j   ],
       [ 1.   +0.j   ,  0.707+0.707j, -0.   +1.j   , -0.707+0.707j,
        -1.   -0.j   , -0.707-0.707j,  0.   -1.j   ,  0.707-0.707j]])
```

```
# normalize
Q8 = F8 / np.sqrt(8)
```

```
# verify the orthogonality (unitarity)
Q8 @ np.conjugate(Q8)
```

```
array([[ 1.+0.j, -0.+0.j, -0.+0.j, -0.+0.j, -0.+0.j,  0.+0.j,  0.+0.j,
         0.+0.j],
       [-0.-0.j,  1.+0.j, -0.+0.j, -0.+0.j, -0.+0.j, -0.+0.j,  0.+0.j,
         0.+0.j],
       [-0.-0.j, -0.-0.j,  1.+0.j, -0.+0.j, -0.+0.j, -0.+0.j,  0.+0.j,
         0.+0.j],
       [-0.-0.j, -0.-0.j, -0.-0.j,  1.+0.j, -0.+0.j, -0.+0.j, -0.+0.j,
        -0.+0.j],
       [-0.-0.j, -0.-0.j, -0.-0.j, -0.-0.j,  1.+0.j, -0.+0.j, -0.+0.j,
```

(continues on next page)

(continued from previous page)

```
        -0.+0.j],
       [ 0.-0.j, -0.-0.j, -0.-0.j, -0.-0.j, -0.-0.j,  1.+0.j, -0.+0.j,
        -0.+0.j],
       [ 0.-0.j,  0.-0.j,  0.-0.j, -0.-0.j, -0.-0.j, -0.-0.j,  1.+0.j,
        -0.+0.j],
       [ 0.-0.j,  0.-0.j,  0.-0.j, -0.-0.j, -0.-0.j, -0.-0.j, -0.-0.j,
         1.+0.j]])
```

Let's verify that kth column of Q_8 is an eigenvector of P_8 with an eigenvalue w^k.

```
P8 = construct_cyclic_shift_matrix(8)
```

```
diff_arr = np.empty(8, dtype=complex)
for j in range(8):
    diff = P8 @ Q8[:, j] - w ** j * Q8[:, j]
    diff_arr[j] = diff @ diff.T
```

```
diff_arr
```

```
array([ 0.+0.j, -0.+0.j, -0.+0.j, -0.+0.j, -0.+0.j, -0.+0.j, -0.+0.j,
       -0.+0.j])
```

4.5 Associated Permutation Matrix

Next, we execute calculations to verify that the circulant matrix C defined in equation (4.1) can be written as

$$C = c_0 I + c_1 P + \cdots + c_{N-1} P^{N-1}$$

and that every eigenvector of P is also an eigenvector of C.

We illustrate this for $N = 8$ case.

```
rng = np.random.default_rng()
c = rng.random(8)
```

```
c
```

```
array([0.233, 0.207, 0.636, 0.702, 0.501, 0.135, 0.601, 0.719])
```

```
C8 = construct_circulant(c)
```

Compute $c_0 I + c_1 P + \cdots + c_{N-1} P^{N-1}$.

```
N = 8

C = np.zeros((N, N))
P = np.eye(N)

for i in range(N):
    C += c[i] * P
    P = P8 @ P
```

```
C
```

```
array([[0.233, 0.207, 0.636, 0.702, 0.501, 0.135, 0.601, 0.719],
       [0.719, 0.233, 0.207, 0.636, 0.702, 0.501, 0.135, 0.601],
       [0.601, 0.719, 0.233, 0.207, 0.636, 0.702, 0.501, 0.135],
       [0.135, 0.601, 0.719, 0.233, 0.207, 0.636, 0.702, 0.501],
       [0.501, 0.135, 0.601, 0.719, 0.233, 0.207, 0.636, 0.702],
       [0.702, 0.501, 0.135, 0.601, 0.719, 0.233, 0.207, 0.636],
       [0.636, 0.702, 0.501, 0.135, 0.601, 0.719, 0.233, 0.207],
       [0.207, 0.636, 0.702, 0.501, 0.135, 0.601, 0.719, 0.233]])
```

```
C8
```

```
array([[0.233, 0.207, 0.636, 0.702, 0.501, 0.135, 0.601, 0.719],
       [0.719, 0.233, 0.207, 0.636, 0.702, 0.501, 0.135, 0.601],
       [0.601, 0.719, 0.233, 0.207, 0.636, 0.702, 0.501, 0.135],
       [0.135, 0.601, 0.719, 0.233, 0.207, 0.636, 0.702, 0.501],
       [0.501, 0.135, 0.601, 0.719, 0.233, 0.207, 0.636, 0.702],
       [0.702, 0.501, 0.135, 0.601, 0.719, 0.233, 0.207, 0.636],
       [0.636, 0.702, 0.501, 0.135, 0.601, 0.719, 0.233, 0.207],
       [0.207, 0.636, 0.702, 0.501, 0.135, 0.601, 0.719, 0.233]])
```

Now let's compute the difference between two circulant matrices that we have constructed in two different ways.

```
np.abs(C - C8).max()
```

```
np.float64(0.0)
```

The jth column of Q_8 is an eigenvector of C_8 associated with eigenvalue $\sum_{k=0}^{7} c_k w^{jk}$.

```
λ_C8 = np.zeros(8, dtype=complex)

for j in range(8):
    for k in range(8):
        λ_C8[j] += c[k] * w ** (j * k)
```

```
λ_C8
```

```
array([ 3.735+0.j   , -0.205-0.075j, -0.502+1.079j, -0.331-0.004j,
        0.208-0.j   , -0.331+0.004j, -0.502-1.079j, -0.205+0.075j])
```

We can verify this by comparing `C8 @ Q8[:, j]` with `λ_C8[j] * Q8[:, j]`.

```
# verify
for j in range(8):
    diff = C8 @ Q8[:, j] - λ_C8[j] * Q8[:, j]
    print(diff)
```

```
[0.+0.j 0.+0.j 0.+0.j 0.+0.j 0.+0.j 0.+0.j 0.+0.j 0.+0.j]
[ 0.+0.j  0.-0.j -0.-0.j -0.-0.j -0.-0.j -0.+0.j -0.+0.j -0.+0.j]
[-0.+0.j  0.-0.j -0.-0.j -0.+0.j  0.+0.j  0.-0.j -0.-0.j -0.+0.j]
[-0.+0.j -0.-0.j -0.-0.j  0.+0.j  0.-0.j -0.-0.j  0.-0.j -0.-0.j]
[ 0.-0.j -0.-0.j  0.+0.j -0.-0.j  0.-0.j -0.-0.j  0.-0.j -0.-0.j]
[ 0.+0.j -0.-0.j  0.-0.j -0.+0.j  0.-0.j  0.-0.j -0.-0.j  0.-0.j]
```

(continues on next page)

(continued from previous page)

```
[ 0.+0.j  0.-0.j  0.-0.j  0.+0.j -0.+0.j -0.-0.j  0.-0.j  0.+0.j]
[0.-0.j 0.-0.j 0.-0.j 0.-0.j 0.-0.j 0.+0.j 0.+0.j 0.+0.j]
```

4.6 Discrete Fourier Transform

The **Discrete Fourier Transform** (DFT) allows us to represent a discrete time sequence as a weighted sum of complex sinusoids.

Consider a sequence of N real numbers $\{x_j\}_{j=0}^{N-1}$.

The **Discrete Fourier Transform** maps $\{x_j\}_{j=0}^{N-1}$ into a sequence of complex numbers $\{X_k\}_{k=0}^{N-1}$

where

$$X_k = \sum_{n=0}^{N-1} x_n e^{-2\pi \frac{kn}{N} i}$$

```
def DFT(x):
    "The discrete Fourier transform."

    N = len(x)
    w = np.e ** (-complex(0, 2*np.pi/N))

    X = np.zeros(N, dtype=complex)
    for k in range(N):
        for n in range(N):
            X[k] += x[n] * w ** (k * n)

    return X
```

Consider the following example.

$$x_n = \begin{cases} 1/2 & n = 0, 1 \\ 0 & \text{otherwise} \end{cases}$$

```
x = np.zeros(10)
x[0:2] = 1/2
```

```
x
```

```
array([0.5, 0.5, 0. , 0. , 0. , 0. , 0. , 0. , 0. , 0. ])
```

Apply a discrete Fourier transform.

```
X = DFT(x)
```

```
X
```

```
array([ 1.   +0.j   ,  0.905-0.294j,  0.655-0.476j,  0.345-0.476j,
        0.095-0.294j, -0.   +0.j   ,  0.095+0.294j,  0.345+0.476j,
        0.655+0.476j,  0.905+0.294j])
```

We can plot magnitudes of a sequence of numbers and the associated discrete Fourier transform.

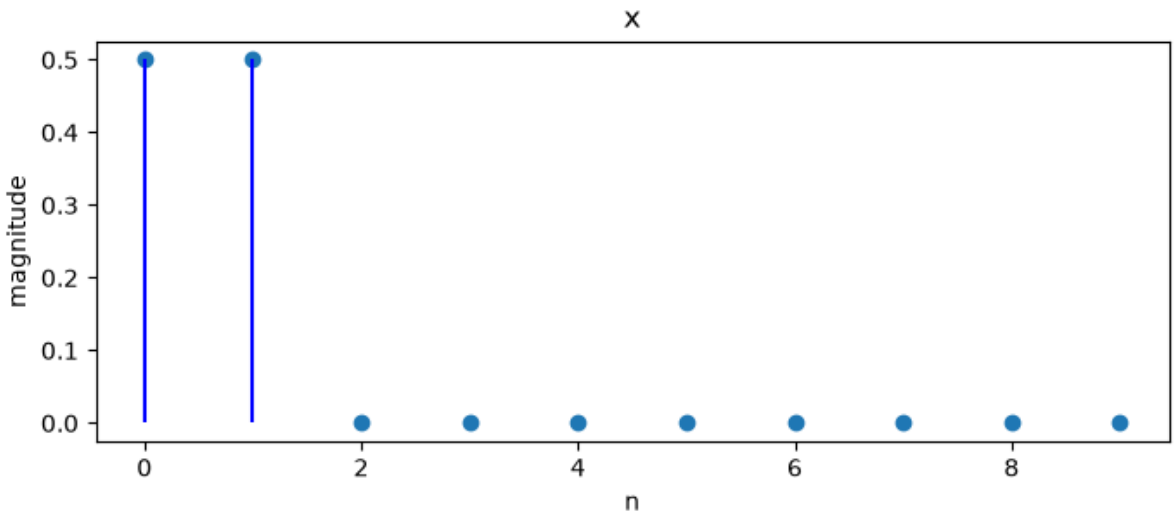
```
def plot_magnitude(x=None, X=None):

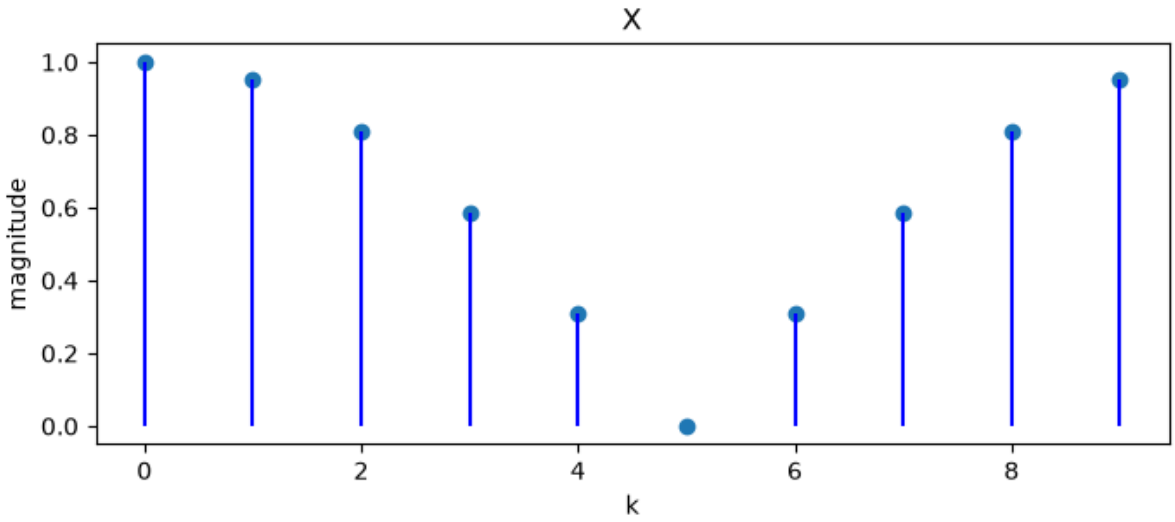
    data = []
    names = []
    xs = []
    if (x is not None):
        data.append(x)
        names.append('x')
        xs.append('n')
    if (X is not None):
        data.append(X)
        names.append('X')
        xs.append('k')

    num = len(data)
    for i in range(num):
        n = data[i].size
        plt.figure(figsize=(8, 3))
        plt.scatter(range(n), np.abs(data[i]))
        plt.vlines(range(n), 0, np.abs(data[i]), color='b')

        plt.xlabel(xs[i])
        plt.ylabel('magnitude')
        plt.title(names[i])
        plt.show()
```

```
plot_magnitude(x=x, X=X)
```





The **inverse Fourier transform** transforms a Fourier transform X of x back to x.

The inverse Fourier transform is defined as

$$x_n = \sum_{k=0}^{N-1} \frac{1}{N} X_k e^{2\pi(\frac{kn}{N})i}, \quad n = 0, 1, \dots, N-1$$

```
def inverse_DFT(X):

    N = len(X)
    w = np.e ** (complex(0, 2*np.pi/N))

    x = np.zeros(N, dtype=complex)
    for n in range(N):
        for k in range(N):
            x[n] += X[k] * w ** (k * n) / N

    return x
```

```
inverse_DFT(X)
```

```
array([ 0.5+0.j,  0.5-0.j, -0. -0.j, -0. -0.j, -0. -0.j, -0. -0.j, -0. -0.j,
       -0. +0.j, -0. +0.j, -0. +0.j, -0. +0.j])
```

Another example is

$$x_n = 2\cos\left(2\pi \frac{11}{40} n\right), \; n = 0, 1, 2, \cdots 19$$

Since $N = 20$, we cannot use an integer multiple of $\frac{1}{20}$ to represent a frequency $\frac{11}{40}$.

To handle this, we shall end up using all N of the available frequencies in the DFT.

Since $\frac{11}{40}$ is in between $\frac{10}{40}$ and $\frac{12}{40}$ (each of which is an integer multiple of $\frac{1}{20}$), the complex coefficients in the DFT have their largest magnitudes at $k = 5, 6, 15, 16$, not just at a single frequency.

```
N = 20
x = np.empty(N)
```

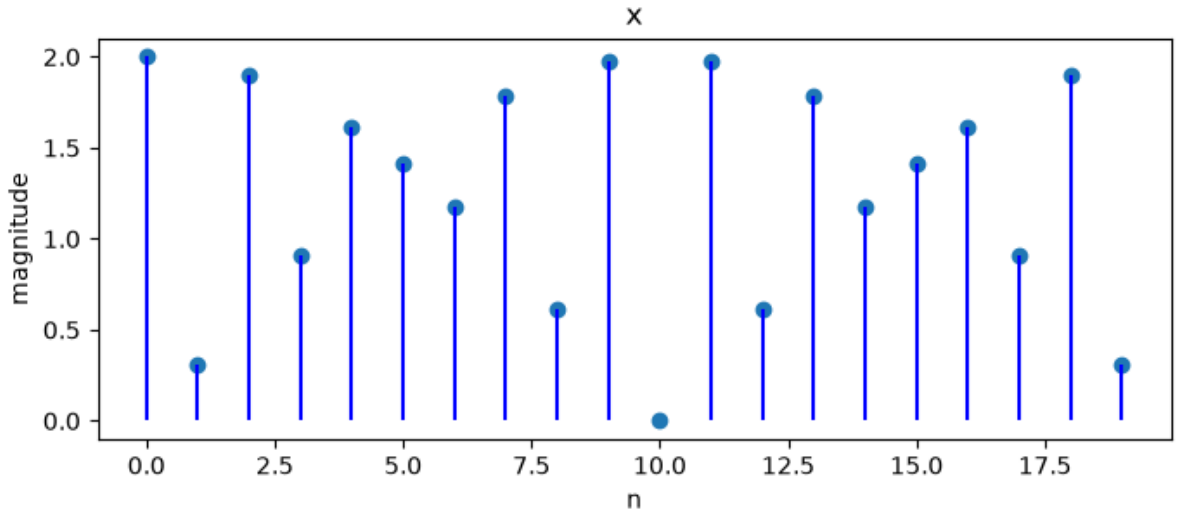
(continues on next page)

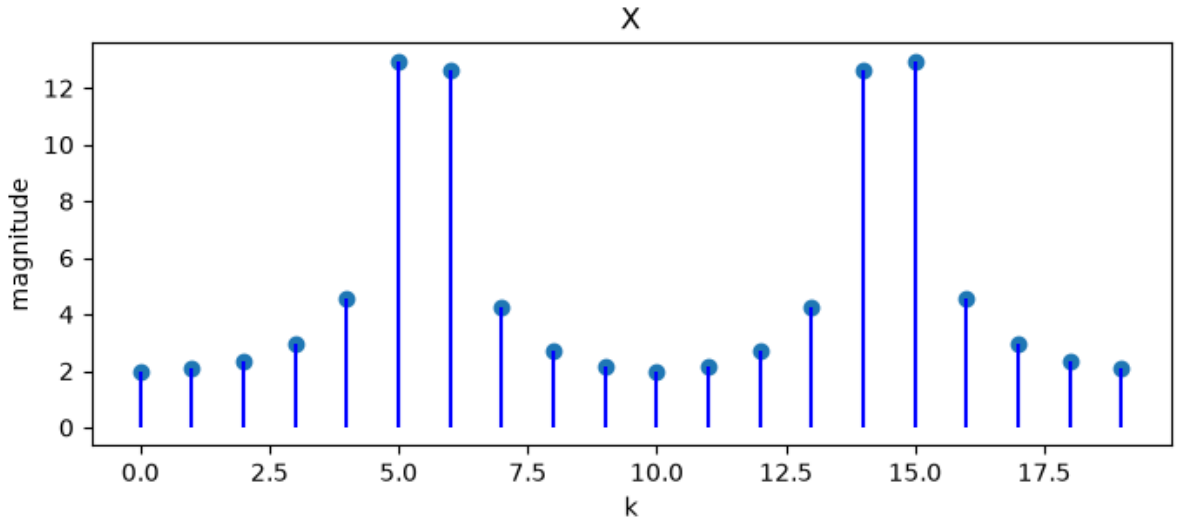
(continued from previous page)

```
for j in range(N):
    x[j] = 2 * np.cos(2 * np.pi * 11 * j / 40)
```

```
X = DFT(x)
```

```
plot_magnitude(x=x, X=X)
```





What happens if we change the last example to $x_n = 2\cos\left(2\pi\frac{10}{40}n\right)$?

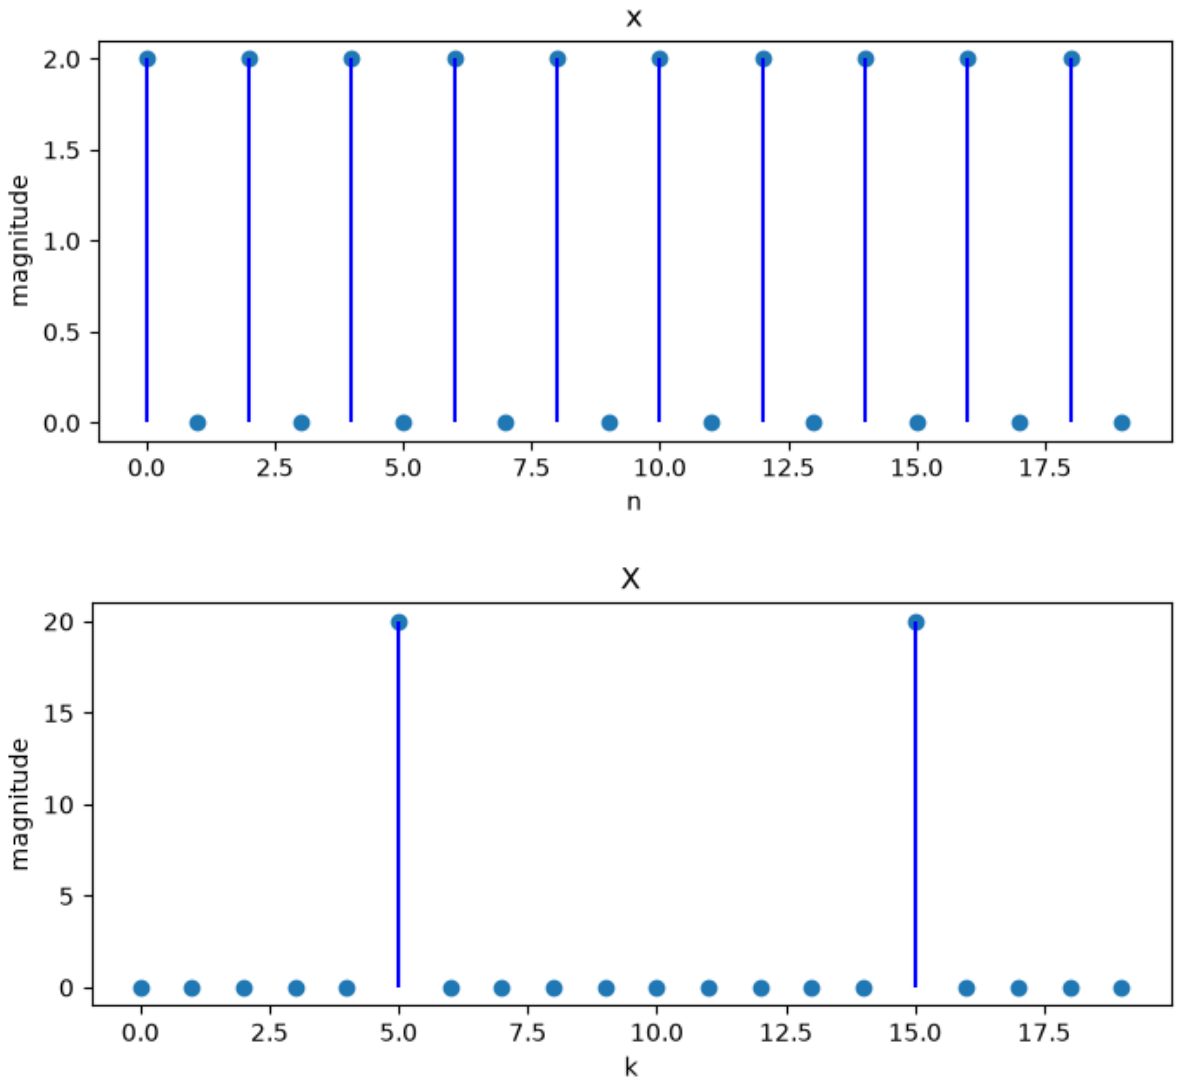
Note that $\frac{10}{40}$ is an integer multiple of $\frac{1}{20}$.

```
N = 20
x = np.empty(N)

for j in range(N):
    x[j] = 2 * np.cos(2 * np.pi * 10 * j / 40)
```

```
X = DFT(x)
```

```
plot_magnitude(x=x, X=X)
```



If we represent the discrete Fourier transform as a matrix, we discover that it equals the matrix F_N of eigenvectors of the permutation matrix P_N.

We can use the example where $x_n = 2\cos\left(2\pi\frac{11}{40}n\right),\ n = 0, 1, 2, \cdots 19$ to illustrate this.

```
N = 20
x = np.empty(N)

for j in range(N):
    x[j] = 2 * np.cos(2 * np.pi * 11 * j / 40)
```

```
x
```

```
array([ 2.   , -0.313, -1.902,  0.908,  1.618, -1.414, -1.176,  1.782,
        0.618, -1.975, -0.   ,  1.975, -0.618, -1.782,  1.176,  1.414,
       -1.618, -0.908,  1.902,  0.313])
```

First use the summation formula to transform x to X.

```
X = DFT(x)
X
```

```
array([2. +0.j   , 2. +0.558j, 2. +1.218j, 2. +2.174j, 2. +4.087j,
       2.+12.785j, 2.-12.466j, 2. -3.751j, 2. -1.801j, 2. -0.778j,
       2. -0.j   , 2. +0.778j, 2. +1.801j, 2. +3.751j, 2.+12.466j,
       2.-12.785j, 2. -4.087j, 2. -2.174j, 2. -1.218j, 2. -0.558j])
```

Now let's evaluate the outcome of postmultiplying the eigenvector matrix F_{20} by the vector x, a product that we claim should equal the Fourier transform of the sequence $\{x_n\}_{n=0}^{N-1}$.

```
F20, _ = construct_F(20)
```

```
F20 @ x
```

```
array([2. +0.j   , 2. +0.558j, 2. +1.218j, 2. +2.174j, 2. +4.087j,
       2.+12.785j, 2.-12.466j, 2. -3.751j, 2. -1.801j, 2. -0.778j,
       2. -0.j   , 2. +0.778j, 2. +1.801j, 2. +3.751j, 2.+12.466j,
       2.-12.785j, 2. -4.087j, 2. -2.174j, 2. -1.218j, 2. -0.558j])
```

Similarly, the inverse DFT can be expressed as an inverse DFT matrix F_{20}^{-1}.

```
F20_inv = np.linalg.inv(F20)
F20_inv @ X
```

```
array([ 2.   -0.j, -0.313-0.j, -1.902+0.j,  0.908-0.j,  1.618-0.j,
       -1.414+0.j, -1.176+0.j,  1.782+0.j,  0.618-0.j, -1.975-0.j,
       -0.   +0.j,  1.975-0.j, -0.618-0.j, -1.782+0.j,  1.176+0.j,
        1.414-0.j, -1.618-0.j, -0.908+0.j,  1.902+0.j,  0.313-0.j])
```

CHAPTER

FIVE

SINGULAR VALUE DECOMPOSITION (SVD)

5.1 Overview

The **singular value decomposition** (SVD) is a work-horse in applications of least squares projection that form foundations for many statistical and machine learning methods.

After defining the SVD, we'll describe how it connects to

- **four fundamental spaces** of linear algebra
- **low-rank approximations** of matrices
- under-determined and over-determined **least squares regressions** (see the exercises)
- **principal components analysis** (PCA)

In a sequel to this lecture, *VARs and DMDs*, we'll describe how SVDs provide ways rapidly to compute reduced-rank approximations to first-order vector autoregressions (VARs) by means of a **dynamic mode decomposition** (DMD).

Like principal components analysis, DMD can be thought of as a data-reduction procedure that represents salient patterns by projecting data onto a limited set of factors.

5.2 The setting

Let X be an $m \times n$ matrix of rank p.

Necessarily, $p \leq \min(m, n)$.

In much of this lecture, we'll think of X as a matrix of data in which

- each column is an **individual** – a time period or person, depending on the application
- each row is a **random variable** describing an attribute of a time period or a person, depending on the application

We'll be interested in two situations

- A **short and fat** case in which $m \ll n$, so that there are many more columns (individuals) than rows (attributes).
- A **tall and skinny** case in which $m \gg n$, so that there are many more rows (attributes) than columns (individuals).

We'll apply a **singular value decomposition** of X in both situations.

In the $m \ll n$ case in which there are many more individuals n than attributes m, we can calculate sample moments of a joint distribution by taking averages across observations of functions of the observations.

In this $m \ll n$ case, we'll look for patterns by using a singular value decomposition to do a principal components analysis (PCA).

In the $m \gg n$ case in which there are many more attributes m than individuals n and when we are in a time-series setting in which n equals the number of time periods covered in the data set X, we'll proceed in a different way.

In the sequel *VARs and DMDs*, we'll again use a singular value decomposition, but now to construct a **dynamic mode decomposition** (DMD).

5.3 Singular value decomposition

A **singular value decomposition** of an $m \times n$ matrix X of rank $p \leq \min(m, n)$ is

$$X = U\Sigma V^\top \tag{5.1}$$

where

$$\begin{aligned} UU^\top &= I \qquad U^\top U = I \\ VV^\top &= I \qquad V^\top V = I \end{aligned}$$

and

- U is an $m \times m$ orthogonal matrix of **left singular vectors** of X
- Columns of U are eigenvectors of $XX^\top$
- V is an $n \times n$ orthogonal matrix of **right singular vectors** of X
- Columns of V are eigenvectors of $X^\top X$
- Σ is an $m \times n$ matrix in which the first p places on its main diagonal are positive numbers $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p$ called **singular values**; remaining entries of Σ are all zero
- The p singular values are positive square roots of the nonzero eigenvalues of the $m \times m$ matrix $XX^\top$ and also of the $n \times n$ matrix $X^\top X$
- We adopt a convention that when U is a complex valued matrix, $U^\top$ denotes the **conjugate-transpose** or **Hermitian-transpose** of U, meaning that $U^\top_{ij}$ is the complex conjugate of U_{ji}.
- Similarly, when V is a complex valued matrix, $V^\top$ denotes the **conjugate-transpose** or **Hermitian-transpose** of V

The matrices U, Σ, V entail linear transformations that reshape vectors in the following ways:

- multiplying vectors by the unitary matrices U and V **rotates** them, but leaves **angles between vectors** and **lengths of vectors** unchanged.
- multiplying vectors by the diagonal matrix Σ **rescales** each coordinate by a singular value (and adds or drops coordinates when $m \neq n$); in general this changes both lengths of vectors and angles between them.

Thus, representation (5.1) asserts that multiplying an $n \times 1$ vector y by the $m \times n$ matrix X amounts to performing the following three multiplications of y sequentially:

- **rotating** y by computing $V^\top y$
- **rescaling** $V^\top y$ by multiplying it by Σ
- **rotating** $\Sigma V^\top y$ by multiplying it by U

This structure of the $m \times n$ matrix X opens the door to constructing systems of data **encoders** and **decoders**.

Thus,

- $V^\top$ is an encoder
- Σ is an operator to be applied to the encoded data

- U is a decoder to be applied to the output from applying operator Σ to the encoded data

We'll apply this circle of ideas in the sequel *VARs and DMDs* when we study dynamic mode decompositions.

Road ahead

What we have described above is called a **full** SVD.

In a **full** SVD, the shapes of U, Σ, and V are (m, m), (m, n), (n, n), respectively.

Later we'll also describe an **economy** or **reduced** SVD.

Before we study a **reduced** SVD we'll say a little more about properties of a **full** SVD.

5.4 Four fundamental subspaces

Let $\mathcal{C}$ denote a column space, $\mathcal{N}$ denote a null space, and $\mathcal{R}$ denote a row space.

Let's start by recalling the four fundamental subspaces of an $m \times n$ matrix X of rank p.

- The **column space** of X, denoted $\mathcal{C}(X)$, is the span of the columns of X, i.e., all vectors y that can be written as linear combinations of columns of X. Its dimension is p.
- The **null space** of X, denoted $\mathcal{N}(X)$ consists of all vectors y that satisfy $Xy = 0$. Its dimension is $n - p$.
- The **row space** of X, denoted $\mathcal{R}(X)$ is the column space of $X^\top$. It consists of all vectors z that can be written as linear combinations of rows of X. Its dimension is p.
- The **left null space** of X, denoted $\mathcal{N}(X^\top)$, consists of all vectors z such that $X^\top z = 0$. Its dimension is $m - p$.

For a full SVD of a matrix X, the matrix U of left singular vectors and the matrix V of right singular vectors contain orthogonal bases for all four subspaces.

They form two pairs of orthogonal subspaces that we'll describe now.

Let $u_i, i = 1, \ldots, m$ be the m column vectors of U and let $v_i, i = 1, \ldots, n$ be the n column vectors of V.

Let's write the full SVD of X as

$$X = \begin{bmatrix} U_L & U_R \end{bmatrix} \begin{bmatrix} \Sigma_p & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} V_L & V_R \end{bmatrix}^\top \tag{5.2}$$

where Σ_p is a $p \times p$ diagonal matrix with the p singular values on the diagonal and

$$\begin{aligned} U_L &= \begin{bmatrix} u_1 & \cdots & u_p \end{bmatrix}, \quad U_R = \begin{bmatrix} u_{p+1} & \cdots & u_m \end{bmatrix} \\ V_L &= \begin{bmatrix} v_1 & \cdots & v_p \end{bmatrix}, \quad V_R = \begin{bmatrix} v_{p+1} & \cdots & v_n \end{bmatrix} \end{aligned}$$

Representation (5.2) implies that

$$X \begin{bmatrix} V_L & V_R \end{bmatrix} = \begin{bmatrix} U_L & U_R \end{bmatrix} \begin{bmatrix} \Sigma_p & 0 \\ 0 & 0 \end{bmatrix}$$

or

$$\begin{aligned} XV_L &= U_L \Sigma_p \\ XV_R &= 0 \end{aligned} \tag{5.3}$$

or

$$\begin{aligned} Xv_i &= \sigma_i u_i, \quad i = 1, \ldots, p \\ Xv_i &= 0, \quad i = p+1, \ldots, n \end{aligned} \tag{5.4}$$

Equations (5.4) tell how the transformation X maps a pair of orthonormal vectors v_i, v_j for i and j both less than or equal to the rank p of X into a pair of orthogonal vectors $\sigma_i u_i, \sigma_j u_j$ with lengths σ_i and σ_j.

Equations (5.3) assert that

$$\begin{aligned}\mathcal{C}(X) &= \mathcal{C}(U_L)\\ \mathcal{N}(X) &= \mathcal{C}(V_R)\end{aligned}$$

Taking transposes on both sides of representation (5.2) implies

$$X^\top \begin{bmatrix} U_L & U_R \end{bmatrix} = \begin{bmatrix} V_L & V_R \end{bmatrix} \begin{bmatrix} \Sigma_p & 0 \\ 0 & 0 \end{bmatrix}$$

or

$$\begin{aligned}X^\top U_L &= V_L \Sigma_p\\ X^\top U_R &= 0\end{aligned} \tag{5.5}$$

or

$$\begin{aligned}X^\top u_i &= \sigma_i v_i, \quad i = 1, \ldots, p\\ X^\top u_i &= 0 \quad i = p+1, \ldots, m\end{aligned} \tag{5.6}$$

Notice how equations (5.6) assert that the transformation $X^\top$ maps a pair of distinct orthonormal vectors u_i, u_j for i and j both less than or equal to the rank p of X into a pair of orthogonal vectors $\sigma_i v_i, \sigma_j v_j$.

Equations (5.5) assert that

$$\begin{aligned}\mathcal{R}(X) \equiv \mathcal{C}(X^\top) &= \mathcal{C}(V_L)\\ \mathcal{N}(X^\top) &= \mathcal{C}(U_R)\end{aligned}$$

Thus, taken together, the systems of equations (5.3) and (5.5) describe the four fundamental subspaces of X in the following ways:

$$\begin{aligned}\mathcal{C}(X) &= \mathcal{C}(U_L)\\ \mathcal{N}(X^\top) &= \mathcal{C}(U_R)\\ \mathcal{R}(X) \equiv \mathcal{C}(X^\top) &= \mathcal{C}(V_L)\\ \mathcal{N}(X) &= \mathcal{C}(V_R)\end{aligned} \tag{5.7}$$

Since U and V are both orthonormal matrices, collection (5.7) asserts that

- U_L is an orthonormal basis for the column space of X
- U_R is an orthonormal basis for the null space of $X^\top$
- V_L is an orthonormal basis for the row space of X
- V_R is an orthonormal basis for the null space of X

We have verified the four claims in (5.7) simply by performing the multiplications called for by the right side of (5.2) and reading them.

The claims in (5.7) and the fact that U and V are both unitary (i.e., orthonormal) matrices imply that

- the column space of X is orthogonal to the null space of $X^\top$
- the null space of X is orthogonal to the row space of X

Sometimes these properties are described with the following two pairs of orthogonal complement subspaces:

- $\mathcal{C}(X)$ is the orthogonal complement of $\mathcal{N}(X^\top)$

- $\mathcal{R}(X)$ is the orthogonal complement of $\mathcal{N}(X)$

Let's do an example.

```
import numpy as np
import numpy.linalg as LA
import matplotlib.pyplot as plt

rng = np.random.default_rng(1234)
```

Having imported these modules, let's do the example.

We use a 4×5 matrix of rank 2, so that the four subspaces have dimensions 2, 2, 2, and 3.

Note that `np.linalg.svd` returns $V^\top$, not V, as its third output.

So the row space and the null space are spanned by rows of that output.

```
np.set_printoptions(precision=2, suppress=True)

# Define the matrix
A = np.array([[1, 2, 3, 4, 5],
              [2, 3, 4, 5, 6],
              [3, 4, 5, 6, 7],
              [4, 5, 6, 7, 8]])

# Compute the SVD of the matrix; the third output is V^T
U, S, VT = np.linalg.svd(A, full_matrices=True)

# Compute the rank of the matrix
rank = np.linalg.matrix_rank(A)

print("Rank of matrix:\n", rank)
print("S: \n", S)

# Orthonormal bases for the four fundamental subspaces
col_space = U[:, :rank]            # C(A),   a subspace of R^4
left_null_space = U[:, rank:]      # N(A^T), a subspace of R^4
row_space = VT[:rank, :].T         # R(A),   a subspace of R^5
null_space = VT[rank:, :].T        # N(A),   a subspace of R^5

print("Column space:\n", col_space)
print("Left null space:\n", left_null_space)
print("Row space:\n", row_space)
print("Null space:\n", null_space)
```

```
Rank of matrix:
 2
S:
 [21.63  1.46  0.    0.  ]
Column space:
 [[ 0.34  0.76]
 [ 0.44  0.33]
 [ 0.54 -0.11]
 [ 0.64 -0.54]]
Left null space:
 [[-0.27  0.48]
 [ 0.71 -0.44]
 [-0.62 -0.56]
```

(continues on next page)

(continued from previous page)

```
 [ 0.18  0.52]]
Row space:
 [[ 0.25 -0.73]
 [ 0.34 -0.43]
 [ 0.43 -0.13]
 [ 0.52  0.18]
 [ 0.61  0.48]]
Null space:
 [[ 0.62 -0.08  0.08]
 [-0.71 -0.05  0.45]
 [-0.24 -0.18 -0.84]
 [ 0.11  0.83  0.03]
 [ 0.21 -0.52  0.29]]
```

Let's verify that the bases have the properties that (5.7) asserts.

```
print("A @ null_space = 0:         ", np.allclose(A @ null_space, 0))
print("A.T @ left_null_space = 0: ", np.allclose(A.T @ left_null_space, 0))
print("col_space ⊥ left_null_space:", np.allclose(col_space.T @ left_null_space, 0))
print("row_space ⊥ null_space:     ", np.allclose(row_space.T @ null_space, 0))
```

```
A @ null_space = 0:         True
A.T @ left_null_space = 0:  True
col_space ⊥ left_null_space: True
row_space ⊥ null_space:     True
```

5.5 Eckart-Young theorem

Suppose that we want to construct the best rank r approximation of an $m \times n$ matrix X.

By best, we mean a matrix X_r of rank $r < p$ that, among all rank r matrices of dimension $m \times n$, minimizes

$$||X - X_r||$$

where $||\cdot||$ denotes a norm of a matrix.

Three popular **matrix norms** of an $m \times n$ matrix X can be expressed in terms of the singular values of X

- the **spectral** or l^2 norm $||X||_2 = \max_{||y||\neq 0} \frac{||Xy||}{||y||} = \sigma_1$
- the **Frobenius** norm $||X||_F = \sqrt{\sigma_1^2 + \cdots + \sigma_p^2}$
- the **nuclear** norm $||X||_N = \sigma_1 + \cdots + \sigma_p$

The Eckart-Young theorem states that for each of these three norms, the same rank r matrix is best and that it equals

$$\hat{X}_r = \sigma_1 u_1 v_1^\top + \sigma_2 u_2 v_2^\top + \cdots + \sigma_r u_r v_r^\top = U_r \Sigma_r V_r^\top \tag{5.8}$$

where U_r and V_r consist of the first r columns of U and V, respectively, and Σ_r is the $r \times r$ diagonal matrix of the r largest singular values.

The resulting approximation errors are

$$\begin{aligned}
||X - \hat{X}_r||_2 &= \sigma_{r+1} \\
||X - \hat{X}_r||_F &= \sqrt{\sigma_{r+1}^2 + \cdots + \sigma_p^2} \\
||X - \hat{X}_r||_N &= \sigma_{r+1} + \cdots + \sigma_p
\end{aligned}$$

This is a very powerful theorem.

It says that we can approximate an $m \times n$ matrix X of rank p by an $m \times n$ matrix of rank $r < p$ built from the r largest singular values of X and their singular vectors, and that the singular values that we discard measure the approximation error.

If a few singular values are much larger than the others, a low-rank approximation captures most of X while requiring us to store only $r(m + n + 1)$ numbers instead of mn.

You can read about the Eckart-Young theorem and some of its uses here.

We'll make use of this theorem when we discuss principal components analysis (PCA) below.

It also underlies the sequel *VARs and DMDs*, where the pseudo-inverse

$$\hat{X}_r^+ = V_r \Sigma_r^{-1} U_r^\top$$

of the approximation $\hat{X}_r$ is a key ingredient of a dynamic mode decomposition.

Exercise 5.12.2 and Exercise 5.12.3 explore these ideas.

5.6 Full and reduced SVD's

Up to now we have described properties of a **full** SVD in which shapes of U, Σ, and V are (m, m), (m, n), (n, n), respectively.

There is an alternative bookkeeping convention called an **economy** or **reduced** SVD in which the shapes of U, Σ and V are different from what they are in a full SVD.

Because we assume that X has rank p, there are only p nonzero singular values, where $p = \text{rank}(X) \leq \min(m, n)$.

A **reduced** SVD uses this fact to express U, Σ, and V as matrices with shapes (m, p), (p, p), (n, p).

You can read about reduced and full SVD here https://numpy.org/doc/stable/reference/generated/numpy.linalg.svd.html

(With the option `full_matrices=False`, `numpy` returns $\min(m, n)$ columns of U and V; when $p < \min(m, n)$, we drop the columns that correspond to zero singular values to obtain a reduced SVD in the sense used here.)

For a full SVD,

$$\begin{aligned} UU^\top = I \qquad & U^\top U = I \\ VV^\top = I \qquad & V^\top V = I \end{aligned}$$

But not all these properties hold for a **reduced** SVD.

For a reduced SVD, $U^\top U = I_{p \times p}$ and $V^\top V = I_{p \times p}$ always hold, but $UU^\top = I_{m \times m}$ holds only if $p = m$ and $VV^\top = I_{n \times n}$ holds only if $p = n$.

Which properties hold thus depends on whether we are in a **tall-skinny** case or a **short-fat** case.

- In a **tall-skinny** case in which $m \gg n$ and X has full column rank $p = n$, for a **reduced** SVD

$$\begin{aligned} UU^\top \neq I \qquad & U^\top U = I \\ VV^\top = I \qquad & V^\top V = I \end{aligned}$$

- In a **short-fat** case in which $m \ll n$ and X has full row rank $p = m$, for a **reduced** SVD

$$\begin{aligned} UU^\top = I \qquad & U^\top U = I \\ VV^\top \neq I \qquad & V^\top V = I \end{aligned}$$

When we study dynamic mode decompositions in the sequel *VARs and DMDs*, we shall want to remember these properties because we will use a reduced SVD in a tall-skinny case.

Let's do an example to compare **full** and **reduced** SVD's.

To review,

- in a **full** SVD
  - U is $m \times m$
  - Σ is $m \times n$
  - V is $n \times n$
- in a **reduced** SVD
  - U is $m \times p$
  - Σ is $p \times p$
  - V is $n \times p$

First, let's study a case in which $m = 5 > n = 2$.

(This is a small example of the **tall-skinny** case that will concern us when we study dynamic mode decompositions in the sequel.)

```
X = rng.random((5, 2))
U, S, V = np.linalg.svd(X, full_matrices=True)  # full SVD
Uhat, Shat, Vhat = np.linalg.svd(X, full_matrices=False) # economy SVD
print('U, S, V =')
U, S, V
```

```
U, S, V =
```

```
(array([[-0.58,  0.17, -0.2 , -0.67, -0.38],
        [-0.53, -0.23, -0.05,  0.68, -0.45],
        [-0.19,  0.03,  0.98, -0.09, -0.04],
        [-0.18,  0.92, -0.03,  0.29,  0.19],
        [-0.55, -0.27, -0.07, -0.01,  0.78]]),
 array([1.8 , 0.24]),
 array([[-0.95, -0.33],
        [-0.33,  0.95]]))
```

```
print('Uhat, Shat, Vhat = ')
Uhat, Shat, Vhat
```

```
Uhat, Shat, Vhat =
```

```
(array([[-0.58,  0.17],
        [-0.53, -0.23],
        [-0.19,  0.03],
        [-0.18,  0.92],
        [-0.55, -0.27]]),
 array([1.8 , 0.24]),
 array([[-0.95, -0.33],
        [-0.33,  0.95]]))
```

```
rr = np.linalg.matrix_rank(X)
print(f'rank of X = {rr}')
```

```
rank of X = 2
```

Properties:

- Where U is constructed via a full SVD, $U^\top U = I_{m\times m}$ and $UU^\top = I_{m\times m}$
- Where $\hat{U}$ is constructed via a reduced SVD, although $\hat{U}^\top\hat{U} = I_{p\times p}$, it happens that $\hat{U}\hat{U}^\top \neq I_{m\times m}$

We illustrate these properties for our example with the following code cells.

```
UTU = U.T@U
UUT = U@U.T
print('UUT, UTU = ')
UUT, UTU
```

```
UUT, UTU =
```

```
(array([[ 1.,  0., -0., -0., -0.],
        [ 0.,  1., -0.,  0., -0.],
        [-0., -0.,  1., -0., -0.],
        [-0.,  0., -0.,  1.,  0.],
        [-0., -0., -0.,  0.,  1.]]),
 array([[ 1., -0., -0.,  0., -0.],
        [-0.,  1., -0.,  0., -0.],
        [-0., -0.,  1.,  0., -0.],
        [ 0.,  0.,  0.,  1.,  0.],
        [-0., -0., -0.,  0.,  1.]]))
```

```
UhatUhatT = Uhat@Uhat.T
UhatTUhat = Uhat.T@Uhat
print('UhatUhatT, UhatTUhat= ')
UhatUhatT, UhatTUhat
```

```
UhatUhatT, UhatTUhat=
```

```
(array([[ 0.37,  0.27,  0.12,  0.26,  0.28],
        [ 0.27,  0.33,  0.09, -0.11,  0.36],
        [ 0.12,  0.09,  0.04,  0.06,  0.1 ],
        [ 0.26, -0.11,  0.06,  0.88, -0.15],
        [ 0.28,  0.36,  0.1 , -0.15,  0.38]]),
 array([[ 1., -0.],
        [-0.,  1.]]))
```

Remarks:

The cells above illustrate the application of the `full_matrices=True` and `full_matrices=False` options.

Using `full_matrices=False` returns a reduced singular value decomposition.

The **full** and **reduced** SVD's both accurately decompose an $m \times n$ matrix X.

When we study dynamic mode decompositions in the sequel, it will be important for us to remember the preceding properties of full and reduced SVD's in such tall-skinny cases.

Now let's turn to a short-fat case.

To illustrate this case, we'll set $m = 2 < 5 = n$ and compute both full and reduced SVD's.

```
X = rng.random((2, 5))
U, S, V = np.linalg.svd(X, full_matrices=True)  # full SVD
Uhat, Shat, Vhat = np.linalg.svd(X, full_matrices=False) # economy SVD
print('U, S, V = ')
U, S, V
```

```
U, S, V =
```

```
(array([[ 0.78, -0.63],
        [ 0.63,  0.78]]),
 array([1.97, 0.68]),
 array([[ 0.39,  0.48,  0.41,  0.4 ,  0.55],
        [ 0.35,  0.28, -0.55, -0.61,  0.37],
        [-0.52,  0.45,  0.49, -0.54, -0.  ],
        [-0.52,  0.51, -0.54,  0.43,  0.02],
        [-0.44, -0.49, -0.02, -0.  ,  0.75]]))
```

```
print('Uhat, Shat, Vhat = ')
Uhat, Shat, Vhat
```

```
Uhat, Shat, Vhat =
```

```
(array([[ 0.78, -0.63],
        [ 0.63,  0.78]]),
 array([1.97, 0.68]),
 array([[ 0.39,  0.48,  0.41,  0.4 ,  0.55],
        [ 0.35,  0.28, -0.55, -0.61,  0.37]]))
```

Let's verify that our reduced SVD accurately represents X

```
SShat=np.diag(Shat)
np.allclose(X, Uhat@SShat@Vhat)
```

```
True
```

5.7 Polar decomposition

A **reduced** singular value decomposition (SVD) of X is related to a **polar decomposition** of X

$$X = SQ$$

where

$$S = U\Sigma U^\top$$
$$Q = UV^\top$$

and in our reduced SVD

- U is an $m \times p$ orthonormal matrix
- Σ is a $p \times p$ diagonal matrix
- V is an $n \times p$ orthonormal matrix

Because $U^\top U = I$, we have $SQ = U\Sigma U^\top UV^\top = U\Sigma V^\top = X$.

Here

- S is an $m \times m$ **symmetric positive semidefinite** matrix
- Q is an $m \times n$ matrix with $QQ^\top = UU^\top$; so when $p = m$, as for a short-fat X with full row rank, $QQ^\top = I$ and Q has orthonormal rows

5.8 Application: principal components analysis (PCA)

Let's begin with a case in which $n \gg m$, so that we have many more individuals n than attributes m.

The matrix X is **short and fat** in an $n \gg m$ case as opposed to a **tall and skinny** case with $m \gg n$ to be discussed in the sequel.

We regard X as an $m \times n$ matrix of **data**:

$$X = \begin{bmatrix} X_1 \mid X_2 \mid \cdots \mid X_n \end{bmatrix}$$

where for $j = 1, \ldots, n$ the column vector $X_j = \begin{bmatrix} x_{1j} \\ x_{2j} \\ \vdots \\ x_{mj} \end{bmatrix}$ is a vector of observations on variables $1, 2, \ldots, m$.

In a **time series** setting, we would think of columns j as indexing different **times** at which random variables are observed, while rows index different random variables.

In a **cross-section** setting, we would think of columns j as indexing different **individuals** for which random variables are observed, while rows index different **attributes**.

As we have seen before, the SVD is a way to decompose a matrix into useful components, just like polar decomposition, eigendecomposition, and many others.

PCA, on the other hand, is a method that builds on the SVD to analyze data.

Its goal is to find a few linear combinations of the variables that capture the most important patterns of variation in the data.

Step 1: Center the data:

Because we are interested in variation of the data around their means, we first subtract sample means.

(Because our data matrix may hold variables of different units and scales, we might also **standardize** the data by dividing each row by its sample standard deviation.)

We first compute the average of each row of X

$$\bar{X}_i = \frac{1}{n} \sum_{j=1}^{n} x_{ij}$$

We then create a matrix of these means:

$$\bar{X} = \begin{bmatrix} \bar{X}_1 \\ \bar{X}_2 \\ \vdots \\ \bar{X}_m \end{bmatrix} \begin{bmatrix} 1 \mid 1 \mid \cdots \mid 1 \end{bmatrix}$$

and subtract it from the original matrix to create a mean-centered matrix:

$$B = X - \bar{X}$$

Step 2: Compute the covariance matrix:

Because we want to extract relationships between variables rather than just their magnitudes – in other words, we want to know how they can explain each other – we compute the sample covariance matrix of B.

$$C = \frac{1}{n} BB^\top$$

Step 3: Decompose the covariance matrix and arrange the singular values:

Since the matrix C is symmetric and positive semidefinite, we can eigendecompose it, find its eigenvalues, and arrange the eigenvalue and eigenvector matrices in decreasing order.

The eigendecomposition of C can be found by decomposing B instead. Since B is not a square matrix, we obtain an SVD of B:

$$\begin{aligned} BB^\top &= U\Sigma V^\top (U\Sigma V^\top)^\top \\ &= U\Sigma V^\top V \Sigma^\top U^\top \\ &= U\Sigma\Sigma^\top U^\top \end{aligned}$$

so that

$$C = \frac{1}{n} U\Sigma\Sigma^\top U^\top$$

Singular values are conventionally arranged in decreasing order, as `numpy` does.

Step 4: Select singular values, (optional) truncate the rest:

We can now decide how many singular values to keep, based on how much variance we want to retain (e.g., retaining 95% of the total variance).

We can obtain the percentage by calculating the variance contained in the leading r factors divided by the total variance:

$$\frac{\sum_{i=1}^{r} \sigma_i^2}{\sum_{i=1}^{p} \sigma_i^2}$$

By the Eckart-Young theorem, one minus this ratio is also the squared relative Frobenius-norm error $||B - \hat{B}_r||_F^2 / ||B||_F^2$ of the best rank r approximation $\hat{B}_r$ of B.

Step 5: Create the score matrix:

The matrix of **principal components** (or **scores**) is

$$\begin{aligned} T &= U^\top B \\ &= U^\top U \Sigma V^\top \\ &= \Sigma V^\top \end{aligned}$$

Its kth row is the sequence, across observations $j = 1, \ldots, n$, of the kth principal component.

Keeping only the first r rows gives the $r \times n$ matrix $T_r = U_r^\top B$ of the first r principal components, and $U_r T_r = U_r \Sigma_r V_r^\top = \hat{B}_r$ is the Eckart-Young approximation of B.

In the sequel *VARs and DMDs*, the vector $U_r^\top X_t$ of the first r principal components of a time t cross section serves as the regressor in a **principal components regression** that forecasts next period's cross section.

5.9 Relationship of PCA to SVD

To relate an SVD to a PCA of data set X, first construct the SVD of the data matrix X.

Let's assume that sample means of all variables are zero, so we don't need to center our matrix.

$$X = U\Sigma V^\top = \sigma_1 U_1 V_1^\top + \sigma_2 U_2 V_2^\top + \cdots + \sigma_p U_p V_p^\top \tag{5.9}$$

where

$$U = [U_1|U_2|...|U_m]$$

$$V^\top = \begin{bmatrix} V_1^\top \\ V_2^\top \\ \vdots \\ V_n^\top \end{bmatrix}$$

In equation (5.9), each of the $m \times n$ matrices $U_j V_j^\top$ is evidently of rank 1.

Thus, we have

$$X = \sigma_1 \begin{bmatrix} U_{11}V_1^\top \\ U_{21}V_1^\top \\ \cdots \\ U_{m1}V_1^\top \end{bmatrix} + \sigma_2 \begin{bmatrix} U_{12}V_2^\top \\ U_{22}V_2^\top \\ \cdots \\ U_{m2}V_2^\top \end{bmatrix} + ... + \sigma_p \begin{bmatrix} U_{1p}V_p^\top \\ U_{2p}V_p^\top \\ \cdots \\ U_{mp}V_p^\top \end{bmatrix} \tag{5.10}$$

Here is how we would interpret the objects in the matrix equation (5.10) in a time series context:

- for each $k = 1, ..., p$, the object $\{V_{jk}\}_{j=1}^n$ (the kth column of V) is a time series for the kth **principal component**, normalized to have unit length; the unnormalized principal component is $\sigma_k V_k^\top = U_k^\top X$
- $U_k = \begin{bmatrix} U_{1k} \\ U_{2k} \\ \vdots \\ U_{mk} \end{bmatrix}$, $k = 1, ..., p$, is a vector of **loadings** of variables X_i on the kth principal component, $i = 1, ..., m$
- σ_k for each $k = 1, ..., p$ is the strength of the kth **principal component**, where strength means contribution to the overall covariance of X: the kth principal component contributes σ_k^2/n to the sum of the sample variances of the m variables

5.10 PCA with eigenvalues and eigenvectors

We now use an eigen decomposition of a sample covariance matrix to do PCA.

Let $X_{m\times n}$ be our $m \times n$ data matrix.

Let's assume that sample means of all variables are zero.

We can assure this by **pre-processing** the data by subtracting sample means.

Define a matrix Ω that is proportional to the sample covariance matrix (it equals n times it) as

$$\Omega = XX^\top$$

Then use an eigen decomposition to represent Ω as follows:

$$\Omega = P\Lambda P^\top$$

Here

- P is an $m \times m$ orthogonal matrix of eigenvectors of Ω
- Λ is a diagonal matrix of eigenvalues of Ω

We can then represent X as

$$X = P\epsilon$$

where

$$\epsilon = P^{-1}X = P^{\top}X$$

and

$$\epsilon\epsilon^{\top} = \Lambda.$$

We can verify that

$$XX^{\top} = P\Lambda P^{\top}. \tag{5.11}$$

It follows that we can represent the data matrix X as

$$X = [P_1|P_2|...|P_m] \begin{bmatrix} \epsilon_1 \\ \epsilon_2 \\ \vdots \\ \epsilon_m \end{bmatrix} = P_1\epsilon_1 + P_2\epsilon_2 + ... + P_m\epsilon_m$$

where P_j is the jth column of P and ϵ_j is the $1 \times n$ jth row of ϵ.

To reconcile the preceding representation with the PCA that we had obtained earlier through the SVD, we first note that $\epsilon_j\epsilon_j^{\top} = \lambda_j \equiv \sigma_j^2$.

Now for each j with $\lambda_j > 0$ define $\tilde{\epsilon}_j = \frac{\epsilon_j}{\sqrt{\lambda_j}}$, which implies that $\tilde{\epsilon}_j\tilde{\epsilon}_j^{\top} = 1$.

Terms with $\lambda_j = 0$ vanish because then $\epsilon_j = 0$.

Therefore

$$\begin{aligned} X &= \sqrt{\lambda_1}P_1\tilde{\epsilon}_1 + \sqrt{\lambda_2}P_2\tilde{\epsilon}_2 + ... + \sqrt{\lambda_p}P_p\tilde{\epsilon}_p \\ &= \sigma_1P_1\tilde{\epsilon}_1 + \sigma_2P_2\tilde{\epsilon}_2 + ... + \sigma_pP_p\tilde{\epsilon}_p, \end{aligned}$$

which agrees with

$$X = \sigma_1U_1{V_1}^{\top} + \sigma_2U_2{V_2}^{\top} + ... + \sigma_pU_p{V_p}^{\top}$$

provided that we set

- $U_j = P_j$ (a vector of loadings of variables on principal component j)
- ${V_k}^{\top} = \tilde{\epsilon}_k$ (the kth principal component)

Because there are alternative algorithms for computing P and U for a given data matrix X, depending on the algorithms used, we might have sign differences or different orders of eigenvectors.

We can resolve such ambiguities about U and P by

1. sorting eigenvalues and singular values in descending order
2. imposing positive diagonals on P and U and adjusting signs in $V^{\top}$ accordingly

5.11 Connections

To pull things together, it is useful to assemble and compare some formulas presented above.

First, consider an SVD of an $m \times n$ matrix:

$$X = U\Sigma V^\top$$

Compute:

$$\begin{aligned} XX^\top &= U\Sigma V^\top V\Sigma^\top U^\top \\ &= U\Sigma\Sigma^\top U^\top \\ &\equiv U\Lambda U^\top \end{aligned} \tag{5.12}$$

Compare representation (5.12) with equation (5.11) above.

Evidently, U in the SVD is the matrix P of eigenvectors of $XX^\top$ and $\Sigma\Sigma^\top$ is the matrix Λ of eigenvalues.

Second, let's compute

$$\begin{aligned} X^\top X &= V\Sigma^\top U^\top U\Sigma V^\top \\ &= V\Sigma^\top\Sigma V^\top \end{aligned}$$

Thus, the matrix V in the SVD is the matrix of eigenvectors of $X^\top X$.

Summarizing and fitting things together, we have the eigen decomposition of the sample covariance matrix

$$XX^\top = P\Lambda P^\top$$

where P is an orthogonal matrix.

Further, from the SVD of X, we know that

$$XX^\top = U\Sigma\Sigma^\top U^\top$$

where U is an orthogonal matrix.

Thus, $P = U$ (up to the ordering and signs of columns) and we have the representation of X

$$X = P\epsilon = U\Sigma V^\top$$

It follows that

$$U^\top X = \Sigma V^\top = \epsilon$$

so that ϵ is the matrix of principal components (scores) that we constructed in Step 5 above.

Note that the preceding implies that

$$\epsilon\epsilon^\top = \Sigma V^\top V\Sigma^\top = \Sigma\Sigma^\top = \Lambda,$$

so that everything fits together.

Below we define a class `DecomAnalysis` that wraps PCA and SVD for a given data matrix `X`.

```
class DecomAnalysis:
    """
    A class for conducting PCA and SVD.
    X: data matrix
    r_component: chosen rank for best approximation
    """

    def __init__(self, X, r_component=None):

        self.X = X

        self.Ω = (X @ X.T)

        self.m, self.n = X.shape
        self.r = LA.matrix_rank(X)

        if r_component:
            self.r_component = r_component
        else:
            self.r_component = min(self.m, self.n)

    def pca(self):

        λ, P = LA.eigh(self.Ω)     # columns of P are eigenvectors

        ind = sorted(range(λ.size), key=lambda x: λ[x], reverse=True)

        # sort by eigenvalues
        self.λ = λ[ind]
        P = P[:, ind]
        self.P = P @ diag_sign(P)

        self.Λ = np.diag(self.λ)

        self.explained_ratio_pca = np.cumsum(self.λ) / self.λ.sum()

        # compute the m by n matrix of principal components
        self.ϵ = self.P.T @ self.X

        P = self.P[:, :self.r_component]
        ϵ = self.ϵ[:self.r_component, :]

        # transform data
        self.X_pca = P @ ϵ

    def svd(self):

        U, σ, VT = LA.svd(self.X)

        ind = sorted(range(σ.size), key=lambda x: σ[x], reverse=True)

        # sort by eigenvalues
        d = min(self.m, self.n)

        self.σ = σ[ind]
        U = U[:, ind]
        D = diag_sign(U)
```

(continues on next page)

(continued from previous page)

```
        self.U = U @ D
        VT[:d, :] = D @ VT[ind, :]
        self.VT = VT

        self.Σ = np.zeros((self.m, self.n))
        self.Σ[:d, :d] = np.diag(self.σ)

        σ_sq = self.σ ** 2
        self.explained_ratio_svd = np.cumsum(σ_sq) / σ_sq.sum()

        # slicing matrices by the number of components to use
        U = self.U[:, :self.r_component]
        Σ = self.Σ[:self.r_component, :self.r_component]
        VT = self.VT[:self.r_component, :]

        # transform data
        self.X_svd = U @ Σ @ VT

    def fit(self, r_component):

        # pca
        P = self.P[:, :r_component]
        ϵ = self.ϵ[:r_component, :]

        # transform data
        self.X_pca = P @ ϵ

        # svd
        U = self.U[:, :r_component]
        Σ = self.Σ[:r_component, :r_component]
        VT = self.VT[:r_component, :]

        # transform data
        self.X_svd = U @ Σ @ VT

def diag_sign(A):
    "Compute the signs of the diagonal of matrix A"

    D = np.diag(np.sign(np.diag(A)))

    return D
```

We also define a function that prints out information so that we can compare decompositions obtained by different algorithms.

```
def compare_pca_svd(da):
    """
    Compare the outcomes of PCA and SVD.
    """

    da.pca()
    da.svd()

    print('Eigenvalues and Singular values\n')
    print(f'λ = {da.λ}\n')
    print(f'σ^2 = {da.σ**2}\n')
    print('\n')
```

(continues on next page)

(continued from previous page)

```
    k = da.r

    # loading matrices
    fig, axs = plt.subplots(1, 2, figsize=(14, 5))
    plt.suptitle('loadings')
    axs[0].plot(da.P[:, :k])
    axs[0].set_title('P')
    axs[0].set_xlabel('variable $i$')
    axs[1].plot(da.U[:, :k])
    axs[1].set_title('U')
    axs[1].set_xlabel('variable $i$')
    plt.show()

    # principal components
    fig, axs = plt.subplots(1, 2, figsize=(14, 5))
    plt.suptitle('principal components')
    axs[0].plot(da.ε[:k, :].T)
    axs[0].set_title('ε')
    axs[0].set_xlabel('n')
    axs[1].plot(da.VT[:k, :].T * np.sqrt(da.λ[:k]))
    axs[1].set_title(r'$V^\top *\sqrt{\lambda}$')
    axs[1].set_xlabel('n')
    plt.show()
```

Let's apply these tools to a short-fat data matrix with $m = 5$ variables and $n = 200$ observations that is generated by two factors plus noise.

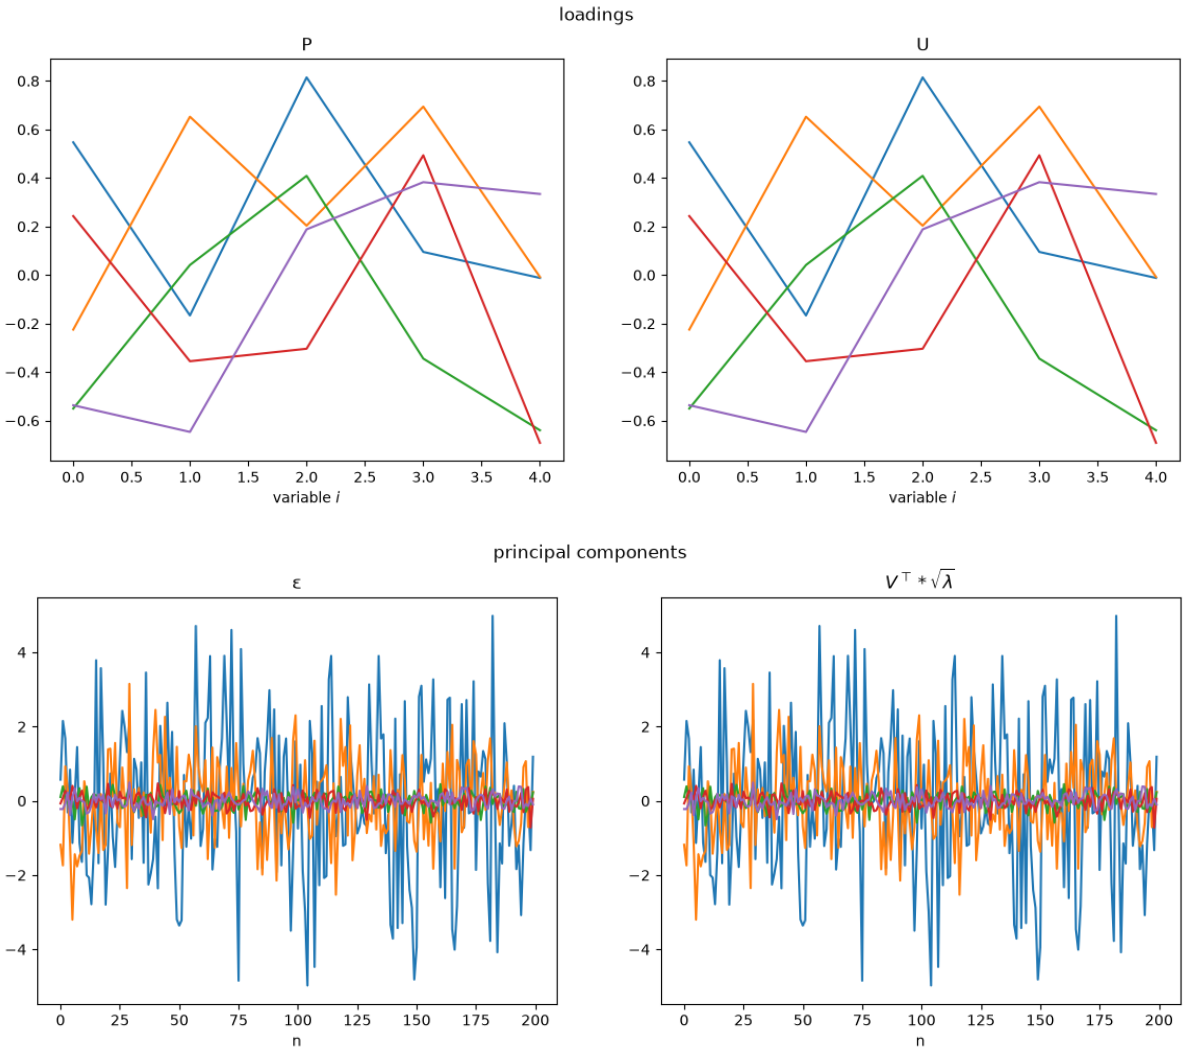
```
m, n = 5, 200
G = rng.normal(size=(m, 2))               # loadings on two factors
Z = rng.normal(size=(2, n))               # factors
X = G @ Z + 0.2 * rng.normal(size=(m, n))
X = X - X.mean(axis=1, keepdims=True)     # subtract sample means

da = DecomAnalysis(X)
compare_pca_svd(da)
```

```
Eigenvalues and Singular values

λ = [884.59 233.75   9.36   8.13   6.41]

σ^2 = [884.59 233.75   9.36   8.13   6.41]
```



The two algorithms produce the same loadings and the same principal components.

The cumulative fractions of variance explained reveal that two principal components account for almost all of the variation in X, which is consistent with how we generated the data.

```
print("explained ratio (PCA):", da.explained_ratio_pca)
print("explained ratio (SVD):", da.explained_ratio_svd)
```

```
explained ratio (PCA): [0.77 0.98 0.99 0.99 1.  ]
explained ratio (SVD): [0.77 0.98 0.99 0.99 1.  ]
```

For an example of PCA applied to analyzing the structure of intelligence tests see the lecture *Multivariate Normal Distribution*.

Look at parts of that lecture that describe and illustrate the classic factor analysis model.

As mentioned earlier, in a sequel to this lecture, *VARs and DMDs*, we'll describe how SVD's provide ways rapidly to compute reduced-order approximations to first-order Vector Autoregressions (VARs).

5.12 Exercises

Exercise 5.12.1

In ordinary least squares (OLS), we learn to compute $\hat{\beta} = (X^\top X)^{-1} X^\top y$, where the matrix X of regressors has one row for each observation and one column for each regressor.

But the matrix $X^\top X$ is not invertible (its determinant is zero) when regressors are perfectly collinear or when there are more regressors than observations, so that X is **short and fat** and the system $X\beta = y$ is underdetermined.

And $X^\top X$ is **ill-conditioned** when regressors are nearly collinear, meaning that the ratio of its largest to its smallest eigenvalue is very large.

What we can do instead is to use what is called a pseudoinverse X^+ of X in place of $(X^\top X)^{-1} X^\top$.

1. Use an SVD of X to build the pseudoinverse matrix X^+ and use it to compute $\hat{\beta} = X^+ y$.
2. Show that when X has linearly independent columns (an **over-determined** system), $X^+ = (X^\top X)^{-1} X^\top$, so that $\hat{\beta}$ is the usual OLS estimator.
3. Show that when X has linearly independent rows (an **under-determined** system), $\hat{\beta}$ solves $X\beta = y$ exactly and has the smallest norm among all solutions.
4. Verify these claims numerically for a 5×8 matrix X of random numbers.

Solution

We can use an SVD to compute the pseudoinverse.

Starting from a reduced SVD

$$X = U \Sigma V^\top$$

in which Σ is the $p \times p$ diagonal matrix of positive singular values, the pseudoinverse is

$$X^+ = V \Sigma^{-1} U^\top$$

Equivalently, in terms of a full SVD, $X^+ = V \Sigma^+ U^\top$, where the $n \times m$ matrix Σ^+ is formed by transposing Σ and replacing each positive singular value σ_j by $1/\sigma_j$:

$$\Sigma^+ = \begin{bmatrix} \frac{1}{\sigma_1} & 0 & \cdots & 0 & 0 \\ 0 & \frac{1}{\sigma_2} & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \frac{1}{\sigma_p} & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$$

and finally:

$$\hat{\beta} = X^+ y = V \Sigma^{-1} U^\top y$$

When X has linearly independent columns, V is a square orthogonal matrix, so

$$(X^\top X)^{-1} X^\top = (V \Sigma^2 V^\top)^{-1} V \Sigma U^\top = V \Sigma^{-2} V^\top V \Sigma U^\top = V \Sigma^{-1} U^\top = X^+.$$

When X has linearly independent rows, U is a square orthogonal matrix, so

$$X \hat{\beta} = U \Sigma V^\top V \Sigma^{-1} U^\top y = U U^\top y = y.$$

Every other solution of $X\beta = y$ has the form $\beta = \hat{\beta} + \eta$ with $X\eta = 0$.

The vector $\hat{\beta} = V(\Sigma^{-1}U^\top y)$ lies in the row space of X, which is spanned by the columns of V, while η lies in the null space of X.

These two subspaces are orthogonal, so $\|\beta\|^2 = \|\hat{\beta}\|^2 + \|\eta\|^2 \geq \|\hat{\beta}\|^2$.

Here is a numerical check.

```
n_obs, k = 5, 8
X = rng.normal(size=(n_obs, k))
y = rng.normal(size=n_obs)

U, σ, VT = np.linalg.svd(X, full_matrices=False)
X_plus = VT.T @ np.diag(1 / σ) @ U.T
β_hat = X_plus @ y

print("X_plus equals np.linalg.pinv(X):", np.allclose(X_plus, np.linalg.pinv(X)))
print("X β_hat = y:                     ", np.allclose(X @ β_hat, y))

# another exact solution: add an element of the null space of X
β_other = β_hat + (np.eye(k) - X_plus @ X) @ rng.normal(size=k)
print("X β_other = y:                   ", np.allclose(X @ β_other, y))
print(f"norms: |β_hat| = {LA.norm(β_hat):.3f}, |β_other| = {LA.norm(β_other):.3f}")

# over-determined case: pseudo-inverse reproduces OLS
X_tall = rng.normal(size=(8, 5))
print("pinv equals (X^T X)^{-1} X^T:   ",
      np.allclose(np.linalg.pinv(X_tall), LA.inv(X_tall.T @ X_tall) @ X_tall.T))
```

```
X_plus equals np.linalg.pinv(X): True
X β_hat = y:                     True
X β_other = y:                   True
norms: |β_hat| = 1.099, |β_other| = 1.175
pinv equals (X^T X)^{-1} X^T:    True
```

Exercise 5.12.2

This exercise verifies the Eckart-Young theorem numerically.

Construct a 60×40 matrix X equal to a rank-4 matrix plus noise, for example

```
X = rng.normal(size=(60, 4)) @ rng.normal(size=(4, 40)) + 0.3 * rng.
normal(size=(60, 40))
```

For $r = 1, 2, 4, 8$

1. compute the Eckart-Young approximation $\hat{X}_r$ and verify the formulas for $||X - \hat{X}_r||$ in the spectral, Frobenius, and nuclear norms
2. compare $||X - \hat{X}_r||_F$ with $||X - QQ^\top X||_F$ for 500 matrices Q with r orthonormal columns that span randomly drawn subspaces (for each Q, $QQ^\top X$ is the best approximation of X whose columns lie in the column space of Q)

Finally, plot the fraction $\sum_{i \leq r} \sigma_i^2 / \sum_i \sigma_i^2$ against r and relate it to Step 4 of the PCA recipe above.

Solution

Here is one solution.

```
m, n = 60, 40
X = rng.normal(size=(m, 4)) @ rng.normal(size=(4, n)) + 0.3 * rng.normal(size=(m,␣
↪n))
U, σ, VT = np.linalg.svd(X, full_matrices=False)

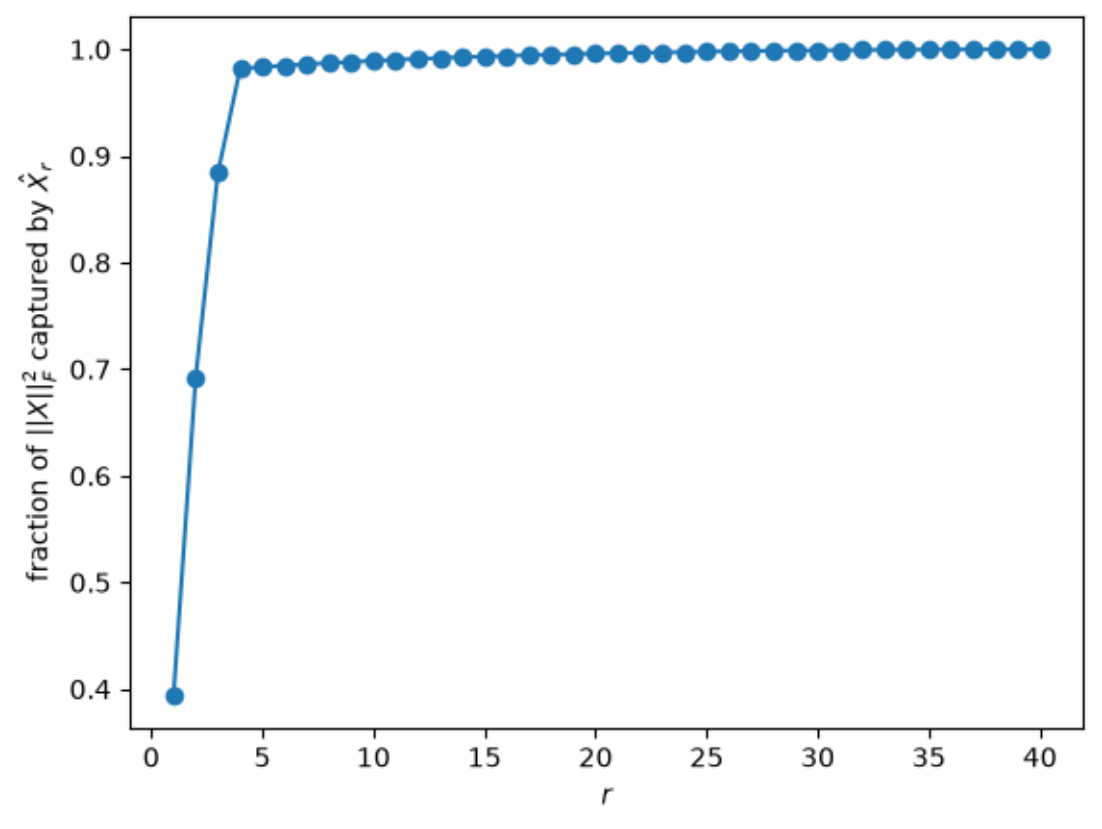
print(" r   spectral  Frobenius  nuclear   |X - X_r|_F   best random |X - QQ'X|_F")
for r in (1, 2, 4, 8):
    X_r = U[:, :r] @ np.diag(σ[:r]) @ VT[:r, :]
    E = X - X_r
    checks = (np.isclose(LA.norm(E, 2), σ[r]),
              np.isclose(LA.norm(E, 'fro'), np.sqrt(np.sum(σ[r:]**2))),
              np.isclose(LA.norm(E, 'nuc'), np.sum(σ[r:])))
    best_random = np.inf
    for _ in range(500):
        Q = LA.qr(rng.normal(size=(m, r)))[0]
        best_random = min(best_random, LA.norm(X - Q @ Q.T @ X, 'fro'))
    print(f"{r:2d}   {str(checks[0]):8s}  {str(checks[1]):9s}  {str(checks[2]):8s}"
          f"  {LA.norm(E, 'fro'):10.2f}   {best_random:10.2f}")
```

```
r   spectral  Frobenius  nuclear   |X - X_r|_F   best random |X - QQ'X|_F
1   True      True       True           77.46        96.00
2   True      True       True           55.16        94.62
4   True      True       True           13.41        91.33
8   True      True       True           11.27        84.75
```

All three error formulas hold, and no randomly drawn rank-r approximation comes close to the Eckart-Young approximation.

```
share = np.cumsum(σ**2) / np.sum(σ**2)

fig, ax = plt.subplots()
ax.plot(np.arange(1, len(σ) + 1), share, 'o-')
ax.set_xlabel('$r$')
ax.set_ylabel('fraction of $||X||_F^2$ captured by $\\hat X_r$')
plt.show()
```



The fraction rises steeply until $r = 4$, the rank of the signal part of X, and then flattens out.

This fraction is exactly the "fraction of variance retained" that Step 4 of the PCA recipe uses to choose how many principal components to keep, and one minus it is the squared relative Frobenius error of $\hat{X}_r$.

Exercise 5.12.3

This exercise builds a bridge to the sequel *VARs and DMDs*.

Simulate a **tall and skinny** time series data set $\tilde{X}$ with $m = 100$ variables and 61 periods in which

$$\tilde{X}_t = Gz_t + 0.5v_t, \qquad z_{t+1} = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.6 \end{bmatrix} z_t + w_{t+1}$$

where G is a 100×2 matrix of standard normal random numbers and v_t and w_{t+1} are standard normal random vectors.

Subtract the sample mean of each variable, then use the first 41 periods to form $X = \left[\tilde{X}_1 \mid \cdots \mid \tilde{X}_{40}\right]$ and $X' = \left[\tilde{X}_2 \mid \cdots \mid \tilde{X}_{41}\right]$.

1. Compute the least squares estimator $\hat{A} = X'X^+$ of the matrix A in the first-order VAR $\tilde{X}_{t+1} = A\tilde{X}_t + \epsilon_{t+1}$ and verify that it fits perfectly: $\hat{A}X = X'$.
2. Let $\hat{X}_2$ be the Eckart-Young rank-2 approximation of X. Verify numerically that $\hat{X}_2^+ = V_2\Sigma_2^{-1}U_2^\top$ and compute the rank-2 estimator $\hat{A}_2 = X'\hat{X}_2^+$.
3. Let $F = U_2^\top X$ be the 2×40 matrix of the first two principal components of X. Verify that $\hat{A}_2 =$

$X' F^\top (FF^\top)^{-1} U_2^\top$, i.e., that $\hat{A}_2 \tilde{X}_t$ equals the fitted value from a regression of $\tilde{X}_{t+1}$ on the first two principal components of $\tilde{X}_t$.

4. Use the remaining 20 periods to compare one-step-ahead out-of-sample forecast errors $\|X'_{\text{out}} - \check{A} X_{\text{out}}\|_F / \|X'_{\text{out}}\|_F$ for $\check{A} = \hat{A}$ and $\check{A} = \hat{A}_2$. Repeat for a few random seeds.

Solution

Here is one solution.

```
def simulate_tall(m=100, T=61, σ_v=0.5, seed=1):
    rng = np.random.default_rng(seed)
    G = rng.normal(size=(m, 2))
    a = np.array([0.9, 0.6])
    z = np.zeros((2, T))
    for t in range(T - 1):
        z[:, t+1] = a * z[:, t] + rng.normal(size=2)
    X_tilde = G @ z + σ_v * rng.normal(size=(m, T))
    return X_tilde - X_tilde.mean(axis=1, keepdims=True)

def rel_error(A, X, X_prime):
    return LA.norm(X_prime - A @ X) / LA.norm(X_prime)

X_tilde = simulate_tall()
X, X_prime = X_tilde[:, :40], X_tilde[:, 1:41]           # estimation sample
X_out, X_prime_out = X_tilde[:, 40:-1], X_tilde[:, 41:]  # hold-out sample

# 1. minimum-norm least squares estimator
A_hat = X_prime @ np.linalg.pinv(X)
print("A_hat X = X':", np.allclose(A_hat @ X, X_prime))

# 2. pseudo-inverse of the Eckart-Young approximation
U, σ, VT = np.linalg.svd(X, full_matrices=False)
r = 2
U_r, σ_r, V_r = U[:, :r], σ[:r], VT[:r, :].T
X_r = U_r @ np.diag(σ_r) @ V_r.T
X_r_plus = V_r @ np.diag(1 / σ_r) @ U_r.T
print("pinv(X_r) = V_r Σ_r^{-1} U_r^T:", np.allclose(np.linalg.pinv(X_r, rcond=1e-
↪10), X_r_plus))
A_r = X_prime @ X_r_plus

# 3. principal components regression
F = U_r.T @ X
A_pcr = X_prime @ F.T @ LA.inv(F @ F.T) @ U_r.T
print("A_r equals principal components regression:", np.allclose(A_r, A_pcr))

# 4. in-sample and out-of-sample fit
print(f"\nin-sample:     A_hat {rel_error(A_hat, X, X_prime):.3f},  A_r {rel_
↪error(A_r, X, X_prime):.3f}")
print(f"out-of-sample: A_hat {rel_error(A_hat, X_out, X_prime_out):.3f},  A_r {rel_
↪error(A_r, X_out, X_prime_out):.3f}")

print("\nout-of-sample errors for other seeds")
for seed in range(2, 7):
    X_tilde = simulate_tall(seed=seed)
    X, X_prime = X_tilde[:, :40], X_tilde[:, 1:41]
```

```
X_out, X_prime_out = X_tilde[:, 40:-1], X_tilde[:, 41:]
U, σ, VT = np.linalg.svd(X, full_matrices=False)
A_r = X_prime @ VT[:r, :].T @ np.diag(1 / σ[:r]) @ U[:, :r].T
A_hat = X_prime @ np.linalg.pinv(X)
print(f"seed {seed}:  A_hat {rel_error(A_hat, X_out, X_prime_out):.3f},"
      f"  A_r {rel_error(A_r, X_out, X_prime_out):.3f}")
```

```
A_hat X = X': True
pinv(X_r) = V_r Σ_r^{-1} U_r^T: True
A_r equals principal components regression: True

in-sample:     A_hat 0.000,  A_r 0.728
out-of-sample: A_hat 0.999,  A_r 0.757

out-of-sample errors for other seeds
seed 2:  A_hat 0.912,  A_r 0.790
seed 3:  A_hat 0.863,  A_r 0.701
seed 4:  A_hat 0.704,  A_r 0.680
seed 5:  A_hat 0.627,  A_r 0.530
seed 6:  A_hat 0.979,  A_r 0.789
```

The minimum-norm estimator $\hat{A}$ fits the estimation sample perfectly but forecasts poorly out of sample, because the pseudo-inverse $X^+ = V\Sigma^{-1}U^\top$ puts large weights $1/\sigma_j$ on directions with small singular values that mostly contain noise.

The rank-2 estimator $\hat{A}_2$ discards those directions and forecasts better out of sample.

The estimator $\hat{A}_2 = X'\hat{X}_2^+$ is exactly the estimator that a **dynamic mode decomposition** computes; the sequel *VARs and DMDs* studies it in detail.

CHAPTER

SIX

VARS AND DMDS

6.1 Overview

This lecture applies computational methods that we learned about in the lecture *Singular Value Decomposition* to

- first-order vector autoregressions (VARs)
- dynamic mode decompositions (DMDs)
- connections between DMDs and first-order VARs

[Sargent *et al.*, 2026] study these connections in detail.

We are especially interested in **tall and skinny** data sets in which the number of variables m exceeds the number of time periods n.

Such data sets are common.

For example, [Sargent *et al.*, 2026] study quarterly Consumer Expenditure Survey data on 100 quantiles of each of three cross sections – private income, post-tax-and-transfer income, and consumption – together with aggregate income growth.

That gives them $m = 301$ variables but only $n = 133$ quarterly observations.

They use a DMD to estimate a first-order VAR of rank 3 for these 301 variables.

Along the way, we'll learn that

- a DMD computes a reduced-rank estimator of the coefficient matrix of a first-order VAR
- that estimator is a **principal components regression**: compress the cross section into a few principal components, then regress next period's data on them
- DMD **modes** are right eigenvectors of the estimated coefficient matrix
- matching **left** eigenvectors can be computed cheaply, and together with the modes they give an exact modal representation of the estimated VAR that is useful for forecasting and for computing long-run responses

We'll use the following imports.

```
import numpy as np
import matplotlib.pyplot as plt
```

6.2 First-order vector autoregressions

We want to fit a **first-order vector autoregression**

$$X_{t+1} = AX_t + C\epsilon_{t+1}, \quad \epsilon_{t+1} \perp X_t \tag{6.1}$$

where ϵ_{t+1} is the time $t+1$ component of a sequence of IID $m \times 1$ random vectors with mean vector zero and identity covariance matrix and where the $m \times 1$ vector X_t is

$$X_t = \begin{bmatrix} X_{1,t} & X_{2,t} & \cdots & X_{m,t} \end{bmatrix}^\top \tag{6.2}$$

Here $\cdot^\top$ denotes matrix transposition and $X_{i,t}$ is variable i at time t.

Equation (6.1) has no constant term.

That is because we assume that each variable has already been transformed – for example, by taking differences, by subtracting a common trend, and by subtracting its time-series mean – so that $\{X_t\}$ can be regarded as a mean-zero, covariance-stationary process.

For example, [Sargent *et al.*, 2026] subtract a cross-section mean of log private income from every log quantile, use the growth rate of that mean as an additional variable, and then subtract time-series means from all variables.

Our data are organized in an $m \times (n+1)$ matrix $\tilde{X}$

$$\tilde{X} = \begin{bmatrix} X_1 \mid X_2 \mid \cdots \mid X_n \mid X_{n+1} \end{bmatrix}$$

where for $t = 1, \ldots, n+1$, the $m \times 1$ vector X_t is given by (6.2).

From $\tilde{X}$, we form two $m \times n$ matrices

$$X = \begin{bmatrix} X_1 \mid X_2 \mid \cdots \mid X_n \end{bmatrix}$$

and

$$X' = \begin{bmatrix} X_2 \mid X_3 \mid \cdots \mid X_{n+1} \end{bmatrix}$$

Here $'$ is part of the name of the matrix X' and does not indicate matrix transposition.

In forming X and X', we have in each case dropped a column from $\tilde{X}$, the last column in the case of X, and the first column in the case of X'.

We want to estimate a system (6.1) that consists of m least squares regressions of **everything** on one lagged value of **everything**.

The ith equation of (6.1) is a regression of $X_{i,t+1}$ on the vector X_t.

So the ith row of an estimator $\hat{A}$ is a $1 \times m$ vector of regression coefficients of $X_{i,t+1}$ on $X_{j,t}, j = 1, \ldots, m$.

A least squares estimator $\hat{A}$ solves

$$\min_{\check{A}} \|X' - \check{A}X\|_F \tag{6.3}$$

where $\|\cdot\|_F$ denotes the Frobenius norm of a matrix

$$\|B\|_F = \sqrt{\sum_i \sum_j |B_{ij}|^2}.$$

We denote the rank of X by $p \leq \min(m, n)$.

The **minimum-norm** solution of (6.3) is

$$\hat{A} = X' X^+ \tag{6.4}$$

where X^+ is the Moore-Penrose pseudo-inverse of X.

Two cases interest us

- $n \gg m$, so that we have many more time series observations n than variables m
- $m \gg n$, so that we have many more variables m than time series observations n

Short-Fat Case:

When $n \gg m$ and X has linearly independent **rows**, $XX^\top$ has an inverse and the pseudo-inverse X^+ is

$$X^+ = X^\top (XX^\top)^{-1}$$

Here X^+ is a **right-inverse** that verifies $XX^+ = I_{m \times m}$.

In this case, problem (6.3) has a unique solution and formula (6.4) becomes

$$\hat{A} = X' X^\top (XX^\top)^{-1} \tag{6.5}$$

This formula for least-squares regression coefficients is widely used in econometrics to estimate vector autoregressions.

The right side of formula (6.5) is proportional to the empirical cross second moment matrix of X_{t+1} and X_t times the inverse of the second moment matrix of X_t.

Tall-Skinny Case:

When $m \gg n$ and X has linearly independent **columns**, $X^\top X$ has an inverse and the pseudo-inverse X^+ is

$$X^+ = (X^\top X)^{-1} X^\top$$

Here X^+ is a **left-inverse** that verifies $X^+ X = I_{n \times n}$.

In this case, formula (6.4) becomes

$$\hat{A} = X' (X^\top X)^{-1} X^\top \tag{6.6}$$

If we use formula (6.6) to calculate $\hat{A}X$ we find that

$$\hat{A} X = X'$$

so that the regression equation **fits perfectly**.

This is a typical outcome in an **underdetermined least-squares** model.

In fact, every matrix $\check{A} = \hat{A} + N$ with $NX = 0$ also fits perfectly, so problem (6.3) has infinitely many solutions.

Formula (6.6) selects the one with the smallest Frobenius norm.

A perfect in-sample fit is a warning sign: with more coefficients than observations, the estimator $\hat{A}$ is tracking noise as well as signal.

We'll respond by constructing a **reduced-rank** estimator that no longer fits perfectly.

6.2.1 Computing X^+ with an SVD

An efficient way to compute the pseudo-inverse X^+ is to start with a **reduced** singular value decomposition

$$X = U\Sigma V^\top \tag{6.7}$$

where U is an $m \times p$ matrix, Σ is a $p \times p$ diagonal matrix of positive singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_p > 0$, and V is an $n \times p$ matrix.

As we saw in *Singular Value Decomposition*, for a reduced SVD

- $U^\top U = I_{p\times p}$ and $V^\top V = I_{p\times p}$, but
- $UU^\top \neq I_{m\times m}$ unless $p = m$, and $VV^\top \neq I_{n\times n}$ unless $p = n$.

For any rank p, the Moore-Penrose pseudo-inverse of X is

$$X^+ = V\Sigma^{-1}U^\top \tag{6.8}$$

When X has linearly independent columns, so that $p = n$ and V is a square orthogonal matrix, we can confirm that formula (6.8) agrees with our earlier formula:

$$\begin{aligned}(X^\top X)^{-1}X^\top &= (V\Sigma U^\top U\Sigma V^\top)^{-1}V\Sigma U^\top \\ &= (V\Sigma^2 V^\top)^{-1}V\Sigma U^\top \\ &= V\Sigma^{-2}V^\top V\Sigma U^\top \\ &= V\Sigma^{-1}U^\top\end{aligned}$$

where the third line uses $V^{-1} = V^\top$.

Substituting formula (6.8) into formula (6.4) gives

$$\hat{A} = X'V\Sigma^{-1}U^\top \tag{6.9}$$

Formula (6.9) shows a weakness of $\hat{A}$.

Small singular values σ_j enter through $1/\sigma_j$.

So directions in which the data X vary very little receive very large weights, which amplifies noise.

6.3 A reduced-rank VAR

A remedy is to keep only the $r < p$ largest singular values.

Let U_r be the $m \times r$ matrix consisting of the first r columns of U, let Σ_r be the $r \times r$ upper-left block of Σ, and let V_r be the $n \times r$ matrix consisting of the first r columns of V.

Our **rank-r estimator** of A is

$$\hat{A}_r = X'V_r\Sigma_r^{-1}U_r^\top \tag{6.10}$$

The $m \times m$ matrix $\hat{A}_r$ has rank at most r.

Notice that $V_r\Sigma_r^{-1}U_r^\top$ is the pseudo-inverse of $\hat{X}_r = U_r\Sigma_r V_r^\top$, the best rank-$r$ approximation of X described by the Eckart-Young theorem in *Singular Value Decomposition*.

So $\hat{A}_r = X'\hat{X}_r^+$: we obtain the rank-$r$ estimator by replacing X with its Eckart-Young approximation in the least squares formula $\hat{A} = X'X^+$.

This is the estimator that a **dynamic mode decomposition** computes.

Dynamic mode decomposition was introduced by [Schmid, 2010].

You can read about it in [Kutz *et al.*, 2016] and [Brunton and Kutz, 2019] (section 7.2).

Associated with $\hat{A}_r$ are fitted residuals and their covariance matrix

$$\hat{a}_{t+1} = X_{t+1} - \hat{A}_r X_t, \qquad \hat{\Omega} = \frac{1}{n}\sum_{t=1}^{n} \hat{a}_{t+1}\hat{a}_{t+1}^\top = \frac{1}{n}(X' - \hat{A}_r X)(X' - \hat{A}_r X)^\top \tag{6.11}$$

Thus, our estimated rank-r VAR is

$$X_{t+1} = \hat{A}_r X_t + \hat{a}_{t+1} \tag{6.12}$$

With $r < p$, the regression no longer fits perfectly and $\hat{\Omega}$ is not zero.

In applications, r is often very small, e.g., three or fewer.

A common way to choose r is to plot the singular values $\sigma_1, \sigma_2, \ldots$ and to truncate where they stop falling rapidly and level off – an "elbow rule."

[Sargent *et al.*, 2026] used this rule to choose $r = 3$.

6.3.1 Interpretation as a principal components regression

We can interpret formula (6.10) statistically.

Define the $r \times 1$ vector

$$f_t = U_r^\top X_t$$

whose components are the first r **principal components** of the cross section X_t (see the discussion of principal components analysis in *Singular Value Decomposition*).

Stacking them gives the $r \times n$ matrix $F = U_r^\top X = \Sigma_r V_r^\top$.

The least-squares regression of X_{t+1} on f_t has coefficient matrix

$$X' F^\top (F F^\top)^{-1} = X' V_r \Sigma_r (\Sigma_r^2)^{-1} = X' V_r \Sigma_r^{-1}$$

so that

$$\hat{A}_r X_t = \left(X' V_r \Sigma_r^{-1}\right) f_t.$$

Similarly, the least-squares regression of $f_{t+1} = U_r^\top X_{t+1}$ on f_t has $r \times r$ coefficient matrix

$$\tilde{A} = U_r^\top X' V_r \Sigma_r^{-1} \tag{6.13}$$

So $\tilde{A}$ is the first-order VAR coefficient matrix for the r principal components.

Note that $\tilde{A} = U_r^\top \hat{A}_r U_r$.

Thus, a DMD proceeds in two steps:

1. it compresses the cross section X_t into r principal components f_t, using only the variation in X
2. it regresses next period's data X_{t+1} on those principal components

Statisticians call this a **principal components regression**.

It resembles the "diffusion index" forecasting method of [Stock and Watson, 2002].

Note

The estimator $\hat{A}_r$ is **not** the matrix of rank r that minimizes $\|X' - \breve{A}X\|_F$.

That minimizer is the **reduced-rank regression** estimator of [Anderson, 1951].

It chooses the r directions that best *predict* X_{t+1}, while $\hat{A}_r$ uses the r directions that best *describe* X_t.

Exercise 6.10.2 compares the two estimators.

6.4 The DMD algorithm

Here is the DMD algorithm.

1. Compute the reduced SVD $X = U\Sigma V^\top$.
2. Retain the r largest singular values to form U_r, Σ_r, and V_r.
3. Form the $r \times r$ matrix $\tilde{A} = U_r^\top X' V_r \Sigma_r^{-1}$ and compute its eigendecomposition

$$\tilde{A}W = W\Lambda \tag{6.14}$$

   where Λ is a diagonal matrix of eigenvalues $\lambda_1, \ldots, \lambda_r$ and the columns of W are corresponding eigenvectors.
4. Form the $m \times r$ matrix Φ and the $r \times m$ matrix Ψ

$$\Phi = X' V_r \Sigma_r^{-1} W, \qquad \Psi = (W\Lambda)^{-1} U_r^\top \tag{6.15}$$

The eigenvalues λ_i on the diagonal of Λ are **DMD eigenvalues** and the columns ϕ_i of Φ are **DMD modes**.

We'll see below that the rows ψ_i of Ψ are matching **left** eigenvectors.

Notice that the heavy lifting consists of an SVD of the $m \times n$ matrix X and an eigendecomposition of the small $r \times r$ matrix $\tilde{A}$.

We never have to form or decompose an $m \times m$ matrix.

Throughout, we assume that $\tilde{A}$ has r linearly independent eigenvectors and nonzero eigenvalues, which is the generic case.

The matrix $\tilde{A}$ is real but not symmetric, so some eigenvalues can come in complex-conjugate pairs.

In that case, the corresponding columns of Φ and rows of Ψ also come in complex-conjugate pairs, while products such as $\Phi\Lambda^j\Psi$ are real.

For complex matrices, we use $\cdot^*$ to denote the conjugate transpose.

Note

[Tu *et al.*, 2014] scale the modes differently, defining them as the columns of $\Phi\Lambda^{-1}$.

With that scaling, the matching left eigenvectors are the rows of $\Lambda\Psi = W^{-1}U_r^\top$.

Rescaling a mode and its left eigenvector in opposite directions changes none of the formulas below.

6.5 Right eigenvectors: DMD modes

The following result of [Tu *et al.*, 2014] explains why the columns of Φ are called modes.

Proposition 6.5.1

The columns of Φ are eigenvectors of $\hat{A}_r$:

$$\hat{A}_r \Phi = \Phi \Lambda$$

Proof. Using formulas (6.10), (6.13), and (6.15),

$$\begin{aligned} \hat{A}_r \Phi &= (X' V_r \Sigma_r^{-1} U_r^\top)(X' V_r \Sigma_r^{-1} W) \\ &= X' V_r \Sigma_r^{-1} \tilde{A} W \\ &= X' V_r \Sigma_r^{-1} W \Lambda \\ &= \Phi \Lambda \end{aligned}$$

Equating columns gives $\hat{A}_r \phi_i = \lambda_i \phi_i$ for $i = 1, \ldots, r$.

We can say more.

Write $\hat{A}_r = B U_r^\top$, where $B = X' V_r \Sigma_r^{-1}$ is $m \times r$.

Because the nonzero eigenvalues of $B U_r^\top$ coincide with those of $U_r^\top B = \tilde{A}$, the eigenvalues of $\hat{A}_r$ are $\lambda_1, \ldots, \lambda_r$ together with $m - r$ zeros.

So the estimated VAR (6.12) is stable if and only if $\max_i |\lambda_i| < 1$, a condition we can check with an $r \times r$ matrix.

(Also see [Brunton and Kutz, 2022], p. 238.)

6.5.1 Projected modes

[Schmid, 2010] originally worked with the **projected modes** $U_r W$ instead of Φ.

Because $U_r^\top \Phi = \tilde{A} W = W \Lambda$, we have

$$U_r U_r^\top \Phi = U_r W \Lambda$$

so, up to scale, the projected modes are orthogonal projections of the modes Φ onto the column space of U_r.

The projected modes are eigenvectors of $\hat{A}_r$ only when $\Phi = U_r W \Lambda$, that is, when the columns of Φ lie in the column space of U_r.

That need not be true, because $\hat{A}_r U_r W = \Phi$.

The construction $\Phi = \hat{A}_r U_r W$ of [Tu *et al.*, 2014] delivers eigenvectors of $\hat{A}_r$ in all cases.

6.6 Left eigenvectors and a modal representation

Now we come to the rows of Ψ.

Proposition 6.6.1

Let $\Psi = (W\Lambda)^{-1}U_r^\top$. Then

1. $\Psi\Phi = I_{r\times r}$
2. $\Psi\hat{A}_r = \Lambda\Psi$, so that the rows of Ψ are left eigenvectors of $\hat{A}_r$
3. $\hat{A}_r = \Phi\Lambda\Psi = \sum_{i=1}^r \lambda_i\phi_i\psi_i$

Proof. For the first claim,

$$\Psi\Phi = \Lambda^{-1}W^{-1}U_r^\top X'V_r\Sigma_r^{-1}W = \Lambda^{-1}W^{-1}\tilde{A}W = \Lambda^{-1}\Lambda = I.$$

For the second claim, since $U_r^\top\hat{A}_r = U_r^\top X'V_r\Sigma_r^{-1}U_r^\top = \tilde{A}U_r^\top$,

$$\Psi\hat{A}_r = \Lambda^{-1}W^{-1}\tilde{A}U_r^\top = \Lambda^{-1}\Lambda W^{-1}U_r^\top = W^{-1}U_r^\top = \Lambda\Psi.$$

For the third claim,

$$\Phi\Lambda\Psi = X'V_r\Sigma_r^{-1}W\Lambda\Lambda^{-1}W^{-1}U_r^\top = X'V_r\Sigma_r^{-1}U_r^\top = \hat{A}_r.$$

Define the $r\times 1$ vector of **modal coordinates**

$$b_t = \Psi X_t = (W\Lambda)^{-1}U_r^\top X_t \tag{6.16}$$

Computing b_t requires only the principal components $f_t = U_r^\top X_t$ and an $r\times r$ matrix.

Because $\Psi\Phi = I$, iterating $\hat{A}_r = \Phi\Lambda\Psi$ gives

$$\hat{A}_r^j = \Phi\Lambda^j\Psi, \quad j\geq 1 \tag{6.17}$$

Consequently, the estimated VAR (6.12) implies

$$b_{t+1} = \Lambda b_t + \Psi\hat{a}_{t+1}$$

and the j-step-ahead forecast of X_{t+j} conditional on X_t is

$$\hat{A}_r^j X_t = \Phi\Lambda^j b_t = \sum_{i=1}^r \phi_i\lambda_i^j b_{i,t} \tag{6.18}$$

Each mode ϕ_i contributes a component whose amplitude decays geometrically at rate λ_i.

The matrix $\Phi\Psi$ is an oblique projection onto the column space of Φ.

It splits X_t into Φb_t plus a remainder $(I-\Phi\Psi)X_t$ that $\hat{A}_r$ sends to zero, because $\hat{A}_r(I-\Phi\Psi) = \Phi\Lambda\Psi - \Phi\Lambda\Psi\Phi\Psi = 0$.

6.6.1 A tempting alternative that does not work

The Moore-Penrose pseudo-inverse $\Phi^+ = (\Phi^*\Phi)^{-1}\Phi^*$ also satisfies $\Phi^+\Phi = I$.

The vector $\check{b}_t = \Phi^+ X_t$ is a vector of least squares regression coefficients of X_t on the columns of Φ, and $\Phi\check{b}_t$ is the orthogonal projection of X_t onto the column space of Φ.

That makes $\check{b}_t$ a natural **descriptive** summary of how much each mode is present in X_t.

But in general

$$\Phi\Lambda\Phi^+ \neq \hat{A}_r$$

so $\check{b}_t$ does **not** evolve according to $\check{b}_{t+1} = \Lambda\check{b}_t$ under the estimated VAR, and $\Phi\Lambda^j\Phi^+ X_t$ is not the VAR forecast $\hat{A}_r^j X_t$.

The reason is that Φ^+ and Ψ are both left inverses of Φ, but their rows span different spaces:

- the rows of Φ^+ span the same space as the rows of Φ^*
- the rows of Ψ span the same space as the rows of $U_r^\top$

The two left inverses coincide if and only if the column spaces of Φ and U_r are the same.

The column space of Φ is the column space of $X'V_r$, which describes where the data go **next** period.

The column space of U_r is spanned by the r directions in which the **current** data X vary most.

These two spaces coincide in special cases, for example when noise-free data are driven by exactly r factors (see Exercise 6.10.1), but not in general.

Note

The formula $(W\Lambda)^{-1}U_r^\top X_1$ appears in [Brunton and Kutz, 2022] (p. 240), where it is presented as a computationally cheap approximation to $\Phi^+ X_1$.

Proposition 6.6.1 shows that it is exactly the vector of coordinates that is consistent with the dynamics of the estimated VAR.

6.6.2 A state-space representation

The modal representation lets us write the estimated VAR as a linear state-space system with an $r \times 1$ state vector.

Define

$$\hat{x}_t = \Lambda\Psi X_{t-1} = W^{-1}U_r^\top X_{t-1} = W^{-1}f_{t-1}$$

Then $X_t = \hat{A}_r X_{t-1} + \hat{a}_t = \Phi\Lambda\Psi X_{t-1} + \hat{a}_t$ and $\hat{x}_{t+1} = \Lambda\Psi X_t$ imply

$$\begin{aligned} \hat{x}_{t+1} &= \Lambda\hat{x}_t + \Lambda\Psi\hat{a}_t \\ X_t &= \Phi\hat{x}_t + \hat{a}_t \end{aligned} \tag{6.19}$$

where we have used $\Psi\Phi = I$.

In representation (6.19)

- the state $\hat{x}_t$ is the vector of lagged principal components expressed in the eigenvector basis of the principal-component VAR $\tilde{A}$
- the transition matrix Λ is diagonal, so the r components of the state evolve independently, apart from correlated shocks

- the columns of Φ are **loadings** of the m observed variables on the r components of the state
- a single $m \times 1$ vector of residuals $\hat{a}_t$ drives both equations; by the least squares normal equations, it is orthogonal in sample to f_{t-1} and therefore to $\hat{x}_t$

[Sargent *et al.*, 2026] connect DMDs to linear state-space models of this form.

6.7 Long-run responses

Modal representation (6.17) also simplifies calculations of long-run responses.

If $|\lambda_i| < 1$ for all i, then

$$\sum_{j=1}^{\infty} \hat{A}_r^j = \Phi\Lambda(I - \Lambda)^{-1}\Psi$$

and

$$(I - \hat{A}_r)^{-1} = I + \Phi\Lambda(I - \Lambda)^{-1}\Psi \tag{6.20}$$

The matrix $\Lambda(I - \Lambda)^{-1}$ is diagonal with entries $\lambda_i/(1 - \lambda_i)$.

So long-run responses are dominated by the most persistent modes.

For example, the rank-3 VAR of [Sargent *et al.*, 2026] has DMD eigenvalues 0.963, 0.865, and 0.547, with long-run weights of roughly 26, 6.4, and 1.2.

Formulas like (6.20) appear in **additive functionals**, which are studied in Additive and Multiplicative Functionals.

Suppose that the first component of X_t is the growth rate of the logarithm Y_t of some aggregate, measured as a deviation from its mean ν, so that $Y_{t+1} - Y_t - \nu = e_1 X_{t+1}$ where $e_1 = \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix}$.

If X_t obeys the VAR (6.12), then Y_t can be decomposed into a deterministic trend $t\nu$, a martingale $\sum_{j=1}^{t} H\hat{a}_j$, and a stationary component $-gX_t$, plus a constant, where

$$g = e_1\hat{A}_r(I - \hat{A}_r)^{-1}, \qquad H = e_1(I - \hat{A}_r)^{-1} = e_1 + g$$

[Sargent *et al.*, 2026] use this decomposition for aggregate income and for all 300 CEX quantiles.

Formula (6.20) implies

$$g = e_1\Phi\Lambda(I - \Lambda)^{-1}\Psi$$

so computing g and H requires no inversion of an $m \times m$ matrix.

6.8 A simulated example

To see these objects at work, we simulate a tall-skinny data set that shares some features with the CEX data of [Sargent *et al.*, 2026].

There are $m = 300$ cross-section variables, which we can think of as quantiles indexed by $q \in [0, 1]$, and $n + 1 = 134$ quarterly observations.

The cross section is driven by three latent factors z_t that obey a first-order VAR with a diagonal transition matrix

$$z_{t+1} = Dz_t + w_{t+1}, \qquad D = \text{diag}(0.95, 0.80, 0.50)$$

and the observed data are

$$X_t = G z_t + \sigma_v v_t$$

where w_{t+1} and v_t are standard normal random vectors.

The three columns of G are shaped like a **level** factor that moves all quantiles together, a **slope** factor that moves low and high quantiles in opposite directions, and a **curvature** factor that moves the tails relative to the middle.

Here is code that simulates the data.

```
def simulate_data(m=300, T=134, D_diag=(0.95, 0.80, 0.50), σ_v=0.5, seed=1234):
    """
    Simulate an m x T data matrix driven by three latent AR(1) factors.
    """
    rng = np.random.default_rng(seed)
    q = np.linspace(0, 1, m)
    G = np.column_stack([np.ones(m),                # level
                         2 * (q - 0.5),             # slope
                         6 * (q - 0.5)**2 - 0.5])   # curvature
    d = np.array(D_diag)
    k = len(d)
    z = np.zeros((k, T))
    z[:, 0] = rng.normal(size=k) / np.sqrt(1 - d**2)
    for t in range(T - 1):
        z[:, t+1] = d * z[:, t] + rng.normal(size=k)
    X_tilde = G @ z + σ_v * rng.normal(size=(m, T))
    return X_tilde, G, z, q

X_tilde, G, z, q = simulate_data()
X_tilde = X_tilde - X_tilde.mean(axis=1, keepdims=True)    # demean each variable
X_tilde.shape
```

```
(300, 134)
```

The next function implements the DMD algorithm.

```
def dmd(X_tilde, r):
    """
    Rank-r DMD of the m x (n+1) data matrix X_tilde.
    """
    X, X_prime = X_tilde[:, :-1], X_tilde[:, 1:]
    U, σ, Vt = np.linalg.svd(X, full_matrices=False)
    U_r, σ_r, V_r = U[:, :r], σ[:r], Vt[:r, :].T

    B = X_prime @ V_r / σ_r          # X' V_r Σ_r^{-1}
    A_tilde = U_r.T @ B              # r x r VAR for principal components
    λ, W = np.linalg.eig(A_tilde)
    order = np.argsort(-np.abs(λ))   # most persistent modes first
    λ, W = λ[order], W[:, order]

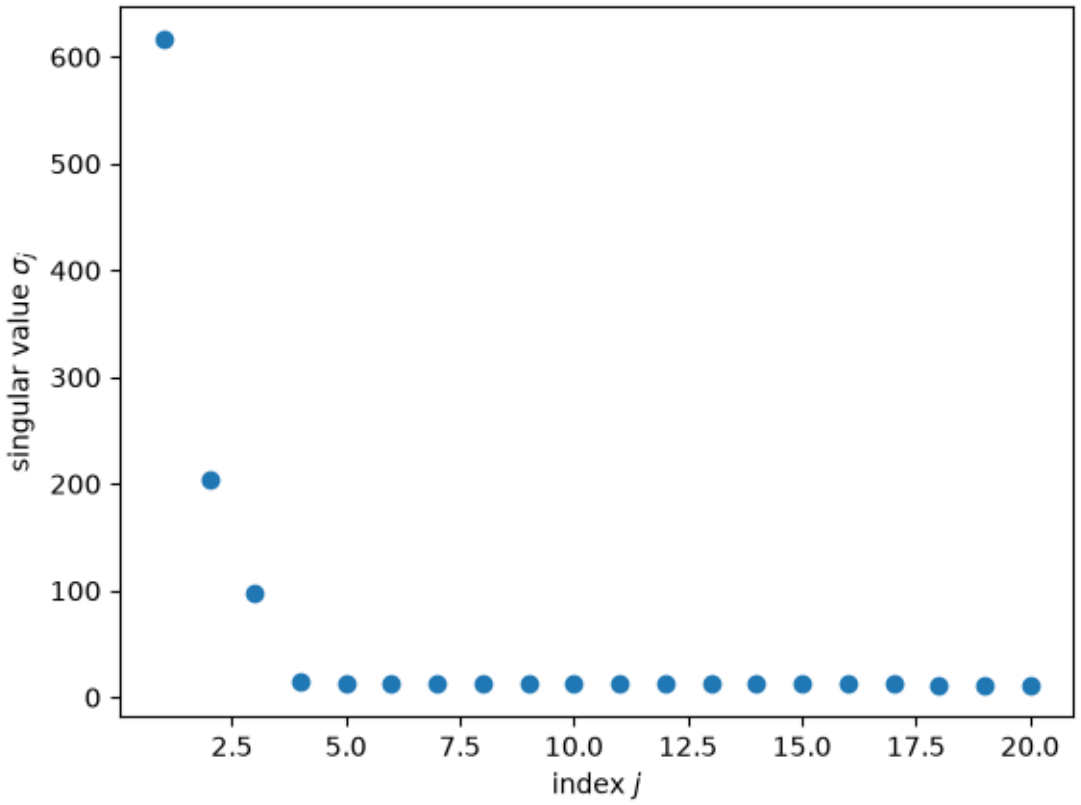
    Φ = B @ W                                # DMD modes (right eigenvectors)
    Ψ = np.linalg.solve(W * λ, U_r.T)        # (W Λ)^{-1} U_r^T (left eigenvectors)
    A_hat = B @ U_r.T                        # rank-r estimator of A
    a_hat = X_prime - A_hat @ X              # residuals
    Ω_hat = a_hat @ a_hat.T / X.shape[1]

    return dict(λ=λ, Φ=Φ, Ψ=Ψ, W=W, A_tilde=A_tilde, A_hat=A_hat,
                Ω_hat=Ω_hat, σ=σ, U_r=U_r)
```

To choose r, we plot the largest singular values of X.

```
σ = np.linalg.svd(X_tilde[:, :-1], compute_uv=False)

fig, ax = plt.subplots()
ax.plot(np.arange(1, 21), σ[:20], 'o')
ax.set_xlabel('index $j$')
ax.set_ylabel(r'singular value $\sigma_j$')
plt.show()
```



The singular values fall rapidly for $j = 1, 2, 3$ and then level off, so the elbow rule tells us to set $r = 3$.

```
r = 3
res = dmd(X_tilde, r)
λ, Φ, Ψ, A_hat = res['λ'], res['Φ'], res['Ψ'], res['A_hat']

print("DMD eigenvalues:  ", np.round(λ, 3))
print("true eigenvalues: ", np.array([0.95, 0.80, 0.50]))
```

```
DMD eigenvalues:   [0.93  0.82  0.337]
true eigenvalues:  [0.95 0.8  0.5 ]
```

The two most persistent eigenvalues are estimated well.

The least persistent one is biased toward zero.

That is because the data are noisy measurements of the factors, so $\{X_t\}$ is not exactly a first-order VAR; errors in the principal components f_t attenuate the estimated coefficients, most visibly for the least persistent mode.

Next we verify *Proposition 6.5.1* and *Proposition 6.6.1*, and check whether $\Phi\Lambda\Phi^+$ equals $\hat{A}_r$.

```
Φ_pinv = np.linalg.pinv(Φ)
rel = lambda M: np.linalg.norm(M) / np.linalg.norm(A_hat)

print("max |A_hat Φ - Φ Λ|         =", np.abs(A_hat @ Φ - Φ * λ).max())
print("max |Ψ Φ - I|               =", np.abs(Ψ @ Φ - np.eye(r)).max())
print("max |Ψ A_hat - Λ Ψ|         =", np.abs(Ψ @ A_hat - λ[:, None] * Ψ).max())
print("relative error, Φ Λ Ψ       =", rel(A_hat - (Φ * λ) @ Ψ))
print("relative error, Φ Λ Φ^+     =", rel(A_hat - (Φ * λ) @ Φ_pinv))
```

```
max |A_hat Φ - Φ Λ|         = 1.1796119636642288e-16
max |Ψ Φ - I|               = 1.7763568394002505e-15
max |Ψ A_hat - Λ Ψ|         = 2.7755575615628914e-16
relative error, Φ Λ Ψ       = 1.9407620402191376e-16
relative error, Φ Λ Φ^+     = 0.07337622086364687
```

The first four numbers are zero up to rounding error, while $\Phi\Lambda\Phi^+$ misses $\hat{A}_r$ by several percent.

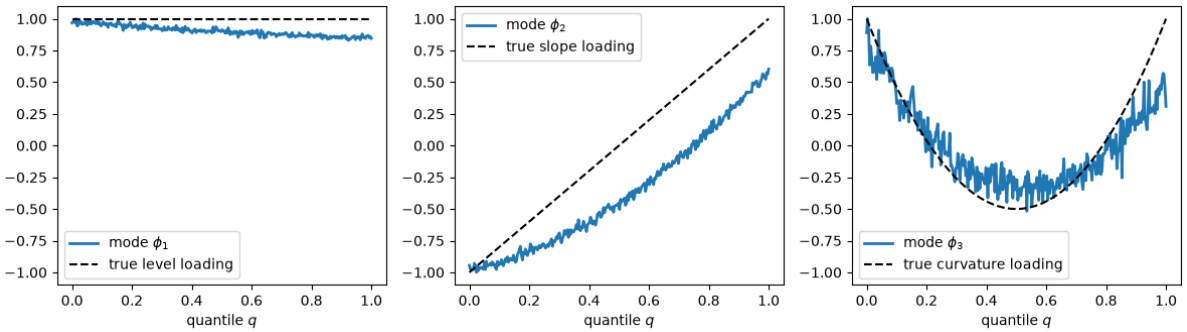
Now let's look at the modes.

Modes are determined only up to scale, so we normalize each mode to have maximum absolute value one and choose its sign to line up with the corresponding column of G, which we normalize in the same way.

```
def normalize(v):
    return v / np.abs(v).max()

labels = ['level', 'slope', 'curvature']
Φ_real = np.real_if_close(Φ)

fig, axes = plt.subplots(1, 3, figsize=(12, 3.5))
for i, ax in enumerate(axes):
    g_i = normalize(G[:, i])
    ψ_i = normalize(Φ_real[:, i])
    ψ_i *= np.sign(ψ_i @ g_i)
    ax.plot(q, ψ_i, lw=2, label=rf'mode $\phi_{i+1}$')
    ax.plot(q, g_i, 'k--', label=f'true {labels[i]} loading')
    ax.set_ylim(-1.1, 1.1)
    ax.set_xlabel('quantile $q$')
    ax.legend()
plt.tight_layout()
plt.show()
```



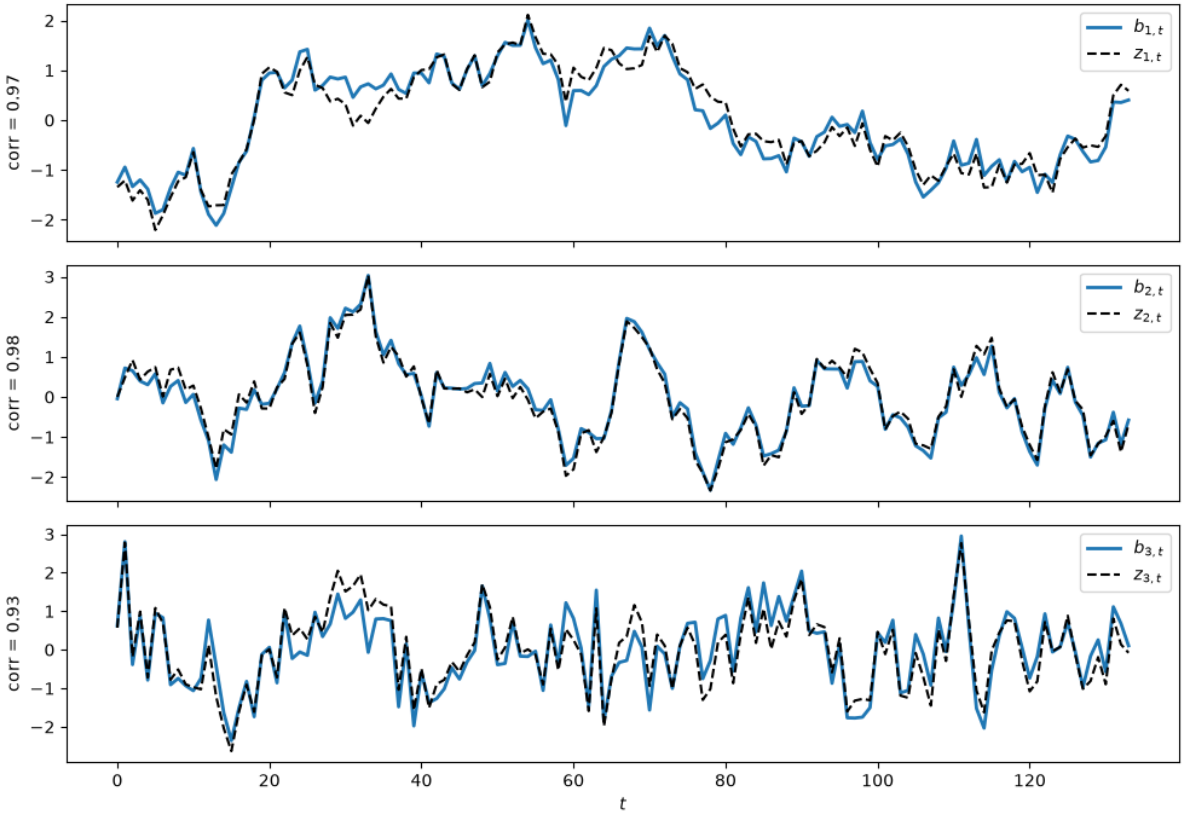
Each mode is close to one of the true loading patterns, but not identical to it.

Because of sampling error, the estimated eigenvectors mix the true factors a little, most visibly for the second mode.

Next we compare the modal coordinates $b_t = \Psi X_t$ with the true factors z_t, after standardizing both.

```
b = np.real_if_close(Ψ @ X_tilde)
standardize = lambda x: (x - x.mean()) / x.std()

fig, axes = plt.subplots(3, 1, figsize=(10, 7), sharex=True)
for i, ax in enumerate(axes):
    corr = np.corrcoef(b[i], z[i])[0, 1]
    ax.plot(np.sign(corr) * standardize(b[i]), lw=2, label=rf'$b_{{{i+1},t}}$')
    ax.plot(standardize(z[i]), 'k--', label=rf'$z_{{{i+1},t}}$')
    ax.set_ylabel(f'corr = {abs(corr):.2f}')
    ax.legend(loc='upper right')
axes[-1].set_xlabel('$t$')
plt.tight_layout()
plt.show()
```



The DMD recovers the latent factors well.

Finally, we compare j-step-ahead forecasts computed in three ways:

- directly from $\hat{A}_r^j X_t$
- from the modal representation (6.18) with $b_t = \Psi X_t$
- from the tempting alternative $\Phi \Lambda^j \Phi^+ X_t$

```
X = X_tilde[:, :-1]
for j in (1, 4, 12):
    direct = np.linalg.matrix_power(A_hat, j) @ X
    via_Ψ = np.real_if_close((Φ * λ**j) @ Ψ @ X)
    via_pinv = np.real_if_close((Φ * λ**j) @ Φ_pinv @ X)
```

(continues on next page)

(continued from previous page)

```
    err = lambda F: np.linalg.norm(F - direct) / np.linalg.norm(direct)
    print(f"j = {j:2d}:  error with Ψ = {err(via_Ψ):.1e},  "
          f"error with Φ^+ = {err(via_pinv):.3f}")
```

```
j =  1:  error with Ψ = 1.7e-16,  error with Φ^+ = 0.004
j =  4:  error with Ψ = 2.8e-15,  error with Φ^+ = 0.001
j = 12:  error with Ψ = 1.0e-14,  error with Φ^+ = 0.001
```

Forecasts built from Ψ reproduce the VAR forecasts exactly, while those built from Φ^+ do not.

In this example the forecast discrepancy from using Φ^+ is small, but, as Exercise 6.10.1 shows, it grows with the amount of noise in the data.

6.9 Source for some Python code

You can find a Python implementation of DMD in the PyDMD package.

6.10 Exercises

Exercise 6.10.1

In this lecture we saw that $\Phi^+ = \Psi$ if and only if the column spaces of Φ and U_r are the same.

1. Simulate data with `simulate_data(σ_v=0.0)`, so that the data are exact linear combinations of three factors, and compute a rank-3 DMD. Verify numerically that $\Phi^+ = \Psi$ and that $\Phi \Lambda \Phi^+ = \hat{A}_r$.
2. Explain why the column spaces of Φ and U_r coincide in this case.
3. Repeat part 1 with $\sigma_v = 0.1, 0.5, 1.0$ and report how far $\Phi \Lambda \Phi^+$ is from $\hat{A}_r$.

Solution

Here is one solution.

```
for σ_v in (0.0, 0.1, 0.5, 1.0):
    X_sim = simulate_data(σ_v=σ_v)[0]
    X_sim = X_sim - X_sim.mean(axis=1, keepdims=True)
    out = dmd(X_sim, 3)
    Φ_s, Ψ_s, λ_s, A_s = out['Φ'], out['Ψ'], out['λ'], out['A_hat']
    Φ_s_pinv = np.linalg.pinv(Φ_s)
    gap_left = np.linalg.norm(Φ_s_pinv - Ψ_s) / np.linalg.norm(Ψ_s)
    gap_A = np.linalg.norm(A_s - (Φ_s * λ_s) @ Φ_s_pinv) / np.linalg.norm(A_s)
    print(f"σ_v = {σ_v:.1f}:  |Φ^+ - Ψ|/|Ψ| = {gap_left:.1e},  "
          f"|A_hat - Φ Λ Φ^+|/|A_hat| = {gap_A:.1e}")
```

```
σ_v = 0.0:  |Φ^+ - Ψ|/|Ψ| = 9.0e-16,  |A_hat - Φ Λ Φ^+|/|A_hat| = 1.1e-15
σ_v = 0.1:  |Φ^+ - Ψ|/|Ψ| = 4.2e-02,  |A_hat - Φ Λ Φ^+|/|A_hat| = 1.5e-02
σ_v = 0.5:  |Φ^+ - Ψ|/|Ψ| = 2.1e-01,  |A_hat - Φ Λ Φ^+|/|A_hat| = 7.3e-02
σ_v = 1.0:  |Φ^+ - Ψ|/|Ψ| = 4.0e-01,  |A_hat - Φ Λ Φ^+|/|A_hat| = 1.4e-01
```

When $\sigma_v = 0$, $X_t = Gz_t$ for every t, so every column of both X and X' lies in the three-dimensional column space of G.

Because X has rank three, the column space of U_3 equals the column space of G.

The columns of $\Phi = X'V_3\Sigma_3^{-1}W$ are linear combinations of columns of X', so they also lie in the column space of G.

Since Φ has three linearly independent columns, its column space is the column space of G, which is the column space of U_3.

Then the rows of Ψ lie in the space spanned by the rows of Φ^*.

So the rows of $\Psi - \Phi^+$ can be written as $C\Phi^*$ for some matrix C.

Because $(\Psi - \Phi^+)\Phi = I - I = 0$, we get $C\Phi^*\Phi = 0$, and since $\Phi^*\Phi$ is invertible, $C = 0$.

With measurement noise, X' has components outside the column space of U_3, the two left inverses differ, and the gap between $\Phi\Lambda\Phi^+$ and $\hat{A}_r$ grows with σ_v.

Exercise 6.10.2

This exercise compares the DMD estimator $\hat{A}_r$ with the reduced-rank regression estimator of [Anderson, 1951].

The reduced-rank regression estimator is the rank-r matrix that minimizes $\|X' - \check{A}X\|_F$.

It can be computed as $\hat{A}_r^{RRR} = P_r\hat{A}$, where $\hat{A} = X'X^+$ is the minimum-norm least squares estimator and $P_r = Q_rQ_r^\top$, where the columns of Q_r are the first r left singular vectors of the fitted values $\hat{A}X$.

Using data from `simulate_data()`, estimate both rank-3 estimators on the first 100 periods, then

1. compute the relative in-sample residual norm $\|X' - \check{A}X\|_F / \|X'\|_F$ for each estimator
2. compute the same statistic out of sample, using one-step-ahead forecasts for the remaining periods
3. repeat for several random seeds and comment on the results

Solution

Here is one solution.

```
def rrr(X, X_prime, r):
    "Reduced-rank regression estimator of rank r."
    A_full = X_prime @ np.linalg.pinv(X)
    Q = np.linalg.svd(A_full @ X, full_matrices=False)[0][:, :r]
    return Q @ Q.T @ A_full

def rel_resid(A, X, X_prime):
    return np.linalg.norm(X_prime - A @ X) / np.linalg.norm(X_prime)

n_est = 100
print("seed   DMD in   DMD out   RRR in   RRR out")
for seed in range(5):
    X_sim = simulate_data(seed=seed)[0]
    X_sim = X_sim - X_sim.mean(axis=1, keepdims=True)
    X_in, Xp_in = X_sim[:, :n_est-1], X_sim[:, 1:n_est]
    X_out, Xp_out = X_sim[:, n_est-1:-1], X_sim[:, n_est:]
```

```
A_dmd = dmd(X_sim[:, :n_est], 3)['A_hat']
A_rrr = rrr(X_in, Xp_in, 3)
print(f"{seed:4d}   {rel_resid(A_dmd, X_in, Xp_in):.3f}     "
      f"{rel_resid(A_dmd, X_out, Xp_out):.3f}      "
      f"{rel_resid(A_rrr, X_in, Xp_in):.3f}     "
      f"{rel_resid(A_rrr, X_out, Xp_out):.3f}")
```

```
seed   DMD in   DMD out   RRR in   RRR out
   0   0.515    0.616     0.185    0.709
   1   0.570    0.578     0.228    0.638
   2   0.463    0.788     0.165    0.901
   3   0.429    0.786     0.155    0.946
   4   0.405    0.384     0.155    0.430
```

By construction, reduced-rank regression has the smaller in-sample residuals.

But in these samples it forecasts worse out of sample than the DMD estimator.

Reduced-rank regression starts from $\hat{A} = X'V\Sigma^{-1}U^\top$, which uses **all** singular values of X, including tiny ones that amplify noise, and then chooses the three directions that best fit X' in sample.

The DMD estimator instead discards the small singular values before it regresses, which acts as a form of regularization.

Exercise 6.10.3

Using the rank-3 DMD of the simulated data from the lecture,

1. prove formula (6.17) and use it to derive formula (6.20)
2. verify formula (6.20) numerically by comparing it with a direct computation of $(I - \hat{A}_r)^{-1}$
3. treating the first component of X_t as the growth-rate variable, compute $g = e_1 \hat{A}_r (I - \hat{A}_r)^{-1}$ both directly and with the modal formula $g = e_1 \Phi \Lambda (I - \Lambda)^{-1} \Psi$, and report the contribution of each mode to g
4. compute the modal formula again with Φ^+ in place of Ψ and report how much the answer changes

Solution

For part 1, *Proposition 6.6.1* gives $\hat{A}_r = \Phi \Lambda \Psi$ and $\Psi \Phi = I$.

So $\hat{A}_r^2 = \Phi \Lambda (\Psi \Phi) \Lambda \Psi = \Phi \Lambda^2 \Psi$, and induction gives $\hat{A}_r^j = \Phi \Lambda^j \Psi$ for $j \geq 1$.

When $|\lambda_i| < 1$ for all i,

$$\sum_{j=1}^{\infty} \hat{A}_r^j = \Phi \left(\sum_{j=1}^{\infty} \Lambda^j \right) \Psi = \Phi \Lambda (I - \Lambda)^{-1} \Psi$$

and since all eigenvalues of $\hat{A}_r$ are inside the unit circle, $(I - \hat{A}_r)^{-1} = I + \sum_{j=1}^{\infty} \hat{A}_r^j$.

Here is code for parts 2–4.

```
m = A_hat.shape[0]
I = np.eye(m)
weights = λ / (1 - λ)

LR_direct = np.linalg.inv(I - A_hat)
```

```
LR_modal = np.real_if_close(I + (Φ * weights) @ Ψ)
print("max |direct - modal| for (I - A_hat)^{-1}:",
      np.abs(LR_direct - LR_modal).max())

e1 = np.zeros(m)
e1[0] = 1.0
g_direct = e1 @ A_hat @ LR_direct
g_modal = np.real_if_close(e1 @ (Φ * weights) @ Ψ)
print("max |g_direct - g_modal|:", np.abs(g_direct - g_modal).max())

print("\nlong-run weights λ/(1-λ):", np.round(weights.real, 2))
for i in range(r):
    g_i = np.real_if_close(Φ[0, i] * weights[i] * Ψ[i, :])
    share = np.linalg.norm(g_i) / np.linalg.norm(g_direct)
    print(f"mode {i+1}: |contribution to g| / |g| = {share:.2f}")

g_pinv = np.real_if_close(e1 @ (Φ * weights) @ Φ_pinv)
print("\nrelative error in g using Φ^+:",
      np.linalg.norm(g_pinv - g_direct) / np.linalg.norm(g_direct))
```

```
max |direct - modal| for (I - A_hat)^{-1}: 1.457167719820518e-15
max |g_direct - g_modal|: 1.3183898417423734e-15

long-run weights λ/(1-λ): [13.37  4.57  0.51]
mode 1: |contribution to g| / |g| = 1.29
mode 2: |contribution to g| / |g| = 0.66
mode 3: |contribution to g| / |g| = 0.10

relative error in g using Φ^+: 0.026761234549937188
```

The modal formula reproduces the direct computation up to rounding error, and requires only $r \times r$ operations beyond forming Φ and Ψ.

The most persistent mode, with the largest weight $\lambda_i/(1-\lambda_i)$, contributes the most to g.

The ratios need not sum to one, and one can exceed one, because contributions of different modes partially offset each other.

Replacing Ψ by Φ^+ gives a different, incorrect answer.

CHAPTER

SEVEN

USING NEWTON'S METHOD TO SOLVE ECONOMIC MODELS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Using Newton's Method to Solve Economic Models*
  - *Overview*
  - *Fixed point computation using Newton's method*
  - *Root-Finding in one dimension*
  - *Multivariate Newton's method*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install jax
```

7.1 Overview

Many economic problems involve finding fixed points or zeros (also called "roots") of functions.

For example, in a simple supply and demand model, an equilibrium price is one that makes excess demand zero.

In other words, an equilibrium is a zero of the excess demand function.

There are various computational techniques for solving for fixed points and zeros.

In this lecture we study an important gradient-based technique called Newton's method.

Newton's method does not always work but, in situations where it does, convergence is often fast when compared to other methods.

The lecture will apply Newton's method in one-dimensional and multidimensional settings to solve fixed-point and zero-finding problems.

- When finding the fixed point of a function f, Newton's method updates an existing guess of the fixed point by solving for the fixed point of a linear approximation to the function f.
- When finding the zero of a function f, Newton's method updates an existing guess by solving for the zero of a linear approximation to the function f.

To build intuition, we first consider an easy, one-dimensional fixed point problem where we know the solution and solve it using both successive approximation and Newton's method.

Then we apply Newton's method to multidimensional settings to solve for market equilibria with multiple goods.

At the end of the lecture, we leverage the power of automatic differentiation in `jax` to solve a very high-dimensional equilibrium problem.

We use the following imports in this lecture

```
import matplotlib.pyplot as plt
from typing import NamedTuple
from scipy.optimize import root
import jax.numpy as jnp
import jax

# Enable 64-bit precision
jax.config.update("jax_enable_x64", True)
```

7.2 Fixed point computation using Newton's method

In this section we solve the fixed point of the law of motion for capital in the setting of the Solow growth model.

We will inspect the fixed point visually, solve it by successive approximation, and then apply Newton's method to achieve faster convergence.

7.2.1 The Solow model

In the Solow growth model, assuming Cobb-Douglas production technology and zero population growth, the law of motion for capital is

$$k_{t+1} = g(k_t) \quad \text{where} \quad g(k) := sAk^\alpha + (1-\delta)k \tag{7.1}$$

Here

- k_t is capital stock per worker,
- $A, \alpha > 0$ are production parameters with $\alpha < 1$
- $s > 0$ is a savings rate, and
- $\delta \in (0, 1)$ is a rate of depreciation

In this example, we wish to calculate the unique strictly positive fixed point of g, the law of motion for capital.

In other words, we seek a $k^* > 0$ such that $g(k^*) = k^*$.

- Such a k^* is called a steady state, since $k_t = k^*$ implies $k_{t+1} = k^*$.

Using pencil and paper to solve $g(k) = k$, you will be able to confirm that

$$k^* = \left(\frac{sA}{\delta}\right)^{1/(1-\alpha)}$$

7.2.2 Implementation

Let's store our parameters in `NamedTuple` to help us keep our code clean and concise.

```
class SolowParameters(NamedTuple):
    A: float
    s: float
    α: float
    δ: float
```

This function creates a suitable `SolowParameters` with default parameter values.

```
def create_solow_params(A=2.0, s=0.3, α=0.3, δ=0.4):
    """Creates a Solow model parameterization with default values."""
    return SolowParameters(A=A, s=s, α=α, δ=δ)
```

The next two functions implement the law of motion *(7.2.1)* and store the true fixed point k^*.

```
def g(k, params):
    A, s, α, δ = params
    return A * s * k**α + (1 - δ) * k


def exact_fixed_point(params):
    A, s, α, δ = params
    return ((s * A) / δ) ** (1 / (1 - α))
```

Here is a function to provide a 45 degree plot of the dynamics.

```
def plot_45(params, ax, fontsize=14):

    k_min, k_max = 0.0, 3.0
    k_grid = jnp.linspace(k_min, k_max, 1200)

    # Plot the functions
    lb = r"$g(k) = sAk^{\alpha} + (1 - \delta)k$"
    ax.plot(k_grid, g(k_grid, params), lw=2, alpha=0.6, label=lb)
    ax.plot(k_grid, k_grid, "k--", lw=1, alpha=0.7, label="45")

    # Show and annotate the fixed point
    kstar = exact_fixed_point(params)
    fps = (kstar,)
    ax.plot(fps, fps, "go", ms=10, alpha=0.6)
    ax.annotate(
        r"$k^* = (sA / \delta)^{\frac{1}{1-\alpha}}$",
        xy=(kstar, kstar),
        xycoords="data",
        xytext=(20, -20),
        textcoords="offset points",
        fontsize=fontsize,
```

(continues on next page)

(continued from previous page)

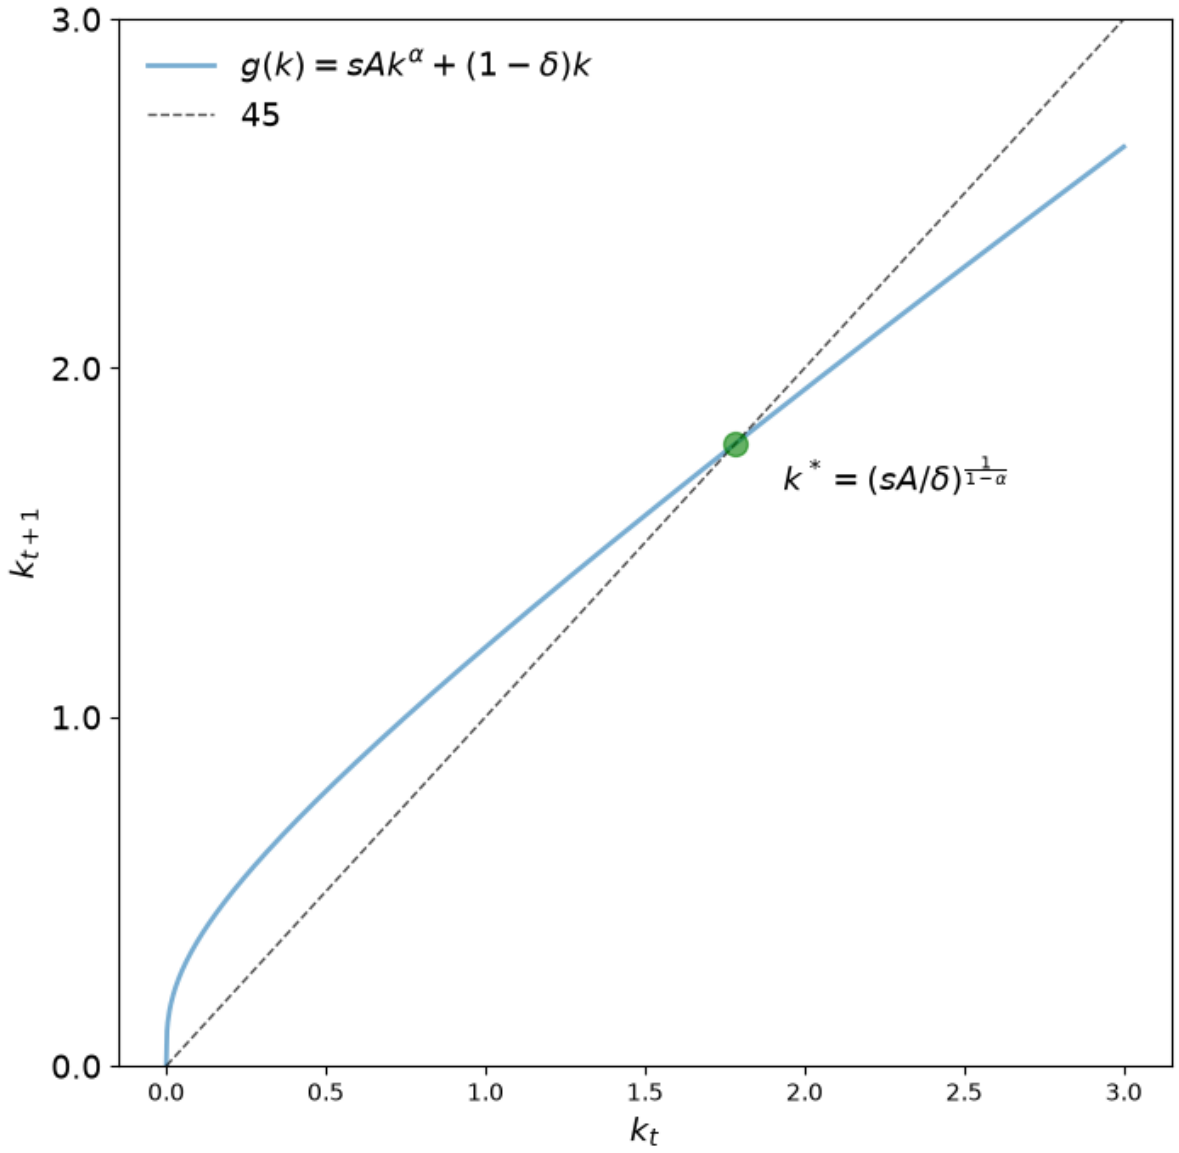
```
    )

    ax.legend(loc="upper left", frameon=False, fontsize=fontsize)

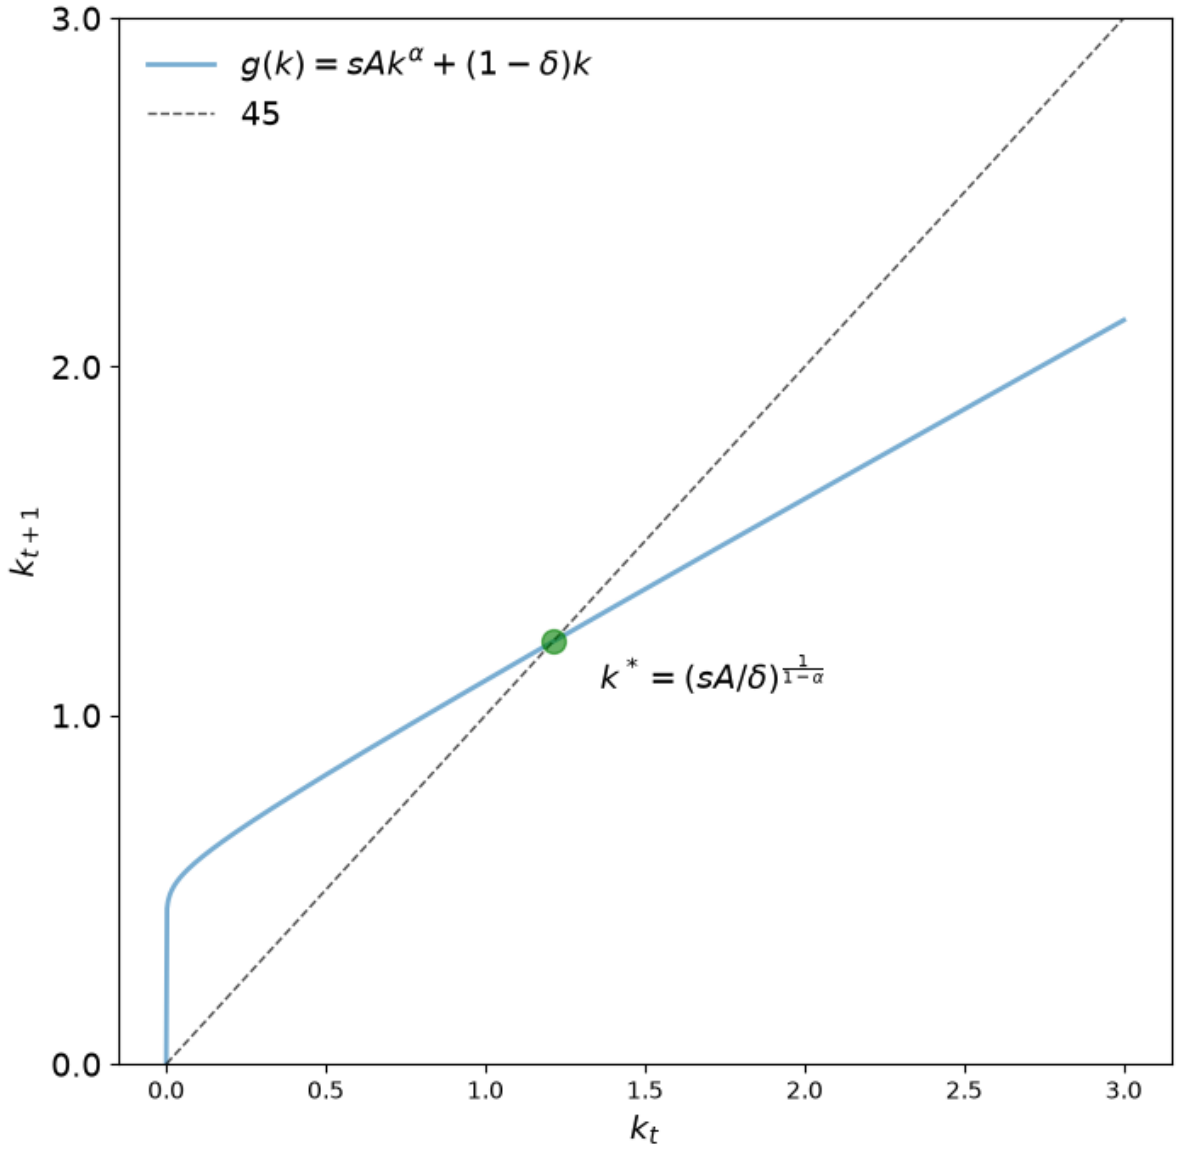
    ax.set_yticks((0, 1, 2, 3))
    ax.set_yticklabels((0.0, 1.0, 2.0, 3.0), fontsize=fontsize)
    ax.set_ylim(0, 3)
    ax.set_xlabel("$k_t$", fontsize=fontsize)
    ax.set_ylabel("$k_{t+1}$", fontsize=fontsize)
```

Let's look at the 45 degree diagram for two parameterizations.

```
params = create_solow_params()
fig, ax = plt.subplots(figsize=(8, 8))
plot_45(params, ax)
plt.show()
```



```
params = create_solow_params(α=0.05, δ=0.5)
fig, ax = plt.subplots(figsize=(8, 8))
plot_45(params, ax)
plt.show()
```



We see that k^* is indeed the unique positive fixed point.

Successive approximation

First let's compute the fixed point using successive approximation.

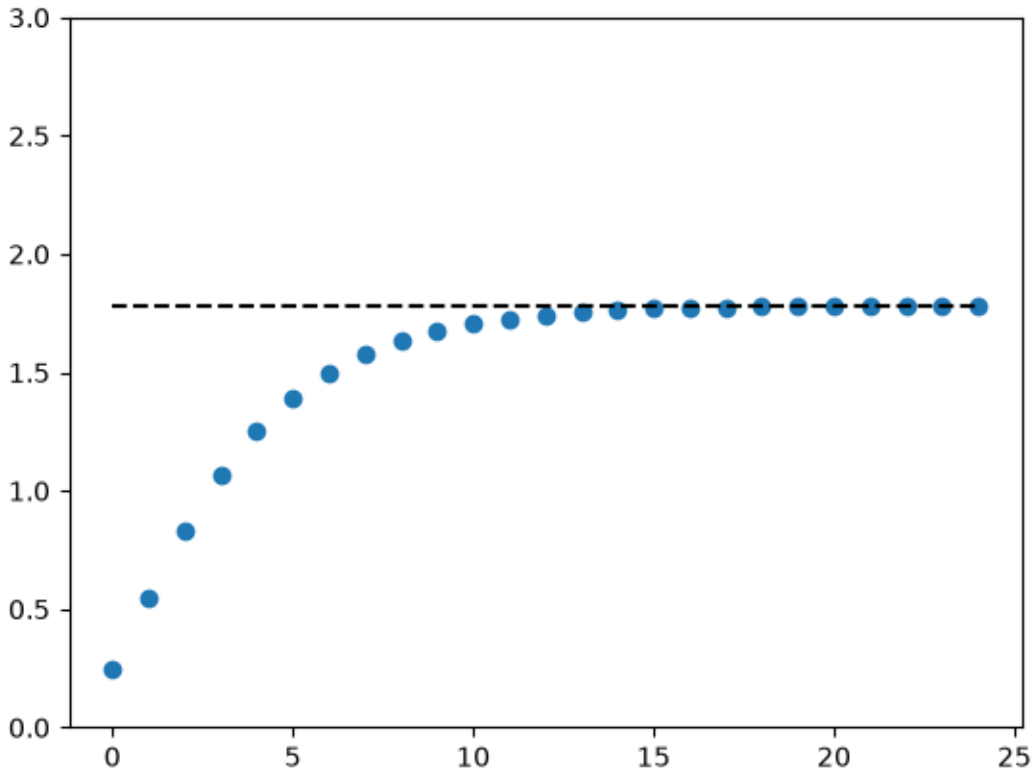
In this case, successive approximation means repeatedly updating capital from some initial state k_0 using the law of motion.

Here's a time series from a particular choice of k_0.

```
def compute_iterates(k_0, f, params, n=25):
    """Compute time series of length n generated by function f."""
    k = k_0
    k_iterates = []
    for t in range(n):
        k_iterates.append(k)
        k = f(k, params)
    return k_iterates
```

```
params = create_solow_params()
k_0 = 0.25
k_series = compute_iterates(k_0, g, params)
k_star = exact_fixed_point(params)

fig, ax = plt.subplots()
ax.plot(k_series, "o")
ax.plot([k_star] * len(k_series), "k--")
ax.set_ylim(0, 3)
plt.show()
```



Let's see the output for a long time series.

```
k_series = compute_iterates(k_0, g, params, n=10_000)
k_star_approx = k_series[-1]
k_star_approx
```

```
1.7846741842265788
```

This is close to the true value.

```
k_star
```

```
1.7846741842265788
```

Newton's method

In general, when applying Newton's fixed point method to some function g, we start with a guess x_0 of the fixed point and then update by solving for the fixed point of a tangent line at x_0.

To begin with, we recall that the first-order approximation of g at x_0 (i.e., the first order Taylor approximation of g at x_0) is the function

$$\hat{g}(x) \approx g(x_0) + g'(x_0)(x - x_0) \tag{7.2}$$

We solve for the fixed point of $\hat{g}$ by calculating the x_1 that solves

$$x_1 = \frac{g(x_0) - g'(x_0)x_0}{1 - g'(x_0)}$$

Generalising the process above, Newton's fixed point method iterates on

$$x_{t+1} = \frac{g(x_t) - g'(x_t)x_t}{1 - g'(x_t)}, \quad x_0 \text{ given} \tag{7.3}$$

To implement Newton's method we observe that the derivative of the law of motion for capital *(7.2.1)* is

$$g'(k) = \alpha s A k^{\alpha - 1} + (1 - \delta) \tag{7.4}$$

Let's define this:

```
def Dg(k, params):
    A, s, α, δ = params
    return α * A * s * k ** (α - 1) + (1 - δ)
```

Here's a function q representing *(7.2.3)*.

```
def q(k, params):
    return (g(k, params) - Dg(k, params) * k) / (1 - Dg(k, params))
```

Now let's plot some trajectories.

```
def plot_trajectories(
    params,
    k0_a=0.8,  # first initial condition
    k0_b=3.1,  # second initial condition
    n=20,      # length of time series
    fs=14,     # fontsize
```

(continues on next page)

(continued from previous page)

```
):

    fig, axes = plt.subplots(2, 1, figsize=(10, 6))
    ax1, ax2 = axes

    ks1 = compute_iterates(k0_a, g, params, n)
    ax1.plot(ks1, "-o", label="successive approximation")

    ks2 = compute_iterates(k0_b, g, params, n)
    ax2.plot(ks2, "-o", label="successive approximation")

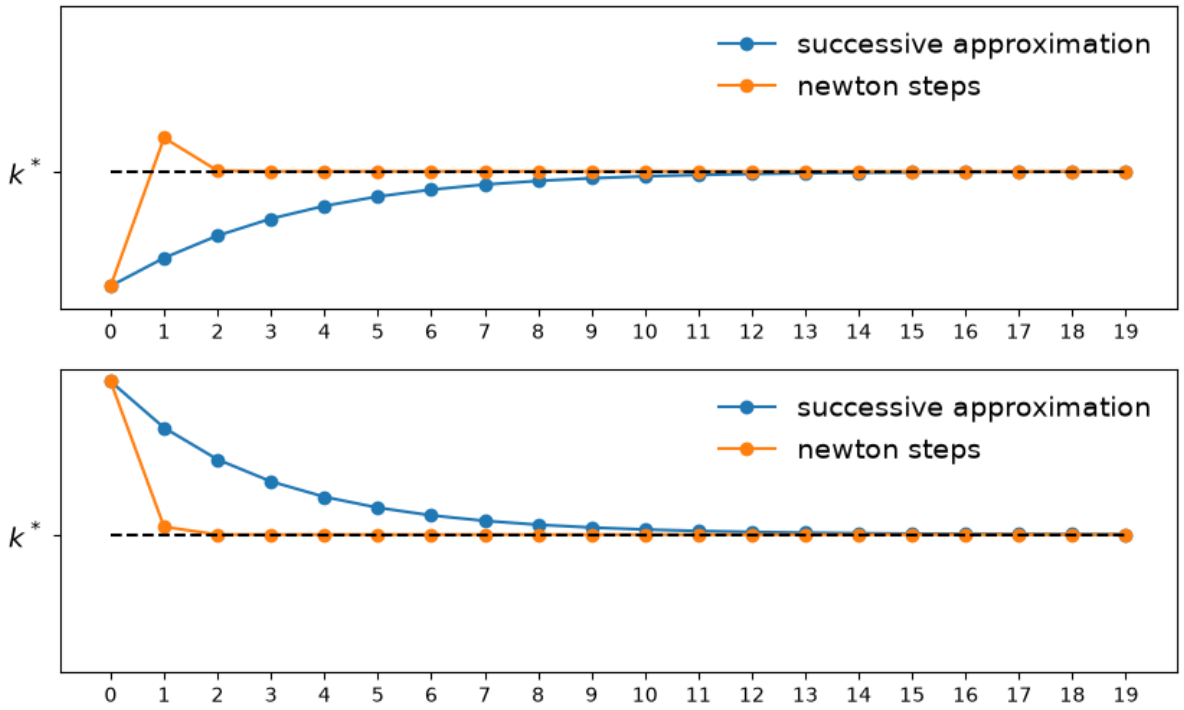
    ks3 = compute_iterates(k0_a, q, params, n)
    ax1.plot(ks3, "-o", label="newton steps")

    ks4 = compute_iterates(k0_b, q, params, n)
    ax2.plot(ks4, "-o", label="newton steps")

    for ax in axes:
        ax.plot(k_star * jnp.ones(n), "k--")
        ax.legend(fontsize=fs, frameon=False)
        ax.set_ylim(0.6, 3.2)
        ax.set_yticks((k_star,))
        ax.set_yticklabels(("$k^*$",), fontsize=fs)
        ax.set_xticks(jnp.linspace(0, 19, 20))

    plt.show()
```

```
params = create_solow_params()
plot_trajectories(params)
```



We can see that Newton's method converges faster than successive approximation.

7.3 Root-Finding in one dimension

In the previous section we computed fixed points.

In fact Newton's method is more commonly associated with the problem of finding zeros of functions.

Let's discuss this "root-finding" problem and then show how it is connected to the problem of finding fixed points.

7.3.1 Newton's method for zeros

Let's suppose we want to find an x such that $f(x) = 0$ for some smooth function f mapping real numbers to real numbers.

Suppose we have a guess x_0 and we want to update it to a new point x_1.

As a first step, we take the first-order approximation of f around x_0:

$$\hat{f}(x) \approx f(x_0) + f'(x_0)(x - x_0)$$

Now we solve for the zero of $\hat{f}$.

In particular, we set $\hat{f}(x_1) = 0$ and solve for x_1 to get

$$x_1 = x_0 - \frac{f(x_0)}{f'(x_0)}, \quad x_0 \text{ given}$$

Generalizing the formula above, for one-dimensional zero-finding problems, Newton's method iterates on

$$x_{t+1} = x_t - \frac{f(x_t)}{f'(x_t)}, \quad x_0 \text{ given} \tag{7.5}$$

The following code implements the iteration *(7.3.1)*

```
def newton(f, x_0, tol=1e-7, max_iter=100_000):
    x = x_0
    Df = jax.grad(f)

    # Implement the zero-finding formula
    @jax.jit
    def q(x):
        return x - f(x) / Df(x)

    error = tol + 1
    n = 0
    while error > tol:
        n += 1
        if n > max_iter:
            raise Exception("Max iteration reached without convergence")
        y = q(x)
        error = jnp.abs(x - y)
        x = y
        print(f"iteration {n}, error = {error:.5f}")
    return x.item()
```

Numerous libraries implement Newton's method in one dimension, including SciPy, so the code is just for illustrative purposes.

(That said, when we want to apply Newton's method using techniques such as automatic differentiation or GPU acceleration, it will be helpful to know how to implement Newton's method ourselves.)

7.3.2 Application to finding fixed points

Now consider again the Solow fixed-point calculation, where we solve for k satisfying $g(k) = k$.

We can convert to this to a zero-finding problem by setting $f(x) := g(x) - x$.

Any zero of f is clearly a fixed point of g.

Let's apply this idea to the Solow problem

```
params = create_solow_params()
k_star_approx_newton = newton(f = lambda x: g(x, params) - x, x_0=0.8)
```

```
iteration 1, error = 1.27209
iteration 2, error = 0.28180
iteration 3, error = 0.00561
iteration 4, error = 0.00000
iteration 5, error = 0.00000
```

```
k_star_approx_newton
```

```
1.7846741842265788
```

The result confirms convergence we saw in the graphs above: a very accurate result is reached with only 5 iterations.

7.4 Multivariate Newton's method

In this section, we introduce a two-good problem, present a visualization of the problem, and solve for the equilibrium of the two-good market using both a zero finder in `SciPy` and Newton's method.

We then expand the idea to a larger market with 5,000 goods and compare the performance of the two methods again.

We will see a significant performance gain when using Newton's method.

7.4.1 A two-goods market equilibrium

Let's start by computing the market equilibrium of a two-good problem.

We consider a market for two related products, good 0 and good 1, with price vector $p = (p_0, p_1)$

Supply of good i at price p is

$$q_i^s(p) = b_i \sqrt{p_i}$$

Demand of good i at price p is

$$q_i^d(p) = \exp(-(a_{i0} p_0 + a_{i1} p_1)) + c_i$$

Here c_i, b_i and a_{ij} are parameters.

For example, the two goods might be computer components that are typically used together, in which case they are complements. Hence demand depends on the price of both components.

The excess demand function is

$$e_i(p) = q_i^d(p) - q_i^s(p), \quad i = 0, 1$$

An equilibrium price vector p^* satisfies $e_i(p^*) = 0$.

We set

$$A = \begin{bmatrix} a_{00} & a_{01} \\ a_{10} & a_{11} \end{bmatrix}, \qquad b = \begin{bmatrix} b_0 \\ b_1 \end{bmatrix} \qquad \text{and} \qquad c = \begin{bmatrix} c_0 \\ c_1 \end{bmatrix}$$

for this particular question.

A graphical exploration

Since our problem is only two-dimensional, we can use graphical analysis to visualize and help understand the problem.

Our first step is to define the excess demand function

$$e(p) = \begin{bmatrix} e_0(p) \\ e_1(p) \end{bmatrix}$$

The function below calculates the excess demand for given parameters

```
@jax.jit
def e(p, A, b, c):
    return jnp.exp(-A @ p) + c - b * jnp.sqrt(p)
```

Our default parameter values will be

$$A = \begin{bmatrix} 0.5 & 0.4 \\ 0.8 & 0.2 \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 1 \end{bmatrix} \qquad \text{and} \qquad c = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$$

```
A = jnp.array([[0.5, 0.4], [0.8, 0.2]])
b = jnp.ones(2)
c = jnp.ones(2)
```

At a price level of $p = (1, 0.5)$, the excess demand is

```
p = jnp.array([1, 0.5])
ex_demand = e(p, A, b, c)

print(
    f"The excess demand for good 0 is {ex_demand[0]:.3f} \n"
    f"The excess demand for good 1 is {ex_demand[1]:.3f}"
)
```

```
The excess demand for good 0 is 0.497
The excess demand for good 1 is 0.699
```

To increase the efficiency of computation, we will use the power of vectorization using `jax.vmap`. This is much faster than the python loops.

```
# Create vectorization on the first axis of p.
e_vectorized_p_1 = jax.vmap(e, in_axes=(0, None, None, None))

# Create vectorization on the second axis of p.
e_vectorized = jax.vmap(e_vectorized_p_1, in_axes=(0, None, None, None))
```

Next we plot the two functions e_0 and e_1 on a grid of (p_0, p_1) values, using contour surfaces and lines.

We will use the following function to build the contour plots

```
def plot_excess_demand(ax, good=0, grid_size=100, grid_max=4, surface=True):
    p_grid = jnp.linspace(0, grid_max, grid_size)

    # Create meshgrid for all combinations of p_1 and p_2
    P1, P2 = jnp.meshgrid(p_grid, p_grid, indexing="ij")

    # Stack to create array of shape (grid_size, grid_size, 2)
    P = jnp.stack([P1, P2], axis=-1)

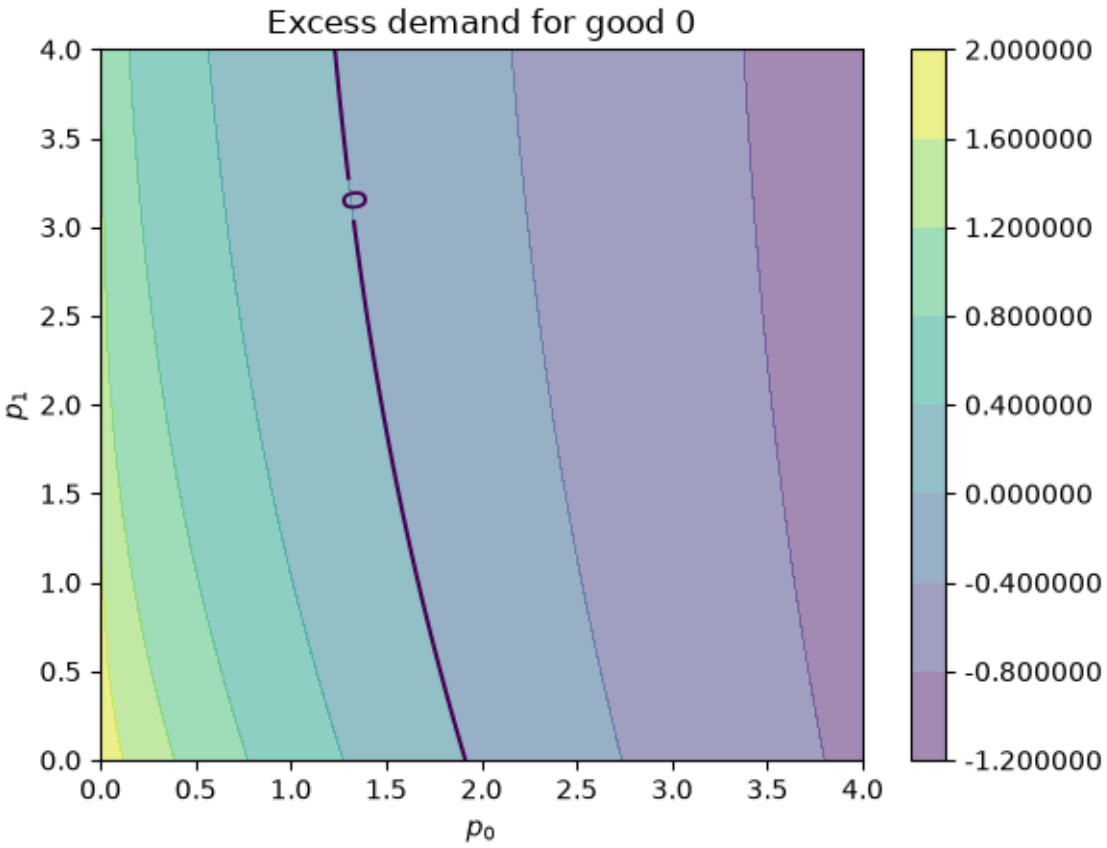
    # Compute all values at once using vectorized function
    z_full = e_vectorized(P, A, b, c)
    z = z_full[:, :, good]

    if surface:
        cs1 = ax.contourf(p_grid, p_grid, z.T, alpha=0.5)
        plt.colorbar(cs1, ax=ax, format="%.6f")

    ctr1 = ax.contour(p_grid, p_grid, z.T, levels=[0.0])
    ax.set_xlabel("$p_0$")
    ax.set_ylabel("$p_1$")
    ax.set_title(f"Excess demand for good {good}")
    plt.clabel(ctr1, inline=1, fontsize=13)
```

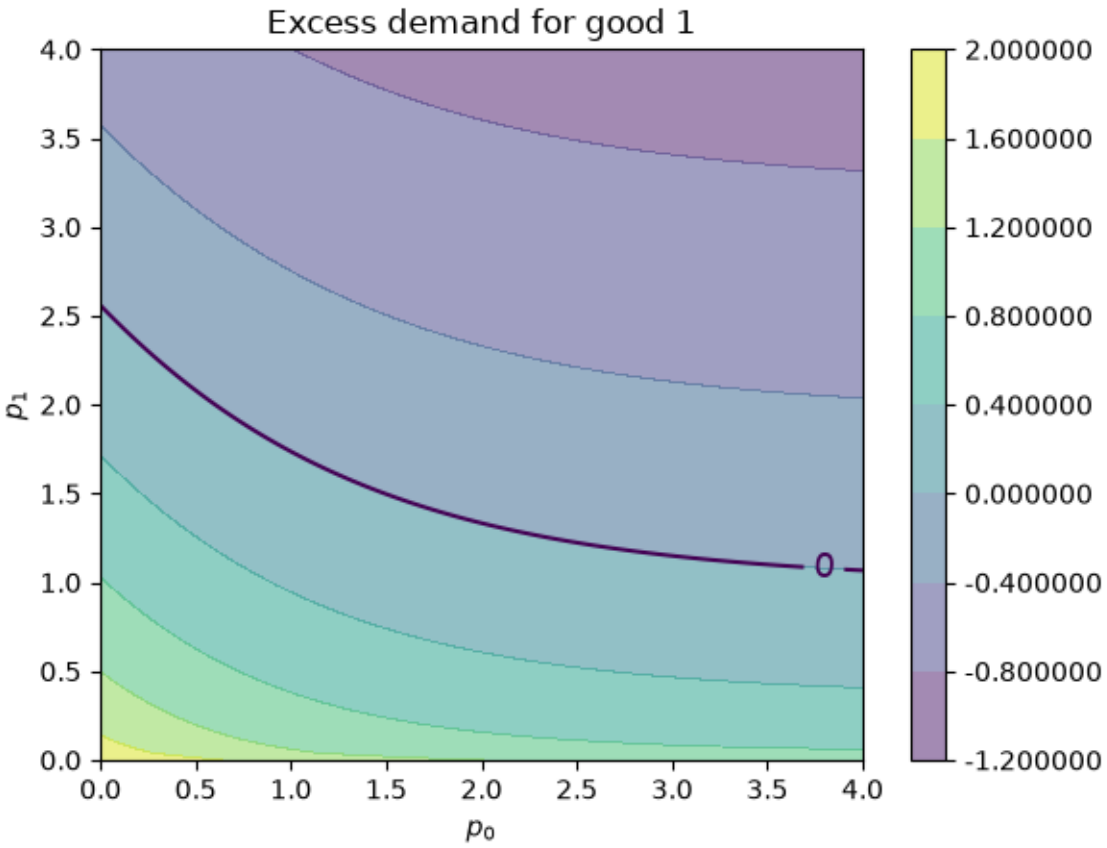
Here's our plot of e_0:

```
fig, ax = plt.subplots()
plot_excess_demand(ax, good=0)
plt.show()
```



Here's our plot of e_1:

```
fig, ax = plt.subplots()
plot_excess_demand(ax, good=1)
plt.show()
```

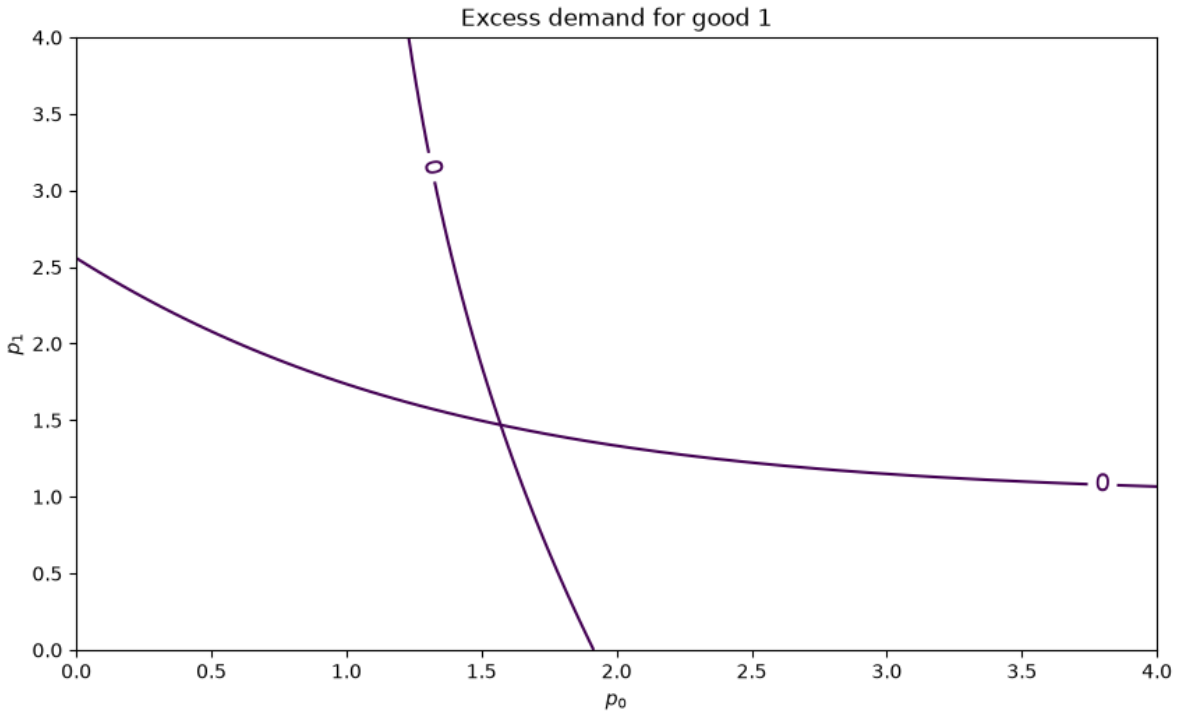


We see the black contour line of zero, which tells us when $e_i(p) = 0$.

For a price vector p such that $e_i(p) = 0$ we know that good i is in equilibrium (demand equals supply).

If these two contour lines cross at some price vector p^*, then p^* is an equilibrium price vector.

```
fig, ax = plt.subplots(figsize=(10, 5.7))
for good in (0, 1):
    plot_excess_demand(ax, good=good, surface=False)
plt.show()
```



It seems there is an equilibrium close to $p = (1.6, 1.5)$.

Using a multidimensional root finder

To solve for p^* more precisely, we use a zero-finding algorithm from `scipy.optimize`.

We supply $p = (1, 1)$ as our initial guess.

```
init_p = jnp.ones(2)
```

This uses the modified Powell method to find the zero

```
%%time
solution = root(lambda p: e(p, A, b, c), init_p, method="hybr")
```

```
CPU times: user 5.27 ms, sys: 2.48 ms, total: 7.75 ms
Wall time: 4.93 ms
```

Here's the resulting value:

```
p = solution.x
p
```

```
array([1.57080182, 1.46928838])
```

This looks close to our guess from observing the figure. We can plug it back into e to test that $e(p) \approx 0$:

```
e_p = jnp.max(jnp.abs(e(p, A, b, c)))
e_p.item()
```

```
2.0383694732117874e-13
```

This is indeed a very small error.

Adding gradient information

In many cases, for zero-finding algorithms applied to smooth functions, supplying the Jacobian of the function leads to better convergence properties.

Here, we manually calculate the elements of the Jacobian

$$J(p) = \begin{bmatrix} \frac{\partial e_0}{\partial p_0}(p) & \frac{\partial e_0}{\partial p_1}(p) \\ \frac{\partial e_1}{\partial p_0}(p) & \frac{\partial e_1}{\partial p_1}(p) \end{bmatrix}$$

```
def jacobian_e(p, A, b, c):
    p_0, p_1 = p
    a_00, a_01 = A[0, :]
    a_10, a_11 = A[1, :]
    j_00 = -a_00 * jnp.exp(-a_00 * p_0) - (b[0] / 2) * p_0 ** (-1 / 2)
    j_01 = -a_01 * jnp.exp(-a_01 * p_1)
    j_10 = -a_10 * jnp.exp(-a_10 * p_0)
    j_11 = -a_11 * jnp.exp(-a_11 * p_1) - (b[1] / 2) * p_1 ** (-1 / 2)
    J = [[j_00, j_01], [j_10, j_11]]
    return jnp.array(J)
```

```
%%time
solution = root(
    lambda p: e(p, A, b, c),
    init_p,
    jac = lambda p: jacobian_e(p, A, b, c),
    method="hybr",
)
```

```
CPU times: user 260 ms, sys: 42.2 ms, total: 302 ms
Wall time: 291 ms
```

Now the solution is even more accurate (although, in this low-dimensional problem, the difference is quite small):

```
p = solution.x
e_p = jnp.max(jnp.abs(e(p, A, b, c)))
e_p.item()
```

```
1.3322676295501878e-15
```

Using Newton's method

Now let's use Newton's method to compute the equilibrium price using the multivariate version of Newton's method

$$p_{n+1} = p_n - J_e(p_n)^{-1} e(p_n) \tag{7.6}$$

This is a multivariate version of *(7.3.1)*

(Here $J_e(p_n)$ is the Jacobian of e evaluated at p_n.)

The iteration starts from some initial guess of the price vector p_0.

Here, instead of coding Jacobian by hand, we use the `jacobian()` function in the `jax` library to auto-differentiate and calculate the Jacobian.

With only slight modification, we can generalize *our previous attempt* to multidimensional problems

```
def newton(f, x_0, tol=1e-5, max_iter=10):
    x = x_0
    f_jac = jax.jacobian(f)

    @jax.jit
    def q(x):
        return x - jnp.linalg.solve(f_jac(x), f(x))

    error = tol + 1
    n = 0
    while error > tol:
        n += 1
        if n > max_iter:
            raise Exception("Max iteration reached without convergence")
        y = q(x)
        if any(jnp.isnan(y)):
            raise Exception("Solution not found with NaN generated")
        error = jnp.linalg.norm(x - y)
        x = y
        print(f"iteration {n}, error = {error:.5f}")
    print("\n" + f"Result = {x} \n")
    return x
```

We find the algorithm terminates in 4 steps

```
%%time
p = newton(lambda p: e(p, A, b, c), init_p)
```

```
iteration 1, error = 0.62515
iteration 2, error = 0.11152
iteration 3, error = 0.00258
iteration 4, error = 0.00000

Result = [1.57080182 1.46928838]

CPU times: user 418 ms, sys: 55.5 ms, total: 473 ms
Wall time: 309 ms
```

```
e_p = jnp.max(jnp.abs(e(p, A, b, c)))
e_p.item()
```

```
1.4632739464559563e-13
```

The result is very accurate.

With the larger overhead, the speed is not better than the optimized `scipy` function.

7.4.2 A high-dimensional problem

Our next step is to investigate a large market with 3,000 goods.

The excess demand function is essentially the same, but now the matrix A is 3000×3000 and the parameter vectors b and c are 3000×1.

```
dim = 3000

# Create JAX random key
key = jax.random.key(0)

# Create a random matrix A and normalize the columns to sum to one
A = jax.random.uniform(key, (dim, dim))
s = jnp.sum(A, axis=0)
A = A / s

# Set up b and c
b = jnp.ones(dim)
c = jnp.ones(dim)
```

Here's our initial condition

```
init_p = jnp.ones(dim)
```

```
%%time
p = newton(lambda p: e(p, A, b, c), init_p)
```

```
iteration 1, error = 23.22262
iteration 2, error = 3.94537
iteration 3, error = 0.08500
iteration 4, error = 0.00004
iteration 5, error = 0.00000

Result = [1.50723773 1.51041603 1.50134795 ... 1.49941629 1.49033692 1.49666807]

CPU times: user 5 s, sys: 1.85 s, total: 6.85 s
Wall time: 6.84 s
```

```
e_p = jnp.max(jnp.abs(e(p, A, b, c)))
e_p.item()
```

```
4.440892098500626e-16
```

With the same tolerance, we compare the runtime and accuracy of Newton's method to SciPy's `root` function

```
%%time
solution = root(
    lambda p: e(p, A, b, c),
    init_p,
    jac = lambda p: jax.jacobian(e)(p, A, b, c),
    method="hybr",
    tol=1e-5,
)
```

```
CPU times: user 34.8 s, sys: 114 ms, total: 34.9 s
Wall time: 35.1 s
```

```
p = solution.x
e_p = jnp.max(jnp.abs(e(p, A, b, c)))
e_p.item()
```

```
9.209231102147442e-07
```

7.5 Exercises

Exercise 7.5.1

Consider a three-dimensional extension of the Solow fixed point problem with

$$A = \begin{bmatrix} 2 & 3 & 3 \\ 2 & 4 & 2 \\ 1 & 5 & 1 \end{bmatrix}, \quad s = 0.2, \quad \alpha = 0.5, \quad \delta = 0.8$$

As before the law of motion is

$$k_{t+1} = g(k_t) \quad \text{where} \quad g(k) := sAk^\alpha + (1 - \delta)k$$

However, k_t is now a 3×1 vector.

Solve for the fixed point using Newton's method with the following initial values:

$$k1_0 = (1, 1, 1)$$
$$k2_0 = (3, 5, 5)$$
$$k3_0 = (50, 50, 50)$$

Hint

- The computation of the fixed point is equivalent to computing k^* such that $g(k^*) - k^* = 0$.
- If you are unsure about your solution, you can start with the solved example:

$$A = \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{bmatrix}$$

with $s = 0.3$, $\alpha = 0.3$, and $\delta = 0.4$ and starting value:

$$k_0 = (1, 1, 1)$$

The result should converge to the *analytical solution*.

Solution

Let's first define the parameters for this problem

```
A = jnp.array([[2.0, 3.0, 3.0], [2.0, 4.0, 2.0], [1.0, 5.0, 1.0]])

s = 0.2
α = 0.5
δ = 0.8

initLs = [jnp.ones(3), jnp.array([3.0, 5.0, 5.0]), jnp.repeat(50.0, 3)]
```

Then define the multivariate version of the formula for the *(7.2.1)*

```
@jax.jit
def multivariate_solow(k, A=A, s=s, α=α, δ=δ):
    return s * jnp.dot(A, k**α) + (1 - δ) * k
```

Let's run through each starting value and see the output

```
attempt = 1
for init in initLs:
    print(f'Attempt {attempt}: Starting value is {init} \n')
    %time k = newton(lambda k: multivariate_solow(k) - k, \
                     init)
    print('-'*64)
    attempt += 1
```

```
Attempt 1: Starting value is [1. 1. 1.]
iteration 1, error = 50.49630
iteration 2, error = 41.10937
iteration 3, error = 4.29413
iteration 4, error = 0.38543
iteration 5, error = 0.00544
iteration 6, error = 0.00000

Result = [3.84058108 3.87071771 3.41091933]

CPU times: user 371 ms, sys: 57.7 ms, total: 428 ms
Wall time: 273 ms
----------------------------------------------------------------
Attempt 2: Starting value is [3. 5. 5.]

iteration 1, error = 2.07011
iteration 2, error = 0.12642
iteration 3, error = 0.00060
iteration 4, error = 0.00000

Result = [3.84058108 3.87071771 3.41091933]

CPU times: user 210 ms, sys: 38.3 ms, total: 248 ms
Wall time: 105 ms
----------------------------------------------------------------
Attempt 3: Starting value is [50. 50. 50.]
iteration 1, error = 73.00943
iteration 2, error = 6.49379
iteration 3, error = 0.68070
iteration 4, error = 0.01620
iteration 5, error = 0.00001
iteration 6, error = 0.00000

Result = [3.84058108 3.87071771 3.41091933]

CPU times: user 495 ms, sys: 68.5 ms, total: 563 ms
Wall time: 264 ms
```

```
--------------------------------------------------------------------
```

We find that the results are invariant to the starting values given the well-defined property of this question.

But the number of iterations it takes to converge is dependent on the starting values.

Let's substitute the output back into the formula to check our last result

```
multivariate_solow(k) - k
```

```
Array([0., 0., 0.], dtype=float64)
```

Note the error is very small.

We can also test our results on the known solution

```
A = jnp.array([[2.0, 0.0, 0.0], [0.0, 2.0, 0.0], [0.0, 0.0, 2.0]])

s = 0.3
α = 0.3
δ = 0.4

init = jnp.repeat(1.0, 3)
%%time

k = newton(lambda k: multivariate_solow(k, A=A, s=s, α=α, δ=δ) - k, init)
```

```
iteration 1, error = 1.57459
iteration 2, error = 0.21345
iteration 3, error = 0.00205
iteration 4, error = 0.00000

Result = [1.78467418 1.78467418 1.78467418]

CPU times: user 449 ms, sys: 61.1 ms, total: 510 ms
Wall time: 225 ms
```

The result is very close to the ground truth but still slightly different.

```
%%time

k = newton(
    lambda k: multivariate_solow(k, A=A, s=s, α=α, δ=δ) - k, init, tol=1e-7
)
```

```
iteration 1, error = 1.57459
iteration 2, error = 0.21345
iteration 3, error = 0.00205
iteration 4, error = 0.00000
iteration 5, error = 0.00000

Result = [1.78467418 1.78467418 1.78467418]

CPU times: user 456 ms, sys: 68.3 ms, total: 524 ms
Wall time: 223 ms
```

We can see it steps towards a more accurate solution.

Exercise 7.5.2

In this exercise, let's try different initial values and check how Newton's method responds to different starting points.

Let's define a three-good problem with the following default values:

$$A = \begin{bmatrix} 0.2 & 0.1 & 0.7 \\ 0.3 & 0.2 & 0.5 \\ 0.1 & 0.8 & 0.1 \end{bmatrix}, \qquad b = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix} \qquad \text{and} \qquad c = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}$$

For this exercise, use the following extreme price vectors as initial values:

$$p1_0 = (5, 5, 5)$$
$$p2_0 = (1, 1, 1)$$
$$p3_0 = (4.5, 0.1, 4)$$

Set the tolerance to $1e - 15$ for more accurate output.

Solution

Define parameters and initial values

```
A = jnp.array([[0.2, 0.1, 0.7], [0.3, 0.2, 0.5], [0.1, 0.8, 0.1]])
b = jnp.array([1.0, 1.0, 1.0])
c = jnp.array([1.0, 1.0, 1.0])

initLs = [jnp.repeat(5.0, 3), jnp.ones(3), jnp.array([4.5, 0.1, 4.0])]
```

Let's run through each initial guess and check the output

```
for attempt, init in enumerate(initLs, start=1):
    print(f"Attempt {attempt}: Starting value is {init} \n")
    %time p = newton(lambda p: e(p, A, b, c), init, tol=1e-15, max_iter=15)
    print("-" * 64)
```

```
Attempt 1: Starting value is [5. 5. 5.]
iteration 1, error = 9.24381
CPU times: user 456 ms, sys: 51.1 ms, total: 507 ms
Wall time: 224 ms
```

```
---------------------------------------------------------------------------
Exception                                 Traceback (most recent call last)
Cell In[52], line 3
      1 for attempt, init in enumerate(initLs, start=1):
      2     print(f"Attempt {attempt}: Starting value is {init} \n")
----> 3     get_ipython().run_line_magic('time', 'p = newton(lambda p: e(p, A,␣
 ↪b, c), init, tol=1e-15, max_iter=15)')
      4     print("-" * 64)

File <timed exec>:1
----> 1 'Could not get source, probably due dynamically evaluated source code.'

Cell In[35], line 17, in newton(f, x_0, tol, max_iter)
     13         if n > max_iter:
     14             raise Exception("Max iteration reached without convergence")
     15         y = q(x)
     16         if any(jnp.isnan(y)):
---> 17             raise Exception("Solution not found with NaN generated")
     18         error = jnp.linalg.norm(x - y)
     19         x = y
     20         print(f"iteration {n}, error = {error:.5f}")

Exception: Solution not found with NaN generated
```

We can see that Newton's method may fail for some starting values.

Sometimes it may take a few initial guesses to achieve convergence.

Substitute the result back to the formula to check our result using the second initial guess which converges

```
p_solution = newton(lambda p: e(p, A, b, c), initLs[1], tol=1e-15, max_iter=15)
e(p_solution, A, b, c)
```

```
iteration 1, error = 0.73419
iteration 2, error = 0.12472
iteration 3, error = 0.00269
iteration 4, error = 0.00000
iteration 5, error = 0.00000
iteration 6, error = 0.00000

Result = [1.49744442 1.49744442 1.49744442]
```

```
Array([0.00000000e+00, 0.00000000e+00, 2.22044605e-16], dtype=float64)
```

We can see the result is very accurate.

Part II

Elementary Statistics

CHAPTER

EIGHT

ELEMENTARY PROBABILITY WITH MATRICES

This lecture uses matrix algebra to illustrate some basic ideas about probability theory.

After introducing underlying objects, we'll use matrices and vectors to describe probability distributions.

Among concepts that we'll be studying include

- a joint probability distribution
- marginal distributions associated with a given joint distribution
- conditional probability distributions
- statistical independence of two random variables
- joint distributions associated with a prescribed set of marginal distributions
  - couplings
  - copulas
- the probability distribution of a sum of two independent random variables
  - convolution of marginal distributions
- parameters that define a probability distribution
- sufficient statistics as data summaries

We'll use a matrix to represent a bivariate or multivariate probability distribution and a vector to represent a univariate probability distribution

This *companion lecture* describes some popular probability distributions and describes how to use Python to sample from them.

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install prettytable
```

As usual, we'll start with some imports

```
import numpy as np
import matplotlib.pyplot as plt
import prettytable as pt
from scipy import stats
from scipy.special import comb
from mpl_toolkits.mplot3d import Axes3D
from matplotlib_inline.backend_inline import set_matplotlib_formats
set_matplotlib_formats('retina')

rng = np.random.default_rng(0)
```

8.1 Sketch of basic concepts

We'll briefly define what we mean by a **probability space**, a **probability measure**, and a **random variable**.

For most of this lecture, we sweep these objects into the background

> **Note**
>
> Nevertheless, they'll be lurking beneath **induced distributions** of random variables that we'll focus on here.
>
> These deeper objects are essential for defining and analysing the concepts of stationarity and ergodicity that underly laws of large numbers.
>
> For a relatively nontechnical presentation of some of these results see this chapter from Lars Peter Hansen and Thomas J. Sargent's online monograph titled *Risk, Uncertainty, and Values*.

Let Ω be a set of possible underlying outcomes and let $\omega \in \Omega$ be a particular underlying outcome.

Let $\mathcal{F}$ be a collection of subsets of Ω that we call **events**.

(Technically, $\mathcal{F}$ is a σ-algebra.)

A **probability measure** μ maps each event $\mathcal{G} \in \mathcal{F}$ into a scalar number $\mu(\mathcal{G})$ between 0 and 1, with $\mu(\Omega) = 1$.

The triple $\Omega, \mathcal{F}, \mu$ forms our **probability space**.

A **random variable** $X(\omega)$ is a function of the underlying outcome $\omega \in \Omega$ that assigns a value in some set of possible values.

If A is a set of possible values of X, then the event that X lies in A is

$$\mathcal{G} = \{\omega \in \Omega : X(\omega) \in A\}.$$

The random variable $X(\omega)$ has a **probability distribution** induced by the probability measure μ:

$$\text{Prob}(X \in A) = \mu(\mathcal{G}).$$

If μ has a density $p(\omega)$, then we can also write

$$\text{Prob}(X \in A) = \int_{\mathcal{G}} p(\omega)\, d\omega \tag{8.1}$$

We call this the induced probability distribution of random variable X.

Instead of working explicitly with an underlying probability space $\Omega, \mathcal{F}$ and probability measure μ, applied statisticians often proceed simply by specifying a form for an induced distribution for a random variable X.

That is how we'll proceed in this lecture and in many subsequent lectures.

8.2 What does probability mean?

Before diving in, we'll say a few words about what probability theory means and how it connects to statistics.

We also touch on these topics in *Two Meanings of Probability* and *Bayesian versus Frequentist Decision Rules*.

For much of this lecture we'll be discussing fixed "population" probabilities.

These are purely mathematical objects.

To appreciate how statisticians connect probabilities to data, the key is to understand the following concepts:

- A single draw from a probability distribution
- Repeated independently and identically distributed (i.i.d.) draws of "samples" or "realizations" from the same probability distribution
- A **statistic** defined as a function of a sequence of samples
- An **empirical distribution** or **histogram** (a binned empirical distribution) that records observed **relative frequencies**
- The idea that a population probability distribution is what we anticipate **relative frequencies** will be in a long sequence of i.i.d. draws. Here the following mathematical machinery makes precise what is meant by **anticipated relative frequencies**
  - **Law of Large Numbers (LLN)**
  - **Central Limit Theorem (CLT)**

8.2.1 A discrete random variable example

Scalar example

Let X be a scalar random variable that takes on the I possible values $0, 1, 2, \ldots, I-1$ with probabilities

$$\text{Prob}(X = i) = f_i,$$

where

$$f_i \geqslant 0, \quad \sum_i f_i = 1.$$

We sometimes write

$$X \sim \{f_i\}_{i=0}^{I-1}$$

as a short-hand way of saying that the random variable X is described by the probability distribution $\{f_i\}_{i=0}^{I-1}$.

Consider drawing a sample $x_0, x_1, \ldots, x_{N-1}$ of N independent and identically distributed draws of X.

What do "identical" and "independent" mean in IID or iid ("identically and independently distributed")?

- "identical" means that each draw is from the same distribution.
- "independent" means that the joint distribution equals the product of marginal distributions, i.e.,

$$\begin{aligned}\text{Prob}\{x_0 = i_0, x_1 = i_1, \ldots, x_{N-1} = i_{N-1}\} &= \text{Prob}\{x_0 = i_0\} \cdot \cdots \cdot \text{Prob}\{x_{N-1} = i_{N-1}\} \\ &= f_{i_0} f_{i_1} \cdot \cdots \cdot f_{i_{N-1}}\end{aligned}$$

We define an **empirical distribution** as follows.

For each $i = 0, \ldots, I-1$, let

$$\begin{aligned}N_i &= \text{number of times } X = i, \\ N &= \sum_{i=0}^{I-1} N_i \quad \text{total number of draws}, \\ \tilde{f}_i &= \frac{N_i}{N} \sim \text{ frequency of draws for which } X = i\end{aligned}$$

Key concepts that connect probability theory with statistics are laws of large numbers and central limit theorems.

A Law of Large Numbers (LLN) states that $\tilde{f}_i \to f_i$ as $N \to \infty$.

A Central Limit Theorem (CLT) describes a **rate** at which $\tilde{f}_i \to f_i$.

See *LLN and CLT* for a detailed treatment of both results.

8.2.2 Understanding probability: frequentist vs. Bayesian

For "frequentist" statisticians, **anticipated relative frequency** is **all** that a probability distribution means.

But for a Bayesian it means something else – something partly subjective and purely personal.

We say "partly" because a Bayesian also pays attention to relative frequencies.

8.3 Representing probability distributions

A probability distribution $\text{Prob}(X \in A)$ can be described by its **cumulative distribution function (CDF)**

$$F_X(x) = \text{Prob}\{X \leq x\}.$$

Sometimes, but not always, a random variable can also be described by a **density function** $f(x)$ that is related to its CDF by

$$\text{Prob}\{X \in B\} = \int_{t \in B} f(t)dt$$

$$F(x) = \int_{-\infty}^{x} f(t)dt$$

Here B is a set of possible X's whose probability of occurring we want to compute.

When a probability density exists, a probability distribution can be characterized either by its CDF or by its density.

For a **discrete-valued** random variable

- the number of possible values of X is finite or countably infinite
- we replace a **density** with a **probability mass function**, a non-negative sequence that sums to one
- when a density exists, we replace integration with summation in formulas like (8.1)

In this lecture, we mostly discuss discrete random variables.

Doing this enables us to confine our tool set basically to linear algebra.

Later we'll briefly discuss how to approximate a continuous random variable with a discrete random variable.

8.4 Univariate probability distributions

We'll devote most of this lecture to discrete-valued random variables, but we'll say a few things about continuous-valued random variables.

8.4.1 Discrete random variable

Let X be a discrete random variable that takes possible values: $i = 0, 1, \ldots, I-1 = \bar{X}$.

Here, we choose the maximum index $I-1$ because of how this aligns nicely with Python's index convention.

Define $f_i \equiv \text{Prob}\{X = i\}$ and assemble the non-negative vector

$$f = \begin{bmatrix} f_0 \\ f_1 \\ \vdots \\ f_{I-1} \end{bmatrix} \tag{8.2}$$

for which $f_i \in [0,1]$ for each i and $\sum_{i=0}^{I-1} f_i = 1$.

This vector defines a **probability mass function**.

The distribution (8.2) has **parameters** $\{f_i\}_{i=0,1,\cdots,I-2}$ since $f_{I-1} = 1 - \sum_{i=0}^{I-2} f_i$.

These parameters pin down the shape of the distribution.

(Sometimes $I = \infty$.)

Such a "non-parametric" distribution has as many "parameters" as there are possible values of the random variable.

We often work with special distributions that are characterized by a small number parameters.

In these special parametric distributions,

$$f_i = g(i;\theta)$$

where θ is a vector of parameters that is of much smaller dimension than I.

A **statistical model** is a joint probability distribution characterized by a list of **parameters**.

The concept of **parameter** is intimately related to the notion of **sufficient statistic**.

A **statistic** is a nonlinear function of a data set.

Sufficient statistics summarize all **information** that a data set contains about parameters of a statistical model.

Note that a sufficient statistic corresponds to a particular statistical model.

Sufficient statistics are key tools that AI uses to summarize or compress a **big data** set.

R. A. Fisher provided a rigorous definition of **information** – see Fisher information.

An example of a parametric probability distribution is a **geometric distribution**.

It is described by

$$f_i = \text{Prob}\{X = i\} = (1-\lambda)\lambda^i, \quad \lambda \in [0,1), \quad i = 0,1,2,...$$

Evidently, $\sum_{i=0}^{\infty} f_i = 1$.

Let θ be a vector of parameters of the distribution described by f, then

$$f_i(\theta) \geq 0, \sum_{i=0}^{\infty} f_i(\theta) = 1$$

8.4.2 Continuous random variable

Let X be a continuous random variable that takes values in a set $\tilde{X} \subseteq \mathbb{R}$ and whose distribution has parameters θ.

$$\text{Prob}\{X \in A\} = \int_{x \in A} f(x;\theta)\, dx; \quad f(x;\theta) \geq 0$$

where A is a subset of $\tilde{X}$ and

$$\text{Prob}\{X \in \tilde{X}\} = 1$$

8.5 Bivariate probability distributions

We'll now discuss a bivariate **joint distribution**.

To begin, we restrict ourselves to two discrete random variables.

Let X, Y be two discrete random variables that take values:

$$X \in \{0, \ldots, I-1\}$$

$$Y \in \{0, \ldots, J-1\}$$

Then their **joint distribution** is described by a matrix

$$F_{I \times J} = [f_{ij}]_{i \in \{0,\ldots,I-1\}, j \in \{0,\ldots,J-1\}}$$

whose elements are

$$f_{ij} = \text{Prob}\{X = i, Y = j\} \geq 0$$

where

$$\sum_i \sum_j f_{ij} = 1$$

8.6 Marginal probability distributions

The joint distribution induces marginal distributions

$$\text{Prob}\{X = i\} = \sum_{j=0}^{J-1} f_{ij} = \mu_i, \quad i = 0, \ldots, I-1$$

$$\text{Prob}\{Y = j\} = \sum_{i=0}^{I-1} f_{ij} = \nu_j, \quad j = 0, \ldots, J-1$$

For example, let a joint distribution over (X, Y) be

$$F = \begin{bmatrix} .25 & .1 \\ .15 & .5 \end{bmatrix} \tag{8.3}$$

The implied marginal distributions are:

$$\begin{aligned} \text{Prob}\{X = 0\} &= .25 + .1 = .35 \\ \text{Prob}\{X = 1\} &= .15 + .5 = .65 \\ \text{Prob}\{Y = 0\} &= .25 + .15 = .4 \\ \text{Prob}\{Y = 1\} &= .1 + .5 = .6 \end{aligned}$$

As a digression, if two random variables X, Y are continuous and have joint density $f(x, y)$, then marginal distributions can be computed by

$$f(x) = \int_{\mathbb{R}} f(x, y) dy$$

$$f(y) = \int_{\mathbb{R}} f(x, y) dx$$

8.7 Conditional probability distributions

Conditional probabilities are defined according to

$$\text{Prob}\{A \mid B\} = \frac{\text{Prob}\{A \cap B\}}{\text{Prob}\{B\}}$$

where A, B are two events.

For a pair of discrete random variables, we have the **conditional distribution**

$$\text{Prob}\{X = i | Y = j\} = \frac{f_{ij}}{\sum_i f_{ij}} = \frac{\text{Prob}\{X = i, Y = j\}}{\text{Prob}\{Y = j\}}$$

where $i = 0, \ldots, I-1, \quad j = 0, \ldots, J-1.$

Note that

$$\sum_i \text{Prob}\{X = i | Y = j\} = \frac{\sum_i f_{ij}}{\sum_i f_{ij}} = 1$$

The mathematics of conditional probability implies:

$$\text{Prob}\{X = i | Y = j\} = \frac{\text{Prob}\{X = i, Y = j\}}{\text{Prob}\{Y = j\}} = \frac{\text{Prob}\{Y = j | X = i\}\text{Prob}\{X = i\}}{\text{Prob}\{Y = j\}} \tag{8.4}$$

Note

Formula (8.4) is also what a Bayesian calls **Bayes' Law**.

A Bayesian statistician regards marginal probability distribution $\text{Prob}(X = i), i = 0, \ldots, I-1$ as a **prior** distribution that describes his personal subjective beliefs about X.

He then interprets formula (8.4) as a procedure for constructing a **posterior** distribution that describes how he would revise his subjective beliefs after observing that Y equals j.

For the joint distribution (8.3)

$$\text{Prob}\{X = 0 | Y = 1\} = \frac{.1}{.1 + .5} = \frac{.1}{.6}$$

8.8 Transition probability matrix

Consider the following joint probability distribution of two random variables.

Let X, Y be discrete random variables with joint distribution

$$\text{Prob}\{X = i, Y = j\} = \rho_{ij}$$

where $i = 0, \ldots, I-1; j = 0, \ldots, J-1$ and

$$\sum_i \sum_j \rho_{ij} = 1, \quad \rho_{ij} \geqslant 0.$$

An associated conditional distribution is

$$\text{Prob}\{Y = j | X = i\} = \frac{\rho_{ij}}{\sum_j \rho_{ij}} = \frac{\text{Prob}\{Y = j, X = i\}}{\text{Prob}\{X = i\}}$$

We can define a transition probability matrix P with i, j component

$$p_{ij} = \text{Prob}\{Y = j | X = i\} = \frac{\rho_{ij}}{\sum_j \rho_{ij}}$$

where

$$\begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix}$$

The first row is the probability that $Y = j, j = 0, 1$ conditional on $X = 0$.

The second row is the probability that $Y = j, j = 0, 1$ conditional on $X = 1$.

Note that

- $\sum_j p_{ij} = \frac{\sum_j \rho_{ij}}{\sum_j \rho_{ij}} = 1$, so each row of the transition matrix P is a probability distribution (not so for each column).

8.9 Application: forecasting a time series

Suppose that there are two time periods.

- $t = 0$ "today"
- $t = 1$ "tomorrow"

Let $X(0)$ be a random variable to be realized at $t = 0$, $X(1)$ be a random variable to be realized at $t = 1$.

Suppose that

$$\text{Prob}\{X(0) = i, X(1) = j\} = f_{ij} \geq 0, \quad i = 0, \cdots, I-1, \quad j = 0, \cdots, J-1$$
$$\sum_i \sum_j f_{ij} = 1$$

f_{ij} is a joint distribution over $[X(0), X(1)]$.

A conditional distribution is

$$\text{Prob}\{X(1) = j | X(0) = i\} = \frac{f_{ij}}{\sum_j f_{ij}}$$

This formula is a workhorse for applied economic forecasters.

8.10 Statistical independence

Random variables X and Y are statistically **independent** if

$$\text{Prob}\{X = i, Y = j\} = f_i g_j$$

where

$$\text{Prob}\{X = i\} = f_i \geq 0, \quad \sum_i f_i = 1$$
$$\text{Prob}\{Y = j\} = g_j \geq 0, \quad \sum_j g_j = 1$$

Conditional distributions are

$$\text{Prob}\{X = i|Y = j\} = \frac{f_i g_j}{\sum_i f_i g_j} = \frac{f_i g_j}{g_j} = f_i$$
$$\text{Prob}\{Y = j|X = i\} = \frac{f_i g_j}{\sum_j f_i g_j} = \frac{f_i g_j}{f_i} = g_j$$

8.11 Means and variances

The mean and variance of a discrete random variable X are

$$\mu_X \equiv \mathbb{E}[X] = \sum_k k\text{Prob}\{X = k\}$$
$$\sigma_X^2 \equiv \mathbb{D}[X] = \sum_k (k - \mathbb{E}[X])^2 \text{Prob}\{X = k\}$$

A continuous random variable having density $f_X(x)$ has mean and variance

$$\mu_X \equiv \mathbb{E}[X] = \int_{-\infty}^{\infty} x f_X(x) dx$$
$$\sigma_X^2 \equiv \mathbb{D}[X] = \text{E}\left[(X - \mu_X)^2\right] = \int_{-\infty}^{\infty} (x - \mu_X)^2 f_X(x) dx$$

8.12 Matrix representations of some bivariate distributions

Let's use matrices to represent a joint distribution, conditional distribution, marginal distribution, and the mean and variance of a bivariate random variable.

The table below illustrates a probability distribution for a bivariate random variable.

$$F = [f_{ij}] = \begin{bmatrix} 0.3 & 0.2 \\ 0.1 & 0.4 \end{bmatrix}$$

Marginal distributions are

$$\text{Prob}(X = i) = \sum_j f_{ij} = u_i$$

$$\text{Prob}(Y = j) = \sum_i f_{ij} = v_j$$

Let's write some Python code that lets us draw some long samples and compute relative frequencies.

The code lets us check whether the "sampling" distribution agrees with the "population" distribution – confirming that the population distribution correctly tells us the relative frequencies that we should expect in a large sample.

```
# specify parameters
xs = np.array([0, 1])
ys = np.array([10, 20])
f = np.array([[0.3, 0.2], [0.1, 0.4]])
f_cum = np.cumsum(f)

# draw random numbers
```

(continues on next page)

(continued from previous page)

```
p = rng.random(1_000_000)
x = np.vstack([xs[1]*np.ones(p.shape), ys[1]*np.ones(p.shape)])
# map to the bivariate distribution

x[0, p < f_cum[2]] = xs[1]
x[1, p < f_cum[2]] = ys[0]

x[0, p < f_cum[1]] = xs[0]
x[1, p < f_cum[1]] = ys[1]

x[0, p < f_cum[0]] = xs[0]
x[1, p < f_cum[0]] = ys[0]
print(x)
```

```
[[ 1.  0.  0. ...  0.  0.  0.]
 [20. 10. 10. ... 10. 10. 20.]]
```

Note

To generate random draws from the joint distribution F, we use the inverse CDF technique described in *this companion lecture*.

```
# marginal distribution
xp = np.sum(x[0, :] == xs[0])/1_000_000
yp = np.sum(x[1, :] == ys[0])/1_000_000

# print output
print("marginal distribution for x")
xmtb = pt.PrettyTable()
xmtb.field_names = ['x_value', 'x_prob']
xmtb.add_row([xs[0], xp])
xmtb.add_row([xs[1], 1-xp])
print(xmtb)

print("\nmarginal distribution for y")
ymtb = pt.PrettyTable()
ymtb.field_names = ['y_value', 'y_prob']
ymtb.add_row([ys[0], yp])
ymtb.add_row([ys[1], 1-yp])
print(ymtb)
```

```
marginal distribution for x
+---------+----------+
| x_value |  x_prob  |
+---------+----------+
|    0    | 0.500194 |
|    1    | 0.499806 |
+---------+----------+

marginal distribution for y
+---------+----------+
| y_value |  y_prob  |
+---------+----------+
|    10   | 0.399741 |
```

(continues on next page)

(continued from previous page)

```
|    20   | 0.600259 |
+---------+----------+
```

```
# conditional distributions
xc1 = x[0, x[1, :] == ys[0]]
xc2 = x[0, x[1, :] == ys[1]]
yc1 = x[1, x[0, :] == xs[0]]
yc2 = x[1, x[0, :] == xs[1]]

xc1p = np.sum(xc1 == xs[0])/len(xc1)
xc2p = np.sum(xc2 == xs[0])/len(xc2)
yc1p = np.sum(yc1 == ys[0])/len(yc1)
yc2p = np.sum(yc2 == ys[0])/len(yc2)

# print output
print("conditional distribution for x")
xctb = pt.PrettyTable()
xctb.field_names = ['y_value', 'prob(x=0)', 'prob(x=1)']
xctb.add_row([ys[0], xc1p, 1-xc1p])
xctb.add_row([ys[1], xc2p, 1-xc2p])
print(xctb)

print("\nconditional distribution for y")
yctb = pt.PrettyTable()
yctb.field_names = ['x_value',  'prob(y=10)', 'prob(y=20)']
yctb.add_row([xs[0], yc1p, 1-yc1p])
yctb.add_row([xs[1], yc2p, 1-yc2p])
print(yctb)
```

```
conditional distribution for x
+---------+--------------------+---------------------+
| y_value |     prob(x=0)      |      prob(x=1)      |
+---------+--------------------+---------------------+
|    10   | 0.7504634250677313 | 0.24953657493226866 |
|    20   | 0.333527693878809  |  0.666472306121191  |
+---------+--------------------+---------------------+

conditional distribution for y
+---------+--------------------+---------------------+
| x_value |      prob(y=10)    |      prob(y=20)     |
+---------+--------------------+---------------------+
|    0    |  0.5997492972726582 | 0.40025070272734176 |
|    1    | 0.19957743604518552 |  0.8004225639548145 |
+---------+--------------------+---------------------+
```

Let's calculate population marginal and conditional probabilities using matrix algebra.

$$
\begin{bmatrix}
 & \vdots & y_1 & y_2 & \vdots & x \\
\cdots & \vdots & \cdots & \cdots & \vdots & \cdots \\
x_1 & \vdots & 0.3 & 0.2 & \vdots & 0.5 \\
x_2 & \vdots & 0.1 & 0.4 & \vdots & 0.5 \\
\cdots & \vdots & \cdots & \cdots & \vdots & \cdots \\
y & \vdots & 0.4 & 0.6 & \vdots & 1
\end{bmatrix}
$$

$\Longrightarrow$

(1) Marginal distribution:

$$\begin{bmatrix} var & \vdots & var_1 & var_2 \\ \cdots & \vdots & \cdots & \cdots \\ x & \vdots & 0.5 & 0.5 \\ \cdots & \vdots & \cdots & \cdots \\ y & \vdots & 0.4 & 0.6 \end{bmatrix}$$

(2) Conditional distribution:

$$\begin{bmatrix} x & \vdots & x_1 & x_2 \\ \cdots\cdots & \vdots & \cdots\cdots & \cdots\cdots \\ y = y_1 & \vdots & \frac{0.3}{0.4} = 0.75 & \frac{0.1}{0.4} = 0.25 \\ \cdots\cdots & \vdots & \cdots\cdots & \cdots\cdots \\ y = y_2 & \vdots & \frac{0.2}{0.6} \approx 0.33 & \frac{0.4}{0.6} \approx 0.67 \end{bmatrix}$$

$$\begin{bmatrix} y & \vdots & y_1 & y_2 \\ \cdots\cdots & \vdots & \cdots\cdots & \cdots\cdots \\ x = x_1 & \vdots & \frac{0.3}{0.5} = 0.6 & \frac{0.2}{0.5} = 0.4 \\ \cdots\cdots & \vdots & \cdots\cdots & \cdots\cdots \\ x = x_2 & \vdots & \frac{0.1}{0.5} = 0.2 & \frac{0.4}{0.5} = 0.8 \end{bmatrix}$$

These population objects closely resemble the sample counterparts computed above.

Let's wrap some of the functions we have used in a Python class that will let us generate and sample from a discrete bivariate joint distribution.

```
class discrete_bijoint:

    def __init__(self, f, xs, ys):
        '''initialization
        -----------------
        parameters:
        f: the bivariate joint probability matrix
        xs: values of x vector
        ys: values of y vector
        '''
        self.f, self.xs, self.ys = f, xs, ys

    def joint_tb(self):
        '''print the joint distribution table'''
        xs = self.xs
        ys = self.ys
        f = self.f
        jtb = pt.PrettyTable()
        jtb.field_names = ['x_value/y_value', *ys, 'marginal sum for x']
        for i in range(len(xs)):
            jtb.add_row([xs[i], *f[i, :], np.sum(f[i, :])])
        jtb.add_row(['marginal_sum for y', *np.sum(f, 0), np.sum(f)])
        print("\nThe joint probability distribution for x and y\n", jtb)
        self.jtb = jtb

    def draw(self, n):
        '''draw random numbers
        ----------------------
        parameters:
        n: number of random numbers to draw
        '''
```

(continues on next page)

(continued from previous page)

```
        xs = self.xs
        ys = self.ys
        f_cum = np.cumsum(self.f)
        p = rng.random(n)
        x = np.empty([2, p.shape[0]])
        lf = len(f_cum)
        lx = len(xs)-1
        ly = len(ys)-1
        for i in range(lf):
            x[0, p < f_cum[lf-1-i]] = xs[lx]
            x[1, p < f_cum[lf-1-i]] = ys[ly]
            if ly == 0:
                lx -= 1
                ly = len(ys)-1
            else:
                ly -= 1
        self.x = x
        self.n = n

    def marg_dist(self):
        '''marginal distribution'''
        x = self.x
        xs = self.xs
        ys = self.ys
        n = self.n
        xmp = [np.sum(x[0, :] == xs[i])/n for i in range(len(xs))]
        ymp = [np.sum(x[1, :] == ys[i])/n for i in range(len(ys))]

        # print output
        xmtb = pt.PrettyTable()
        ymtb = pt.PrettyTable()
        xmtb.field_names = ['x_value', 'x_prob']
        ymtb.field_names = ['y_value', 'y_prob']
        for i in range(max(len(xs), len(ys))):
            if i < len(xs):
                xmtb.add_row([xs[i], xmp[i]])
            if i < len(ys):
                ymtb.add_row([ys[i], ymp[i]])
        xmtb.add_row(['sum', np.sum(xmp)])
        ymtb.add_row(['sum', np.sum(ymp)])
        print("\nmarginal distribution for x\n", xmtb)
        print("\nmarginal distribution for y\n", ymtb)

        self.xmp = xmp
        self.ymp = ymp

    def cond_dist(self):
        '''conditional distribution'''
        x = self.x
        xs = self.xs
        ys = self.ys
        n = self.n
        xcp = np.empty([len(ys), len(xs)])
        ycp = np.empty([len(xs), len(ys)])
        for i in range(max(len(ys), len(xs))):
            if i < len(ys):
                xi = x[0, x[1, :] == ys[i]]
```

(continues on next page)

(continued from previous page)

```
                idx = xi.reshape(len(xi), 1) == xs.reshape(1, len(xs))
                xcp[i, :] = np.sum(idx, 0)/len(xi)
            if i < len(xs):
                yi = x[1, x[0, :] == xs[i]]
                idy = yi.reshape(len(yi), 1) == ys.reshape(1, len(ys))
                ycp[i, :] = np.sum(idy, 0)/len(yi)

        # print output
        xctb = pt.PrettyTable()
        yctb = pt.PrettyTable()
        xctb.field_names = ['x_value', *xs, 'sum']
        yctb.field_names = ['y_value', *ys, 'sum']
        for i in range(max(len(xs), len(ys))):
            if i < len(ys):
                xctb.add_row([ys[i], *xcp[i], np.sum(xcp[i])])
            if i < len(xs):
                yctb.add_row([xs[i], *ycp[i], np.sum(ycp[i])])
        print("\nconditional distribution for x\n", xctb)
        print("\nconditional distribution for y\n", yctb)

        self.xcp = xcp
        self.xyp = ycp
```

Let's apply our code to some examples.

8.12.1 Numerical examples

Example 1

```
# joint
d = discrete_bijoint(f, xs, ys)
d.joint_tb()
```

```
The joint probability distribution for x and y
 +--------------------+-----+--------------------+--------------------+
|  x_value/y_value   |  10 |         20         | marginal sum for x |
+--------------------+-----+--------------------+--------------------+
|         0          | 0.3 |        0.2         |        0.5         |
|         1          | 0.1 |        0.4         |        0.5         |
| marginal_sum for y | 0.4 | 0.6000000000000001 |        1.0         |
+--------------------+-----+--------------------+--------------------+
```

```
# sample marginal
d.draw(1_000_000)
d.marg_dist()
```

```
marginal distribution for x
 +---------+----------+
| x_value |  x_prob  |
+---------+----------+
|    0    | 0.499395 |
|    1    | 0.500605 |
|   sum   |   1.0    |
```

(continues on next page)

(continued from previous page)

```
+---------+----------+

marginal distribution for y
 +---------+----------+
| y_value |  y_prob  |
+---------+----------+
|    10   | 0.400501 |
|    20   | 0.599499 |
|   sum   |   1.0    |
+---------+----------+
```

```
# sample conditional
d.cond_dist()
```

```
conditional distribution for x
 +---------+--------------------+---------------------+-----+
| x_value |         0          |          1          | sum |
+---------+--------------------+---------------------+-----+
|    10   | 0.7491766562380618 | 0.25082334376193816 | 1.0 |
|    20   | 0.3325259925371018 |  0.6674740074628982 | 1.0 |
+---------+--------------------+---------------------+-----+

conditional distribution for y
 +---------+---------------------+--------------------+-----+
| y_value |          10         |         20         | sum |
+---------+---------------------+--------------------+-----+
|    0    |  0.6008189909790846 | 0.3991810090209153 | 1.0 |
|    1    | 0.20066719269683683 | 0.7993328073031631 | 1.0 |
+---------+---------------------+--------------------+-----+
```

Example 2

```
xs_new = np.array([10, 20, 30])
ys_new = np.array([1, 2])
f_new = np.array([[0.2, 0.1], [0.1, 0.3], [0.15, 0.15]])
d_new = discrete_bijoint(f_new, xs_new, ys_new)
d_new.joint_tb()
```

```
The joint probability distribution for x and y
 +--------------------+---------------------+------+---------------------+
|  x_value/y_value   |          1          |  2   |  marginal sum for x |
+--------------------+---------------------+------+---------------------+
|         10         |         0.2         | 0.1  | 0.30000000000000004 |
|         20         |         0.1         | 0.3  |         0.4         |
|         30         |         0.15        | 0.15 |         0.3         |
| marginal_sum for y | 0.45000000000000007 | 0.55 |         1.0         |
+--------------------+---------------------+------+---------------------+
```

```
d_new.draw(1_000_000)
d_new.marg_dist()
```

```
marginal distribution for x
 +---------+----------+
```

(continues on next page)

(continued from previous page)

```
| x_value |  x_prob  |
+---------+----------+
|    10   | 0.299848 |
|    20   | 0.401015 |
|    30   | 0.299137 |
|   sum   |   1.0    |
+---------+----------+

marginal distribution for y
 +---------+----------+
| y_value |  y_prob  |
+---------+----------+
|    1    | 0.449195 |
|    2    | 0.550805 |
|   sum   |   1.0    |
+---------+----------+
```

```
d_new.cond_dist()
```

```
conditional distribution for x
 +---------+---------------------+---------------------+--------------------+-----+
| x_value |          10         |          20         |         30         | sum |
+---------+---------------------+---------------------+--------------------+-----+
|    1    |  0.4445886530348735 | 0.22256258417836353 | 0.332848762786763  | 1.0 |
|    2    | 0.18180844400468404 |  0.5465473261862184 | 0.2716442298090976 | 1.0 |
+---------+---------------------+---------------------+--------------------+-----+

conditional distribution for y
 +---------+---------------------+--------------------+-----+
| y_value |          1          |         2          | sum |
+---------+---------------------+--------------------+-----+
|    10   |  0.6660274539099811 | 0.3339725460900189 | 1.0 |
|    20   | 0.24930239517225042 | 0.7506976048277496 | 1.0 |
|    30   |  0.4998178092312218 | 0.5001821907687782 | 1.0 |
+---------+---------------------+--------------------+-----+
```

8.13 A continuous bivariate random vector

A two-dimensional Gaussian distribution has joint density

$$f(x,y) = (2\pi\sigma_1\sigma_2\sqrt{1-\rho^2})^{-1} \exp\left[-\frac{1}{2(1-\rho^2)}\left(\frac{(x-\mu_1)^2}{\sigma_1^2} - \frac{2\rho(x-\mu_1)(y-\mu_2)}{\sigma_1\sigma_2} + \frac{(y-\mu_2)^2}{\sigma_2^2}\right)\right]$$

We start with a bivariate normal distribution pinned down by

$$\mu = \begin{bmatrix} 0 \\ 5 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 5 & .2 \\ .2 & 1 \end{bmatrix}$$

```
# define the joint probability density function
def func(x, y, μ1=0, μ2=5, σ1=np.sqrt(5), σ2=np.sqrt(1), ρ=.2/np.sqrt(5*1)):
    A = (2 * np.pi * σ1 * σ2 * np.sqrt(1 - ρ**2))**(-1)
    B = -1 / 2 / (1 - ρ**2)
    C1 = (x - μ1)**2 / σ1**2
```

(continues on next page)

(continued from previous page)

```
    C2 = 2 * ρ * (x - μ1) * (y - μ2) / σ1 / σ2
    C3 = (y - μ2)**2 / σ2**2
    return A * np.exp(B * (C1 - C2 + C3))
```

```
μ1 = 0
μ2 = 5
σ1 = np.sqrt(5)
σ2 = np.sqrt(1)
ρ = .2 / np.sqrt(5 * 1)
```

```
x = np.linspace(-10, 10, 1_000)
y = np.linspace(-10, 10, 1_000)
x_mesh, y_mesh = np.meshgrid(x, y, indexing="ij")
```

8.13.1 Joint, marginal, and conditional distributions

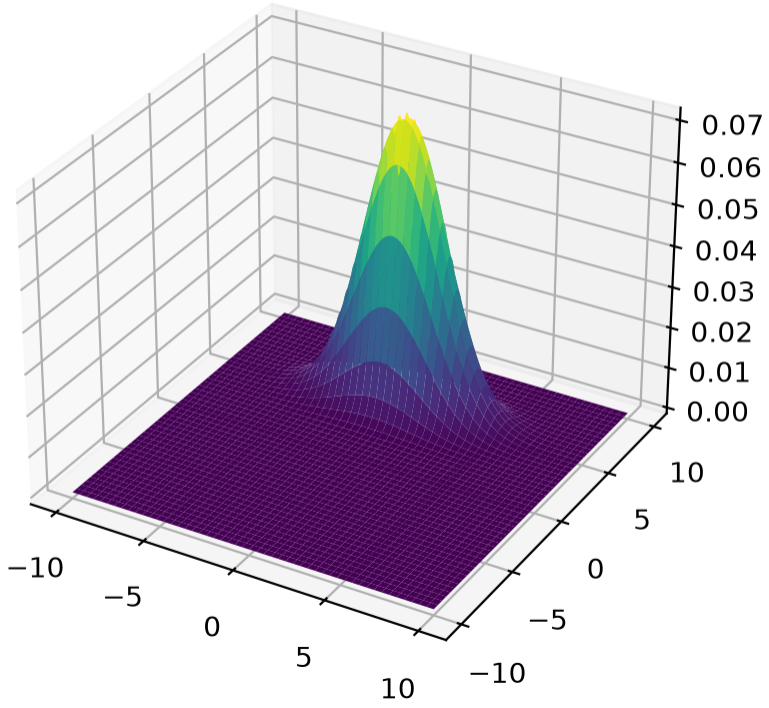
Joint distribution

Let's plot the **population** joint density.

```
# %matplotlib notebook

fig = plt.figure()
ax = plt.axes(projection='3d')

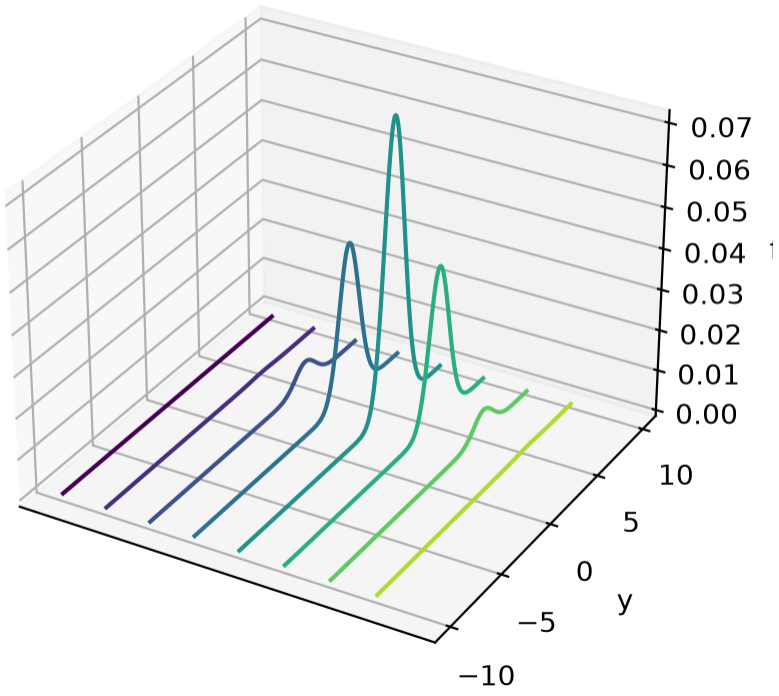
surf = ax.plot_surface(x_mesh, y_mesh, func(x_mesh, y_mesh), cmap='viridis')
plt.show()
```




```
# %matplotlib notebook

fig = plt.figure()
ax = plt.axes(projection='3d')

curve = ax.contour(x_mesh, y_mesh, func(x_mesh, y_mesh), zdir='x')
plt.ylabel('y')
ax.set_zlabel('f')
ax.set_xticks([])
plt.show()
```



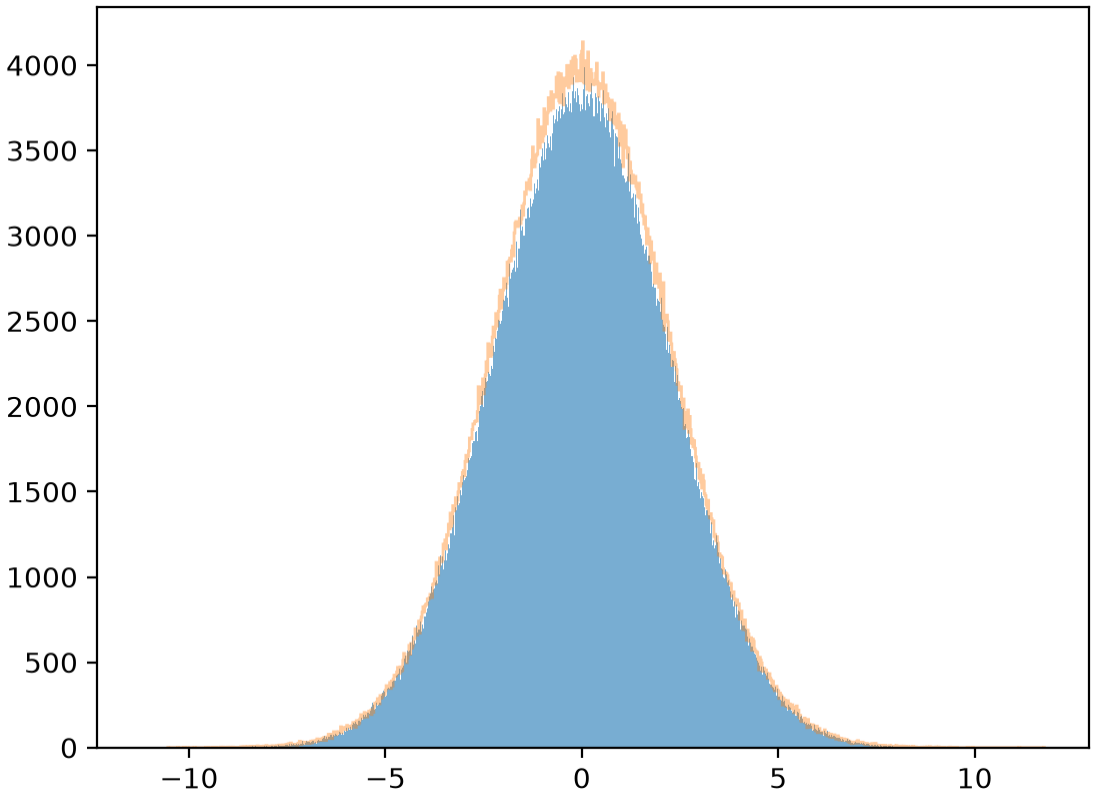
Next we can use a built-in `numpy` function to draw random samples, then calculate a **sample** marginal distribution from the sample mean and variance.

```
μ= np.array([0, 5])
σ= np.array([[5, .2], [.2, 1]])
n = 1_000_000
data = rng.multivariate_normal(μ, σ, n)
x = data[:, 0]
y = data[:, 1]
```

Marginal distribution

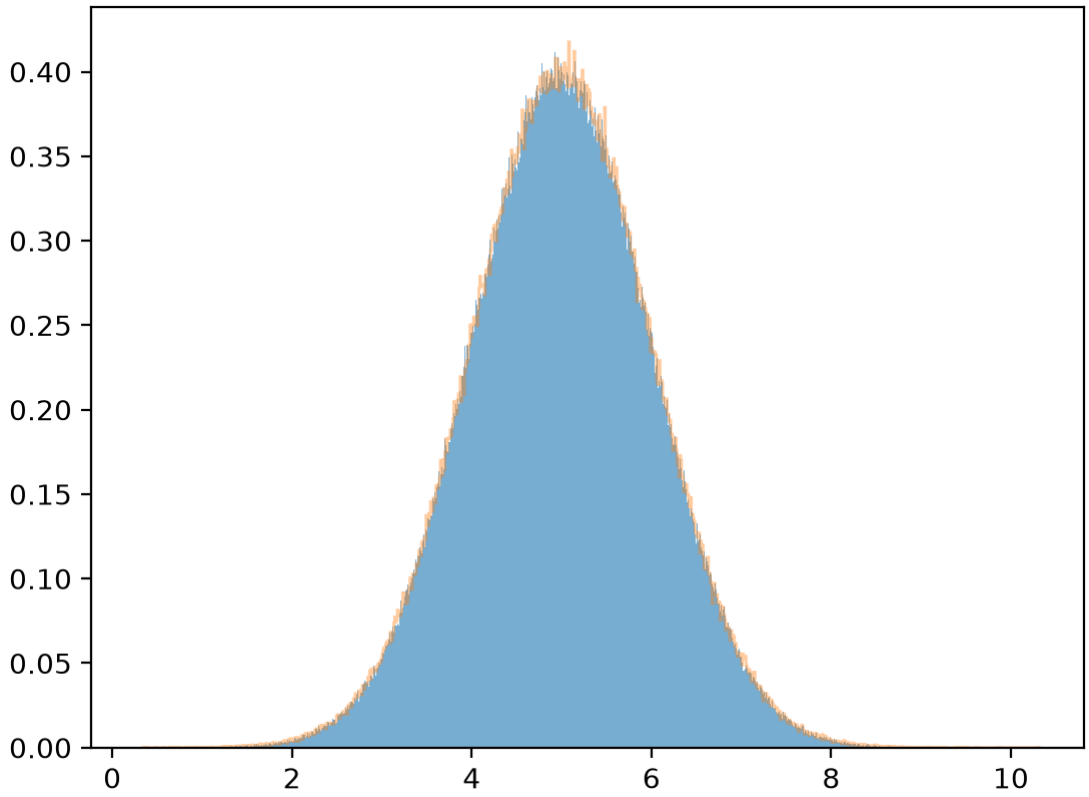
```
plt.hist(x, bins=1_000, alpha=0.6)
μx_hat, σx_hat = np.mean(x), np.std(x)
print(μx_hat, σx_hat)
x_sim = rng.normal(μx_hat, σx_hat, 1_000_000)
plt.hist(x_sim, bins=1_000, alpha=0.4, histtype="step")
plt.show()
```

```
-0.002835018476282638 2.237281153907461
```




```
plt.hist(y, bins=1_000, density=True, alpha=0.6)
μy_hat, σy_hat = np.mean(y), np.std(y)
print(μy_hat, σy_hat)
y_sim = rng.normal(μy_hat, σy_hat, 1_000_000)
plt.hist(y_sim, bins=1_000, density=True, alpha=0.4, histtype="step")
plt.show()
```

```
4.998242637632656 0.9993565018250069
```



Conditional distribution

For a bivariate normal population distribution, the conditional distributions are also normal:

$$X \mid Y = y \sim \mathbb{N}\left[\mu_X + \rho\sigma_X \frac{y - \mu_Y}{\sigma_Y}, \sigma_X^2(1-\rho^2)\right]$$
$$Y \mid X = x \sim \mathbb{N}\left[\mu_Y + \rho\sigma_Y \frac{x - \mu_X}{\sigma_X}, \sigma_Y^2(1-\rho^2)\right]$$

Note

Please see this *quantecon lecture* for more details.

Let's approximate the joint density by discretizing and mapping the approximating joint density into a matrix.

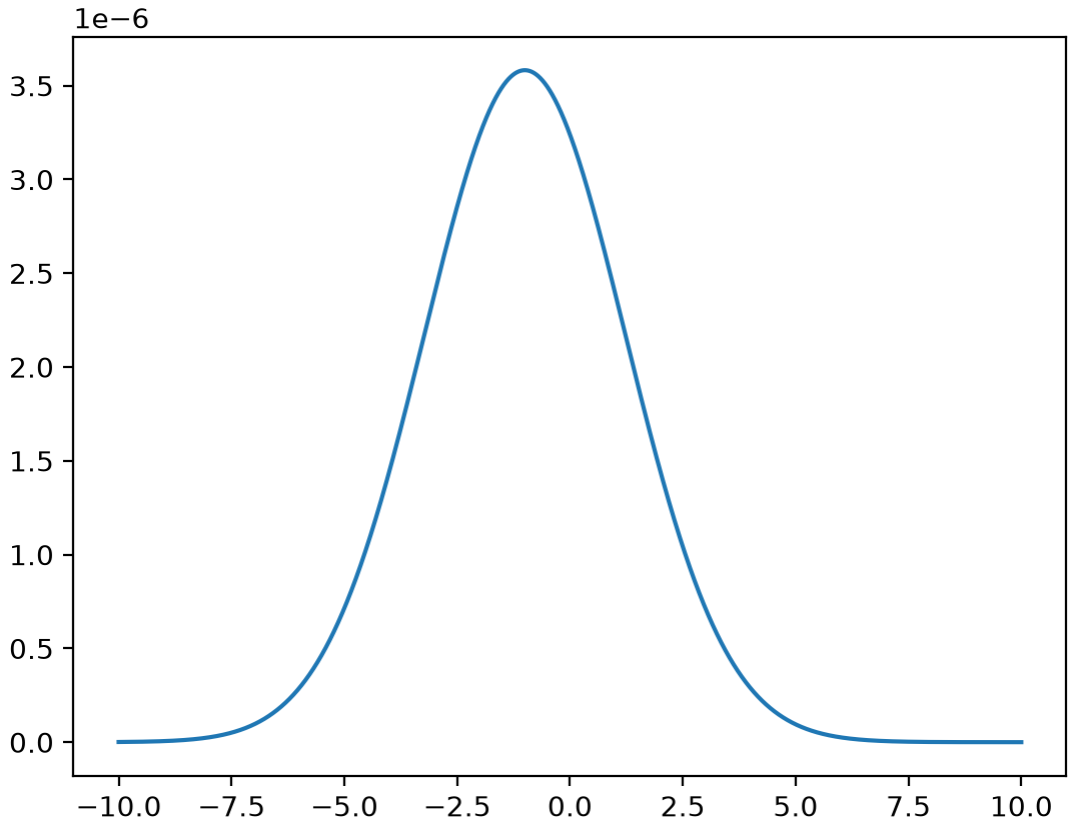
On an evenly spaced grid, we can approximate the conditional distribution by assigning probability weights proportional to a slice of the joint density.

For fixed y, this means that

$$z_i \equiv \frac{f(x_i, y)}{\sum_k f(x_k, y)}$$

Fix $y = 0$.

```
# discretized conditional distribution of X given Y = 0
x = np.linspace(-10, 10, 1_000_000)
z = func(x, y=0) / np.sum(func(x, y=0))
plt.plot(x, z)
plt.show()
```



The conditional mean and variance are then approximated by

$$\mathbb{E}\left[X|Y=y\right] \approx \sum_i x_i z_i$$

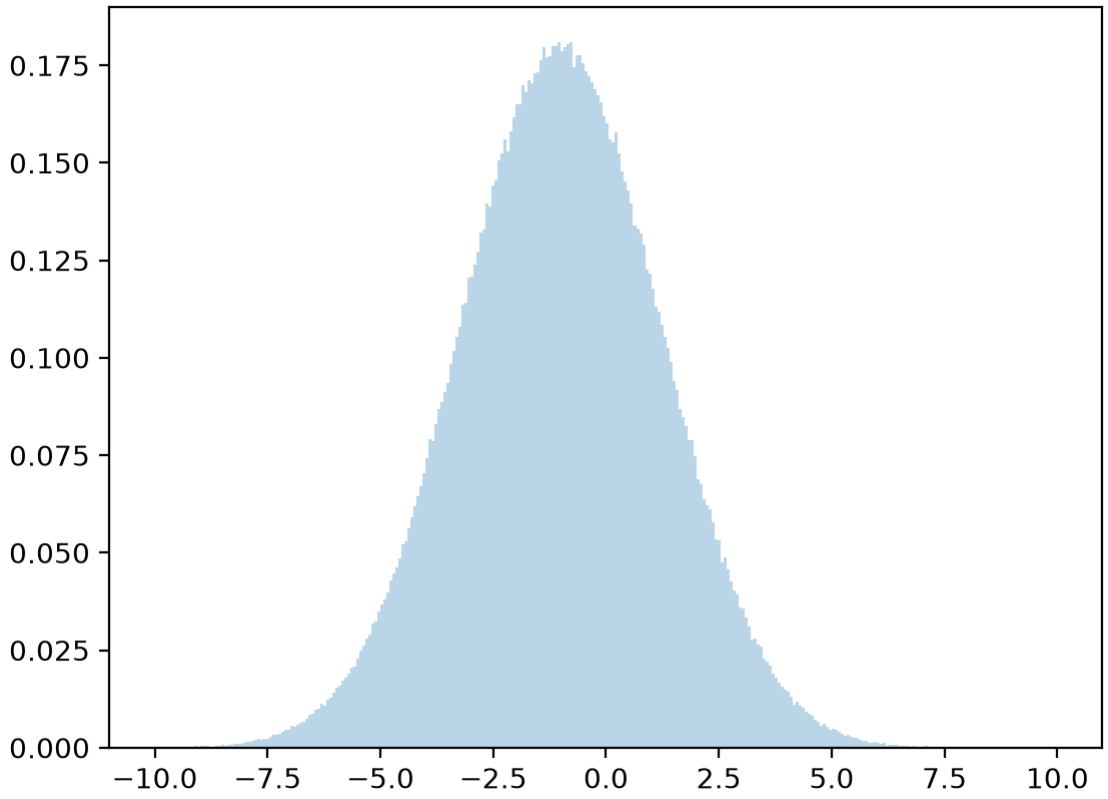
$$\mathbb{D}\left[X|Y=y\right] \approx \sum_i \left(x_i - \mu_{X|Y=y}\right)^2 z_i$$

Let's draw from a normal distribution with above mean and variance and check how accurate our approximation is.

```
# discretized mean
μx = np.dot(x, z)

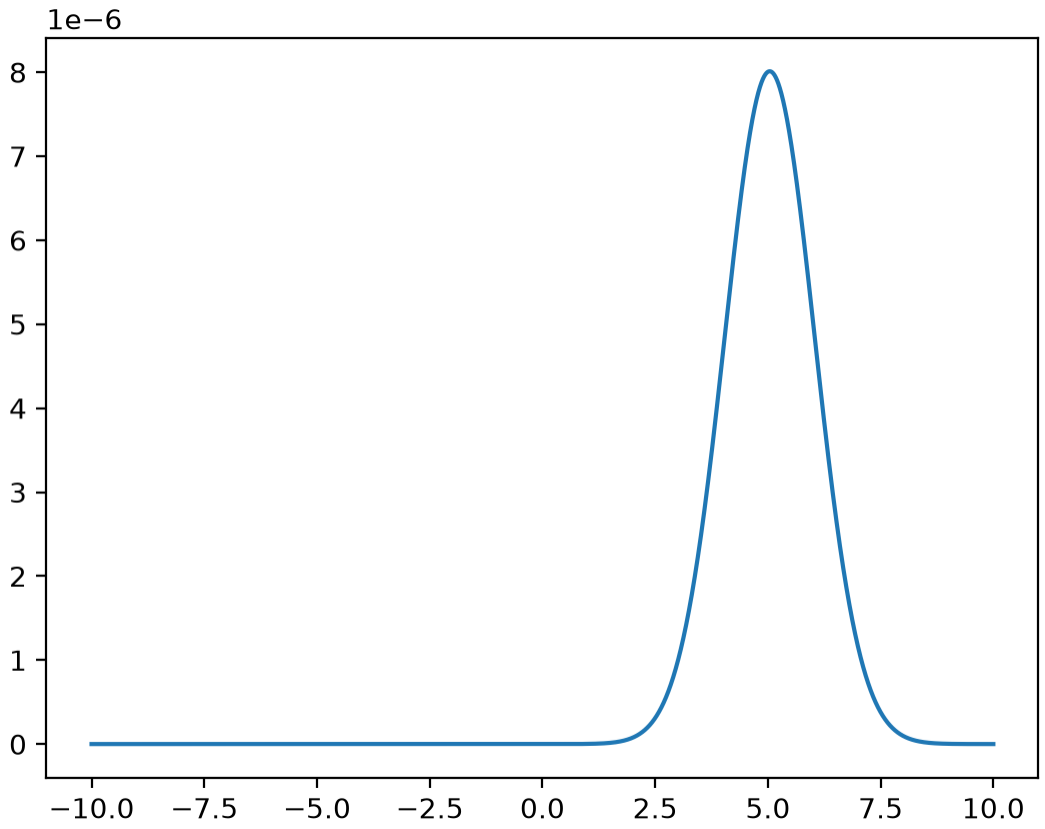
# discretized standard deviation
σx = np.sqrt(np.dot((x - μx)**2, z))

# sample
zz = rng.normal(μx, σx, 1_000_000)
plt.hist(zz, bins=300, density=True, alpha=0.3, range=[-10, 10])
plt.show()
```



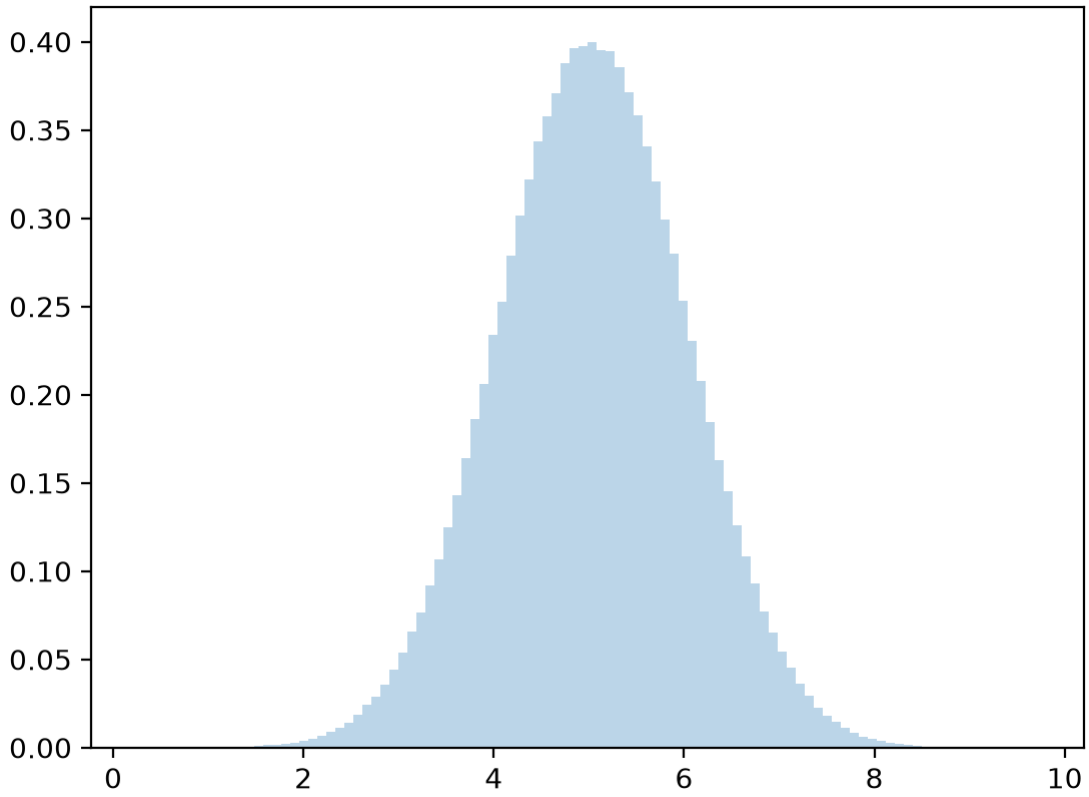
Fix $x = 1$.

```
y = np.linspace(-10, 10, 1_000_000)
z = func(x=1, y=y) / np.sum(func(x=1, y=y))
plt.plot(y,z)
plt.show()
```




```
# discretized conditional mean and standard deviation
μy = np.dot(y,z)
σy = np.sqrt(np.dot((y - μy)**2, z))

# sample
zz = rng.normal(μy, σy, 1_000_000)
plt.hist(zz, bins=100, density=True, alpha=0.3)
plt.show()
```



We compare with the analytically computed parameters and note that they are close.

```
print(μx, σx)
print(μ1 + ρ * σ1 * (0 - μ2) / σ2, np.sqrt(σ1**2 * (1 - ρ**2)))

print(μy, σy)
print(μ2 + ρ * σ2 * (1 - μ1) / σ1, np.sqrt(σ2**2 * (1 - ρ**2)))
```

```
-0.9997518414498444 2.2265841331697698
-1.0 2.227105745132009
5.039998363587647 0.9959878931873033
5.04 0.9959919678390986
```

8.14 Sum of two independently distributed random variables

Let X, Y be two independent discrete random variables that take values in $\bar{X}, \bar{Y}$, respectively.

Define a new random variable $Z = X + Y$.

Evidently, Z takes values from $\bar{Z}$ defined as follows:

$$\begin{aligned} \bar{X} &= \{0, 1, \ldots, I-1\}; \qquad & f_i &= \text{Prob}\{X = i\} \\ \bar{Y} &= \{0, 1, \ldots, J-1\}; \qquad & g_j &= \text{Prob}\{Y = j\} \\ \bar{Z} &= \{0, 1, \ldots, I+J-2\}; \qquad & h_k &= \text{Prob}\{X + Y = k\} \end{aligned}$$

Independence of X and Y implies that

$$h_k = \text{Prob}\{X = 0, Y = k\} + \text{Prob}\{X = 1, Y = k-1\} + ... + \text{Prob}\{X = k, Y = 0\}$$
$$h_k = f_0 g_k + f_1 g_{k-1} + ... + f_{k-1} g_1 + f_k g_0 \qquad \text{for} \quad k = 0, 1, ..., I + J - 2$$

Thus, we have:

$$h_k = \sum_{i=0}^{k} f_i g_{k-i} \equiv f * g$$

where $f * g$ denotes the **convolution** of the f and g sequences.

Similarly, for two random variables X, Y with densities f_X, g_Y, the density of $Z = X + Y$ is

$$f_Z(z) = \int_{-\infty}^{\infty} f_X(x) f_Y(z - x) dx \equiv f_X * g_Y$$

where $f_X * g_Y$ denotes the **convolution** of the f_X and g_Y functions.

8.15 Coupling

Start with a joint distribution

$$f_{ij} = \text{Prob}\{X = i, Y = j\}$$
$$i = 0, \cdots, I - 1$$
$$j = 0, \cdots, J - 1$$
$$\text{stacked to an } I \times J \text{ matrix}$$
$$e.g. \quad I = 2, J = 2$$

where

$$\begin{bmatrix} f_{00} & f_{01} \\ f_{10} & f_{11} \end{bmatrix}$$

From the joint distribution, we have shown above that we obtain **unique** marginal distributions.

Now we'll try to go in a reverse direction.

We'll find that from two marginal distributions we can usually construct more than one joint distribution that satisfies these marginals.

Each of these joint distributions is called a **coupling** of the two marginal distributions.

Let's start with marginal distributions

$$\text{Prob}\{X = i\} = \sum_j f_{ij} = \mu_i, i = 0, \cdots, I - 1$$
$$\text{Prob}\{Y = j\} = \sum_i f_{ij} = \nu_j, j = 0, \cdots, J - 1$$

Given two marginal distributions, μ for X and ν for Y, a joint distribution f_{ij} is said to be a **coupling** of μ and ν.

Consider the following bivariate example.

$$\text{Prob}\{X = 0\} = 1 - q = \mu_0$$
$$\text{Prob}\{X = 1\} = q = \mu_1$$
$$\text{Prob}\{Y = 0\} = 1 - r = \nu_0$$
$$\text{Prob}\{Y = 1\} = r = \nu_1$$
$$\text{where } 0 \leq q \leq r \leq 1$$

We construct two couplings.

The first coupling of our two marginal distributions is the joint distribution

$$f_{ij} = \begin{bmatrix} (1-q)(1-r) & (1-q)r \\ q(1-r) & qr \end{bmatrix}$$

To verify that it is a coupling, we check that

$$\begin{aligned} (1-q)(1-r) + (1-q)r + q(1-r) + qr &= 1 \\ \mu_0 = (1-q)(1-r) + (1-q)r &= 1-q \\ \mu_1 = q(1-r) + qr &= q \\ \nu_0 = (1-q)(1-r) + (1-r)q &= 1-r \\ \nu_1 = r(1-q) + qr &= r \end{aligned}$$

A second coupling of our two marginal distributions is the joint distribution

$$f_{ij} = \begin{bmatrix} (1-r) & r-q \\ 0 & q \end{bmatrix}$$

To verify that this is a coupling, note that

$$\begin{aligned} 1 - r + r - q + q &= 1 \\ \mu_0 &= 1-q \\ \mu_1 &= q \\ \nu_0 &= 1-r \\ \nu_1 &= r \end{aligned}$$

Thus, our two proposed joint distributions have the same marginal distributions.

But the joint distributions differ.

Thus, multiple joint distributions $[f_{ij}]$ can have the same marginals.

Couplings are important in optimal transport problems and in Markov processes. Please see this *lecture about optimal transport*.

8.16 Copula functions

Suppose that $X_1, X_2, \ldots, X_N$ are N random variables and that

- their marginal distributions are $F_1(x_1), F_2(x_2), \ldots, F_N(x_N)$, and
- their joint distribution is $H(x_1, x_2, \ldots, x_N)$

Then there exists a **copula function** $C(\cdot)$ that verifies

$$H(x_1, x_2, \ldots, x_N) = C(F_1(x_1), F_2(x_2), \ldots, F_N(x_N)).$$

If the marginal distributions are continuous, then the copula is unique.

In that case, we can recover it from the marginal inverses:

$$C(u_1, u_2, \ldots, u_N) = H(F_1^{-1}(u_1), F_2^{-1}(u_2), \ldots, F_N^{-1}(u_N))$$

When marginal distributions are not continuous, one uses generalized inverses, and the copula is uniquely determined only on $\text{Ran}(F_1) \times \cdots \times \text{Ran}(F_N)$.

In a reverse direction of logic, given univariate **marginal distributions** $F_1(x_1), F_2(x_2), \ldots, F_N(x_N)$ and a copula function $C(\cdot)$, the function $H(x_1, x_2, \ldots, x_N) = C(F_1(x_1), F_2(x_2), \ldots, F_N(x_N))$ is a **coupling** of $F_1(x_1), F_2(x_2), \ldots, F_N(x_N)$.

Thus, for given marginal distributions, we can use a copula function to determine a joint distribution when the associated univariate random variables are not independent.

Copula functions are often used to characterize **dependence** of random variables.

8.16.1 Bivariate examples with discrete and continuous distributions

Discrete marginal distribution

As mentioned above, for two given marginal distributions there can be more than one coupling.

For example, consider two random variables X, Y with distributions

$$\begin{aligned}\text{Prob}(X = 0) &= 0.6,\\ \text{Prob}(X = 1) &= 0.4,\\ \text{Prob}(Y = 0) &= 0.3,\\ \text{Prob}(Y = 1) &= 0.7,\end{aligned}$$

For these two random variables there can be more than one coupling.

Let's first generate X and Y.

```
μ = np.array([0.6, 0.4])
ν = np.array([0.3, 0.7])

# number of draws
draws = 1_000_000

# generate independent draws from uniform distribution for X and Y
p_x = rng.random(draws)
p_y = rng.random(draws)

# generate draws of X and Y via independent uniform draws
x = np.ones(draws)
y = np.ones(draws)
x[p_x <= μ[0]] = 0
x[p_x > μ[0]] = 1
y[p_y <= ν[0]] = 0
y[p_y > ν[0]] = 1
```

```
# calculate parameters from draws
q_hat = sum(x[x == 1])/draws
r_hat = sum(y[y == 1])/draws

# print output
print("distribution for x")
xmtb = pt.PrettyTable()
xmtb.field_names = ['x_value', 'x_prob']
xmtb.add_row([0, 1-q_hat])
xmtb.add_row([1, q_hat])
print(xmtb)
```

(continues on next page)

(continued from previous page)

```
print("distribution for y")
ymtb = pt.PrettyTable()
ymtb.field_names = ['y_value', 'y_prob']
ymtb.add_row([0, 1-r_hat])
ymtb.add_row([1, r_hat])
print(ymtb)
```

```
distribution for x
+---------+----------+
| x_value |  x_prob  |
+---------+----------+
|    0    | 0.599645 |
|    1    | 0.400355 |
+---------+----------+
distribution for y
+---------+---------------------+
| y_value |        y_prob       |
+---------+---------------------+
|    0    | 0.30042599999999997 |
|    1    |       0.699574      |
+---------+---------------------+
```

Let's now take our two marginal distributions, one for X, the other for Y, and construct two distinct couplings.

For the first joint distribution:

$$\text{Prob}(X = i, Y = j) = f_{ij}$$

where

$$[f_{ij}] = \begin{bmatrix} 0.18 & 0.42 \\ 0.12 & 0.28 \end{bmatrix}$$

Let's use Python to construct this joint distribution and then verify that its marginal distributions are what we want.

```
# define parameters
f1 = np.array([[0.18, 0.42], [0.12, 0.28]])
f1_cum = np.cumsum(f1)

# number of draws
draws1 = 1_000_000

# generate draws from uniform distribution
p = rng.random(draws1)

# generate draws of first coupling via uniform distribution
c1 = np.vstack([np.ones(draws1), np.ones(draws1)])
# X=0, Y=0
c1[0, p <= f1_cum[0]] = 0
c1[1, p <= f1_cum[0]] = 0
# X=0, Y=1
c1[0, (p > f1_cum[0])*(p <= f1_cum[1])] = 0
c1[1, (p > f1_cum[0])*(p <= f1_cum[1])] = 1
# X=1, Y=0
c1[0, (p > f1_cum[1])*(p <= f1_cum[2])] = 1
c1[1, (p > f1_cum[1])*(p <= f1_cum[2])] = 0
# X=1, Y=1
```

(continues on next page)

(continued from previous page)

```
c1[0, (p > f1_cum[2])*(p <= f1_cum[3])] = 1
c1[1, (p > f1_cum[2])*(p <= f1_cum[3])] = 1
```

```
# calculate parameters from draws
f1_00 = sum((c1[0, :] == 0)*(c1[1, :] == 0))/draws1
f1_01 = sum((c1[0, :] == 0)*(c1[1, :] == 1))/draws1
f1_10 = sum((c1[0, :] == 1)*(c1[1, :] == 0))/draws1
f1_11 = sum((c1[0, :] == 1)*(c1[1, :] == 1))/draws1

# print output of first joint distribution
print("first joint distribution for c1")
c1_mtb = pt.PrettyTable()
c1_mtb.field_names = ['c1_x_value', 'c1_y_value', 'c1_prob']
c1_mtb.add_row([0, 0, f1_00])
c1_mtb.add_row([0, 1, f1_01])
c1_mtb.add_row([1, 0, f1_10])
c1_mtb.add_row([1, 1, f1_11])
print(c1_mtb)
```

```
first joint distribution for c1
+------------+------------+----------+
| c1_x_value | c1_y_value | c1_prob  |
+------------+------------+----------+
|     0      |     0      | 0.18002  |
|     0      |     1      | 0.419939 |
|     1      |     0      | 0.120069 |
|     1      |     1      | 0.279972 |
+------------+------------+----------+
```

```
# calculate parameters from draws
c1_q_hat = sum(c1[0, :] == 1)/draws1
c1_r_hat = sum(c1[1, :] == 1)/draws1

# print output
print("marginal distribution for x")
c1_x_mtb = pt.PrettyTable()
c1_x_mtb.field_names = ['c1_x_value', 'c1_x_prob']
c1_x_mtb.add_row([0, 1-c1_q_hat])
c1_x_mtb.add_row([1, c1_q_hat])
print(c1_x_mtb)

print("marginal distribution for y")
c1_ymtb = pt.PrettyTable()
c1_ymtb.field_names = ['c1_y_value', 'c1_y_prob']
c1_ymtb.add_row([0, 1-c1_r_hat])
c1_ymtb.add_row([1, c1_r_hat])
print(c1_ymtb)
```

```
marginal distribution for x
+------------+-----------+
| c1_x_value | c1_x_prob |
+------------+-----------+
|     0      |  0.599959 |
|     1      |  0.400041 |
+------------+-----------+
```

(continues on next page)

(continued from previous page)

```
marginal distribution for y
+------------+---------------------+
| c1_y_value |      c1_y_prob      |
+------------+---------------------+
|     0      | 0.30008900000000005 |
|     1      |       0.699911      |
+------------+---------------------+
```

Now, let's construct another joint distribution that is also a coupling of X and Y

$$[f_{ij}] = \begin{bmatrix} 0.3 & 0.3 \\ 0 & 0.4 \end{bmatrix}$$

```
# define parameters
f2 = np.array([[0.3, 0.3], [0, 0.4]])
f2_cum = np.cumsum(f2)

# number of draws
draws2 = 1_000_000

# generate draws from uniform distribution
p = rng.random(draws2)

# generate draws of second coupling via uniform distribution
c2 = np.vstack([np.ones(draws2), np.ones(draws2)])
# X=0, Y=0
c2[0, p <= f2_cum[0]] = 0
c2[1, p <= f2_cum[0]] = 0
# X=0, Y=1
c2[0, (p > f2_cum[0])*(p <= f2_cum[1])] = 0
c2[1, (p > f2_cum[0])*(p <= f2_cum[1])] = 1
# X=1, Y=0
c2[0, (p > f2_cum[1])*(p <= f2_cum[2])] = 1
c2[1, (p > f2_cum[1])*(p <= f2_cum[2])] = 0
# X=1, Y=1
c2[0, (p > f2_cum[2])*(p <= f2_cum[3])] = 1
c2[1, (p > f2_cum[2])*(p <= f2_cum[3])] = 1
```

```
# calculate parameters from draws
f2_00 = sum((c2[0, :] == 0)*(c2[1, :] == 0))/draws2
f2_01 = sum((c2[0, :] == 0)*(c2[1, :] == 1))/draws2
f2_10 = sum((c2[0, :] == 1)*(c2[1, :] == 0))/draws2
f2_11 = sum((c2[0, :] == 1)*(c2[1, :] == 1))/draws2

# print output of second joint distribution
print("second joint distribution for c2")
c2_mtb = pt.PrettyTable()
c2_mtb.field_names = ['c2_x_value', 'c2_y_value', 'c2_prob']
c2_mtb.add_row([0, 0, f2_00])
c2_mtb.add_row([0, 1, f2_01])
c2_mtb.add_row([1, 0, f2_10])
c2_mtb.add_row([1, 1, f2_11])
print(c2_mtb)
```

```
second joint distribution for c2
+------------+------------+----------+
```

(continues on next page)

(continued from previous page)

```
| c2_x_value | c2_y_value | c2_prob  |
+------------+------------+----------+
|     0      |     0      | 0.300568 |
|     0      |     1      | 0.299249 |
|     1      |     0      |   0.0    |
|     1      |     1      | 0.400183 |
+------------+------------+----------+
```

```
# calculate parameters from draws
c2_q_hat = sum(c2[0, :] == 1)/draws2
c2_r_hat = sum(c2[1, :] == 1)/draws2

# print output
print("marginal distribution for x")
c2_x_mtb = pt.PrettyTable()
c2_x_mtb.field_names = ['c2_x_value', 'c2_x_prob']
c2_x_mtb.add_row([0, 1-c2_q_hat])
c2_x_mtb.add_row([1, c2_q_hat])
print(c2_x_mtb)

print("marginal distribution for y")
c2_ymtb = pt.PrettyTable()
c2_ymtb.field_names = ['c2_y_value', 'c2_y_prob']
c2_ymtb.add_row([0, 1-c2_r_hat])
c2_ymtb.add_row([1, c2_r_hat])
print(c2_ymtb)
```

```
marginal distribution for x
+------------+-----------+
| c2_x_value | c2_x_prob |
+------------+-----------+
|     0      |  0.599817 |
|     1      |  0.400183 |
+------------+-----------+
marginal distribution for y
+------------+--------------------+
| c2_y_value |      c2_y_prob     |
+------------+--------------------+
|     0      | 0.30056799999999995 |
|     1      |       0.699432     |
+------------+--------------------+
```

We have verified that both joint distributions, c_1 and c_2, have identical marginal distributions of X and Y, respectively.

So they are both couplings of X and Y.

8.16.2 Gaussian copula example

A **Gaussian copula** uses the bivariate normal distribution to induce dependence between arbitrary marginal distributions.

The construction has three steps:

1. Draw (Z_1, Z_2) from a bivariate standard normal with correlation ρ.
2. Apply the standard normal CDF: $U_k = \Phi(Z_k)$.
    - The pair (U_1, U_2) has uniform marginals but retains the dependence structure of (Z_1, Z_2) — this is the copula.
3. Apply the inverse CDF of any desired marginal: $X_k = F_k^{-1}(U_k)$.

The following code illustrates this with exponential marginals.

```
# Gaussian copula parameters
ρ_cop = 0.8
n_cop = 100_000

# Draw from bivariate standard normal with correlation ρ_cop
z = rng.multivariate_normal(
    [0, 0], [[1, ρ_cop], [ρ_cop, 1]], n_cop
)

# Apply normal CDF -> uniform marginals (the copula itself)
u1 = stats.norm.cdf(z[:, 0])
u2 = stats.norm.cdf(z[:, 1])

# Apply inverse CDFs of desired marginals (here: Exponential)
x1 = stats.expon.ppf(u1, scale=1.0)    # Exp with mean 1
x2 = stats.expon.ppf(u2, scale=0.5)    # Exp with mean 0.5

fig, axes = plt.subplots(1, 2, figsize=(10, 4))
axes[0].scatter(u1[:3000], u2[:3000], alpha=0.2, s=2)
axes[0].set_xlabel('$u_1$')
axes[0].set_ylabel('$u_2$')
axes[1].scatter(x1[:3000], x2[:3000], alpha=0.2, s=2)
axes[1].set_xlabel('$x_1$ (Exp, mean=1)')
axes[1].set_ylabel('$x_2$ (Exp, mean=0.5)')
plt.show()

print(f"Sample correlation of (x1, x2): {np.corrcoef(x1, x2)[0, 1]:.3f}")
print(f"Sample correlation of (u1, u2): {np.corrcoef(u1, u2)[0, 1]:.3f}")
```

```
Sample correlation of (x1, x2): 0.768
Sample correlation of (u1, u2): 0.786
```

The left panel shows the copula itself – the dependence structure in uniform coordinates, drawn from a bivariate normal with correlation $\rho = 0.8$.

The right panel shows the same dependence translated to exponential marginals.

Changing ρ controls the strength of dependence while the marginals remain unchanged.

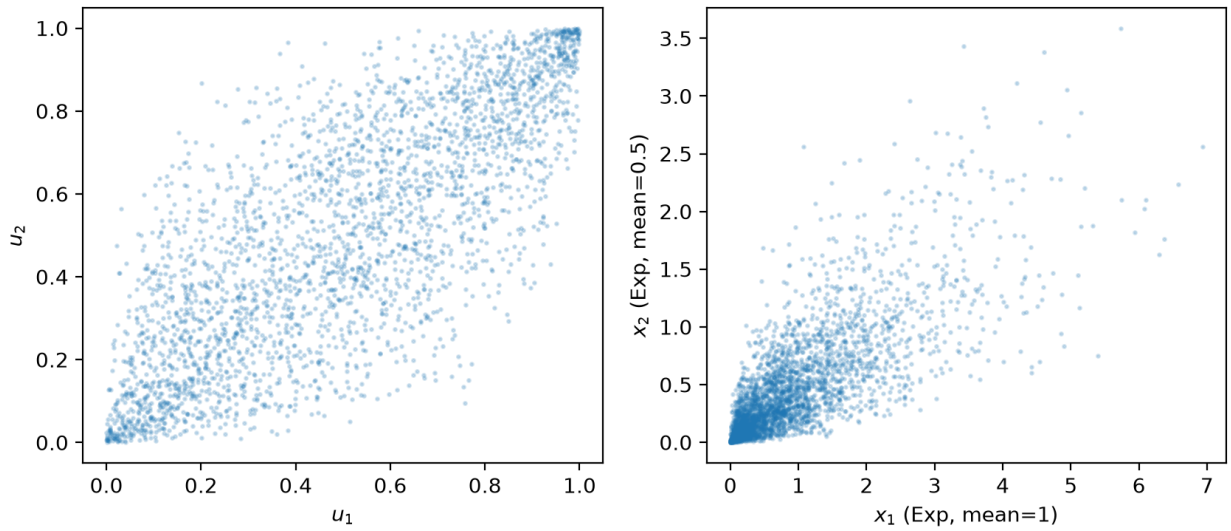


Fig. 8.1: gaussian copula with exponential marginals

8.17 Exercises

Exercise 8.17.1

Independence Test

Consider the joint distribution

$$F = \begin{bmatrix} 0.3 & 0.2 \\ 0.1 & 0.4 \end{bmatrix}$$

where $X \in \{0, 1\}$ and $Y \in \{10, 20\}$.

1. Compute the marginal distributions $\mu_i = \text{Prob}\{X = i\}$ and $\nu_j = \text{Prob}\{Y = j\}$.
2. Form the independence matrix $f_{ij}^{\perp} = \mu_i \nu_j$ (the outer product of the two marginal vectors).
3. Compare F with $f^{\perp}$ and determine whether X and Y are independent.
4. Verify your conclusion by computing $\text{Prob}\{X = 0 | Y = 10\}$ and checking whether it equals $\text{Prob}\{X = 0\}$.

Solution

Here is one solution:

```
F = np.array([[0.3, 0.2],
              [0.1, 0.4]])

μ = F.sum(axis=1)
ν = F.sum(axis=0)
print("μ (marginal of X):", μ)
print("ν (marginal of Y):", ν)

F_indep = np.outer(μ, ν)
```

```
print("\nIndependence matrix (outer product):\n", F_indep)
print("\nActual joint F:\n", F)

print("\nIndependent (F == μ times ν)?", np.allclose(F, F_indep))

prob_X0_given_Y10 = F[0, 0] / ν[0]
print(f"\nProb(X=0 | Y=10) = {prob_X0_given_Y10:.4f}")
print(f"Prob(X=0)          = {μ[0]:.4f}")
```

```
μ (marginal of X): [0.5 0.5]
ν (marginal of Y): [0.4 0.6]

Independence matrix (outer product):
 [[0.2 0.3]
 [0.2 0.3]]

Actual joint F:
 [[0.3 0.2]
 [0.1 0.4]]

Independent (F == μ times ν)? False

Prob(X=0 | Y=10) = 0.7500
Prob(X=0)          = 0.5000
```

Exercise 8.17.2

Covariance and Correlation

Using the same joint distribution F and values $X \in \{0, 1\}$, $Y \in \{10, 20\}$ as in Exercise 1:

1. Compute $\mathbb{E}[X]$, $\mathbb{E}[Y]$, and $\mathbb{E}[XY] = \sum_i \sum_j x_i y_j f_{ij}$.
2. Compute $\text{Cov}(X, Y) = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y]$.
3. Compute $\text{Cor}(X, Y) = \text{Cov}(X, Y)/(\sigma_X \sigma_Y)$.
4. Show analytically that $X \perp Y$ implies $\text{Cov}(X, Y) = 0$.

Solution

Here is one solution:

```
xs = np.array([0, 1])
ys = np.array([10, 20])
F  = np.array([[0.3, 0.2],
               [0.1, 0.4]])

μ = F.sum(axis=1)
ν = F.sum(axis=0)

E_X  = xs @ μ
E_Y  = ys @ ν
E_XY = sum(xs[i] * ys[j] * F[i, j] for i in range(2) for j in range(2))
print(f"E[X] = {E_X}, E[Y] = {E_Y}, E[XY] = {E_XY}")

cov_XY = E_XY - E_X * E_Y
```

```
print(f"Cov(X,Y) = {cov_XY:.4f}")

var_X  = ((xs - E_X)**2) @ μ
var_Y  = ((ys - E_Y)**2) @ ν
cor_XY = cov_XY / np.sqrt(var_X * var_Y)
print(f"Cor(X,Y) = {cor_XY:.4f}")
```

```
E[X] = 0.5, E[Y] = 16.0, E[XY] = 9.0
Cov(X,Y) = 1.0000
Cor(X,Y) = 0.4082
```

For part 4: if $X \perp Y$ then $f_{ij} = \mu_i \nu_j$, so

$$\mathbb{E}[XY] = \sum_i \sum_j x_i y_j \mu_i \nu_j = \left(\sum_i x_i \mu_i\right)\left(\sum_j y_j \nu_j\right) = \mathbb{E}[X]\,\mathbb{E}[Y]$$

and therefore $\text{Cov}(X,Y) = \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y] = 0$.

Exercise 8.17.3

Sum of Two Dice

Let X and Y be **independent** random variables, each uniformly distributed on $\{1,2,3,4,5,6\}$, and let $Z = X + Y$.

1. Use the convolution formula $h_k = \sum_i f_i g_{k-i}$ to compute the distribution of Z.
2. Plot the result generated by the formula.
3. Simulate 10^6 rolls and overlay the empirical histogram on the plot.
4. Compute $\mathbb{E}[Z]$ and $\text{Var}(Z)$ from the two calculations

Solution

Here is one solution:

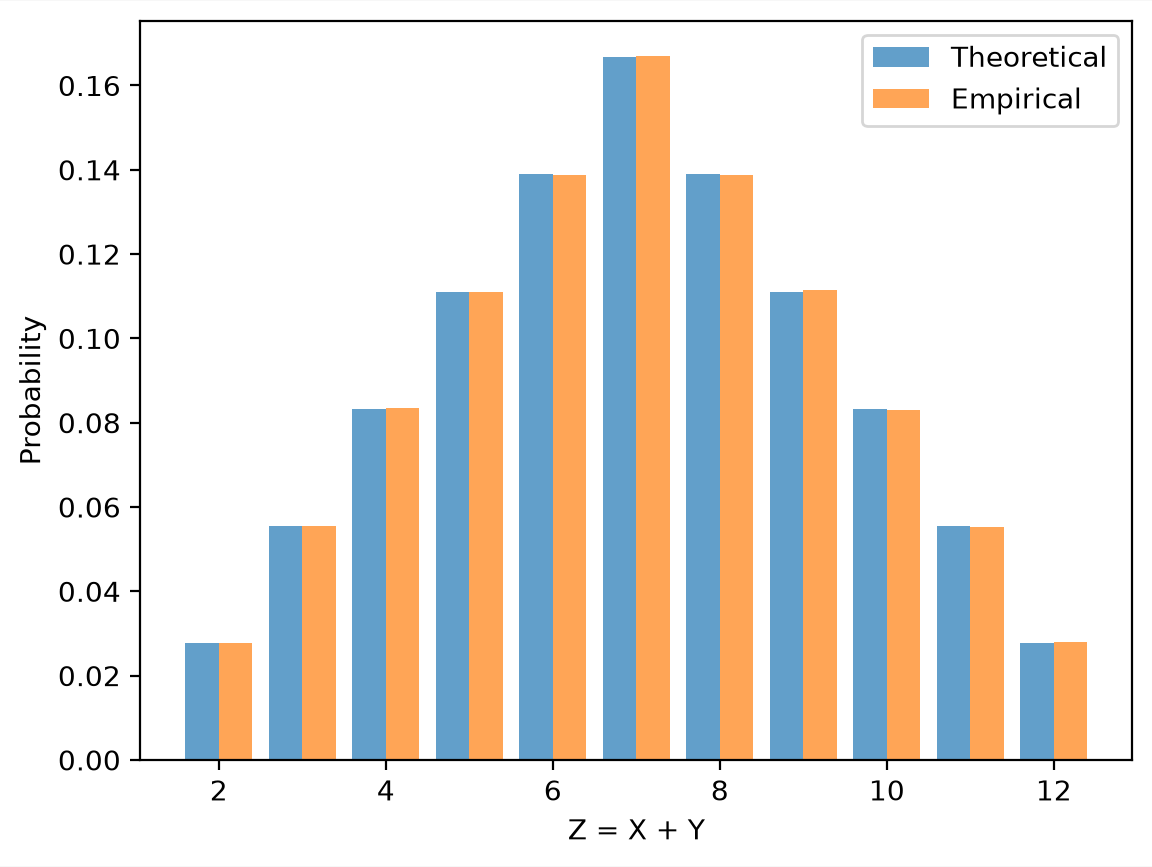
```
f = np.ones(6) / 6
g = np.ones(6) / 6
h = [
    sum(f[i]*g[k-i] for i in range(
        max(0, k-len(g)+1), # f_i exists
        min(len(f), k+1))   # g_{k-i} exists
        )
        for k in range(len(f) + len(g) - 1)]
z_vals = np.arange(2, 13)

n = 1_000_000
z_sim = rng.integers(1, 7, n) + rng.integers(1, 7, n)
counts = np.bincount(z_sim, minlength=13)[2:]

fig, ax = plt.subplots()
ax.bar(z_vals - 0.2, h,          0.4, alpha=0.7, label='Theoretical')
ax.bar(z_vals + 0.2, counts / n, 0.4, alpha=0.7, label='Empirical')
ax.set_xlabel('Z = X + Y')
ax.set_ylabel('Probability')
ax.legend()
```

```
plt.show()

E_Z   = z_vals @ h
Var_Z = ((z_vals - E_Z)**2) @ h
print(f"Theory:     E[Z] = {E_Z:.2f}, Var(Z) = {Var_Z:.4f}")
print(f"Simulation: E[Z] = {np.mean(z_sim):.2f}, Var(Z) = {np.var(z_sim):.4f}")
```




```
Theory:     E[Z] = 7.00, Var(Z) = 5.8333
Simulation: E[Z] = 7.00, Var(Z) = 5.8271
```

Exercise 8.17.4

Multi-Step Transition Probabilities

Consider a two-state Markov chain with transition matrix

$$P = \begin{bmatrix} 0.9 & 0.1 \\ 0.2 & 0.8 \end{bmatrix}$$

where $p_{ij} = \text{Prob}\{X(t+1) = j \mid X(t) = i\}$.

1. Starting from $\psi_0 = [1, 0]$, compute $\psi_n = \psi_0 P^n$ for $n = 1, 5, 20, 100$.
2. Find the stationary distribution ψ^* satisfying $\psi^* P = \psi^*$ and $\sum_i \psi_i^* = 1$.
3. Verify numerically that $\psi_n \to \psi^*$ as n grows.

Solution

Here is one solution:

```
P = np.array([[0.9, 0.1],
              [0.2, 0.8]])
ψ0 = np.array([1.0, 0.0])

for n in [1, 5, 20, 100]:
    print(f"ψ_{n:3d} = {ψ0 @ np.linalg.matrix_power(P, n)}")

A = np.vstack([P.T - np.eye(2), np.ones(2)])
b = np.array([0.0, 0.0, 1.0])
ψ_star, *_ = np.linalg.lstsq(A, b, rcond=None)
print(f"\nStationary distribution: {ψ_star}")

ψ_100 = ψ0 @ np.linalg.matrix_power(P, 100)
print(f"ψ_100 close to stationary? {np.allclose(ψ_100, ψ_star, atol=1e-6)}")
```

```
ψ_  1 = [0.9 0.1]
ψ_  5 = [0.72269 0.27731]
ψ_ 20 = [0.66693264 0.33306736]
ψ_100 = [0.66666667 0.33333333]

Stationary distribution: [0.66666667 0.33333333]
ψ_100 close to stationary? True
```

Exercise 8.17.5

Bayes' Law with a Discrete Prior

A coin has unknown bias $\theta \in \{0.2,\ 0.5,\ 0.8\}$ with prior $\pi = [0.25,\ 0.50,\ 0.25]$.

Assume that, conditional on θ, the coin flips are i.i.d. Bernoulli(θ).

1. After observing $k = 7$ heads in $n = 10$ flips, compute the likelihood

$$\mathcal{L}(\theta \mid \text{data}) = \binom{10}{7} \theta^7 (1-\theta)^3$$

   for each θ.

2. Apply equation (8.4) to compute the posterior $\pi(\theta \mid \text{data})$.
3. Plot the prior and posterior side by side.
4. Repeat for $k = 3$ heads and describe how the posterior shifts.

Solution

Here is one solution:

```
θ_vals = np.array([0.2, 0.5, 0.8])
π = np.array([0.25, 0.50, 0.25])

def compute_posterior(k, n, θ_vals, π):
    likelihood = comb(n, k) * θ_vals**k * (1 - θ_vals)**(n - k)
```

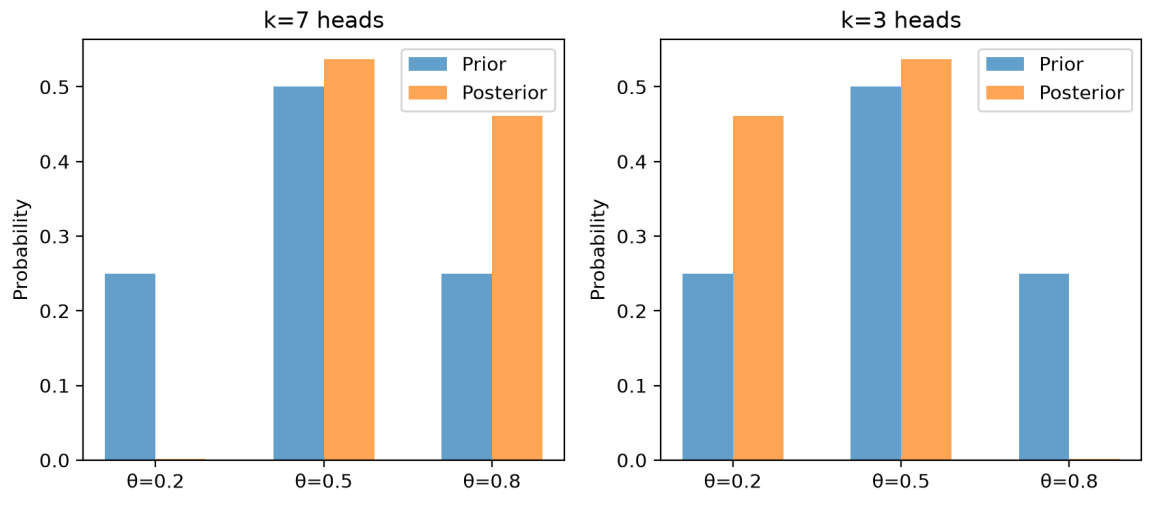
```
    unnorm = likelihood * π
    return unnorm / unnorm.sum(), likelihood

post7, lik7 = compute_posterior(7, 10, θ_vals, π)
post3, lik3 = compute_posterior(3, 10, θ_vals, π)

print("k=7:  likelihood =", lik7.round(4),
      " posterior =", post7.round(4))
print("k=3:  likelihood =", lik3.round(4),
      " posterior =", post3.round(4))

x = np.arange(len(θ_vals))
w = 0.3
fig, axes = plt.subplots(1, 2, figsize=(10, 4))
for ax, post, title in zip(
    axes, [post7, post3], ['k=7 heads', 'k=3 heads']):
    ax.bar(x - w/2, π, w, label='Prior',     alpha=0.7)
    ax.bar(x + w/2, post,  w, label='Posterior', alpha=0.7)
    ax.set_xticks(x)
    ax.set_xticklabels([f'θ={t}' for t in θ_vals])
    ax.set_ylabel('Probability')
    ax.set_title(title)
    ax.legend()
plt.show()
```

```
k=7:  likelihood = [0.0008 0.1172 0.2013]  posterior = [0.0018 0.537  0.4612]
k=3:  likelihood = [0.2013 0.1172 0.0008]  posterior = [0.4612 0.537  0.0018]
```



CHAPTER

NINE

SOME PROBABILITY DISTRIBUTIONS

This lecture is a supplement to *this lecture on statistics with matrices*.

It describes some popular distributions and uses Python to sample from them.

It also describes a way to sample from an arbitrary probability distribution that you make up by transforming a sample from a uniform probability distribution.

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install prettytable
```

As usual, we'll start with some imports

```
import numpy as np
import matplotlib.pyplot as plt
import prettytable as pt
from mpl_toolkits.mplot3d import Axes3D
from matplotlib_inline.backend_inline import set_matplotlib_formats
set_matplotlib_formats('retina')
```

9.1 Some Discrete Probability Distributions

Let's write some Python code to compute means and variances of some univariate random variables.

We'll use our code to

- compute population means and variances from the probability distribution
- generate a sample of N independently and identically distributed draws and compute sample means and variances
- compare population and sample means and variances

9.2 Geometric distribution

A discrete geometric distribution has probability mass function

$$\text{Prob}(X = k) = (1-p)^{k-1}p, k = 1, 2, \ldots, \quad p \in (0, 1)$$

where $k = 1, 2, \ldots$ is the number of trials before the first success.

The mean and variance of this one-parameter probability distribution are

$$\mathbb{E}(X) = \frac{1}{p}$$
$$\mathbb{V}\text{ar}(X) = \frac{1-p}{p^2}$$

Let's use Python draw observations from the distribution and compare the sample mean and variance with the theoretical results.

```
rng = np.random.default_rng()

# specify parameters
p, n = 0.3, 1_000_000

# draw observations from the distribution
x = rng.geometric(p, n)

# compute sample mean and variance
μ_hat = np.mean(x)
σ2_hat = np.var(x)

print("The sample mean is: ", μ_hat, "\nThe sample variance is: ", σ2_hat)

# compare with theoretical results
print("\nThe population mean is: ", 1/p)
print("The population variance is: ", (1-p)/(p**2))
```

```
The sample mean is:  3.336426
The sample variance is:  7.805587546524

The population mean is:  3.3333333333333335
The population variance is:  7.777777777777778
```

9.3 Pascal (negative binomial) distribution

Consider a sequence of independent Bernoulli trials.

Let p be the probability of success.

Let X be a random variable that represents the number of failures before we get r successes.

Its distribution is

$$X \sim NB(r, p)$$
$$\text{Prob}(X = k; r, p) = \begin{bmatrix} k+r-1 \\ r-1 \end{bmatrix} p^r (1-p)^k$$

Here, we choose from among $k + r - 1$ possible outcomes because the last draw is by definition a success.

We compute the mean and variance to be

$$\mathbb{E}(X) = \frac{k(1-p)}{p}$$
$$\mathbb{V}(X) = \frac{k(1-p)}{p^2}$$

```
# specify parameters
r, p, n = 10, 0.3, 1_000_000

# draw observations from the distribution
x = rng.negative_binomial(r, p, n)

# compute sample mean and variance
μ_hat = np.mean(x)
σ2_hat = np.var(x)

print("The sample mean is: ", μ_hat, "\nThe sample variance is: ", σ2_hat)
print("\nThe population mean is: ", r*(1-p)/p)
print("The population variance is: ", r*(1-p)/p**2)
```

```
The sample mean is:  23.323736
The sample variance is:  77.687107002304

The population mean is:  23.333333333333336
The population variance is:  77.77777777777779
```

9.4 Newcomb–Benford distribution

The **Newcomb–Benford law** fits many data sets, e.g., reports of incomes to tax authorities, in which the leading digit is more likely to be small than large.

See https://en.wikipedia.org/wiki/Benford's_law

A Benford probability distribution is

$$\text{Prob}\{X = d\} = \log_{10}(d+1) - \log_{10}(d) = \log_{10}\left(1 + \frac{1}{d}\right)$$

where $d \in \{1, 2, \cdots, 9\}$ can be thought of as a **first digit** in a sequence of digits.

This is a well defined discrete distribution since we can verify that probabilities are nonnegative and sum to 1.

$$\log_{10}\left(1 + \frac{1}{d}\right) \geq 0, \quad \sum_{d=1}^{9} \log_{10}\left(1 + \frac{1}{d}\right) = 1$$

The mean and variance of a Benford distribution are

$$\mathbb{E}[X] = \sum_{d=1}^{9} d \log_{10}\left(1 + \frac{1}{d}\right) \simeq 3.4402$$

$$\mathbb{V}[X] = \sum_{d=1}^{9} (d - \mathbb{E}[X])^2 \log_{10}\left(1 + \frac{1}{d}\right) \simeq 6.0565$$

We verify the above and compute the mean and variance using `numpy`.

```
Benford_pmf = np.array([np.log10(1+1/d) for d in range(1,10)])
k = np.arange(1, 10)

# mean
mean = k @ Benford_pmf
```

(continues on next page)

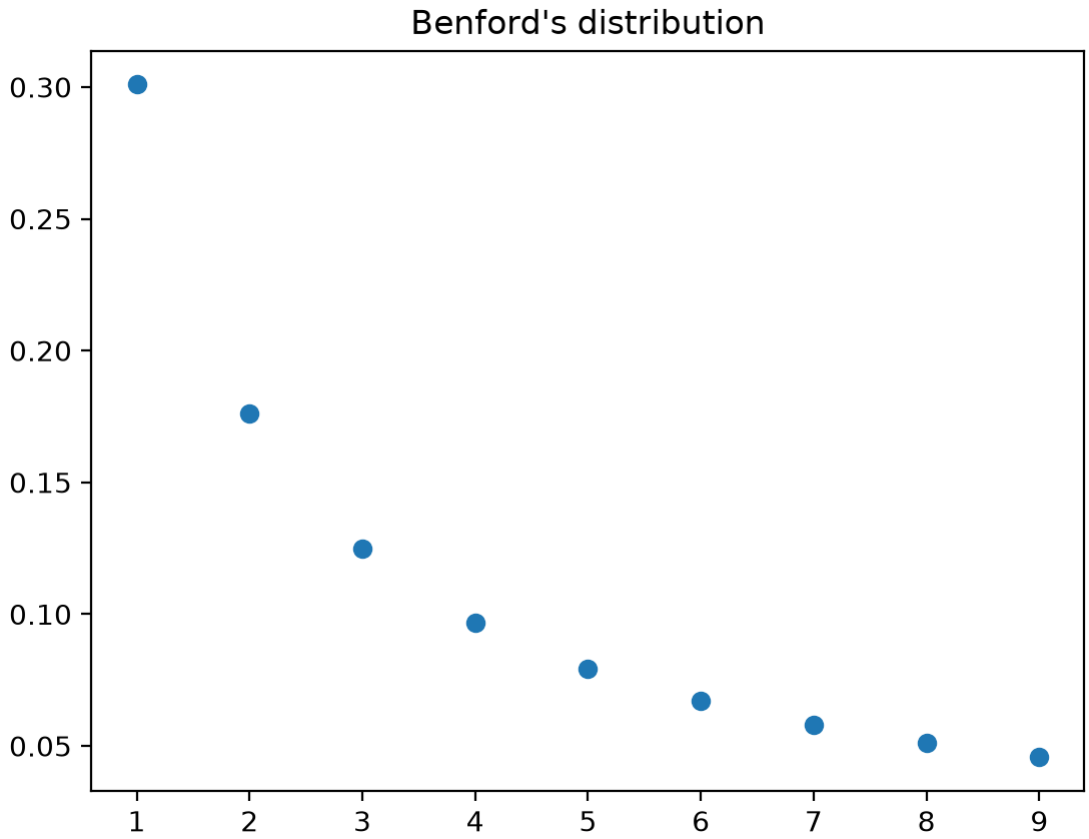
(continued from previous page)

```
# variance
var = ((k - mean) ** 2) @ Benford_pmf

# verify sum to 1
print(np.sum(Benford_pmf))
print(mean)
print(var)
```

```
0.9999999999999999
3.4402369671232065
6.056512631375667
```

```
# plot distribution
plt.plot(range(1,10), Benford_pmf, 'o')
plt.title('Benford\'s distribution')
plt.show()
```



Now let's turn to some continuous random variables.

9.5 Univariate Gaussian distribution

We write

$$X \sim N(\mu, \sigma^2)$$

to indicate the probability distribution

$$f(x|u, \sigma^2) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{[-\frac{1}{2\sigma^2}(x-u)^2]}$$

In the below example, we set $\mu = 0, \sigma = 0.1$.

```
# specify parameters
μ, σ = 0, 0.1

# specify number of draws
n = 1_000_000

# draw observations from the distribution
x = rng.normal(μ, σ, n)

# compute sample mean and variance
μ_hat = np.mean(x)
σ_hat = np.std(x)

print("The sample mean is: ", μ_hat)
print("The sample standard deviation is: ", σ_hat)
```

```
The sample mean is:  0.00014265112036251942
The sample standard deviation is:  0.10015467061449704
```

```
# compare
print(μ-μ_hat < 1e-3)
print(σ-σ_hat < 1e-3)
```

```
True
True
```

9.6 Uniform Distribution

$$X \sim U[a, b]$$

$$f(x) = \begin{cases} \frac{1}{b-a}, & a \leq x \leq b \\ 0, & \text{otherwise} \end{cases}$$

The population mean and variance are

$$\mathbb{E}(X) = \frac{a+b}{2}$$

$$\mathbb{V}(X) = \frac{(b-a)^2}{12}$$

```
# specify parameters
a, b = 10, 20

# specify number of draws
n = 1_000_000

# draw observations from the distribution
x = a + (b-a)*rng.random(n)

# compute sample mean and variance
μ_hat = np.mean(x)
σ2_hat = np.var(x)

print("The sample mean is: ", μ_hat, "\nThe sample variance is: ", σ2_hat)
print("\nThe population mean is: ", (a+b)/2)
print("The population variance is: ", (b-a)**2/12)
```

```
The sample mean is:  14.997280297913704
The sample variance is:  8.329932292674247

The population mean is:  15.0
The population variance is:  8.333333333333334
```

9.7 A Mixed Discrete-Continuous Distribution

We'll motivate this example with a little story.

Suppose that to apply for a job you take an interview and either pass or fail it.

You have 5% chance to pass an interview and you know your salary will uniformly distributed in the interval 300~400 a day only if you pass.

We can describe your daily salary as a discrete-continuous variable with the following probabilities:

$$P(X = 0) = 0.95$$

$$P(300 \leq X \leq 400) = \int_{300}^{400} f(x)\, dx = 0.05$$

$$f(x) = 0.0005$$

Let's start by generating a random sample and computing sample moments.

```
x = rng.random(1_000_000)
# x[x > 0.95] = 100*x[x > 0.95]+300
x[x > 0.95] = 100*rng.random(len(x[x > 0.95]))+300
x[x <= 0.95] = 0

μ_hat = np.mean(x)
σ2_hat = np.var(x)

print("The sample mean is: ", μ_hat, "\nThe sample variance is: ", σ2_hat)
```

```
The sample mean is:  17.48805262165106
The sample variance is:  5853.555467182657
```

The analytical mean and variance can be computed:

$$\begin{aligned}\mu &= \int_{300}^{400} xf(x)dx \\ &= 0.0005 \int_{300}^{400} xdx \\ &= 0.0005 \times \frac{1}{2}x^2\Big|_{300}^{400}\end{aligned}$$

$$\begin{aligned}\sigma^2 &= 0.95 \times (0-17.5)^2 + \int_{300}^{400} (x-17.5)^2 f(x)dx \\ &= 0.95 \times 17.5^2 + 0.0005 \int_{300}^{400} (x-17.5)^2 dx \\ &= 0.95 \times 17.5^2 + 0.0005 \times \frac{1}{3}(x-17.5)^3\Big|_{300}^{400}\end{aligned}$$

```
mean = 0.0005*0.5*(400**2 - 300**2)
var = 0.95*17.5**2+0.0005/3*((400-17.5)**3-(300-17.5)**3)
print("mean: ", mean)
print("variance: ", var)
```

```
mean:  17.5
variance:  5860.416666666666
```

9.8 Drawing a Random Number from a Particular Distribution

Suppose we have at our disposal a pseudo random number that draws a uniform random variable, i.e., one with probability distribution

$$\text{Prob}\{\tilde{X} = i\} = \frac{1}{I}, \quad i = 0, \dots, I-1$$

How can we transform $\tilde{X}$ to get a random variable X for which $\text{Prob}\{X = i\} = f_i, \quad i = 0, \dots, I-1$, where f_i is an arbitary discrete probability distribution on $i = 0, 1, \dots, I-1$?

The key tool is the inverse of a cumulative distribution function (CDF).

Observe that the CDF of a distribution is monotone and non-decreasing, taking values between 0 and 1.

We can draw a sample of a random variable X with a known CDF as follows:

- draw a random variable u from a uniform distribution on $[0, 1]$
- pass the sample value of u into the **"inverse"** target CDF for X
- X has the target CDF

Thus, knowing the **"inverse"** CDF of a distribution is enough to simulate from this distribution.

Note

The "inverse" CDF needs to exist for this method to work.

The inverse CDF is

$$F^{-1}(u) \equiv \inf\{x \in \mathbb{R} : F(x) \geq u\} \quad (0 < u < 1)$$

Here we use infimum because a CDF is a non-decreasing and right-continuous function.

Thus, suppose that

- U is a uniform random variable $U \in [0, 1]$
- We want to sample a random variable X whose CDF is F.

It turns out that if we use draw uniform random numbers U and then compute X from

$$X = F^{-1}(U),$$

then X is a random variable with CDF $F_X(x) = F(x) = \text{Prob}\{X \leq x\}$.

We'll verify this in the special case in which F is continuous and bijective so that its inverse function exists and can be denoted by F^{-1}.

Note that

$$\begin{aligned} F_X(x) &= \text{Prob}\{X \leq x\} \\ &= \text{Prob}\left\{F^{-1}(U) \leq x\right\} \\ &= \text{Prob}\{U \leq F(x)\} \\ &= F(x) \end{aligned}$$

where the last equality occurs because U is distributed uniformly on $[0, 1]$ while $F(x)$ is a constant given x that also lies on $[0, 1]$.

Let's use `numpy` to compute some examples.

Example: A continuous geometric (exponential) distribution

Let X follow a geometric distribution, with parameter $\lambda > 0$.

Its density function is

$$f(x) = \lambda e^{-\lambda x}$$

Its CDF is

$$F(x) = \int_0^\infty \lambda e^{-\lambda x} = 1 - e^{-\lambda x}$$

Let U follow a uniform distribution on $[0, 1]$.

X is a random variable such that $U = F(X)$.

The distribution X can be deduced from

$$\begin{aligned} U = F(X) &= 1 - e^{-\lambda X} \\ \implies \quad -U &= e^{-\lambda X} \\ \implies \quad \log(1-U) &= -\lambda X \\ \implies \quad X &= \frac{(1-U)}{-\lambda} \end{aligned}$$

Let's draw u from $U[0, 1]$ and calculate $x = \frac{log(1-U)}{-\lambda}$.

We'll check whether X seems to follow a **continuous geometric** (exponential) distribution.

Let's check with `numpy`.

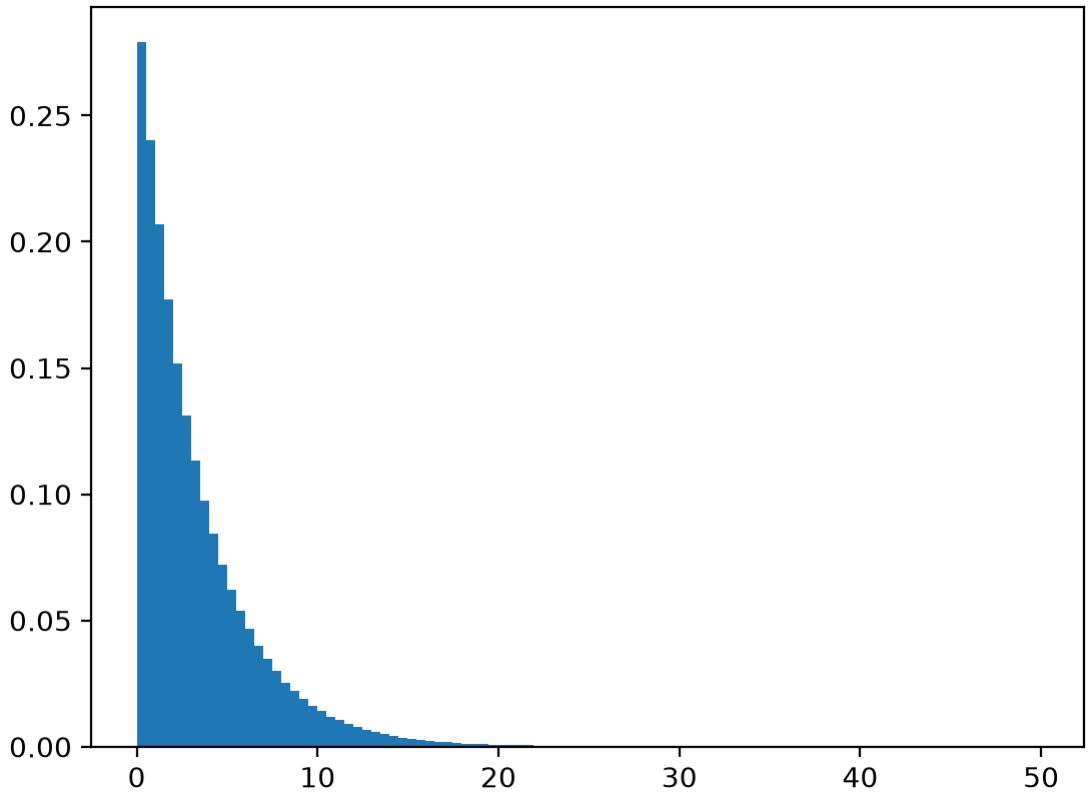
```
n, λ = 1_000_000, 0.3

# draw uniform numbers
u = rng.random(n)

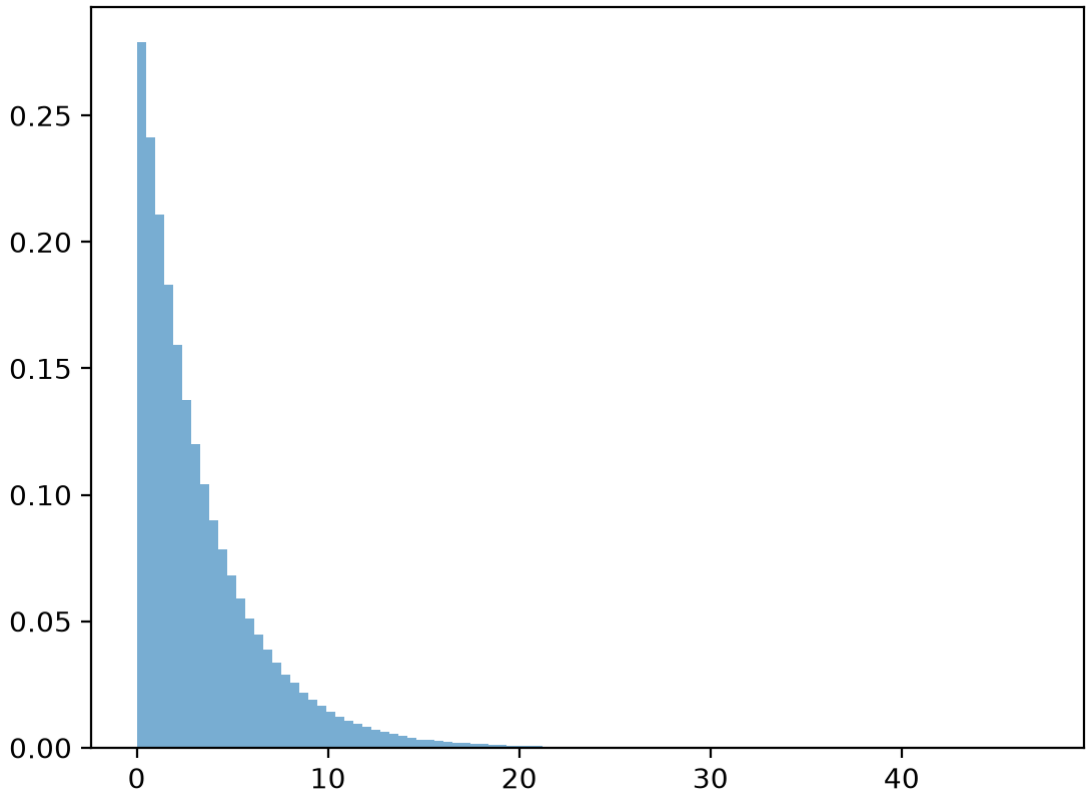
# transform
x = -np.log(1-u)/λ

# draw geometric distributions
x_g = rng.exponential(1 / λ, n)

# plot and compare
plt.hist(x, bins=100, density=True)
plt.show()
```




```
plt.hist(x_g, bins=100, density=True, alpha=0.6)
plt.show()
```



Geometric distribution

Let X distributed geometrically, that is

$$\text{Prob}(X = i) = (1 - \lambda)\lambda^i, \quad \lambda \in (0, 1), \quad i = 0, 1, ...$$

$$\sum_{i=0}^{\infty} \text{Prob}(X = i) = 1 \longleftrightarrow (1 - \lambda) \sum_{i=0}^{\infty} \lambda^i = \frac{1 - \lambda}{1 - \lambda} = 1$$

Its CDF is given by

$$\begin{aligned} \text{Prob}(X \leq i) &= (1 - \lambda) \sum_{j=0}^{i} \lambda^i \\ &= (1 - \lambda)[\frac{1 - \lambda^{i+1}}{1 - \lambda}] \\ &= 1 - \lambda^{i+1} \\ &= F(X) = F_i \end{aligned}$$

Again, let $\tilde{U}$ follow a uniform distribution and we want to find X such that $F(X) = \tilde{U}$.

Let's deduce the distribution of X from

$$\begin{aligned} \tilde{U} = F(X) = 1 - \lambda^{x+1} \\ 1 - \tilde{U} = \lambda^{x+1} \\ \log(1 - \tilde{U}) = (x + 1) \log \lambda \\ \frac{\log(1 - \tilde{U})}{\log \lambda} = x + 1 \\ \frac{\log(1 - \tilde{U})}{\log \lambda} - 1 = x \end{aligned}$$

However, $\tilde{U} = F^{-1}(X)$ may not be an integer for any $x \geq 0$.

So let

$$x = \lceil \frac{\log(1 - \tilde{U})}{\log \lambda} - 1 \rceil$$

where $\lceil . \rceil$ is the ceiling function.

Thus x is the smallest integer such that the discrete geometric CDF is greater than or equal to $\tilde{U}$.

We can verify that x is indeed geometrically distributed by the following `numpy` program.

Note

The exponential distribution is the continuous analog of geometric distribution.

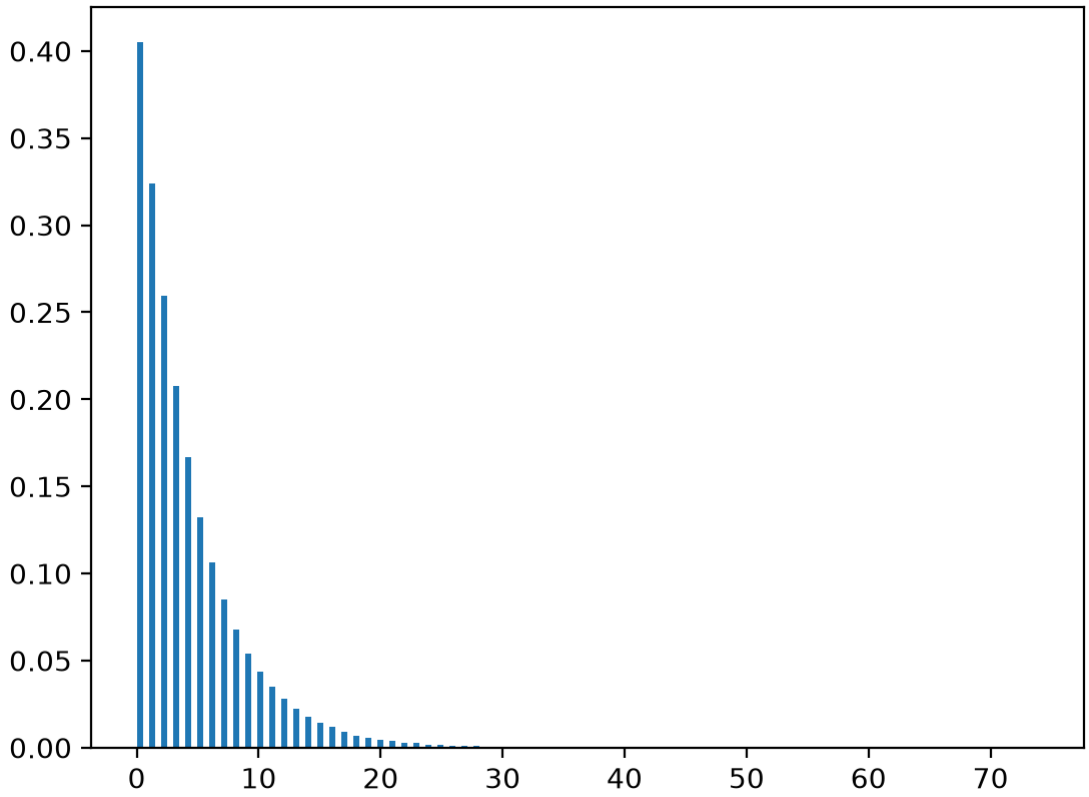
```
n, λ = 1_000_000, 0.8

# draw uniform numbers
u = rng.random(n)

# transform
x = np.ceil(np.log(1-u)/np.log(λ) - 1)

# draw geometric distributions
x_g = rng.geometric(1-λ, n)

# plot and compare
plt.hist(x, bins=150, density=True)
plt.show()
```



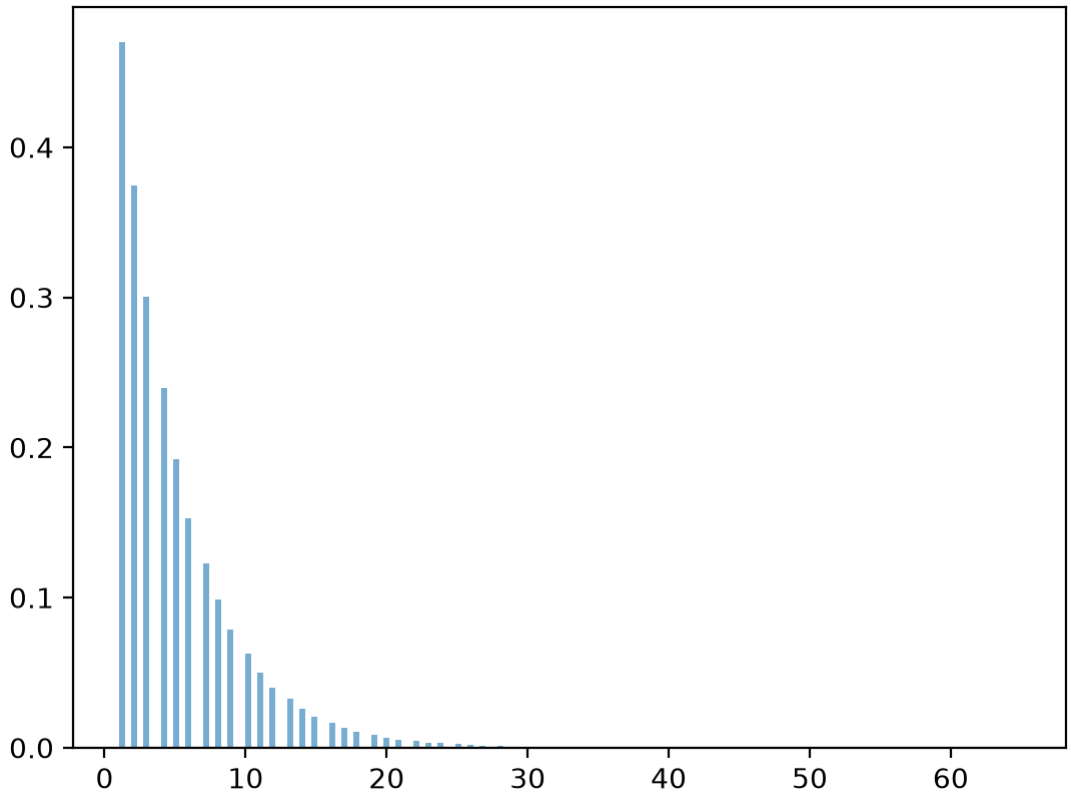

```
rng.geometric(1-λ, n).max()
```

```
np.int64(65)
```

```
np.log(0.4)/np.log(0.3)
```

```
np.float64(0.7610560044063083)
```

```
plt.hist(x_g, bins=150, density=True, alpha=0.6)
plt.show()
```


0.4
0.3
0.2
0.1
0.0
0
10
20
30
40
50
60

CHAPTER

TEN

LLN AND CLT

Contents

- *LLN and CLT*
  - *Overview*
  - *Relationships*
  - *LLN*
  - *CLT*
  - *Exercises*

10.1 Overview

This lecture illustrates two of the most important theorems of probability and statistics: The law of large numbers (LLN) and the central limit theorem (CLT).

These beautiful theorems lie behind many of the most fundamental results in econometrics and quantitative economic modeling.

The lecture is based around simulations that show the LLN and CLT in action.

We also demonstrate how the LLN and CLT break down when the assumptions they are based on do not hold.

In addition, we examine several useful extensions of the classical theorems, such as

- The delta method, for smooth functions of random variables, and
- the multivariate case.

Some of these extensions are presented as exercises.

We'll need the following imports:

```
import matplotlib.pyplot as plt
import random
import numpy as np
from scipy.stats import t, beta, lognorm, expon, gamma, uniform
from scipy.stats import gaussian_kde, poisson, binom, norm, chi2
from mpl_toolkits.mplot3d import Axes3D
```

(continues on next page)

(continued from previous page)

```
from matplotlib.collections import PolyCollection
from scipy.linalg import inv, sqrtm
```

10.2 Relationships

The CLT refines the LLN.

The LLN gives conditions under which sample moments converge to population moments as sample size increases.

The CLT provides information about the rate at which sample moments converge to population moments as sample size increases.

10.3 LLN

We begin with the law of large numbers, which tells us when sample averages will converge to their population means.

10.3.1 The Classical LLN

The classical law of large numbers concerns independent and identically distributed (IID) random variables.

Here is the strongest version of the classical LLN, known as **Kolmogorov's strong law**.

Let $X_1, \ldots, X_n$ be independent and identically distributed scalar random variables, with common distribution F.

When it exists, let μ denote the common mean of this sample:

$$\mu := \mathbb{E}X = \int x F(dx)$$

In addition, let

$$\bar{X}_n := \frac{1}{n} \sum_{i=1}^{n} X_i$$

Kolmogorov's strong law states that, if $\mathbb{E}|X|$ is finite, then

$$\mathbb{P}\left\{\bar{X}_n \to \mu \text{ as } n \to \infty\right\} = 1 \tag{10.1}$$

What does this last expression mean?

Let's think about it from a simulation perspective, imagining for a moment that our computer can generate perfect random samples (which of course it can't).

Let's also imagine that we can generate infinite sequences so that the statement $\bar{X}_n \to \mu$ can be evaluated.

In this setting, (10.1) should be interpreted as meaning that the probability of the computer producing a sequence where $\bar{X}_n \to \mu$ fails to occur is zero.

10.3.2 Proof

The proof of Kolmogorov's strong law is nontrivial – see, for example, theorem 8.3.5 of [Dudley, 2002].

On the other hand, we can prove a weaker version of the LLN very easily and still get most of the intuition.

The version we prove is as follows: If $X_1, \ldots, X_n$ is IID with $\mathbb{E}X_i^2 < \infty$, then, for any $\epsilon > 0$, we have

$$\mathbb{P}\left\{|\bar{X}_n - \mu| \geq \epsilon\right\} \to 0 \quad \text{as} \quad n \to \infty \tag{10.2}$$

(This version is weaker because we claim only convergence in probability rather than almost sure convergence, and assume a finite second moment)

To see that this is so, fix $\epsilon > 0$, and let σ^2 be the variance of each X_i.

Recall the Chebyshev inequality, which tells us that

$$\mathbb{P}\left\{|\bar{X}_n - \mu| \geq \epsilon\right\} \leq \frac{\mathbb{E}[(\bar{X}_n - \mu)^2]}{\epsilon^2} \tag{10.3}$$

Now observe that

$$\begin{aligned}
\mathbb{E}[(\bar{X}_n - \mu)^2] &= \mathbb{E}\left\{\left[\frac{1}{n}\sum_{i=1}^{n}(X_i - \mu)\right]^2\right\} \\
&= \frac{1}{n^2}\sum_{i=1}^{n}\sum_{j=1}^{n}\mathbb{E}(X_i - \mu)(X_j - \mu) \\
&= \frac{1}{n^2}\sum_{i=1}^{n}\mathbb{E}(X_i - \mu)^2 \\
&= \frac{\sigma^2}{n}
\end{aligned}$$

Here the crucial step is at the third equality, which follows from independence.

Independence means that if $i \neq j$, then the covariance term $\mathbb{E}(X_i - \mu)(X_j - \mu)$ drops out.

As a result, $n^2 - n$ terms vanish, leading us to a final expression that goes to zero in n.

Combining our last result with (10.3), we come to the estimate

$$\mathbb{P}\left\{|\bar{X}_n - \mu| \geq \epsilon\right\} \leq \frac{\sigma^2}{n\epsilon^2} \tag{10.4}$$

The claim in (10.2) is now clear.

Of course, if the sequence $X_1, \ldots, X_n$ is correlated, then the cross-product terms $\mathbb{E}(X_i - \mu)(X_j - \mu)$ are not necessarily zero.

While this doesn't mean that the same line of argument is impossible, it does mean that if we want a similar result then the covariances should be "almost zero" for "most" of these terms.

In a long sequence, this would be true if, for example, $\mathbb{E}(X_i - \mu)(X_j - \mu)$ approached zero when the difference between i and j became large.

In other words, the LLN can still work if the sequence $X_1, \ldots, X_n$ has a kind of "asymptotic independence", in the sense that correlation falls to zero as variables become further apart in the sequence.

This idea is very important in time series analysis, and we'll come across it again soon enough.

10.3.3 Illustration

Let's now illustrate the classical IID law of large numbers using simulation.

In particular, we aim to generate some sequences of IID random variables and plot the evolution of $\bar{X}_n$ as n increases.

Below is a figure that does just this (as usual, you can click on it to expand it).

It shows IID observations from three different distributions and plots $\bar{X}_n$ against n in each case.

The dots represent the underlying observations X_i for $i = 1, \ldots, 100$.

In each of the three cases, convergence of $\bar{X}_n$ to μ occurs as predicted

```
n = 100

# Arbitrary collection of distributions
distributions = {"student's t with 10 degrees of freedom": t(10),
                 "β(2, 2)": beta(2, 2),
                 "lognormal LN(0, 1/2)": lognorm(0.5),
                 "γ(5, 1/2)": gamma(5, scale=2),
                 "poisson(4)": poisson(4),
                 "exponential with λ = 1": expon(1)}

# Create a figure and some axes
num_plots = 3
fig, axes = plt.subplots(num_plots, 1, figsize=(10, 20))

# Set some plotting parameters to improve layout
bbox = (0., 1.02, 1., .102)
legend_args = {'ncol': 2,
               'bbox_to_anchor': bbox,
               'loc': 3,
               'mode': 'expand'}
plt.subplots_adjust(hspace=0.5)

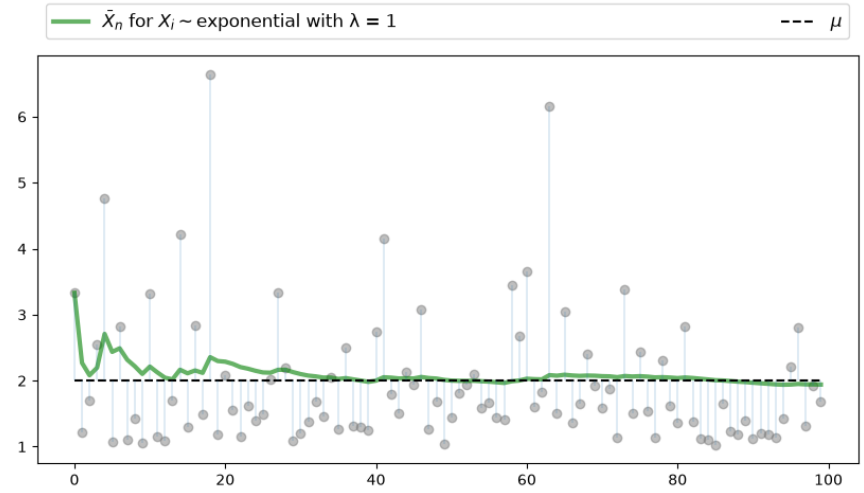
for ax in axes:
    # Choose a randomly selected distribution
    name = random.choice(list(distributions.keys()))
    distribution = distributions.pop(name)

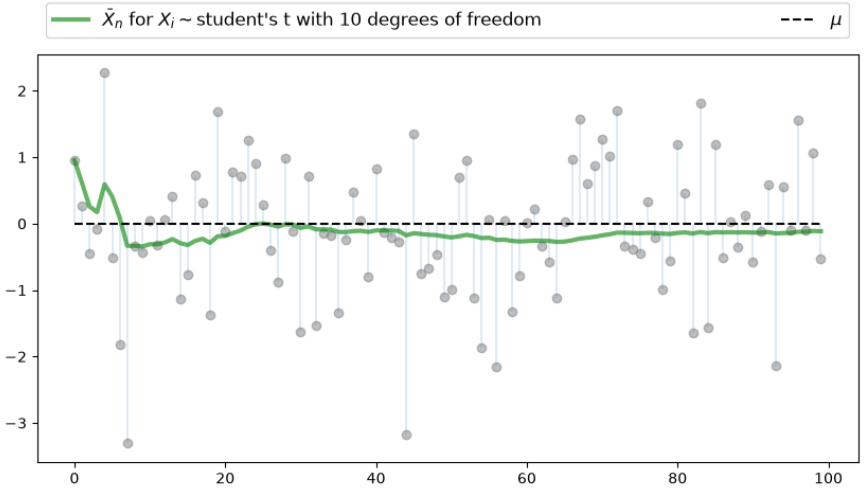
    # Generate n draws from the distribution
    data = distribution.rvs(n)

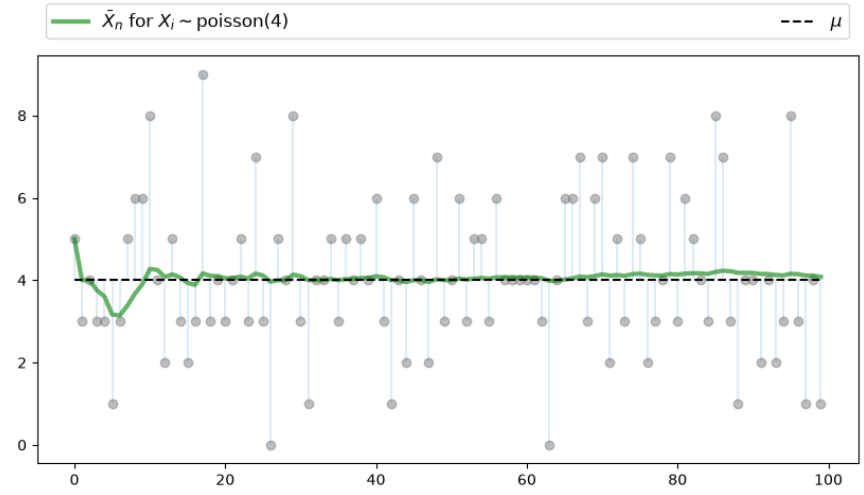
    # Compute sample mean at each n
    sample_mean = np.empty(n)
    for i in range(n):
        sample_mean[i] = np.mean(data[:i+1])

    # Plot
    ax.plot(list(range(n)), data, 'o', color='grey', alpha=0.5)
    axlabel = r'$\bar{X}_n$ for $X_i \sim$' + name
    ax.plot(list(range(n)), sample_mean, 'g-', lw=3, alpha=0.6, label=axlabel)
    m = distribution.mean()
    ax.plot(list(range(n)), [m] * n, 'k--', lw=1.5, label=r'$\mu$')
    ax.vlines(list(range(n)), m, data, lw=0.2)
    ax.legend(**legend_args, fontsize=12)

plt.show()
```


$\bar{X}_n$ for $X_i \sim$ exponential with $\lambda = 1$
μ
1
2
3
4
5
6
0
20
40
60
80
100


$\bar{X}_n$ for $X_i \sim$ student's t with 10 degrees of freedom
μ
−3
−2
−1
0
1
2
0
20
40
60
80
100


$\bar{X}_n$ for $X_i \sim$ poisson(4)
μ
0
2
4
6
8
0
20
40
60
80
100

The three distributions are chosen at random from a selection stored in the dictionary `distributions`.

10.4 CLT

Next, we turn to the central limit theorem, which tells us about the distribution of the deviation between sample averages and population means.

10.4.1 Statement of the Theorem

The central limit theorem is one of the most remarkable results in all of mathematics.

In the classical IID setting, it tells us the following:

If the sequence $X_1, \ldots, X_n$ is IID, with common mean μ and common variance $\sigma^2 \in (0, \infty)$, then

$$\sqrt{n}(\bar{X}_n - \mu) \xrightarrow{d} N(0, \sigma^2) \quad \text{as} \quad n \to \infty \tag{10.5}$$

Here $\xrightarrow{d} N(0, \sigma^2)$ indicates convergence in distribution to a centered (i.e, zero mean) normal with standard deviation σ.

10.4.2 Intuition

The striking implication of the CLT is that for **any** distribution with finite second moment, the simple operation of adding independent copies **always** leads to a Gaussian curve.

A relatively simple proof of the central limit theorem can be obtained by working with characteristic functions (see, e.g., theorem 9.5.6 of [Dudley, 2002]).

The proof is elegant but almost anticlimactic, and it provides surprisingly little intuition.

In fact, all of the proofs of the CLT that we know are similar in this respect.

Why does adding independent copies produce a bell-shaped distribution?

Part of the answer can be obtained by investigating the addition of independent Bernoulli random variables.

In particular, let X_i be binary, with $\mathbb{P}\{X_i = 0\} = \mathbb{P}\{X_i = 1\} = 0.5$, and let $X_1, \ldots, X_n$ be independent.

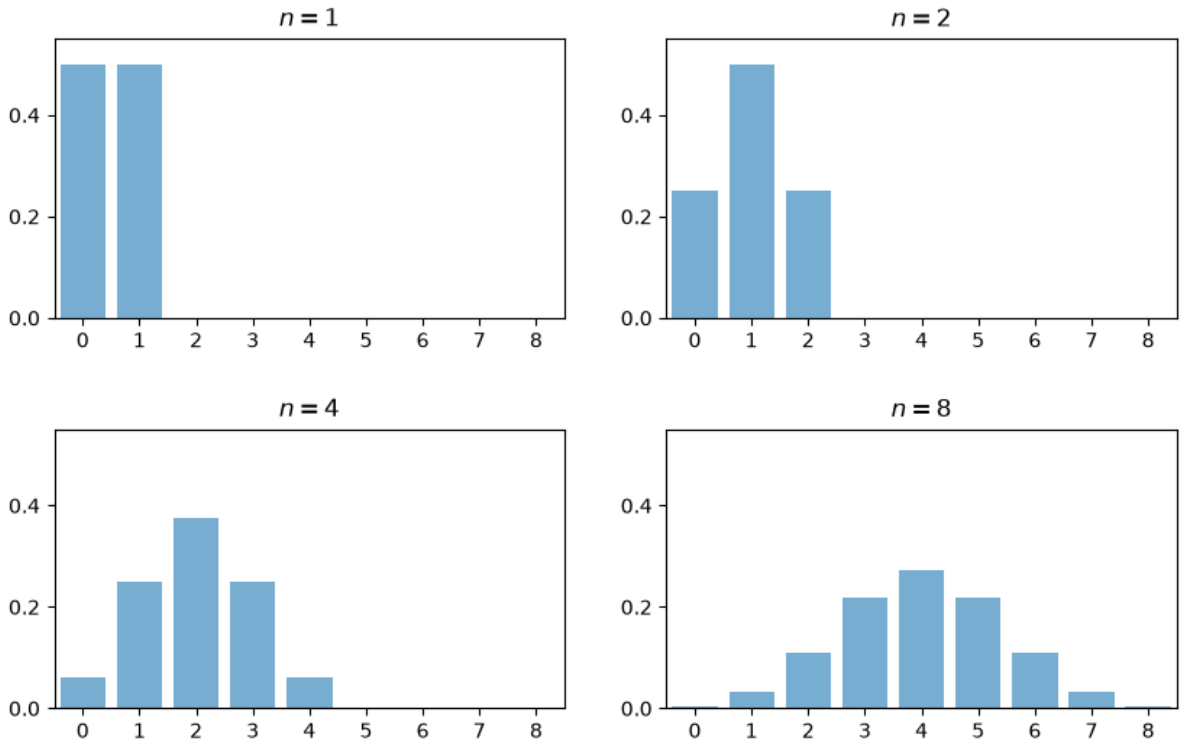
Think of $X_i = 1$ as a "success", so that $Y_n = \sum_{i=1}^n X_i$ is the number of successes in n trials.

The next figure plots the probability mass function of Y_n for $n = 1, 2, 4, 8$

```
fig, axes = plt.subplots(2, 2, figsize=(10, 6))
plt.subplots_adjust(hspace=0.4)
axes = axes.flatten()
ns = [1, 2, 4, 8]
dom = list(range(9))

for ax, n in zip(axes, ns):
    b = binom(n, 0.5)
    ax.bar(dom, b.pmf(dom), alpha=0.6, align='center')
    ax.set(xlim=(-0.5, 8.5), ylim=(0, 0.55),
           xticks=list(range(9)), yticks=(0, 0.2, 0.4),
           title=f'$n = {n}$')

plt.show()
```



When $n = 1$, the distribution is flat — one success or no successes have the same probability.

When $n = 2$ we can either have 0, 1 or 2 successes.

Notice the peak in probability mass at the mid-point $k = 1$.

The reason is that there are more ways to get 1 success ("fail then succeed" or "succeed then fail") than to get zero or two successes.

Moreover, the two trials are independent, so the outcomes "fail then succeed" and "succeed then fail" are just as likely as the outcomes "fail then fail" and "succeed then succeed".

(If there was positive correlation, say, then "succeed then fail" would be less likely than "succeed then succeed")

Here, already we have the essence of the CLT: addition under independence leads probability mass to pile up in the middle and thin out at the tails.

For $n = 4$ and $n = 8$ we again get a peak at the "middle" value (halfway between the minimum and the maximum possible value).

The intuition is the same — there are simply more ways to get these middle outcomes.

If we continue, the bell-shaped curve becomes even more pronounced.

We are witnessing the binomial approximation of the normal distribution.

10.4.3 Simulation 1

Since the CLT seems almost magical, running simulations that verify its implications is one good way to build intuition.

To this end, we now perform the following simulation

1. Choose an arbitrary distribution F for the underlying observations X_i.
2. Generate independent draws of $Y_n := \sqrt{n}(\bar{X}_n - \mu)$.
3. Use these draws to compute some measure of their distribution — such as a histogram.
4. Compare the latter to $N(0, \sigma^2)$.

Here's some code that does exactly this for the exponential distribution $F(x) = 1 - e^{-\lambda x}$.

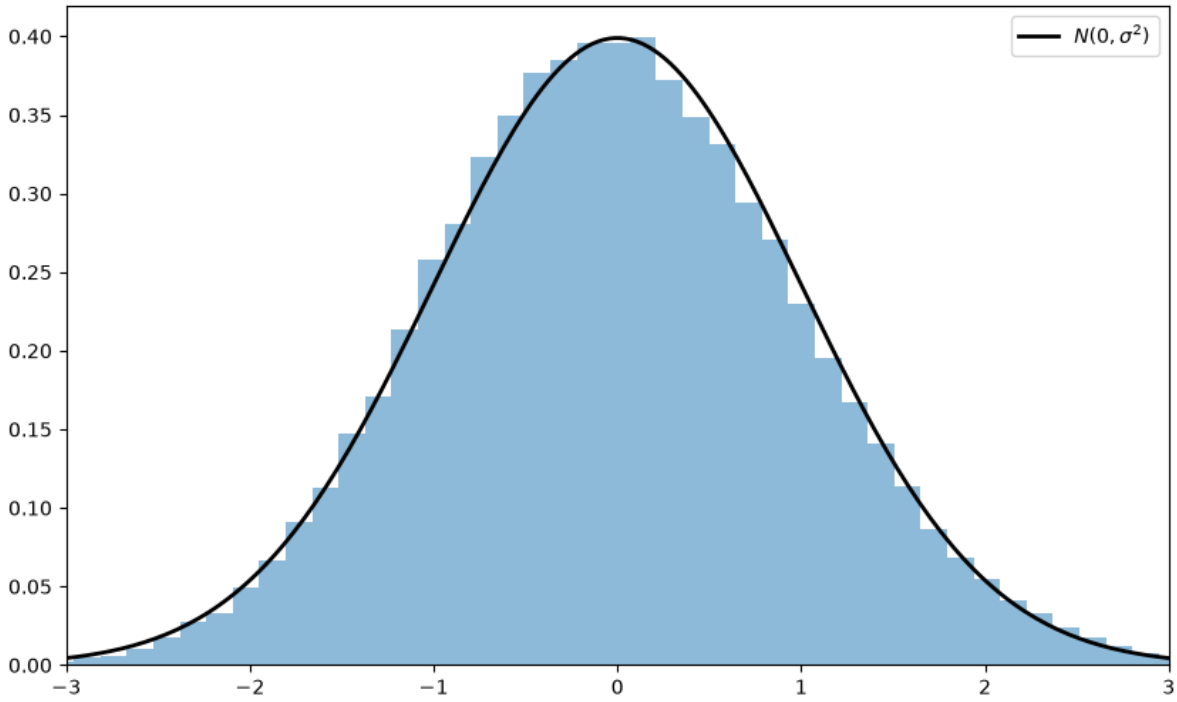
(Please experiment with other choices of F, but remember that, to conform with the conditions of the CLT, the distribution must have a finite second moment.)

```
# Set parameters
n = 250                  # Choice of n
k = 100000               # Number of draws of Y_n
distribution = expon(2)  # Exponential distribution, λ = 1/2
μ, s = distribution.mean(), distribution.std()

# Draw underlying RVs. Each row contains a draw of X_1,..,X_n
data = distribution.rvs((k, n))
# Compute mean of each row, producing k draws of \bar X_n
sample_means = data.mean(axis=1)
# Generate observations of Y_n
Y = np.sqrt(n) * (sample_means - μ)

# Plot
fig, ax = plt.subplots(figsize=(10, 6))
xmin, xmax = -3 * s, 3 * s
ax.set_xlim(xmin, xmax)
ax.hist(Y, bins=60, alpha=0.5, density=True)
xgrid = np.linspace(xmin, xmax, 200)
ax.plot(xgrid, norm.pdf(xgrid, scale=s), 'k-', lw=2, label=r'$N(0, \sigma^2)$')
ax.legend()

plt.show()
```



Notice the absence of for loops — every operation is vectorized, meaning that the major calculations are all shifted to highly optimized C code.

The fit to the normal density is already tight and can be further improved by increasing `n`.

You can also experiment with other specifications of F.

10.4.4 Simulation 2

Our next simulation is somewhat like the first, except that we aim to track the distribution of $Y_n := \sqrt{n}(\bar{X}_n - \mu)$ as n increases.

In the simulation, we'll be working with random variables having $\mu = 0$.

Thus, when $n = 1$, we have $Y_1 = X_1$, so the first distribution is just the distribution of the underlying random variable.

For $n = 2$, the distribution of Y_2 is that of $(X_1 + X_2)/\sqrt{2}$, and so on.

What we expect is that, regardless of the distribution of the underlying random variable, the distribution of Y_n will smooth out into a bell-shaped curve.

The next figure shows this process for $X_i \sim f$, where f was specified as the convex combination of three different beta densities.

(Taking a convex combination is an easy way to produce an irregular shape for f.)

In the figure, the closest density is that of Y_1, while the furthest is that of Y_5

```
beta_dist = beta(2, 2)

def gen_x_draws(k, rng):
    """
    Returns a flat array containing k independent draws from the
    distribution of X, the underlying random variable.  This distribution
```

(continues on next page)

(continued from previous page)

```
    is itself a convex combination of three beta distributions.
    """
    bdraws = beta_dist.rvs((3, k))
    # Transform rows, so each represents a different distribution
    bdraws[0, :] -= 0.5
    bdraws[1, :] += 0.6
    bdraws[2, :] -= 1.1
    # Set X[i] = bdraws[j, i], where j is a random draw from {0, 1, 2}
    js = rng.integers(0, 2, size=k)
    X = bdraws[js, np.arange(k)]
    # Rescale, so that the random variable is zero mean
    m, sigma = X.mean(), X.std()
    return (X - m) / sigma

nmax = 5
reps = 100000
ns = list(range(1, nmax + 1))
rng = np.random.default_rng()

# Form a matrix Z such that each column is reps independent draws of X
Z = np.empty((reps, nmax))
for i in range(nmax):
    Z[:, i] = gen_x_draws(reps, rng)
# Take cumulative sum across columns
S = Z.cumsum(axis=1)
# Multiply j-th column by sqrt j
Y = (1 / np.sqrt(ns)) * S

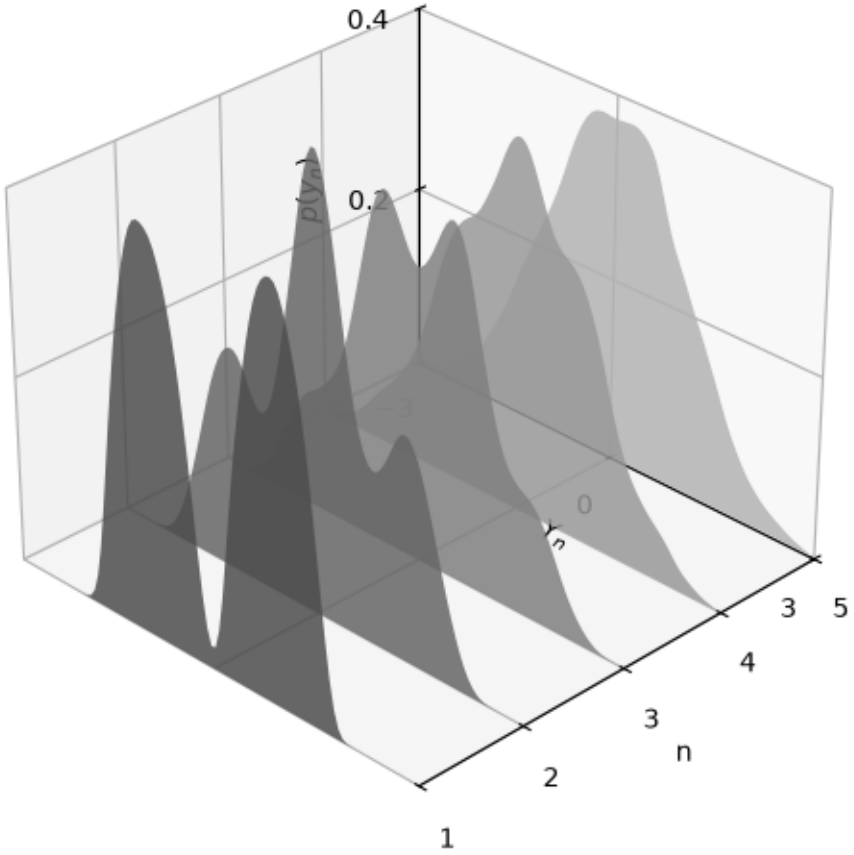
# Plot
ax = plt.figure(figsize = (10, 6)).add_subplot(projection='3d')

a, b = -3, 3
gs = 100
xs = np.linspace(a, b, gs)

# Build verts
greys = np.linspace(0.3, 0.7, nmax)
verts = []
for n in ns:
    density = gaussian_kde(Y[:, n-1])
    ys = density(xs)
    verts.append(list(zip(xs, ys)))

poly = PolyCollection(verts, facecolors=[str(g) for g in greys])
poly.set_alpha(0.85)
ax.add_collection3d(poly, zs=ns, zdir='x')

ax.set(xlim3d=(1, nmax), xticks=(ns), ylabel='$Y_n$', zlabel='$p(y_n)$',
       xlabel=("n"), yticks=((-3, 0, 3)), ylim3d=(a, b),
       zlim3d=(0, 0.4), zticks=((0.2, 0.4)))
ax.invert_xaxis()
# Rotates the plot 30 deg on z axis and 45 deg on x axis
ax.view_init(30, 45)
plt.show()
```



As expected, the distribution smooths out into a bell curve as n increases.

We leave you to investigate its contents if you wish to know more.

If you run the file from the ordinary IPython shell, the figure should pop up in a window that you can rotate with your mouse, giving different views on the density sequence.

10.4.5 The Multivariate Case

The law of large numbers and central limit theorem work just as nicely in multidimensional settings.

To state the results, let's recall some elementary facts about random vectors.

A random vector $\mathbf{X}$ is just a sequence of k random variables $(X_1, \ldots, X_k)$.

Each realization of $\mathbf{X}$ is an element of $\mathbb{R}^k$.

A collection of random vectors $\mathbf{X}_1, \ldots, \mathbf{X}_n$ is called independent if, given any n vectors $\mathbf{x}_1, \ldots, \mathbf{x}_n$ in $\mathbb{R}^k$, we have

$$\mathbb{P}\{\mathbf{X}_1 \leq \mathbf{x}_1, \ldots, \mathbf{X}_n \leq \mathbf{x}_n\} = \mathbb{P}\{\mathbf{X}_1 \leq \mathbf{x}_1\} \times \cdots \times \mathbb{P}\{\mathbf{X}_n \leq \mathbf{x}_n\}$$

(The vector inequality $\mathbf{X} \leq \mathbf{x}$ means that $X_j \leq x_j$ for $j = 1, \ldots, k$)

Let $\mu_j := \mathbb{E}[X_j]$ for all $j = 1, \ldots, k$.

The expectation $\mathbb{E}[\mathbf{X}]$ of $\mathbf{X}$ is defined to be the vector of expectations:

$$\mathbb{E}[\mathbf{X}] := \begin{pmatrix} \mathbb{E}[X_1] \\ \mathbb{E}[X_2] \\ \vdots \\ \mathbb{E}[X_k] \end{pmatrix} = \begin{pmatrix} \mu_1 \\ \mu_2 \\ \vdots \\ \mu_k \end{pmatrix} =: \mu$$

The **variance-covariance matrix** of random vector $\mathbf{X}$ is defined as

$$\text{Var}[\mathbf{X}] := \mathbb{E}[(\mathbf{X} - \mu)(\mathbf{X} - \mu)']$$

Expanding this out, we get

$$\text{Var}[\mathbf{X}] = \begin{pmatrix} \mathbb{E}[(X_1 - \mu_1)(X_1 - \mu_1)] & \cdots & \mathbb{E}[(X_1 - \mu_1)(X_k - \mu_k)] \\ \mathbb{E}[(X_2 - \mu_2)(X_1 - \mu_1)] & \cdots & \mathbb{E}[(X_2 - \mu_2)(X_k - \mu_k)] \\ \vdots & \vdots & \vdots \\ \mathbb{E}[(X_k - \mu_k)(X_1 - \mu_1)] & \cdots & \mathbb{E}[(X_k - \mu_k)(X_k - \mu_k)] \end{pmatrix}$$

The j, k-th term is the scalar covariance between X_j and X_k.

With this notation, we can proceed to the multivariate LLN and CLT.

Let $\mathbf{X}_1, \ldots, \mathbf{X}_n$ be a sequence of independent and identically distributed random vectors, each one taking values in $\mathbb{R}^k$.

Let μ be the vector $\mathbb{E}[\mathbf{X}_i]$, and let Σ be the variance-covariance matrix of $\mathbf{X}_i$.

Interpreting vector addition and scalar multiplication in the usual way (i.e., pointwise), let

$$\bar{\mathbf{X}}_n := \frac{1}{n} \sum_{i=1}^{n} \mathbf{X}_i$$

In this setting, the LLN tells us that

$$\mathbb{P}\left\{\bar{\mathbf{X}}_n \to \mu \text{ as } n \to \infty\right\} = 1 \tag{10.6}$$

Here $\bar{\mathbf{X}}_n \to \mu$ means that $\|\bar{\mathbf{X}}_n - \mu\| \to 0$, where $\|\cdot\|$ is the standard Euclidean norm.

The CLT tells us that, provided Σ is finite,

$$\sqrt{n}(\bar{\mathbf{X}}_n - \mu) \xrightarrow{d} N(\mathbf{0}, \Sigma) \quad \text{as} \quad n \to \infty \tag{10.7}$$

10.5 Exercises

Exercise 10.5.1

One very useful consequence of the central limit theorem is as follows.

Assume the conditions of the CLT as *stated above*.

If $g\colon \mathbb{R} \to \mathbb{R}$ is differentiable at μ and $g'(\mu) \neq 0$, then

$$\sqrt{n}\{g(\bar{X}_n) - g(\mu)\} \xrightarrow{d} N(0, g'(\mu)^2 \sigma^2) \quad \text{as} \quad n \to \infty \tag{10.8}$$

This theorem is used frequently in statistics to obtain the asymptotic distribution of estimators — many of which can be expressed as functions of sample means.

(These kinds of results are often said to use the "delta method".)

The proof is based on a Taylor expansion of g around the point μ.

Taking the result as given, let the distribution F of each X_i be uniform on $[0, \pi/2]$ and let $g(x) = \sin(x)$.

Derive the asymptotic distribution of $\sqrt{n}\{g(\bar{X}_n) - g(\mu)\}$ and illustrate convergence in the same spirit as the program discussed *above*.

What happens when you replace $[0, \pi/2]$ with $[0, \pi]$?

What is the source of the problem?

Solution

Here is one solution

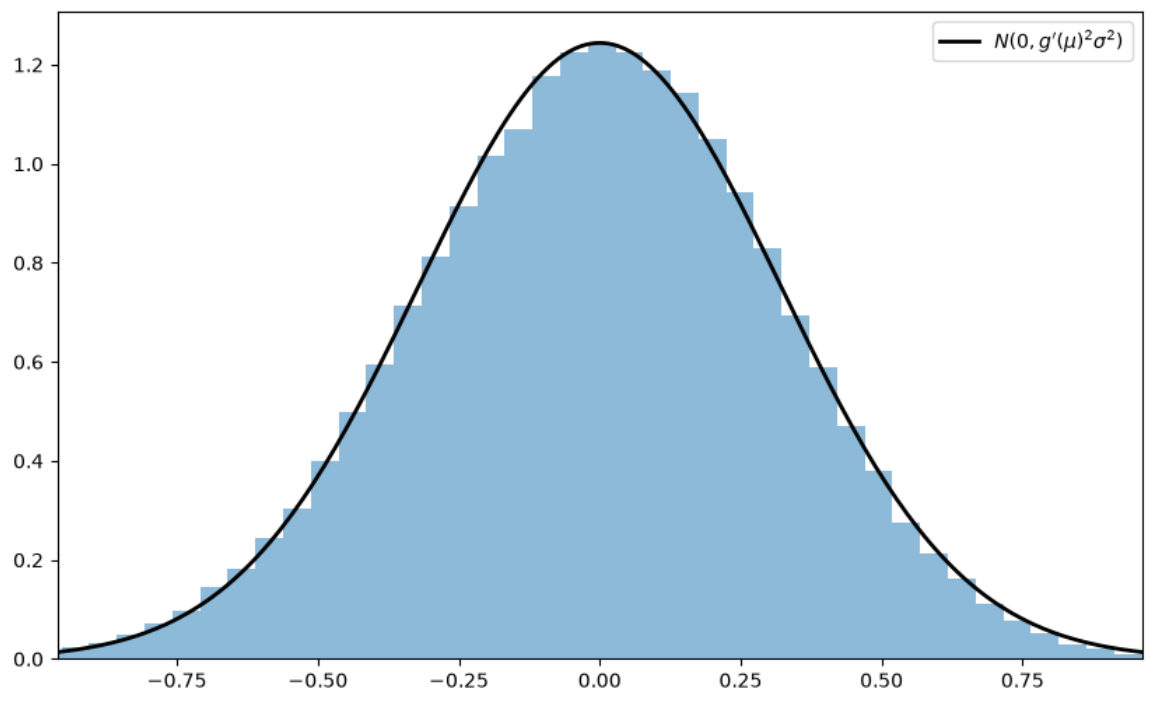
```
"""
Illustrates the delta method, a consequence of the central limit theorem.
"""

# Set parameters
n = 250
replications = 100000
distribution = uniform(loc=0, scale=(np.pi / 2))
μ, s = distribution.mean(), distribution.std()

g = np.sin
g_prime = np.cos

# Generate obs of sqrt{n} (g(X_n) - g(μ))
data = distribution.rvs((replications, n))
sample_means = data.mean(axis=1)  # Compute mean of each row
error_obs = np.sqrt(n) * (g(sample_means) - g(μ))

# Plot
asymptotic_sd = g_prime(μ) * s
fig, ax = plt.subplots(figsize=(10, 6))
xmin = -3 * g_prime(μ) * s
xmax = -xmin
ax.set_xlim(xmin, xmax)
ax.hist(error_obs, bins=60, alpha=0.5, density=True)
xgrid = np.linspace(xmin, xmax, 200)
lb = r"$N(0, g'(\mu)^2  \sigma^2)$"
ax.plot(xgrid, norm.pdf(xgrid, scale=asymptotic_sd), 'k-', lw=2, label=lb)
ax.legend()
plt.show()
```



What happens when you replace $[0, \pi/2]$ with $[0, \pi]$?

In this case, the mean μ of this distribution is $\pi/2$, and since $g' = \cos$, we have $g'(\mu) = 0$.

Hence the conditions of the delta theorem are not satisfied.

Exercise 10.5.2

Here's a result that's often used in developing statistical tests, and is connected to the multivariate central limit theorem.

If you study econometric theory, you will see this result used again and again.

Assume the setting of the multivariate CLT *discussed above*, so that

1. $\mathbf{X}_1, \ldots, \mathbf{X}_n$ is a sequence of IID random vectors, each taking values in $\mathbb{R}^k$.
2. $\mu := \mathbb{E}[\mathbf{X}_i]$, and Σ is the variance-covariance matrix of $\mathbf{X}_i$.
3. The convergence

$$\sqrt{n}(\bar{\mathbf{X}}_n - \mu) \xrightarrow{d} N(\mathbf{0}, \Sigma) \tag{10.9}$$

is valid.

In a statistical setting, one often wants the right-hand side to be **standard** normal so that confidence intervals are easily computed.

This normalization can be achieved on the basis of three observations.

First, if $\mathbf{X}$ is a random vector in $\mathbb{R}^k$ and $\mathbf{A}$ is constant and $k \times k$, then

$$\operatorname{Var}[\mathbf{AX}] = \mathbf{A} \operatorname{Var}[\mathbf{X}] \mathbf{A}'$$

Second, by the continuous mapping theorem, if $\mathbf{Z}_n \xrightarrow{d} \mathbf{Z}$ in $\mathbb{R}^k$ and $\mathbf{A}$ is constant and $k \times k$, then

$$\mathbf{A}\mathbf{Z}_n \xrightarrow{d} \mathbf{A}\mathbf{Z}$$

Third, if $\mathbf{S}$ is a $k \times k$ symmetric positive definite matrix, then there exists a symmetric positive definite matrix $\mathbf{Q}$, called the inverse square root of $\mathbf{S}$, such that

$$\mathbf{Q}\mathbf{S}\mathbf{Q}' = \mathbf{I}$$

Here $\mathbf{I}$ is the $k \times k$ identity matrix.

Putting these things together, your first exercise is to show that if $\mathbf{Q}$ is the inverse square root of [?], then

$$\mathbf{Z}_n := \sqrt{n}\mathbf{Q}(\bar{\mathbf{X}}_n - \mu) \xrightarrow{d} \mathbf{Z} \sim N(\mathbf{0}, \mathbf{I})$$

Applying the continuous mapping theorem one more time tells us that

$$\|\mathbf{Z}_n\|^2 \xrightarrow{d} \|\mathbf{Z}\|^2$$

Given the distribution of $\mathbf{Z}$, we conclude that

$$n\|\mathbf{Q}(\bar{\mathbf{X}}_n - \mu)\|^2 \xrightarrow{d} \chi^2(k) \tag{10.10}$$

where $\chi^2(k)$ is the chi-squared distribution with k degrees of freedom.

(Recall that k is the dimension of $\mathbf{X}_i$, the underlying random vectors.)

Your second exercise is to illustrate the convergence in (10.10) with a simulation.

In doing so, let

$$\mathbf{X}_i := \begin{pmatrix} W_i \\ U_i + W_i \end{pmatrix}$$

where

- each W_i is an IID draw from the uniform distribution on $[-1, 1]$.
- each U_i is an IID draw from the uniform distribution on $[-2, 2]$.
- U_i and W_i are independent of each other.

Hint

1. `scipy.linalg.sqrtm(A)` computes the square root of `A`. You still need to invert it.
2. You should be able to work out Σ from the preceding information.

Solution

First we want to verify the claim that

$$\sqrt{n}\mathbf{Q}(\bar{\mathbf{X}}_n - \mu) \xrightarrow{d} N(\mathbf{0}, \mathbf{I})$$

This is straightforward given the facts presented in the exercise.

Let

$$\mathbf{Y}_n := \sqrt{n}(\bar{\mathbf{X}}_n - \mu) \quad \text{and} \quad \mathbf{Y} \sim N(\mathbf{0}, \Sigma)$$

By the multivariate CLT and the continuous mapping theorem, we have

$$\mathbf{QY}_n \xrightarrow{d} \mathbf{QY}$$

Since linear combinations of normal random variables are normal, the vector $\mathbf{QY}$ is also normal.

Its mean is clearly $\mathbf{0}$, and its variance-covariance matrix is

$$\text{Var}[\mathbf{QY}] = \mathbf{Q}\text{Var}[\mathbf{Y}]\mathbf{Q}' = \mathbf{Q}\Sigma\mathbf{Q}' = \mathbf{I}$$

In conclusion, $\mathbf{QY}_n \xrightarrow{d} \mathbf{QY} \sim N(\mathbf{0}, \mathbf{I})$, which is what we aimed to show.

Now we turn to the simulation exercise.

Our solution is as follows

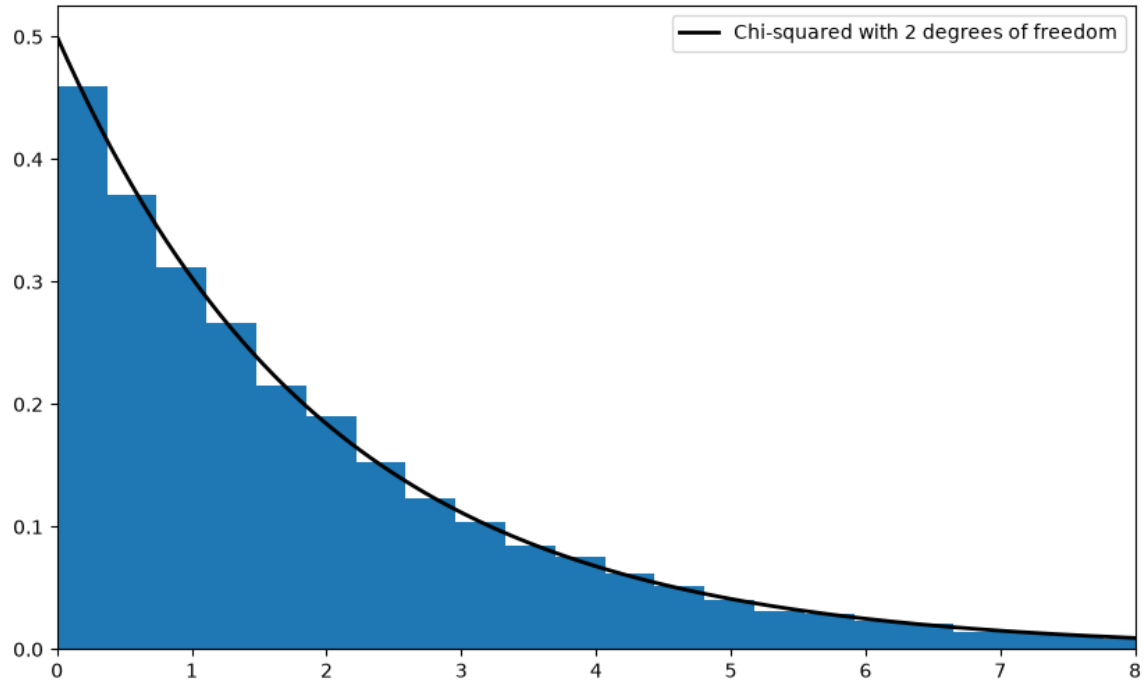
```
# Set parameters
n = 250
replications = 50000
dw = uniform(loc=-1, scale=2)  # Uniform(-1, 1)
du = uniform(loc=-2, scale=4)  # Uniform(-2, 2)
sw, su = dw.std(), du.std()
vw, vu = sw**2, su**2
Σ = ((vw, vw), (vw, vw + vu))
Σ = np.array(Σ)

# Compute Σ^{-1/2}
Q = inv(sqrtm(Σ))

# Generate observations of the normalized sample mean
error_obs = np.empty((2, replications))
for i in range(replications):
    # Generate one sequence of bivariate shocks
    X = np.empty((2, n))
    W = dw.rvs(n)
    U = du.rvs(n)
    # Construct the n observations of the random vector
    X[0, :] = W
    X[1, :] = W + U
    # Construct the i-th observation of Y_n
    error_obs[:, i] = np.sqrt(n) * X.mean(axis=1)

# Premultiply by Q and then take the squared norm
temp = Q @ error_obs
chisq_obs = np.sum(temp**2, axis=0)

# Plot
fig, ax = plt.subplots(figsize=(10, 6))
xmax = 8
ax.set_xlim(0, xmax)
xgrid = np.linspace(0, xmax, 200)
lb = "Chi-squared with 2 degrees of freedom"
ax.plot(xgrid, chi2.pdf(xgrid, 2), 'k-', lw=2, label=lb)
ax.legend()
ax.hist(chisq_obs, bins=50, density=True)
plt.show()
```


Chi-squared with 2 degrees of freedom
0.5
0.4
0.3
0.2
0.1
0.0
0
1
2
3
4
5
6
7
8

CHAPTER

ELEVEN

TWO MEANINGS OF PROBABILITY

11.1 Overview

This lecture illustrates two distinct interpretations of a **probability distribution**

- A frequentist interpretation as **relative frequencies** anticipated to occur in a large IID sample
- A Bayesian interpretation as a **personal opinion** (about a parameter or list of parameters) after seeing a collection of observations

We recommend watching the following two videos before proceeding:

- Hypothesis testing within the frequentist approach
- The Bayesian approach to constructing coverage intervals

After you are familiar with the material in these videos, this lecture uses the Socratic method to help consolidate your understanding of these two meanings of probability.

Along the way we construct a **Bayesian coverage interval** and contrast the question it answers with the frequentist relative-frequency reasoning that opens the lecture.

We do this by inviting you to write some Python code.

It would be especially useful if you tried doing this after each question that we pose for you, before proceeding to read the rest of the lecture.

We provide our own answers as the lecture unfolds, but you'll learn more if you try writing your own code before reading and running ours.

11.1.1 Code for answering questions

To answer our coding questions, we'll start with some imports

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from scipy.stats import binom
import scipy.stats as st
```

Empowered with these Python tools, we'll now explore the two meanings described above.

11.2 Frequentist Interpretation

Consider the following classic example.

The random variable X takes on possible values $k = 0, 1, 2, \ldots, n$ with probabilities

$$p(k \mid \theta) := \mathbb{P}\{X = k \mid \theta\} = \binom{n}{k} \theta^k (1-\theta)^{n-k}$$

where the fixed parameter $\theta \in (0, 1)$.

This is called the binomial distribution.

Here

- θ is the probability that one toss of a coin will be a head, an outcome that we encode as $Y = 1$.
- $1 - \theta$ is the probability that one toss of the coin will be a tail, an outcome that we denote $Y = 0$.
- X is the total number of heads that came up after flipping the coin n times.

Consider the following experiment:

Take I **independent** sequences of n **independent** flips of the coin

Notice the repeated use of the adjective **independent**:

- we use it once to describe that we are drawing n independent times from a **Bernoulli** distribution with parameter θ to arrive at one draw from a **Binomial** distribution with parameters θ, n.
- we use it again to describe that we are then drawing I sequences of n coin draws.

Let $y_h^i \in \{0, 1\}$ be the realized value of Y on the hth flip during the ith sequence of flips.

Let $\sum_{h=1}^n y_h^i$ denote the total number of times heads come up during the ith sequence of n independent coin flips.

Let f_k record the fraction of samples of length n for which $\sum_{h=1}^n y_h^i = k$:

$$f_k^I = \frac{1}{I} \sum_{i=1}^{I} \mathbb{1}\left\{ \sum_{h=1}^{n} y_h^i = k \right\}$$

The probability $p(k \mid \theta)$ answers the following question:

- As I becomes large, in what fraction of I independent draws of n coin flips should we anticipate k heads to occur?

As usual, a law of large numbers justifies this answer.

Exercise 11.2.1

1. Write Python code to compute f_k^I
2. Use your code to compute $f_k^I, k = 0, \ldots, n$ and compare them to $p(k \mid \theta)$ for various values of θ, n and I
3. With the Law of Large Numbers in mind, use your code to describe the relationship between f_k^I and $p(k \mid \theta)$ as I grows

Solution

Here is one solution.

We simulate the coin flips with one function and assemble the comparison table with another.

```
def simulate_head_counts(θ, n, I, seed=1234):
    "Simulate I sequences of n coin flips; return the heads count of each sequence.
 ↪"
    rng = np.random.default_rng(seed)
    Y = (rng.random((I, n)) <= θ).astype(int)
    return Y.sum(axis=1)
def compare_frequencies(θ, n, I, seed=1234):
    "Tabulate theoretical binomial probabilities against simulated frequencies."
    head_counts = simulate_head_counts(θ, n, I, seed)
    rows = [
        (k, binom.pmf(k, n, θ), np.mean(head_counts == k))
        for k in range(n + 1)
    ]
    return pd.DataFrame(
        rows, columns=['k', 'Theoretical', 'Frequentist']
    ).set_index('k')
θ, n, k, I = 0.7, 20, 10, 1_000_000

compare_frequencies(θ, n, I)
```

k	Theoretical	Frequentist
0	3.486784e-11	0.000000
1	1.627166e-09	0.000000
2	3.606885e-08	0.000000
3	5.049639e-07	0.000000
4	5.007558e-06	0.000002
5	3.738977e-05	0.000047
6	2.181070e-04	0.000247
7	1.017833e-03	0.001029
8	3.859282e-03	0.003721
9	1.200665e-02	0.011986
10	3.081708e-02	0.030698
11	6.536957e-02	0.065110
12	1.143967e-01	0.114892
13	1.642620e-01	0.164639
14	1.916390e-01	0.191724
15	1.788631e-01	0.178623
16	1.304210e-01	0.129870
17	7.160367e-02	0.072004
18	2.784587e-02	0.027818
19	6.839337e-03	0.006836
20	7.979227e-04	0.000754

From the table above, can you see the law of large numbers at work?

Let's do some more calculations.

11.2.1 Comparison with different θ

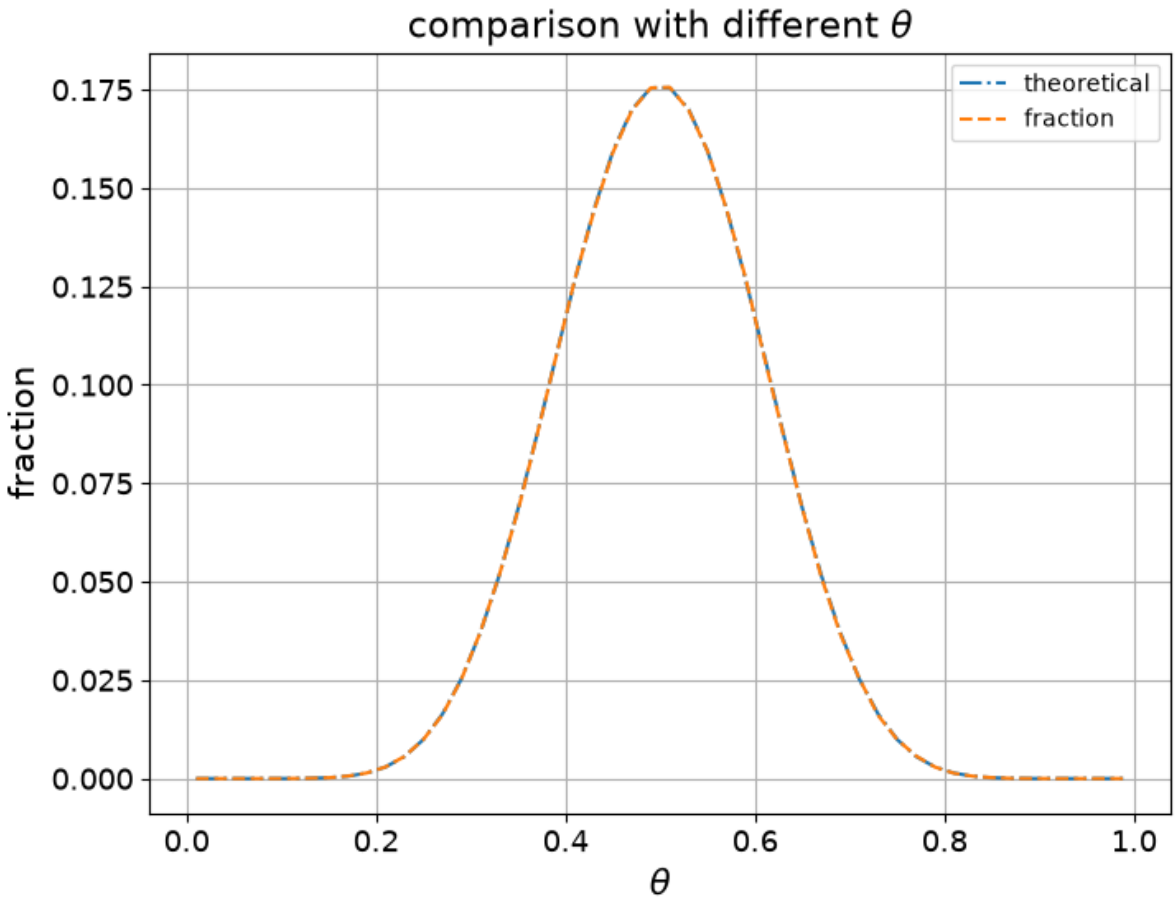
Now we fix

$$n = 20, k = 10, I = 1,000,000$$

We'll vary θ from 0.01 to 0.99 and plot outcomes against θ.

```
θ_low, θ_high, n_thetas = 0.01, 0.99, 50
thetas = np.linspace(θ_low, θ_high, n_thetas)
P = []
f_kI = []
for i in range(n_thetas):
    P.append(binom.pmf(k, n, thetas[i]))
    head_counts = simulate_head_counts(thetas[i], n, I, seed=i)
    f_kI.append(np.mean(head_counts == k))
```

```
fig, ax = plt.subplots(figsize=(8, 6))
ax.grid()
ax.plot(thetas, P, '-.', label='theoretical')
ax.plot(thetas, f_kI, '--', label='fraction')
ax.set_title(r'comparison with different $\theta$',
             fontsize=16)
ax.set_xlabel(r'$\theta$', fontsize=15)
ax.set_ylabel('fraction', fontsize=15)
ax.tick_params(labelsize=13)
ax.legend()
plt.show()
```



11.2.2 Comparison with different n

Now we fix $\theta = 0.7, k = 10, I = 1,000,000$ and vary n from 1 to 100.

Then we'll plot outcomes.

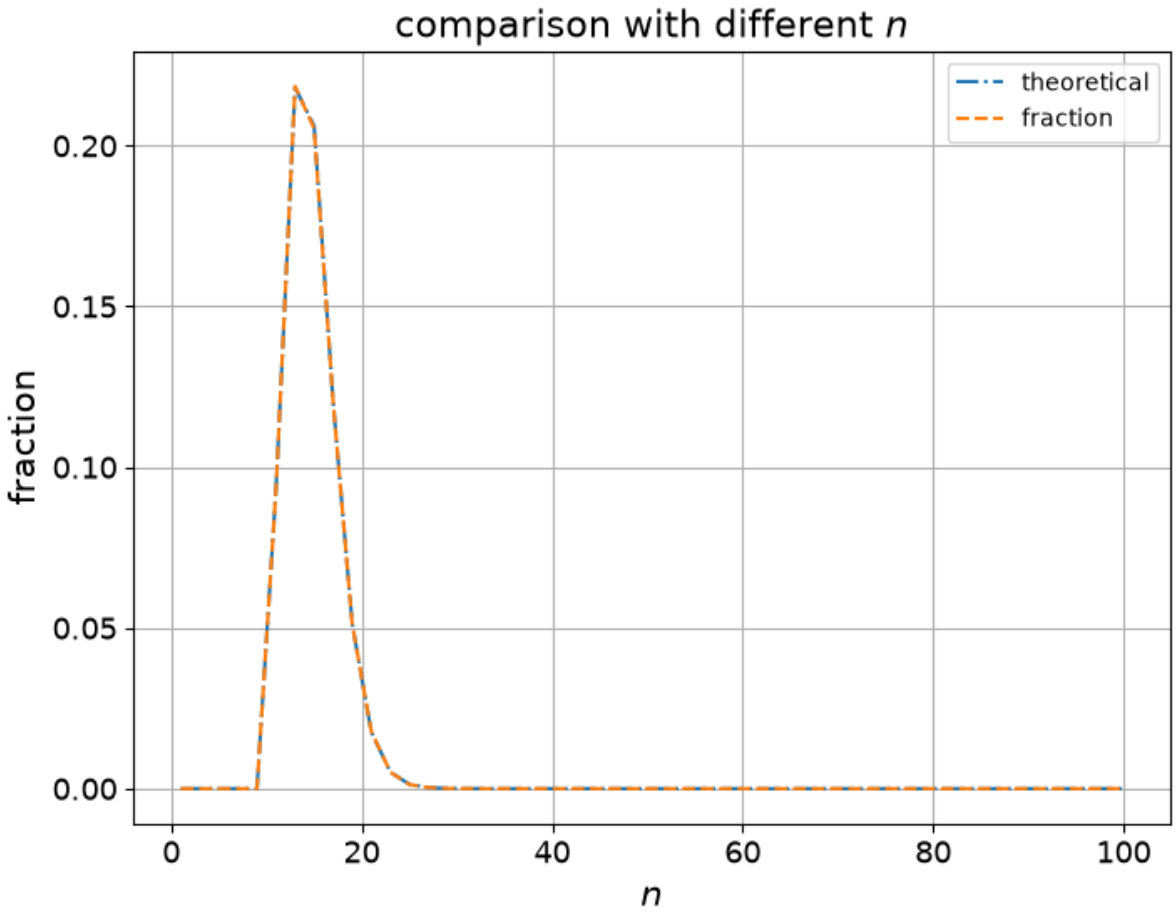
```
n_low, n_high, n_ns = 1, 100, 50
ns = np.linspace(n_low, n_high, n_ns, dtype='int')
P = []
f_kI = []
for i in range(n_ns):
    P.append(binom.pmf(k, ns[i], θ))
    head_counts = simulate_head_counts(θ, ns[i], I, seed=i)
    f_kI.append(np.mean(head_counts == k))
```

```
fig, ax = plt.subplots(figsize=(8, 6))
ax.grid()
ax.plot(ns, P, '-.', label='theoretical')
ax.plot(ns, f_kI, '--', label='fraction')
ax.set_title(r'comparison with different $n$',
             fontsize=16)
ax.set_xlabel(r'$n$', fontsize=15)
ax.set_ylabel('fraction', fontsize=15)
ax.tick_params(labelsize=13)
```

(continues on next page)

(continued from previous page)

```
ax.legend()
plt.show()
```



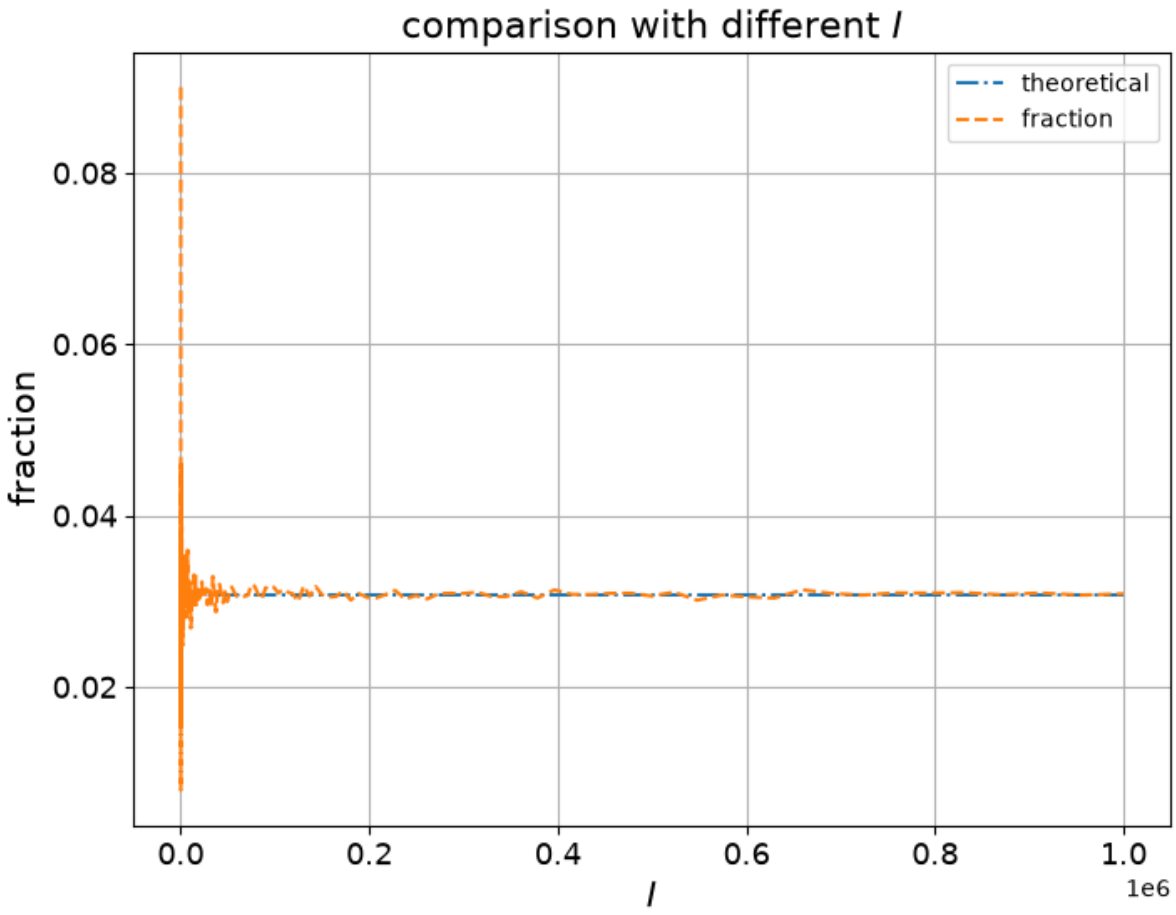
Because $k = 10$ is held fixed, $p(k \mid \theta)$ is zero whenever $n < 10$ — we cannot obtain 10 heads in fewer than 10 flips — and it is largest near $n \approx 14$, where the expected number of heads $n\theta$ is closest to k. The simulated fraction tracks this shape.

11.2.3 Comparison with different I

Now we fix $\theta = 0.7, n = 20, k = 10$ and vary $\log(I)$ from 2 to 6.

```
I_log_low, I_log_high, n_Is = 2, 6, 200
log_Is = np.linspace(I_log_low, I_log_high, n_Is)
Is = np.power(10, log_Is).astype(int)
P = []
f_kI = []
for i in range(n_Is):
    P.append(binom.pmf(k, n, θ))
    head_counts = simulate_head_counts(θ, n, Is[i], seed=i)
    f_kI.append(np.mean(head_counts == k))
```

```
fig, ax = plt.subplots(figsize=(8, 6))
ax.grid()
ax.plot(Is, P, '-.', label='theoretical')
ax.plot(Is, f_kI, '--', label='fraction')
ax.set_title(r'comparison with different $I$',
             fontsize=16)
ax.set_xlabel(r'$I$', fontsize=15)
ax.set_ylabel('fraction', fontsize=15)
ax.tick_params(labelsize=13)
ax.legend()
plt.show()
```



From the above graphs, we can see that I**, the number of independent sequences,** plays an important role.

When I becomes larger, the difference between theoretical probability and frequentist estimate becomes smaller.

Also, as long as I is large enough, changing θ or n does not substantially change the accuracy of the observed fraction as an approximation of $p(k \mid \theta)$.

The Law of Large Numbers is at work here.

For each independent sequence i, define the indicator $\rho_{k,i} = \mathbb{1}\{X_i = k\}$ — that is, $\rho_{k,i}$ equals 1 if the i-th sequence produces exactly k heads and 0 otherwise.

The $\rho_{k,i}$ are IID across i, each with mean $p(k \mid \theta)$ and variance

$$p(k \mid \theta) \cdot (1 - p(k \mid \theta)).$$

So, by the LLN, the average of $\rho_{k,i}$ converges to:

$$\mathbb{E}[\rho_{k,i}] = p(k \mid \theta) = \binom{n}{k} \theta^k (1-\theta)^{n-k}$$

as I goes to infinity.

11.3 Bayesian Interpretation

The likelihood remains binomial, but now we treat θ as a **random variable** rather than a fixed parameter.

So θ is described by a probability distribution.

But now this probability distribution means something different than a relative frequency that we can anticipate to occur in a large IID sample.

Instead, the probability distribution of θ is now a summary of our views about likely values of θ either

- **before** we have seen **any** data at all, or
- **before** we have seen **more** data, after we have seen **some** data

11.3.1 The beta prior and its posterior

Thus, suppose that, before seeing any data, you have a personal prior probability distribution with density

$$p(\theta) = \frac{\theta^{\alpha-1}(1-\theta)^{\beta-1}}{B(\alpha, \beta)}$$

where $B(\alpha, \beta)$ is a **beta function** , so that $p(\theta)$ is the density of a **beta distribution** with parameters α, β.

We can update this prior after observing data using Bayes' Law (see *Probability with Matrices* for an introduction).

For a sample of n coin flips that yields k heads, the **likelihood function** is the binomial PMF $p(k \mid \theta)$ introduced above.

Applying Bayes' Law with our beta prior, the **posterior density** is

$$p(\theta \mid k) = \frac{p(k \mid \theta) \cdot p(\theta)}{\int_0^1 p(k \mid \theta) \cdot p(\theta)\, d\theta}$$

Because the beta prior is conjugate to the binomial likelihood, this integral evaluates to (the kernel of) another beta density, so that

$$\theta \mid k \sim \text{Beta}(\alpha + k,\ \beta + n - k) \tag{11.1}$$

The first exercise below asks you to derive this closed form.

Exercise 11.3.1

a) Write down the **likelihood function** for a single coin flip with outcome $Y \in \{0, 1\}$.

b) Write down the **posterior** distribution for θ after observing that single flip.

c) Derive the closed-form posterior (11.1) for a sample of n flips that yields k heads.

Solution

a) The **likelihood function** for a single coin flip with outcome $Y \in \{0, 1\}$ is

$$p(Y \mid \theta) = \theta^Y (1-\theta)^{1-Y}$$

b) By Bayes' Law, the posterior density for θ after observing a single flip Y is

$$p(\theta \mid Y) = \frac{p(Y \mid \theta) \cdot p(\theta)}{\int_0^1 p(Y \mid \theta) \cdot p(\theta)\, d\theta}$$

Substituting the likelihood from (a) and the beta prior density, this becomes

$$p(\theta \mid Y) = \frac{\theta^Y (1-\theta)^{1-Y} \cdot \theta^{\alpha-1}(1-\theta)^{\beta-1}/B(\alpha,\beta)}{\int_0^1 \theta^Y (1-\theta)^{1-Y} \cdot \theta^{\alpha-1}(1-\theta)^{\beta-1}/B(\alpha,\beta)\, d\theta}$$

Collecting powers of θ and $(1-\theta)$, we recognize the kernel of a beta density:

$$p(\theta \mid Y) = \frac{\theta^{Y+\alpha-1}(1-\theta)^{1-Y+\beta-1}}{\int_0^1 \theta^{Y+\alpha-1}(1-\theta)^{1-Y+\beta-1}\, d\theta}$$

which means that

$$\theta \mid Y \sim \text{Beta}(\alpha + Y,\ \beta + (1-Y))$$

c) The same calculation, with the binomial likelihood in place of the Bernoulli likelihood, generalizes the result to a sample of n flips that yields k heads.

The beta prior contributes the factor $\theta^{\alpha-1}(1-\theta)^{\beta-1}$, and the binomial likelihood contributes $\theta^k(1-\theta)^{n-k}$, so the posterior is proportional to

$$\theta^{\alpha+k-1}(1-\theta)^{\beta+n-k-1},$$

which is the kernel of a beta density. Hence

$$\theta \mid k \sim \text{Beta}(\alpha + k,\ \beta + n - k),$$

as stated in (11.1).

11.3.2 Exploring the posterior numerically

The next exercise puts this posterior to work.

Exercise 11.3.2

a) Now pretend that the true value of $\theta = 0.4$ and that someone who doesn't know this has a beta prior distribution with parameters $\beta = \alpha = 0.5$. Write Python code to simulate this person's personal posterior distribution for θ for a *single* sequence of n draws.

b) Plot the posterior distribution for θ as a function of θ as n grows as $1, 2, \ldots$.

c) For various n's, describe and compute a Bayesian coverage interval for the interval $[0.45, 0.55]$.

d) Tell what question a Bayesian coverage interval answers.

e) Compute the posterior probability that $\theta \in [0.45, 0.55]$ for various values of sample size n.

f) Use your Python code to study what happens to the posterior distribution as $n \to +\infty$, again assuming that the true value of $\theta = 0.4$, though it is unknown to the person doing the updating via Bayes' Law.

Solution

a)

We use one function to simulate a sequence of coin flips and another to form the Beta posterior from the first `n_obs` of those flips.

```
def simulate_flips(θ=0.4, n=1_000_000, seed=1234):
    "Simulate n coin flips, each landing heads (1) with probability θ."
    rng = np.random.default_rng(seed)
    return (rng.random(n) < θ).astype(int)
```

```
def form_posterior(draws, n_obs, α=0.5, β=0.5):
    "Beta posterior for θ from the first n_obs flips, given a Beta(α, β) prior."
    heads = draws[:n_obs].sum()
    return st.beta(α + heads, β + n_obs - heads)
```

b)

```
draws = simulate_flips()

n_obs_list = [1, 2, 3, 4, 5, 10, 20, 50,
              100, 1000,
              5000, 10_000, 50_000, 100_000,
              200_000, 300_000]

posterior_list = [form_posterior(draws, n_obs) for n_obs in n_obs_list]

θ_values = np.linspace(0.01, 1, 1000)

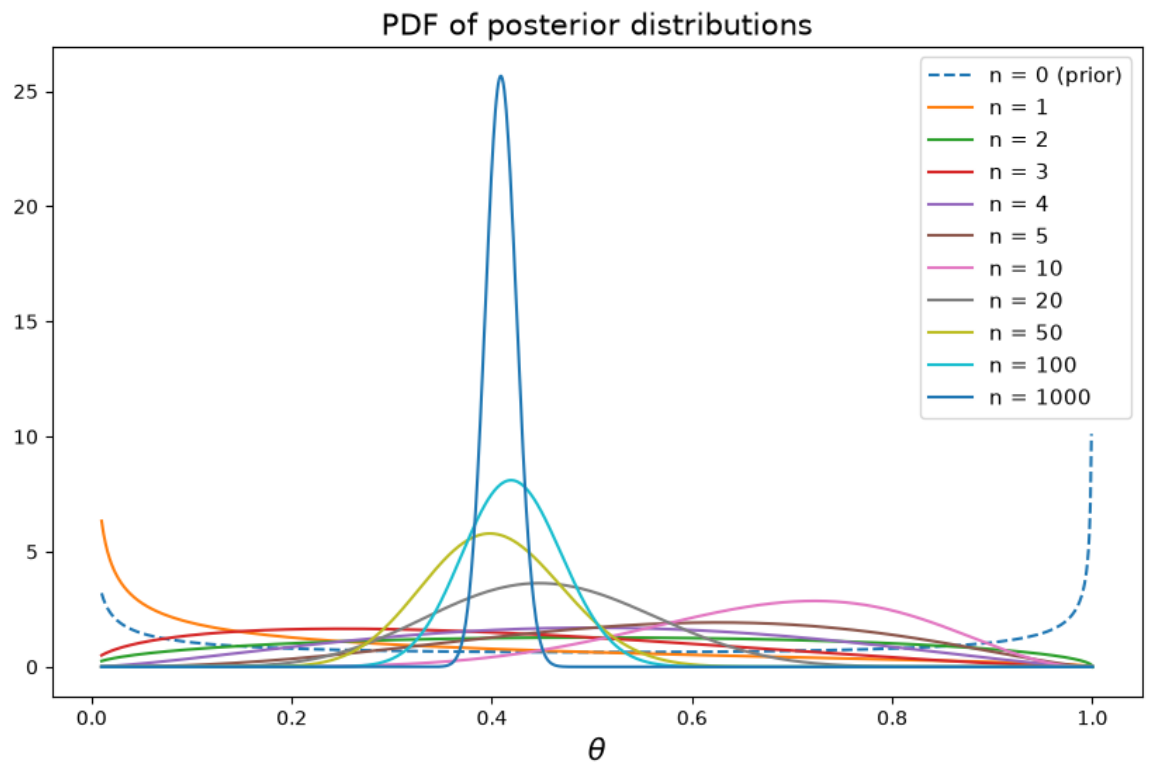
fig, ax = plt.subplots(figsize=(10, 6))

prior = st.beta(0.5, 0.5)
ax.plot(θ_values, prior.pdf(θ_values),
        label='n = 0 (prior)', linestyle='--')

for i, n_obs in enumerate(n_obs_list[:10]):
    posterior = posterior_list[i]
    ax.plot(θ_values, posterior.pdf(θ_values),
            label=f'n = {n_obs}')

ax.set_title('PDF of posterior distributions',
             fontsize=15)
ax.set_xlabel(r"$\theta$", fontsize=15)

ax.legend(fontsize=11)
plt.show()
```



c)

```
lower_bound = [post.ppf(0.05) for post in posterior_list[:10]]
upper_bound = [post.ppf(0.95) for post in posterior_list[:10]]

interval_df = pd.DataFrame()
interval_df['upper'] = upper_bound
interval_df['lower'] = lower_bound
interval_df.index = n_obs_list[:10]
interval_df = interval_df.T
interval_df
```

```
            1         2         3        4         5         10        20    \
upper  0.771480  0.902692  0.764466  0.83472  0.872224  0.882671  0.629953
lower  0.001543  0.097308  0.062413  0.16528  0.260634  0.441873  0.280091

            50        100       1000
upper  0.516104  0.502138  0.434744
lower  0.292234  0.341239  0.383644
```

As n increases, we can see that Bayesian coverage intervals narrow and move toward 0.4.

d) The Bayesian coverage interval tells the range of θ that corresponds to the $[q_1, q_2]$ quantiles of the cumulative distribution function (CDF) of the posterior distribution.

To construct the coverage interval we first compute a posterior distribution of the unknown parameter θ.

If the CDF is $F(\theta)$, then the Bayesian coverage interval $[a, b]$ for the interval $[q_1, q_2]$ is described by

$$F(a) = q_1, F(b) = q_2$$

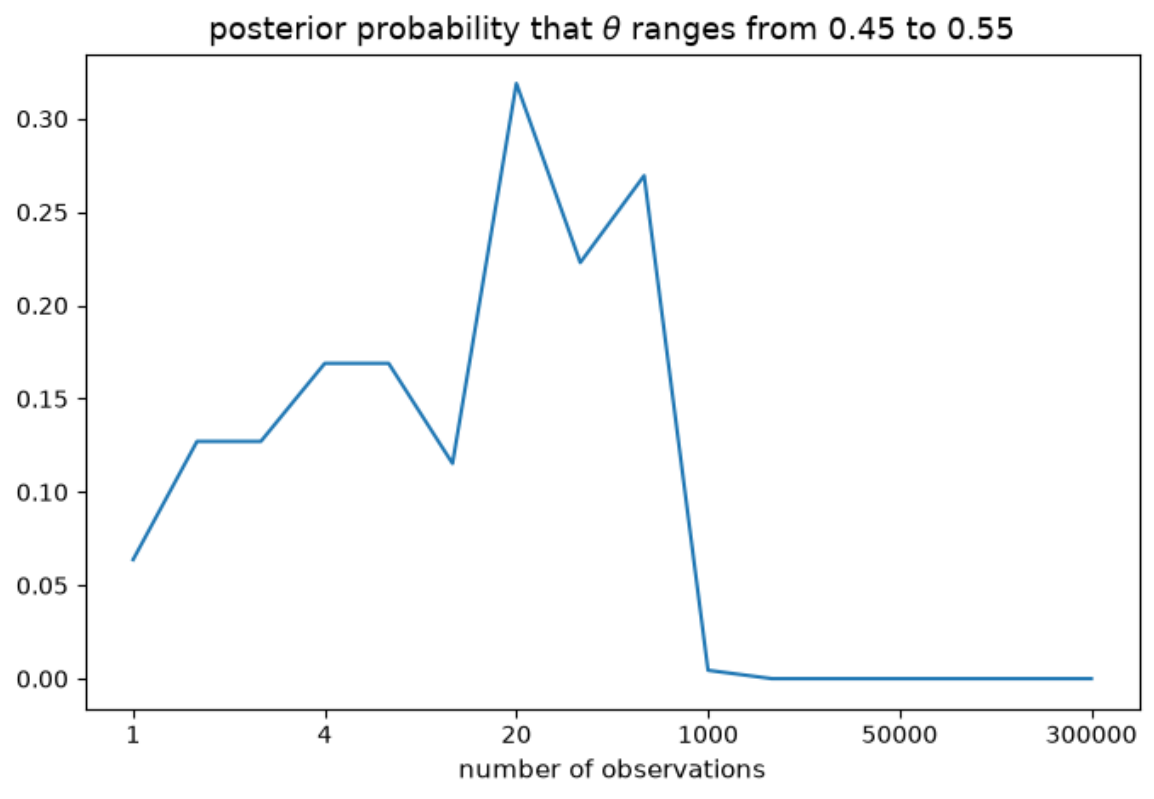
e)

```
left_value, right_value = 0.45, 0.55

posterior_prob_list = [
    post.cdf(right_value) - post.cdf(left_value)
    for post in posterior_list
]

fig, ax = plt.subplots(figsize=(8, 5))
ax.plot(posterior_prob_list)
ax.set_title(
    r'posterior probability that $\theta$'
    f' ranges from {left_value:.2f}'
    f' to {right_value:.2f}',
    fontsize=13)
ax.set_xticks(np.arange(0, len(posterior_prob_list), 3))
ax.set_xticklabels(n_obs_list[::3])
ax.set_xlabel('number of observations', fontsize=11)

plt.show()
```



Notice that in the graph above the posterior probability that $\theta \in [0.45, 0.55]$ exhibits a hump shape as n increases.

Two opposing forces are at work.

The first force is that the individual adjusts his belief as he observes new outcomes, so his posterior probability distribution becomes more and more realistic, which explains the rise of the posterior probability.

However, $[0.45, 0.55]$ actually excludes the true $\theta = 0.4$ that generates the data.

As a result, the posterior probability drops as larger and larger samples refine his posterior probability distribution of θ.

The descent seems precipitous only because of the scale of the graph that has the number of observations increasing disproportionately.

When the number of observations becomes large enough, our Bayesian becomes so confident about θ that he considers $\theta \in [0.45, 0.55]$ very unlikely.

That is why we see a nearly horizontal line when the number of observations exceeds 1000.

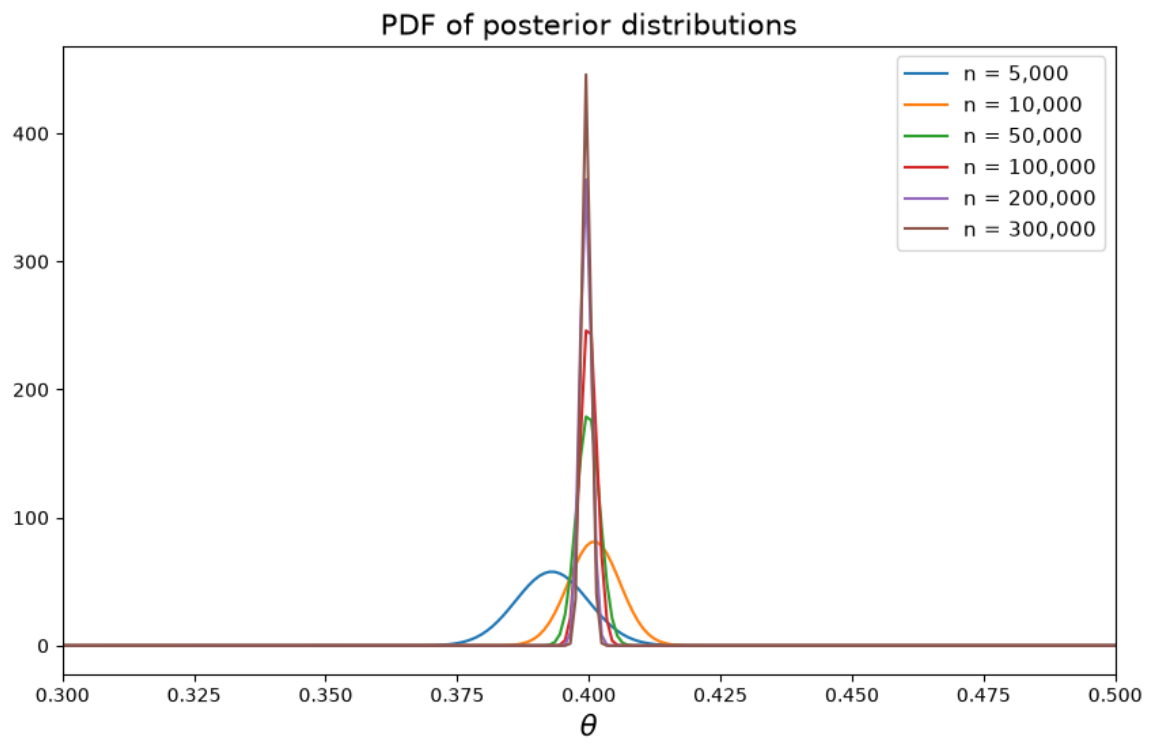
f) Using the functions we wrote above, we can see the evolution of posterior distributions as n approaches infinity.

```
fig, ax = plt.subplots(figsize=(10, 6))

for i, n_obs in enumerate(n_obs_list[10:]):
    posterior = posterior_list[i + 10]
    ax.plot(θ_values, posterior.pdf(θ_values),
            label=f'n = {n_obs:,}')

ax.set_title('PDF of posterior distributions', fontsize=15)
ax.set_xlabel(r"$\theta$", fontsize=15)
ax.set_xlim(0.3, 0.5)

ax.legend(fontsize=11)
plt.show()
```



As n increases, we can see that the probability density functions *concentrate* on 0.4, the true value of θ.

The next section explains *why* this concentration occurs and how fast it happens.

11.3.3 Why the posterior concentrates

In the solution to Exercise 11.3.2 we watched the posterior distribution concentrate ever more tightly around the true value $\theta = 0.4$ as the sample grew. Why does this happen?

The answer is encoded in the posterior we derived.

Recall that after observing k heads in n flips, the posterior is $\text{Beta}(\alpha + k,\ \beta + n - k)$.

A beta distribution with parameters a and b has

- mean $\dfrac{a}{a+b}$,
- variance $\dfrac{a\,b}{(a+b)^2\,(a+b+1)}$.

Substituting the *posterior* parameters $a = \alpha + k$ and $b = \beta + n - k$, so that $a + b = \alpha + \beta + n$, gives

$$\mathbb{E}[\theta \mid k] = \frac{\alpha + k}{\alpha + \beta + n}, \qquad \text{Var}[\theta \mid k] = \frac{(\alpha + k)(\beta + n - k)}{(\alpha + \beta + n)^2\,(\alpha + \beta + n + 1)}.$$

As n grows, the fixed prior counts α and β become negligible beside the data.

Since the data are generated with $\theta = 0.4$, the Law of Large Numbers gives $k/n \to 0.4$ (see Exercise 11.2.1), so the posterior mean

$$\frac{\alpha + k}{\alpha + \beta + n} \;\approx\; \frac{k}{n} \;\to\; 0.4.$$

In the variance, the numerator grows like n^2 while the denominator grows like n^3, so

$$\text{Var}[\theta \mid k] \;\approx\; \frac{\theta(1-\theta)}{n} \;\longrightarrow\; 0.$$

The posterior mean therefore homes in on the truth while its spread vanishes at rate $1/n$.

The next figure confirms both claims: the posterior mean settles on 0.4 and the standard deviation decays toward zero.

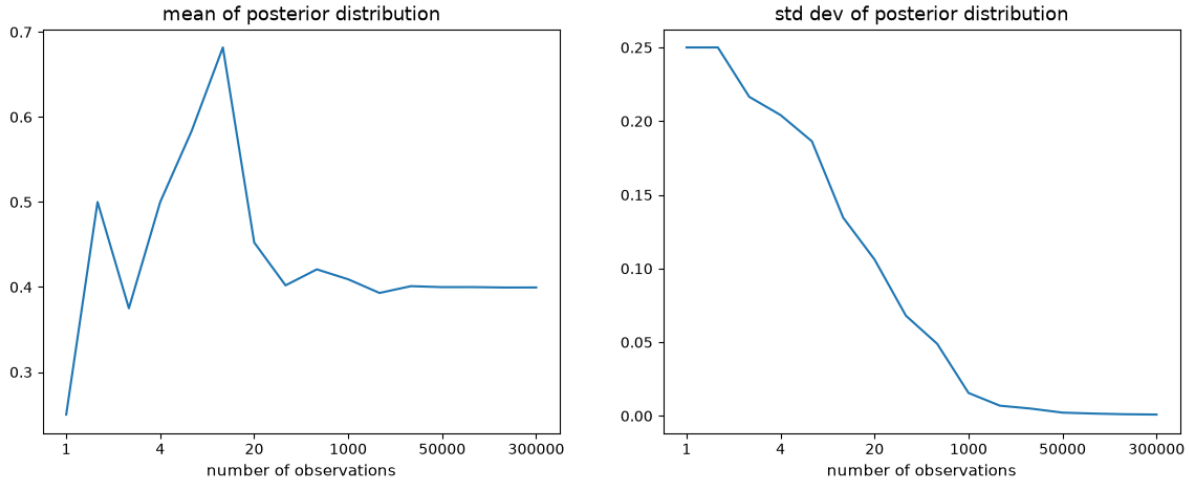
```
mean_list = [post.mean() for post in posterior_list]
std_list = [post.std() for post in posterior_list]

fig, ax = plt.subplots(1, 2, figsize=(14, 5))

ax[0].plot(mean_list)
ax[0].set_title('mean of posterior distribution',
                fontsize=13)
ax[0].set_xticks(np.arange(0, len(mean_list), 3))
ax[0].set_xticklabels(n_obs_list[::3])
ax[0].set_xlabel('number of observations', fontsize=11)

ax[1].plot(std_list)
ax[1].set_title('std dev of posterior distribution',
                fontsize=13)
ax[1].set_xticks(np.arange(0, len(std_list), 3))
ax[1].set_xticklabels(n_obs_list[::3])
ax[1].set_xlabel('number of observations', fontsize=11)

plt.show()
```

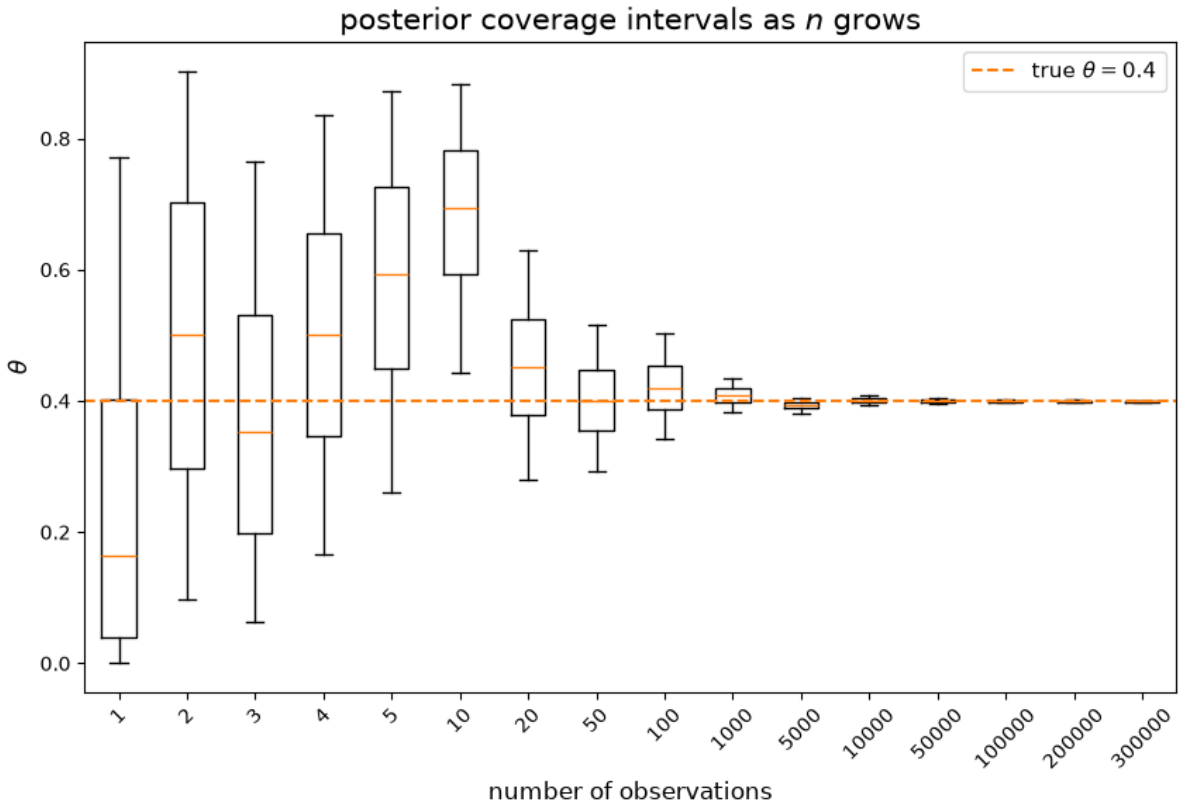


We can also display the Bayesian coverage intervals directly.

The box-and-whisker plot below summarizes each posterior by its median (central line), interquartile range (box), and 5th–95th percentile range (whiskers), with the true value $\theta = 0.4$ marked.

```
quantiles = [0.05, 0.25, 0.5, 0.75, 0.95]
box_stats = []
for post in posterior_list:
    lo, q1, med, q3, hi = post.ppf(quantiles)
    box_stats.append({'med': med, 'q1': q1, 'q3': q3,
                      'whislo': lo, 'whishi': hi, 'fliers': []})

fig, ax = plt.subplots(figsize=(10, 6))
ax.bxp(box_stats, positions=np.arange(len(box_stats)), showfliers=False)
ax.axhline(0.4, color='C1', linestyle='--', label=r'true $\theta = 0.4$')
ax.set_xticks(np.arange(len(box_stats)))
ax.set_xticklabels(n_obs_list, rotation=45)
ax.set_xlabel('number of observations', fontsize=12)
ax.set_ylabel(r'$\theta$', fontsize=12)
ax.set_title('posterior coverage intervals as $n$ grows', fontsize=15)
ax.legend(fontsize=11)
plt.show()
```



After observing a large number of outcomes, the posterior distribution collapses around 0.4.

Thus, the Bayesian statistician comes to believe that θ is near 0.4.

As shown in the figure above, as the number of observations grows, the Bayesian coverage intervals (BCIs) become narrower and narrower around 0.4.

However, if you take a closer look, you will find that the centers of the BCIs are not exactly 0.4, due to the persistent influence of the prior distribution and the randomness of the simulation path.

11.4 Comparing the two interpretations

We can now place the two meanings of probability side by side.

In the frequentist calculations, θ was a *fixed* number and a probability was a *relative frequency*: $p(k \mid \theta)$ is the fraction of a large number of independent length-n sequences in which exactly k heads occur, and we confirmed by simulation that $f_k^I \to p(k \mid \theta)$ as $I \to \infty$.

In the Bayesian calculations, θ was itself a *random variable* describing our beliefs, and a probability was a *degree of belief*. The Bayesian coverage interval summarizes the posterior $p(\theta \mid k)$: it reports the range of θ lying between two posterior quantiles, and so answers the question "given the data I have seen and my prior, where do I now believe θ lies?"

These are genuinely different questions, even though both rest on the same binomial likelihood. The frequentist statement describes the data-generating mechanism for a fixed θ; the Bayesian statement describes our uncertainty about θ given the particular data we happened to observe.

11.5 Role of a Conjugate Prior

We have made assumptions that link functional forms of our likelihood function and our prior in a way that has eased our calculations considerably.

In particular, our assumptions that the likelihood function is **binomial** and that the prior distribution is a **beta distribution** have the consequence that the posterior distribution implied by Bayes' Law is also a **beta distribution**.

So posterior and prior are both beta distributions, albeit ones with different parameters.

When a likelihood function and prior fit together like hand and glove in this way, we can say that the prior and posterior are **conjugate distributions**.

In this situation, we also sometimes say that we have **conjugate prior** for the likelihood function $p(k \mid \theta)$.

Typically, the functional form of the likelihood function determines the functional form of a **conjugate prior**.

A natural question to ask is why should a person's personal prior about a parameter θ be restricted to be described by a conjugate prior?

Why not some other functional form that more sincerely describes the person's beliefs?

To be argumentative, one could ask, why should the form of the likelihood function have *anything* to say about my personal beliefs about θ?

A dignified response to that question is, well, it shouldn't, but if you want to compute a posterior easily you'll just be happier if your prior is conjugate to your likelihood.

Otherwise, your posterior won't have a convenient analytical form and you'll be in the situation of wanting to apply the Markov chain Monte Carlo techniques deployed in *Non-Conjugate Priors*.

We also apply these powerful methods to approximating Bayesian posteriors for non-conjugate priors in *Posterior Distributions for AR(1) Parameters* and *Forecasting an AR(1) Process*.

CHAPTER

TWELVE

MULTIVARIATE HYPERGEOMETRIC DISTRIBUTION

Contents

- *Multivariate Hypergeometric Distribution*
  - *Overview*
  - *The administrator's problem*
  - *Usage*

12.1 Overview

This lecture describes how an administrator deployed a **multivariate hypergeometric distribution** in order to access the fairness of a procedure for awarding research grants.

In the lecture we'll learn about

- properties of the multivariate hypergeometric distribution
- first and second moments of a multivariate hypergeometric distribution
- using a Monte Carlo simulation of a multivariate normal distribution to evaluate the quality of a normal approximation
- the administrator's problem and why the multivariate hypergeometric distribution is the right tool

12.2 The administrator's problem

An administrator in charge of allocating research grants is in the following situation.

To help us forget details that are none of our business here and to protect the anonymity of the administrator and the subjects, we call research proposals **balls** and continents of residence of authors of a proposal a **color**.

There are K_i balls (proposals) of color i.

There are c distinct colors (continents of residence).

Thus, $i = 1, 2, \ldots, c$

So there is a total of $N = \sum_{i=1}^{c} K_i$ balls.

All N of these balls are placed in an urn.

Then n balls are drawn randomly.

The selection procedure is supposed to be **color blind** meaning that **ball quality**, a random variable that is supposed to be independent of **ball color**, governs whether a ball is drawn.

Thus, the selection procedure is supposed randomly to draw n balls from the urn.

The n balls drawn represent successful proposals and are awarded research funds.

The remaining $N - n$ balls receive no research funds.

12.2.1 Details of the awards procedure under study

Let k_i be the number of balls of color i that are drawn.

Things have to add up so $\sum_{i=1}^{c} k_i = n$.

Under the hypothesis that the selection process judges proposals on their quality and that quality is independent of continent of the author's continent of residence, the administrator views the outcome of the selection procedure as a random vector

$$X = \begin{bmatrix} k_1 \\ k_2 \\ \vdots \\ k_c \end{bmatrix}.$$

To evaluate whether the selection procedure is **color blind** the administrator wants to study whether the particular realization of X drawn can plausibly be said to be a random draw from the probability distribution that is implied by the **color blind** hypothesis.

The appropriate probability distribution is the one described here.

Let's now instantiate the administrator's problem, while continuing to use the colored balls metaphor.

The administrator has an urn with $N = 238$ balls.

157 balls are blue, 11 balls are green, 46 balls are yellow, and 24 balls are black.

So $(K_1, K_2, K_3, K_4) = (157, 11, 46, 24)$ and $c = 4$.

15 balls are drawn without replacement.

So $n = 15$.

The administrator wants to know the probability distribution of outcomes

$$X = \begin{bmatrix} k_1 \\ k_2 \\ \vdots \\ k_4 \end{bmatrix}.$$

In particular, he wants to know whether a particular outcome - in the form of a 4×1 vector of integers recording the numbers of blue, green, yellow, and black balls, respectively, - contains evidence against the hypothesis that the selection process is *fair*, which here means *color blind* and truly are random draws without replacement from the population of N balls.

The right tool for the administrator's job is the **multivariate hypergeometric distribution**.

12.2.2 Multivariate hypergeometric distribution

Let's start with some imports.

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.special import comb
from scipy.stats import normaltest
from numba import jit, prange
```

To recapitulate, we assume there are in total c types of objects in an urn.

If there are K_i type i object in the urn and we take n draws at random without replacement, then the numbers of type i objects in the sample $(k_1, k_2, \ldots, k_c)$ has the multivariate hypergeometric distribution.

Note again that $N = \sum_{i=1}^{c} K_i$ is the total number of objects in the urn and $n = \sum_{i=1}^{c} k_i$.

Notation

We use the following notation for **binomial coefficients**: $\binom{m}{q} = \frac{m!}{(m-q)!}$.

The multivariate hypergeometric distribution has the following properties:

Probability mass function:

$$\Pr\{X_i = k_i \ \forall i\} = \frac{\prod_{i=1}^{c} \binom{K_i}{k_i}}{\binom{N}{n}}$$

Mean:

$$\mathrm{E}(X_i) = n\frac{K_i}{N}$$

Variances and covariances:

$$\mathrm{Var}(X_i) = n\frac{N-n}{N-1}\ \frac{K_i}{N}\left(1-\frac{K_i}{N}\right)$$

$$\mathrm{Cov}(X_i, X_j) = -n\frac{N-n}{N-1}\ \frac{K_i}{N}\frac{K_j}{N}$$

To do our work for us, we'll write an `Urn` class.

```
class Urn:

    def __init__(self, K_arr):
        """
        Initialization given the number of each type i object in the urn.

        Parameters
        ----------
        K_arr: ndarray(int)
            number of each type i object.
        """

        self.K_arr = np.array(K_arr)
        self.N = np.sum(K_arr)
        self.c = len(K_arr)

    def pmf(self, k_arr):
        """
```

(continues on next page)

(continued from previous page)

```
        Probability mass function.

        Parameters
        ----------
        k_arr: ndarray(int)
            number of observed successes of each object.
        """

        K_arr, N = self.K_arr, self.N

        k_arr = np.atleast_2d(k_arr)
        n = np.sum(k_arr, 1)

        num = np.prod(comb(K_arr, k_arr), 1)
        denom = comb(N, n)

        pr = num / denom

        return pr

    def moments(self, n):
        """
        Compute the mean and variance-covariance matrix for
        multivariate hypergeometric distribution.

        Parameters
        ----------
        n: int
            number of draws.
        """

        K_arr, N = self.K_arr, self.N

        # mean
        μ = n * K_arr / N

        # variance-covariance matrix
        p = K_arr / N
        Σ = n * (N - n) / (N - 1) * (np.diag(p) - np.outer(p, p))

        return μ, Σ

    def simulate(self, n, size=1, seed=None):
        """
        Simulate a sample from multivariate hypergeometric
        distribution where at each draw we take n objects
        from the urn without replacement.

        Parameters
        ----------
        n: int
            number of objects for each draw.
        size: int(optional)
            sample size.
        seed: int(optional)
            random seed.
        """
```

(continues on next page)

(continued from previous page)

```
        K_arr = self.K_arr

        gen = np.random.Generator(np.random.PCG64(seed))
        sample = gen.multivariate_hypergeometric(K_arr, n, size=size)

        return sample
```

12.3 Usage

12.3.1 First example

Apply this to an example from wiki:

Suppose there are 5 black, 10 white, and 15 red marbles in an urn. If six marbles are chosen without replacement, the probability that exactly two of each color are chosen is

$$P(2 \text{ black}, 2 \text{ white}, 2 \text{ red}) = \frac{\binom{5}{2}\binom{10}{2}\binom{15}{2}}{\binom{30}{6}} = 0.079575596816976$$

```
# construct the urn
K_arr = [5, 10, 15]
urn = Urn(K_arr)
```

Now use the Urn Class method `pmf` to compute the probability of the outcome $X = \begin{bmatrix} 2 & 2 & 2 \end{bmatrix}$

```
k_arr = [2, 2, 2] # array of number of observed successes
urn.pmf(k_arr)
```

```
array([0.0795756])
```

We can use the code to compute probabilities of a list of possible outcomes by constructing a 2-dimensional array `k_arr` and `pmf` will return an array of probabilities for observing each case.

```
k_arr = [[2, 2, 2], [1, 3, 2]]
urn.pmf(k_arr)
```

```
array([0.0795756, 0.1061008])
```

Now let's compute the mean vector and variance-covariance matrix.

```
n = 6
μ, Σ = urn.moments(n)
```

```
μ
```

```
array([1., 2., 3.])
```

```
Σ
```

```
array([[ 0.68965517, -0.27586207, -0.4137931 ],
       [-0.27586207,  1.10344828, -0.82758621],
       [-0.4137931 , -0.82758621,  1.24137931]])
```

12.3.2 Back to the administrator's problem

Now let's turn to the grant administrator's problem.

Here the array of numbers of i objects in the urn is $(157, 11, 46, 24)$.

```
K_arr = [157, 11, 46, 24]
urn = Urn(K_arr)
```

Let's compute the probability of the outcome $(10, 1, 4, 0)$.

```
k_arr = [10, 1, 4, 0]
urn.pmf(k_arr)
```

```
array([0.01547738])
```

We can compute probabilities of three possible outcomes by constructing a 3-dimensional arrays `k_arr` and utilizing the method `pmf` of the `Urn` class.

```
k_arr = [[5, 5, 4 ,1], [10, 1, 2, 2], [13, 0, 2, 0]]
urn.pmf(k_arr)
```

```
array([6.21412534e-06, 2.70935969e-02, 1.61839976e-02])
```

Now let's compute the mean and variance-covariance matrix of X when $n = 6$.

```
n = 6 # number of draws
μ, Σ = urn.moments(n)
```

```
# mean
μ
```

```
array([3.95798319, 0.27731092, 1.15966387, 0.60504202])
```

```
# variance-covariance matrix
Σ
```

```
array([[ 1.31862604, -0.17907267, -0.74884935, -0.39070401],
       [-0.17907267,  0.25891399, -0.05246715, -0.02737417],
       [-0.74884935, -0.05246715,  0.91579029, -0.11447379],
       [-0.39070401, -0.02737417, -0.11447379,  0.53255196]])
```

We can simulate a large sample and verify that sample means and covariances closely approximate the population means and covariances.

```
size = 10_000_000
sample = urn.simulate(n, size=size)
```

```
# mean
np.mean(sample, 0)
```

```
array([3.9579607, 0.2775512, 1.1594838, 0.6050043])
```

```
# variance covariance matrix
np.cov(sample.T)
```

```
array([[ 1.31848053, -0.17917906, -0.74915369, -0.39014778],
       [-0.17917906,  0.25923836, -0.05247573, -0.02758357],
       [-0.74915369, -0.05247573,  0.91616741, -0.114538  ],
       [-0.39014778, -0.02758357, -0.114538  ,  0.53226935]])
```

Evidently, the sample means and covariances approximate their population counterparts well.

12.3.3 Quality of normal approximation

To judge the quality of a multivariate normal approximation to the multivariate hypergeometric distribution, we draw a large sample from a multivariate normal distribution with the mean vector and covariance matrix for the corresponding multivariate hypergeometric distribution and compare the simulated distribution with the population multivariate hypergeometric distribution.

```
sample_normal = np.random.multivariate_normal(μ, Σ, size=size)
```

```
def bivariate_normal(x, y, μ, Σ, i, j):

    μ_x, μ_y = μ[i], μ[j]
    σ_x, σ_y = np.sqrt(Σ[i, i]), np.sqrt(Σ[j, j])
    σ_xy = Σ[i, j]

    x_μ = x - μ_x
    y_μ = y - μ_y

    ρ = σ_xy / (σ_x * σ_y)
    z = x_μ**2 / σ_x**2 + y_μ**2 / σ_y**2 - 2 * ρ * x_μ * y_μ / (σ_x * σ_y)
    denom = 2 * np.pi * σ_x * σ_y * np.sqrt(1 - ρ**2)

    return np.exp(-z / (2 * (1 - ρ**2))) / denom
```

```
@jit
def count(vec1, vec2, n):
    size = sample.shape[0]

    count_mat = np.zeros((n+1, n+1))
    for i in prange(size):
        count_mat[vec1[i], vec2[i]] += 1

    return count_mat
```

```
c = urn.c
fig, axs = plt.subplots(c, c, figsize=(14, 14))

# grids for ploting the bivariate Gaussian
```

(continues on next page)

(continued from previous page)

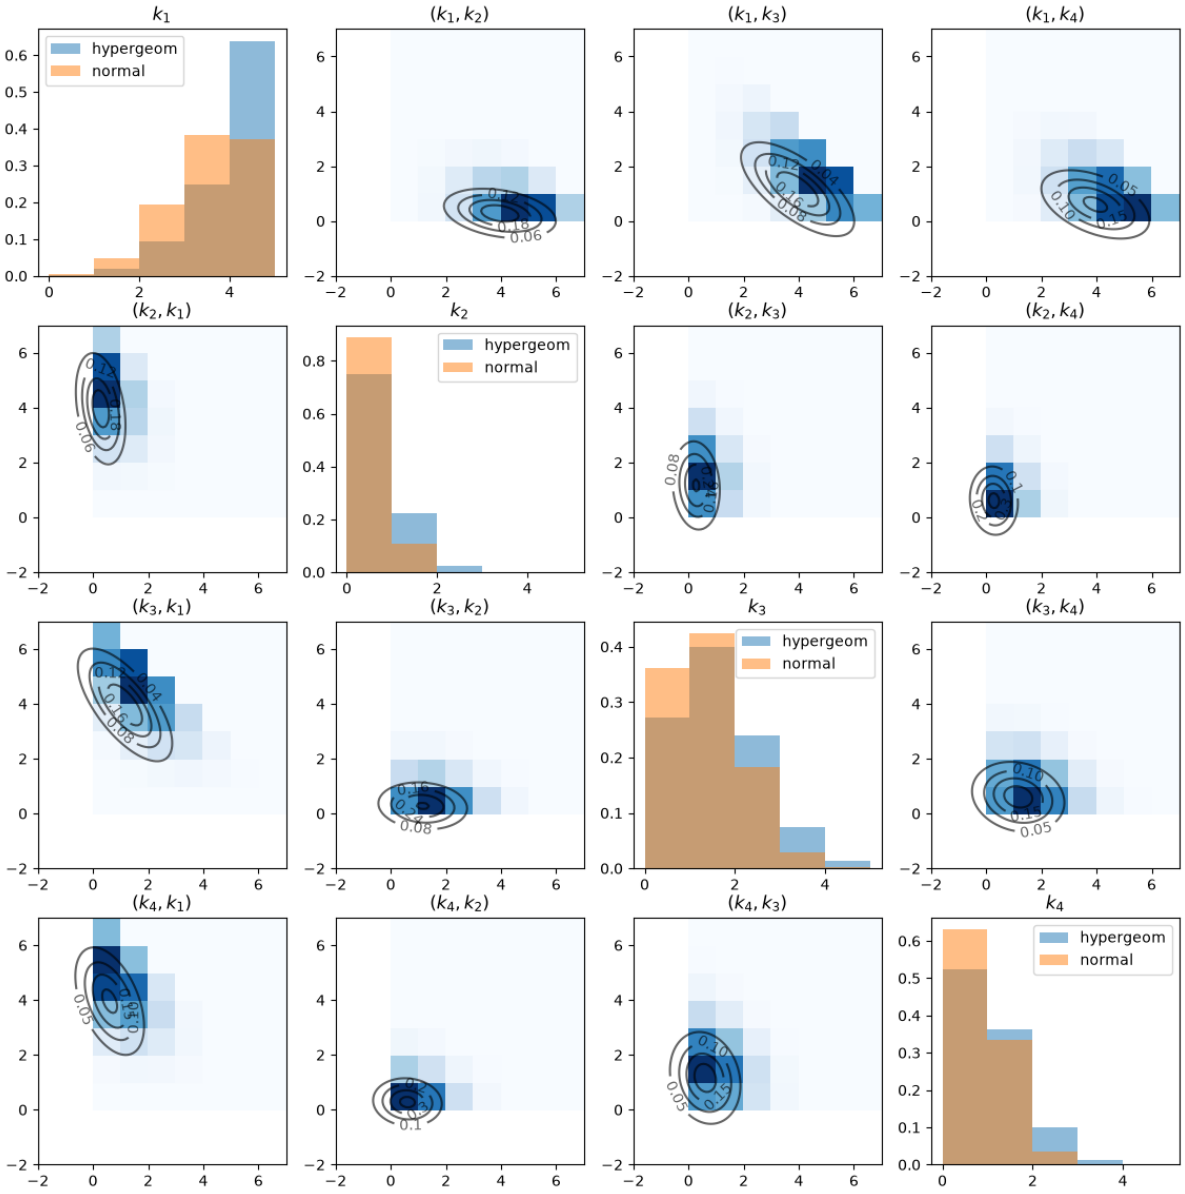
```
x_grid = np.linspace(-2, n+1, 100)
y_grid = np.linspace(-2, n+1, 100)
X, Y = np.meshgrid(x_grid, y_grid)

for i in range(c):
    axs[i, i].hist(sample[:, i], bins=np.arange(0, n, 1), alpha=0.5, density=True,␣
↪label='hypergeom')
    axs[i, i].hist(sample_normal[:, i], bins=np.arange(0, n, 1), alpha=0.5,␣
↪density=True, label='normal')
    axs[i, i].legend()
    axs[i, i].set_title('$k_{' +str(i+1) +'}$')
    for j in range(c):
        if i == j:
            continue

        # bivariate Gaussian density function
        Z = bivariate_normal(X, Y, μ, Σ, i, j)
        cs = axs[i, j].contour(X, Y, Z, 4, colors="black", alpha=0.6)
        axs[i, j].clabel(cs, inline=1, fontsize=10)

        # empirical multivariate hypergeometric distrbution
        count_mat = count(sample[:, i], sample[:, j], n)
        axs[i, j].pcolor(count_mat.T/size, cmap='Blues')
        axs[i, j].set_title('$(k_{' +str(i+1) +'}, k_{' + str(j+1) + '})$')

plt.show()
```



The diagonal graphs plot the marginal distributions of k_i for each i using histograms.

Note the substantial differences between hypergeometric distribution and the approximating normal distribution.

The off-diagonal graphs plot the empirical joint distribution of k_i and k_j for each pair (i, j).

The darker the blue, the more data points are contained in the corresponding cell. (Note that k_i is on the x-axis and k_j is on the y-axis).

The contour maps plot the bivariate Gaussian density function of $\left(k_i, k_j\right)$ with the population mean and covariance given by slices of μ and Σ that we computed above.

Let's also test the normality for each k_i using `scipy.stats.normaltest` that implements D'Agostino and Pearson's test that combines skew and kurtosis to form an omnibus test of normality.

The null hypothesis is that the sample follows normal distribution.

> `normaltest` returns an array of p-values associated with tests for each k_i sample.

```
test_multihyper = normaltest(sample)
test_multihyper.pvalue
```

```
array([0., 0., 0., 0.])
```

As we can see, all the p-values are almost 0 and the null hypothesis is soundly rejected.

By contrast, the sample from normal distribution does not reject the null hypothesis.

```
test_normal = normaltest(sample_normal)
test_normal.pvalue
```

```
array([0.51251396, 0.79865032, 0.68407419, 0.01887718])
```

The lesson to take away from this is that the normal approximation is imperfect.

CHAPTER

THIRTEEN

MULTIVARIATE NORMAL DISTRIBUTION

Contents

- *Multivariate Normal Distribution*
  - *Overview*
  - *The multivariate normal distribution*
  - *Bivariate example*
  - *Trivariate example*
  - *One dimensional intelligence (IQ)*
  - *Information as surprise*
  - *Cholesky factor magic*
  - *Math and verbal intelligence*
  - *Univariate time series analysis*
  - *Stochastic difference equation*
  - *Application to stock price model*
  - *Filtering foundations*
  - *Classic factor analysis model*
  - *PCA and factor analysis*
  - *Exercises*

13.1 Overview

This lecture describes a workhorse in probability theory, statistics, and economics, namely, the **multivariate normal distribution**.

In this lecture, you will learn formulas for

- the joint distribution of a random vector x of length N
- marginal distributions for all subvectors of x
- conditional distributions for subvectors of x conditional on other subvectors of x

We will use the multivariate normal distribution to formulate some useful models:

- a factor analytic model of an intelligence quotient, i.e., IQ
- a factor analytic model of two independent inherent abilities, say, mathematical and verbal.
- a more general factor analytic model
- Principal Components Analysis (PCA) as an approximation to a factor analytic model
- time series generated by linear stochastic difference equations
- optimal linear filtering theory

13.2 The multivariate normal distribution

This lecture defines a Python class `MultivariateNormal` to be used to generate **marginal** and **conditional** distributions associated with a multivariate normal distribution.

For a multivariate normal distribution it is very convenient that

- conditional expectations equal linear least squares projections
- conditional distributions are characterized by multivariate linear regressions

We apply our Python class to some examples.

We use the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
from numba import jit
import statsmodels.api as sm

rng = np.random.default_rng(0)
```

Assume that an $N \times 1$ random vector z has a multivariate normal probability density.

This means that the probability density takes the form

$$f\left(z;\mu,\Sigma\right)=\left(2\pi\right)^{-\left(\frac{N}{2}\right)}\det\left(\Sigma\right)^{-\frac{1}{2}}\exp\left(-.5\left(z-\mu\right)^{\top}\Sigma^{-1}\left(z-\mu\right)\right)$$

where $\mu=Ez$ is the mean of the random vector z and $\Sigma=E\left(z-\mu\right)\left(z-\mu\right)^{\top}$ is the covariance matrix of z.

The covariance matrix Σ is symmetric and positive definite.

```
@jit
def f(z, μ, Σ):
    """
    The density function of multivariate normal distribution.

    Parameters
    ---------------
    z: ndarray(float, dim=2)
        random vector, N by 1
    μ: ndarray(float, dim=1 or 2)
        the mean of z, N by 1
    Σ: ndarray(float, dim=2)
        the covariance matrix of z, N by N
    """
```

(continues on next page)

(continued from previous page)

```
    z = np.atleast_2d(z)
    μ = np.atleast_2d(μ)
    Σ = np.atleast_2d(Σ)

    N = z.size

    temp1 = np.linalg.det(Σ) ** (-1/2)
    temp2 = np.exp(-.5 * (z - μ).T @ np.linalg.inv(Σ) @ (z - μ))

    return (2 * np.pi) ** (-N/2) * temp1 * temp2
```

For some integer $k \in \{1, \ldots, N-1\}$, partition z as

$$z = \begin{bmatrix} z_1 \\ z_2 \end{bmatrix},$$

where z_1 is an $(N-k) \times 1$ vector and z_2 is a $k \times 1$ vector.

Let

$$\mu = \begin{bmatrix} \mu_1 \\ \mu_2 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{bmatrix}$$

be corresponding partitions of μ and Σ.

The **marginal** distribution of z_1 is

- multivariate normal with mean μ_1 and covariance matrix Σ_{11}.

The **marginal** distribution of z_2 is

- multivariate normal with mean μ_2 and covariance matrix Σ_{22}.

The distribution of z_1 **conditional** on z_2 is

- multivariate normal with mean

$$\hat{\mu}_1 = \mu_1 + \beta (z_2 - \mu_2)$$

and covariance matrix

$$\hat{\Sigma}_{11} = \Sigma_{11} - \Sigma_{12}\Sigma_{22}^{-1}\Sigma_{21} = \Sigma_{11} - \beta\Sigma_{22}\beta^\top$$

where

$$\beta = \Sigma_{12}\Sigma_{22}^{-1}$$

is an $(N-k) \times k$ matrix of **population regression coefficients** of the $(N-k) \times 1$ random vector $z_1 - \mu_1$ on the $k \times 1$ random vector $z_2 - \mu_2$.

The following class constructs a multivariate normal distribution instance with two methods.

- a method `partition` computes β, taking k as an input
- a method `cond_dist` computes either the distribution of z_1 conditional on z_2 or the distribution of z_2 conditional on z_1

```
class MultivariateNormal:
    """
    Class of multivariate normal distribution.

    Parameters
    ----------
    μ: ndarray(float, dim=1)
        the mean of z, N by 1
    Σ: ndarray(float, dim=2)
        the covariance matrix of z, N by N

    Arguments
    ---------
    μ, Σ:
        see parameters
    μs: list(ndarray(float, dim=1))
        list of mean vectors μ1 and μ2 in order
    Σs: list(list(ndarray(float, dim=2)))
        2 dimensional list of covariance matrices
        Σ11, Σ12, Σ21, Σ22 in order
    βs: list(ndarray(float, dim=1))
        list of regression coefficients β1 and β2 in order
    """

    def __init__(self, μ, Σ):
        "initialization"
        self.μ = np.array(μ)
        self.Σ = np.atleast_2d(Σ)

    def partition(self, k):
        """
        Given k, partition the random vector z into a size k vector z1
        and a size N-k vector z2. Partition the mean vector μ into
        μ1 and μ2, and the covariance matrix Σ into Σ11, Σ12, Σ21, Σ22
        correspondingly. Compute the regression coefficients β1 and β2
        using the partitioned arrays.
        """
        μ = self.μ
        Σ = self.Σ

        self.μs = [μ[:k], μ[k:]]
        self.Σs = [[Σ[:k, :k], Σ[:k, k:]],
                   [Σ[k:, :k], Σ[k:, k:]]]

        self.βs = [self.Σs[0][1] @ np.linalg.inv(self.Σs[1][1]),
                   self.Σs[1][0] @ np.linalg.inv(self.Σs[0][0])]

    def cond_dist(self, ind, z):
        """
        Compute the conditional distribution of z1 given z2, or reversely.
        Argument ind determines whether we compute the conditional
        distribution of z1 (ind=0) or z2 (ind=1).

        Returns
        ---------
        μ_hat: ndarray(float, ndim=1)
            The conditional mean of z1 or z2.
```

(continues on next page)

(continued from previous page)

```
        Σ_hat: ndarray(float, ndim=2)
            The conditional covariance matrix of z1 or z2.
        """
        β = self.βs[ind]
        μs = self.μs
        Σs = self.Σs

        μ_hat = μs[ind] + β @ (z - μs[1-ind])
        Σ_hat = Σs[ind][ind] - β @ Σs[1-ind][1-ind] @ β.T

        return μ_hat, Σ_hat
```

Let's put this code to work on a suite of examples.

We begin with a simple bivariate example; after that we'll turn to a trivariate example.

We'll compute population moments of some conditional distributions using our `MultivariateNormal` class.

For fun we'll also compute sample analogs of the associated population regressions by generating simulations and then computing linear least squares regressions.

We'll compare those linear least squares regressions for the simulated data to their population counterparts.

13.3 Bivariate example

We start with a bivariate normal distribution pinned down by

$$\mu = \begin{bmatrix} .5 \\ 1.0 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 1 & .5 \\ .5 & 1 \end{bmatrix}$$

```
μ = np.array([.5, 1.])
Σ = np.array([[1., .5], [.5 ,1.]])

# construction of the multivariate normal instance
multi_normal = MultivariateNormal(μ, Σ)
```

```
k = 1 # choose partition

# partition and compute regression coefficients
multi_normal.partition(k)
multi_normal.βs[0],multi_normal.βs[1]
```

```
(array([[0.5]]), array([[0.5]]))
```

Let's illustrate the fact that you *can regress anything on anything else*.

We have computed everything we need to compute two regression lines, one of z_2 on z_1, the other of z_1 on z_2.

We'll represent these regressions as

$$z_1 = a_1 + b_1 z_2 + \epsilon_1$$

and

$$z_2 = a_2 + b_2 z_1 + \epsilon_2$$

where we have the population least squares orthogonality conditions

$$E\epsilon_1 z_2 = 0$$

and

$$E\epsilon_2 z_1 = 0$$

Let's compute a_1, a_2, b_1, b_2.

```
β = multi_normal.βs

a1 = μ[0] - β[0]*μ[1]
b1 = β[0]

a2 = μ[1] - β[1]*μ[0]
b2 = β[1]
```

Let's print out the intercepts and slopes.

For the regression of z_1 on z_2 we have

```
print ("a1 = ", a1)
print ("b1 = ", b1)
```

```
a1 =  [[0.]]
b1 =  [[0.5]]
```

For the regression of z_2 on z_1 we have

```
print ("a2 = ", a2)
print ("b2 = ", b2)
```

```
a2 =  [[0.75]]
b2 =  [[0.5]]
```

Now let's plot the two regression lines and stare at them.

```
z2 = np.linspace(-4,4,100)


a1 = np.squeeze(a1)
b1 = np.squeeze(b1)

a2 = np.squeeze(a2)
b2 = np.squeeze(b2)

z1  = b1*z2 + a1


z1h = z2/b2 - a2/b2


fig = plt.figure(figsize=(12,12))
ax = fig.add_subplot(1, 1, 1)
ax.set(xlim=(-4, 4), ylim=(-4, 4))
ax.spines['left'].set_position('center')
```

(continues on next page)

(continued from previous page)

```
ax.spines['bottom'].set_position('zero')
ax.spines['right'].set_color('none')
ax.spines['top'].set_color('none')
ax.xaxis.set_ticks_position('bottom')
ax.yaxis.set_ticks_position('left')
plt.ylabel('$z_1$', loc = 'top')
plt.xlabel('$z_2$,', loc = 'right')
plt.plot(z2,z1, 'r', label = "$z_1$ on $z_2$")
plt.plot(z2,z1h, 'b', label = "$z_2$ on $z_1$")
plt.legend()
plt.show()
```

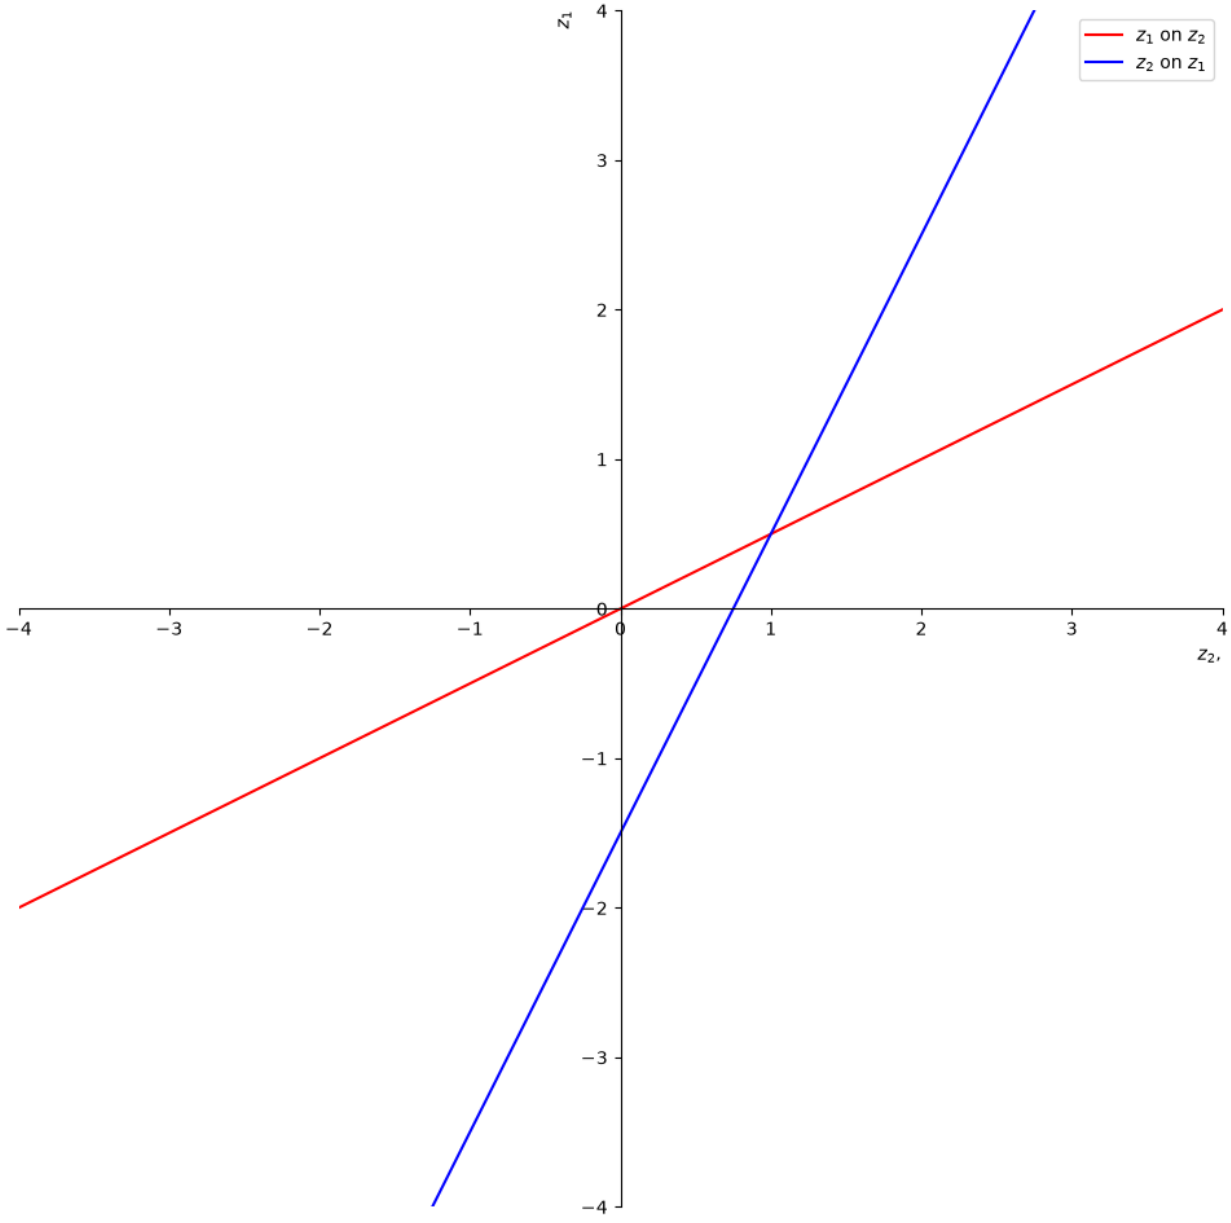


Fig. 13.1: two regressions

The red line is the expectation of z_1 conditional on z_2.

The intercept and slope of the red line are

```
print("a1 = ", a1)
print("b1 = ", b1)
```

```
a1 =  0.0
b1 =  0.5
```

The blue line is the expectation of z_2 conditional on z_1.

The intercept and slope of the blue line are

```
print("-a2/b2 = ", - a2/b2)
print("1/b2 = ", 1/b2)
```

```
-a2/b2 =  -1.5
1/b2 =  2.0
```

We can use these regression lines or our code to compute conditional expectations.

Let's compute the mean and variance of the distribution of z_2 conditional on $z_1 = 5$.

After that we'll reverse what are on the left and right sides of the regression.

```
# compute the cond. dist. of z1
ind = 1
z1 = np.array([5.]) # given z1

μ2_hat, Σ2_hat = multi_normal.cond_dist(ind, z1)
print('μ2_hat, Σ2_hat = ', μ2_hat, Σ2_hat)
```

```
μ2_hat, Σ2_hat =  [3.25] [[0.75]]
```

Now let's compute the mean and variance of the distribution of z_1 conditional on $z_2 = 5$.

```
# compute the cond. dist. of z1
ind = 0
z2 = np.array([5.]) # given z2

μ1_hat, Σ1_hat = multi_normal.cond_dist(ind, z2)
print('μ1_hat, Σ1_hat = ', μ1_hat, Σ1_hat)
```

```
μ1_hat, Σ1_hat =  [2.5] [[0.75]]
```

Let's compare the preceding population mean and variance with outcomes from drawing a large sample and then regressing $z_1 - \mu_1$ on $z_2 - \mu_2$.

We know that

$$Ez_1|z_2 = (\mu_1 - \beta\mu_2) + \beta z_2$$

which can be arranged to

$$z_1 - \mu_1 = \beta\left(z_2 - \mu_2\right) + \epsilon,$$

We anticipate that for larger and larger sample sizes, estimated OLS coefficients will converge to β and the estimated variance of ϵ will converge to $\hat{\Sigma}_1$.

```
n = 1_000_000 # sample size

# simulate multivariate normal random vectors
data = rng.multivariate_normal(μ, Σ, size=n)
z1_data = data[:, 0]
z2_data = data[:, 1]

# OLS regression
μ1, μ2 = multi_normal.μs
results = sm.OLS(z1_data - μ1, z2_data - μ2).fit()
```

Let's compare the preceding population β with the OLS sample estimate on $z_2 - \mu_2$

```
multi_normal.βs[0], results.params
```

```
(array([[0.5]]), array([0.49918333]))
```

Let's compare our population $\hat{\Sigma}_1$ with the degrees-of-freedom adjusted estimate of the variance of ϵ

```
Σ1_hat, results.resid @ results.resid.T / (n - 1)
```

```
(array([[0.75]]), np.float64(0.7489203195072748))
```

Lastly, let's compute the estimate of $\hat{Ez_1|z_2}$ and compare it with $\hat{\mu}_1$

```
μ1_hat, results.predict(z2 - μ2) + μ1
```

```
(array([2.5]), array([2.49673331]))
```

Thus, in each case, for our very large sample size, the sample analogues closely approximate their population counterparts.

A Law of Large Numbers explains why sample analogues approximate population objects.

13.4 Trivariate example

Let's apply our code to a trivariate example.

We'll specify the mean vector and the covariance matrix as follows.

```
μ = rng.random(3)
C = rng.random((3, 3))
Σ = C @ C.T # positive semi-definite

multi_normal = MultivariateNormal(μ, Σ)
```

```
μ, Σ
```

```
(array([0.4201905 , 0.04279279, 0.32149359]),
 array([[0.95070417, 0.48194854, 0.38025012],
        [0.48194854, 1.34029754, 0.43476327],
        [0.38025012, 0.43476327, 0.20864566]]))
```

```
k = 1
multi_normal.partition(k)
```

Let's compute the distribution of z_1 conditional on $z_2 = \begin{bmatrix} 2 \\ 5 \end{bmatrix}$.

```
ind = 0
z2 = np.array([2., 5.])

μ1_hat, Σ1_hat = multi_normal.cond_dist(ind, z2)
```

```
n = 1_000_000
data = rng.multivariate_normal(μ, Σ, size=n)
z1_data = data[:, :k]
z2_data = data[:, k:]
```

```
μ1, μ2 = multi_normal.μs
results = sm.OLS(z1_data - μ1, z2_data - μ2).fit()
```

As above, we compare population and sample regression coefficients, the conditional covariance matrix, and the conditional mean vector in that order.

```
multi_normal.βs[0], results.params
```

```
(array([[-0.71459379,  3.311496  ]]), array([-0.71383888,  3.31026328]))
```

```
Σ1_hat, results.resid @ results.resid.T / (n - 1)
```

```
(array([[0.03590485]]), np.float64(0.035947932906627136))
```

```
μ1_hat, results.predict(z2 - μ2) + μ1
```

```
(array([14.5144376]), array([14.51014783]))
```

Once again, sample analogues do a good job of approximating their populations counterparts.

13.5 One dimensional intelligence (IQ)

Let's move closer to a real-life example, namely, inferring a one-dimensional measure of intelligence called IQ from a list of test scores.

The ith test score y_i equals the sum of an unknown scalar IQ θ and a random variable w_i.

$$y_i = \theta + \sigma_y w_i, \quad i = 1, \ldots, n$$

The distribution of IQ's for a cross-section of people is a normal random variable described by

$$\theta = \mu_\theta + \sigma_\theta w_{n+1}.$$

We assume that the noises $\{w_i\}_{i=1}^N$ in the test scores are IID and not correlated with IQ.

We also assume that $\{w_i\}_{i=1}^{n+1}$ are i.i.d. standard normal:

$$w = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \\ w_{n+1} \end{bmatrix} \sim N\left(0, I_{n+1}\right)$$

The following system describes the $(n+1) \times 1$ random vector X that interests us:

$$X = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \\ \theta \end{bmatrix} = \begin{bmatrix} \mu_\theta \\ \mu_\theta \\ \vdots \\ \mu_\theta \\ \mu_\theta \end{bmatrix} + \begin{bmatrix} \sigma_y & 0 & \cdots & 0 & \sigma_\theta \\ 0 & \sigma_y & \cdots & 0 & \sigma_\theta \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \sigma_y & \sigma_\theta \\ 0 & 0 & \cdots & 0 & \sigma_\theta \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_n \\ w_{n+1} \end{bmatrix},$$

or equivalently,

$$X = \mu_\theta 1_{n+1} + Dw$$

where $X = \begin{bmatrix} y \\ \theta \end{bmatrix}$, 1_{n+1} is a vector of 1s of size $n+1$, and D is an $n+1$ by $n+1$ matrix.

Let's define a Python function that constructs the mean μ and covariance matrix Σ of the random vector X that we know is governed by a multivariate normal distribution.

As arguments, the function takes the number of tests n, the mean μ_θ and the standard deviation σ_θ of the IQ distribution, and the standard deviation of the randomness in test scores σ_y.

```
def construct_moments_IQ(n, μθ, σθ, σy):

    μ_IQ = np.full(n+1, μθ)

    D_IQ = np.zeros((n+1, n+1))
    D_IQ[range(n), range(n)] = σy
    D_IQ[:, n] = σθ

    Σ_IQ = D_IQ @ D_IQ.T

    return μ_IQ, Σ_IQ, D_IQ
```

Now let's consider a specific instance of this model.

Assume we have recorded 50 test scores and we know that $\mu_\theta = 100$, $\sigma_\theta = 10$, and $\sigma_y = 10$.

We can compute the mean vector and covariance matrix of X easily with our `construct_moments_IQ` function as follows.

```
n = 50
μθ, σθ, σy = 100., 10., 10.

μ_IQ, Σ_IQ, D_IQ = construct_moments_IQ(n, μθ, σθ, σy)
μ_IQ, Σ_IQ, D_IQ
```

```
(array([100., 100., 100., 100., 100., 100., 100., 100., 100., 100., 100.,
        100., 100., 100., 100., 100., 100., 100., 100., 100., 100., 100.,
        100., 100., 100., 100., 100., 100., 100., 100., 100., 100., 100.,
        100., 100., 100., 100., 100., 100., 100., 100., 100., 100., 100.,
```

(continues on next page)

(continued from previous page)

```
       100., 100., 100., 100., 100., 100., 100.]),
array([[200., 100., 100., ..., 100., 100., 100.],
       [100., 200., 100., ..., 100., 100., 100.],
       [100., 100., 200., ..., 100., 100., 100.],
       ...,
       [100., 100., 100., ..., 200., 100., 100.],
       [100., 100., 100., ..., 100., 200., 100.],
       [100., 100., 100., ..., 100., 100., 100.]], shape=(51, 51)),
array([[10.,  0.,  0., ...,  0.,  0., 10.],
       [ 0., 10.,  0., ...,  0.,  0., 10.],
       [ 0.,  0., 10., ...,  0.,  0., 10.],
       ...,
       [ 0.,  0.,  0., ..., 10.,  0., 10.],
       [ 0.,  0.,  0., ...,  0., 10., 10.],
       [ 0.,  0.,  0., ...,  0.,  0., 10.]], shape=(51, 51)))
```

We can now use our `MultivariateNormal` class to construct an instance, then partition the mean vector and covariance matrix as we wish.

We want to regress IQ, the random variable θ (*what we don't know*), on the vector y of test scores (*what we do know*).

We choose `k=n` so that $z_1 = y$ and $z_2 = \theta$.

```
multi_normal_IQ = MultivariateNormal(μ_IQ, Σ_IQ)

k = n
multi_normal_IQ.partition(k)
```

Using the generator `multivariate_normal`, we can make one draw of the random vector from our distribution and then compute the distribution of θ conditional on our test scores.

Let's do that and then print out some pertinent quantities.

```
x = rng.multivariate_normal(μ_IQ, Σ_IQ)
y = x[:-1] # test scores
θ = x[-1]  # IQ
```

```
# the true value
θ
```

```
np.float64(98.2495327983163)
```

The method `cond_dist` takes test scores y as input and returns the conditional normal distribution of the IQ θ.

In the following code, `ind` sets the variables on the right side of the regression.

Given the way we have defined the vector X, we want to set `ind=1` in order to make θ the left side variable in the population regression.

```
ind = 1
multi_normal_IQ.cond_dist(ind, y)
```

```
(array([97.20269774]), array([[1.96078431]]))
```

The first number is the conditional mean $\hat{\mu}_\theta$ and the second is the conditional variance $\hat{\Sigma}_\theta$.

How do additional test scores affect our inferences?

To shed light on this, we compute a sequence of conditional distributions of θ by varying the number of test scores in the conditioning set from 1 to n.

We'll make a pretty graph showing how our judgment of the person's IQ change as more test results come in.

```
# array for containing moments
μθ_hat_arr = np.empty(n)
Σθ_hat_arr = np.empty(n)

# loop over number of test scores
for i in range(1, n+1):
    # construction of multivariate normal distribution instance
    μ_IQ_i, Σ_IQ_i, D_IQ_i = construct_moments_IQ(i, μθ, σθ, σy)
    multi_normal_IQ_i = MultivariateNormal(μ_IQ_i, Σ_IQ_i)

    # partition and compute conditional distribution
    multi_normal_IQ_i.partition(i)
    scores_i = y[:i]
    μθ_hat_i, Σθ_hat_i = multi_normal_IQ_i.cond_dist(1, scores_i)

    # store the results
    μθ_hat_arr[i-1] = μθ_hat_i[0]
    Σθ_hat_arr[i-1] = Σθ_hat_i[0, 0]

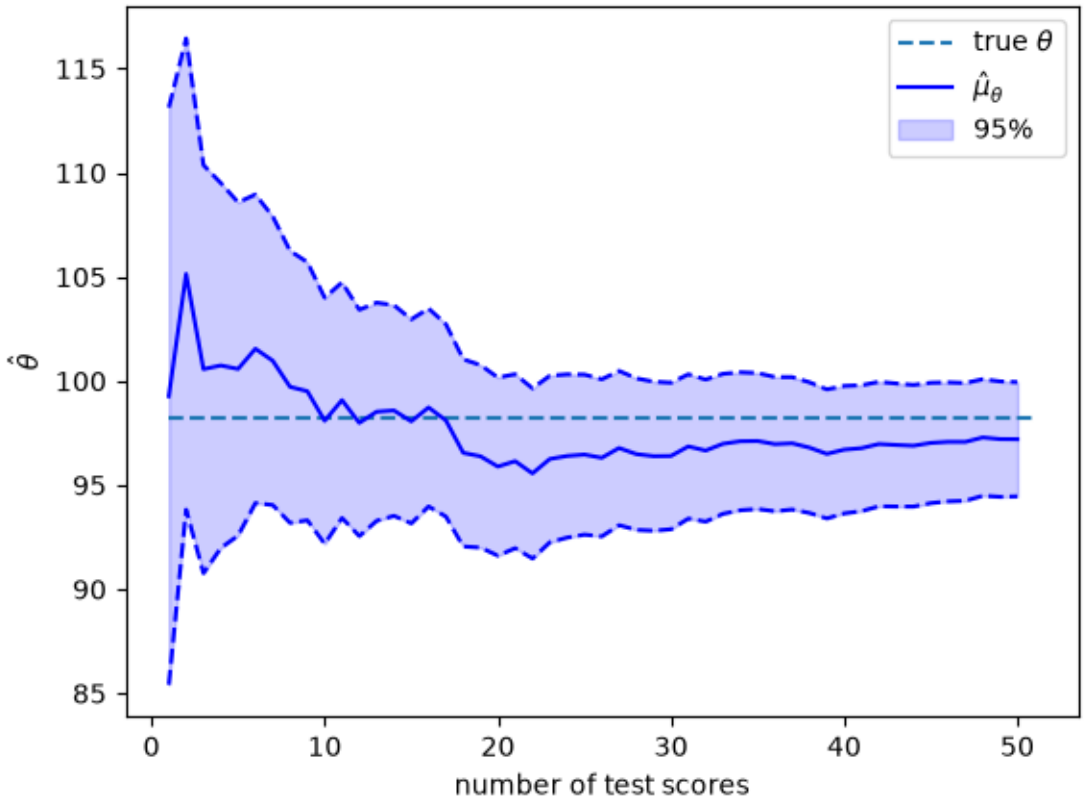
# transform variance to standard deviation
σθ_hat_arr = np.sqrt(Σθ_hat_arr)
```

```
μθ_hat_lower = μθ_hat_arr - 1.96 * σθ_hat_arr
μθ_hat_higher = μθ_hat_arr + 1.96 * σθ_hat_arr

plt.hlines(θ, 1, n+1, ls='--', label='true $θ$')
plt.plot(range(1, n+1), μθ_hat_arr, color='b', label=r'$\hat{μ}_{θ}$')
plt.plot(range(1, n+1), μθ_hat_lower, color='b', ls='--')
plt.plot(range(1, n+1), μθ_hat_higher, color='b', ls='--')
plt.fill_between(range(1, n+1), μθ_hat_lower, μθ_hat_higher,
                 color='b', alpha=0.2, label='95%')

plt.xlabel('number of test scores')
plt.ylabel(r'$\hat{θ}$')
plt.legend()

plt.show()
```



The solid blue line in the plot above shows $\hat{\mu}_\theta$ as a function of the number of test scores that we have recorded and conditioned on.

The blue area shows the span that comes from adding or subtracting $1.96\hat{\sigma}_\theta$ from $\hat{\mu}_\theta$.

Therefore, 95% of the probability mass of the conditional distribution falls in this range.

The value of the random θ that we drew is shown by the black dotted line.

As more and more test scores come in, our estimate of the person's θ become more and more reliable.

By staring at the changes in the conditional distributions, we see that adding more test scores makes $\hat{\theta}$ settle down and approach θ.

Thus, each y_i adds information about θ.

If we were to drive the number of tests $n \to +\infty$, the conditional standard deviation $\hat{\sigma}_\theta$ would converge to 0 at rate $\frac{1}{n^{.5}}$.

13.6 Information as surprise

By using a different representation, let's look at things from a different perspective.

We can represent the random vector X defined above as

$$X = \mu_\theta 1_{n+1} + C\epsilon, \quad \epsilon \sim N(0, I)$$

where C is a lower triangular **Cholesky factor** of Σ so that

$$\Sigma \equiv DD^\top = CC^\top$$

and

$$E\epsilon\epsilon^\top = I.$$

It follows that

$$\epsilon \sim N(0, I).$$

Let $G = C^{-1}$

G is also lower triangular.

We can compute ϵ from the formula

$$\epsilon = G\left(X - \mu_\theta 1_{n+1}\right)$$

This formula confirms that the orthonormal vector ϵ contains the same information as the non-orthogonal vector $(X - \mu_\theta 1_{n+1})$.

We can say that ϵ is an orthogonal basis for $(X - \mu_\theta 1_{n+1})$.

Let c_i be the ith element in the last row of C.

Then we can write

$$\theta = \mu_\theta + c_1\epsilon_1 + c_2\epsilon_2 + \cdots + c_n\epsilon_n + c_{n+1}\epsilon_{n+1} \tag{13.1}$$

The mutual orthogonality of the ϵ_i's provides us with an informative way to interpret them in light of equation (13.1).

Thus, relative to what is known from tests $i = 1, \ldots, n-1$, $c_i\epsilon_i$ is the amount of **new information** about θ brought by the test number i.

Here **new information** means **surprise** or what could not be predicted from earlier information.

Formula (13.1) also provides us with an enlightening way to express conditional means and conditional variances that we computed earlier.

In particular,

$$E\left[\theta \mid y_1, \ldots, y_k\right] = \mu_\theta + c_1\epsilon_1 + \cdots + c_k\epsilon_k$$

and

$$Var\left(\theta \mid y_1, \ldots, y_k\right) = c_{k+1}^2 + c_{k+2}^2 + \cdots + c_{n+1}^2.$$

```
C = np.linalg.cholesky(Σ_IQ)
G = np.linalg.inv(C)

ε = G @ (x - μθ)
```

```
cε = C[n, :] * ε

# compute the sequence of μθ and Σθ conditional on y1, y2, ..., yk
μθ_hat_arr_C = np.array([np.sum(cε[:k+1]) for k in range(n)]) + μθ
Σθ_hat_arr_C = np.array([C[n, i+1:n+1] @ C[n, i+1:n+1] for i in range(n)])
```

To confirm that these formulas give the same answers that we computed earlier, we can compare the means and variances of θ conditional on $\{y_i\}_{i=1}^k$ with what we obtained above using the formulas implemented in the class `MultivariateNormal` built on our original representation of conditional distributions for multivariate normal distributions.

```
# conditional mean
np.max(np.abs(μθ_hat_arr - μθ_hat_arr_C)) < 1e-10
```

```
np.True_
```

```
# conditional variance
np.max(np.abs(Σθ_hat_arr - Σθ_hat_arr_C)) < 1e-10
```

```
np.True_
```

13.7 Cholesky factor magic

Evidently, the Cholesky factorizations automatically computes the population **regression coefficients** and associated statistics that are produced by our `MultivariateNormal` class.

The Cholesky factorization computes these things **recursively**.

Indeed, in formula (13.1),

- the random variable $c_i \epsilon_i$ is information about θ that is not contained by the information in $\epsilon_1, \epsilon_2, \ldots, \epsilon_{i-1}$
- the coefficient c_i is the simple population regression coefficient of $\theta - \mu_\theta$ on ϵ_i

13.8 Math and verbal intelligence

We can alter the preceding example to be more realistic.

There is ample evidence that IQ is not a scalar.

Some people are good in math skills but poor in language skills.

Other people are good in language skills but poor in math skills.

So now we shall assume that there are two dimensions of IQ, θ and η.

These determine average performances in math and language tests, respectively.

We observe math scores $\{y_i\}_{i=1}^n$ and language scores $\{y_i\}_{i=n+1}^{2n}$.

When $n = 2$, we assume that outcomes are draws from a multivariate normal distribution with representation

$$X = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ \theta \\ \eta \end{bmatrix} = \begin{bmatrix} \mu_\theta \\ \mu_\theta \\ \mu_\eta \\ \mu_\eta \\ \mu_\theta \\ \mu_\eta \end{bmatrix} + \begin{bmatrix} \sigma_y & 0 & 0 & 0 & \sigma_\theta & 0 \\ 0 & \sigma_y & 0 & 0 & \sigma_\theta & 0 \\ 0 & 0 & \sigma_y & 0 & 0 & \sigma_\eta \\ 0 & 0 & 0 & \sigma_y & 0 & \sigma_\eta \\ 0 & 0 & 0 & 0 & \sigma_\theta & 0 \\ 0 & 0 & 0 & 0 & 0 & \sigma_\eta \end{bmatrix} \begin{bmatrix} w_1 \\ w_2 \\ w_3 \\ w_4 \\ w_5 \\ w_6 \end{bmatrix}$$

where $w = \begin{bmatrix} w_1 \\ w_2 \\ \vdots \\ w_6 \end{bmatrix}$ is a standard normal random vector.

We construct a Python function `construct_moments_IQ2d` to construct the mean vector and covariance matrix of the joint normal distribution.

```
def construct_moments_IQ2d(n, μθ, σθ, μη, ση, σy):

    μ_IQ2d = np.empty(2*(n+1))
    μ_IQ2d[:n] = μθ
    μ_IQ2d[2*n] = μθ
    μ_IQ2d[n:2*n] = μη
    μ_IQ2d[2*n+1] = μη


    D_IQ2d = np.zeros((2*(n+1), 2*(n+1)))
    D_IQ2d[range(2*n), range(2*n)] = σy
    D_IQ2d[:n, 2*n] = σθ
    D_IQ2d[2*n, 2*n] = σθ
    D_IQ2d[n:2*n, 2*n+1] = ση
    D_IQ2d[2*n+1, 2*n+1] = ση

    Σ_IQ2d = D_IQ2d @ D_IQ2d.T

    return μ_IQ2d, Σ_IQ2d, D_IQ2d
```

Let's put the function to work.

```
n = 2
# mean and variance of θ, η, and y
μθ, σθ, μη, ση, σy = 100., 10., 100., 10, 10

μ_IQ2d, Σ_IQ2d, D_IQ2d = construct_moments_IQ2d(n, μθ, σθ, μη, ση, σy)
μ_IQ2d, Σ_IQ2d, D_IQ2d
```

```
(array([100., 100., 100., 100., 100., 100.]),
 array([[200., 100.,   0.,   0., 100.,   0.],
        [100., 200.,   0.,   0., 100.,   0.],
        [  0.,   0., 200., 100.,   0., 100.],
        [  0.,   0., 100., 200.,   0., 100.],
        [100., 100.,   0.,   0., 100.,   0.],
        [  0.,   0., 100., 100.,   0., 100.]]),
 array([[10.,  0.,  0.,  0., 10.,  0.],
        [ 0., 10.,  0.,  0., 10.,  0.],
        [ 0.,  0., 10.,  0.,  0., 10.],
        [ 0.,  0.,  0., 10.,  0., 10.],
        [ 0.,  0.,  0.,  0., 10.,  0.],
        [ 0.,  0.,  0.,  0.,  0., 10.]]))
```

```
# take one draw
x = rng.multivariate_normal(μ_IQ2d, Σ_IQ2d)
y1 = x[:n]
y2 = x[n:2*n]
θ = x[2*n]
η = x[2*n+1]

# the true values
θ, η
```

```
(np.float64(86.84580129140488), np.float64(98.5057398320253))
```

We first compute the joint normal distribution of (θ, η).

```
multi_normal_IQ2d = MultivariateNormal(μ_IQ2d, Σ_IQ2d)

k = 2*n # the length of data vector
multi_normal_IQ2d.partition(k)

multi_normal_IQ2d.cond_dist(1, [*y1, *y2])
```

```
(array([ 87.15090791, 104.56448792]),
 array([[33.33333333,  0.        ],
        [ 0.        , 33.33333333]]))
```

Now let's compute distributions of θ and η separately conditional on various subsets of test scores.

It will be fun to compare outcomes with the help of an auxiliary function `cond_dist_IQ2d` that we now construct.

```
def cond_dist_IQ2d(μ, Σ, data):

    n = len(μ)

    multi_normal = MultivariateNormal(μ, Σ)
    multi_normal.partition(n-1)
    μ_hat, Σ_hat = multi_normal.cond_dist(1, data)

    return μ_hat, Σ_hat
```

Let's see how things work for an example.

```
for indices, IQ, conditions in [([*range(2*n), 2*n], 'θ', 'y1, y2, y3, y4'),
                                ([*range(n), 2*n], 'θ', 'y1, y2'),
                                ([*range(n, 2*n), 2*n], 'θ', 'y3, y4'),
                                ([*range(2*n), 2*n+1], 'η', 'y1, y2, y3, y4'),
                                ([*range(n), 2*n+1], 'η', 'y1, y2'),
                                ([*range(n, 2*n), 2*n+1], 'η', 'y3, y4')]:

    μ_hat, Σ_hat = cond_dist_IQ2d(μ_IQ2d[indices], Σ_IQ2d[indices][:, indices],␣
↪x[indices[:-1]])
    print(f'The mean and variance of {IQ} conditional on {conditions: <15} are ' +
          f'{μ_hat[0]:1.2f} and {Σ_hat[0, 0]:1.2f} respectively')
```

```
The mean and variance of θ conditional on y1, y2, y3, y4  are 87.15 and 33.33␣
↪respectively
The mean and variance of θ conditional on y1, y2          are 87.15 and 33.33␣
↪respectively
The mean and variance of θ conditional on y3, y4          are 100.00 and 100.00␣
↪respectively
The mean and variance of η conditional on y1, y2, y3, y4  are 104.56 and 33.33␣
↪respectively
The mean and variance of η conditional on y1, y2          are 100.00 and 100.00␣
↪respectively
The mean and variance of η conditional on y3, y4          are 104.56 and 33.33␣
↪respectively
```

Evidently, math tests provide no information about η and language tests provide no information about θ.

13.9 Univariate time series analysis

We can use the multivariate normal distribution and a little matrix algebra to present foundations of univariate linear time series analysis.

Let x_t, y_t, v_t, w_{t+1} each be scalars for $t \geq 0$.

Consider the following model:

$$\begin{aligned} x_0 &\sim N\left(0, \sigma_0^2\right) \\ x_{t+1} &= ax_t + bw_{t+1}, \quad w_{t+1} \sim N\left(0,1\right), t \geq 0 \\ y_t &= cx_t + dv_t, \quad v_t \sim N\left(0,1\right), t \geq 0 \end{aligned}$$

We can compute the moments of x_t

1. $Ex_{t+1}^2 = a^2 Ex_t^2 + b^2, t \geq 0$, where $Ex_0^2 = \sigma_0^2$
2. $Ex_{t+j}x_t = a^j Ex_t^2, \forall t \; \forall j$

Given some T, we can formulate the sequence $\{x_t\}_{t=0}^T$ as a random vector

$$X = \begin{bmatrix} x_0 \\ x_1 \\ \vdots \\ x_T \end{bmatrix}$$

and the covariance matrix Σ_x can be constructed using the moments we have computed above.

Similarly, we can define

$$Y = \begin{bmatrix} y_0 \\ y_1 \\ \vdots \\ y_T \end{bmatrix}, \quad v = \begin{bmatrix} v_0 \\ v_1 \\ \vdots \\ v_T \end{bmatrix}$$

and therefore

$$Y = CX + DV$$

where C and D are both diagonal matrices with constant c and d as diagonal respectively.

Consequently, the covariance matrix of Y is

$$\Sigma_y = EYY^\top = C\Sigma_x C^\top + DD^\top$$

By stacking X and Y, we can write

$$Z = \begin{bmatrix} X \\ Y \end{bmatrix}$$

and

$$\Sigma_z = EZZ^\top = \begin{bmatrix} \Sigma_x & \Sigma_x C^\top \\ C\Sigma_x & \Sigma_y \end{bmatrix}$$

Thus, the stacked sequences $\{x_t\}_{t=0}^T$ and $\{y_t\}_{t=0}^T$ jointly follow the multivariate normal distribution $N\left(0, \Sigma_z\right)$.

```
# as an example, consider the case where T = 3
T = 3
```

```
# variance of the initial distribution x_0
σ0 = 1.

# parameters of the equation system
a = .9
b = 1.
c = 1.0
d = .05
```

```
# construct the covariance matrix of X
Σx = np.empty((T+1, T+1))

Σx[0, 0] = σ0 ** 2
for i in range(T):
    Σx[i, i+1:] = Σx[i, i] * a ** np.arange(1, T+1-i)
    Σx[i+1:, i] = Σx[i, i+1:]

    Σx[i+1, i+1] = a ** 2 * Σx[i, i] + b ** 2
```

```
Σx
```

```
array([[1.      , 0.9     , 0.81    , 0.729   ],
       [0.9     , 1.81    , 1.629   , 1.4661  ],
       [0.81    , 1.629   , 2.4661  , 2.21949 ],
       [0.729   , 1.4661  , 2.21949 , 2.997541]])
```

```
# construct the covariance matrix of Y
C = np.eye(T+1) * c
D = np.eye(T+1) * d

Σy = C @ Σx @ C.T + D @ D.T
```

```
# construct the covariance matrix of Z
Σz = np.empty((2*(T+1), 2*(T+1)))

Σz[:T+1, :T+1] = Σx
Σz[:T+1, T+1:] = Σx @ C.T
Σz[T+1:, :T+1] = C @ Σx
Σz[T+1:, T+1:] = Σy
```

```
Σz
```

```
array([[1.      , 0.9     , 0.81    , 0.729   , 1.      , 0.9     ,
        0.81    , 0.729   ],
       [0.9     , 1.81    , 1.629   , 1.4661  , 0.9     , 1.81    ,
        1.629   , 1.4661  ],
       [0.81    , 1.629   , 2.4661  , 2.21949 , 0.81    , 1.629   ,
        2.4661  , 2.21949 ],
       [0.729   , 1.4661  , 2.21949 , 2.997541, 0.729   , 1.4661  ,
        2.21949 , 2.997541],
       [1.      , 0.9     , 0.81    , 0.729   , 1.0025  , 0.9     ,
        0.81    , 0.729   ],
       [0.9     , 1.81    , 1.629   , 1.4661  , 0.9     , 1.8125  ,
        1.629   , 1.4661  ],
       [0.81    , 1.629   , 2.4661  , 2.21949 , 0.81    , 1.629   ,
```

(continues on next page)

(continued from previous page)

```
        2.4686  , 2.21949 ],
       [0.729   , 1.4661  , 2.21949 , 2.997541, 0.729   , 1.4661  ,
        2.21949 , 3.000041]])
```

```
# construct the mean vector of Z
μz = np.zeros(2*(T+1))
```

The following Python code lets us sample random vectors X and Y.

This is going to be very useful for doing the conditioning to be used in the fun exercises below.

```
z = rng.multivariate_normal(μz, Σz)

x = z[:T+1]
y = z[T+1:]
```

13.9.1 Smoothing example

This is an instance of a classic `smoothing` calculation whose purpose is to compute $EX \mid Y$.

An interpretation of this example is

- X is a random sequence of hidden Markov state variables x_t
- Y is a sequence of observed signals y_t bearing information about the hidden state

```
# construct a MultivariateNormal instance
multi_normal_ex1 = MultivariateNormal(μz, Σz)
x = z[:T+1]
y = z[T+1:]
```

```
# partition Z into X and Y
multi_normal_ex1.partition(T+1)
```

```
# compute the conditional mean and covariance matrix of X given Y=y

print("X = ", x)
print("Y = ", y)
print(" E [ X | Y] = ", )

multi_normal_ex1.cond_dist(0, y)
```

```
X =  [ 0.241748   -0.48167052 -0.2751467   0.62696378]
Y =  [ 0.2465054  -0.55890629 -0.28918582  0.57124943]
 E [ X | Y] =
```

```
(array([ 0.24414853, -0.55648653, -0.28785472,  0.56917881]),
 array([[2.48875094e-03, 5.57449314e-06, 1.24861727e-08, 2.80242496e-11],
        [5.57449314e-06, 2.48876343e-03, 5.57452116e-06, 1.25113948e-08],
        [1.24861727e-08, 5.57452116e-06, 2.48876346e-03, 5.58575339e-06],
        [2.80242496e-11, 1.25113948e-08, 5.58575339e-06, 2.49377812e-03]]))
```

13.9.2 Filtering exercise

Compute $E\left[x_t \mid y_{t-1}, y_{t-2}, \dots, y_0\right]$.

To do so, we need to first construct the mean vector and the covariance matrix of the subvector $\left[x_t, y_0, \dots, y_{t-2}, y_{t-1}\right]$.

For example, let's say that we want the conditional distribution of x_3.

```
t = 3
```

```
# mean of the subvector
sub_μz = np.zeros(t+1)

# covariance matrix of the subvector
sub_Σz = np.empty((t+1, t+1))

sub_Σz[0, 0] = Σz[t, t] # x_t
sub_Σz[0, 1:] = Σz[t, T+1:T+t+1]
sub_Σz[1:, 0] = Σz[T+1:T+t+1, t]
sub_Σz[1:, 1:] = Σz[T+1:T+t+1, T+1:T+t+1]
```

```
sub_Σz
```

```
array([[2.997541, 0.729   , 1.4661  , 2.21949 ],
       [0.729   , 1.0025  , 0.9     , 0.81    ],
       [1.4661  , 0.9     , 1.8125  , 1.629   ],
       [2.21949 , 0.81    , 1.629   , 2.4686  ]])
```

```
multi_normal_ex2 = MultivariateNormal(sub_μz, sub_Σz)
multi_normal_ex2.partition(1)
```

```
sub_y = y[:t]

multi_normal_ex2.cond_dist(0, sub_y)
```

```
(array([-0.26074227]), array([[1.00201996]]))
```

13.9.3 Prediction exercise

Compute $E\left[y_t \mid y_{t-j}, \dots, y_0\right]$.

As what we did in exercise 2, we will construct the mean vector and covariance matrix of the subvector $\left[y_t, y_0, \dots, y_{t-j-1}, y_{t-j}\right]$.

For example, we take a case in which $t = 3$ and $j = 2$.

```
t = 3
j = 2
```

```
sub_μz = np.zeros(t-j+2)
sub_Σz = np.empty((t-j+2, t-j+2))

sub_Σz[0, 0] = Σz[T+t+1, T+t+1]
sub_Σz[0, 1:] = Σz[T+t+1, T+1:T+t-j+2]
```

(continues on next page)

(continued from previous page)

```
sub_Σz[1:, 0] = Σz[T+1:T+t-j+2, T+t+1]
sub_Σz[1:, 1:] = Σz[T+1:T+t-j+2, T+1:T+t-j+2]
```

```
sub_Σz
```

```
array([[3.000041, 0.729   , 1.4661  ],
       [0.729   , 1.0025  , 0.9     ],
       [1.4661  , 0.9     , 1.8125  ]])
```

```
multi_normal_ex3 = MultivariateNormal(sub_μz, sub_Σz)
multi_normal_ex3.partition(1)
```

```
sub_y = y[:t-j+1]

multi_normal_ex3.cond_dist(0, sub_y)
```

```
(array([-0.45114129]), array([[1.81413617]]))
```

13.9.4 Constructing a Wold representation

Now we'll apply Cholesky decomposition to decompose $\Sigma_y = HH^\top$ and form

$$\epsilon = H^{-1}Y.$$

Then we can represent y_t as

$$y_t = h_{t,t}\epsilon_t + h_{t,t-1}\epsilon_{t-1} + \cdots + h_{t,0}\epsilon_0.$$

```
H = np.linalg.cholesky(Σy)

H
```

```
array([[1.00124922, 0.        , 0.        , 0.        ],
       [0.8988771 , 1.00225743, 0.        , 0.        ],
       [0.80898939, 0.89978675, 1.00225743, 0.        ],
       [0.72809046, 0.80980808, 0.89978676, 1.00225743]])
```

```
ε = np.linalg.inv(H) @ y

ε
```

```
array([ 0.24619784, -0.7784506 ,  0.2116046 ,  0.83011777])
```

```
y
```

```
array([ 0.2465054 , -0.55890629, -0.28918582,  0.57124943])
```

This example is an instance of what is known as a **Wold representation** in time series analysis.

13.10 Stochastic difference equation

Consider the stochastic second-order linear difference equation

$$y_t = \alpha_0 + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + u_t$$

where $u_t \sim N\left(0, \sigma_u^2\right)$ and

$$\begin{bmatrix} y_{-1} \\ y_0 \end{bmatrix} \sim N\left(\mu_{\tilde{y}}, \Sigma_{\tilde{y}}\right)$$

It can be written as a stacked system

$$\underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ -\alpha_1 & 1 & 0 & 0 & \cdots & 0 & 0 & 0 \\ -\alpha_2 & -\alpha_1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & -\alpha_2 & -\alpha_1 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & -\alpha_2 & -\alpha_1 & 1 \end{bmatrix}}_{\equiv A} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ \vdots \\ y_T \end{bmatrix} = \underbrace{\begin{bmatrix} \alpha_0 + \alpha_1 y_0 + \alpha_2 y_{-1} \\ \alpha_0 + \alpha_2 y_0 \\ \alpha_0 \\ \alpha_0 \\ \vdots \\ \alpha_0 \end{bmatrix}}_{\equiv b} + \underbrace{\begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \\ \vdots \\ u_T \end{bmatrix}}_{\equiv u}$$

We can compute y by solving the system

$$y = A^{-1}\left(b + u\right)$$

We have

$$\begin{aligned} \mu_y &= A^{-1}\mu_b \\ \Sigma_y &= A^{-1}E\left[\left(b - \mu_b + u\right)\left(b - \mu_b + u\right)^\top\right]\left(A^{-1}\right)^\top \\ &= A^{-1}\left(\Sigma_b + \Sigma_u\right)\left(A^{-1}\right)^\top \end{aligned}$$

where

$$\mu_b = \begin{bmatrix} \alpha_0 + \alpha_1 \mu_{y_0} + \alpha_2 \mu_{y_{-1}} \\ \alpha_0 + \alpha_2 \mu_{y_0} \\ \alpha_0 \\ \vdots \\ \alpha_0 \end{bmatrix}$$

$$\Sigma_b = \begin{bmatrix} C\Sigma_{\tilde{y}} C^\top & 0_{N-2\times N-2} \\ 0_{N-2\times 2} & 0_{N-2\times N-2} \end{bmatrix}, \quad C = \begin{bmatrix} \alpha_2 & \alpha_1 \\ 0 & \alpha_2 \end{bmatrix}$$

$$\Sigma_u = \begin{bmatrix} \sigma_u^2 & 0 & \cdots & 0 \\ 0 & \sigma_u^2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \sigma_u^2 \end{bmatrix}$$

```
# set parameters
T = 160
# coefficients of the second order difference equation
a0 = 10
a1 = 1.53
a2 = -.9

# variance of u
σu = 10.

# distribution of y_{-1} and y_{0}
μy_tilde = np.array([1., 0.5])
Σy_tilde = np.array([[2., 1.], [1., 0.5]])
```

```
# construct A and A^\top
A = np.zeros((T, T))

for i in range(T):
    A[i, i] = 1

    if i-1 >= 0:
        A[i, i-1] = -α1

    if i-2 >= 0:
        A[i, i-2] = -α2

A_inv = np.linalg.inv(A)
```

```
# compute the mean vectors of b and y
μb = np.full(T, α0)
μb[0] += α1 * μy_tilde[1] + α2 * μy_tilde[0]
μb[1] += α2 * μy_tilde[1]

μy = A_inv @ μb
```

```
# compute the covariance matrices of b and y
Σu = np.eye(T) * σu ** 2

Σb = np.zeros((T, T))

C = np.array([[α2, α1], [0, α2]])
Σb[:2, :2] = C @ Σy_tilde @ C.T

Σy = A_inv @ (Σb + Σu) @ A_inv.T
```

13.11 Application to stock price model

Let

$$p_t = \sum_{j=0}^{T-t} \beta^j y_{t+j}$$

Form

$$\underbrace{\begin{bmatrix} p_1 \\ p_2 \\ p_3 \\ \vdots \\ p_T \end{bmatrix}}_{\equiv p} = \underbrace{\begin{bmatrix} 1 & \beta & \beta^2 & \cdots & \beta^{T-1} \\ 0 & 1 & \beta & \cdots & \beta^{T-2} \\ 0 & 0 & 1 & \cdots & \beta^{T-3} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{bmatrix}}_{\equiv B} \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ \vdots \\ y_T \end{bmatrix}$$

we have

$$\mu_p = B\mu_y$$
$$\Sigma_p = B\Sigma_y B^\top$$

```
β = .96
```

```
# construct B
B = np.zeros((T, T))

for i in range(T):
    B[i, i:] = β ** np.arange(0, T-i)
```

Denote

$$z = \begin{bmatrix} y \\ p \end{bmatrix} = \underbrace{\begin{bmatrix} I \\ B \end{bmatrix}}_{\equiv D} y$$

Thus, $\{y_t\}_{t=1}^T$ and $\{p_t\}_{t=1}^T$ jointly follow the multivariate normal distribution $N(\mu_z, \Sigma_z)$, where

$$\mu_z = D\mu_y$$

$$\Sigma_z = D\Sigma_y D^\top$$

```
D = np.vstack([np.eye(T), B])
```

```
μz = D @ μy
Σz = D @ Σy @ D.T
```

We can simulate paths of y_t and p_t and compute the conditional mean $E[p_t \mid y_{t-1}, y_t]$ using the `MultivariateNormal` class.

```
z = rng.multivariate_normal(μz, Σz)
y, p = z[:T], z[T:]
```

```
cond_Ep = np.empty(T-1)

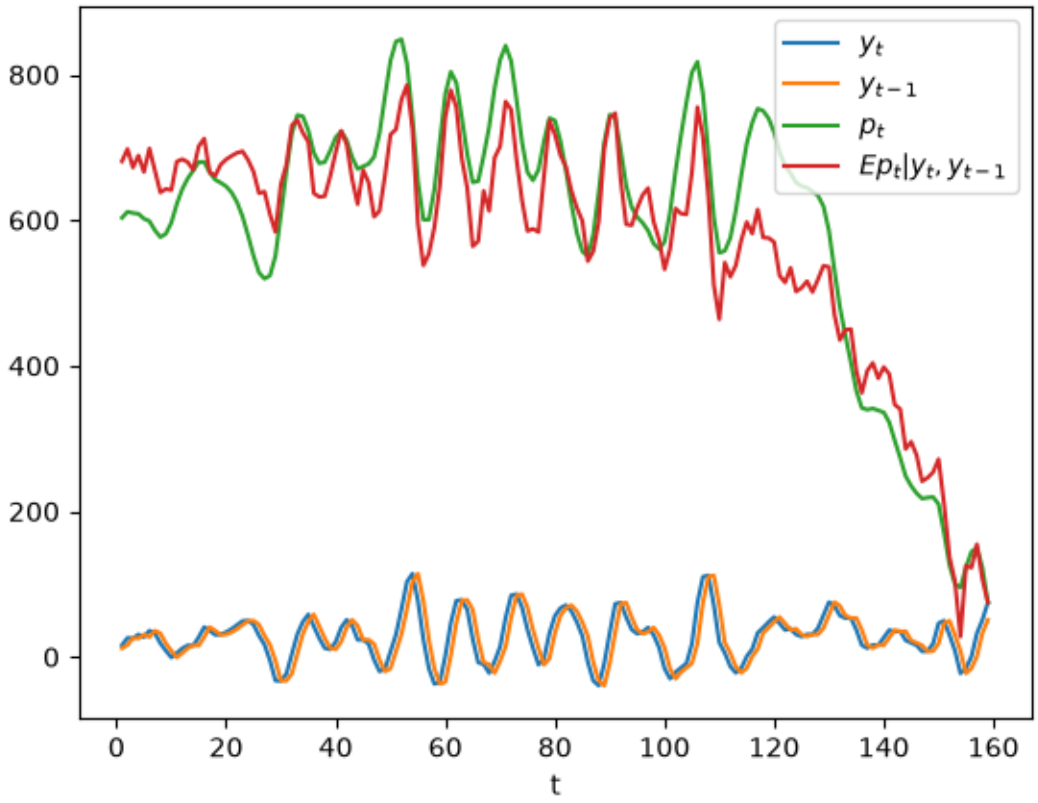
sub_μ = np.empty(3)
sub_Σ = np.empty((3, 3))
for t in range(2, T+1):
    sub_μ[:] = μz[[t-2, t-1, T-1+t]]
    sub_Σ[:, :] = Σz[[t-2, t-1, T-1+t], :][:, [t-2, t-1, T-1+t]]

    multi_normal = MultivariateNormal(sub_μ, sub_Σ)
    multi_normal.partition(2)

    cond_Ep[t-2] = multi_normal.cond_dist(1, y[t-2:t])[0][0]
```

```
plt.plot(range(1, T), y[1:], label='$y_{t}$')
plt.plot(range(1, T), y[:-1], label='$y_{t-1}$')
plt.plot(range(1, T), p[1:], label='$p_{t}$')
plt.plot(range(1, T), cond_Ep, label='$Ep_{t}|y_{t}, y_{t-1}$')

plt.xlabel('t')
plt.legend(loc=1)
plt.show()
```



In the above graph, the green line is what the price of the stock would be if people had perfect foresight about the path of dividends while the red line is the conditional expectation $Ep_t|y_t, y_{t-1}$, which is what the price would be if people did not have perfect foresight but were optimally predicting future dividends on the basis of the information y_t, y_{t-1} at time t.

13.12 Filtering foundations

Assume that x_0 is an $n \times 1$ random vector and that y_0 is a $p \times 1$ random vector determined by the *observation equation*

$$y_0 = Gx_0 + v_0, \quad x_0 \sim \mathcal{N}(\hat{x}_0, \Sigma_0), \quad v_0 \sim \mathcal{N}(0, R)$$

where v_0 is orthogonal to x_0, G is a $p \times n$ matrix, and R is a $p \times p$ positive definite matrix.

We consider the problem of someone who

- *observes* y_0
- does not observe x_0,
- knows $\hat{x}_0, \Sigma_0, G, R$ and therefore the joint probability distribution of the vector $\begin{bmatrix} x_0 \\ y_0 \end{bmatrix}$
- wants to infer x_0 from y_0 in light of what he knows about that joint probability distribution.

Therefore, the person wants to construct the probability distribution of x_0 conditional on the random vector y_0.

The joint distribution of $\begin{bmatrix} x_0 \\ y_0 \end{bmatrix}$ is multivariate normal $\mathcal{N}(\mu, \Sigma)$ with

$$\mu = \begin{bmatrix} \hat{x}_0 \\ G\hat{x}_0 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} \Sigma_0 & \Sigma_0 G' \\ G\Sigma_0 & G\Sigma_0 G' + R \end{bmatrix}$$

By applying an appropriate instance of the above formulas for the mean vector $\hat{\mu}_1$ and covariance matrix $\hat{\Sigma}_{11}$ of z_1 conditional on z_2, we find that the probability distribution of x_0 conditional on y_0 is $\mathcal{N}(\tilde{x}_0, \tilde{\Sigma}_0)$ where

$$\begin{aligned}\beta_0 &= \Sigma_0 G'(G\Sigma_0 G' + R)^{-1} \\ \tilde{x}_0 &= \hat{x}_0 + \beta_0(y_0 - G\hat{x}_0) \\ \tilde{\Sigma}_0 &= \Sigma_0 - \Sigma_0 G'(G\Sigma_0 G' + R)^{-1} G\Sigma_0\end{aligned}$$

We can express our finding that the probability distribution of x_0 conditional on y_0 is $\mathcal{N}(\tilde{x}_0, \tilde{\Sigma}_0)$ by representing x_0 as

$$x_0 = \tilde{x}_0 + \zeta_0 \tag{13.2}$$

where ζ_0 is a Gaussian random vector that is orthogonal to $\tilde{x}_0$ and y_0 and that has mean vector 0 and conditional covariance matrix $E[\zeta_0 \zeta_0' | y_0] = \tilde{\Sigma}_0$.

13.12.1 Step toward dynamics

Now suppose that we are in a time series setting and that we have the one-step state transition equation

$$x_1 = Ax_0 + Cw_1, \quad w_1 \sim \mathcal{N}(0, I)$$

where A is an $n \times n$ matrix and C is an $n \times m$ matrix.

Using equation (13.2), we can also represent x_1 as

$$x_1 = A(\tilde{x}_0 + \zeta_0) + Cw_1$$

It follows that

$$Ex_1 | y_0 = A\tilde{x}_0$$

and that the corresponding conditional covariance matrix $E(x_1 - Ex_1|y_0)(x_1 - Ex_1|y_0)' \equiv \Sigma_1$ is

$$\Sigma_1 = A\tilde{\Sigma}_0 A' + CC'$$

or

$$\Sigma_1 = A\Sigma_0 A' - A\Sigma_0 G'(G\Sigma_0 G' + R)^{-1} G\Sigma_0 A'$$

We can write the mean of x_1 conditional on y_0 as

$$\hat{x}_1 = A\hat{x}_0 + A\Sigma_0 G'(G\Sigma_0 G' + R)^{-1}(y_0 - G\hat{x}_0)$$

or

$$\hat{x}_1 = A\hat{x}_0 + K_0(y_0 - G\hat{x}_0)$$

where

$$K_0 = A\Sigma_0 G'(G\Sigma_0 G' + R)^{-1}$$

13.12.2 Dynamic version

Suppose now that for $t \geq 0$, $\{x_{t+1}, y_t\}_{t=0}^{\infty}$ are governed by the equations

$$\begin{aligned} x_{t+1} &= Ax_t + Cw_{t+1} \\ y_t &= Gx_t + v_t \end{aligned}$$

where as before $x_0 \sim \mathcal{N}(\hat{x}_0, \Sigma_0)$, w_{t+1} is the $t+1$th component of an i.i.d. stochastic process distributed as $w_{t+1} \sim \mathcal{N}(0, I)$, and v_t is the tth component of an i.i.d. process distributed as $v_t \sim \mathcal{N}(0, R)$ and the $\{w_{t+1}\}_{t=0}^{\infty}$ and $\{v_t\}_{t=0}^{\infty}$ processes are orthogonal at all pairs of dates.

The logic and formulas that we applied above imply that the probability distribution of x_t conditional on $y_0, y_1, \ldots, y_{t-1} = y^{t-1}$ is

$$x_t | y^{t-1} \sim \mathcal{N}(A\tilde{x}_{t-1}, A\tilde{\Sigma}_{t-1}A' + CC')$$

where $\{\tilde{x}_t, \tilde{\Sigma}_t\}_{t=1}^{\infty}$ can be computed by iterating on the following equations starting from $t = 1$ and initial conditions for $\tilde{x}_0, \tilde{\Sigma}_0$ computed as we have above:

$$\begin{aligned} \Sigma_t &= A\tilde{\Sigma}_{t-1}A' + CC' \\ \hat{x}_t &= A\tilde{x}_{t-1} \\ \beta_t &= \Sigma_t G'(G\Sigma_t G' + R)^{-1} \\ \tilde{x}_t &= \hat{x}_t + \beta_t(y_t - G\hat{x}_t) \\ \tilde{\Sigma}_t &= \Sigma_t - \Sigma_t G'(G\Sigma_t G' + R)^{-1} G\Sigma_t \end{aligned}$$

If we shift the first equation forward one period and then substitute the expression for $\tilde{\Sigma}_t$ on the right side of the fifth equation into it we obtain

$$\Sigma_{t+1} = CC' + A\Sigma_t A' - A\Sigma_t G'(G\Sigma_t G' + R)^{-1} G\Sigma_t A'.$$

This is a matrix Riccati difference equation that is closely related to another matrix Riccati difference equation that appears in a quantecon lecture on the basics of linear quadratic control theory.

That equation has the form

$$P_{t-1} = R + A'P_t A - A'P_t B(B'P_t B + Q)^{-1} B'P_t A.$$

Stare at the two preceding equations for a moment or two, the first being a matrix difference equation for a conditional covariance matrix, the second being a matrix difference equation in the matrix appearing in a quadratic form for an intertemporal cost of value function.

Although the two equations are not identical, they display striking family resemblances.

- the first equation tells dynamics that work **forward** in time
- the second equation tells dynamics that work **backward** in time
- while many of the terms are similar, one equation seems to apply matrix transformations to some matrices that play similar roles in the other equation

The family resemblences of these two equations reflects a transcendent **duality** that prevails between control theory and filtering theory.

13.12.3 An example

We can use the Python class *MultivariateNormal* to construct examples.

Here is an example for a single period problem at time 0

```
G = np.array([[1., 3.]])
R = np.array([[1.]])

x0_hat = np.array([0., 1.])
Σ0 = np.array([[1., .5], [.5, 2.]])

μ = np.hstack([x0_hat, G @ x0_hat])
Σ = np.block([[Σ0, Σ0 @ G.T], [G @ Σ0, G @ Σ0 @ G.T + R]])
```

```
# construction of the multivariate normal instance
multi_normal = MultivariateNormal(μ, Σ)
```

```
multi_normal.partition(2)
```

```
# the observation of y
y0 = 2.3

# conditional distribution of x0
μ1_hat, Σ11 = multi_normal.cond_dist(0, y0)
μ1_hat, Σ11
```

```
(array([-0.07608696,  0.80217391]),
 array([[ 0.72826087, -0.20652174],
        [-0.20652174,  0.16304348]]))
```

```
A = np.array([[0.5, 0.2], [-0.1, 0.3]])
C = np.array([[2.], [1.]])

# conditional distribution of x1
x1_cond = A @ μ1_hat
Σ1_cond = C @ C.T + A @ Σ11 @ A.T
x1_cond, Σ1_cond
```

```
(array([0.1223913 , 0.24826087]),
 array([[4.14728261, 1.94652174],
        [1.94652174, 1.03434783]]))
```

13.12.4 Code for iterating

Here is code for solving a dynamic filtering problem by iterating on our equations, followed by an example.

```
def iterate(x0_hat, Σ0, A, C, G, R, y_seq):

    p, n = G.shape

    T = len(y_seq)
    x_hat_seq = np.empty((T+1, n))
    Σ_hat_seq = np.empty((T+1, n, n))
```

(continues on next page)

(continued from previous page)

```
    x_hat_seq[0] = x0_hat
    Σ_hat_seq[0] = Σ0

    for t in range(T):
        xt_hat = x_hat_seq[t]
        Σt = Σ_hat_seq[t]
        μ = np.hstack([xt_hat, G @ xt_hat])
        Σ = np.block([[Σt, Σt @ G.T], [G @ Σt, G @ Σt @ G.T + R]])

        # filtering
        multi_normal = MultivariateNormal(μ, Σ)
        multi_normal.partition(n)
        x_tilde, Σ_tilde = multi_normal.cond_dist(0, y_seq[t])

        # forecasting
        x_hat_seq[t+1] = A @ x_tilde
        Σ_hat_seq[t+1] = C @ C.T + A @ Σ_tilde @ A.T

    return x_hat_seq, Σ_hat_seq
```

```
iterate(x0_hat, Σ0, A, C, G, R, [2.3, 1.2, 3.2])
```

```
(array([[0.        , 1.        ],
        [0.1223913 , 0.24826087],
        [0.18729801, 0.06881254],
        [0.75600369, 0.05547207]]),
 array([[[1.        , 0.5       ],
         [0.5       , 2.        ]],

        [[4.14728261, 1.94652174],
         [1.94652174, 1.03434783]],

        [[4.08853181, 1.98917687],
         [1.98917687, 1.00757802]],

        [[4.06650435, 1.99952946],
         [1.99952946, 1.00271293]]]))
```

The iterative algorithm just described is a version of the celebrated **Kalman filter**.

We describe the Kalman filter and some applications of it in *A First Look at the Kalman Filter*

13.13 Classic factor analysis model

The factor analysis model can be represented as

$$Y = \Lambda f + U$$

where

1. Y is $n \times 1$ random vector, $EUU^\top = D$ is a diagonal matrix,
2. Λ is $n \times k$ coefficient matrix,
3. f is $k \times 1$ random vector, $Eff^\top = I$,

4. U is $n \times 1$ random vector, and $U \perp f$ (i.e., $EUf^\top = 0$)
5. It is presumed that k is small relative to n; often k is only 1 or 2, as in our IQ examples.

This implies that

$$\begin{aligned} \Sigma_y = EYY^\top &= \Lambda\Lambda^\top + D \\ EYf^\top &= \Lambda \\ EfY^\top &= \Lambda^\top \end{aligned}$$

Thus, the covariance matrix Σ_Y is the sum of a diagonal matrix D and a positive semi-definite matrix $\Lambda\Lambda^\top$ of rank k.

This means that all covariances among the n components of the Y vector are intermediated by their common dependencies on the k factors.

Form

$$Z = \begin{pmatrix} f \\ Y \end{pmatrix}$$

the covariance matrix of the expanded random vector Z can be computed as

$$\Sigma_z = EZZ^\top = \begin{pmatrix} I & \Lambda^\top \\ \Lambda & \Lambda\Lambda^\top + D \end{pmatrix}$$

In the following, we first construct the mean vector and the covariance matrix for the case where $N = 10$ and $k = 2$.

```
N = 10
k = 2
```

We set the coefficient matrix Λ and the covariance matrix of U to be

$$\Lambda = \begin{pmatrix} 1 & 0 \\ \vdots & \vdots \\ 1 & 0 \\ 0 & 1 \\ \vdots & \vdots \\ 0 & 1 \end{pmatrix}, \quad D = \begin{pmatrix} \sigma_u^2 & 0 & \cdots & 0 \\ 0 & \sigma_u^2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \cdots & \sigma_u^2 \end{pmatrix}$$

where the first half of the first column of Λ is filled with 1s and 0s for the rest half, and symmetrically for the second column.

D is a diagonal matrix with parameter σ_u^2 on the diagonal.

```
Λ = np.zeros((N, k))
Λ[:N//2, 0] = 1
Λ[N//2:, 1] = 1

σu = .5
D = np.eye(N) * σu ** 2
```

```
# compute Σy
Σy = Λ @ Λ.T + D
```

We can now construct the mean vector and the covariance matrix for Z.

```
μz = np.zeros(k+N)

Σz = np.empty((k+N, k+N))
```

(continues on next page)

(continued from previous page)

```
Σz[:k, :k] = np.eye(k)
Σz[:k, k:] = Λ.T
Σz[k:, :k] = Λ
Σz[k:, k:] = Σy
```

```
z = rng.multivariate_normal(μz, Σz)

f = z[:k]
y = z[k:]
```

```
multi_normal_factor = MultivariateNormal(μz, Σz)
multi_normal_factor.partition(k)
```

Let's compute the conditional distribution of the hidden factor f on the observations Y, namely, $f \mid Y = y$.

```
multi_normal_factor.cond_dist(0, y)
```

```
(array([0.14356426, 1.84615504]),
 array([[0.04761905, 0.        ],
        [0.        , 0.04761905]]))
```

We can verify that the conditional mean $E[f \mid Y = y] = BY$ where $B = \Lambda^\top \Sigma_y^{-1}$.

```
B = Λ.T @ np.linalg.inv(Σy)

B @ y
```

```
array([0.14356426, 1.84615504])
```

Similarly, we can compute the conditional distribution $Y \mid f$.

```
multi_normal_factor.cond_dist(1, f)
```

```
(array([0.41858317, 0.41858317, 0.41858317, 0.41858317, 0.41858317,
        2.04367341, 2.04367341, 2.04367341, 2.04367341, 2.04367341]),
 array([[0.25, 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
        [0.  , 0.25, 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
        [0.  , 0.  , 0.25, 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
        [0.  , 0.  , 0.  , 0.25, 0.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
        [0.  , 0.  , 0.  , 0.  , 0.25, 0.  , 0.  , 0.  , 0.  , 0.  ],
        [0.  , 0.  , 0.  , 0.  , 0.  , 0.25, 0.  , 0.  , 0.  , 0.  ],
        [0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.25, 0.  , 0.  , 0.  ],
        [0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.25, 0.  , 0.  ],
        [0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.25, 0.  ],
        [0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.  , 0.25]]))
```

It can be verified that the mean is $\Lambda I^{-1} f = \Lambda f$.

```
Λ @ f
```

```
array([0.41858317, 0.41858317, 0.41858317, 0.41858317, 0.41858317,
       2.04367341, 2.04367341, 2.04367341, 2.04367341, 2.04367341])
```

13.14 PCA and factor analysis

To learn about Principal Components Analysis (PCA), please see this lecture *Singular Value Decompositions*.

For fun, let's apply a PCA decomposition to a covariance matrix Σ_y that in fact is governed by our factor-analytic model.

Technically, this means that the PCA model is misspecified.

(Can you explain why?)

Nevertheless, this exercise will let us study how well the first two principal components from a PCA can approximate the conditional expectations $Ef_i|Y$ for our two factors $f_i, i = 1, 2$ for the factor analytic model that we have assumed truly governs the data on Y we have generated.

So we compute the PCA decomposition

$$\Sigma_y = P\tilde{\Lambda}P^\top$$

where $\tilde{\Lambda}$ is a diagonal matrix.

We have

$$Y = P\epsilon$$

and

$$\epsilon = P^\top Y$$

Note that we will arrange the eigenvectors in P in the *descending* order of eigenvalues.

```
λ_tilde, P = np.linalg.eigh(Σy)

# arrange the eigenvectors by eigenvalues
ind = sorted(range(N), key=lambda x: λ_tilde[x], reverse=True)

P = P[:, ind]
λ_tilde = λ_tilde[ind]
Λ_tilde = np.diag(λ_tilde)

print('λ_tilde =', λ_tilde)
```

```
λ_tilde = [5.25 5.25 0.25 0.25 0.25 0.25 0.25 0.25 0.25 0.25]
```

```
# verify the orthogonality of eigenvectors
np.abs(P @ P.T - np.eye(N)).max()
```

```
np.float64(4.440892098500626e-16)
```

```
# verify the eigenvalue decomposition is correct
P @ Λ_tilde @ P.T
```

```
array([[1.25, 1.  , 1.  , 1.  , 1.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
       [1.  , 1.25, 1.  , 1.  , 1.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
       [1.  , 1.  , 1.25, 1.  , 1.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
       [1.  , 1.  , 1.  , 1.25, 1.  , 0.  , 0.  , 0.  , 0.  , 0.  ],
       [1.  , 1.  , 1.  , 1.  , 1.25, 0.  , 0.  , 0.  , 0.  , 0.  ],
```

(continues on next page)

(continued from previous page)

```
       [0.  , 0.  , 0.  , 0.  , 0.  , 1.25, 1.  , 1.  , 1.  , 1.  ],
       [0.  , 0.  , 0.  , 0.  , 0.  , 1.  , 1.25, 1.  , 1.  , 1.  ],
       [0.  , 0.  , 0.  , 0.  , 0.  , 1.  , 1.  , 1.25, 1.  , 1.  ],
       [0.  , 0.  , 0.  , 0.  , 0.  , 1.  , 1.  , 1.  , 1.25, 1.  ],
       [0.  , 0.  , 0.  , 0.  , 0.  , 1.  , 1.  , 1.  , 1.  , 1.25]])
```

```
ε = P.T @ y

print("ε = ", ε)
```

```
ε =  [ 0.33707041  4.33453458  0.26386722 -0.39706296 -0.30020843  0.45473871
  0.19640695  0.07004059  0.35348891  0.81330557]
```

```
# print the values of the two factors

print('f = ', f)
```

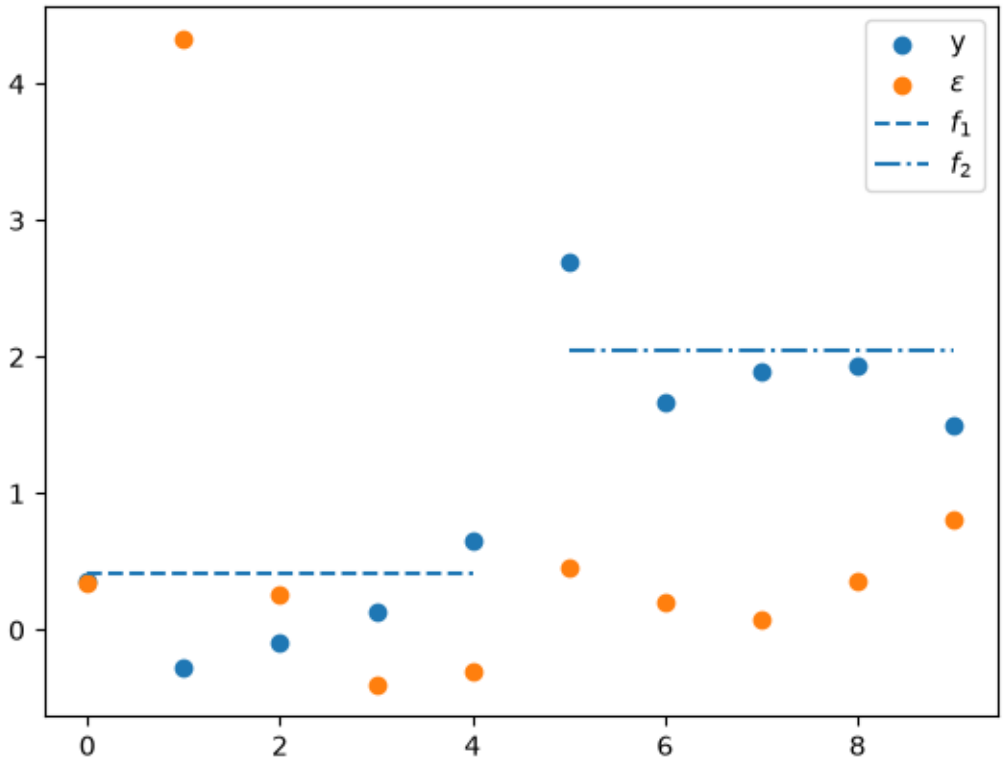
```
f =  [0.41858317 2.04367341]
```

Below we'll plot several things

- the N values of y
- the N values of the principal components ϵ
- the value of the first factor f_1 plotted only for the first $N/2$ observations of y for which it receives a non-zero loading in Λ
- the value of the second factor f_2 plotted only for the final $N/2$ observations for which it receives a non-zero loading in Λ

```
plt.scatter(range(N), y, label='y')
plt.scatter(range(N), ε, label=r'$\epsilon$')
plt.hlines(f[0], 0, N//2-1, ls='--', label='$f_{1}$')
plt.hlines(f[1], N//2, N-1, ls='-.', label='$f_{2}$')
plt.legend()

plt.show()
```



Consequently, the first two ϵ_j correspond to the largest two eigenvalues.

Let's look at them, after which we'll look at $Ef|y = By$

```
ε[:2]
```

```
array([0.33707041, 4.33453458])
```

```
# compare with Ef|y
B @ y
```

```
array([0.14356426, 1.84615504])
```

Note

The two largest eigenvalues are both 5.25 in this example.

When an eigenvalue is repeated, the associated principal components are not individually pinned down: any orthonormal basis for the same two-dimensional eigenspace is valid.

For that reason, it is not meaningful to compare ϵ_1 and ϵ_2 component-by-component with $E[f \mid Y]$.

The PC scores live in a PCA coordinate system, while $E[f \mid Y]$ lives in factor space.

Even within the common two-dimensional subspace, the PCA basis can be rotated or sign-flipped, and its coordinates need not use the same scaling as the factor coordinates.

What is uniquely determined is the two-dimensional subspace spanned by the first two columns of P.

In this symmetric example, that subspace is exactly the column space of Λ.

The fraction of variance in y_t explained by the first two principal components is

```
λ_tilde[:2].sum() / λ_tilde.sum()
```

```
np.float64(0.84)
```

To compare PCA with the factor model in observation space, compute

$$\hat{Y} = P_j \epsilon_j + P_k \epsilon_k$$

where P_j and P_k are the eigenvectors associated with the two largest eigenvalues.

```
y_hat = P[:, :2] @ ε[:2]
```

$\hat{Y}$ is the rank-2 PCA approximation to Y in observation space, so it is a 10-vector rather than a 2-vector.

The natural observation-space counterpart from the factor model is $\Lambda E[f \mid Y]$, which is also a 10-vector.

In this symmetric example, both vectors lie in the same two-dimensional subspace, namely the column space of Λ.

They are therefore close, but not identical.

The PCA reconstruction uses the block means directly, while $\Lambda E[f \mid Y]$ shrinks those block means toward zero by the factor $5/(5 + \sigma_u^2) \approx 0.952$.

The next plot makes this comparison concrete.

The two scatter plots, $E[Y \mid f] = \Lambda f$ and $\hat{Y}$, are both 10-vectors in observation space, so they can be compared directly.

The horizontal lines show the factor values f_1 and f_2, together with their posterior means $E[f_i \mid Y]$.

These are 2-dimensional factor-space quantities, drawn over the relevant half of the index set to match the block structure of Λ.

This uses the same idea as the earlier formula $E[Y \mid f] = \Lambda f$: the matrix Λ maps a 2-vector in factor space into a 10-vector in observation space.

In our example,

$$\Lambda a = \begin{bmatrix} a_1 \\ \vdots \\ a_1 \\ a_2 \\ \vdots \\ a_2 \end{bmatrix} \quad \text{for any } a = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix},$$

because the first five rows of Λ are $(1, 0)$ and the last five rows are $(0, 1)$.

Therefore, once we observe $Y = y$, the posterior mean $E[f \mid Y = y] = \begin{bmatrix} E[f_1 \mid y] \\ E[f_2 \mid y] \end{bmatrix}$ is converted into the observation-space vector

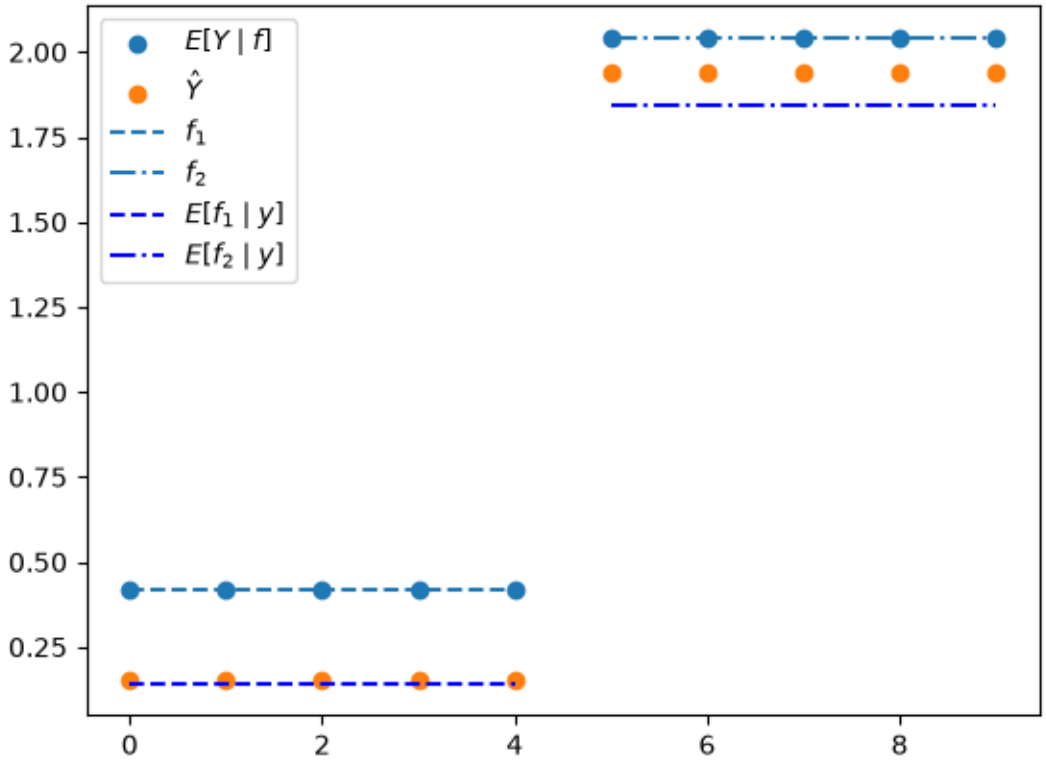
$$\Lambda E[f \mid Y = y] = \begin{bmatrix} E[f_1 \mid y] \\ \vdots \\ E[f_1 \mid y] \\ E[f_2 \mid y] \\ \vdots \\ E[f_2 \mid y] \end{bmatrix}.$$

So the horizontal line at height $E[f_1 \mid y]$ over the first five indices, together with the horizontal line at height $E[f_2 \mid y]$ over the last five indices, is exactly a picture of $\Lambda E[f \mid Y = y]$.

```
plt.scatter(range(N), Λ @ f, label=r'$E[Y \mid f]$')
plt.scatter(range(N), y_hat, label=r'$\hat{Y}$')
plt.hlines(f[0], 0, N//2-1, ls='--', label='$f_{1}$')
plt.hlines(f[1], N//2, N-1, ls='-.', label='$f_{2}$')

Efy = B @ y
plt.hlines(Efy[0], 0, N//2-1, ls='--', color='b', label=r'$E[f_1 \mid y]$')
plt.hlines(Efy[1], N//2, N-1, ls='-.', color='b', label=r'$E[f_2 \mid y]$')
plt.legend()

plt.show()
```



To compute the covariance matrix of $\hat{Y}$, first form the covariance matrix of ϵ and then extract the upper-left block corresponding to ϵ_1 and ϵ_2.

```
Σεjk = (P.T @ Σy @ P)[:2, :2]

Pjk = P[:, :2]

Σy_hat = Pjk @ Σεjk @ Pjk.T
print('Σy_hat = \n', Σy_hat)
```

```
Σy_hat =
 [[1.05 1.05 1.05 1.05 1.05 0.   0.   0.   0.   0.  ]
 [1.05 1.05 1.05 1.05 1.05 0.   0.   0.   0.   0.  ]
 [1.05 1.05 1.05 1.05 1.05 0.   0.   0.   0.   0.  ]
 [1.05 1.05 1.05 1.05 1.05 0.   0.   0.   0.   0.  ]
 [1.05 1.05 1.05 1.05 1.05 0.   0.   0.   0.   0.  ]
 [0.   0.   0.   0.   0.   1.05 1.05 1.05 1.05 1.05]
```

(continues on next page)

(continued from previous page)

```
[0.   0.   0.   0.   0.   1.05 1.05 1.05 1.05 1.05]
[0.   0.   0.   0.   0.   1.05 1.05 1.05 1.05 1.05]
[0.   0.   0.   0.   0.   1.05 1.05 1.05 1.05 1.05]
[0.   0.   0.   0.   0.   1.05 1.05 1.05 1.05 1.05]]
```

13.15 Exercises

Exercise 13.15.1

Verify conditional mean and variance by simulation

For the bivariate normal with

$$\mu = \begin{bmatrix} 0.5 \\ 1.0 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 1 & 0.5 \\ 0.5 & 1 \end{bmatrix}$$

fix $z_2 = 2$.

1. Use `MultivariateNormal` to compute the analytical conditional mean $\hat{\mu}_1$ and variance $\hat{\Sigma}_{11}$ of $z_1 \mid z_2 = 2$.
2. Draw 10^6 samples from the joint distribution.

   Retain only those for which $|z_2 - 2| < 0.05$.

   Compute the sample mean and variance of the retained z_1 values.
3. Confirm that the sample estimates are close to the analytical values.

Solution

Here is one solution:

```
μ = np.array([.5, 1.])
Σ = np.array([[1., .5], [.5, 1.]])

mn = MultivariateNormal(μ, Σ)
mn.partition(1)
μ1_hat, Σ11_hat = mn.cond_dist(0, np.array([2.]))
print(f"Analytical  μ1_hat = {μ1_hat[0]:.4f},  Σ11_hat = {Σ11_hat[0,0]:.4f}")

n = 1_000_000
data = rng.multivariate_normal(μ, Σ, size=n)
z1_all, z2_all = data[:, 0], data[:, 1]

mask = np.abs(z2_all - 2.) < 0.05
z1_cond = z1_all[mask]
print(f"Sample size in band: {mask.sum()}")
print(f"Sample      μ1_hat = {np.mean(z1_cond):.4f},  Σ11_hat = {np.var(z1_cond,␣
↪ddof=1):.4f}")
```

```
Analytical  μ1_hat = 1.0000,  Σ11_hat = 0.7500
Sample size in band: 24314
Sample      μ1_hat = 1.0050,  Σ11_hat = 0.7458
```

Exercise 13.15.2

Product of regression slopes equals squared correlation

For a bivariate normal with standard deviations $\sigma_1 = \sigma_2 = 1$ and correlation ρ, show analytically that $b_1 b_2 = \rho^2$, where b_1 is the slope of z_1 on z_2 and b_2 is the slope of z_2 on z_1.

Then verify numerically for $\rho \in \{0.2, 0.5, 0.9\}$ that `βs[0] * βs[1]` $= \rho^2$ by constructing the appropriate `MultivariateNormal` instances.

Solution

The regression slopes are

$$b_1 = \frac{\Sigma_{12}}{\Sigma_{22}} = \frac{\rho \sigma_1 \sigma_2}{\sigma_2^2} = \rho \frac{\sigma_1}{\sigma_2}, \qquad b_2 = \frac{\Sigma_{21}}{\Sigma_{11}} = \rho \frac{\sigma_2}{\sigma_1}$$

so $b_1 b_2 = \rho^2$.

```
for ρ in [0.2, 0.5, 0.9]:
    Σ = np.array([[1., ρ], [ρ, 1.]])
    mn = MultivariateNormal(np.zeros(2), Σ)
    mn.partition(1)
    product = mn.βs[0].item() * mn.βs[1].item()
    print(f"ρ={ρ:.1f}:  b1*b2 = {product:.4f}")
    print(f"ρ^2 = {ρ**2:.4f},  match: {np.isclose(product, ρ**2)}")
```

```
ρ=0.2:  b1*b2 = 0.0400
ρ^2 = 0.0400,  match: True
ρ=0.5:  b1*b2 = 0.2500
ρ^2 = 0.2500,  match: True
ρ=0.9:  b1*b2 = 0.8100
ρ^2 = 0.8100,  match: True
```

Exercise 13.15.3

IQ inference: effect of the signal-to-noise ratio

Using the one-dimensional IQ model with $n = 50$ test scores and $\mu_\theta = 100$, $\sigma_\theta = 10$:

1. Vary the test-score noise $\sigma_y \in \{1, 5, 10, 20, 50\}$.
   - For each value, plot the posterior standard deviation $\hat{\sigma}_\theta$ as a function of the number of test scores included (from 1 to 50), with all curves on the same axes.
2. Explain intuitively why a larger σ_y leads to a slower decline of posterior uncertainty.

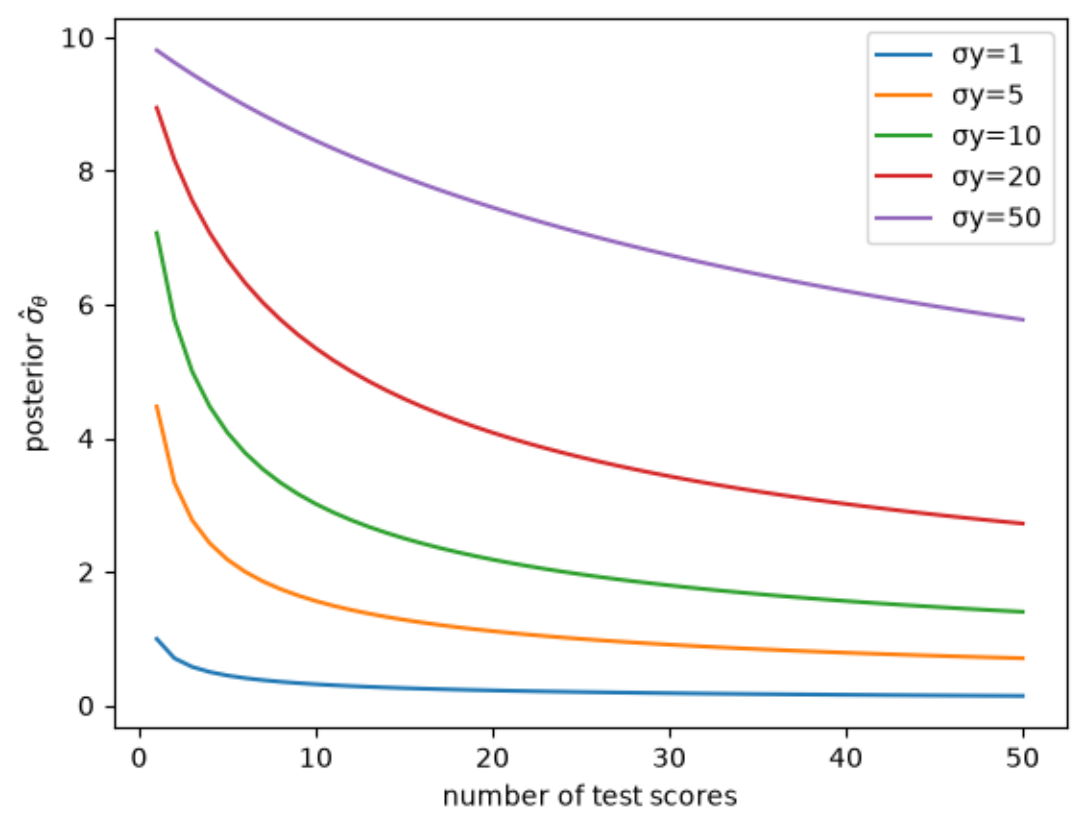
Solution

Here is one solution:

```
n_max = 50
μθ_val, σθ_val = 100., 10.

fig, ax = plt.subplots()
for σy_val in [1., 5., 10., 20., 50.]:
    σθ_hat_arr = np.empty(n_max)
    for i in range(1, n_max + 1):
        μ_i, Σ_i, _ = construct_moments_IQ(i, μθ_val, σθ_val, σy_val)
        mn_i = MultivariateNormal(μ_i, Σ_i)
        mn_i.partition(i)
        _, Σθ_i = mn_i.cond_dist(1, np.zeros(i))
        σθ_hat_arr[i - 1] = np.sqrt(Σθ_i[0, 0])
    ax.plot(range(1, n_max + 1), σθ_hat_arr, label=f'σy={σy_val:.0f}')

ax.set_xlabel('number of test scores')
ax.set_ylabel(r'posterior $\hat{\sigma}_\theta$')
ax.legend()
plt.show()
```



When σ_y is large each test score is a noisy signal about θ, so many more observations are required before the posterior variance falls appreciably.

In the limit $\sigma_y \to 0$ a single observation pins down θ exactly.

Exercise 13.15.4

Prior vs. likelihood in IQ inference

Using the one-dimensional IQ model with $n = 20$ test scores and $\mu_\theta = 100$, $\sigma_y = 10$:

1. Fix $\sigma_y = 10$ and vary the prior spread $\sigma_\theta \in \{1, 5, 10, 50, 500\}$.
   - For each value compute the posterior mean $\hat{\mu}_\theta$ given the same set of $n = 20$ test scores and plot $\hat{\mu}_\theta$ against σ_θ.
2. Show analytically (or verify numerically) that
   - as $\sigma_\theta \to \infty$ the posterior mean converges to the sample mean $\bar{y}$ (the data dominate the prior), and
   - as $\sigma_\theta \to 0$ the posterior mean converges to the prior mean μ_θ (the prior dominates the data).

Hint

The posterior mean formula is $\hat{\mu}_\theta = (\mu_\theta/\sigma_\theta^2 + n\bar{y}/\sigma_y^2)\big/(1/\sigma_\theta^2 + n/\sigma_y^2)$.

Examine each limit by letting σ_θ go to ∞ or 0.

Solution

Here is one solution:

```
n_scores = 20
μθ_val, σy_val = 100., 10.

rng = np.random.default_rng(42)
true_θ = 108.
y_obs = true_θ + σy_val * rng.standard_normal(n_scores)
y_bar = np.mean(y_obs)

σθ_vals = [1., 5., 10., 50., 500.]
μθ_hat_vals = []

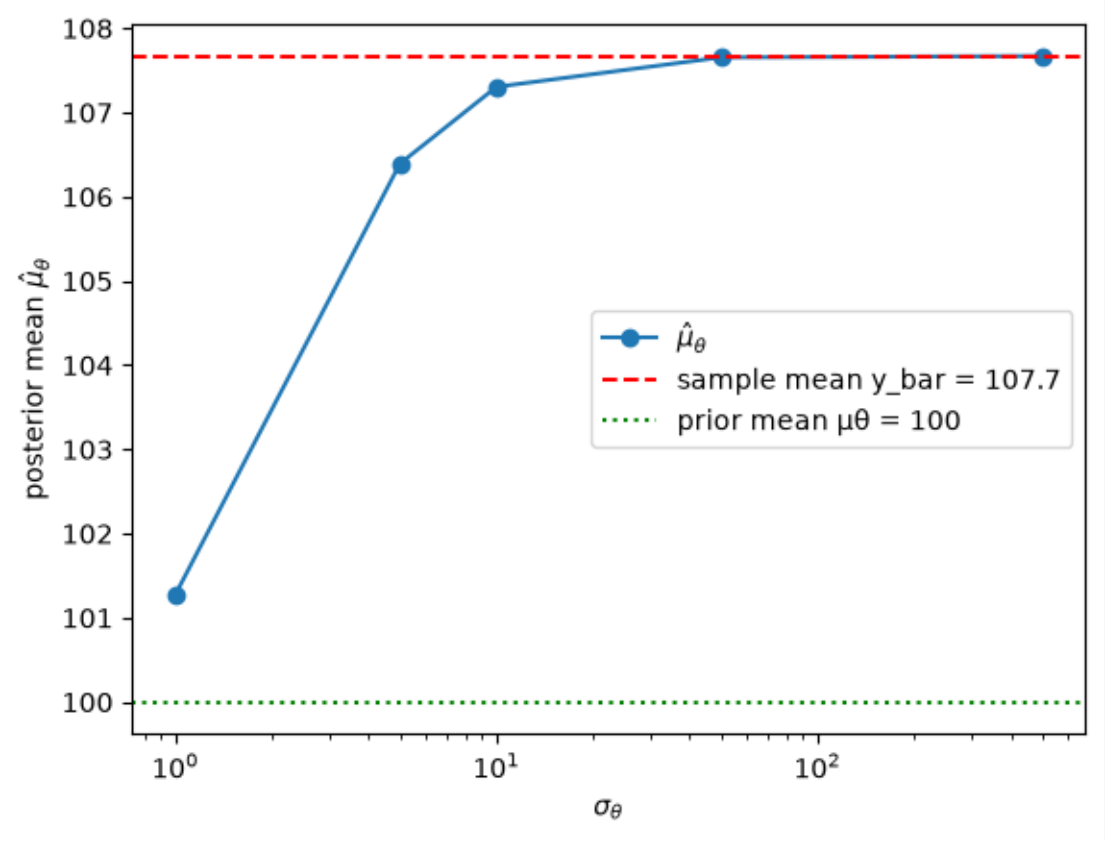
for σθ_val in σθ_vals:
    μ_i, Σ_i, _ = construct_moments_IQ(n_scores, μθ_val, σθ_val, σy_val)
    mn_i = MultivariateNormal(μ_i, Σ_i)
    mn_i.partition(n_scores)
    μθ_hat, _ = mn_i.cond_dist(1, y_obs)
    μθ_hat_vals.append(μθ_hat.item())

def posterior_mean(σθ_val):
    μ_i, Σ_i, _ = construct_moments_IQ(n_scores, μθ_val, σθ_val, σy_val)
    mn_i = MultivariateNormal(μ_i, Σ_i)
    mn_i.partition(n_scores)
    μθ_hat, _ = mn_i.cond_dist(1, y_obs)
    return μθ_hat.item()

fig, ax = plt.subplots()
ax.semilogx(σθ_vals, μθ_hat_vals, 'o-',
            label=r'$\hat{\mu}_\theta$')
ax.axhline(y_bar,  ls='--', color='r',
            label=f'sample mean y_bar = {y_bar:.1f}')
ax.axhline(μθ_val, ls=':',  color='g',
            label=f'prior mean μθ = {μθ_val:.0f}')
ax.set_xlabel(r'$\sigma_\theta$')
ax.set_ylabel(r'posterior mean $\hat{\mu}_\theta$')
ax.legend()
plt.show()

σθ_small = 1e-2
σθ_large = 1e4

print(f"y_bar = {y_bar:.4f}")
print(f"Posterior mean with σθ={σθ_large:.0e}: {posterior_mean(σθ_large):.4f}")
print(f"Posterior mean with σθ={σθ_small:.0e}: {posterior_mean(σθ_small):.4f}")
print(f"Prior mean μθ = {μθ_val:.4f}")
```




```
y_bar = 107.6707
Posterior mean with σθ=1e+04: 107.6707
Posterior mean with σθ=1e-02: 100.0002
Prior mean μθ = 100.0000
```

Exercise 13.15.5

Kalman filter convergence

Using the `iterate` function from the Filtering Foundations section with

$$A = \begin{bmatrix} 0.9 & 0 \\ 0 & 0.5 \end{bmatrix}, \quad C = \begin{bmatrix} 1 \\ 1 \end{bmatrix}, \quad G = \begin{bmatrix} 1 & 0 \end{bmatrix}, \quad R = [1]$$

and initial conditions $\hat{x}_0 = [0, 0]'$, $\Sigma_0 = I_2$:

1. Simulate $T = 60$ periods of $\{x_t, y_t\}$ and run the filter.
2. Plot the sequences of conditional variances $\Sigma_t[0, 0]$ and $\Sigma_t[1, 1]$ over time.

   Verify that they converge to a steady state.
3. Plot the filtered state estimates $\tilde{x}_t[0]$ together with the true $x_t[0]$ and the raw observations y_t on a single figure.

Solution

Here is one solution:

```
A_ex = np.array([[0.9, 0.], [0., 0.5]])
C_ex = np.array([[1.], [1.]])
G_ex = np.array([[1., 0.]])
R_ex = np.array([[1.]])

T_ex = 60
x0_hat_ex = np.zeros(2)
Σ0_ex = np.eye(2)

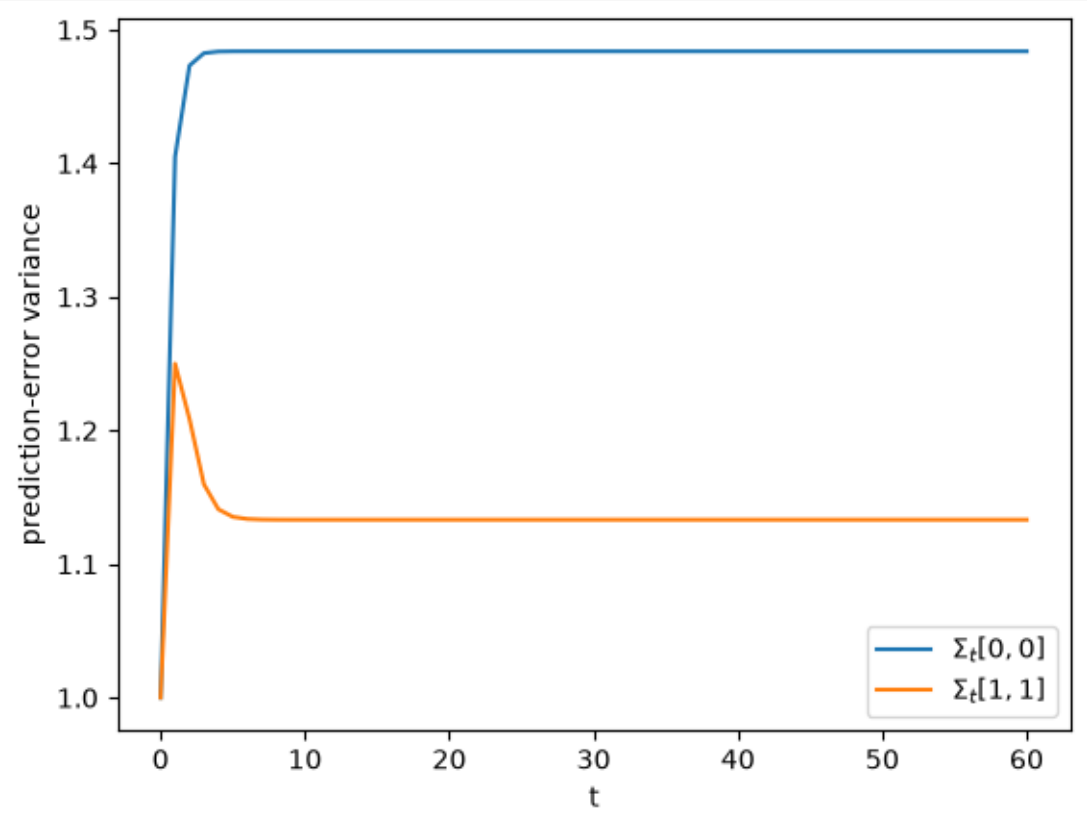
rng = np.random.default_rng(7)
x_true = np.zeros((T_ex + 1, 2))
y_seq_ex = np.zeros(T_ex)
for t in range(T_ex):
    x_true[t + 1] = A_ex @ x_true[t] + C_ex[:, 0] * rng.standard_normal()
    y_seq_ex[t] = (G_ex @ x_true[t]).item() + rng.standard_normal()

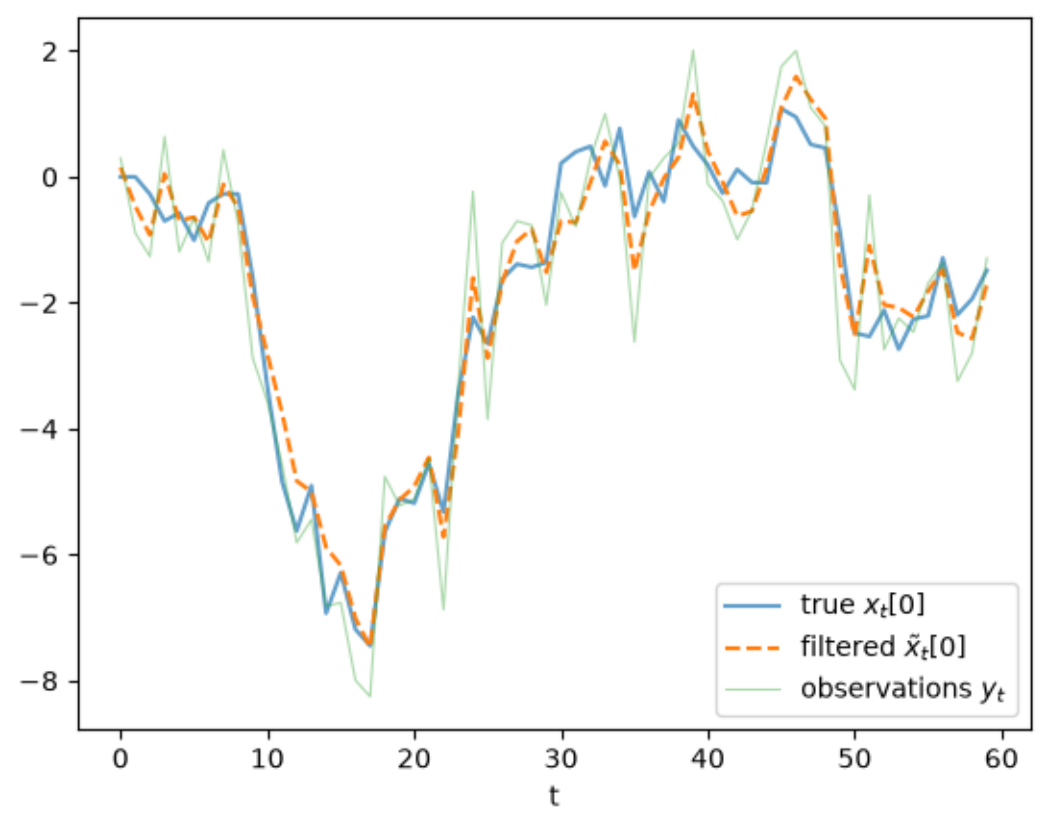
x_hat_seq, Σ_hat_seq = iterate(
    x0_hat_ex, Σ0_ex, A_ex, C_ex, G_ex, R_ex, y_seq_ex)

# x_hat_seq[t] = E[x_t | y^{t-1}] (one-step-ahead prediction)
# Σ_hat_seq[t] = corresponding prediction-error covariance
fig, ax = plt.subplots()
ax.plot(Σ_hat_seq[:, 0, 0], label=r'$\Sigma_t[0,0]$')
ax.plot(Σ_hat_seq[:, 1, 1], label=r'$\Sigma_t[1,1]$')
ax.set_xlabel('t')
ax.set_ylabel('prediction-error variance')
ax.legend()
plt.show()

# The `iterate` function stores one-step-ahead predictions.
# We recover the filtered estimates E[x_t | y^t] by re-applying
# the measurement-update step at each t.
n_state = 2
x_filt_seq = np.empty((T_ex, n_state))
for t in range(T_ex):
    xt_hat = x_hat_seq[t]
    Σt     = Σ_hat_seq[t]
    μ_k = np.hstack([xt_hat, G_ex @ xt_hat])
    Σ_k = np.block([[Σt,          Σt  @ G_ex.T          ],
                    [G_ex @ Σt,   G_ex @ Σt @ G_ex.T + R_ex]])
    mn_k = MultivariateNormal(μ_k, Σ_k)
    mn_k.partition(n_state)
    x_filt_seq[t], _ = mn_k.cond_dist(0, y_seq_ex[t:t+1])

fig, ax = plt.subplots()
ax.plot(x_true[:-1, 0], label='true $x_t[0]$', alpha=0.7)
ax.plot(x_filt_seq[:, 0], label=r'filtered $\tilde{x}_t[0]$', ls='--')
ax.plot(y_seq_ex, label='observations $y_t$', alpha=0.4, lw=0.8)
ax.set_xlabel('t')
ax.legend()
plt.show()
```


prediction-error variance
t
$\Sigma_t[0,0]$
$\Sigma_t[1,1]$


t
true $x_t[0]$
filtered $\tilde{x}_t[0]$
observations y_t

Exercise 13.15.6

PCA vs. factor analysis

In the classic factor analysis model at the end of the lecture the true covariance is $\Sigma_y = \Lambda\Lambda' + D$.

1. Set $\sigma_u = 2$ (instead of 0.5).
   - Recompute the fraction of variance explained by the first two principal components and compare it with the $\sigma_u = 0.5$ result.
   - Explain the change.
2. Show that the observation-space factor-analytic posterior $\Lambda E[f \mid Y] = \Lambda BY$ (an N-vector) is **not** equal to the two-component PCA reconstruction $\hat{Y} = P_{:,1:2}\,\epsilon_{1:2}$ (also an N-vector).
   - Plot both on the same axes.

   Note: $E[f \mid Y] = BY$ is a k-vector and $\hat{Y}$ is an N-vector, so they cannot be compared directly; the comparison must be made in observation space via $\Lambda E[f \mid Y]$.
3. In one or two sentences, explain why PCA is misspecified for factor-analytic data.

Solution

Here is one solution:

```
rng = np.random.default_rng(42)

N_fa = 10
k_fa = 2

Λ_fa = np.zeros((N_fa, k_fa))
Λ_fa[:N_fa//2, 0] = 1
Λ_fa[N_fa//2:, 1] = 1

for σu_val in [0.5, 2.0]:
    D_fa = np.eye(N_fa) * σu_val ** 2
    Σy_fa = Λ_fa @ Λ_fa.T + D_fa

    λ_fa, P_fa = np.linalg.eigh(Σy_fa)
    ind_fa = sorted(range(N_fa), key=lambda x: λ_fa[x], reverse=True)
    P_fa   = P_fa[:, ind_fa]
    λ_fa   = λ_fa[ind_fa]

    frac = λ_fa[:2].sum() / λ_fa.sum()
    print(f"σu={σu_val}: fraction explained by first 2 PCs = {frac:.4f}")

σu_b = 0.5
D_b = np.eye(N_fa) * σu_b ** 2
Σy_b = Λ_fa @ Λ_fa.T + D_b

μz_b = np.zeros(k_fa + N_fa)
Σz_b = np.block([[np.eye(k_fa), Λ_fa.T], [Λ_fa, Σy_b]])
z_b = rng.multivariate_normal(μz_b, Σz_b)
f_b = z_b[:k_fa]
y_b = z_b[k_fa:]

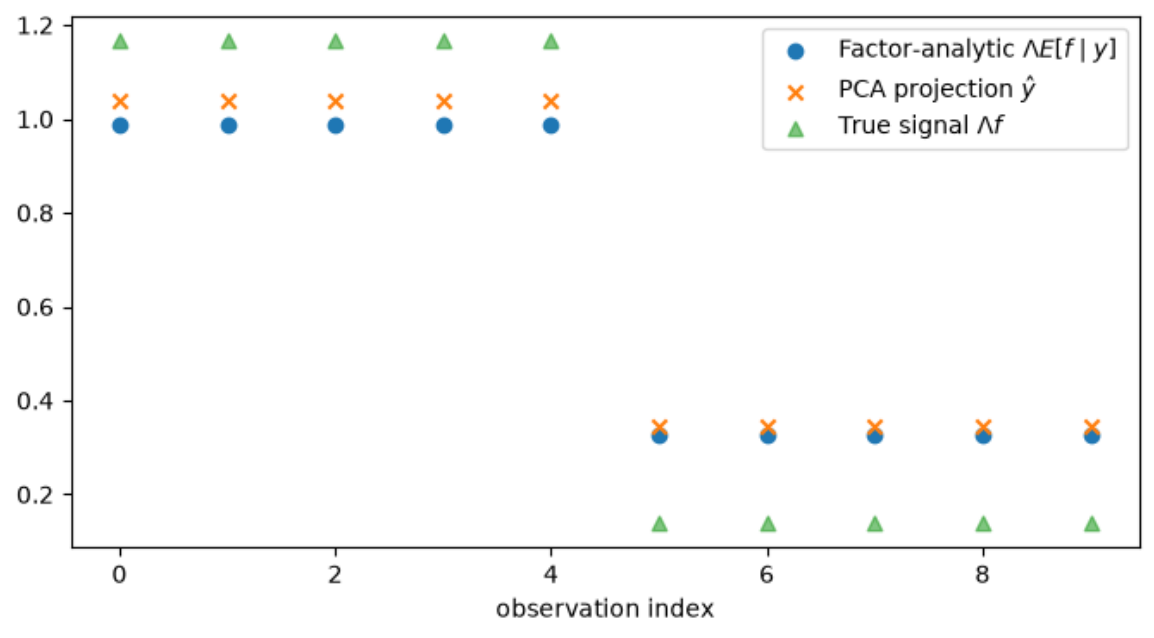
B_b = Λ_fa.T @ np.linalg.inv(Σy_b)
```

```
Efy_b  = B_b @ y_b

λ_b, P_b = np.linalg.eigh(Σy_b)
ind_b = sorted(range(N_fa), key=lambda x: λ_b[x], reverse=True)
P_b = P_b[:, ind_b]
ε_b = P_b.T @ y_b
y_hat_b = P_b[:, :2] @ ε_b[:2]

fig, ax = plt.subplots(figsize=(8, 4))
ax.scatter(range(N_fa),
        Λ_fa @ Efy_b, label=r'Factor-analytic $\Lambda E[f\mid y]$')
ax.scatter(range(N_fa),
        y_hat_b, marker='x', label=r'PCA projection $\hat{y}$')
ax.scatter(range(N_fa),
        Λ_fa @ f_b, marker='^', alpha=0.6, label=r'True signal $\Lambda f$')
ax.set_xlabel('observation index')
ax.legend()
plt.show()
```

```
σu=0.5: fraction explained by first 2 PCs = 0.8400
σu=2.0: fraction explained by first 2 PCs = 0.3600
```



In this symmetric example, PCA does recover the same two-dimensional observation-space subspace as the factor model, namely the column space of Λ. But PCA is still misspecified for factor-analytic data, because it treats the covariance matrix as an arbitrary matrix to be approximated and does not use the special decomposition $\Sigma_y = \Lambda\Lambda^\top + D$ into a common part and an idiosyncratic noise part.

So the two methods are solving different problems. PCA forms $\hat{Y}$ as the best rank-2 approximation to the observed data vector Y, which in this example amounts to using the block means. The factor model instead computes $\Lambda E[f \mid Y]$, the conditional mean of the latent common component Λf given the data, and because it accounts for noise it shrinks those block means toward zero.

CHAPTER

FOURTEEN

FAULT TREE UNCERTAINTIES

Contents

- *Fault Tree Uncertainties*
  - *Overview*
  - *The lognormal distribution*
  - *The convolution theorem*
  - *Approximating continuous distributions*
  - *Discretizing the lognormal distribution*
  - *Convolving probability mass functions*
  - *Fault tree analysis*
  - *Failure rates unknown*
  - *Application: waste hoist failure rate*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon tabulate
```

14.1 Overview

This lecture puts elementary tools to work to approximate probability distributions of the annual failure rates of a system consisting of a number of critical parts.

We'll use log normal distributions to approximate probability distributions of critical component parts.

To approximate the probability distribution of the *sum* of n lognormal random variables (representing the system's total failure rate), we compute the convolution of these distributions.

We'll use the following concepts and tools:

- lognormal distributions
- the convolution theorem that describes the probability distribution of the sum of independent random variables
- fault tree analysis for approximating a failure rate of a multi-component system

- a hierarchical probability model for describing uncertain probabilities
- Fourier transforms and inverse Fourier transforms as efficient ways of computing convolutions of sequences

See also

For more on Fourier transforms, see *Circulant Matrices* as well as Covariance Stationary Processes and Estimation of Spectra.

El-Shanawany *et al.* [2018] and Greenfield and Sargent [1993] applied these methods to approximate failure probabilities of safety systems in nuclear facilities.

These techniques respond to recommendations by Apostolakis [1990] for quantifying uncertainty in safety system reliability.

We will use the following imports and settings throughout this lecture:

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.signal import fftconvolve
from tabulate import tabulate
import quantecon as qe
```

14.2 The lognormal distribution

If random variable x follows a normal distribution with mean μ and variance σ^2, then $y = \exp(x)$ follows a **lognormal distribution** with parameters μ, σ^2.

Note

We refer to μ and σ^2 as *parameters* rather than mean and variance because:

- μ and σ^2 are the mean and variance of $x = \log(y)$
- They are *not* the mean and variance of y
- The mean of y is $\exp(\mu + \frac{1}{2}\sigma^2)$ and the variance is $(e^{\sigma^2} - 1)e^{2\mu+\sigma^2}$

A lognormal random variable y is always nonnegative.

The probability density function for y is

$$f(y) = \frac{1}{y\sigma\sqrt{2\pi}} \exp\left(\frac{-(\log y - \mu)^2}{2\sigma^2}\right), \quad y \geq 0 \tag{14.1}$$

Important properties of a lognormal random variable are:

$$\begin{aligned}
\text{Mean:} &\quad e^{\mu+\frac{1}{2}\sigma^2} \\
\text{Variance:} &\quad (e^{\sigma^2} - 1)e^{2\mu+\sigma^2} \\
\text{Median:} &\quad e^{\mu} \\
\text{Mode:} &\quad e^{\mu-\sigma^2} \\
\text{0.95 quantile:} &\quad e^{\mu+1.645\sigma} \\
\text{0.95/0.05 quantile ratio:} &\quad e^{3.29\sigma}
\end{aligned} \tag{14.2}$$

14.2.1 Stability properties

Recall that independent normally distributed random variables have the following stability property:

If $x_1 \sim N(\mu_1, \sigma_1^2)$ and $x_2 \sim N(\mu_2, \sigma_2^2)$ are independent, then $x_1 + x_2 \sim N(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$.

Independent lognormal distributions have a different stability property: the *product* of independent lognormal random variables is also lognormal.

Specifically, if y_1 is lognormal with parameters (μ_1, σ_1^2) and y_2 is lognormal with parameters (μ_2, σ_2^2), then $y_1 y_2$ is lognormal with parameters $(\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2)$.

Warning

While the product of two lognormal distributions is lognormal, the *sum* of two lognormal distributions is *not* lognormal.

This observation motivates the central challenge of this lecture: approximating the probability distribution of *sums* of independent lognormal random variables.

14.3 The convolution theorem

Let x and y be independent random variables with probability densities $f(x)$ and $g(y)$, where $x, y \in \mathbb{R}$.

Let $z = x + y$.

Then the probability density of z is

$$h(z) = (f * g)(z) \equiv \int_{-\infty}^{\infty} f(\tau) g(z - \tau) d\tau \tag{14.3}$$

where $(f * g)$ denotes the **convolution** of f and g.

For nonnegative random variables, this specializes to

$$h(z) = (f * g)(z) \equiv \int_{0}^{z} f(\tau) g(z - \tau) d\tau \tag{14.4}$$

14.3.1 Discrete convolution

We will use a discretized version of the convolution formula.

We replace both f and g with discretized counterparts, normalized to sum to 1.

The discrete convolution formula is

$$h_n = (f * g)_n = \sum_{m=0}^{n} f_m g_{n-m}, \quad n \geq 0 \tag{14.5}$$

This computes the probability mass function of the sum of two discrete random variables.

14.3.2 Example: discrete distributions

Consider two probability mass functions:

$$f_j = \mathbb{P}\{X = j\}, \quad j = 0, 1$$

and

$$g_j = \mathbb{P}\{Y = j\}, \quad j = 0, 1, 2, 3$$

The distribution of $Z = X + Y$ is given by the convolution $h = f * g$.

```
# Define probability mass functions
f = [0.75, 0.25]
g = [0.0, 0.6, 0.0, 0.4]

# Compute convolution using two methods
h = np.convolve(f, g)
hf = fftconvolve(f, g)

print(f"f = {f}, sum = {np.sum(f):.3f}")
print(f"g = {g}, sum = {np.sum(g):.3f}")
print(f"h = {h}, sum = {np.sum(h):.3f}")
print(f"hf = {hf}, sum = {np.sum(hf):.3f}")
```

```
f = [0.75, 0.25], sum = 1.000
g = [0.0, 0.6, 0.0, 0.4], sum = 1.000
h = [0.   0.45 0.15 0.3  0.1 ], sum = 1.000
hf = [0.   0.45 0.15 0.3  0.1 ], sum = 1.000
```

Both `numpy.convolve` and `scipy.signal.fftconvolve` produce the same result, but `fftconvolve` is much faster for long sequences.

We will use `fftconvolve` throughout this lecture for efficiency.

14.4 Approximating continuous distributions

We now verify that discretized distributions can accurately approximate samples from underlying continuous distributions.

We generate samples of size 25,000 from three independent lognormal random variables and compute their pairwise and triple-wise sums.

We then compare histograms of the samples with histograms of the discretized distributions.

```
# Set parameters for lognormal distributions
μ, σ = 5.0, 1.0
n_samples = 25000

# Generate samples
rng = np.random.default_rng(1234)
s1 = rng.lognormal(μ, σ, n_samples)
s2 = rng.lognormal(μ, σ, n_samples)
s3 = rng.lognormal(μ, σ, n_samples)

# Compute sums
ssum2 = s1 + s2
```

(continues on next page)

(continued from previous page)

```
ssum3 = s1 + s2 + s3

# Plot histogram of s1
fig, ax = plt.subplots()
ax.hist(s1, 1000, density=True, alpha=0.6)
ax.set_xlabel('value')
ax.set_ylabel('density')
plt.show()
```

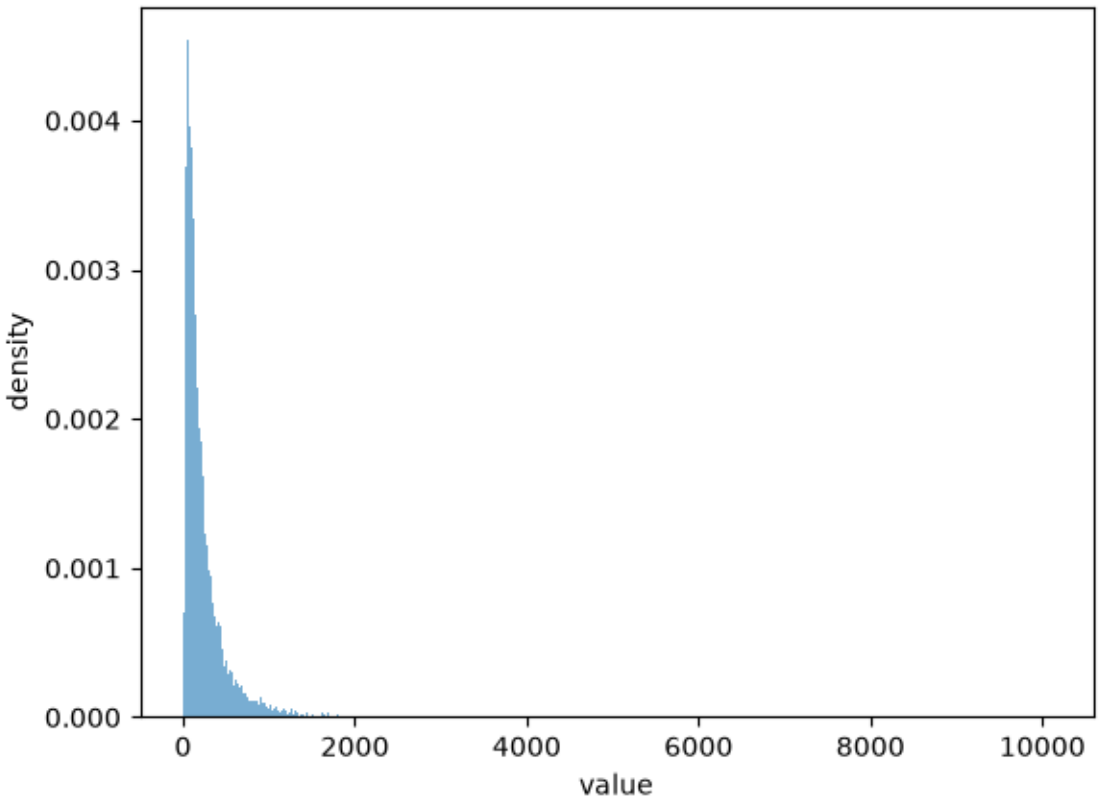


Fig. 14.1: Sample histogram of one lognormal

```
# Plot histogram of sum of two lognormal distributions
fig, ax = plt.subplots()
ax.hist(ssum2, 1000, density=True, alpha=0.6)
ax.set_xlabel('value')
ax.set_ylabel('density')
plt.show()
```

```
# Plot histogram of sum of three lognormal distributions
fig, ax = plt.subplots()
ax.hist(ssum3, 1000, density=True, alpha=0.6)
ax.set_xlabel('value')
```

(continues on next page)

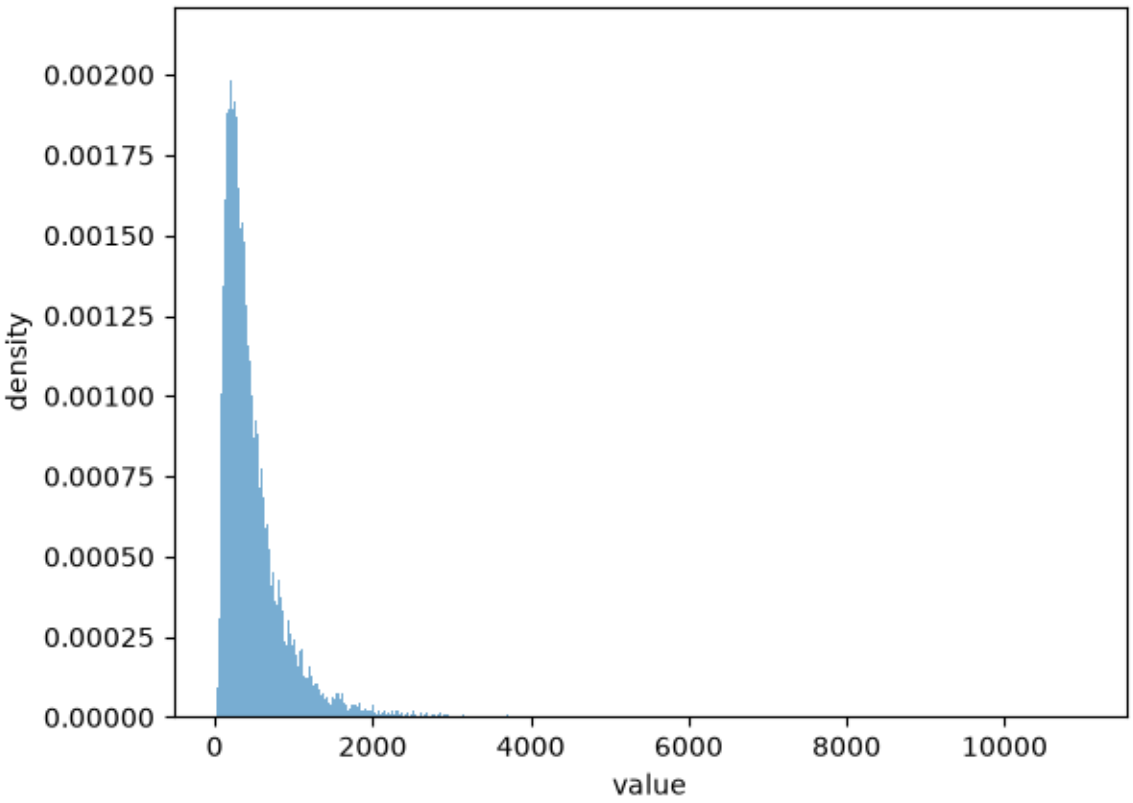


Fig. 14.2: Histogram of a sum of two lognormals

(continued from previous page)

```
ax.set_ylabel('density')
plt.show()
```

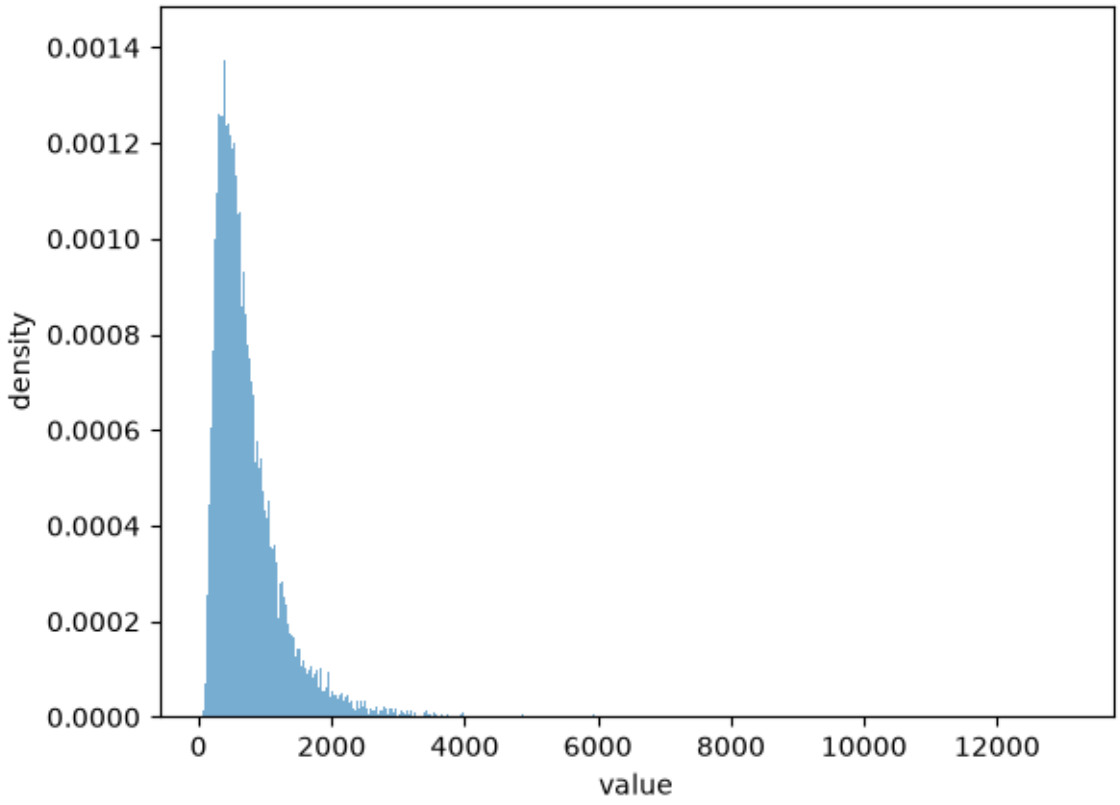


Fig. 14.3: Histogram of a sum of three lognormals

Let's verify that the sample mean matches the theoretical mean:

```
samp_mean = np.mean(s2)
theoretical_mean = np.exp(μ + σ**2 / 2)

print(f"Theoretical mean: {theoretical_mean:.3f}")
print(f"Sample mean: {samp_mean:.3f}")
```

```
Theoretical mean: 244.692
Sample mean: 243.197
```

14.5 Discretizing the lognormal distribution

We define helper functions to create discretized versions of lognormal probability density functions.

We write out the density by hand to keep the formula (14.1) in view; `scipy.stats.lognorm(s=σ, scale=np.exp(μ)).pdf(x)` computes the same thing.

```
def lognormal_pdf(x, μ, σ):
    """
    Compute lognormal probability density function.
    """
    p = 1 / (σ * x * np.sqrt(2 * np.pi)) \
            * np.exp(-0.5 * ((np.log(x) - μ) / σ)**2)
    return p


def discretize_lognormal(μ, σ, I, m):
    """
    Discretize a lognormal distribution on the grid 0, m, 2m, ..., up to I.

    Parameters
    ----------
    μ, σ : parameters of the lognormal distribution
    I    : upper end of the grid, which truncates the right tail
    m    : spacing between grid points, which sets the resolution

    Returns
    -------
    p_array      : the density evaluated on the grid
    p_array_norm : the implied probability mass function, summing to one
    x            : the grid itself, with I / m points
    """
    x = np.arange(1e-7, I, m)
    p_array = lognormal_pdf(x, μ, σ)
    p_array_norm = p_array / np.sum(p_array)
    return p_array, p_array_norm, x
```

Two separate choices govern the quality of this approximation, and it pays to keep them straight.

- I fixes where the grid *stops*, so it controls how much of the right tail we throw away
- m fixes the *spacing* between grid points, so it controls resolution

The grid has I/m points, so raising I at fixed m buys range, while lowering m at fixed I buys accuracy.

Once I is large enough that almost no probability mass lies beyond it, further increases change nothing, and only m matters.

Exercise 14.10.1 asks you to verify this.

Note

`scipy.signal.fftconvolve` pads its inputs to a convenient length internally, so there is no need to choose I/m to be a power of two.

```
# Set grid parameters
p = 15
```

(continues on next page)

(continued from previous page)

```
I = 2**p  # where the grid stops: truncates the right tail
m = 0.1   # spacing between grid points: sets the resolution
```

Let's visualize how well the discretized distribution approximates the continuous lognormal distribution:

```
# Compute discretized PDF
pdf, pdf_norm, x = discretize_lognormal(μ, σ, I, m)

# Plot discretized PDF against histogram
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(x, pdf, 'r-', lw=2, label='discretized PDF')
ax.hist(s1, 1000, density=True, alpha=0.6, label='sample histogram')
ax.set_xlim(0, 2500)
ax.set_xlabel('value')
ax.set_ylabel('density')
ax.legend()
plt.show()
```

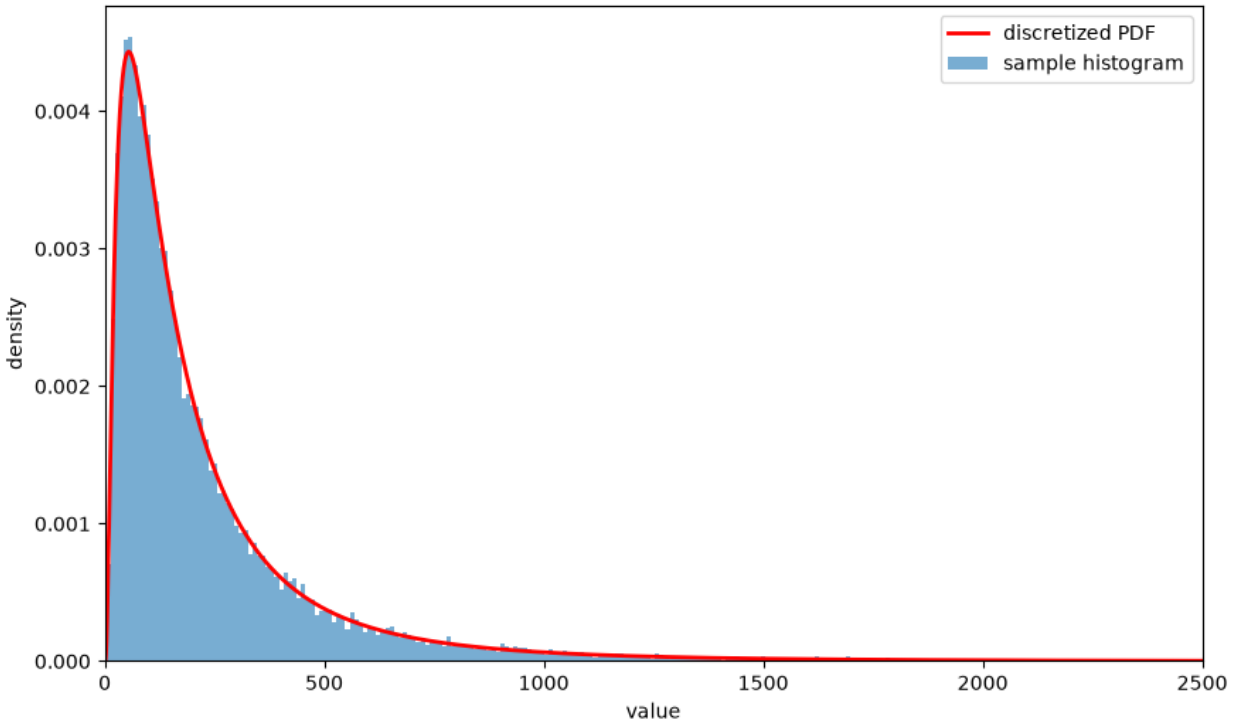


Fig. 14.4: Discretized density against sample

Now let's verify that the discretized distribution has the correct mean:

```
# Compute mean from discretized PDF
mean_discrete = np.sum(x * pdf_norm)
mean_theory = np.exp(μ + 0.5 * σ**2)

print(f"Theoretical mean: {mean_theory:.3f}")
print(f"Discretized mean: {mean_discrete:.3f}")
```

```
Theoretical mean: 244.692
Discretized mean: 244.691
```

14.6 Convolving probability mass functions

Now let's use the convolution theorem to compute the probability distribution of a sum of the two lognormal random variables we have parameterized above.

We'll also compute the probability distribution of a sum of three log normal distributions constructed above.

For long sequences, `scipy.signal.fftconvolve` is much faster than `numpy.convolve` because it uses fast Fourier transforms.

Let's define the Fourier transform and the inverse Fourier transform first

14.6.1 The fast Fourier transform

The **Fourier transform** of a sequence $\{x_t\}_{t=0}^{T-1}$ is

$$x(\omega_j) = \sum_{t=0}^{T-1} x_t \exp(-i\omega_j t) \tag{14.6}$$

where $\omega_j = \frac{2\pi j}{T}$ for $j = 0, 1, \ldots, T-1$.

The **inverse Fourier transform** of the sequence $\{x(\omega_j)\}_{j=0}^{T-1}$ is

$$x_t = T^{-1} \sum_{j=0}^{T-1} x(\omega_j) \exp(i\omega_j t) \tag{14.7}$$

The sequences $\{x_t\}_{t=0}^{T-1}$ and $\{x(\omega_j)\}_{j=0}^{T-1}$ contain the same information.

The pair of equations (14.6) and (14.7) tell how to recover one series from its Fourier partner.

The program `scipy.signal.fftconvolve` deploys the theorem that a convolution of two sequences $\{f_k\}, \{g_k\}$ can be computed in the following way:

- Compute Fourier transforms $F(\omega), G(\omega)$ of the $\{f_k\}$ and $\{g_k\}$ sequences, respectively
- Form the product $H(\omega) = F(\omega)G(\omega)$
- The convolution of $f * g$ is the inverse Fourier transform of $H(\omega)$

The **fast Fourier transform** and the associated **inverse fast Fourier transform** execute these calculations very quickly.

This is the algorithm used by `fftconvolve`.

Let's do a warmup calculation that compares the times taken by `numpy.convolve` and `scipy.signal.fftconvolve`

Our three components are identically distributed, so a single discretization serves for all of them.

```
# Discretize the lognormal distribution; the three components are IID
_, pmf1, x = discretize_lognormal(μ, σ, I, m)
pmf2 = pmf3 = pmf1

# Direct convolution costs O(N²), so we time it on a short prefix
```

(continues on next page)

(continued from previous page)

```
short = pmf1[:20_000]

with qe.Timer() as timer_numpy:
    np.convolve(short, short)
time_numpy = timer_numpy.elapsed

with qe.Timer() as timer_fft:
    fftconvolve(short, short)
time_fft = timer_fft.elapsed

print(f"On {len(short):,} points:")
print(f"  np.convolve: {time_numpy:.4f} seconds")
print(f"  fftconvolve: {time_fft:.4f} seconds")
print(f"  speedup:     {time_numpy / time_fft:.0f}x")
```

```
0.0663 seconds elapsed
0.0027 seconds elapsed
On 20,000 points:
  np.convolve: 0.0663 seconds
  fftconvolve: 0.0027 seconds
  speedup:     25x
```

The gap widens rapidly with the length of the sequences, because direct convolution costs $O(N^2)$ operations while the FFT approach costs $O(N \log N)$.

On the full grid used below, the direct method would be far slower still.

```
# The full calculation, done the fast way
conv_fft = fftconvolve(fftconvolve(pmf1, pmf2), pmf3)
print(f"grid points per component: {len(pmf1):,}")
```

```
grid points per component: 327,680
```

Now let's plot our computed probability mass function approximation for the sum of two log normal random variables against the histogram of the sample that we formed above

```
# Compute convolution of two distributions for comparison
conv2 = fftconvolve(pmf1, pmf2)

fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(x, conv2[:len(x)] / m, 'r-', lw=2, label='convolution (FFT)')
ax.hist(ssum2, 1000, density=True, alpha=0.6, label='sample histogram')
ax.set_xlim(0, 5000)
ax.set_xlabel('value')
ax.set_ylabel('density')
ax.legend()
plt.show()
```

Now we present the plot for the sum of three lognormal random variables:

```
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(x, conv_fft[:len(x)] / m, 'r-', lw=2, label='convolution (FFT)')
ax.hist(ssum3, 1000, density=True, alpha=0.6, label='sample histogram')
ax.set_xlim(0, 5000)
```

(continues on next page)

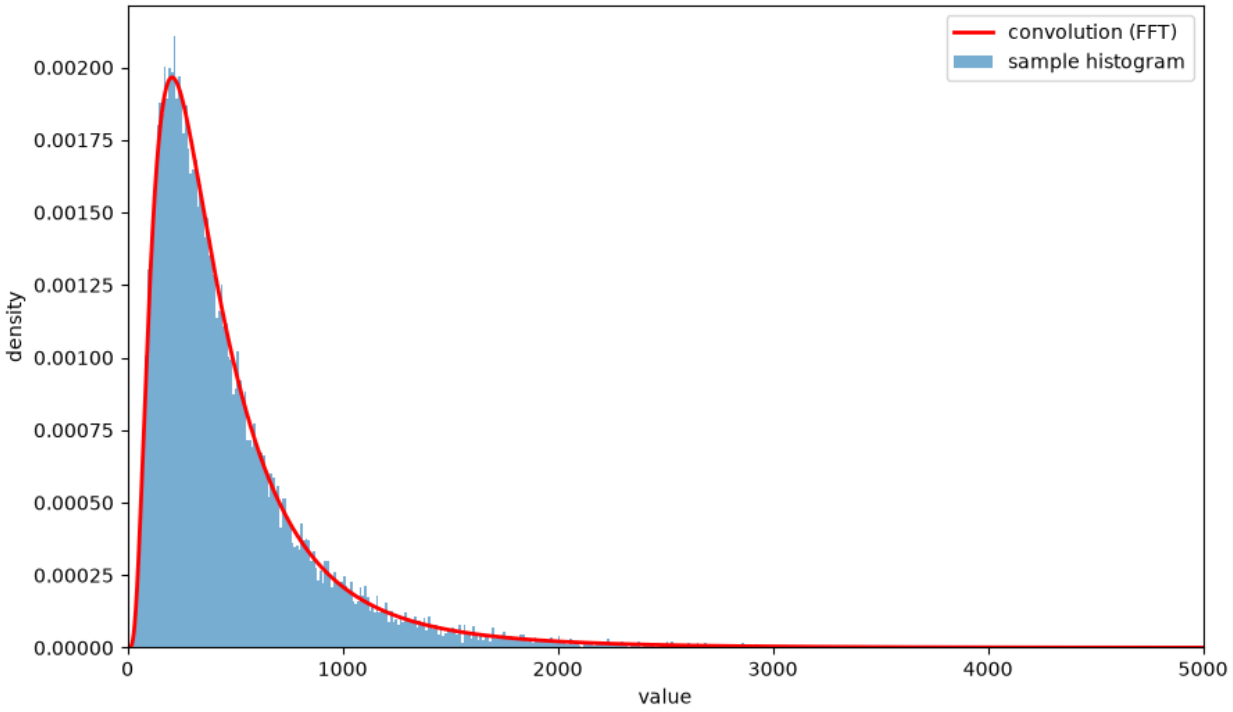


Fig. 14.5: Convolution against sample, two components

(continued from previous page)

```
ax.set_xlabel('value')
ax.set_ylabel('density')
ax.legend()
plt.show()
```

Let's verify that the means are correct

```
# Mean of sum of two distributions
mean_conv2 = np.sum(x * conv2[:len(x)])
mean_theory2 = 2 * np.exp(μ + 0.5 * σ**2)

print(f"Sum of two distributions:")
print(f"  Theoretical mean: {mean_theory2:.3f}")
print(f"  Computed mean: {mean_conv2:.3f}")
```

```
Sum of two distributions:
  Theoretical mean: 489.384
  Computed mean: 489.381
```

```
# Mean of sum of three distributions
mean_conv3 = np.sum(x * conv_fft[:len(x)])
mean_theory3 = 3 * np.exp(μ + 0.5 * σ**2)

print(f"Sum of three distributions:")
print(f"  Theoretical mean: {mean_theory3:.3f}")
print(f"  Computed mean: {mean_conv3:.3f}")
```

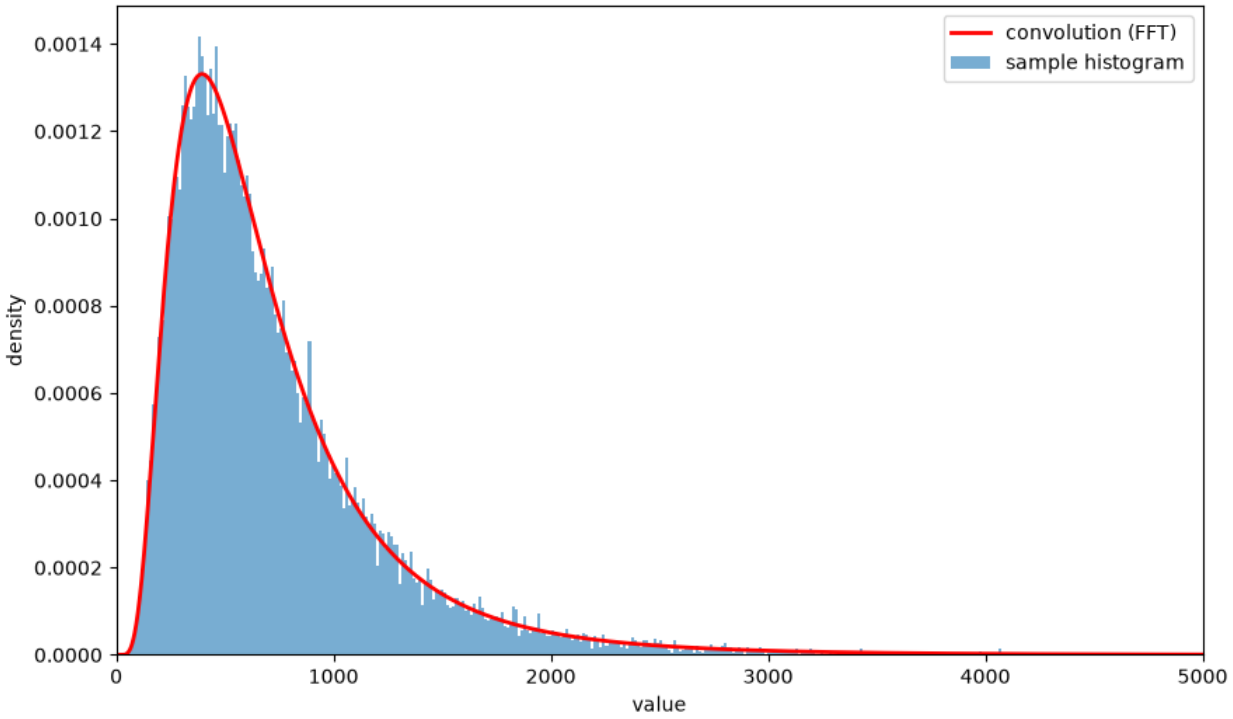


Fig. 14.6: Convolution against sample, three components

```
Sum of three distributions:
  Theoretical mean: 734.076
  Computed mean: 734.071
```

14.7 Fault tree analysis

We shall soon apply the convolution theorem to compute the probability of a **top event** in a failure tree analysis.

Before applying the convolution theorem, we first describe the model that connects constituent events to the *top event* whose failure rate we seek to quantify.

Fault tree analysis is a widely used technique for assessing system reliability, as described by El-Shanawany *et al.* [2018].

To construct the statistical model, we repeatedly use what is called the **rare event approximation**.

14.7.1 The rare event approximation

We want to compute the probability of an event $A \cup B$.

For events A and B, the probability of the union is

$$P(A \cup B) = P(A) + P(B) - P(A \cap B)$$

where $A \cup B$ is the event that A **or** B occurs, and $A \cap B$ is the event that A **and** B both occur.

If A and B are independent, then $P(A \cap B) = P(A)P(B)$.

When $P(A)$ and $P(B)$ are both small, $P(A)P(B)$ is even smaller.

The **rare event approximation** is

$$P(A \cup B) \approx P(A) + P(B)$$

This approximation is widely used in system failure analysis.

14.7.2 System failure probability

Consider a system with n critical components where system failure occurs when *any* component fails.

We assume:

- The failure probability $P(A_i)$ of each component A_i is small
- Component failures are statistically independent

We repeatedly apply a **rare event approximation** to obtain the following formula for the probability of a system failure:

$$P(F) \approx P(A_1) + P(A_2) + \cdots + P(A_n)$$

or

$$P(F) \approx \sum_{i=1}^{n} P(A_i) \tag{14.8}$$

where $P(F)$ is the system failure probability.

Probabilities for each event are recorded as failure rates per year.

> **Note**
>
> Strictly speaking, a failure *rate* per year and a failure *probability* within a year are different objects.
>
> For rare events they nearly coincide, because $1 - e^{-\lambda} \approx \lambda$ when λ is small.
>
> The same approximation that lets us add probabilities across components also lets us move between rates and probabilities, so we follow the reliability literature in using the two words interchangeably here.

14.8 Failure rates unknown

Now we come to the problem that really interests us, following El-Shanawany *et al.* [2018] and Greenfield and Sargent [1993] in the spirit of Apostolakis [1990].

The component failure rates $P(A_i)$ are not known precisely and must be estimated.

We address this problem by specifying **probabilities of probabilities** that capture one notion of not knowing the constituent probabilities that are inputs into a failure tree analysis.

Thus, we assume that a system analyst is uncertain about the failure rates $P(A_i), i = 1, \ldots, n$ for components of a system.

The analyst copes with this situation by regarding the system's failure probability $P(F)$ and each of the component probabilities $P(A_i)$ as random variables.

- dispersions of the probability distribution of $P(A_i)$ characterizes the analyst's uncertainty about the failure probability $P(A_i)$

- the dispersion of the implied probability distribution of $P(F)$ characterizes his uncertainty about the probability of a system's failure.

This leads to what is sometimes called a **hierarchical** model in which the analyst has probabilities about the probabilities $P(A_i)$.

Note

Two distinct kinds of randomness appear in this model, and it is worth keeping them apart.

Aleatory uncertainty is the randomness in whether a component fails during a given year; it is described by the failure rate $P(A_i)$.

Epistemic uncertainty is the analyst's ignorance about the value of that rate; it is described by the lognormal distribution that he places over $P(A_i)$.

The distribution that we compute below is an epistemic object: it describes what the analyst knows about a failure rate, not how often the system fails.

Separating the two is the central recommendation of Apostolakis [1990].

The analyst formalizes his uncertainty by assuming that

- the failure probability $P(A_i)$ is itself a log normal random variable with parameters (μ_i, σ_i).
- failure rates $P(A_i)$ and $P(A_j)$ are statistically independent for all pairs with $i \neq j$.

The analyst calibrates the parameters (μ_i, σ_i) for the failure events $i = 1, \ldots, n$ by reading reliability studies in engineering papers that have studied historical failure rates of components that are as similar as possible to the components being used in the system under study.

The analyst assumes that such information about the observed dispersion of annual failure rates, or times to failure, can inform him of what to expect about parts' performances in his system.

The analyst assumes that the random variables $P(A_i)$ are statistically mutually independent.

Warning

Independence is a strong assumption and it is the one that reliability analysts worry about most.

A design flaw, a shared power supply, a common maintenance crew, or a single environmental shock can push many components toward failure at once.

Such **common-cause** failures make the upper tail of the distribution of $P(F)$ much fatter than the independent calculation suggests, which is precisely the region that a safety regulator cares about.

Exercise 14.10.5 quantifies how much difference this makes.

The analyst wants to approximate a probability mass function and cumulative distribution function of the system's failure probability $P(F)$.

- We say probability mass function because of how we discretize each random variable, as described earlier.

The analyst calculates the probability mass function for the *top event* F, i.e., a *system failure*, by repeatedly applying the convolution theorem to compute the probability distribution of a sum of independent log normal random variables, as described in equation (14.8).

14.9 Application: waste hoist failure rate

We now analyze a real-world example with $n = 14$ components.

The application estimates the annual failure rate of a critical hoist at a nuclear waste facility.

A regulatory agency requires the system to be designed so that the top event failure rate is small with high probability.

14.9.1 Model specification

This example is Design Option B-2 (Case I) described in Table 10 on page 27 of Greenfield and Sargent [1993].

The table describes parameters μ_i, σ_i for fourteen log normal random variables that consist of **seven pairs** of random variables that are identically and independently distributed.

- Within a pair, parameters μ_i, σ_i are the same
- As described in table 10 of Greenfield and Sargent [1993] p. 27, parameters of log normal distributions for the seven unique probabilities $P(A_i)$ have been calibrated to be the values in the following Python code:

```
# Component failure rate parameters
# (see Table 10 of Greenfield & Sargent 1993)
params = [
    (4.28, 1.1947),    # Component type 1
    (3.39, 1.1947),    # Component type 2
    (2.795, 1.1947),   # Component type 3
    (2.717, 1.1947),   # Component type 4
    (2.717, 1.1947),   # Component type 5
    (1.444, 1.4632),   # Component type 6
    (-0.040, 1.4632), # Component type 7 (appears 8 times)
]
```

Note

Since failure rates are very small, these lognormal distributions actually describe $P(A_i) \times 10^{-9}$.

So the probabilities that we'll put on the x axis of the probability mass function and associated cumulative distribution function should be multiplied by 10^{-09}

We define a helper function to find array indices:

```
def find_nearest(array, value):
    """
    Index of the array element nearest to the given value.

    Applied to a cumulative distribution function, this returns the grid point
    whose cumulative probability is closest to a target, which for a finely
    discretized distribution is indistinguishable from the usual definition of
    a quantile as the smallest x with CDF(x) >= q.
    """
    array = np.asarray(array)
    idx = (np.abs(array - value)).argmin()
    return idx
```

We compute the required thirteen convolutions in the following code.

(Please feel free to try different values of the power parameter p that we use to set the number of points in our grid for constructing the probability mass functions that discretize the continuous log normal distributions.)

```
# Set grid parameters
p = 15
I = 2**p
m = 0.05

# Discretize all component failure rate distributions
# First 6 components use unique parameters, last 8 share the same parameters
component_pmfs = []
for μ, σ in params[:6]:
    _, pmf, x = discretize_lognormal(μ, σ, I, m)
    component_pmfs.append(pmf)

# Add 8 copies of component type 7
μ7, σ7 = params[6]
_, pmf7, x = discretize_lognormal(μ7, σ7, I, m)
component_pmfs.extend([pmf7] * 8)

# Compute system failure distribution via sequential convolution
with qe.Timer() as timer:
    system_pmf = component_pmfs[0]
    for pmf in component_pmfs[1:]:
        system_pmf = fftconvolve(system_pmf, pmf)

print(f"Time for 13 convolutions: {timer.elapsed:.4f} seconds")

# the convolution lives on the same grid spacing, but extends much further
system_grid = np.arange(len(system_pmf)) * m
print(f"grid points in the answer: {len(system_pmf):,}")
```

```
7.7865 seconds elapsed
Time for 13 convolutions: 7.7865 seconds
grid points in the answer: 9,175,027
```

Before plotting the cumulative distribution function, let's look at the density itself.

```
fig, ax = plt.subplots(figsize=(10, 6))
upper = 2000
ax.plot(system_grid[:int(upper/m)], system_pmf[:int(upper/m)] / m, 'b-', lw=2)
ax.set_xlabel(r'failure rate ($\times 10^{-9}$ per year)')
ax.set_ylabel('density')
plt.show()
```

The density is strongly skewed to the right: a long upper tail stretches far beyond the bulk of the distribution.

This asymmetry is what makes a single point estimate of a failure rate a poor summary, and it is why the analyst reports quantiles instead.

We now plot a counterpart to the cumulative distribution function (CDF) in figure 5 on page 29 of Greenfield and Sargent [1993]

```
# Compute cumulative distribution function
cdf = np.cumsum(system_pmf)
```

(continues on next page)

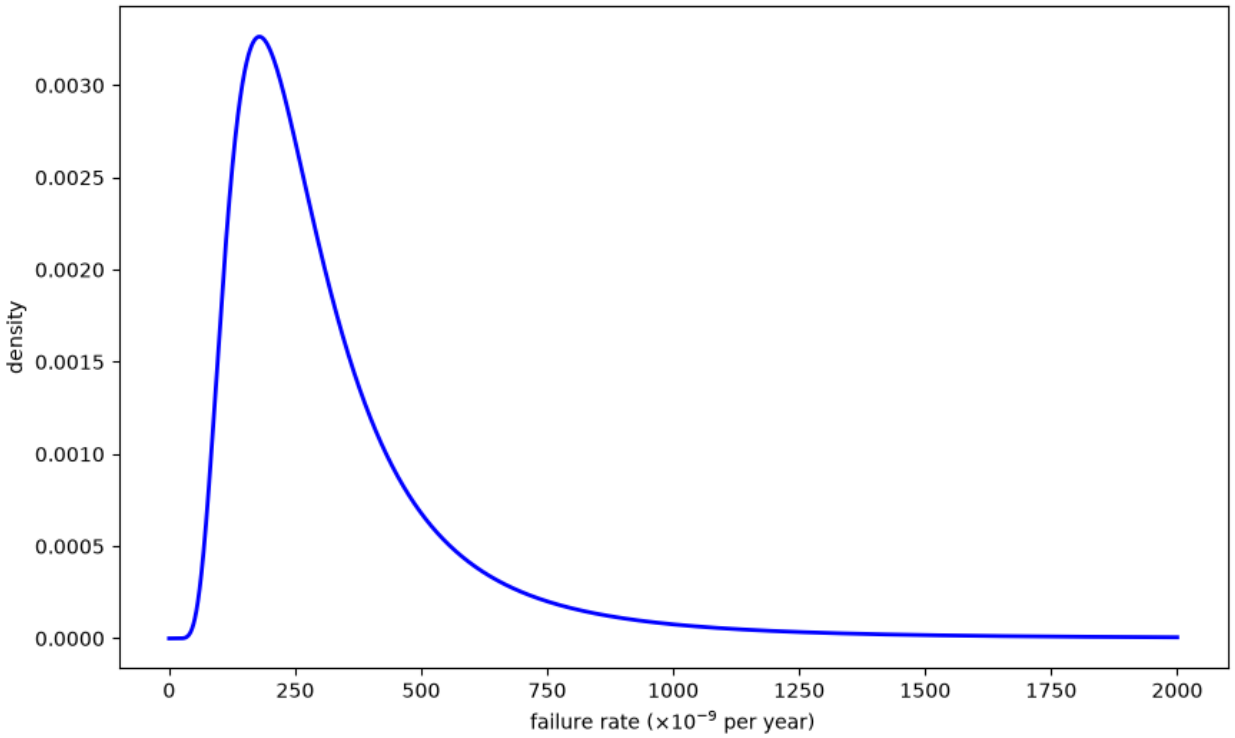


Fig. 14.7: Density of the system failure rate

(continued from previous page)

```
# Plot CDF
Nx = 1400
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(x[:int(Nx / m)], cdf[:int(Nx / m)], 'b-', lw=2)

# Add reference lines for key quantiles
quantile_levels = [0.05, 0.10, 0.50, 0.90, 0.95]
for q in quantile_levels:
    ax.axhline(q, color='gray', linestyle='--', alpha=0.5)

ax.set_xlim(0, Nx)
ax.set_ylim(0, 1)
ax.set_xlabel(r'failure rate ($\times 10^{-9}$ per year)')
ax.set_ylabel('cumulative probability')
plt.show()
```

We also present a counterpart to their Table 11 on page 28 of Greenfield and Sargent [1993], which lists key quantiles of the system failure rate distribution

```
# Percentiles reported in Table 11 of Greenfield and Sargent (1993),
# together with their published values, in units of 10^-9 per year
reference = {1.0: 77, 10.0: 130, 50.0: 263, 66.5: 341,
             85.0: 513, 95.0: 811, 99.0: 1480, 99.78: 2490}

table_data = []
```

(continues on next page)

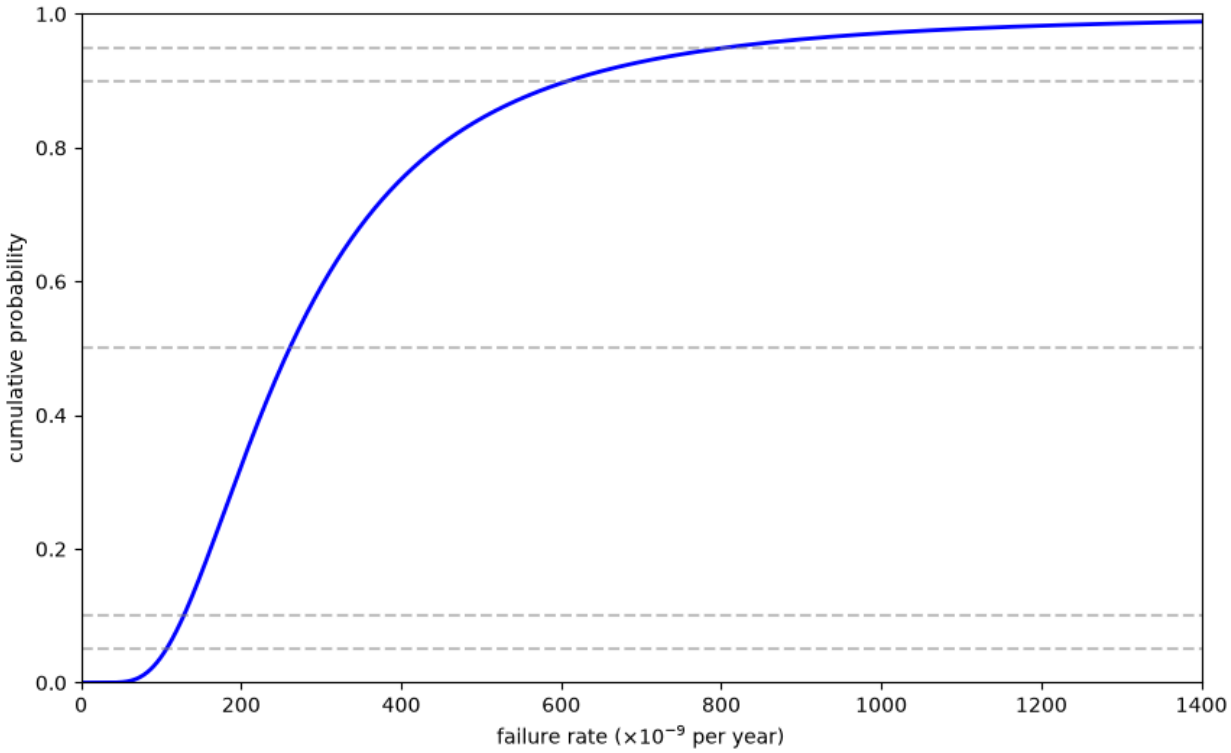


Fig. 14.8: CDF of the system failure rate

(continued from previous page)

```
for pc, published in reference.items():
    ours = system_grid[find_nearest(cdf, pc/100)]
    table_data.append([f"{pc}%", f"{ours:.1f}", published,
                       f"{100*(ours - published)/published:+.1f}%"])

print("\nSystem failure rate quantiles (×10^-9 per year):")
print(tabulate(table_data,
      headers=['Percentile', 'Computed here', 'Greenfield-Sargent', 'Difference'],
      tablefmt='grid'))
```

```
System failure rate quantiles (×10^-9 per year):
+--------------+-----------------+----------------------+--------------+
| Percentile   |   Computed here |   Greenfield-Sargent | Difference   |
+==============+=================+======================+==============+
| 1.0%         |            76.2 |                   77 | -1.1%        |
+--------------+-----------------+----------------------+--------------+
| 10.0%        |           128.2 |                  130 | -1.4%        |
+--------------+-----------------+----------------------+--------------+
| 50.0%        |           260.6 |                  263 | -0.9%        |
+--------------+-----------------+----------------------+--------------+
| 66.5%        |           338.6 |                  341 | -0.7%        |
+--------------+-----------------+----------------------+--------------+
| 85.0%        |           509.4 |                  513 | -0.7%        |
+--------------+-----------------+----------------------+--------------+
| 95.0%        |           807.6 |                  811 | -0.4%        |
+--------------+-----------------+----------------------+--------------+
```

(continues on next page)

(continued from previous page)

```
| 99.0%        |          1470.2 |                  1480 | -0.7%        |
+--------------+-----------------+-----------------------+--------------+
| 99.78%       |          2474.9 |                  2490 | -0.6%        |
+--------------+-----------------+-----------------------+--------------+
```

Our quantiles reproduce the published ones to within one and a half per cent, and slightly understate each of them.

The small discrepancies reflect the precision of the reported parameters μ_i, σ_i, the grid spacing m, and the point at which the grid is truncated.

14.9.2 Reading the answer

The numbers in this table, rather than any single one of them, are the output of the analysis.

The median failure rate is about 261×10^{-9} per year, while the 95th percentile is about 808×10^{-9}, three times larger.

That spread is not a statement about how often the hoist fails; it is a statement about how little the analyst knows about how often the hoist fails.

Notice also where the *mean* of the distribution falls.

```
mean_rate = np.sum(system_grid * system_pmf)
mean_percentile = 100 * cdf[find_nearest(system_grid, mean_rate)]

print(f"mean failure rate: {mean_rate:.1f} × 10⁻⁹ per year")
print(f"the mean sits at the {mean_percentile:.1f}th percentile")
```

```
mean failure rate: 338.1 × 10⁻⁹ per year
the mean sits at the 66.4th percentile
```

Because the distribution is skewed, the mean lies well above the median, at about the 66th percentile.

This is why Table 11 of Greenfield and Sargent [1993] records the mean alongside the 66.5th percentile.

The practical implication is the one that motivated the original study.

```
# The point estimate used in the U.S. Department of Energy's 1990 risk
# assessment, expressed in the same units
doe_estimate = 220    # 2.2 × 10^-7 per year

pct = 100 * cdf[find_nearest(system_grid, doe_estimate)]
print(f"the DOE point estimate of {doe_estimate} × 10⁻⁹ lies at "
      f"the {pct:.0f}th percentile")
print(f"so the analyst assigns probability {100-pct:.0f}% to the "
      f"true rate exceeding it")
```

```
the DOE point estimate of 220 × 10⁻⁹ lies at the 39th percentile
so the analyst assigns probability 61% to the true rate exceeding it
```

An analysis that reports a single number in place of a distribution conveys none of this.

Greenfield and Sargent [1993] made exactly this point: reading their figure, they put the Department of Energy's point estimate at the 36th percentile and concluded that there was roughly a 64 per cent chance that the true failure rate was higher.

14.10 Exercises

Exercise 14.10.1

Our discretization involves two separate choices: where the grid stops, $I = 2^p$, and how finely it is spaced, m.

Investigate what each one controls.

1. Holding $m = 0.05$ fixed, compute the median, the 95th percentile and the 99.78th percentile of the system failure rate for $p = 10, 11, \ldots, 15$. For each p, also compute how much probability mass the truncation discards, using $\sum_i \mathbb{P}\{P(A_i) > I\}$.
2. Holding $p = 14$ fixed, repeat for $m = 0.4, 0.2, 0.1, 0.05, 0.025$.
3. Which statistic is sensitive to which choice, and why? Are the values $p = 15$, $m = 0.05$ used in the lecture well chosen?

Solution

```
from scipy.stats import norm

def system_distribution(p_grid, m_grid):
    "Failure rate distribution of the whole system on a given grid."
    I_grid = 2**p_grid
    pmfs = []
    for μ_i, σ_i in params[:6]:
        _, pmf_i, _ = discretize_lognormal(μ_i, σ_i, I_grid, m_grid)
        pmfs.append(pmf_i)
    μ7, σ7 = params[6]
    _, pmf7, _ = discretize_lognormal(μ7, σ7, I_grid, m_grid)
    pmfs.extend([pmf7] * 8)

    total = pmfs[0]
    for pmf_i in pmfs[1:]:
        total = fftconvolve(total, pmf_i)
    return total, np.arange(len(total)) * m_grid


def quantiles_of(pmf, grid, levels=(0.5, 0.95, 0.9978)):
    cdf_local = np.cumsum(pmf)
    return [grid[find_nearest(cdf_local, q)] for q in levels]


def discarded_mass(I_grid):
    "Probability that a component's rate exceeds the end of the grid."
    lost = sum(norm.sf((np.log(I_grid) - μ_i)/σ_i) for μ_i, σ_i in params[:6])
    μ7, σ7 = params[6]
    return lost + 8 * norm.sf((np.log(I_grid) - μ7)/σ7)


rows = []
for p_test in range(10, 16):
    pmf_t, grid_t = system_distribution(p_test, 0.05)
    med, q95, q9978 = quantiles_of(pmf_t, grid_t)
    rows.append([p_test, 2**p_test, f"{discarded_mass(2**p_test):.1e}",
                 f"{med:.2f}", f"{q95:.2f}", f"{q9978:.2f}"])
```

```
print(tabulate(rows, headers=['p', 'I', 'mass discarded',
                              'median', '95th', '99.78th'], tablefmt='grid'))
```

```
+-----+-------+------------------+----------+--------+------------+
|   p |     I |   mass discarded |   median |   95th |   99.78th |
+=====+=======+==================+==========+========+============+
|  10 |  1024 |          0.016   |   257.55 | 728.8  |   1247.95 |
+-----+-------+------------------+----------+--------+------------+
|  11 |  2048 |          0.0028  |   260    | 791.1  |   1870.75 |
+-----+-------+------------------+----------+--------+------------+
|  12 |  4096 |          0.00039 |   260.5  | 805.3  |   2344.85 |
+-----+-------+------------------+----------+--------+------------+
|  13 |  8192 |          3.9e-05 |   260.55 | 807.35 |   2460.5  |
+-----+-------+------------------+----------+--------+------------+
|  14 | 16384 |          2.9e-06 |   260.55 | 807.6  |   2473.85 |
+-----+-------+------------------+----------+--------+------------+
|  15 | 32768 |          1.6e-07 |   260.55 | 807.6  |   2474.85 |
+-----+-------+------------------+----------+--------+------------+
```

```
rows = []
for m_test in (0.4, 0.2, 0.1, 0.05, 0.025):
    pmf_t, grid_t = system_distribution(14, m_test)
    med, q95, q9978 = quantiles_of(pmf_t, grid_t)
    rows.append([m_test, len(grid_t), f"{med:.3f}", f"{q95:.2f}", f"{q9978:.2f}"])

print(tabulate(rows, headers=['m', 'grid points', 'median', '95th', '99.78th'],
               tablefmt='grid'))
```

```
+-------+---------------+----------+--------+-----------+
|     m |   grid points |   median |   95th |   99.78th |
+=======+===============+==========+========+===========+
| 0.4   |        573427 |  264.4   |  811.6 |   2477.6  |
+-------+---------------+----------+--------+-----------+
| 0.2   |       1146867 |  261.8   |  809   |   2475.2  |
+-------+---------------+----------+--------+-----------+
| 0.1   |       2293747 |  260.9   |  807.9 |   2474.2  |
+-------+---------------+----------+--------+-----------+
| 0.05  |       4587507 |  260.55  |  807.6 |   2473.85 |
+-------+---------------+----------+--------+-----------+
| 0.025 |       9175027 |  260.475 |  807.5 |   2473.75 |
+-------+---------------+----------+--------+-----------+
```

The two choices do quite different jobs.

Truncation governs the *far tail*.

At $p = 10$ the 99.78th percentile is badly understated, and it keeps rising until about $p = 14$, by which point the discarded mass has fallen to roughly 10^{-6}.

The median, by contrast, has settled by $p = 12$: throwing away the extreme right tail of each component hardly moves the middle of the distribution of their sum.

Resolution governs *overall precision*.

Halving m shifts every quantile slightly and uniformly, and the shifts are small: going from $m = 0.4$ to $m = 0.025$ moves the median by about 1.5 per cent.

The lecture's choices are sensible.

With $p = 15$ the discarded mass is around 10^{-7}, so even the 99.78th percentile is accurate, and $m = 0.05$ is fine enough that further refinement changes little.

The moral is that a grid that looks adequate for the median can be badly inadequate for the upper tail, which is exactly the region a safety regulator cares about.

Exercise 14.10.2

The rare event approximation replaces $P(A \cup B)$ by $P(A) + P(B)$, discarding $P(A \cap B)$.

Assess how good it is here.

1. Using the *mean* failure rate of each of the fourteen components as a representative value, compare $\sum_i p_i$ with the exact probability that at least one component fails, $1 - \prod_i (1 - p_i)$.
2. Repeat with all the rates multiplied by 10^3, 10^6 and 10^7, and report the relative error in each case.
3. At what order of magnitude does the approximation start to matter?

Solution

```
# representative rate for each of the 14 components
component_means = [np.exp(μ_i + 0.5*σ_i**2) for μ_i, σ_i in params[:6]]
μ7, σ7 = params[6]
component_means.extend([np.exp(μ7 + 0.5*σ7**2)] * 8)
component_means = np.array(component_means)

rows = []
for factor, label in ((1e-9, 'as calibrated'), (1e-6, '× 10³'),
                      (1e-3, '× 10⁶'), (1e-2, '× 10⁷')):
    probs = component_means * factor
    approx = probs.sum()
    exact = 1 - np.prod(1 - probs)
    rows.append([label, f"{approx:.6e}", f"{exact:.6e}",
                 f"{100*(approx - exact)/exact:.4f}%"])

print(tabulate(rows, headers=['failure rates', 'Σ pᵢ', '1 - Π(1-pᵢ)',
                              'relative error'], tablefmt='grid'))
```

```
+-----------------+-------------+---------------+------------------+
| failure rates   |        Σ pᵢ |   1 - Π(1-pᵢ) | relative error   |
+=================+=============+===============+==================+
| as calibrated   | 3.38012e-07 |   3.38012e-07 | 0.0000%          |
+-----------------+-------------+---------------+------------------+
| × 10³           | 0.000338012 |   0.000337969 | 0.0127%          |
+-----------------+-------------+---------------+------------------+
| × 10⁶           | 0.338012    |   0.297883    | 13.4715%         |
+-----------------+-------------+---------------+------------------+
| × 10⁷           | 3.38012     |   1.04157     | 224.5216%        |
+-----------------+-------------+---------------+------------------+
```

At the calibrated magnitudes, around 3×10^{-7} per year in total, the approximation is exact to the precision shown: the neglected term is of order $p_i p_j \approx 10^{-14}$.

Multiplying every rate by a thousand still leaves an error of about one part in ten thousand.

The approximation only becomes consequential once individual failure probabilities reach the order of a per cent, where it overstates the system failure probability by more than ten per cent, and it fails completely when $\sum_i p_i$ approaches or exceeds one, where it can return a "probability" greater than one.

Note that a naive check of this approximation – comparing the mean of the computed distribution of $\sum_i P(A_i)$ with the sum of the component means – reveals nothing, because those two quantities are equal by linearity of expectation whatever the quality of the approximation.

Exercise 14.10.3

A regulator who learns that the 95th percentile of the failure rate is too high will ask which components to improve.

Answer that question by computing, for each of the seven component types, the 95th percentile of the system failure rate when that type is removed from the system entirely.

Rank the component types by how much they contribute to the upper tail, and compare the ranking with the components' mean failure rates.

Solution

```
def system_without(drop):
    "System failure rate distribution with component type `drop` removed."
    pmfs = []
    for k, (μ_i, σ_i) in enumerate(params[:6]):
        if k == drop:
            continue
        _, pmf_i, _ = discretize_lognormal(μ_i, σ_i, I, m)
        pmfs.append(pmf_i)
    if drop != 6:
        μ7, σ7 = params[6]
        _, pmf7, _ = discretize_lognormal(μ7, σ7, I, m)
        pmfs.extend([pmf7] * 8)

    total = pmfs[0]
    for pmf_i in pmfs[1:]:
        total = fftconvolve(total, pmf_i)
    return total, np.arange(len(total)) * m


base_q95 = system_grid[find_nearest(cdf, 0.95)]

rows = []
for k in range(7):
    pmf_k, grid_k = system_without(k)
    q95 = grid_k[find_nearest(np.cumsum(pmf_k), 0.95)]
    μ_k, σ_k = params[k]
    n_units = 8 if k == 6 else 1
    rows.append([f"type {k+1}", n_units, f"{np.exp(μ_k + 0.5*σ_k**2):.1f}",
                 f"{q95:.1f}", f"{100*(base_q95 - q95)/base_q95:.1f}%"])

rows.sort(key=lambda r: -float(r[4].rstrip('%')))
print(f"95th percentile with all components: {base_q95:.1f}\n")
print(tabulate(rows, headers=['removed', 'units', 'mean rate each',
                              '95th pct without it', 'reduction'],
               tablefmt='grid'))
```

```
95th percentile with all components: 807.6

+-----------+---------+------------------+-----------------------+-------------+
| removed   |   units |   mean rate each |   95th pct without it | reduction   |
+===========+=========+==================+=======================+=============+
| type 1    |       1 |            147.5 |                 439.3 | 45.6%       |
+-----------+---------+------------------+-----------------------+-------------+
| type 2    |       1 |             60.6 |                 699.4 | 13.4%       |
+-----------+---------+------------------+-----------------------+-------------+
| type 3    |       1 |             33.4 |                 760.1 | 5.9%        |
+-----------+---------+------------------+-----------------------+-------------+
| type 4    |       1 |             30.9 |                 764.7 | 5.3%        |
+-----------+---------+------------------+-----------------------+-------------+
| type 5    |       1 |             30.9 |                 764.7 | 5.3%        |
+-----------+---------+------------------+-----------------------+-------------+
| type 7    |       8 |              2.8 |                 783.7 | 3.0%        |
+-----------+---------+------------------+-----------------------+-------------+
| type 6    |       1 |             12.4 |                 791.1 | 2.0%        |
+-----------+---------+------------------+-----------------------+-------------+
```

Component type 1 dominates: removing that single unit cuts the 95th percentile by almost half, far more than any other change available to the designer.

The ranking follows the components' mean rates closely here, because all seven types have similar dispersions.

It need not do so in general: a component with a modest mean but a large σ contributes disproportionately to the upper tail, which is why the analyst works with the whole distribution rather than with means.

Note also that type 7, which appears eight times, matters less than type 1, which appears once.

Counting components is no guide to where the risk lies.

Exercise 14.10.4

We could have computed the distribution of the system failure rate by simulation instead of by convolution.

Draw samples of all fourteen component rates, sum them, and compare the resulting quantiles with those from the convolution, for sample sizes 10^4, 10^5 and 10^6.

Compare the median, the 95th percentile and the 99.78th percentile.

Which method would you prefer, and why?

Solution

```
all_params = list(params[:6]) + [params[6]] * 8
rng_mc = np.random.default_rng(0)
levels = (50, 95, 99.78)

rows = []
for N in (10_000, 100_000, 1_000_000):
    draws = sum(rng_mc.lognormal(μ_i, σ_i, N) for μ_i, σ_i in all_params)
    rows.append([f"{N:,}"] + [f"{np.percentile(draws, pc):.1f}" for pc in levels])

rows.append(['convolution'] +
            [f"{system_grid[find_nearest(cdf, pc/100)]:.1f}" for pc in levels])
```

```
print(tabulate(rows, headers=['method', 'median', '95th', '99.78th'],
               tablefmt='grid'))
```

```
+-------------+----------+--------+-----------+
| method      |   median |   95th |   99.78th |
+=============+==========+========+===========+
| 10,000      |    262.7 |  811.8 |    2240.6 |
+-------------+----------+--------+-----------+
| 100,000     |    261.1 |  807.1 |    2458.4 |
+-------------+----------+--------+-----------+
| 1,000,000   |    260.4 |  807.5 |    2447   |
+-------------+----------+--------+-----------+
| convolution |    260.6 |  807.6 |    2474.9 |
+-------------+----------+--------+-----------+
```

Simulation converges to the same answer, which is a useful check on both calculations.

The two methods differ in where their errors lie.

Monte Carlo error is largest exactly where the analysis matters most: the 99.78th percentile is pinned down by roughly one draw in five hundred, so with 10^4 draws only about twenty observations inform it, and the estimate is visibly off.

The convolution, by contrast, computes the whole distribution at once and its error comes from the grid rather than from sampling noise, so it is equally accurate in the tail as in the middle.

It is also deterministic: the answer does not change when you rerun it with a different seed.

Exercise 14.10.5

The entire calculation assumes that the fourteen component failure rates are statistically independent.

Investigate what happens when they are not.

Suppose that

$$\log P(A_i) = \mu_i + \sigma_i \left(\sqrt{\rho}\, z_0 + \sqrt{1-\rho}\, z_i \right),$$

where z_0 is a shock common to all components and $z_1, \ldots, z_{14}$ are idiosyncratic, all standard normal.

Each component still has exactly its original marginal distribution, but any two of them now have correlation ρ in logs.

Simulate the system failure rate for $\rho = 0, 0.2, 0.5, 0.8$ and report the median, the 95th, the 99th and the 99.9th percentiles.

Explain what happens and why it matters for a safety analysis.

Solution

```
μ_vec = np.array([q[0] for q in all_params])
σ_vec = np.array([q[1] for q in all_params])

N_sim = 400_000
rng_cc = np.random.default_rng(1)

rows = []
for ρ in (0.0, 0.2, 0.5, 0.8):
```

```
    z0 = rng_cc.normal(size=(N_sim, 1))
    zi = rng_cc.normal(size=(N_sim, len(all_params)))
    logs = μ_vec + σ_vec * (np.sqrt(ρ)*z0 + np.sqrt(1-ρ)*zi)
    totals = np.exp(logs).sum(axis=1)
    rows.append([ρ] + [f"{np.percentile(totals, pc):.0f}"
                       for pc in (50, 95, 99, 99.9)])

print(tabulate(rows, headers=['ρ', 'median', '95th', '99th', '99.9th'],
               tablefmt='grid'))
```

```
+-----+----------+--------+--------+----------+
|   ρ |   median |   95th |   99th |   99.9th |
+=====+==========+========+========+==========+
| 0   |      261 |    811 |   1481 |     3235 |
+-----+----------+--------+--------+----------+
| 0.2 |      242 |    910 |   1650 |     3416 |
+-----+----------+--------+--------+----------+
| 0.5 |      209 |   1043 |   2074 |     4396 |
+-----+----------+--------+--------+----------+
| 0.8 |      179 |   1136 |   2444 |     5763 |
+-----+----------+--------+--------+----------+
```

Correlation leaves each component's marginal distribution untouched, and it leaves the mean of the sum untouched as well.

What it changes is the shape of the distribution of the sum.

With independent components, a high draw for one is typically offset by ordinary draws for the others, and the fourteen-fold averaging produces a relatively concentrated total.

A common shock removes that diversification: when z_0 is large every component is bad at once.

The result is a distribution with a *lower* median and a much *heavier* upper tail.

At $\rho = 0.8$ the median falls by about a third while the 99.9th percentile rises by roughly three quarters.

For a safety analysis this is the dangerous direction of error.

Assuming independence when a common cause is present makes the system look both typically safer and much less likely to suffer a very bad year than it really is.

This is why reliability studies devote so much attention to identifying shared power supplies, shared maintenance procedures, common design faults, and other mechanisms that defeat the independence assumption.

CHAPTER

FIFTEEN

INTRODUCTION TO ARTIFICIAL NEURAL NETWORKS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's included in base Anaconda, we need to install the following packages

```
!pip install -U kaleido plotly
!conda install -y -c plotly plotly-orca

# kaleido needs chrome to build images
import kaleido
kaleido.get_chrome_sync()
```

Note

If you are running this on Google Colab the above cell will present an error. This is because Google Colab doesn't use Anaconda to manage the Python packages. However this lecture will still execute as Google Colab has `plotly` installed.

We also need to install JAX to run this lecture

```
!pip install --upgrade jax
```

```
import jax
print(f"JAX backend: {jax.devices()[0].platform}") # to check that gpu is activated␣
↪in environment
```

```
JAX backend: gpu
```

15.1 Overview

Substantial parts of **machine learning** and **artificial intelligence** are about

- approximating an unknown function with a known function
- estimating the known function from a set of data on the left- and right-hand variables

This lecture describes the structure of a plain vanilla **artificial neural network** (ANN) of a type that is widely used to approximate a function f that maps x in a space X into y in a space Y.

To introduce elementary concepts, we study an example in which x and y are scalars.

We'll describe the following concepts that are brick and mortar for neural networks:

- a neuron
- an activation function
- a network of neurons
- A neural network as a composition of functions
- back-propagation and its relationship to the chain rule of differential calculus

15.2 A Deep (but not Wide) Artificial Neural Network

We describe a "deep" neural network of "width" one.

Deep means that the network composes a large number of functions organized into nodes of a graph.

Width refers to the number of right hand side variables on the right hand side of the function being approximated.

Setting "width" to one means that the network composes just univariate functions.

Let $x \in \mathbb{R}$ be a scalar and $y \in \mathbb{R}$ be another scalar.

We assume that y is a nonlinear function of x:

$$y = f(x)$$

We want to approximate $f(x)$ with another function that we define recursively.

For a network of depth $N \geq 1$, each **layer** $i = 1, \ldots N$ consists of

- an input x_i
- an **affine function** $w_i x_i + bI$, where w_i is a scalar **weight** placed on the input x_i and b_i is a scalar **bias**
- an **activation function** h_i that takes $(w_i x_i + b_i)$ as an argument and produces an output x_{i+1}

An example of an activation function h is the **sigmoid** function

$$h(z) = \frac{1}{1 + e^{-z}}$$

Another popular activation function is the **rectified linear unit** (ReLU) function

$$h(z) = \max(0, z)$$

Yet another activation function is the identity function

$$h(z) = z$$

As activation functions below, we'll use the sigmoid function for layers 1 to $N-1$ and the identity function for layer N.

To approximate a function $f(x)$ we construct $\hat{f}(x)$ by proceeding as follows.

Let

$$l_i(x) = w_i x + b_i.$$

We construct $\hat{f}$ by iterating on compositions of functions $h_i \circ l_i$:

$$f(x) \approx \hat{f}(x) = h_N \circ l_N \circ h_{N-1} \circ l_1 \circ \cdots \circ h_1 \circ l_1(x)$$

If $N > 1$, we call the right side a "deep" neural net.

The larger is the integer N, the "deeper" is the neural net.

Evidently, if we know the parameters $\{w_i, b_i\}_{i=1}^N$, then we can compute $\hat{f}(x)$ for a given $x = \tilde{x}$ by iterating on the recursion

$$x_{i+1} = h_i \circ l_i(x_i), \quad , i = 1, \ldots N \tag{15.1}$$

starting from $x_1 = \tilde{x}$.

The value of x_{N+1} that emerges from this iterative scheme equals $\hat{f}(\tilde{x})$.

15.3 Calibrating Parameters

We now consider a neural network like the one describe above with width 1, depth N, and activation functions h_i for $1 \leqslant i \leqslant N$ that map $\mathbb{R}$ into itself.

Let $\{(w_i, b_i)\}_{i=1}^N$ denote a sequence of weights and biases.

As mentioned above, for a given input x_1, our approximating function $\hat{f}$ evaluated at x_1 equals the "output" x_{N+1} from our network that can be computed by iterating on $x_{i+1} = h_i(w_i x_i + b_i)$.

For a given **prediction** $\hat{y}(x)$ and **target** $y = f(x)$, consider the loss function

$$\mathcal{L}(\hat{y}, y)(x) = \frac{1}{2}(\hat{y} - y)^2(x).$$

This criterion is a function of the parameters $\{(w_i, b_i)\}_{i=1}^N$ and the point x.

We're interested in solving the following problem:

$$\min_{\{(w_i, b_i)\}_{i=1}^N} \int \mathcal{L}(x_{N+1}, y)(x) d\mu(x)$$

where $\mu(x)$ is some measure of points $x \in \mathbb{R}$ over which we want a good approximation $\hat{f}(x)$ to $f(x)$.

Stack weights and biases into a vector of parameters p:

$$p = \begin{bmatrix} w_1 \\ b_1 \\ w_2 \\ b_2 \\ \vdots \\ w_N \\ b_N \end{bmatrix}$$

Applying a "poor man's version" of a **stochastic gradient descent** algorithm for finding a zero of a function leads to the following update rule for parameters:

$$p_{k+1} = p_k - \alpha \frac{d\mathcal{L}}{dx_{N+1}} \frac{dx_{N+1}}{dp_k} \tag{15.2}$$

where $\frac{d\mathcal{L}}{dx_{N+1}} = -\left(x_{N+1} - y\right)$ and $\alpha > 0$ is a step size.

(See this and this to gather insights about how stochastic gradient descent relates to Newton's method.)

To implement one step of this parameter update rule, we want the vector of derivatives $\frac{dx_{N+1}}{dp_k}$.

In the neural network literature, this step is accomplished by what is known as **back propagation**.

15.4 Back Propagation and the Chain Rule

Thanks to properties of

- the chain and product rules for differentiation from differential calculus, and
- lower triangular matrices

back propagation can actually be accomplished in one step by

- inverting a lower triangular matrix, and
- matrix multiplication

(This idea is from the last 7 minutes of this great youtube video by MIT's Alan Edelman)

https://youtu.be/rZS2LGiurKY

Here goes.

Define the derivative of $h(z)$ with respect to z evaluated at $z = z_i$ as δ_i:

$$\delta_i = \frac{d}{dz} h(z)|_{z=z_i}$$

or

$$\delta_i = h'\left(w_i x_i + b_i\right).$$

Repeated application of the chain rule and product rule to our recursion (15.1) allows us to obtain:

$$dx_{i+1} = \delta_i \left(dw_i x_i + w_i dx_i + b_i\right)$$

After imposing $dx_1 = 0$, we get the following system of equations:

$$\begin{pmatrix} dx_2 \\ \vdots \\ dx_{N+1} \end{pmatrix} = \underbrace{\begin{pmatrix} \delta_1 w_1 & \delta_1 & 0 & 0 & 0 \\ 0 & 0 & \ddots & 0 & 0 \\ 0 & 0 & 0 & \delta_N w_N & \delta_N \end{pmatrix}}_{D} \begin{pmatrix} dw_1 \\ db_1 \\ \vdots \\ dw_N \\ db_N \end{pmatrix} + \underbrace{\begin{pmatrix} 0 & 0 & 0 & 0 \\ w_2 & 0 & 0 & 0 \\ 0 & \ddots & 0 & 0 \\ 0 & 0 & w_N & 0 \end{pmatrix}}_{L} \begin{pmatrix} dx_2 \\ \vdots \\ dx_{N+1} \end{pmatrix}$$

or

$$dx = Ddp + Ldx$$

which implies that

$$dx = (I - L)^{-1} D dp$$

which in turn implies

$$\begin{pmatrix} dx_{N+1}/dw_1 \\ dx_{N+1}/db_1 \\ \vdots \\ dx_{N+1}/dw_N \\ dx_{N+1}/db_N \end{pmatrix} = e_N (I - L)^{-1} D.$$

We can then solve the above problem by applying our update for p multiple times for a collection of input-output pairs $\{(x_1^i, y^i)\}_{i=1}^M$ that we'll call our "training set".

15.5 Training Set

Choosing a training set amounts to a choice of measure μ in the above formulation of our function approximation problem as a minimization problem.

In this spirit, we shall use a uniform grid of, say, 50 or 200 points.

There are many possible approaches to the minimization problem posed above:

- batch gradient descent in which you use an average gradient over the training set
- stochastic gradient descent in which you sample points randomly and use individual gradients
- something in-between (so-called "mini-batch gradient descent")

The update rule (15.2) described above amounts to a stochastic gradient descent algorithm.

```
from IPython.display import Image
import jax.numpy as jnp
from jax import grad, jit, jacfwd, vmap
from jax import random
import jax
import plotly.graph_objects as go
```

```
# A helper function to randomly initialize weights and biases
# for a dense neural network layer
def random_layer_params(m, n, key, scale=1.):
    w_key, b_key = random.split(key)
    return scale * random.normal(w_key, (n, m)), scale * random.normal(b_key, (n,))

# Initialize all layers for a fully-connected neural network with sizes "sizes"
def init_network_params(sizes, key):
    keys = random.split(key, len(sizes))
    return [random_layer_params(m, n, k) for m, n, k in zip(sizes[:-1], sizes[1:],␣
 ↪keys)]
```

```
def compute_xδw_seq(params, x):
    # Initialize arrays
    δ = jnp.zeros(len(params))
    xs = jnp.zeros(len(params) + 1)
    ws = jnp.zeros(len(params))
```

(continues on next page)

(continued from previous page)

```
    bs = jnp.zeros(len(params))

    h = jax.nn.sigmoid

    xs = xs.at[0].set(x)
    for i, (w, b) in enumerate(params[:-1]):
        output = w * xs[i] + b
        activation = h(output[0, 0])

        # Store elements
        δ = δ.at[i].set(grad(h)(output[0, 0]))
        ws = ws.at[i].set(w[0, 0])
        bs = bs.at[i].set(b[0])
        xs = xs.at[i+1].set(activation)

    final_w, final_b = params[-1]
    preds = final_w * xs[-2] + final_b

    # Store elements
    δ = δ.at[-1].set(1.)
    ws = ws.at[-1].set(final_w[0, 0])
    bs = bs.at[-1].set(final_b[0])
    xs = xs.at[-1].set(preds[0, 0])

    return xs, δ, ws, bs


def loss(params, x, y):
    xs, δ, ws, bs = compute_xδw_seq(params, x)
    preds = xs[-1]

    return 1 / 2 * (y - preds) ** 2
```

```
# Parameters
N = 3  # Number of layers
layer_sizes = [1, ] * (N + 1)
param_scale = 0.1
step_size = 0.01
params = init_network_params(layer_sizes, random.key(1))
```

```
x = 5
y = 3
xs, δ, ws, bs = compute_xδw_seq(params, x)
```

```
dxs_ad = jacfwd(lambda params, x: compute_xδw_seq(params, x)[0], argnums=0)(params, x)
dxs_ad_mat = jnp.block([dx.reshape((-1, 1)) for dx_tuple in dxs_ad for dx in dx_tuple␣
↪])[1:]
```

```
jnp.block([[δ * xs[:-1]], [δ]])
```

```
Array([[1.0165801 , 0.06087969, 0.09382247],
       [0.20331602, 0.08501981, 1.        ]], dtype=float32)
```

```
L = jnp.diag(δ * ws, k=-1)
```

(continues on next page)

(continued from previous page)

```
L = L[1:, 1:]

D = jax.scipy.linalg.block_diag(*[row.reshape((1, 2)) for row in jnp.block([[δ * xs[:-
↪1]], [δ]]).T])

dxs_la = jax.scipy.linalg.solve_triangular(jnp.eye(N) - L, D, lower=True)
```

```
# Check that the `dx` generated by the linear algebra method
# are the same as the ones generated using automatic differentiation
jnp.max(jnp.abs(dxs_ad_mat - dxs_la))
```

```
Array(0., dtype=float32)
```

```
grad_loss_ad = jnp.block([dx.reshape((-1, 1)) for dx_tuple in grad(loss)(params, x,␣
↪y) for dx in dx_tuple ])
```

```
# Check that the gradient of the loss is the same for both approaches
jnp.max(jnp.abs(-(y - xs[-1]) * dxs_la[-1] - grad_loss_ad))
```

```
Array(5.9604645e-08, dtype=float32)
```

```
@jit
def update_ad(params, x, y):
    grads = grad(loss)(params, x, y)
    return [(w - step_size * dw, b - step_size * db)
          for (w, b), (dw, db) in zip(params, grads)]

@jit
def update_la(params, x, y):
    xs, δ, ws, bs = compute_xδw_seq(params, x)
    N = len(params)
    L = jnp.diag(δ * ws, k=-1)
    L = L[1:, 1:]

    D = jax.scipy.linalg.block_diag(*[row.reshape((1, 2)) for row in jnp.block([[δ *␣
↪xs[:-1]], [δ]]).T])

    dxs_la = jax.scipy.linalg.solve_triangular(jnp.eye(N) - L, D, lower=True)

    grads = -(y - xs[-1]) * dxs_la[-1]

    return [(w - step_size * dw, b - step_size * db)
            for (w, b), (dw, db) in zip(params, grads.reshape((-1, 2)))]
```

```
# Check that both updates are the same
update_la(params, x, y)
```

```
[(Array([[-0.00826643]], dtype=float32), Array([0.94700736], dtype=float32)),
 (Array([[-2.0638916]], dtype=float32), Array([-0.7872697], dtype=float32)),
 (Array([[1.6248171]], dtype=float32), Array([1.5765371], dtype=float32))]
```

```
update_ad(params, x, y)
```

```
[(Array([[-0.00826644]], dtype=float32), Array([0.94700736], dtype=float32)),
 (Array([[-2.0638916]], dtype=float32), Array([-0.7872697], dtype=float32)),
 (Array([[1.6248171]], dtype=float32), Array([1.5765371], dtype=float32))]
```

15.6 Example 1

Consider the function

$$f(x) = -3x + 2$$

on $[0.5, 3]$.

We use a uniform grid of 200 points and update the parameters for each point on the grid 300 times.

h_i is the sigmoid activation function for all layers except the final one for which we use the identity function and $N = 3$.

Weights are initialized randomly.

```
def f(x):
    return -3 * x + 2

M = 200
grid = jnp.linspace(0.5, 3, num=M)
f_val = f(grid)
```

```
indices = jnp.arange(M)
key = random.key(0)

def train(params, grid, f_val, key, num_epochs=300):
    for epoch in range(num_epochs):
        key, _ = random.split(key)
        random_permutation = random.permutation(random.key(1), indices)
        for x, y in zip(grid[random_permutation], f_val[random_permutation]):
            params = update_la(params, x, y)

    return params
```

```
# Parameters
N = 3  # Number of layers
layer_sizes = [1, ] * (N + 1)
params_ex1 = init_network_params(layer_sizes, key)
```

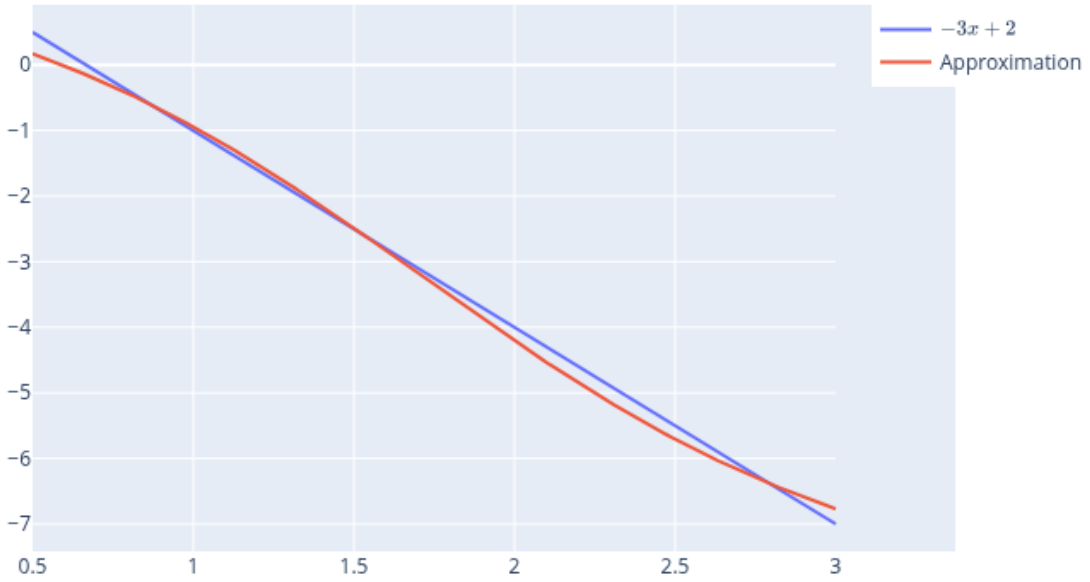
```
%%time
params_ex1 = train(params_ex1, grid, f_val, key, num_epochs=500)
```

```
CPU times: user 20.8 s, sys: 4.27 s, total: 25.1 s
Wall time: 18 s
```

```
predictions = vmap(compute_xδw_seq, in_axes=(None, 0))(params_ex1, grid)[0][:, -1]
```

```
fig = go.Figure()
fig.add_trace(go.Scatter(x=grid, y=f_val, name=r'$-3x+2$'))
fig.add_trace(go.Scatter(x=grid, y=predictions, name='Approximation'))

# Export to PNG file
Image(fig.to_image(format="png"))
# fig.show() will provide interactive plot when running
# notebook locally
```



15.7 How Deep?

It is fun to think about how deepening the neural net for the above example affects the quality of approximation

- If the network is too deep, you'll run into the vanishing gradient problem
- Other parameters such as the step size and the number of epochs can be as important or more important than the number of layers in the situation considered in this lecture.
- Indeed, since f is a linear function of x, a one-layer network with the identity map as an activation would probably work best.

15.8 Example 2

We use the same setup as for the previous example with

$$f(x) = \log(x)$$

```
def f(x):
    return jnp.log(x)

grid = jnp.linspace(0.5, 3, num=M)
f_val = f(grid)
```

```
# Parameters
N = 1  # Number of layers
layer_sizes = [1, ] * (N + 1)
params_ex2_1 = init_network_params(layer_sizes, key)
```

```
# Parameters
N = 2  # Number of layers
layer_sizes = [1, ] * (N + 1)
params_ex2_2 = init_network_params(layer_sizes, key)
```

```
# Parameters
N = 3  # Number of layers
layer_sizes = [1, ] * (N + 1)
params_ex2_3 = init_network_params(layer_sizes, key)
```

```
params_ex2_1 = train(params_ex2_1, grid, f_val, key, num_epochs=300)
```

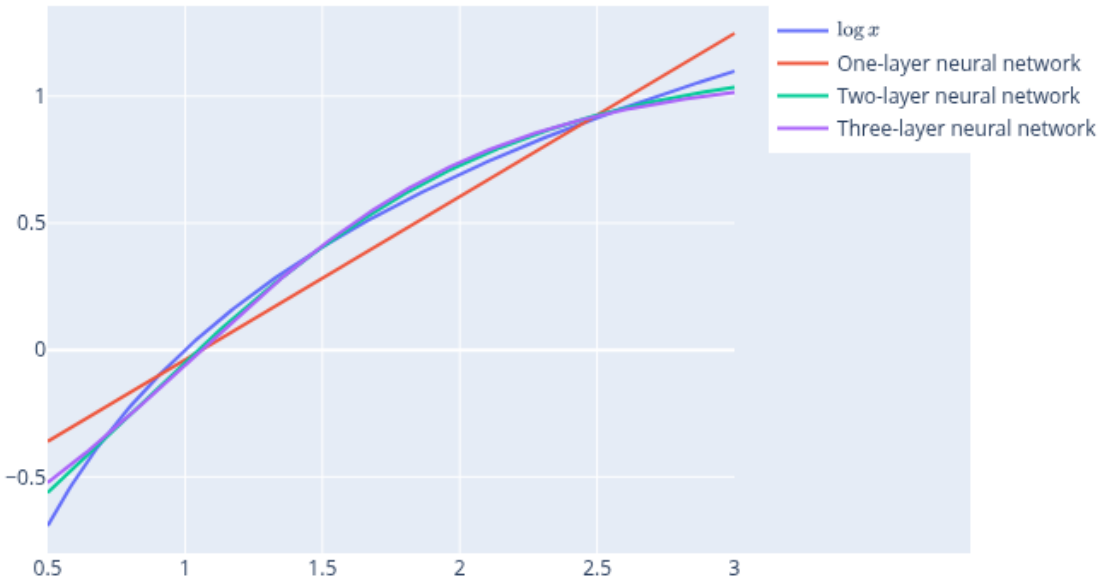
```
params_ex2_2 = train(params_ex2_2, grid, f_val, key, num_epochs=300)
```

```
params_ex2_3 = train(params_ex2_3, grid, f_val, key, num_epochs=300)
```

```
predictions_1 = vmap(compute_xδw_seq, in_axes=(None, 0))(params_ex2_1, grid)[0][:, -1]
predictions_2 = vmap(compute_xδw_seq, in_axes=(None, 0))(params_ex2_2, grid)[0][:, -1]
predictions_3 = vmap(compute_xδw_seq, in_axes=(None, 0))(params_ex2_3, grid)[0][:, -1]
```

```
fig = go.Figure()
fig.add_trace(go.Scatter(x=grid, y=f_val, name=r'$\log{x}$'))
fig.add_trace(go.Scatter(x=grid, y=predictions_1, name='One-layer neural network'))
fig.add_trace(go.Scatter(x=grid, y=predictions_2, name='Two-layer neural network'))
fig.add_trace(go.Scatter(x=grid, y=predictions_3, name='Three-layer neural network'))

# Export to PNG file
Image(fig.to_image(format="png"))
# fig.show() will provide interactive plot when running
# notebook locally
```


log x
One-layer neural network
Two-layer neural network
Three-layer neural network
1
0.5
0
−0.5
0.5
1
1.5
2
2.5
3

CHAPTER

SIXTEEN

RANDOMIZED RESPONSE SURVEYS

16.1 Overview

Social stigmas can inhibit people from confessing potentially embarrassing activities or opinions.

When people are reluctant to participate a sample survey about personally sensitive issues, they might decline to participate, and even if they do participate, they might choose to provide incorrect answers to sensitive questions.

These problems induce **selection** biases that present challenges to interpreting and designing surveys.

To illustrate how social scientists have thought about estimating the prevalence of such embarrassing activities and opinions, this lecture describes a classic approach of S. L. Warner [Warner, 1965].

Warner used elementary probability to construct a way to protect the privacy of **individual** respondents to surveys while still estimating the fraction of a **collection** of individuals who have a socially stigmatized characteristic or who engage in a socially stigmatized activity.

Warner's idea was to add **noise** between the respondent's answer and the **signal** about that answer that the survey maker ultimately receives.

Knowing about the structure of the noise assures the respondent that the survey maker does not observe his answer.

Statistical properties of the noise injection procedure provide the respondent **plausible deniability**.

Related ideas underlie modern **differential privacy** systems.

(See https://en.wikipedia.org/wiki/Differential_privacy)

16.2 Warner's Strategy

As usual, let's bring in the Python modules we'll be using.

```
import numpy as np
import pandas as pd
```

Suppose that every person in population either belongs to Group A or Group B.

We want to estimate the proportion π who belong to Group A while protecting individual respondents' privacy.

Warner [Warner, 1965] proposed and analyzed the following procedure.

- A random sample of n people is drawn with replacement from the population and each person is interviewed.
- Draw n random samples from the population with replacement and interview each person.
- Prepare a **random spinner** that with p probability points to the Letter A and with $(1-p)$ probability points to the Letter B.

- Each subject spins a random spinner and sees an outcome (A or B) that the interviewer does **not observe**.
- The subject states whether he belongs to the group to which the spinner points.
- If the spinner points to the group that the spinner belongs, the subject reports "yes"; otherwise he reports "no".
- The subject answers the question truthfully.

Warner constructed a maximum likelihood estimators of the proportion of the population in set A.

Let

- π : True probability of A in the population
- p : Probability that the spinner points to A
- $X_i = \begin{cases} 1, \text{ if the } i\text{th subject says yes} \\ 0, \text{ if the } i\text{th subject says no} \end{cases}$

Index the sample set so that the first n_1 report "yes", while the second $n - n_1$ report "no".

The likelihood function of a sample set is

$$L = [\pi p + (1-\pi)(1-p)]^{n_1} [(1-\pi)p + \pi(1-p)]^{n-n_1} \tag{16.1}$$

The log of the likelihood function is:

$$\log(L) = n_1 \log [\pi p + (1-\pi)(1-p)] + (n-n_1) \log [(1-\pi)p + \pi(1-p)] \tag{16.2}$$

The first-order necessary condition for maximizing the log likelihood function with respect to π is:

$$\frac{(n-n_1)(2p-1)}{(1-\pi)p + \pi(1-p)} = \frac{n_1(2p-1)}{\pi p + (1-\pi)(1-p)}$$

or

$$\pi p + (1-\pi)(1-p) = \frac{n_1}{n} \tag{16.3}$$

If $p \neq \frac{1}{2}$, then the maximum likelihood estimator (MLE) of π is:

$$\hat{\pi} = \frac{p-1}{2p-1} + \frac{n_1}{(2p-1)n} \tag{16.4}$$

We compute the mean and variance of the MLE estimator $\hat{\pi}$ to be:

$$\begin{aligned} \mathbb{E}(\hat{\pi}) &= \frac{1}{2p-1} \left[p - 1 + \frac{1}{n} \sum_{i=1}^{n} \mathbb{E} X_i \right] \\ &= \frac{1}{2p-1} [p - 1 + \pi p + (1-\pi)(1-p)] \\ &= \pi \end{aligned} \tag{16.5}$$

and

$$\begin{aligned} Var(\hat{\pi}) &= \frac{nVar(X_i)}{(2p-1)^2 n^2} \\ &= \frac{[\pi p + (1-\pi)(1-p)] [(1-\pi)p + \pi(1-p)]}{(2p-1)^2 n^2} \\ &= \frac{\frac{1}{4} + (2p^2 - 2p + \frac{1}{2})(-2\pi^2 + 2\pi - \frac{1}{2})}{(2p-1)^2 n^2} \\ &= \frac{1}{n} \left[\frac{1}{16(p - \frac{1}{2})^2} - (\pi - \frac{1}{2})^2 \right] \end{aligned} \tag{16.6}$$

Equation (16.5) indicates that $\hat{\pi}$ is an **unbiased estimator** of π while equation (16.6) tell us the variance of the estimator.

To compute a confidence interval, first rewrite (16.6) as:

$$Var(\hat{\pi}) = \frac{\frac{1}{4} - (\pi - \frac{1}{2})^2}{n} + \frac{\frac{1}{16(p-\frac{1}{2})^2} - \frac{1}{4}}{n} \tag{16.7}$$

This equation indicates that the variance of $\hat{\pi}$ can be represented as a sum of the variance due to sampling plus the variance due to the random device.

From the expressions above we can find that:

- When p is $\frac{1}{2}$, expression (16.1) degenerates to a constant.
- When p is 1 or 0, the randomized estimate degenerates to an estimator without randomized sampling.

We shall only discuss situations in which $p \in (\frac{1}{2}, 1)$

(a situation in which $p \in (0, \frac{1}{2})$ is symmetric).

From expressions (16.5) and (16.7) we can deduce that:

- The MSE of $\hat{\pi}$ decreases as p increases.

16.3 Comparing Two Survey Designs

Let's compare the preceding randomized-response method with a stylized non-randomized response method.

In our non-randomized response method, we suppose that:

- Members of Group A tells the truth with probability T_a while the members of Group B tells the truth with probability T_b
- Y_i is 1 or 0 according to whether the sample's ith member's report is in Group A or not.

Then we can estimate π as:

$$\hat{\pi} = \frac{\sum_{i=1}^{n} Y_i}{n} \tag{16.8}$$

We calculate the expectation, bias, and variance of the estimator to be:

$$\mathbb{E}(\hat{\pi}) = \pi T_a + [(1-\pi)(1-T_b)] \tag{16.9}$$

$$\begin{aligned} Bias(\hat{\pi}) &= \mathbb{E}(\hat{\pi} - \pi) \\ &= \pi[T_a + T_b - 2] + [1 - T_b] \end{aligned} \tag{16.10}$$

$$Var(\hat{\pi}) = \frac{[\pi T_a + (1-\pi)(1-T_b)]\,[1 - \pi T_a - (1-\pi)(1-T_b)]}{n} \tag{16.11}$$

It is useful to define a

$$\text{MSE Ratio} = \frac{\text{Mean Square Error Randomized}}{\text{Mean Square Error Regular}}$$

We can compute MSE Ratios for different survey designs associated with different parameter values.

The following Python code computes objects we want to stare at in order to make comparisons under different values of π_A and n:

```
class Comparison:
    def __init__(self, A, n):
        self.A = A
        self.n = n
        TaTb = np.array([[0.95,  1], [0.9,   1], [0.7,    1],
                         [0.5,   1], [1,  0.95], [1,    0.9],
                         [1,   0.7], [1,   0.5], [0.95, 0.95],
                         [0.9, 0.9], [0.7, 0.7], [0.5,  0.5]])
        self.p_arr = np.array([0.6, 0.7, 0.8, 0.9])
        self.p_map = dict(zip(self.p_arr, [f"MSE Ratio: p = {x}" for x in self.p_
↪arr]))
        self.template = pd.DataFrame(columns=self.p_arr)
        self.template[['T_a','T_b']] = TaTb
        self.template['Bias'] = None

    def theoretical(self):
        A = self.A
        n = self.n
        df = self.template.copy()
        df['Bias'] = A * (df['T_a'] + df['T_b'] - 2) + (1 - df['T_b'])
        for p in self.p_arr:
            df[p] = (1 / (16 * (p - 1/2)**2) - (A - 1/2)**2) / n / \
                    (df['Bias']**2 + ((A * df['T_a'] + (1 - A) * (1 - df['T_b'])) *␣
↪(1 - A * df['T_a'] - (1 - A) * (1 - df['T_b'])) / n))
            df[p] = df[p].round(2)
        df = df.set_index(["T_a", "T_b", "Bias"]).rename(columns=self.p_map)
        return df

    def MCsimulation(self, size=1000, seed=123456):
        A = self.A
        n = self.n
        df = self.template.copy()
        rng = np.random.default_rng(seed)
        sample = rng.random((size, self.n)) <= A
        random_device = rng.random((size, n))
        mse_rd = {}
        for p in self.p_arr:
            spinner = random_device <= p
            rd_answer = sample * spinner + (1 - sample) * (1 - spinner)
            n1 = rd_answer.sum(axis=1)
            pi_hat = (p - 1) / (2 * p - 1) + n1 / n / (2 * p - 1)
            mse_rd[p] = np.sum((pi_hat - A)**2)
        for inum, irow in df.iterrows():
            truth_a = rng.random((size, self.n)) <= irow.T_a
            truth_b = rng.random((size, self.n)) <= irow.T_b
            trad_answer = sample * truth_a + (1 - sample) * (1 - truth_b)
            pi_trad = trad_answer.sum(axis=1) / n
            df.loc[inum, 'Bias'] = pi_trad.mean() - A
            mse_trad = np.sum((pi_trad - A)**2)
            for p in self.p_arr:
                df.loc[inum, p] = (mse_rd[p] / mse_trad).round(2)
        df = df.set_index(["T_a", "T_b", "Bias"]).rename(columns=self.p_map)
        return df
```

Let's put the code to work for parameter values

- $\pi_A = 0.6$
- $n = 1000$

We can generate MSE Ratios theoretically using the above formulas.

We can also perform Monte Carlo simulations of a MSE Ratio.

```
cp1 = Comparison(0.6, 1000)
df1_theoretical = cp1.theoretical()
df1_theoretical
```

```
                MSE Ratio: p = 0.6  MSE Ratio: p = 0.7  MSE Ratio: p = 0.8  \
T_a  T_b  Bias
0.95 1.00 -0.03               5.45                1.36                0.60
0.90 1.00 -0.06               1.62                0.40                0.18
0.70 1.00 -0.18               0.19                0.05                0.02
0.50 1.00 -0.30               0.07                0.02                0.01
1.00 0.95  0.02               9.82                2.44                1.08
     0.90  0.04               3.41                0.85                0.37
     0.70  0.12               0.43                0.11                0.05
     0.50  0.20               0.16                0.04                0.02
0.95 0.95 -0.01              18.25                4.54                2.00
0.90 0.90 -0.02               9.70                2.41                1.06
0.70 0.70 -0.06               1.62                0.40                0.18
0.50 0.50 -0.10               0.61                0.15                0.07

                MSE Ratio: p = 0.9
T_a  T_b  Bias
0.95 1.00 -0.03               0.33
0.90 1.00 -0.06               0.10
0.70 1.00 -0.18               0.01
0.50 1.00 -0.30               0.00
1.00 0.95  0.02               0.60
     0.90  0.04               0.21
     0.70  0.12               0.03
     0.50  0.20               0.01
0.95 0.95 -0.01               1.11
0.90 0.90 -0.02               0.59
0.70 0.70 -0.06               0.10
0.50 0.50 -0.10               0.04
```

```
df1_mc = cp1.MCsimulation()
df1_mc
```

```
                   MSE Ratio: p = 0.6 MSE Ratio: p = 0.7 MSE Ratio: p = 0.8  \
T_a  T_b  Bias
0.95 1.00 -0.030191               5.37               1.34               0.62
0.90 1.00 -0.060081               1.61                0.4               0.19
0.70 1.00 -0.180543               0.19               0.05               0.02
0.50 1.00 -0.29977                0.07               0.02               0.01
1.00 0.95  0.019297              10.18               2.54               1.18
     0.90  0.03985                3.41               0.85                0.4
     0.70  0.119465               0.43               0.11               0.05
     0.50  0.199398               0.16               0.04               0.02
0.95 0.95 -0.010326              16.91               4.21               1.96
0.90 0.90 -0.019838               9.67               2.41               1.12
0.70 0.70 -0.06076                1.58               0.39               0.18
0.50 0.50 -0.100091               0.61               0.15               0.07

                   MSE Ratio: p = 0.9
```

(continues on next page)

(continued from previous page)

```
T_a  T_b  Bias
0.95 1.00 -0.030191              0.35
0.90 1.00 -0.060081              0.11
0.70 1.00 -0.180543              0.01
0.50 1.00 -0.29977                0.0
1.00 0.95  0.019297              0.67
     0.90  0.03985               0.22
     0.70  0.119465              0.03
     0.50  0.199398              0.01
0.95 0.95 -0.010326               1.1
0.90 0.90 -0.019838              0.63
0.70 0.70 -0.06076                0.1
0.50 0.50 -0.100091              0.04
```

The theoretical calculations do a good job of predicting Monte Carlo results.

We see that in many situations, especially when the bias is not small, the MSE of the randomized-sampling methods is smaller than that of the non-randomized sampling method.

These differences become larger as p increases.

By adjusting parameters π_A and n, we can study outcomes in different situations.

For example, for another situation described in Warner [Warner, 1965]:

- $\pi_A = 0.5$
- $n = 1000$

we can use the code

```
cp2 = Comparison(0.5, 1000)
df2_theoretical = cp2.theoretical()
df2_theoretical
```

```
                  MSE Ratio: p = 0.6  MSE Ratio: p = 0.7  MSE Ratio: p = 0.8  \
T_a  T_b  Bias
0.95 1.00 -0.025                7.15                1.79                0.79
0.90 1.00 -0.050                2.27                0.57                0.25
0.70 1.00 -0.150                0.27                0.07                0.03
0.50 1.00 -0.250                0.10                0.02                0.01
1.00 0.95  0.025                7.15                1.79                0.79
     0.90  0.050                2.27                0.57                0.25
     0.70  0.150                0.27                0.07                0.03
     0.50  0.250                0.10                0.02                0.01
0.95 0.95  0.000               25.00                6.25                2.78
0.90 0.90  0.000               25.00                6.25                2.78
0.70 0.70  0.000               25.00                6.25                2.78
0.50 0.50  0.000               25.00                6.25                2.78

                  MSE Ratio: p = 0.9
T_a  T_b  Bias
0.95 1.00 -0.025                0.45
0.90 1.00 -0.050                0.14
0.70 1.00 -0.150                0.02
0.50 1.00 -0.250                0.01
1.00 0.95  0.025                0.45
     0.90  0.050                0.14
     0.70  0.150                0.02
```

(continues on next page)

(continued from previous page)

```
     0.50  0.250                0.01
0.95 0.95  0.000                1.56
0.90 0.90  0.000                1.56
0.70 0.70  0.000                1.56
0.50 0.50  0.000                1.56
```

```
df2_mc = cp2.MCsimulation()
df2_mc
```

```
                    MSE Ratio: p = 0.6 MSE Ratio: p = 0.7 MSE Ratio: p = 0.8  \
T_a  T_b  Bias
0.95 1.00 -0.025243               7.17               1.78               0.83
0.90 1.00 -0.050271               2.29               0.57               0.27
0.70 1.00 -0.150607               0.28               0.07               0.03
0.50 1.00 -0.250029                0.1               0.03               0.01
1.00 0.95  0.024097               7.63               1.89               0.89
     0.90  0.049794               2.33               0.58               0.27
     0.70  0.149336               0.28               0.07               0.03
     0.50  0.249073                0.1               0.03               0.01
0.95 0.95 -0.000357              23.53               5.84               2.74
0.90 0.90  0.000068              25.04               6.21               2.91
0.70 0.70 -0.00079               26.04               6.46               3.03
0.50 0.50 -0.000077              24.26               6.02               2.82

                    MSE Ratio: p = 0.9
T_a  T_b  Bias
0.95 1.00 -0.025243                0.47
0.90 1.00 -0.050271                0.15
0.70 1.00 -0.150607                0.02
0.50 1.00 -0.250029                0.01
1.00 0.95  0.024097                 0.5
     0.90  0.049794                0.15
     0.70  0.149336                0.02
     0.50  0.249073                0.01
0.95 0.95 -0.000357                1.54
0.90 0.90  0.000068                1.64
0.70 0.70 -0.00079                  1.7
0.50 0.50 -0.000077                1.59
```

We can also revisit a calculation in the concluding section of Warner [Warner, 1965] in which

- $\pi_A = 0.6$
- $n = 2000$

We use the code

```
cp3 = Comparison(0.6, 2000)
df3_theoretical = cp3.theoretical()
df3_theoretical
```

```
                  MSE Ratio: p = 0.6  MSE Ratio: p = 0.7  MSE Ratio: p = 0.8  \
T_a  T_b  Bias
0.95 1.00 -0.03                 3.05                0.76                0.33
0.90 1.00 -0.06                 0.84                0.21                0.09
0.70 1.00 -0.18                 0.10                0.02                0.01
0.50 1.00 -0.30                 0.03                0.01                0.00
```

(continues on next page)

(continued from previous page)

```
1.00 0.95  0.02                6.03                1.50                0.66
     0.90  0.04                1.82                0.45                0.20
     0.70  0.12                0.22                0.05                0.02
     0.50  0.20                0.08                0.02                0.01
0.95 0.95 -0.01               14.12                3.51                1.55
0.90 0.90 -0.02                5.98                1.49                0.66
0.70 0.70 -0.06                0.84                0.21                0.09
0.50 0.50 -0.10                0.31                0.08                0.03

                 MSE Ratio: p = 0.9
T_a  T_b  Bias
0.95 1.00 -0.03                0.19
0.90 1.00 -0.06                0.05
0.70 1.00 -0.18                0.01
0.50 1.00 -0.30                0.00
1.00 0.95  0.02                0.37
     0.90  0.04                0.11
     0.70  0.12                0.01
     0.50  0.20                0.00
0.95 0.95 -0.01                0.86
0.90 0.90 -0.02                0.36
0.70 0.70 -0.06                0.05
0.50 0.50 -0.10                0.02
```

```
df3_mc = cp3.MCsimulation()
df3_mc
```

```
                   MSE Ratio: p = 0.6 MSE Ratio: p = 0.7 MSE Ratio: p = 0.8  \
T_a  T_b  Bias
0.95 1.00 -0.029752               3.04               0.74               0.33
0.90 1.00 -0.059862               0.82                0.2               0.09
0.70 1.00 -0.179691               0.09               0.02               0.01
0.50 1.00 -0.299452               0.03               0.01                0.0
1.00 0.95  0.020192               5.85               1.43               0.64
     0.90  0.040388               1.75               0.43               0.19
     0.70  0.120138               0.21               0.05               0.02
     0.50  0.199962               0.08               0.02               0.01
0.95 0.95 -0.009622              14.35               3.51               1.57
0.90 0.90 -0.019695               6.05               1.48               0.66
0.70 0.70 -0.059774               0.83                0.2               0.09
0.50 0.50 -0.1                     0.3               0.07               0.03

                   MSE Ratio: p = 0.9
T_a  T_b  Bias
0.95 1.00 -0.029752               0.18
0.90 1.00 -0.059862               0.05
0.70 1.00 -0.179691               0.01
0.50 1.00 -0.299452                0.0
1.00 0.95  0.020192               0.35
     0.90  0.040388                0.1
     0.70  0.120138               0.01
     0.50  0.199962                0.0
0.95 0.95 -0.009622               0.85
0.90 0.90 -0.019695               0.36
0.70 0.70 -0.059774               0.05
0.50 0.50 -0.1                    0.02
```

Evidently, as n increases, the randomized response method does better performance in more situations.

16.4 Concluding Remarks

This QuantEcon lecture describes some alternative randomized response surveys.

That lecture presents a utilitarian analysis of those alternatives conducted by Lars Ljungqvist [Ljungqvist, 1993].

CHAPTER

SEVENTEEN

EXPECTED UTILITIES OF RANDOM RESPONSES

17.1 Overview

This QuantEcon lecture describes randomized response surveys in the tradition of Warner [1965] that are designed to protect respondents' privacy.

Lars Ljungqvist [Ljungqvist, 1993] analyzed how a respondent's decision about whether to answer truthfully depends on **expected utility**.

The lecture tells how Ljungqvist used his framework to shed light on alternative randomized response survey techniques proposed, for example, by [Lanke, 1975], [Lanke, 1976], [Leysieffer and Warner, 1976], [Anderson, 1976], [Fligner *et al.*, 1977], [Greenberg *et al.*, 1977], [Greenberg *et al.*, 1969].

We use the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

17.2 Privacy measures

We consider randomized response models with only two possible answers, "yes" and "no."

The design determines probabilities

$$\Pr(\text{yes}|A) = 1 - \Pr(\text{no}|A)$$
$$\Pr(\text{yes}|A^{'}) = 1 - \Pr(\text{no}|A^{'})$$

These design probabilities in turn can be used to compute the conditional probability of belonging to the sensitive group A for a given response, say r:

$$\Pr(A|r) = \frac{\pi_A \Pr(r|A)}{\pi_A \Pr(r|A) + (1 - \pi_A)\Pr(r|A^{'})} \tag{17.1}$$

17.3 Zoo of concepts

At this point, we describe some concepts proposed by various researchers.

17.3.1 Leysieffer and Warner (1976)

The response r is regarded as jeopardizing with respect to A or $A^{'}$ if

$$\begin{aligned} &\Pr(A|r) > \pi_A \\ &\text{or} \\ &\Pr(A^{'}|r) > 1 - \pi_A \end{aligned} \tag{17.2}$$

From Bayes's rule:

$$\frac{\Pr(A|r)}{\Pr(A^{'}|r)} \times \frac{(1-\pi_A)}{\pi_A} = \frac{\Pr(r|A)}{\Pr(r|A^{'})} \tag{17.3}$$

If this expression is greater (less) than unity, it follows that r is jeopardizing with respect to A ($A^{'}$).

Then, the natural measure of jeopardy will be:

$$\begin{aligned} g(r|A) &= \frac{\Pr(r|A)}{\Pr(r|A^{'})} \\ &\text{and} \\ g(r|A^{'}) &= \frac{\Pr(r|A^{'})}{\Pr(r|A)} \end{aligned} \tag{17.4}$$

Suppose, without loss of generality, that $\Pr(\text{yes}|A) > \Pr(\text{yes}|A^{'})$, then a yes (no) answer is jeopardizing with respect to A ($A^{'}$), that is,

$$\begin{aligned} &g(\text{yes}|A) > 1 \\ &\text{and} \\ &g(\text{no}|A^{'}) > 1 \end{aligned}$$

Leysieffer and Warner proved that the variance of the estimate can only be decreased through an increase in one or both of these two measures of jeopardy.

An efficient randomized response model is, therefore, any model that attains the maximum acceptable levels of jeopardy that are consistent with cooperation of the respondents.

As a special example, Leysieffer and Warner considered "a problem in which there is no jeopardy in a no answer"; that is, $g(\text{no}|A^{'})$ can be of unlimited magnitude.

Evidently, an optimal design must have

$$\Pr(\text{yes}|A) = 1$$

which implies that

$$\Pr(A|\text{no}) = 0$$

17.3.2 Lanke (1976)

Lanke [1975] argued that "it is membership in Group A that people may want to hide, not membership in the complementary Group A'."

For that reason, Lanke [1976] argued that an appropriate measure of protection is to minimize

$$\max \{\Pr(A|\text{yes}), \Pr(A|\text{no})\} \tag{17.5}$$

Holding this measure constant, he explained under what conditions the smallest variance of the estimate was achieved with the unrelated question model or the original model of Warner [1965].

17.3.3 Fligner et al. (1977)

Fligner *et al.* [1977] reached a conclusion similar to that of Lanke [1976].

They measured "private protection" as

$$\frac{1 - \max \{\Pr(A|\text{yes}), \Pr(A|\text{no})\}}{1 - \pi_A} \tag{17.6}$$

17.3.4 Greenberg et al. (1977)

Greenberg *et al.* [1977] stressed the importance of examining the risk to respondents who do not belong to A as well as the risk to those who do belong to the sensitive group.

They defined the hazard for an individual in A as the probability that he or she is perceived as belonging to A:

$$\Pr(\text{yes}|A) \times \Pr(A|\text{yes}) + \Pr(\text{no}|A) \times \Pr(A|\text{no}) \tag{17.7}$$

Similarly, the hazard for an individual who does not belong to A would be

$$\Pr(\text{yes}|A^{'}) \times \Pr(A|\text{yes}) + \Pr(\text{no}|A^{'}) \times \Pr(A|\text{no}) \tag{17.8}$$

Greenberg *et al.* [1977] also considered an alternative related measure of hazard that "is likely to be closer to the actual concern felt by a respondent."

The "limited hazard" for an individual in A and $A^{'}$ is

$$\Pr(\text{yes}|A) \times \Pr(A|\text{yes}) \tag{17.9}$$

and

$$\Pr(\text{yes}|A^{'}) \times \Pr(A|\text{yes}) \tag{17.10}$$

This measure is just the first term in (17.7), i.e., the probability that an individual answers "yes" and is perceived to belong to A.

17.4 Respondent's expected utility

17.4.1 Truth border

Key assumptions that underlie a randomized response technique for estimating the fraction of a population that belongs to A are

Assumption 17.4.1 (discomfort)

Respondents feel discomfort from being thought of as belonging to A.

Assumption 17.4.2 (preference for truth)

Respondents prefer to answer questions truthfully than to lie, so long as the cost of doing so is not too high.

The cost is taken to be the discomfort in *Assumption 17.4.1*.

Let r_i denote individual i's response to the randomized question.

r_i can only take values "yes" or "no".

For a given design of a randomized response interview and a given belief about the fraction of the population that belongs to A, the respondent's answer is associated with a conditional probability $\Pr(A|r_i)$ that the individual belongs to A.

Given r_i and complete privacy, the individual's utility is higher if r_i represents a truthful answer rather than a lie.

In terms of a respondent's expected utility as a function of $\Pr(A|r_i)$ and r_i

- The higher is $\Pr(A|r_i)$, the lower is individual i's expected utility.
- Expected utility is higher if r_i represents a truthful answer rather than a lie.

Define:

- $\phi_i \in \{\text{truth}, \text{lie}\}$, a dichotomous variable that indicates whether or not r_i is a truthful statement.
- $U_i\left(\Pr(A|r_i), \phi_i\right)$, a utility function that is differentiable in its first argument, summarizes individual i's expected utility.

Then there is an r_i such that

$$\frac{\partial U_i\left(\Pr(A|r_i), \phi_i\right)}{\partial \Pr(A|r_i)} < 0, \text{ for } \phi_i \in \{\text{truth}, \text{lie}\} \tag{17.11}$$

and

$$U_i\left(\Pr(A|r_i), \text{truth}\right) > U_i\left(\Pr(A|r_i), \text{lie}\right), \text{ for } \Pr(A|r_i) \in [0, 1] \tag{17.12}$$

Suppose now that the correct answer for individual i is "yes".

Individual i would choose to answer truthfully if

$$U_i\left(\Pr(A|\text{yes}), \text{truth}\right) \geq U_i\left(\Pr(A|\text{no}), \text{lie}\right) \tag{17.13}$$

If the correct answer is "no", individual i would volunteer the correct answer only if

$$U_i\left(\Pr(A|\text{no}), \text{truth}\right) \geq U_i\left(\Pr(A|\text{yes}), \text{lie}\right) \tag{17.14}$$

Assume that

$$\Pr(A|\text{yes}) > \pi_A > \Pr(A|\text{no})$$

so that a "yes" answer increases the odds that an individual belongs to A.

Constraint (17.14) holds for sure.

Consequently, constraint (17.13) becomes the single necessary condition for individual i always to answer truthfully.

At equality, constraint (17.13) determines the conditional probabilities that make the individual indifferent between telling the truth and lying when the correct answer is "yes":

$$U_i\left(\Pr(A|\text{yes}), \text{truth}\right) = U_i\left(\Pr(A|\text{no}), \text{lie}\right) \tag{17.15}$$

Equation (17.15) defines a "truth border".

Differentiating (17.15) with respect to the conditional probabilities shows that the truth border has a positive slope in the space of conditional probabilities:

$$\frac{\partial \Pr(A|\text{no})}{\partial \Pr(A|\text{yes})} = \frac{\frac{\partial U_i(\Pr(A|\text{yes}),\text{truth})}{\partial \Pr(A|\text{yes})}}{\frac{\partial U_i(\Pr(A|\text{no}),\text{lie})}{\partial \Pr(A|\text{no})}} > 0 \tag{17.16}$$

The source of the positive relationship is:

- The individual is willing to volunteer a truthful "yes" answer so long as the utility from doing so (i.e., the left side of (17.15)) is at least as high as the utility of lying on the right side of (17.15).
- Suppose now that $\Pr(A|\text{yes})$ increases.
  - This reduces the utility of telling the truth.
  - To preserve indifference between a truthful answer and a lie, $\Pr(A|\text{no})$ must increase to reduce the utility of lying.

17.4.2 Drawing a truth border

We can deduce two things about the truth border:

- The truth border divides the space of conditional probabilities into two subsets: "truth telling" and "lying".
  - Thus, sufficient privacy elicits a truthful answer, whereas insufficient privacy results in a lie.
  - The truth border depends on a respondent's utility function.
- Assumptions in (17.11) and (17.12) are sufficient only to guarantee a positive slope of the truth border.
  - The truth border can have either a concave or a convex shape.

We can draw some truth borders in Fig. 17.1 with the following Python code, writing p for $\Pr(A|r_i)$:

```
x1 = np.arange(0, 1, 0.001)
y1 = x1 - 0.4
x2 = np.arange(0.4**2, 1, 0.001)
y2 = (pow(x2, 0.5) - 0.4)**2
x3 = np.arange(0.4**0.5, 1, 0.001)
y3 = pow(x3**2 - 0.4, 0.5)

fig, ax = plt.subplots()
ax.plot(x1, y1, 'r-', lw=2, label=r'$U_i = -p + f(\phi_i)$')
ax.fill_between(x1, 0, y1, facecolor='red', alpha=0.05)
ax.plot(x2, y2, 'b-', lw=2, label=r'$U_i = -p^2 + f(\phi_i)$')
ax.fill_between(x2, 0, y2, facecolor='blue', alpha=0.05)
ax.plot(x3, y3, 'y-', lw=2, label=r'$U_i = -\sqrt{p} + f(\phi_i)$')
ax.fill_between(x3, 0, y3, facecolor='green', alpha=0.05)
ax.plot(x1, x1, ':', lw=2)
ax.set(xlim=(0, 1), ylim=(0, 1),
       xlabel='Pr(A|yes)', ylabel='Pr(A|no)')
ax.text(0.42, 0.3, "truth telling")
ax.text(0.8, 0.1, "lying")
```

(continues on next page)

(continued from previous page)

```
ax.legend(loc='upper left')
plt.show()
```

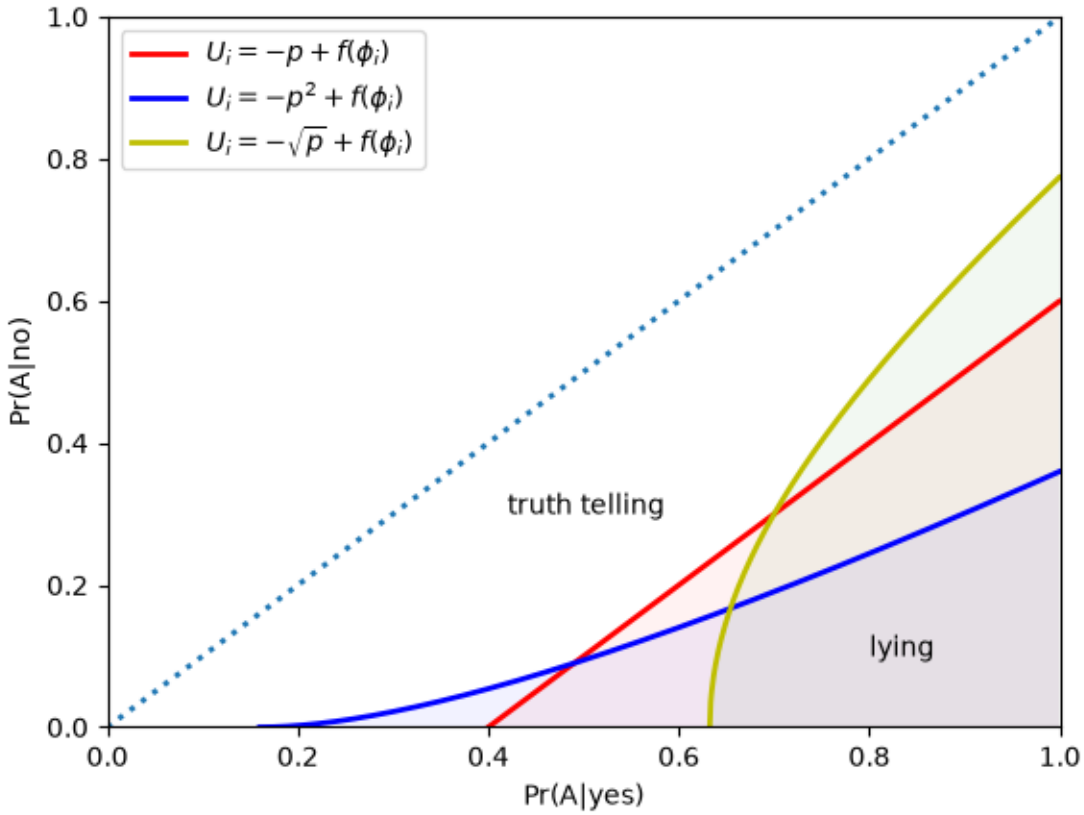


Fig. 17.1: Three types of truth border

Without loss of generality, we consider the truth border:

$$U_i(\Pr(A|r_i), \phi_i) = -\Pr(A|r_i) + f(\phi_i)$$

and plot the "truth telling" and "lying area" of individual i in Fig. 17.2:

```
x1 = np.arange(0, 1, 0.001)
y1 = x1 - 0.4
z1 = x1
z2 = 0

fig, ax = plt.subplots()
ax.plot(x1, y1, 'r-', lw=2, label='truth border')
ax.plot(x1, x1, ':', lw=2)
ax.fill_between(x1, y1, z1, facecolor='blue', alpha=0.05,
                label='truth telling')
ax.fill_between(x1, z2, y1, facecolor='green', alpha=0.05,
                label='lying')
ax.set(xlim=(0, 1), ylim=(0, 1),
       xlabel='Pr(A|yes)', ylabel='Pr(A|no)')
```

(continues on next page)

(continued from previous page)

```
ax.text(0.5, 0.4, "truth telling")
ax.text(0.8, 0.2, "lying")
ax.legend(loc='upper left')
plt.show()
```

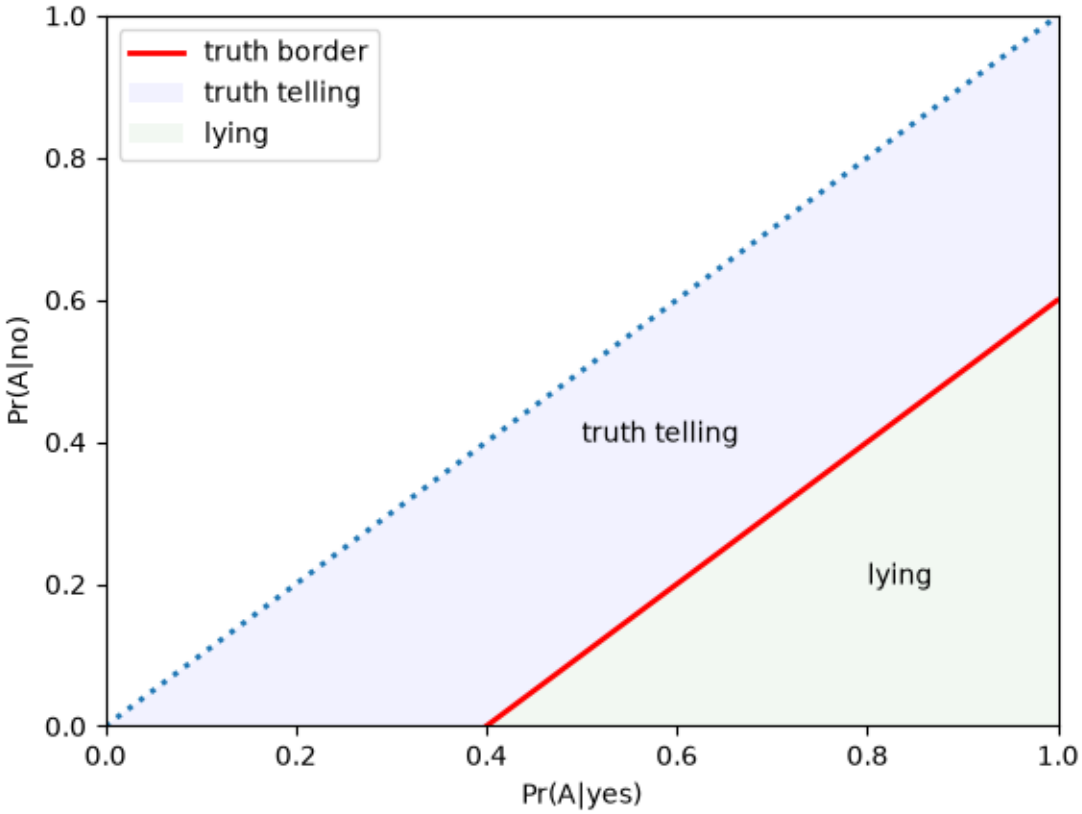


Fig. 17.2: Truth telling and lying areas

17.5 Utilitarian view of survey design

17.5.1 Iso-variance curves

A statistician's objective is

- to find a randomized response survey design that minimizes the bias and the variance of the estimator.

Given a design that ensures truthful answers by all respondents, Anderson [1976] (Theorem 1) showed that the minimum variance estimate in the two-response model has variance

$$V(\Pr(A|\text{yes}), \Pr(A|\text{no})) = \frac{{\pi_A}^2(1-\pi_A)^2}{n} \times \frac{1}{\Pr(A|\text{yes}) - \pi_A} \times \frac{1}{\pi_A - \Pr(A|\text{no})} \tag{17.17}$$

where the random sample with replacement consists of n individuals.

We can use Expression (17.17) to draw iso-variance curves.

The following inequalities restrict the shapes of iso-variance curves:

$$\left.\frac{d\,\Pr(A|\text{no})}{d\,\Pr(A|\text{yes})}\right|_{\text{constant variance}} = \frac{\pi_A - \Pr(A|\text{no})}{\Pr(A|\text{yes}) - \pi_A} > 0 \tag{17.18}$$

$$\left.\frac{d^2\,\Pr(A|\text{no})}{d\,\Pr(A|\text{yes})^2}\right|_{\text{constant variance}} = -\frac{2\left[\pi_A - \Pr(A|\text{no})\right]}{\left[\Pr(A|\text{yes}) - \pi_A\right]^2} < 0 \tag{17.19}$$

From expression (17.17), (17.18) and (17.19) we can see that:

- Variance can be reduced only by increasing the distance of $\Pr(A|\text{yes})$ and/or $\Pr(A|\text{no})$ from r_A.
- Iso-variance curves are always upward-sloping and concave.

17.5.2 Drawing iso-variance curves

We use Python code to draw iso-variance curves.

The pairs of conditional probabilities can be attained using the model of Warner [1965].

Note that:

- Any point on the iso-variance curves can be attained with the unrelated question model as long as the statistician can completely control the model design.
- The original randomized response model of Warner [1965] is less flexible than the unrelated question model.

```
def plot_iso_variance_curves(π, n):
    nv = np.array([0.27, 0.34, 0.49, 0.74, 0.92, 1.1, 1.47, 2.94, 14.7])
    x = np.arange(0, 1, 0.001)
    x0 = np.arange(π, 1, 0.001)
    x2 = np.arange(0, π, 0.001)
    y1 = [π for i in x0]
    y2 = [π for i in x2]
    y0 = 1 / (1 + (x0 * (1 - π)**2) / ((1 - x0) * π**2))

    fig, ax = plt.subplots()
    ax.plot(x0, y0, 'm-', lw=2, label='Warner')
    ax.plot(x, x, 'c:', lw=2)
    ax.plot(x0, y1, 'c:', lw=2)
    ax.plot(y2, x2, 'c:', lw=2)
    for i in range(len(nv)):
        y = π - (π**2 * (1 - π)**2) / (n * (nv[i] / n) * (x0 - π + 1e-8))
        ax.plot(x0, y, 'k--', alpha=1 - 0.07 * i,
                label='iso-variance curves' if i == 0 else None)
    ax.set(xlim=(0, 1), ylim=(0, 0.5),
           xlabel='Pr(A|yes)', ylabel='Pr(A|no)')
    ax.text(0.42, 0.32, "high variance")
    ax.text(0.83, 0.01, "low variance")
    ax.legend(loc='upper left')
    plt.show()
```

Properties of iso-variance curves are:

- All points on one iso-variance curve share the same variance
- Variance increases as the curves move up and to the left, and the curves lighten as variance increases

Suppose the parameters of the iso-variance model follow those in Ljungqvist [Ljungqvist, 1993], which are:

- $\pi = 0.3$
- $n = 100$

Then we can plot the iso-variance curve in Fig. 17.3:

```
plot_iso_variance_curves(π=0.3, n=100)
```

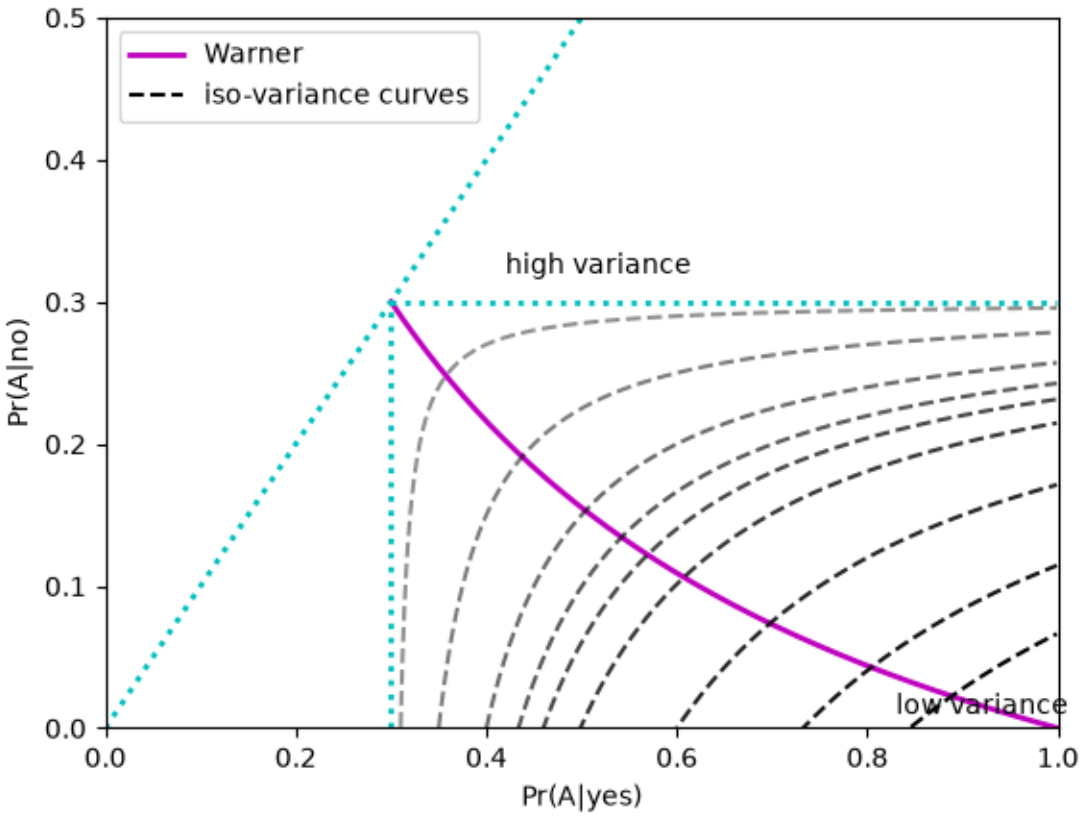


Fig. 17.3: Iso-variance curves

17.5.3 Optimal survey

A point on iso-variance curves can be attained with the unrelated question design.

We now focus on finding an "optimal survey design" that

- Minimizes the variance of the estimator subject to privacy restrictions.

To obtain an optimal design, we first superimpose all individuals' truth borders on the iso-variance mapping.

To construct an optimal design

- The statistician should find the intersection of areas above all truth borders; that is, the set of conditional probabilities ensuring truthful answers from all respondents.
- The point where this set touches the lowest possible iso-variance curve determines an optimal survey design.

Consequently, a minimum variance unbiased estimator is pinned down by an individual who is the least willing to volunteer a truthful answer.

Here are some comments about the model design:

- An individual's decision of whether or not to answer truthfully depends on his or her belief about other respondents' behavior, because this determines the individual's calculation of $\Pr(A|\text{yes})$ and $\Pr(A|\text{no})$.
- An equilibrium of the optimal design model is a Nash equilibrium of a noncooperative game.
- Assumption (17.12) is sufficient to guarantee the existence of an optimal model design.
  - By choosing $\Pr(A|\text{yes})$ and $\Pr(A|\text{no})$ sufficiently close to each other, all respondents will find it optimal to answer truthfully.
  - The closer are these probabilities, the higher the variance of the estimator becomes.
- If respondents experience a large enough increase in expected utility from telling the truth, then there is no need to use a randomized response model.

The smallest possible variance of the estimate is then obtained at $\Pr(A|\text{yes}) = 1$ and $\Pr(A|\text{no}) = 0$; that is, when respondents answer truthfully to direct questioning.

- A more general design problem would be to minimize some weighted sum of the estimator's variance and bias.
  - It would be optimal to accept some lies from the most "reluctant" respondents.

17.6 Criticisms of proposed privacy measures

We can use a utilitarian approach to analyze some privacy measures.

We'll enlist Python code to help us.

17.6.1 Analysis of the method of Lanke (1976)

Lanke [1976] recommends a privacy protection criterion that minimizes:

$$\max\{\Pr(A|\text{yes}), \Pr(A|\text{no})\} \tag{17.20}$$

Following Lanke's suggestion, the statistician should find the highest possible $\Pr(A|\text{yes})$ consistent with truth telling while $\Pr(A|\text{no})$ is fixed at 0.

The variance is then minimized at point X in Fig. 17.4.

However, as shown in Fig. 17.4, point Z offers a smaller variance that still allows cooperation of the respondents, and it is achievable following our earlier discussion of the truth border:

```
π = 0.3
n = 100
nv = [0.27, 0.34, 0.49, 0.74, 0.92, 1.1, 1.47, 2.94, 14.7]
x = np.arange(0, 1, 0.001)
y = x - 0.4
z = x
x0 = np.arange(π, 1, 0.001)
x2 = np.arange(0, π, 0.001)
y1 = [π for i in x0]
y2 = [π for i in x2]
```

(continues on next page)

(continued from previous page)

```
fig, ax = plt.subplots()
ax.plot(x, x, 'c:', lw=2)
ax.plot(x0, y1, 'c:', lw=2)
ax.plot(y2, x2, 'c:', lw=2)
ax.plot(x, y, 'r-', lw=2, label='truth border')
ax.fill_between(x, y, z, facecolor='blue', alpha=0.05,
                label='truth telling')
ax.fill_between(x, 0, y, facecolor='green', alpha=0.05,
                label='lying')
for i in range(len(nv)):
    y = π - (π**2 * (1 - π)**2) / (n * (nv[i] / n) * (x0 - π + 1e-8))
    ax.plot(x0, y, 'k--', alpha=1 - 0.07 * i,
            label='iso-variance curves' if i == 0 else None)
ax.scatter(0.498, 0.1, c='b', marker='*', s=150, label='Z')
ax.scatter(0.4, 0, c='y', s=150, label='X')
ax.set(xlim=(0, 1), ylim=(0, 0.5),
       xlabel='Pr(A|yes)', ylabel='Pr(A|no)')
ax.text(0.45, 0.4, "truth telling")
ax.text(0.85, 0.4, "lying")
ax.text(0.515, 0.095, "optimal design", color='b')
ax.legend(loc='upper left')
plt.show()
```

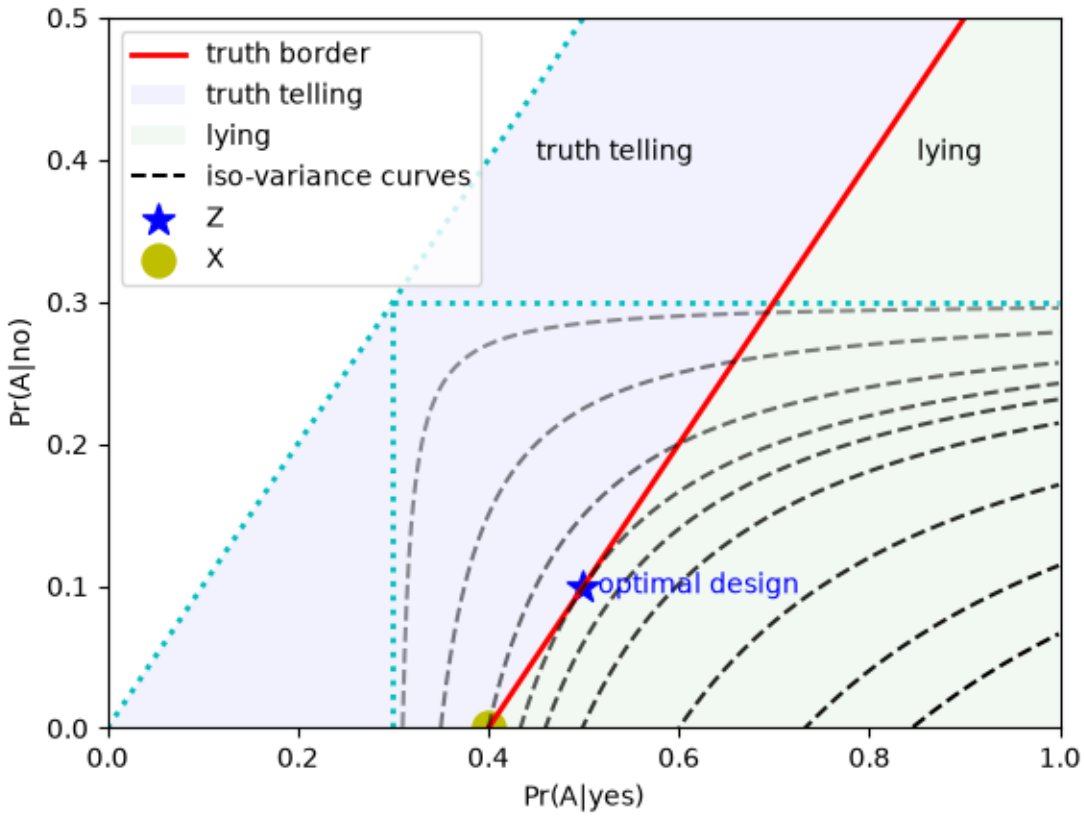


Fig. 17.4: Optimal design under Lanke's criterion

17.6.2 Method of Leysieffer and Warner (1976)

Leysieffer and Warner [1976] recommend a two-dimensional measure of jeopardy that reduces to a single dimension when there is no jeopardy in a 'no' answer, which means that

$$\Pr(\text{yes}|A) = 1$$

and

$$\Pr(A|\text{no}) = 0$$

This is not an optimal choice under a utilitarian approach.

17.6.3 Analysis of the method of Chaudhuri and Mukerjee (1988)

Chaudhuri and Mukerjee [1988] argued that since "yes" may sometimes relate to the sensitive group A, a clever respondent may falsely but safely always be inclined to respond "no".

In this situation, the truth border is such that individuals choose to lie whenever the truthful answer is "yes" and

$$\Pr(A|\text{no}) = 0$$

Here the gain from lying is too high for someone to volunteer a "yes" answer.

This means that

$$U_i\left(\Pr(A|\text{yes}), \text{truth}\right) < U_i\left(\Pr(A|\text{no}), \text{lie}\right)$$

always holds.

As a result, there is no attainable model design.

However, under a utilitarian approach there should exist other survey designs that are consistent with truthful answers.

In particular, respondents will choose to answer truthfully if the relative advantage from lying is eliminated.

We can use Python to show the optimal model design.

```
def f(x):
    "Truth border under the Chaudhuri and Mukerjee assumption."
    return np.where(x < 0.16, 0.0, (np.sqrt(x) - 0.4)**2)
```

```
π = 0.3
n = 100
nv = [0.27, 0.34, 0.49, 0.74, 0.92, 1.1, 1.47, 2.94, 14.7]
x = np.arange(0, 1, 0.001)
y = f(x)
z = x
x0 = np.arange(π, 1, 0.001)
x2 = np.arange(0, π, 0.001)
y1 = [π for i in x0]
y2 = [π for i in x2]
x3 = np.arange(0.16, 1, 0.001)

fig, ax = plt.subplots()
ax.plot(x, x, 'c:', lw=2)
ax.plot(x0, y1, 'c:', lw=2)
```

(continues on next page)

(continued from previous page)

```
ax.plot(y2, x2, 'c:', lw=2)
ax.plot(x3, f(x3), 'b-', lw=2, label='truth border')
ax.fill_between(x, y, z, facecolor='blue', alpha=0.05,
                label='truth telling')
ax.fill_between(x3, 0, f(x3), facecolor='green', alpha=0.05,
                label='lying')
for i in range(len(nv)):
    y = π - (π**2 * (1 - π)**2) / (n * (nv[i] / n) * (x0 - π + 1e-8))
    ax.plot(x0, y, 'k--', alpha=1 - 0.07 * i,
            label='iso-variance curves' if i == 0 else None)
ax.scatter(0.61, 0.146, c='r', marker='*', s=150, label='Q')
ax.set(xlim=(0, 1), ylim=(0, 0.5),
       xlabel='Pr(A|yes)', ylabel='Pr(A|no)')
ax.text(0.45, 0.4, "truth telling")
ax.text(0.85, 0.1, "lying")
ax.text(0.63, 0.141, "optimal design", color='r')
ax.legend(loc='upper left')
plt.show()
```

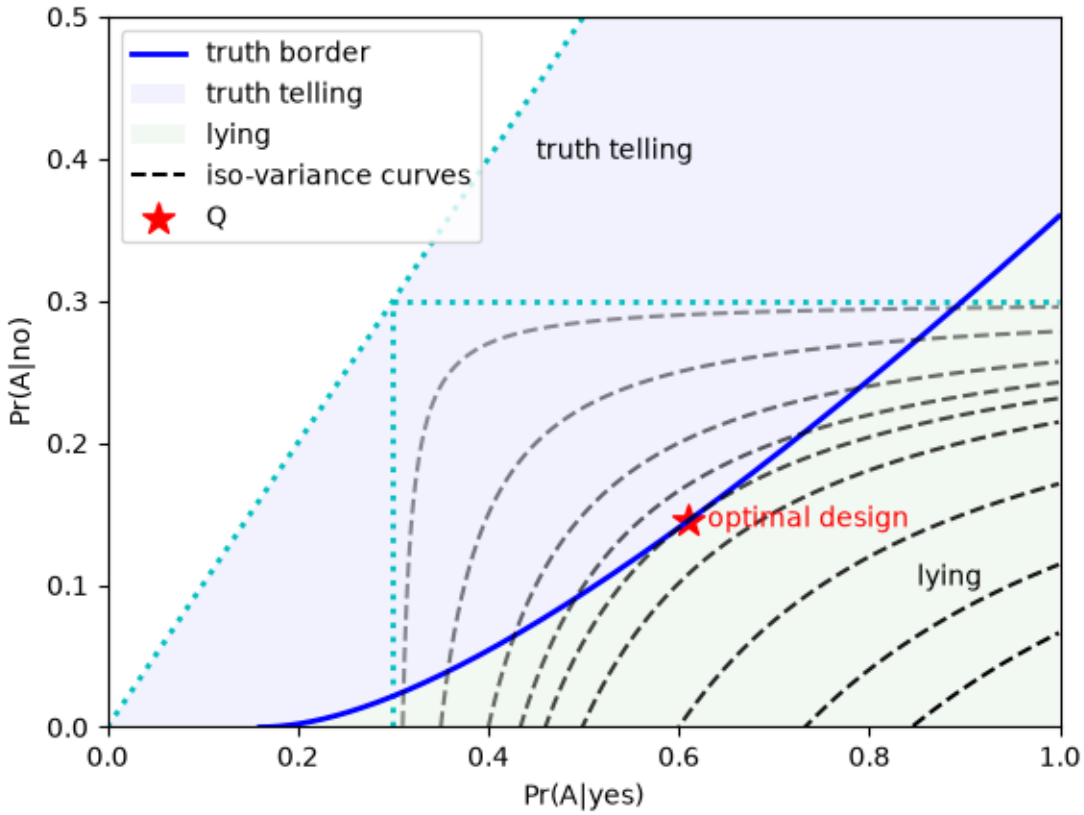


Fig. 17.5: Optimal design under a utilitarian approach

Here the optimal model design corresponds to point Q in Fig. 17.5.

17.6.4 Method of Greenberg et al. (1977)

Greenberg *et al.* [1977] defined the hazard for an individual in A as the probability that he or she is perceived as belonging to A:

$$\Pr(\text{yes}|A) \times \Pr(A|\text{yes}) + \Pr(\text{no}|A) \times \Pr(A|\text{no}) \tag{17.21}$$

The hazard for an individual who does not belong to A is

$$\Pr(\text{yes}|A^{'}) \times \Pr(A|\text{yes}) + \Pr(\text{no}|A^{'}) \times \Pr(A|\text{no}) \tag{17.22}$$

They also considered an alternative related measure of hazard that they said "is likely to be closer to the actual concern felt by a respondent."

Their "limited hazard" for an individual in A and $A^{'}$ is

$$\Pr(\text{yes}|A) \times \Pr(A|\text{yes}) \tag{17.23}$$

and

$$\Pr(\text{yes}|A^{'}) \times \Pr(A|\text{yes}) \tag{17.24}$$

According to Greenberg *et al.* [1977], a respondent commits himself or herself to answer truthfully on the basis of a probability in (17.21) or (17.23) **before** randomly selecting the question to be answered.

Suppose that the appropriate privacy measure is captured by the notion of "limited hazard" in (17.23) and (17.24).

Consider an unrelated question model where the unrelated question is replaced by the instruction "Say the word 'no'", which implies that

$$\Pr(A|\text{yes}) = 1$$

and it follows that:

- Hazard for an individual in $A^{'}$ is 0.
- Hazard for an individual in A can also be made arbitrarily small by choosing a sufficiently small $\Pr(\text{yes}|A)$.

Even though this hazard can be set arbitrarily close to 0, an individual in A will completely reveal his or her identity whenever truthfully answering the sensitive question.

However, under a utilitarian framework, this is obviously contradictory.

If the individuals are willing to volunteer this information, it seems that the randomized response design was not necessary in the first place.

It ignores the fact that respondents retain the option of lying until they have seen the question to be answered.

17.7 Concluding remarks

The justifications for a randomized response procedure are that

- Respondents are thought to feel discomfort from being perceived as belonging to the sensitive group.
- Respondents prefer to answer questions truthfully than to lie, unless it is too revealing.

If a privacy measure is not completely consistent with the rational behavior of the respondents, all efforts to derive an optimal model design are futile.

A utilitarian approach provides a systematic way to model respondents' behavior under the assumption that they maximize their expected utilities.

In a utilitarian analysis:

- A truth border divides the space of conditional probabilities of being perceived as belonging to the sensitive group, $\Pr(A|\text{yes})$ and $\Pr(A|\text{no})$, into the truth-telling region and the lying region.
- The optimal model design is obtained at the point where the truth border touches the lowest possible iso-variance curve.

A practical implication of the analysis of [Ljungqvist, 1993] is that uncertainty about respondents' demands for privacy can be acknowledged by *choosing Pr*$(A|yes)$ *and Pr*$(A|no)$ *sufficiently close to each other*.

Part III

Bayes Law

CHAPTER

EIGHTEEN

NON-CONJUGATE PRIORS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install numpyro jax arviz
```

18.1 Overview

This lecture is a sequel to *Two Meanings of Probability*.

In that lecture we adopted a **beta** prior for the unknown probability θ of a coin landing heads, together with a **binomial** likelihood.

That prior and likelihood form a **conjugate pair**: applying Bayes' law returns a posterior of the *same* family as the prior — again a beta distribution.

Conjugacy is convenient because it delivers a posterior in closed form.

But a person's prior beliefs are their own business, and in general they will not happen to be conjugate to the likelihood.

When the prior and likelihood are **not** conjugate, the posterior usually has no closed form, and we must approximate it numerically.

This lecture introduces two widely used ways to do that, both implemented in the probabilistic programming library NumPyro:

- **Markov chain Monte Carlo (MCMC)** — construct a Markov chain whose stationary distribution is the posterior, then sample from it. We use the **No-U-Turn Sampler (NUTS)**, a state-of-the-art form of Hamiltonian Monte Carlo.
- **Variational inference (VI)** — replace sampling with optimization: search within a tractable family of distributions for the member closest to the posterior.

> **Note**
>
> We treat NUTS as a black box in this lecture.
>
> In brief, it is a form of **Hamiltonian Monte Carlo**, which is itself a version of the **Metropolis–Hastings** algorithm: it proposes candidate draws and accepts or rejects them so that the resulting Markov chain has the posterior as its stationary distribution.
>
> What distinguishes it from a basic Metropolis–Hastings sampler is that its proposals are built from *gradient* (derivative) information about the log-posterior, which lets the chain move efficiently through the parameter space; NUTS additionally tunes the length of each proposed move automatically.
>
> For a more advanced introduction to MCMC and the Metropolis–Hastings algorithm, see this lecture.

Our plan is:

1. Confirm that MCMC reproduces the *conjugate* beta posterior that we can compute analytically — this validates the machinery on a problem whose answer we already know.
2. Replace the beta prior with several **non-conjugate** priors and approximate each posterior with MCMC.
3. Introduce variational inference and compare it with MCMC.

Let us start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
import scipy.stats as st

import jax.numpy as jnp
from jax import random

import numpyro
import numpyro.distributions as dist
from numpyro.infer import MCMC, NUTS, SVI, Trace_ELBO
from numpyro.infer.autoguide import AutoNormal
from numpyro.optim import Adam

import arviz as az
```

18.2 The coin-flipping model

As in *Two Meanings of Probability*, a coin lands heads ($Y = 1$) with probability θ and tails ($Y = 0$) with probability $1 - \theta$.

If we flip the coin n times, the number of heads k has the **binomial** distribution

$$p(k \mid \theta) = \binom{n}{k} \theta^k (1-\theta)^{n-k}.$$

We treat θ as a random variable with a prior density $p(\theta)$, and we want the posterior

$$p(\theta \mid k) \propto p(k \mid \theta)\, p(\theta).$$

18.2.1 Generating data

We simulate a sequence of coin flips from a coin whose true (but unknown to the analyst) probability of heads is $\theta = 0.4$.

```
def simulate_coin_flips(θ=0.4, n=20, seed=1234):
    "Flip a coin n times; return an array of 0s (tails) and 1s (heads)."
    rng = np.random.default_rng(seed)
    return (rng.random(n) < θ).astype(int)

data = simulate_coin_flips()
k, n = int(data.sum()), len(data)
k, n
```

```
(9, 20)
```

We deliberately use a **small** sample ($n = 20$).

The reason is that the prior matters most when data are scarce.

With a large sample the likelihood dominates and almost any reasonable prior leads to the same posterior — exactly the concentration we saw in *Two Meanings of Probability*.

A modest n keeps the influence of the prior visible, which is what we want to study here.

18.2.2 Specifying the model in NumPyro

For most readers this will be a first encounter with NumPyro, whose style takes some getting used to.

To use it we describe our probability model as a Python function — which, a little confusingly, NumPyro calls a **model**.

Such a function does not *compute* anything when called, and it does not return the posterior.

Instead it is a *declaration* of the generative story for the data: which quantities are random, how they are distributed, and how the data depend on them.

An inference algorithm — such as the NUTS sampler below — then *reads* this declaration and works out the posterior for us.

Inside a model, every random quantity is introduced by a call to `numpyro.sample`, and the keyword `obs` decides its role:

- `numpyro.sample("θ", prior)` introduces a **latent** (unobserved) variable named `"θ"`, drawn from `prior` — a quantity we wish to infer.
- `numpyro.sample("k", dist.Binomial(n, θ), obs=k)` introduces an **observed** variable: the keyword `obs=k` pins it to the data, which is how the likelihood $p(k \mid \theta)$ enters.

The string names (`"θ"` and `"k"`) are the labels NumPyro uses to keep track of the variables; we will use them later to pull the posterior draws back out.

We write a *single* model that takes the prior distribution as an argument, so we can reuse it unchanged for every prior we consider — conjugate or not.

```
def binomial_model(prior, k, n):
    "Binomial likelihood with a caller-supplied prior on θ."
    θ = numpyro.sample("θ", prior)
    numpyro.sample("k", dist.Binomial(n, θ), obs=k)
```

Notice that `binomial_model` returns nothing, and that we never call it ourselves.

Instead we hand it to an inference algorithm, which supplies the arguments and traces the two `sample` statements to assemble the posterior.

We also write a small helper that runs NUTS on a given model and returns the fitted sampler.

We request four chains so that we can check convergence below, and run them with `chain_method="vectorized"`, which evaluates all chains together on a single device — so the same code runs unchanged on a CPU or a GPU.

```
def run_nuts(model, *args, seed=0, num_warmup=1000, num_samples=4000, num_chains=4):
    "Sample a NumPyro model with the NUTS sampler."
    mcmc = MCMC(
        NUTS(model),
        num_warmup=num_warmup,
        num_samples=num_samples,
        num_chains=num_chains,
        chain_method="vectorized",
        progress_bar=False,
    )
    mcmc.run(random.key(seed), *args)
    return mcmc
```

NumPyro is built on JAX, which treats randomness explicitly: rather than relying on a global random state, each run needs its own **PRNG key**, created here with `random.key(seed)`.

(This is why we used NumPy's generator to make the data above but JAX keys here.)

`run_nuts` is deliberately generic: it samples whatever model we pass and forwards the extra arguments (`*args`) on to that model through `mcmc.run`. We always call it as `run_nuts(binomial_model, prior, k, n)`, so `prior`, `k`, and `n` reach `binomial_model` unchanged — there is only ever the one prior.

18.3 MCMC reproduces the conjugate posterior

Before trusting MCMC on hard problems, let us check it on an easy one.

With a Beta(α_0, β_0) prior the posterior is known analytically (see *Two Meanings of Probability*):

$$\theta \mid k \sim \text{Beta}(\alpha_0 + k,\ \beta_0 + n - k).$$

We take $\alpha_0 = \beta_0 = 2$ and sample the posterior with NUTS.

```
α0, β0 = 2.0, 2.0
mcmc = run_nuts(binomial_model, dist.Beta(α0, β0), k, n)
```

Before looking at the posterior we should check that the sampler has done its job.

Unlike the independent draws we are used to, MCMC returns a *dependent* sequence — a Markov chain — whose early draws still remember where the chain started.

We can trust the output only once the chain has "forgotten" its starting point and settled into its stationary distribution, which by construction is the posterior we want.

As a safeguard we ran **four** chains from different random starting points (`num_chains=4` in `run_nuts`) and now check that they agree with one another.

ArviZ is a companion library for examining the output of Bayesian samplers.

The function `az.from_numpyro` repackages our NumPyro results into ArviZ's standard data structure, and `az.summary` prints a table of per-parameter summaries and convergence diagnostics.

```
idata = az.from_numpyro(mcmc)
az.summary(idata, var_names=["θ"])
```

```
     mean    sd eti89_lb eti89_ub ess_bulk ess_tail r_hat mcse_mean  mcse_sd
θ   0.461   0.1       0.3      0.63     5944     7270  1.00     0.0013  0.00084
```

Two columns of this table are convergence diagnostics worth understanding.

- **`r_hat`** (the Gelman–Rubin statistic) compares the spread of the draws *within* each chain to the spread *between* chains. If the chains have all converged to the same distribution these two match and `r_hat` is close to 1.0; values above roughly 1.01 warn that the chains disagree and the draws cannot yet be trusted.
- **`ess_bulk`** and **`ess_tail`** report the *effective sample size*. Because consecutive MCMC draws are correlated, a chain of length N carries less information than N independent draws would; the effective sample size estimates how many independent draws it is worth (in the bulk and in the tails of the distribution respectively). Larger is better.

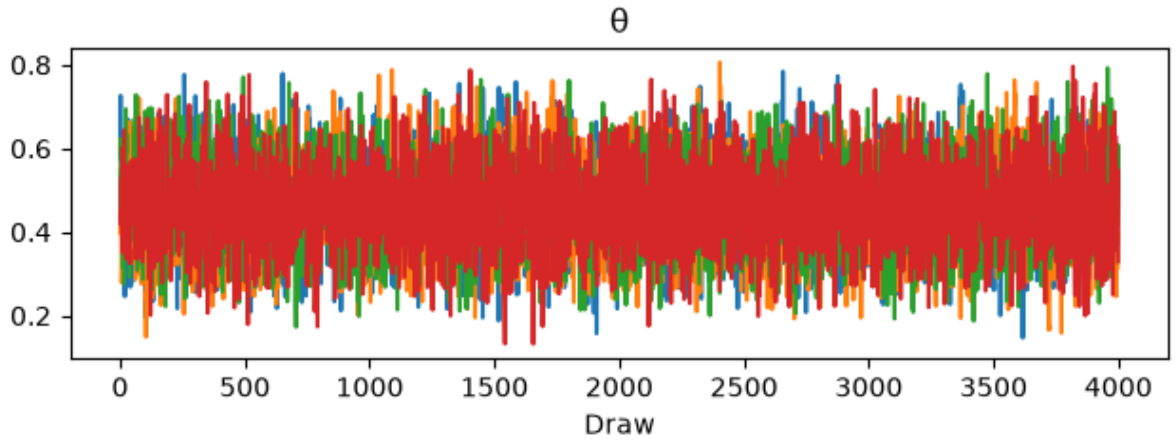
Here `r_hat` is essentially 1.0 and the effective sample sizes run into the thousands, so the chains have mixed well.

A **trace plot** gives a visual check of the same thing.

ArviZ's `plot_trace` draws two panels for each parameter: on the right, the sampled value against the iteration number (one coloured line per chain); on the left, a density estimate of the draws from each chain.

Well-mixed chains look like stationary noise on the right — a fuzzy, flat band, with the chains overlapping rather than drifting or wandering — and their densities on the left lie almost on top of one another.

```
az.plot_trace(idata, var_names=["θ"])
plt.tight_layout()
plt.show()
```



Our chains pass both checks, so we can trust the draws and turn to the posterior itself.

Now we compare the MCMC posterior with the analytical beta posterior.

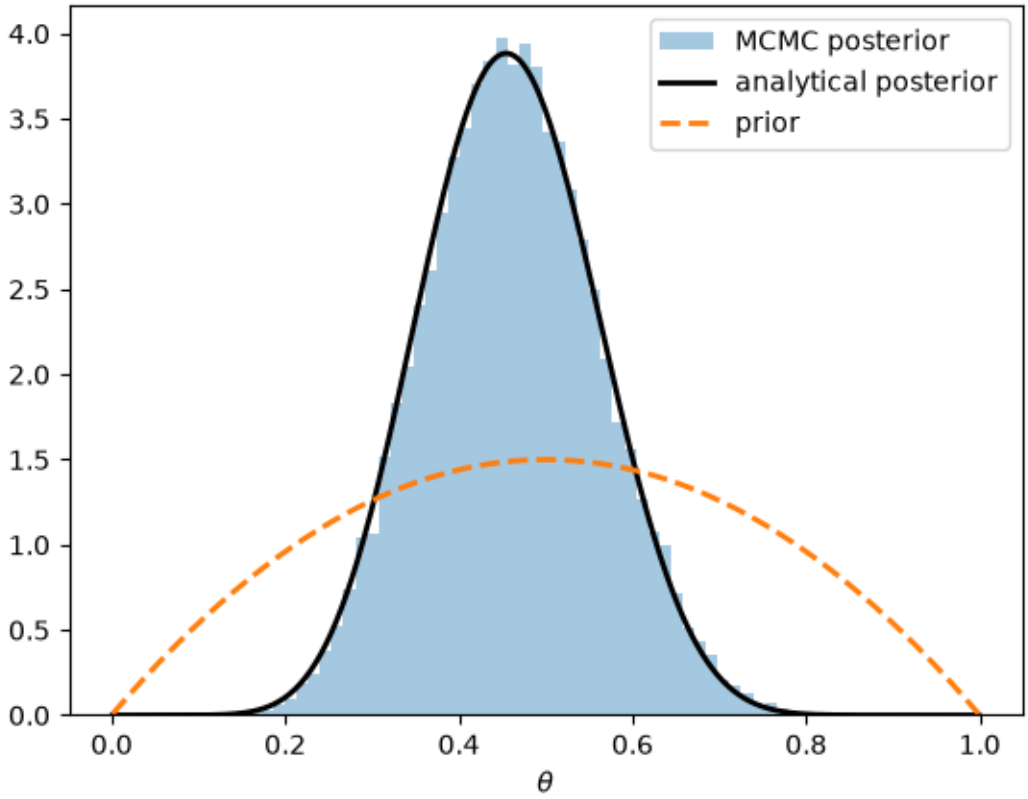
```
θ_grid = np.linspace(0.001, 0.999, 500)
samples = np.asarray(mcmc.get_samples()["θ"])

fig, ax = plt.subplots()
ax.hist(samples, bins=50, density=True, alpha=0.4,
        label="MCMC posterior")
ax.plot(θ_grid, st.beta(α0 + k, β0 + n - k).pdf(θ_grid),
```

(continues on next page)

(continued from previous page)

```
        'k-', lw=2, label="analytical posterior")
ax.plot(θ_grid, st.beta(α0, β0).pdf(θ_grid),
        'C1--', lw=2, label="prior")
ax.set_xlabel(r"$\theta$")
ax.legend()
plt.show()
```



The histogram of MCMC draws sits right on top of the analytical posterior density.

The sampler works, so we can rely on it for priors that have no closed-form posterior.

18.4 Non-conjugate priors

We now keep the binomial likelihood and the same data, but replace the beta prior with priors that are **not** conjugate to it.

For each prior the recipe is identical:

1. describe the prior and build it as a NumPyro distribution,
2. pass it to `binomial_model` and run NUTS,
3. plot the prior against the resulting posterior.

The following helper draws a prior density and the posterior samples on the same axes.

```
def plot_prior_posterior(prior, samples, title=""):
    "Overlay a prior density and posterior MCMC draws for θ on [0, 1]."
    grid = jnp.linspace(0.001, 0.999, 500)
    # mask the density to the prior's support: dist.Uniform.log_prob
    # returns its constant value even outside [low, high]
    in_support = np.asarray(prior.support(grid))
    prior_pdf = np.where(in_support, np.exp(np.asarray(prior.log_prob(grid))), 0.0)

    fig, ax = plt.subplots()
    ax.hist(np.asarray(samples), bins=50, density=True, alpha=0.4,
            label="posterior (MCMC)")
    ax.plot(np.asarray(grid), prior_pdf, 'C1--', lw=2, label="prior")
    ax.set_xlabel(r"$\theta$")
    ax.set_xlim(0, 1)
    ax.legend()
    if title:
        ax.set_title(title)
    plt.show()
```

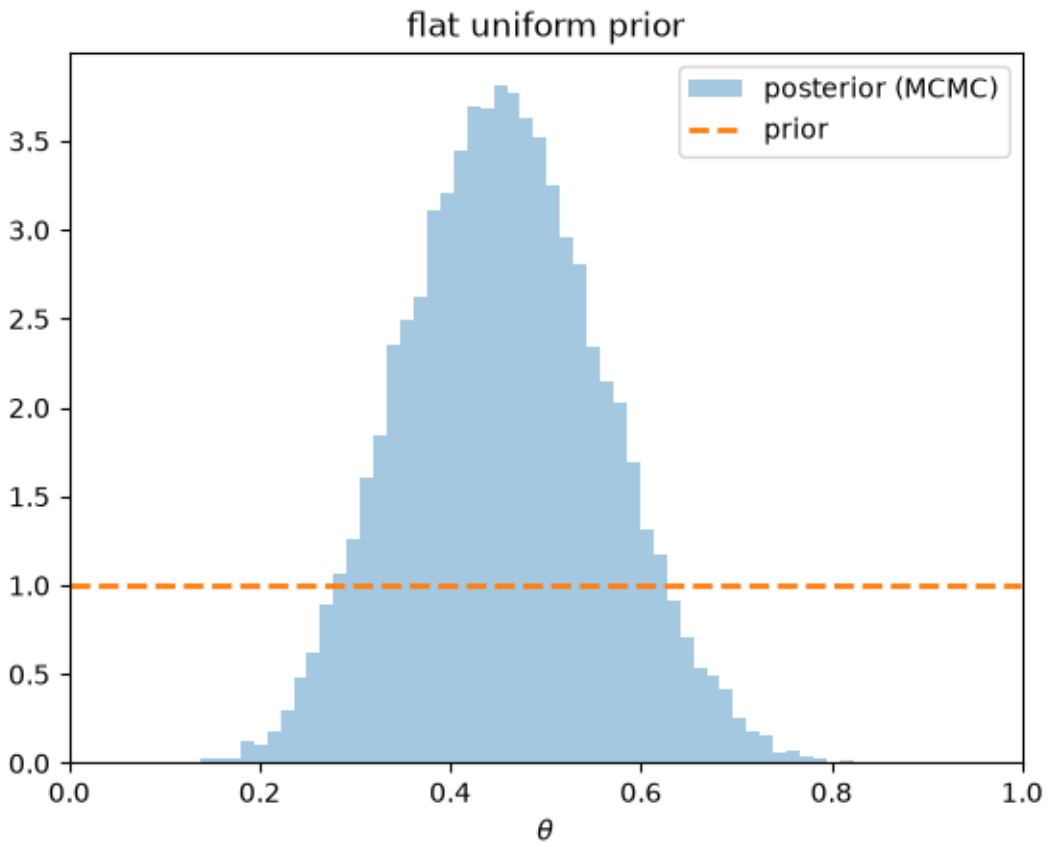
18.4.1 A uniform prior

The simplest non-conjugate prior is **uniform**: the analyst regards every value of θ in some interval as equally likely.

A uniform prior on all of $[0, 1]$ expresses indifference.

Because its density is constant, the posterior is then proportional to the likelihood alone.

```
mcmc_flat = run_nuts(binomial_model, dist.Uniform(0.0, 1.0), k, n)
plot_prior_posterior(dist.Uniform(0.0, 1.0),
                     mcmc_flat.get_samples()["θ"],
                     title="flat uniform prior")
```



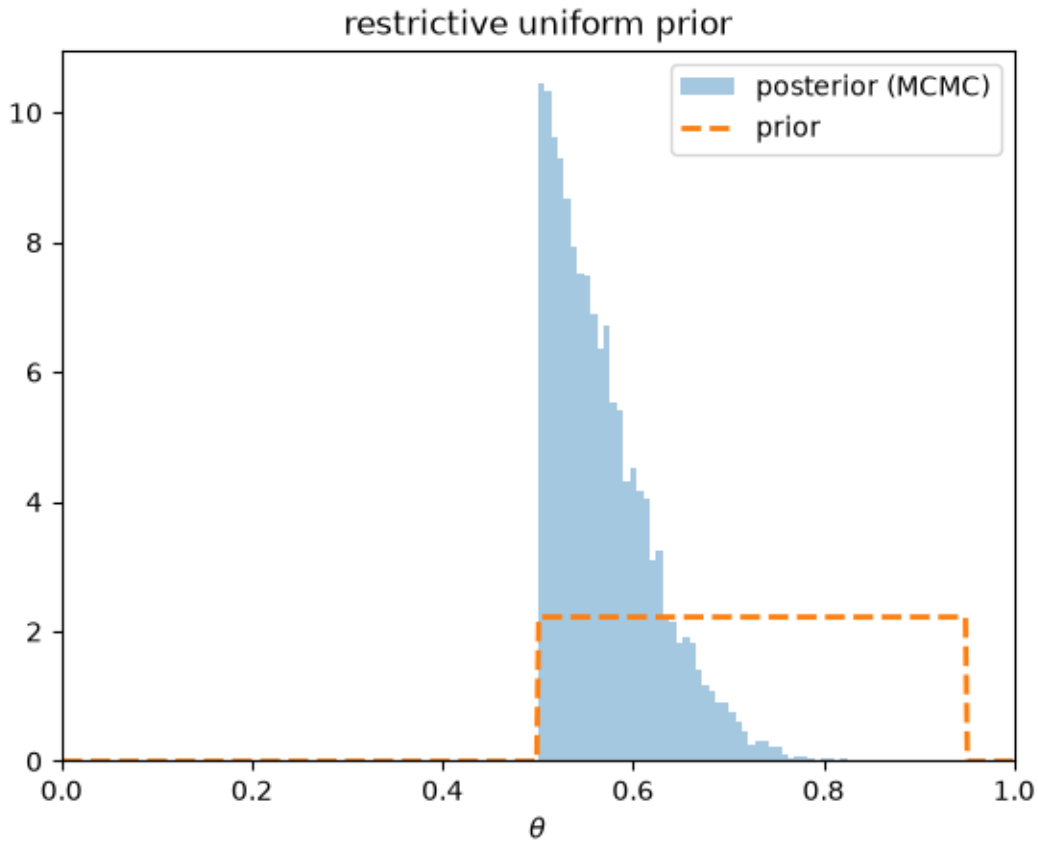
The posterior is centered near the sample frequency k/n, just as the likelihood is.

Now suppose instead that the analyst is convinced the coin favors heads, and places a uniform prior on $[0.5, 0.95]$.

This prior assigns *zero* density to the region around the true value $\theta = 0.4$.

```
mcmc_restr = run_nuts(binomial_model, dist.Uniform(0.5, 0.95), k, n)
plot_prior_posterior(dist.Uniform(0.5, 0.95),
                     mcmc_restr.get_samples()["θ"],
                     title="restrictive uniform prior")
```

```
/tmp/ipykernel_12164/411337751.py:7: UserWarning: Out-of-support values provided␣
↪to log prob method. The value argument should be within the support.
  prior_pdf = np.where(in_support, np.exp(np.asarray(prior.log_prob(grid))), 0.0)
```



The posterior cannot put mass where the prior is zero, so it piles up against the lower boundary 0.5 — as close to the data as the prior permits.

This is a vivid warning: a prior that rules out the truth can never be overturned by data, no matter how much we collect.

18.4.2 A truncated log-normal prior

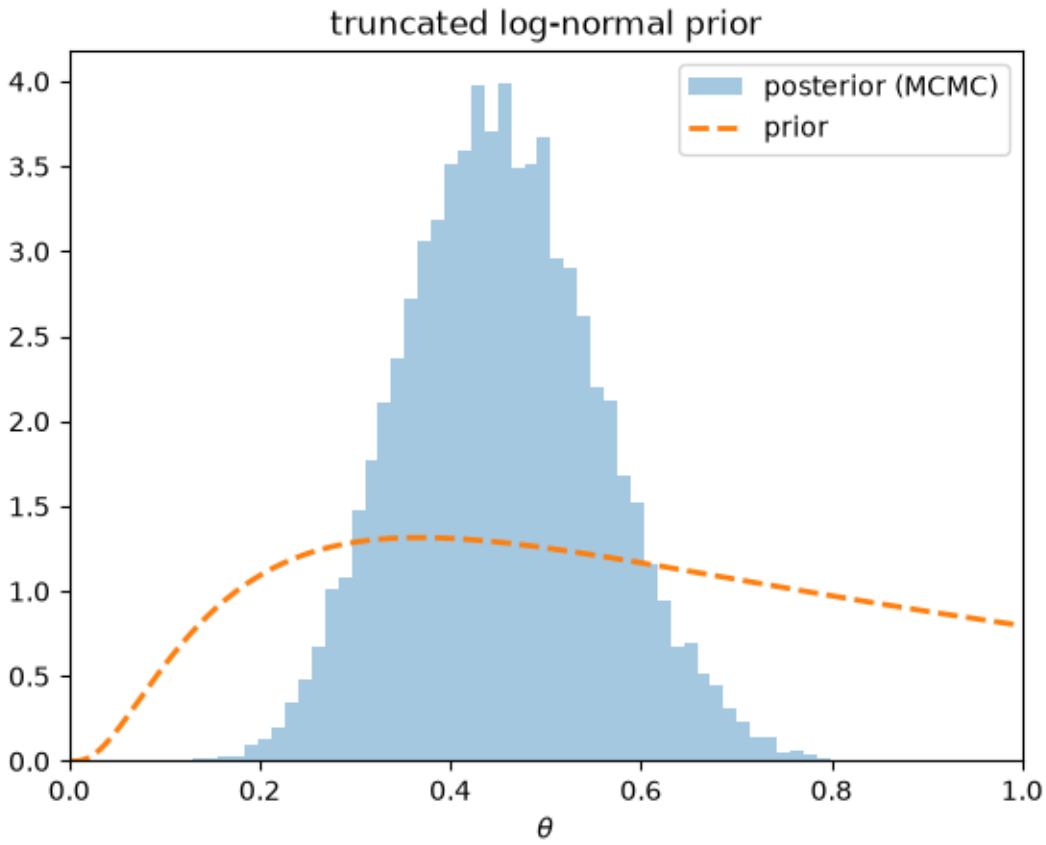
A uniform prior is flat. A more realistic prior is smooth and asymmetric.

A convenient choice on $[0, 1]$ is a **truncated log-normal**: take $Z \sim N(\mu, \sigma)$ truncated to $Z \leq 0$, and set $\theta = e^Z$, which then lies in $(0, 1]$.

NumPyro builds this by feeding a `TruncatedNormal` through an `ExpTransform`.

```
def truncated_lognormal(μ, σ):
    "Log-normal distribution truncated to the unit interval (0, 1]."
    base = dist.TruncatedNormal(loc=μ, scale=σ, low=-jnp.inf, high=0.0)
    # Declare the (0, 1] support: ExpTransform alone advertises (0, ∞),
    # which would let the sampler propose θ > 1.
    class _UnitLogNormal(dist.TransformedDistribution):
        support = dist.constraints.interval(0.0, 1.0)
    return _UnitLogNormal(base, dist.transforms.ExpTransform())

prior_ln = truncated_lognormal(0.0, 1.0)
mcmc_ln = run_nuts(binomial_model, prior_ln, k, n)
plot_prior_posterior(prior_ln, mcmc_ln.get_samples()["θ"],
                     title="truncated log-normal prior")
```



The prior favors smaller values of θ, but with $\sigma = 1$ it is diffuse, so the likelihood pulls the posterior toward the sample frequency.

We keep `mcmc_ln` — we will compare it with variational inference below.

18.4.3 A truncated Laplace prior

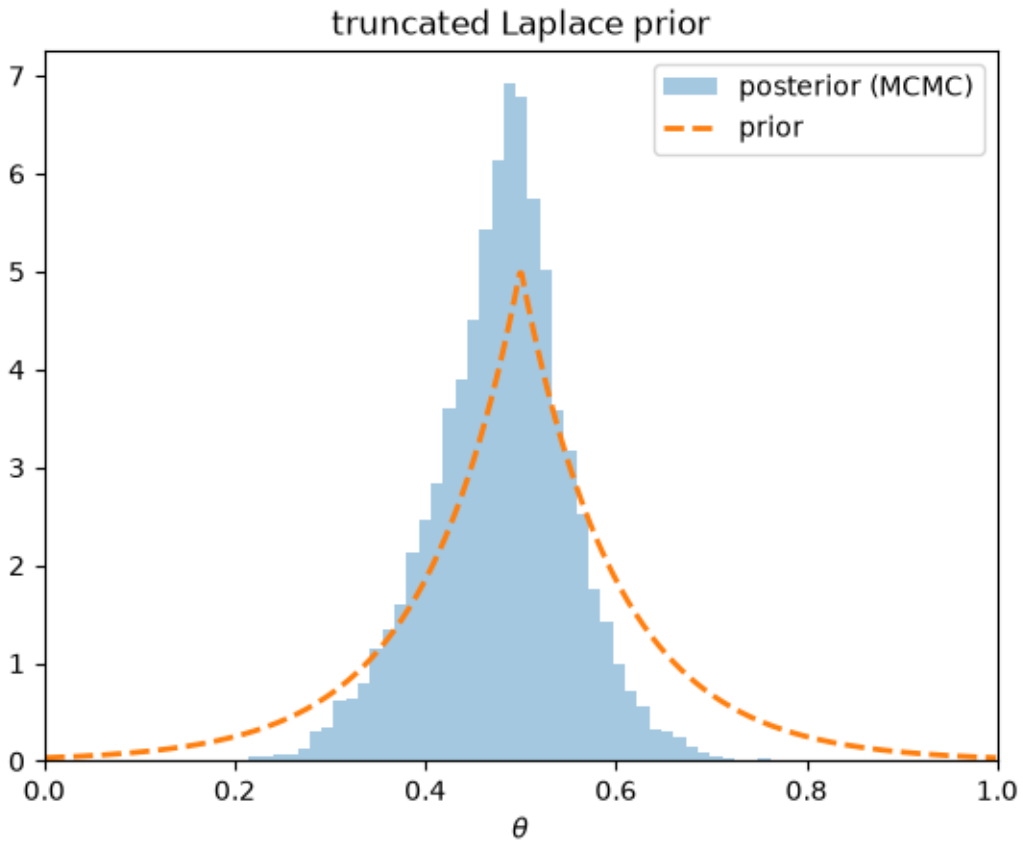
Our final prior has a sharp, non-smooth peak.

A **Laplace** density $\propto e^{-|\theta-\mu|/b}$ has a kink at its center μ, expressing a strong belief that θ sits near μ while still allowing for surprises in the tails.

We truncate it to $[0, 1]$ and center it at 0.5.

```
def truncated_laplace(μ, b):
    "Laplace distribution truncated to the unit interval [0, 1]."
    return dist.TruncatedDistribution(dist.Laplace(μ, b), low=0.0, high=1.0)

prior_lp = truncated_laplace(0.5, 0.1)
mcmc_lp = run_nuts(binomial_model, prior_lp, k, n)
plot_prior_posterior(prior_lp, mcmc_lp.get_samples()["θ"],
                     title="truncated Laplace prior")
```



The spiked prior tugs the posterior toward 0.5, away from the sample frequency near 0.4.

The pull is gentle here because the prior, though peaked, is not very tight; with a smaller b it would dominate the modest sample.

NUTS handles the kink in the prior without any special tuning — a practical advantage of gradient-based samplers paired with automatic differentiation.

18.5 Variational inference

MCMC approximates the posterior by *sampling* from it.

Variational inference (VI) takes a different route: it turns posterior approximation into an *optimization* problem.

We restrict attention to a tractable family of densities $q_\phi(\theta)$ — the **guide** — indexed by parameters ϕ, and we search for the member of that family closest to the posterior.

18.5.1 Why variational inference?

If NUTS already returns accurate posteriors, why introduce another method?

The answer is **scale**.

MCMC evaluates the likelihood over the entire dataset at every step, and the number of steps it needs tends to grow with the dimension of the parameter.

For large datasets or high-dimensional models — for instance the hierarchical models and neural networks common in machine learning — this can become too slow to be practical.

Variational inference scales much better, because the objective (the ELBO, introduced below) can be maximized with *stochastic* gradients computed on small random subsets of the data — the same machinery that trains deep learning models.

It also yields a compact parametric approximation that is cheap to store and to draw from afterwards.

The price is accuracy: VI returns only the best fit *within the guide family*, and it can understate uncertainty.

As a rule of thumb, prefer MCMC when you need an accurate posterior and the problem is small enough to afford it, and VI when the model is too large for MCMC or a fast, approximate answer is good enough.

18.5.2 The evidence lower bound

Let the prior be $p(\theta)$ and the likelihood be $p(Y \mid \theta)$, where Y denotes the observed data (here the head count k).

By Bayes' rule,

$$p(\theta \mid Y) = \frac{p(Y, \theta)}{p(Y)} = \frac{p(Y \mid \theta)\, p(\theta)}{p(Y)},$$

where

$$p(Y) = \int p(Y \mid \theta)\, p(\theta)\, d\theta. \tag{18.1}$$

The integral in (18.1) is the troublesome one: in the non-conjugate case it has no closed form.

We measure the discrepancy between the guide $q_\phi(\theta)$ and the posterior with the **Kullback–Leibler (KL) divergence**

$$D_{KL}(q_\phi(\theta) \,\|\, p(\theta \mid Y)) = -\int q_\phi(\theta)\, \log \frac{p(\theta \mid Y)}{q_\phi(\theta)}\, d\theta,$$

and we choose ϕ to minimize it.

The KL divergence still involves the intractable posterior, but we can rearrange it. Using $p(\theta \mid Y) = p(\theta, Y)/p(Y)$,

$$\begin{aligned} D_{KL}(q_\phi \,\|\, p(\theta \mid Y)) &= -\int q_\phi(\theta)\, \log \frac{p(\theta, Y)/p(Y)}{q_\phi(\theta)}\, d\theta \\ &= -\int q_\phi(\theta) \left[\log \frac{p(\theta, Y)}{q_\phi(\theta)} - \log p(Y)\right] d\theta \\ &= -\int q_\phi(\theta)\, \log \frac{p(\theta, Y)}{q_\phi(\theta)}\, d\theta + \log p(Y), \end{aligned}$$

where the last line uses $\int q_\phi(\theta)\, d\theta = 1$. Rearranging,

$$\log p(Y) = D_{KL}(q_\phi \,\|\, p(\theta \mid Y)) + \underbrace{\int q_\phi(\theta)\, \log \frac{p(\theta, Y)}{q_\phi(\theta)}\, d\theta}_{\text{ELBO}}.$$

The marginal likelihood $\log p(Y)$ on the left does not depend on ϕ.

Hence **minimizing** the KL divergence is equivalent to **maximizing** the second term, the **evidence lower bound (ELBO)**:

$$\text{ELBO}(\phi) \equiv \int q_\phi(\theta)\, \log \frac{p(\theta, Y)}{q_\phi(\theta)}\, d\theta = \mathbb{E}_{q_\phi(\theta)}\big[\log p(\theta, Y) - \log q_\phi(\theta)\big]. \tag{18.2}$$

Because $D_{KL} \geq 0$, the ELBO is a lower bound on $\log p(Y)$ — hence its name.

Crucially, (18.2) involves only the *joint* density $p(\theta, Y) = p(Y \mid \theta)\, p(\theta)$, which we can evaluate, not the intractable normalizing constant $p(Y)$.

The expectation can be estimated by sampling from q_ϕ, and ϕ improved by gradient ascent — this is **stochastic variational inference (SVI)**.

18.5.3 Implementing SVI in NumPyro

We need a guide q_ϕ.

The simplest choice is an **autoguide**: NumPyro inspects the model and automatically constructs a guide for us.

`AutoNormal` places an independent normal distribution on each latent variable, transformed to respect its support — here, to keep θ inside $(0, 1)$.

We apply SVI to the truncated log-normal model from above and maximize the ELBO with the Adam optimizer.

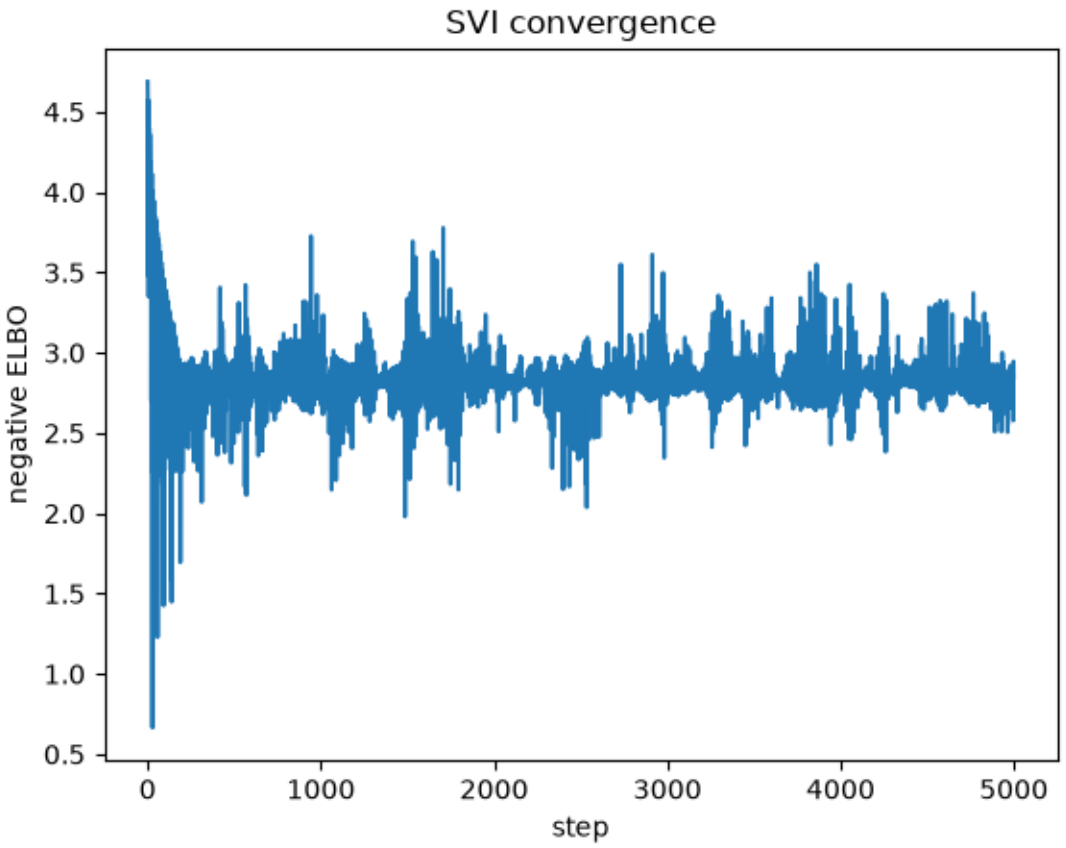
```
guide = AutoNormal(binomial_model)
optimizer = Adam(step_size=0.01)
svi = SVI(binomial_model, guide, optimizer, loss=Trace_ELBO())

svi_result = svi.run(random.key(0), 5000, prior_ln, k, n, progress_bar=False)
```

SVI maximizes the ELBO; equivalently, it minimizes its negative, which is the reported loss.

A loss curve that flattens out indicates convergence.

```
fig, ax = plt.subplots()
ax.plot(svi_result.losses)
ax.set_xlabel("step")
ax.set_ylabel("negative ELBO")
ax.set_title("SVI convergence")
plt.show()
```

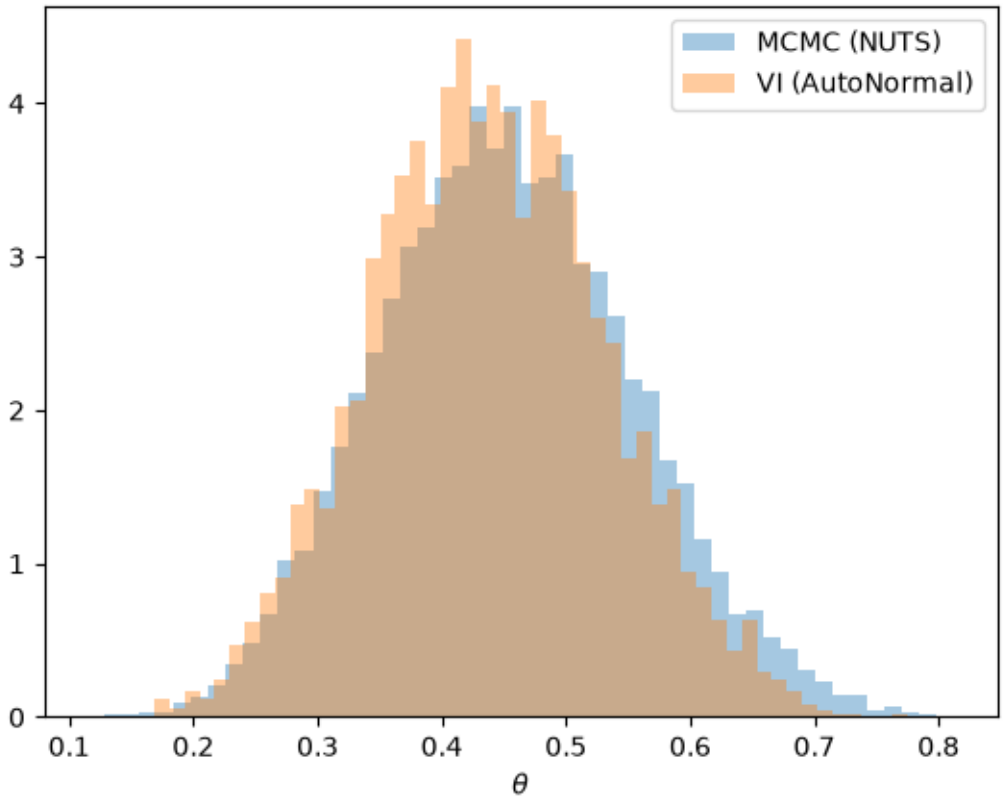


18.5.4 Comparing VI with MCMC

To assess the approximation, we draw samples from the fitted guide and compare them with the NUTS posterior for the same (log-normal-prior) model.

```
vi_samples = guide.sample_posterior(
    random.key(1), svi_result.params, sample_shape=(4000,)
)["θ"]
nuts_samples = mcmc_ln.get_samples()["θ"]

fig, ax = plt.subplots()
ax.hist(np.asarray(nuts_samples), bins=50, density=True, alpha=0.4,
        label="MCMC (NUTS)")
ax.hist(np.asarray(vi_samples), bins=50, density=True, alpha=0.4,
        label="VI (AutoNormal)")
ax.set_xlabel(r"$\theta$")
ax.legend()
plt.show()
```



The two approximations broadly agree on the location and spread of the posterior.

They need not agree perfectly.

MCMC samples the true posterior (up to Monte Carlo error), whereas VI reports the best fit *within its guide family*.

A mean-field normal guide is symmetric on the transformed scale and can miss skewness or heavy tails in the true posterior.

The trade-off is one of cost against fidelity: VI replaces sampling with optimization and is often much faster in high dimensions, but it delivers an approximation whose quality is capped by the flexibility of the guide.

18.6 Where to next

This lecture showed how to compute posteriors when prior and likelihood are not conjugate, using NUTS and stochastic variational inference in NumPyro.

The same tools carry over to richer models.

The lectures *Posterior Distributions for AR(1) Parameters* and *Forecasting an AR(1) Process* apply NumPyro to Bayesian estimation and forecasting of autoregressive time series, where the parameter is a vector and conjugate analysis is unavailable.

CHAPTER

NINETEEN

POSTERIOR DISTRIBUTIONS FOR AR(1) PARAMETERS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's included in base Anaconda, this lecture requires the following libraries.

We first install `numpyro` and `jax`:

```
!pip install numpyro jax
```

We also install `arviz` and `pymc`:

```
!pip install arviz pymc
```

19.1 Overview

This lecture uses Bayesian methods offered by pymc and numpyro to make statistical inferences about two parameters of a univariate first-order autoregression.

The model is a good laboratory for illustrating consequences of alternative ways of modeling the distribution of the initial y_0:

- As a fixed number
- As a random variable drawn from the stationary distribution of the $\{y_t\}$ stochastic process

19.1.1 Setting

The first component of the statistical model is

$$y_{t+1} = \rho y_t + \sigma_x \epsilon_{t+1}, \quad t \geq 0 \tag{19.1}$$

where

- the scalars ρ and σ_x satisfy $|\rho| < 1$ and $\sigma_x > 0$
- $\{\epsilon_{t+1}\}$ is a sequence of IID normal random variables with mean 0 and variance 1.

The second component of the statistical model is

$$y_0 \sim N(\mu_0, \sigma_0^2) \tag{19.2}$$

Unknown parameters are ρ, σ_x.

We have independent *prior probability distributions* for ρ, σ_x and want to compute a posterior probability distribution after observing a sample $\{y_t\}_{t=0}^T$.

We want to study how inferences about the unknown parameters (ρ, σ_x) depend on what is assumed about the parameters μ_0, σ_0 of the distribution of y_0.

We study two assumptions about the initial value y_0, and we refer to them by name throughout the lecture.

Under the **conditioning assumption** we take the observed y_0 as given.

Formally we set $(\mu_0, \sigma_0) = (y_0, 0)$, so the density of y_0 is a spike at its observed value.

This density does not depend on ρ or σ_x, so y_0 carries no information about the parameters; in effect we **condition on** y_0 and model only what happens after it.

Under the **stationary assumption** we treat y_0 as a draw from the stationary distribution of the process,

$$y_0 \sim N\left(0, \frac{\sigma_x^2}{1-\rho^2}\right).$$

This density *does* depend on ρ and σ_x, so now the observed y_0 carries information about the parameters.

The whole lecture is about how this one difference affects our estimates.

> **Note**
>
> We do not treat a third possible case in which μ_0, σ_0 are free parameters to be estimated.

19.1.2 Libraries

We use PyMC and NumPyro to compute the posterior distribution of ρ, σ_x.

We use two libraries because they make different trade-offs.

PyMC offers a mature and highly readable modeling syntax together with a rich set of diagnostic tools, which makes it convenient for learning and prototyping.

NumPyro is built on JAX, so it compiles to fast machine code and can run on a GPU, which helps it scale to larger models and datasets.

Because both libraries fit the same model, running them side by side also lets us check that they agree.

Both libraries support the NUTS sampler, which we use to draw samples from the posterior.

We treat NUTS as a black box here; see the *introduction to NUTS* in *Non-Conjugate Priors* for a brief account of how it works.

19.1.3 Imports

Let's start with some Python imports.

```
import arviz as az
import pymc as pm
import numpyro
from numpyro import distributions as dist

import numpy as np
import jax.numpy as jnp
from jax import random
import matplotlib.pyplot as plt

import logging
logging.basicConfig()
logger = logging.getLogger('pymc')
logger.setLevel(logging.CRITICAL)
```

19.2 Estimation

Let's turn to estimation, starting with the likelihood function.

19.2.1 Likelihood function

For a sample $\{y_t\}_{t=0}^T$ from the AR(1) model, the likelihood function can be *factored* as follows:

$$f(y_T, y_{T-1}, \dots, y_0) = f(y_T|y_{T-1}) f(y_{T-1}|y_{T-2}) \cdots f(y_1|y_0) f(y_0)$$

(We use f to denote a generic probability density.)

The statistical model (19.1)-(19.2) implies

$$\begin{aligned} f(y_t|y_{t-1}) &\sim N(\rho y_{t-1}, \sigma_x^2) \\ f(y_0) &\sim N(\mu_0, \sigma_0^2) \end{aligned}$$

We shall use Bayes' law to construct a posterior distribution under the alternative assumptions.

As discussed above, the way that we select the initial value y_0 matters.

- If we believe y_0 really is a draw from the stationary distribution, the stationary assumption is a good choice, because then y_0 carries useful information about ρ and σ_x.
- If we suspect y_0 is far out in the tail — so that early observations carry a large **transient component** — the conditioning assumption is better.

To illustrate the issue, we'll begin by choosing an initial y_0 that is far out in a tail of the stationary distribution.

19.2.2 Simulation code

We will work with simulated data, fixing parameters ρ and σ_x.

Then we will pretend that we don't know these parameters and try to estimate them under the two assumptions discussed above (conditioning and stationary).

The following function simulates a path of the AR(1) process from a given initial condition.

```
def ar1_simulate(ρ, σ, y0, T, rng):

    # Allocate space and draw shocks
    y = np.empty(T)
    ε = rng.normal(0, σ, T)

    # Initial condition and step forward
    y[0] = y0
    for t in range(1, T):
        y[t] = ρ * y[t-1] + ε[t]

    return y
```

We use the following parameterization.

```
σ_true = 1.0   # Fix a value for σ_x
ρ_true = 0.5   # Fix a value for ρ
```

Our simulated time series will be relatively short, so that the priors matter:

```
T = 50   # Length of time series
```

As mentioned above, we choose an initial y_0 that is far out in a tail of the stationary distribution:

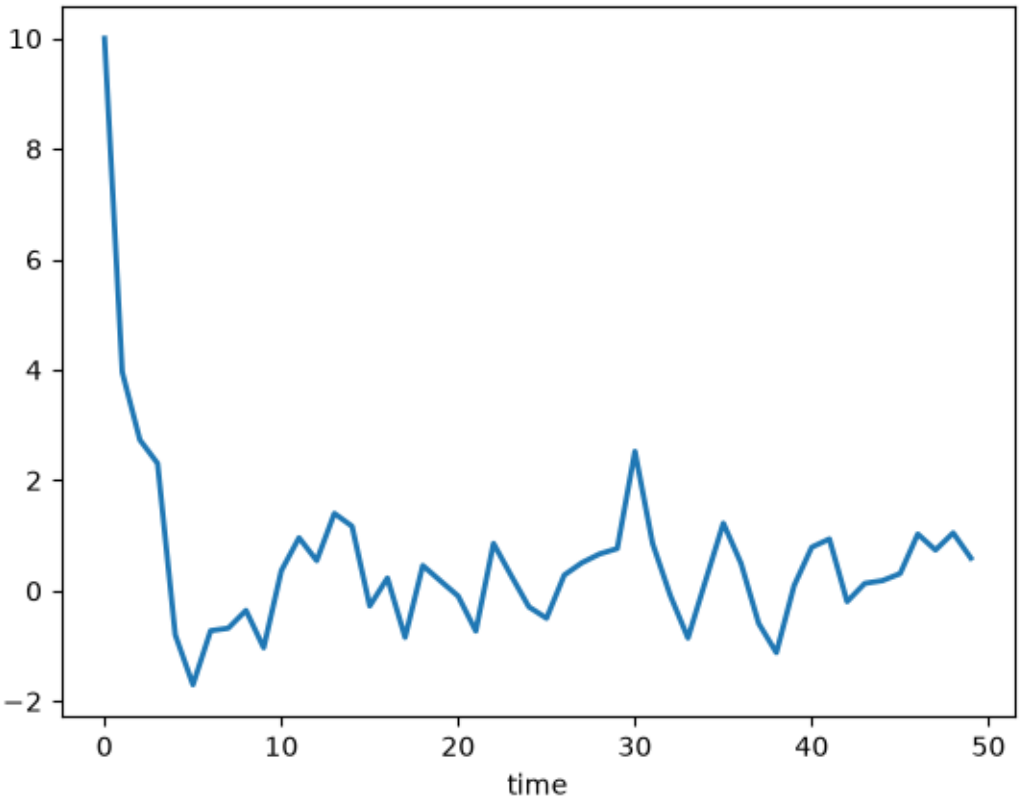
```
y_0 = 10
```

Now let's simulate and generate our data:

```
rng = np.random.default_rng(42)
y = ar1_simulate(ρ_true, σ_true, y_0, T, rng)
```

Here's a plot of the simulated series.

```
fig, ax = plt.subplots()
ax.plot(y, lw=2)
ax.set_xlabel('time')
plt.show()
```



You can see that the initial condition is unusually large — the series moves away from it quickly and fluctuates in a lower band.

19.2.3 PyMC implementation

In this section we use PyMC to compute the posterior under each assumption about y_0 — first the conditioning assumption, then the stationary assumption.

We parameterize each normal distribution in PyMC by its standard deviation σ, via the `sigma` argument.

```
AR1_model = pm.Model()

with AR1_model:

    # Start with priors
    ρ = pm.Uniform('rho', lower=-1., upper=1.)  # Assume stable ρ
    σ = pm.HalfNormal('sigma', sigma=np.sqrt(10))

    # Expected value of y at the next period (ρ * y)
    yhat = ρ * y[:-1]

    # Likelihood of the actual realization
    y_like = pm.Normal('y_obs', mu=yhat, sigma=σ, observed=y[1:])
```

Let's unpack what this model declares.

Inside the `with AR1_model:` block, each `pm` statement adds a random variable to the model.

The first two lines are the **priors** — a uniform prior on ρ over $(-1, 1)$, which builds in stationarity, and a half-normal prior on σ.

The line `yhat = ρ * y[:-1]` is just the vector of conditional means ρy_{t-1} for $t = 1, \ldots, T$.

The last line is the **likelihood**: the keyword `observed=y[1:]` tells PyMC that the values `y[1:]` are data, drawn from $N(\rho y_{t-1}, \sigma^2)$.

Because `yhat` and `y[1:]` are whole vectors, this single line encodes the entire product $\prod_{t=1}^{T} f(y_t \mid y_{t-1})$ from the factorization above.

PyMC multiplies this likelihood by the priors to form the posterior, which we sample from below.

Note

Notice what is *absent*: we never write down a density for y_0 itself.

It enters the model only inside `y[:-1]`, as the conditioning value for the first transition $f(y_1 \mid y_0)$ — never as something the model has to explain.

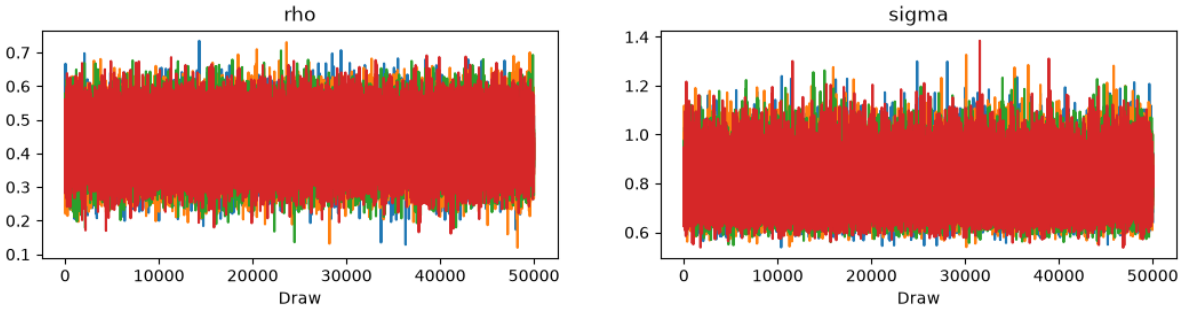
Leaving out the $f(y_0)$ factor in this way is precisely the **conditioning assumption**.

pm.sample by default uses the NUTS sampler to generate samples, as shown in the cell below:

```
with AR1_model:
    trace = pm.sample(50000, tune=10000, return_inferencedata=True)
```

We plot the trace and the posterior densities for the two parameters.

```
with AR1_model:
    az.plot_trace(trace)
```



Recall that we generated the data with $\rho = 0.5$ and $\sigma_x = 1$.

The posteriors concentrate near these values, so conditioning on y_0 recovers the parameters reasonably well, even from this short sample.

Note

The fit is good but not exact — the posterior for ρ sits a little below its true value.

This is partly the classic **Hurwicz bias**: in a first-order autoregression the least-squares (and hence posterior) estimate of ρ is biased downward in small samples, with the bias shrinking as the sample grows (see Hurwicz [1950] and [Orcutt and Winokur, 1969]).

Here is a numerical summary of the posterior.

```
with AR1_model:
    summary = az.summary(trace, round_to=4)

summary
```

```
          mean      sd  eti89_lb  eti89_ub     ess_bulk     ess_tail  r_hat  \
rho     0.4371  0.0634    0.3364    0.5375  181856.2850  133110.0006    1.0
sigma   0.7927  0.0838    0.6716    0.9366  174140.4497  131015.8252    1.0

        mcse_mean  mcse_sd
rho        0.0001   0.0002
sigma      0.0002   0.0002
```

Now we switch to the stationary assumption, using the same data.

Recall that this means

$$y_0 \sim N\left(0, \frac{\sigma_x^2}{1-\rho^2}\right).$$

We alter the code as follows:

```
AR1_model_y0 = pm.Model()

with AR1_model_y0:

    # Start with priors
    ρ = pm.Uniform('rho', lower=-1., upper=1.)  # Assume stable ρ
    σ = pm.HalfNormal('sigma', sigma=np.sqrt(10))

    # Standard deviation of the stationary distribution
    y_sd = σ / np.sqrt(1 - ρ**2)

    # Expected value of y at the next period (ρ * y)
    yhat = ρ * y[:-1]
    y_data = pm.Normal('y_obs', mu=yhat, sigma=σ, observed=y[1:])

    # Density for y0 -- this term imposes the stationary assumption
    y0_data = pm.Normal('y0_obs', mu=0., sigma=y_sd, observed=y[0])
```

The only change from the first model is that final line.

Note

The new line adds a density for y_0 — the stationary density $N(0, \sigma_x^2/(1-\rho^2))$ — through a second `observed` term.

This restores the $f(y_0)$ factor we dropped before, and so *is* the **stationary assumption**.

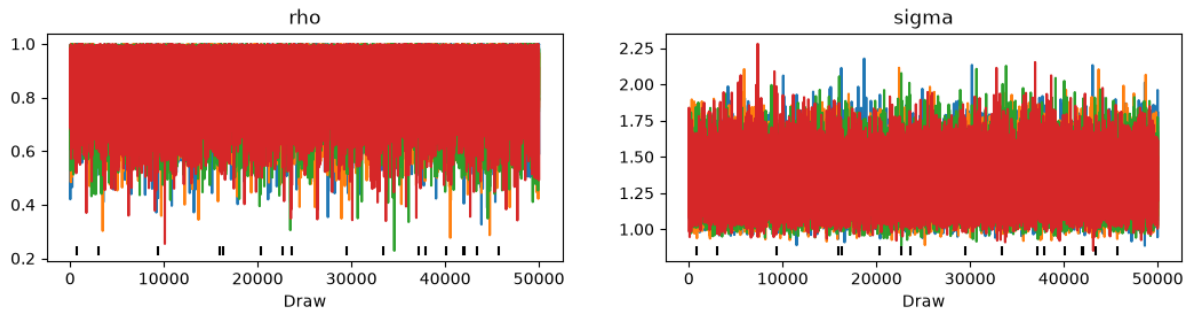
Everything else is identical, so any difference between the two posteriors comes entirely from this one term.

As before, we sample from the posterior.

```
with AR1_model_y0:
    trace_y0 = pm.sample(50000, tune=10000, return_inferencedata=True)
```

In the trace plot below, any grey vertical lines mark sampler divergences.

```
with AR1_model_y0:
    az.plot_trace(trace_y0)
```



Here is a summary of the posterior.

```
with AR1_model_y0:
    summary_y0 = az.summary(trace_y0, round_to=4)

summary_y0
```

```
          mean      sd  eti89_lb  eti89_ub     ess_bulk     ess_tail  r_hat  \
rho     0.8794  0.0853    0.7207    0.9820  112648.2773   90432.7315    1.0
sigma   1.3165  0.1384    1.1154    1.5531  134350.7091  112158.8294    1.0

        mcse_mean  mcse_sd
rho        0.0002   0.0002
sigma      0.0004   0.0004
```

The posterior for ρ has clearly moved when we changed our assumption about y_0.

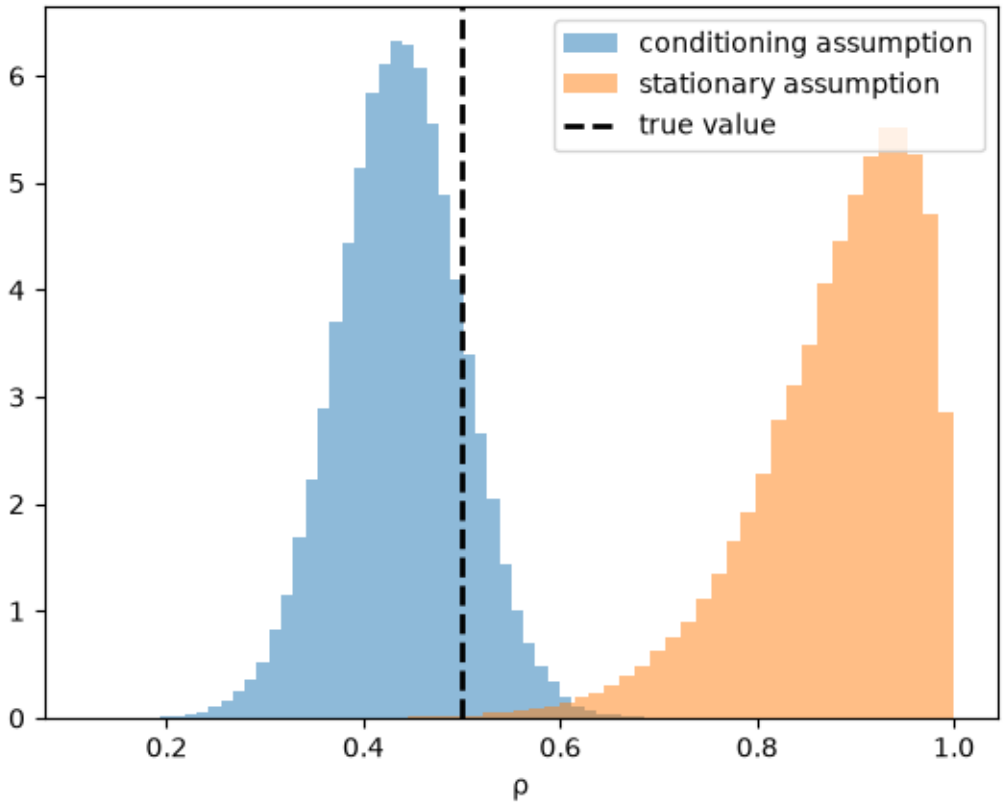
19.2.4 Comparing the two posteriors

Let's put the two posteriors for ρ side by side to see what changed.

The figure below overlays the posterior for ρ under each assumption, with the true value marked by a dashed line.

```
ρ_cond = trace.posterior['rho'].values.flatten()
ρ_stat = trace_y0.posterior['rho'].values.flatten()

fig, ax = plt.subplots()
ax.hist(ρ_cond, bins=50, density=True, alpha=0.5,
        label='conditioning assumption')
ax.hist(ρ_stat, bins=50, density=True, alpha=0.5,
        label='stationary assumption')
ax.axvline(ρ_true, color='k', linestyle='--', lw=2, label='true value')
ax.set_xlabel('ρ')
ax.legend()
plt.show()
```



The posterior from the conditioning assumption sits close to the true value of 0.5.

It is centred a little below 0.5, which is the small-sample Hurwicz bias mentioned above — a mild downward pull that shrinks as the sample grows.

The posterior from the stationary assumption is different: it is pushed well to the right, toward $\rho = 1$.

Here is why.

We chose a starting value $y_0 = 10$ that lies far out in the tail of the stationary distribution.

If y_0 really were a draw from that distribution, such an extreme value would be very unlikely.

So when we force the model to explain y_0 this way, Bayes' law looks for parameter values that make the extreme y_0 plausible.

It does this by pushing ρ toward 1, because a larger ρ inflates the stationary variance $\sigma_x^2/(1-\rho^2)$ and makes a large y_0 less surprising.

A single unusual starting value therefore drags the whole posterior away from the truth.

This is the sense in which the conditioning assumption is more accurate here: it does not let one atypical observation distort our view of ρ and σ_x.

19.2.5 NumPyro implementation

We now redo both computations with NumPyro.

Because it fits the same two models, we expect its posteriors to match those from PyMC.

The models are the same; only the syntax differs.

NumPyro describes a model as an ordinary Python function rather than a `with` block, each `numpyro.sample('name', distribution)` plays the role that a `pm` random variable did above, and the keyword `obs=` is NumPyro's counterpart of PyMC's `observed=`.

Everything we said about priors, the vectorized likelihood, and how the two assumptions are imposed carries over unchanged; see *Non-Conjugate Priors* for a fuller introduction to NumPyro models.

We start with a helper function that plots the trace and posterior histogram for the sampled parameters.

```
def plot_posterior(sample):
    """
    Plot trace and histogram
    """
    # To np array
    ρs = np.array(sample['rho'])
    σs = np.array(sample['sigma'])

    fig, axs = plt.subplots(2, 2, figsize=(17, 6))
    # Plot trace
    axs[0, 0].plot(ρs, lw=2)
    axs[1, 0].plot(σs, lw=2)

    # Plot posterior
    axs[0, 1].hist(ρs, bins=50, density=True, alpha=0.7)
    axs[0, 1].set_xlim([0, 1])
    axs[1, 1].hist(σs, bins=50, density=True, alpha=0.7)

    axs[0, 0].set_ylabel('ρ')
    axs[1, 0].set_ylabel('σ')
    axs[0, 1].set_xlabel('ρ')
    axs[1, 1].set_xlabel('σ')
    plt.show()
```

The first model uses the conditioning assumption.

```
def AR1_model(data):
    # Set prior
    ρ = numpyro.sample('rho', dist.Uniform(low=-1., high=1.))
    σ = numpyro.sample('sigma', dist.HalfNormal(scale=np.sqrt(10)))

    # Expected value of y at the next period (ρ * y)
    yhat = ρ * data[:-1]

    # Likelihood of the actual realization
    y_data = numpyro.sample('y_obs', dist.Normal(loc=yhat, scale=σ), obs=data[1:])
```

We convert the data to a JAX array, build the NUTS sampler, and run MCMC.

```
# Make jnp array
y = jnp.array(y)
```

(continues on next page)

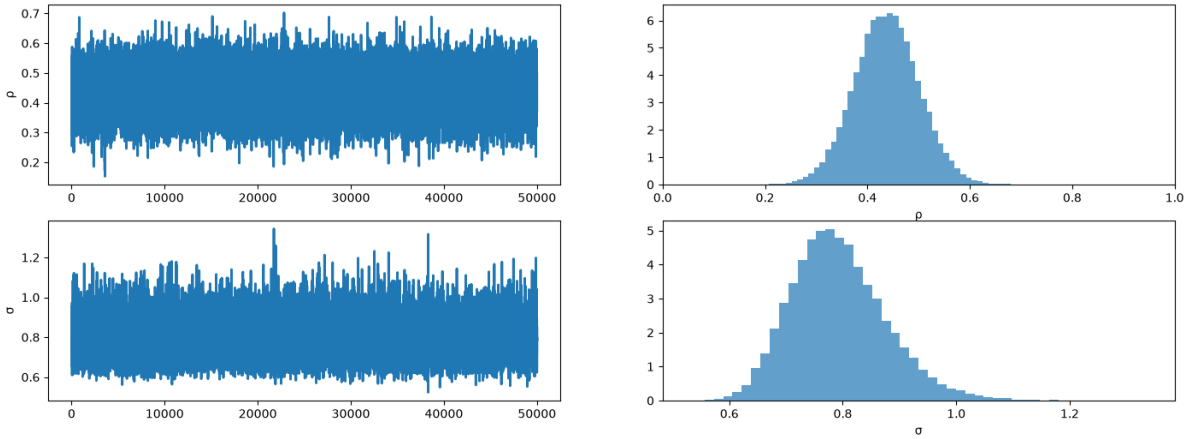
(continued from previous page)

```
# Set NUTS kernel
NUTS_kernel = numpyro.infer.NUTS(AR1_model)

# Run MCMC
mcmc = numpyro.infer.MCMC(NUTS_kernel, num_samples=50000, num_warmup=10000, progress_
↪bar=False)
mcmc.run(rng_key=random.key(1), data=y)
```

We plot the trace and posterior.

```
plot_posterior(mcmc.get_samples())
```



Here is the posterior summary.

```
mcmc.print_summary()
```

```
                mean       std    median      5.0%      95.0%      n_eff      r_hat
       rho      0.44      0.06      0.44      0.33       0.54   32818.11       1.00
     sigma      0.79      0.08      0.79      0.66       0.93   31929.88       1.00

Number of divergences: 0
```

Next we use the stationary assumption, where

$$y_0 \sim N\left(0, \frac{\sigma_x^2}{1-\rho^2}\right).$$

Here's the new code to achieve this.

```
def AR1_model_y0(data):
    # Set prior
    ρ = numpyro.sample('rho', dist.Uniform(low=-1., high=1.))
    σ = numpyro.sample('sigma', dist.HalfNormal(scale=np.sqrt(10)))

    # Standard deviation of the stationary distribution
    y_sd = σ / jnp.sqrt(1 - ρ**2)

    # Expected value of y at the next period (ρ * y)
    yhat = ρ * data[:-1]
```

(continues on next page)

(continued from previous page)

```
    # Likelihood of the actual realization
    y_data = numpyro.sample('y_obs', dist.Normal(loc=yhat, scale=σ), obs=data[1:])

    # Density for y0 -- this term imposes the stationary assumption
    y0_data = numpyro.sample('y0_obs', dist.Normal(loc=0., scale=y_sd), obs=data[0])
```

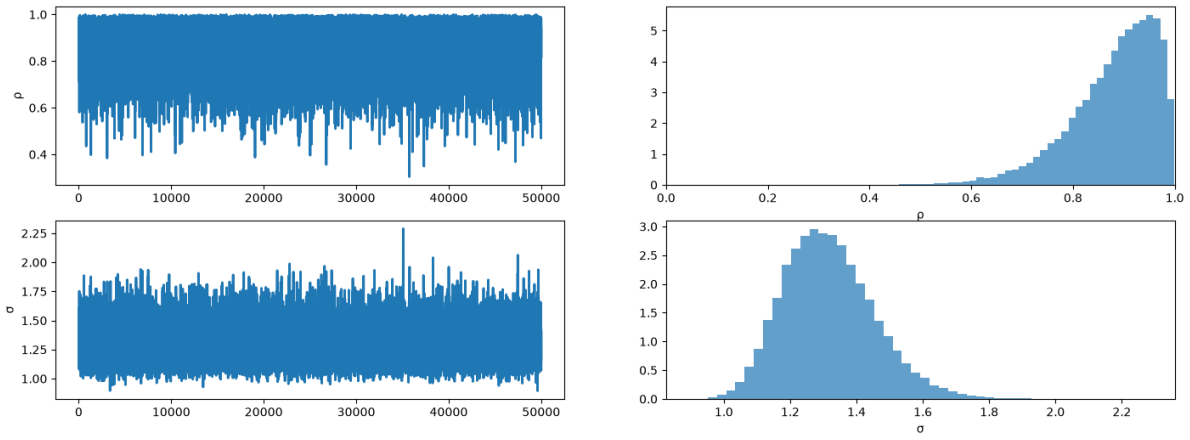
We build the sampler for this model and run MCMC.

```
# Set NUTS kernel
NUTS_kernel = numpyro.infer.NUTS(AR1_model_y0)

# Run MCMC
mcmc2 = numpyro.infer.MCMC(NUTS_kernel, num_samples=50000, num_warmup=10000, progress_
↪bar=False)
mcmc2.run(rng_key=random.key(1), data=y)
```

Again we plot the trace and posterior.

```
plot_posterior(mcmc2.get_samples())
```



And here is the posterior summary.

```
mcmc2.print_summary()
```

```
                mean       std    median      5.0%      95.0%     n_eff     r_hat
       rho       0.88      0.09      0.90      0.76       1.00  32252.54      1.00
     sigma       1.32      0.14      1.31      1.09       1.53  24007.40      1.00

Number of divergences: 0
```

As with PyMC, the posterior for ρ shifts toward 1 once we switch to the stationary assumption.

To confirm that the two libraries agree, we overlay their posteriors for ρ under the conditioning assumption.

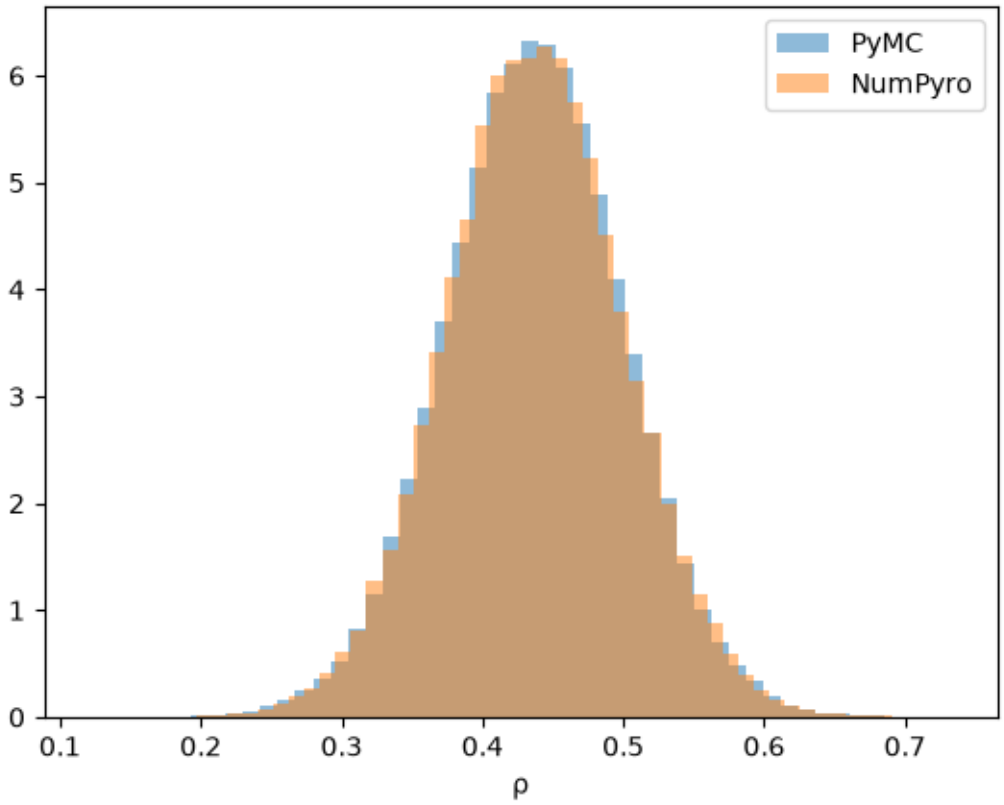
```
ρ_pymc = trace.posterior['rho'].values.flatten()
ρ_numpyro = np.array(mcmc.get_samples()['rho'])

fig, ax = plt.subplots()
ax.hist(ρ_pymc, bins=50, density=True, alpha=0.5, label='PyMC')
ax.hist(ρ_numpyro, bins=50, density=True, alpha=0.5, label='NumPyro')
ax.set_xlabel('ρ')
```

(continues on next page)

(continued from previous page)

```
ax.legend()
plt.show()
```



The two posteriors line up, as they should.

19.3 Conclusion

This lecture showed that what we assume about the initial value y_0 can have a large effect on our estimates of an AR(1) process.

When the sample is short and y_0 might be unusual, the conditioning assumption is the safer choice.

It lets the data speak about ρ and σ_x without forcing the model to explain the starting value.

The stationary assumption adds information, and that information is valuable when the assumption is true.

But when y_0 is in fact far from typical, the same assumption misleads us — here it pushed ρ toward 1 and away from the truth.

A simple rule of thumb:

- use the conditioning assumption when early observations look transient or the starting point may be atypical;
- use the stationary assumption when you are confident the process has been running near its long-run behaviour.

See also

Forecasting an AR(1) Process builds on the posterior computed here to forecast nonlinear sample-path statistics of an AR(1) process, such as the time until the next turning point.

CHAPTER

TWENTY

FORECASTING AN AR(1) PROCESS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

```
!pip install numpyro jax arviz
```

This lecture describes methods for forecasting statistics that are functions of future values of a univariate autoregressive process.

The methods are designed to take into account two possible sources of uncertainty about these statistics:

- random shocks that impinge on the transition law
- uncertainty about the parameter values of the AR(1) process

We consider two sorts of statistics:

- prospective values y_{t+j} of a random process $\{y_t\}$ that is governed by the AR(1) process
- sample path properties that are defined as non-linear functions of future values $\{y_{t+j}\}_{j\geq 1}$ at time t

Sample path properties are things like "time to next turning point" or "time to next recession".

To investigate sample path properties we'll use a simulation procedure recommended by Wecker [Wecker, 1979].

To acknowledge uncertainty about parameters, we'll deploy `numpyro` to construct a Bayesian joint posterior distribution for unknown parameters.

This lecture builds on *Posterior Distributions for AR(1) Parameters*, which studies Bayesian inference for the parameters of exactly this AR(1) model in detail.

We recommend reading that lecture first, since here we treat the construction of the posterior more briefly.

Let's start with some imports.

```
import matplotlib.pyplot as plt
import seaborn as sns
from typing import NamedTuple

# numpyro
```

(continues on next page)

(continued from previous page)

```
import numpyro
import numpyro.distributions as dist
from numpyro.infer import MCMC, NUTS

# jax
import jax
import jax.random as random
import jax.numpy as jnp
from jax import lax

# arviz
import arviz as az

sns.set_style('white')
key = random.PRNGKey(0)
```

20.1 A Univariate First-Order Autoregressive Process

Consider the univariate AR(1) model, which is the same model studied in *Posterior Distributions for AR(1) Parameters*:

$$y_{t+1} = \rho y_t + \sigma \epsilon_{t+1}, \quad t \geq 0 \tag{20.1}$$

where

- the scalars ρ and σ satisfy $|\rho| < 1$ and $\sigma > 0$
- $\{\epsilon_{t+1}\}$ is a sequence of IID normal random variables with mean 0 and variance 1.

The initial condition y_0 is a known number.

Equation (20.1) implies that for $t \geq 0$, the conditional density of y_{t+1} is

$$f(y_{t+1}|y_t; \rho, \sigma) \sim \mathcal{N}(\rho y_t, \sigma^2) \tag{20.2}$$

Further, equation (20.1) also implies that for $t \geq 0$, the conditional density of y_{t+j} for $j \geq 1$ is

$$f(y_{t+j}|y_t; \rho, \sigma) \sim \mathcal{N}\left(\rho^j y_t, \sigma^2 \frac{1-\rho^{2j}}{1-\rho^2}\right) \tag{20.3}$$

The predictive distribution (20.3) assumes that the parameters ρ, σ are known, which we express by conditioning on them.

We also want to compute a predictive distribution that does not condition on ρ, σ but instead takes account of our uncertainty about them.

We form this predictive distribution by integrating (20.3) with respect to a joint posterior distribution $\pi_t(\rho, \sigma|y^t)$ that conditions on an observed history $y^t = \{y_s\}_{s=0}^t$:

$$\begin{aligned} f(y_{t+j}|y^t) &= \int f(y_{t+j}|y^t, \rho, \sigma)\pi_t(\rho, \sigma|y^t)d\rho d\sigma \\ &= \int f(y_{t+j}|y_t, \rho, \sigma)\pi_t(\rho, \sigma|y^t)d\rho d\sigma \end{aligned} \tag{20.4}$$

Predictive distribution (20.3) assumes that parameters (ρ, σ) are known.

Predictive distribution (20.4) assumes that parameters (ρ, σ) are uncertain, but have a known probability distribution $\pi_t(\rho, \sigma|y^t)$. Notice the second equality follows because $\{y_t\}$ is an AR(1) process when (ρ, σ) are given.

We also want to compute some predictive distributions of "sample path statistics" that might include, for example

- the time until the next "recession",
- the minimum value of Y over the next 8 periods,
- "severe recession", and
- the time until the next turning point (positive or negative).

To accomplish that for situations in which we are uncertain about parameter values, we shall extend Wecker's [Wecker, 1979] approach in the following way.

- first simulate an initial path of length T_0;
- for a given prior, draw a sample of size N from the posterior joint distribution of parameters (ρ, σ) after observing the initial path;
- for each draw $n = 0, 1, ..., N$, simulate a "future path" of length T_1 with parameters (ρ_n, σ_n) and compute our "sample path statistics";
- finally, plot the desired statistics from the N samples as an empirical distribution.

20.2 Implementation

First, we'll simulate a sample path from which to launch our forecasts.

In addition to plotting the sample path, under the assumption that the true parameter values are known, we'll plot 0.9 and 0.95 coverage intervals using conditional distribution (20.3) described above.

We'll also plot a bunch of samples of sequences of future values and watch where they fall relative to the coverage interval.

```
class AR1(NamedTuple):
    """
    Represents a univariate first-order autoregressive (AR(1)) process.

    Parameters
    ----------
    ρ : float
        Autoregressive coefficient, must satisfy |ρ| < 1 for stationarity.
    σ : float
        Standard deviation of the error term.
    y0 : float
        Initial value of the process at time t=0.
    T0 : int, optional
        Length of the initial observed path (default is 100).
    T1 : int, optional
        Length of the future path to simulate (default is 100).
    """
    ρ: float
    σ: float
    y0: float
    T0: int = 100
    T1: int = 100
```

We create instances through a factory function that validates the parameters.

```
def create_ar1(ρ=0.9, σ=1.0, y0=10.0, T0=100, T1=100):
    """Create an AR(1) instance, checking the parameter restrictions."""
    if not abs(ρ) < 1:
        raise ValueError("ρ must satisfy |ρ| < 1 for stationarity")
```

(continues on next page)

(continued from previous page)

```
    if not σ > 0:
        raise ValueError("σ must be positive")
    return AR1(ρ=ρ, σ=σ, y0=y0, T0=T0, T1=T1)
```

Using the `AR1` class, we can simulate paths more conveniently. The following function simulates an initial path of length T_0.

```
def AR1_simulate_past(ar1: AR1, key=key):
    """
    Simulate a realization of the AR(1) process for T0 periods.

    Parameters
    ----------
    ar1 : AR1
        AR1 named tuple containing parameters (ρ, σ, y0, T0, T1).
    key : jax.random.PRNGKey
        JAX random key for generating random noise.

    Returns
    -------
    initial_path : jax.numpy.ndarray
        Simulated path of the AR(1) process and the initial y0.
    """
    ρ, σ, y0, T0 = ar1.ρ, ar1.σ, ar1.y0, ar1.T0
    # Draw εs
    ε = σ * random.normal(key, (T0,))

    # Set step function
    def ar1_step(y_prev, t_ρ_ε):
        ρ, ε_t = t_ρ_ε
        y_t = ρ * y_prev + ε_t
        return y_t, y_t

    # Scan over time steps
    _, y_seq = lax.scan(ar1_step, y0, (jnp.full(T0, ρ), ε))

    # Concatenate initial value
    initial_path = jnp.concatenate([jnp.array([y0]), y_seq])

    return initial_path
```

Now we define the simulation function that generates a realization of the AR(1) process for the future T_1 periods.

```
def AR1_simulate_future(ar1: AR1, y_T0, N=10, key=key):
    """
    Simulate a realization of the AR(1) process for T1 periods.

    Parameters
    ----------
    ar1 : AR1
        AR1 named tuple containing parameters (ρ, σ, y0, T0, T1).
    y_T0 : float
        Value of the process at time T0.
    N: int
        Number of paths to simulate.
    key : jax.random.PRNGKey
        JAX random key for generating random noise.
```

(continues on next page)

(continued from previous page)

```
    Returns
    -------
    future_path : jax.numpy.ndarray
        Simulated N paths of the AR(1) process of length T1.
    """
    ρ, σ, T1 = ar1.ρ, ar1.σ, ar1.T1

    def single_path_scan(y_T0, subkey):
        ε = σ * random.normal(subkey, (T1,))

        def ar1_step(y_prev, t_ρ_ε):
            ρ, ε_t = t_ρ_ε
            y_t = ρ * y_prev + ε_t
            return y_t, y_t
        _, y = lax.scan(ar1_step, y_T0, (jnp.full(T1, ρ), ε))
        return y

    # Split key to generate different paths
    subkeys = random.split(key, num=N)

    # Simulate N paths
    future_path = jax.vmap(single_path_scan, in_axes=(None, 0))(y_T0, subkeys)

    return future_path
```

The following function plots the initial observed AR(1) path and simulated future paths along with predictive confidence intervals.

```
def plot_path(ar1, initial_path, future_path, ax, key=key):
    """
    Plot the initial observed AR(1) path and simulated future paths,
    along with predictive confidence intervals.

    Parameters
    ----------
    ar1 : AR1
        AR1 named tuple containing process parameters (ρ, σ, T0, T1).
    initial_path : array-like
        Simulated initial path of the AR(1) process, shape(T0+1,).
    future_path : array-like
        Simulated future paths of the AR(1) process, shape (N, T1).
    ax : matplotlib.axes.Axes
        Matplotlib axes object to plot on.
    key : jax.random.PRNGKey, optional
        JAX random key for reproducible sampling.

    Plots
    -----
    - Initial path (historical data)
    - Multiple simulated future paths
    - 90% and 95% predictive confidence intervals
    - Expected future path
    """
    ρ, σ, T0, T1 = ar1.ρ, ar1.σ, ar1.T0, ar1.T1

    # Compute moments and confidence intervals
```

(continues on next page)

(continued from previous page)

```
    y_T0 = initial_path[-1]
    j = jnp.arange(1, T1+1)
    center = ρ**j * y_T0
    vars = σ**2 * (1 - ρ**(2 * j)) / (1 - ρ**2)

    # 95% CI
    y_upper_c95 = center + 1.96 * jnp.sqrt(vars)
    y_lower_c95 = center - 1.96 * jnp.sqrt(vars)

    # 90% CI
    y_upper_c90 = center + 1.65 * jnp.sqrt(vars)
    y_lower_c90 = center - 1.65 * jnp.sqrt(vars)

    # Plot
    ax.plot(jnp.arange(-T0, 1), initial_path, lw=2)
    ax.axvline(0, linestyle='--', alpha=.4, color='k', lw=1)

    # Choose 10 future paths to plot
    index = random.choice(
        key, jnp.arange(future_path.shape[0]), (10,), replace=False
        )
    for i in index:
        ax.plot(jnp.arange(1, T1+1), future_path[i, :], color='grey', alpha=.5)

    # Plot 90% and 95% CI
    ax.fill_between(
        jnp.arange(1, T1+1), y_upper_c95, y_lower_c95, alpha=.3, label='95% CI'
        )
    ax.fill_between(
        jnp.arange(1, T1+1), y_upper_c90, y_lower_c90, alpha=.35, label='90% CI'
        )
    ax.plot(
        jnp.arange(1, T1+1), center, color='red', alpha=.7, lw=2,
        label='expectation'
        )
    ax.set_xlim([-T0, T1])
    ax.set_xlabel("time", fontsize=13)
    ax.set_ylabel("y", fontsize=13)
    ax.legend(fontsize=8)
```

```
ar1 = create_ar1(ρ=0.9, σ=1.0, y0=10.0)

# Simulate
initial_path = AR1_simulate_past(ar1)
future_path = AR1_simulate_future(ar1, initial_path[-1])

# Plot
fig, ax = plt.subplots(1, 1)
plot_path(ar1, initial_path, future_path, ax)
plt.show()
```

As functions of forecast horizon, the coverage intervals have shapes like those described in *Permanent Income II: LQ Techniques*

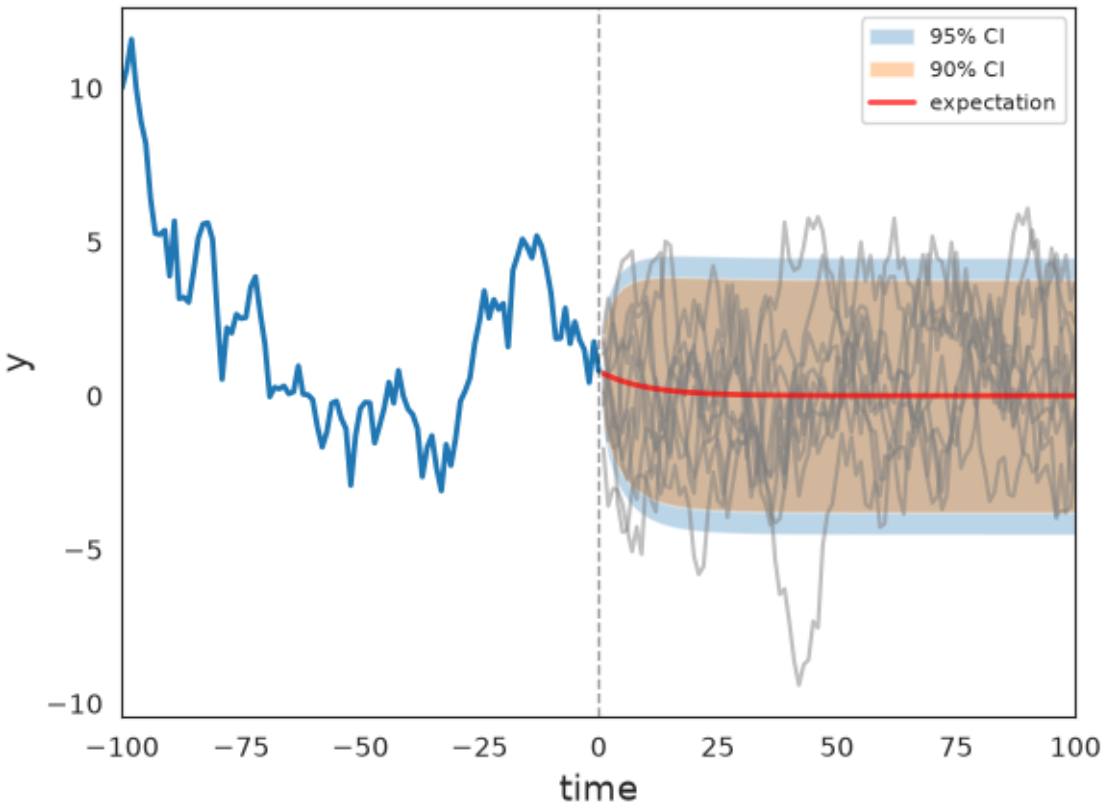


Fig. 20.1: Initial and predictive future paths

20.3 Predictive Distributions of Path Properties

Wecker [Wecker, 1979] proposed using simulation techniques to characterize the predictive distribution of some statistics that are non-linear functions of y.

He called these functions "path properties" to contrast them with properties of single data points.

He studied two special prospective path properties of a given series $\{y_t\}$.

The first was *time until the next turning point*.

- he defined a **"turning point"** to be the date of the second of two successive declines in y.

For example, if $y_t(\omega) < y_{t-1}(\omega) < y_{t-2}(\omega)$, then period t is a turning point.

To examine the *time until the next turning point*, let Z be an indicator process

$$Z_t(\omega) := \begin{cases} 1 & \text{if } Y_t(\omega) < Y_{t-1}(\omega) < Y_{t-2}(\omega) \geq Y_{t-3}(\omega) \\ 0 & \text{otherwise} \end{cases}$$

Here $\omega \in \Omega$ is a sequence of events, and $Y_t : \Omega \to \mathbb{R}$ gives y_t according to ω and the AR(1) process.

By Wecker's definition, period t is a turning point, and $Y_{t-2}(\omega) \geq Y_{t-3}(\omega)$ excludes the possibility that period $t-1$ is a turning point.

Then the random variable **time until the next turning point** is defined as the following *stopping time* with respect to Z:

$$W_t(\omega) := \inf\{k \geq 1 \mid Z_{t+k}(\omega) = 1\}$$

In the following code, we name this statistic *time until the next recession* to distinguish it from another concept of *turning point*.

Moreover, the statistic *time until the next severe recession* is defined in a similar way, except the decline between periods is greater than 0.02.

Wecker [Wecker, 1979] also studied **the minimum value of Y over the next 8 quarters**, which can be defined as the random variable

$$M_t(\omega) := \min\{Y_{t+1}(\omega); Y_{t+2}(\omega); \dots; Y_{t+8}(\omega)\}$$

It is interesting to study yet another possible concept of a **turning point**.

Thus, let

$$T_t(\omega) := \begin{cases} 1 & \text{if } Y_{t-2}(\omega) > Y_{t-1}(\omega) > Y_t(\omega) \text{ and } Y_t(\omega) < Y_{t+1}(\omega) < Y_{t+2}(\omega) \\ -1 & \text{if } Y_{t-2}(\omega) < Y_{t-1}(\omega) < Y_t(\omega) \text{ and } Y_t(\omega) > Y_{t+1}(\omega) > Y_{t+2}(\omega) \\ 0 & \text{otherwise} \end{cases}$$

Define a **positive turning point today or tomorrow** statistic as

$$P_t(\omega) := \begin{cases} 1 & \text{if } T_t(\omega) = 1 \text{ or } T_{t+1}(\omega) = 1 \\ 0 & \text{otherwise} \end{cases}$$

This is designed to express the event

- "after one or two decrease(s), Y will grow for two consecutive quarters"

The **negative turning point today or tomorrow** N_t is defined in the same way.

Following [Wecker, 1979], we can use simulations to calculate probabilities of P_t and N_t for each period t.

However, in the following code, we only use $T_{t+1}(\omega) = 1$ to determine $P_t(\omega)$ and $N_t(\omega)$, because we only want to find the first positive turning point.

20.4 A Wecker-Like Algorithm

The procedure consists of the following steps:

- index a sample path by ω_i
- from a given date t, simulate I sample paths of length N

$$Y(\omega_i) = \{Y_{t+1}(\omega_i), Y_{t+2}(\omega_i), \dots, Y_{t+N}(\omega_i)\}_{i=1}^{I}$$

- for each path ω_i, compute the associated value of $W_t(\omega_i), W_{t+1}(\omega_i), \dots, W_{t+N}$
- consider the sets $\{W_t(\omega_i)\}_{i=1}^{I}$, $\{W_{t+1}(\omega_i)\}_{i=1}^{I}$, ..., $\{W_{t+N}(\omega_i)\}_{i=1}^{I}$ as samples from the predictive distributions $f(W_{t+1} \mid y_t, y_{t-1}, \dots, y_0)$, $f(W_{t+2} \mid y_t, y_{t-1}, \dots, y_0)$, ..., $f(W_{t+N} \mid y_t, y_{t-1}, \dots, y_0)$.

20.5 Using Simulations to Approximate a Posterior Distribution

The next code cells use `numpyro` to compute the time t posterior distribution of ρ, σ.

We construct this posterior just as in *Posterior Distributions for AR(1) Parameters*, to which we refer for a detailed discussion.

As there, in defining the likelihood function we condition on the initial value y_0.

This is the **conditioning assumption** of *Posterior Distributions for AR(1) Parameters*, and it is the appropriate choice here because our initial path starts from an atypical value $y_0 = 10$.

```
def draw_from_posterior(data, size=10000, dis_plot=True, key=key):
    """Draw a sample of the given size from the posterior distribution."""

    def model(data):
        # Start with priors
        ρ = numpyro.sample('ρ', dist.Uniform(-1, 1))  # Assume stable ρ
        σ = numpyro.sample('σ', dist.HalfNormal(jnp.sqrt(10)))

        # Expected value of y at the next period (ρ * y), conditioning on y_0
        yhat = ρ * data[:-1]

        # Likelihood of the actual realization
        numpyro.sample('y_obs', dist.Normal(yhat, σ), obs=data[1:])

    # Compute posterior distribution of parameters
    nuts_kernel = NUTS(model)

    # Define MCMC class to compute the posteriors
    mcmc = MCMC(
        nuts_kernel,
        num_warmup=5000,
        num_samples=size,
        num_chains=4,  # plot 4 chains in the trace
        progress_bar=False,
        chain_method='vectorized'
    )

    # Run MCMC
    mcmc.run(key, data=data)
```

(continues on next page)

(continued from previous page)

```
    # Get posterior samples
    post_sample = {
        'ρ': mcmc.get_samples()['ρ'],
        'σ': mcmc.get_samples()['σ'],
    }

    # Plot posterior distributions and trace plots
    if dis_plot:
        plot_data = az.from_numpyro(posterior=mcmc)
        az.plot_trace(plot_data, var_names=['ρ', 'σ'])

    return post_sample

post_samples = draw_from_posterior(initial_path)
```

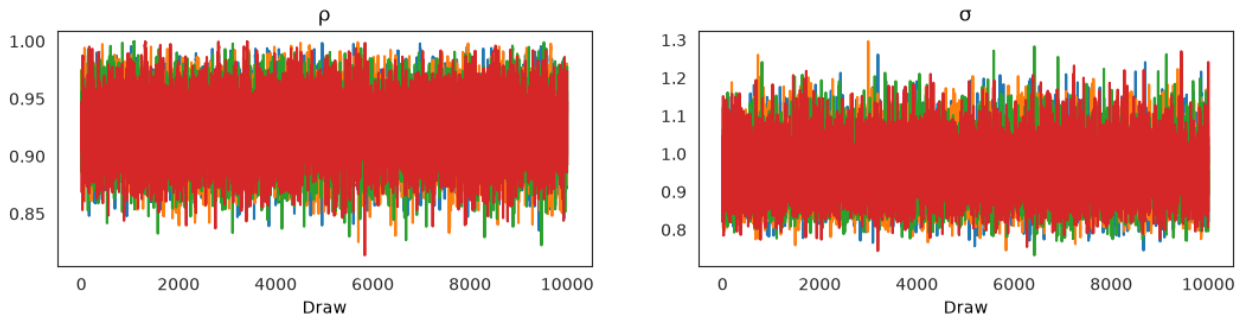


Fig. 20.2: Posterior and trace plots

The graphs above portray posterior distributions and trace plots. The posterior distributions (left column) show the marginal distributions of the parameters after observing the data, while the trace plots (right column) help diagnose MCMC convergence by showing how the sampler explored the parameter space over iterations.

20.6 Calculating Sample Path Statistics

Our next step is to prepare Python code to compute our sample path statistics.

These statistics were originally defined as random variables with respect to ω, but here we use $\{Y_t\}$ as the argument because ω is implicit.

These two kinds of definitions are equivalent because ω determines path statistics only through $\{Y_t\}$.

Moreover, we ignore all equality in the definitions, as equality occurs with zero probability for continuous random variables.

```
@jax.jit
def compute_path_statistics(initial_path, future_path):
    """Compute path statistics for the AR(1) process."""
    # Prepend the last three observed values so we can identify a turn at t=0
    y = jnp.concatenate([initial_path[-3:], future_path])
    n = y.shape[0]
```

(continues on next page)

(continued from previous page)

```
    def step(carry, i):
        # identify recession
        rec_cond = (y[i] < y[i-1]) & (y[i-1] < y[i-2]) & (y[i-2] > y[i-3])
        # identify severe recession
        sev_cond = (
            (y[i] - y[i-1] < -0.02) & (y[i-1] - y[i-2] < -0.02) & (y[i-2] > y[i-3])
            )

        # identify positive turning point
        up_cond = (
            (y[i-2] > y[i-1]) & (y[i-1] > y[i]) & (y[i] < y[i+1]) & (y[i+1] < y[i+2])
        )

        # identify negative turning point
        down_cond = (
            (y[i-2] < y[i-1]) & (y[i-1] < y[i]) & (y[i] > y[i+1]) & (y[i+1] > y[i+2])
        )

        # Convert to int
        rec = jnp.where(rec_cond, 1, 0)
        sev = jnp.where(sev_cond, 1, 0)
        up = jnp.where(up_cond, 1, 0)
        down = jnp.where(down_cond, 1, 0)
        return carry, (rec, sev, up, down)

    _, (rec_seq, sev_seq, up_seq, down_seq) = lax.scan(step, None, jnp.arange(3, n-2))

    # Get the time until the first recession
    next_recession = jnp.where(
        jnp.any(rec_seq == 1), jnp.argmax(rec_seq == 1) + 1, len(y)
    )
    next_severe_recession = jnp.where(
        jnp.any(sev_seq == 1), jnp.argmax(sev_seq == 1) + 1, len(y)
    )

    # Minimum value in the next 8 periods
    min_val_8q = jnp.min(future_path[:8])

    # Get the time until the first turning point
    next_up_turn = jnp.where(
        jnp.any(up_seq == 1),
        jnp.maximum(jnp.argmax(up_seq == 1), 1), # Exclude 0 return
        len(y)
    )
    next_down_turn = jnp.where(
        jnp.any(down_seq == 1),
        jnp.maximum(jnp.argmax(down_seq == 1), 1),
        len(y)
    )
    path_stats = (
        next_recession, next_severe_recession, min_val_8q,
        next_up_turn, next_down_turn
        )
    return path_stats
```

The following function creates visualizations of the path statistics in a subplot grid.

```
def plot_path_stats(next_recession, next_severe_recession, min_val_8q,
                    next_up_turn, next_down_turn, ax, label=None):
    """Plot the path statistics in subplots(3,2)"""
    # ax[0, 0] is for paths of y
    sns.histplot(next_recession, kde=True, stat='density', ax=ax[0, 1],
                 alpha=.8, label=label)
    ax[0, 1].set_xlabel("time until the next recession", fontsize=13)

    sns.histplot(next_severe_recession, kde=True, stat='density', ax=ax[1, 0],
                 alpha=.8, label=label)
    ax[1, 0].set_xlabel("time until the next severe recession", fontsize=13)

    sns.histplot(min_val_8q, kde=True, stat='density', ax=ax[1, 1],
                 alpha=.8, label=label)
    ax[1, 1].set_xlabel("minimum value in next 8 periods", fontsize=13)

    sns.histplot(next_up_turn, kde=True, stat='density', ax=ax[2, 0],
                 alpha=.8, label=label)
    ax[2, 0].set_xlabel("time until the next positive turn", fontsize=13)

    sns.histplot(next_down_turn, kde=True, stat='density', ax=ax[2, 1],
                 alpha=.8, label=label)
    ax[2, 1].set_xlabel("time until the next negative turn", fontsize=13)
```

20.7 Original Wecker Method

Now we apply Wecker's original method by simulating future paths and compute predictive distributions, conditioning on the true parameters associated with the data-generating model.

```
def plot_Wecker(ar1: AR1, initial_path, ax, N=1000, label='true parameters'):
    """
    Plot the predictive distributions from "pure" Wecker's method.

    Parameters
    ----------
    ar1 : AR1
        An AR1 named tuple containing the process parameters (ρ, σ, T0, T1).
    initial_path : array-like
        The initial observed path of the AR(1) process.
    N : int
        Number of future sample paths to simulate for predictive distributions.
    label : str
        Legend label for the plotted distributions.
    """
    # Plot simulated initial and future paths
    y_T0 = initial_path[-1]
    future_path = AR1_simulate_future(ar1, y_T0, N=N)
    plot_path(ar1, initial_path, future_path, ax[0, 0])

    next_reces = jnp.zeros(N)
    severe_rec = jnp.zeros(N)
    min_val_8q = jnp.zeros(N)
    next_up_turn = jnp.zeros(N)
    next_down_turn = jnp.zeros(N)
```

(continues on next page)

(continued from previous page)

```
    # Simulate future paths and compute statistics
    for n in range(N):
        future_temp = future_path[n, :]
        (next_reces_val, severe_rec_val, min_val_8q_val,
        next_up_turn_val, next_down_turn_val
         ) = compute_path_statistics(initial_path, future_temp)

        next_reces = next_reces.at[n].set(next_reces_val)
        severe_rec = severe_rec.at[n].set(severe_rec_val)
        min_val_8q = min_val_8q.at[n].set(min_val_8q_val)
        next_up_turn = next_up_turn.at[n].set(next_up_turn_val)
        next_down_turn = next_down_turn.at[n].set(next_down_turn_val)

    # Plot path statistics
    plot_path_stats(next_reces, severe_rec, min_val_8q,
                    next_up_turn, next_down_turn, ax, label=label)


fig, ax = plt.subplots(3, 2, figsize=(15, 12))
plot_Wecker(ar1, initial_path, ax)
plt.show()
```

Apart from the top-left panel, which repeats the initial path and coverage intervals, each panel shows the predictive distribution of one sample path statistic.

These are computed by simulating many future paths with the parameters held fixed at their true values.

The distributions therefore reflect only the uncertainty that comes from future shocks.

20.8 Extended Wecker Method

Now we apply our "extended" Wecker method based on predictive densities of y defined by (20.4) that acknowledge posterior uncertainty in the parameters ρ, σ.

To approximate the integration on the right side of (20.4), we repeatedly draw parameters from the joint posterior distribution each time we simulate a sequence of future values from model (20.1).

```
def plot_extended_Wecker(
    ar1: AR1, post_samples, initial_path, ax, N=1000,
    label='sampling from posterior'
    ):
    """Plot the extended Wecker's predictive distribution"""
    y0, T1 = ar1.y0, ar1.T1
    y_T0 = initial_path[-1]

    # Select a parameter sample
    index = random.choice(
        key, jnp.arange(len(post_samples['ρ'])), (N,), replace=False
        )
    ρ_sample = post_samples['ρ'][index]
    σ_sample = post_samples['σ'][index]

    # Compute path statistics
```

(continues on next page)

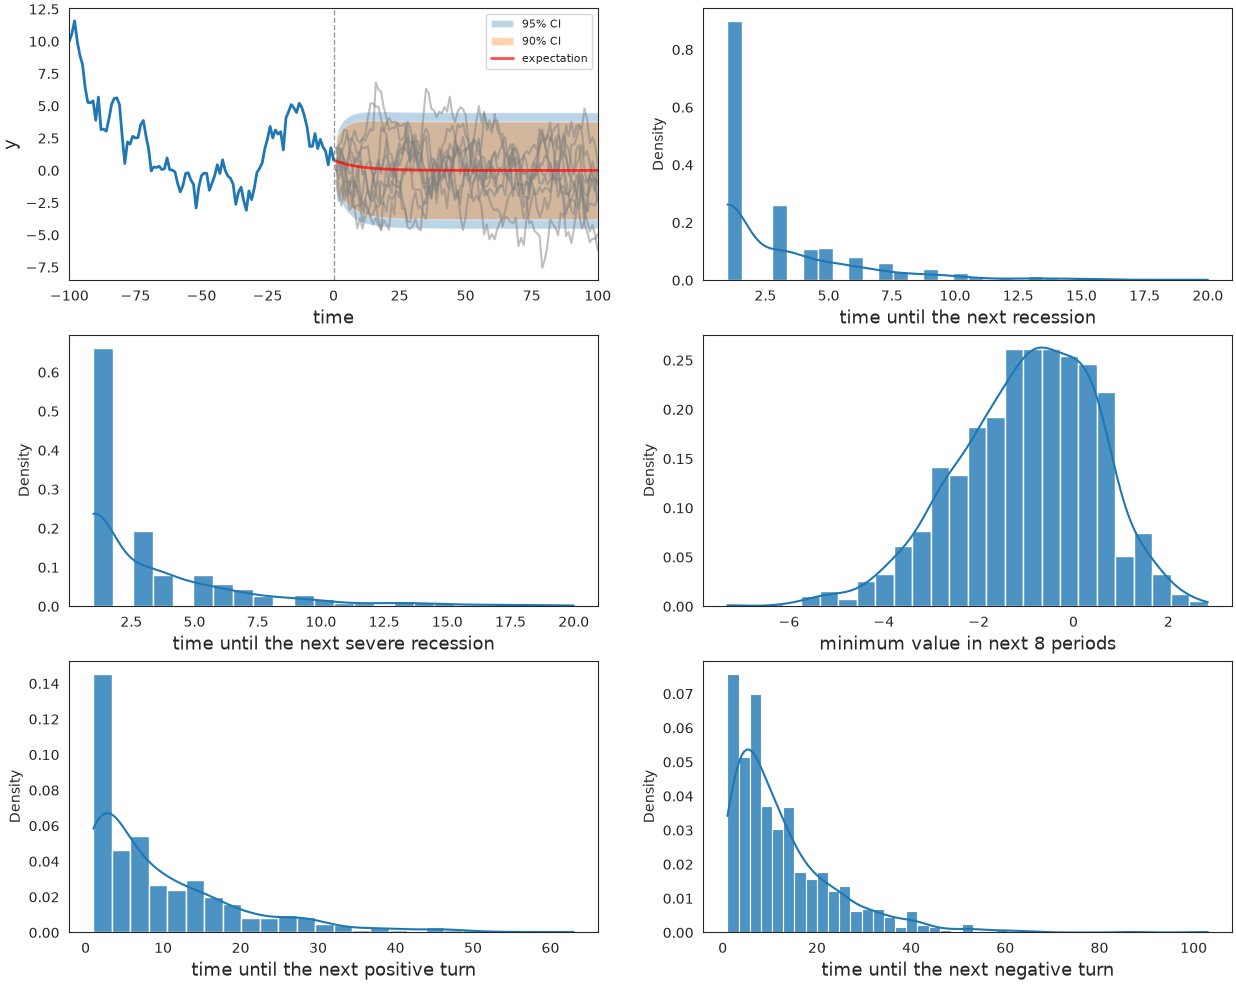


Fig. 20.3: Distributions from Wecker's method

(continued from previous page)

```
    next_reces = jnp.zeros(N)
    severe_rec = jnp.zeros(N)
    min_val_8q = jnp.zeros(N)
    next_up_turn = jnp.zeros(N)
    next_down_turn = jnp.zeros(N)

    subkeys = random.split(key, num=N)

    # Simulate one future path per posterior draw
    future_paths = []
    for n in range(N):
        ar1_n = AR1(ρ=ρ_sample[n], σ=σ_sample[n], y0=y0, T1=T1)
        future_temp = AR1_simulate_future(
            ar1_n, y_T0, N=1, key=subkeys[n]
            ).reshape(-1)
        future_paths.append(future_temp)
        (next_reces_val, severe_rec_val, min_val_8q_val,
        next_up_turn_val, next_down_turn_val
         ) = compute_path_statistics(initial_path, future_temp)

        next_reces = next_reces.at[n].set(next_reces_val)
        severe_rec = severe_rec.at[n].set(severe_rec_val)
        min_val_8q = min_val_8q.at[n].set(min_val_8q_val)
        next_up_turn = next_up_turn.at[n].set(next_up_turn_val)
        next_down_turn = next_down_turn.at[n].set(next_down_turn_val)

    # Plot the initial path and the future paths drawn from the posterior
    future_paths = jnp.stack(future_paths)
    plot_path(ar1, initial_path, future_paths, ax[0, 0])

    # Plot path statistics
    plot_path_stats(next_reces, severe_rec, min_val_8q,
                    next_up_turn, next_down_turn, ax, label=label)

fig, ax = plt.subplots(3, 2, figsize=(12, 15))
plot_extended_Wecker(ar1, post_samples, initial_path, ax)
plt.show()
```

The panels show the same statistics, but now each future path is simulated with parameters drawn from the posterior.

These distributions therefore combine two sources of uncertainty: randomness in future shocks and our uncertainty about (ρ, σ).

20.9 Comparison

Finally, we plot both the original Wecker method and the extended method with parameter values drawn from the posterior together to compare the differences that emerge from pretending to know parameter values when they are actually uncertain.

```
fig, ax = plt.subplots(3, 2, figsize=(12, 15))
plot_Wecker(ar1, initial_path, ax)
ax[0, 0].clear()
plot_extended_Wecker(ar1, post_samples, initial_path, ax)
```

(continues on next page)

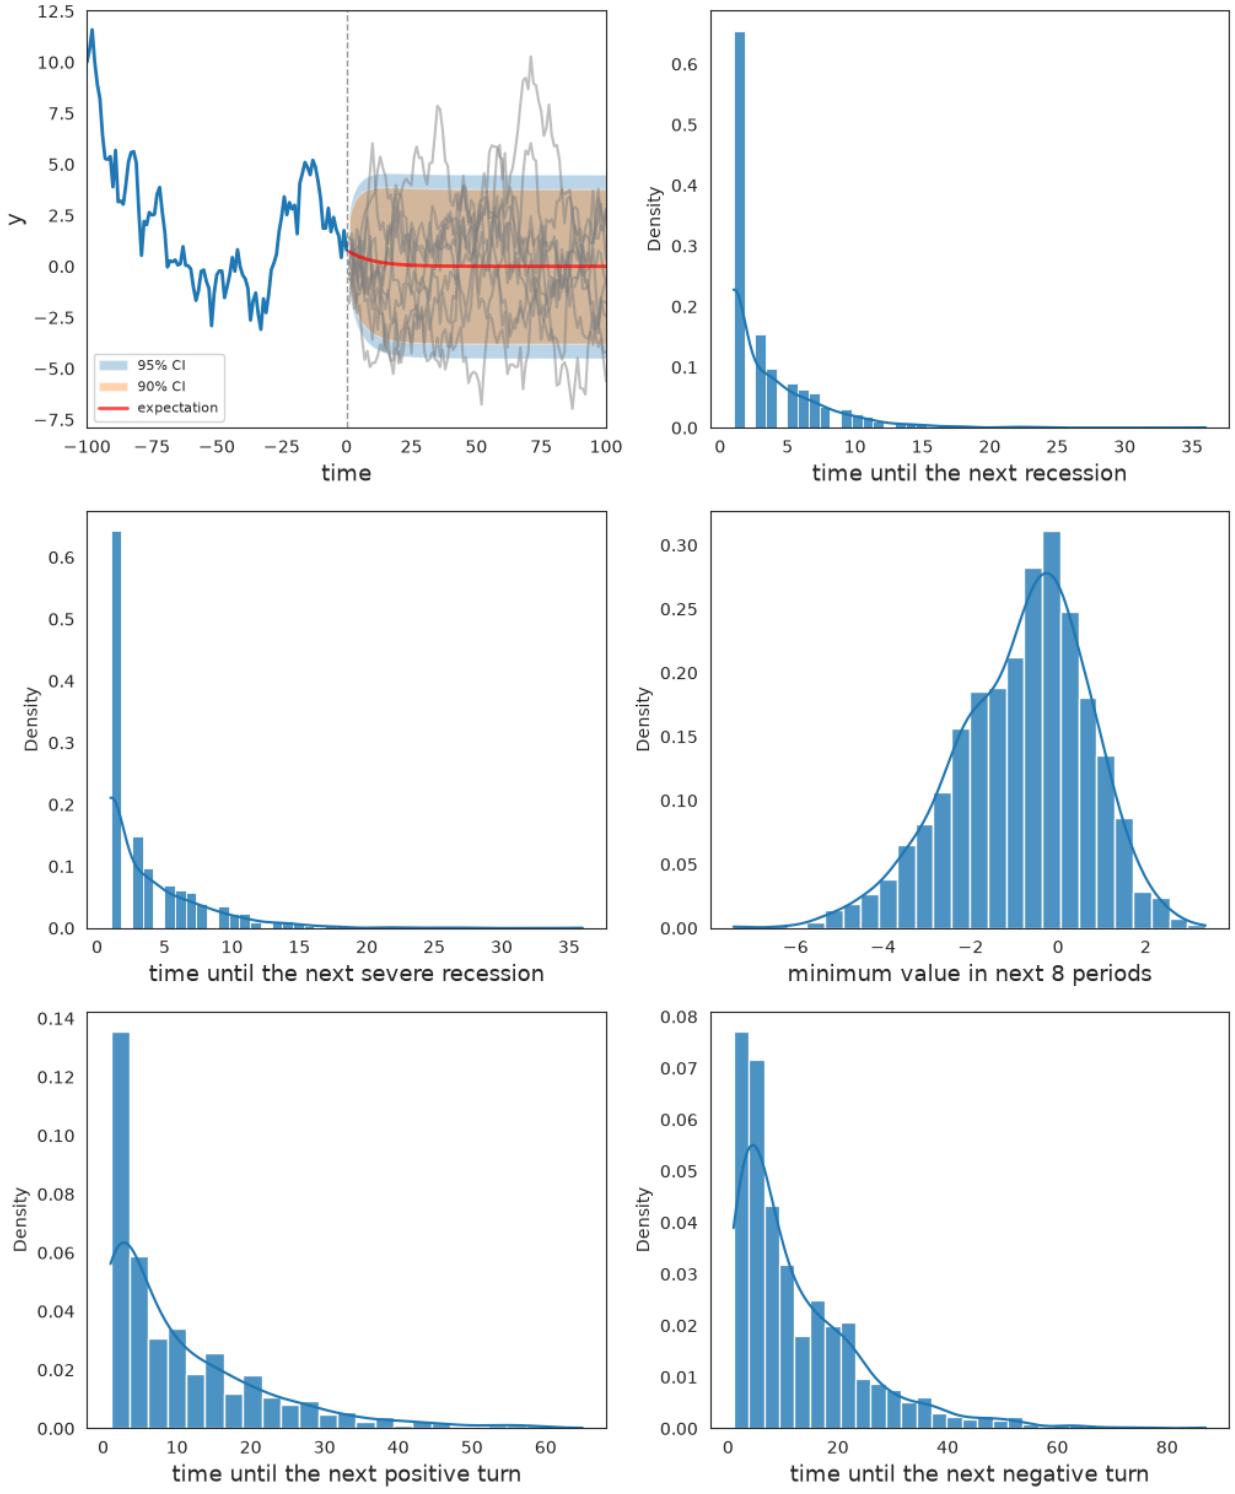


Fig. 20.4: Distributions from extended Wecker's method

(continued from previous page)

```
ax[0, 1].legend(fontsize=10)
plt.show()
```

The two sets of predictive distributions reveal the cost of pretending that we know the parameters.

The extended Wecker method draws (ρ, σ) from their posterior each time it simulates a future path.

It therefore layers parameter uncertainty on top of the shock uncertainty already present in the original method.

As a result, its predictive distributions are more dispersed than those computed with the parameters fixed at their true values.

20.10 Conclusion

This lecture combined two tools to forecast nonlinear functions of the future path of an AR(1) process.

Wecker's simulation method let us approximate predictive distributions of sample path statistics, such as the time until the next turning point.

A Bayesian posterior over (ρ, σ), computed as in *Posterior Distributions for AR(1) Parameters*, let us also account for our uncertainty about the parameters that govern the process.

Putting these together, the extended Wecker method produces predictive distributions that reflect both sources of uncertainty identified at the outset.

Ignoring parameter uncertainty, as the original Wecker method does, yields tighter distributions that overstate how much we really know about these future statistics.

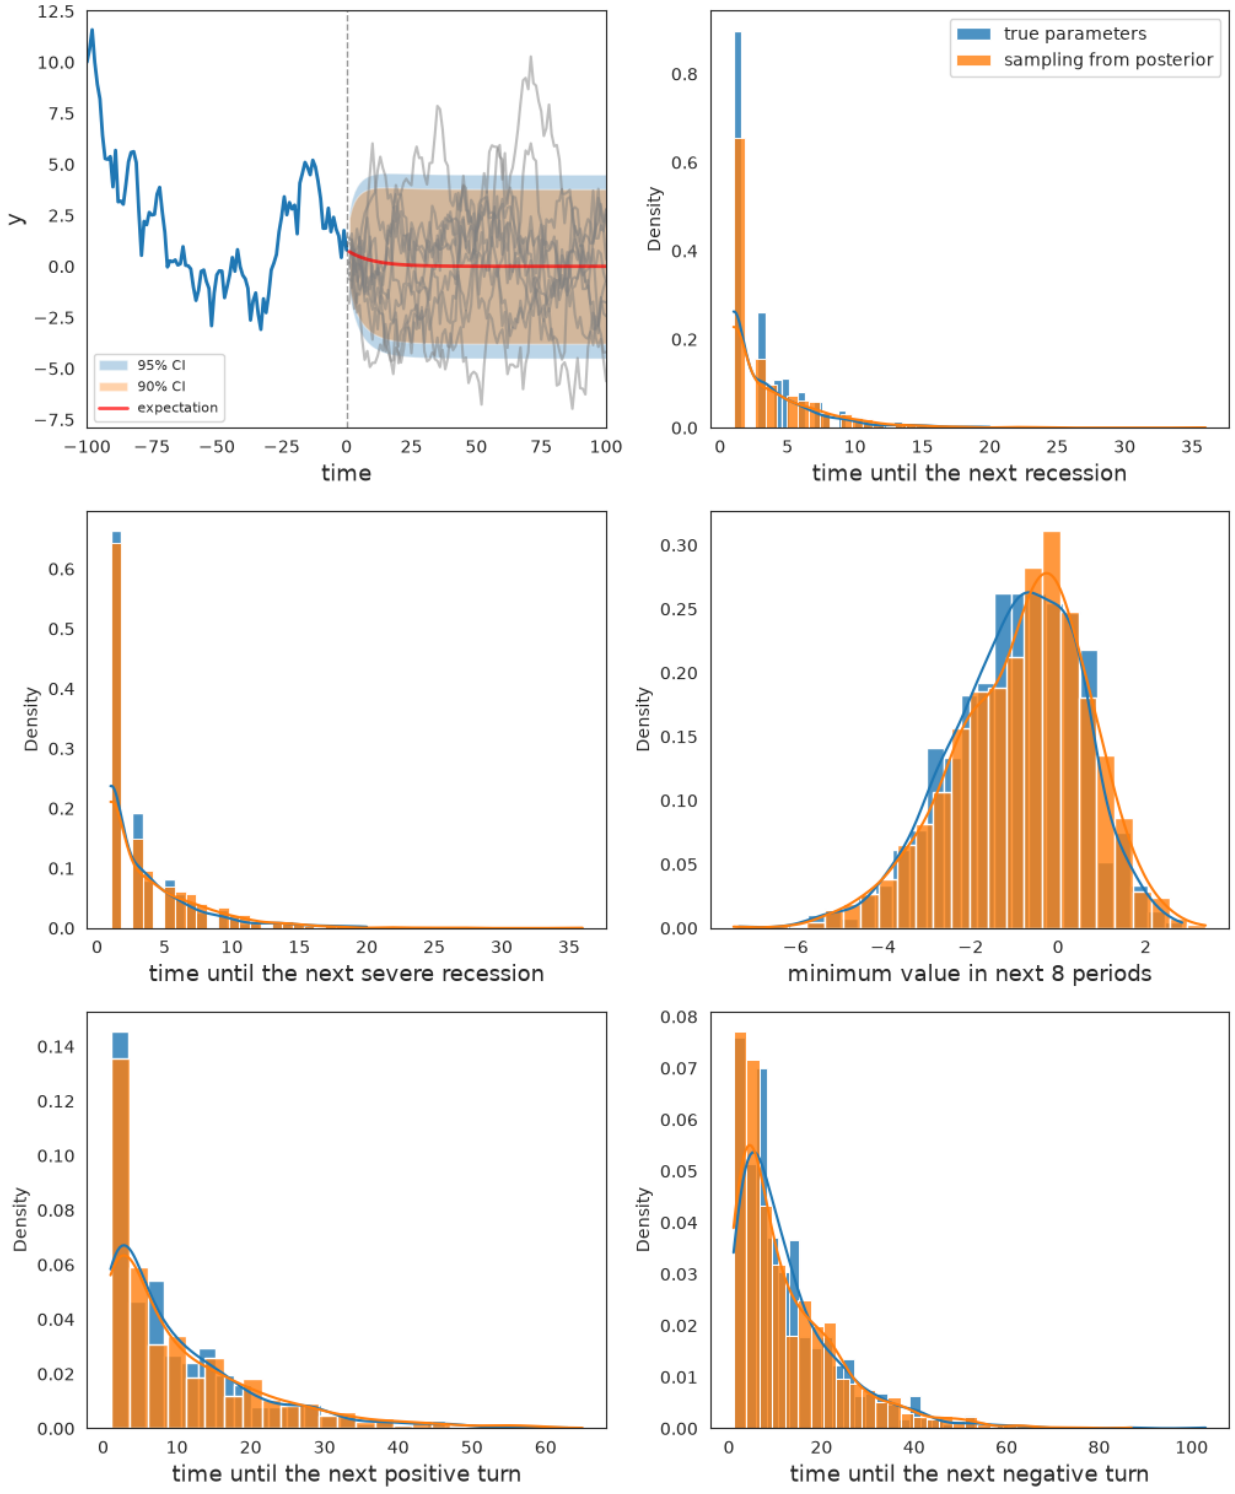


Fig. 20.5: Comparison of the two methods

Part IV

Statistics and Information

CHAPTER

TWENTYONE

STATISTICAL DIVERGENCE MEASURES

Contents

- *Statistical Divergence Measures*
  - *Overview*
  - *Primer on entropy, cross-entropy, KL divergence*
  - *Two Beta distributions: running example*
  - *Kullback–Leibler divergence*
  - *Jensen-Shannon divergence*
  - *Chernoff entropy*
  - *Comparing divergence measures*
  - *KL divergence and maximum-likelihood estimation*
  - *Related lectures*
  - *Exercises*

21.1 Overview

A statistical divergence quantifies discrepancies between two distinct probability distributions that can be challenging to distinguish for the following reason:

- every event that has positive probability under one of the distributions also has positive probability under the other distribution
- this means that there is no "smoking gun" event whose occurrence tells a statistician that one of the probability distributions surely governs the data

A statistical divergence is a **function** that maps two probability distributions into a number in $[0, +\infty]$.

The number is zero if and only if the two distributions are equal.

Statistical divergence functions play important roles in statistics, information theory, and what many people now call "machine learning".

This lecture describes three divergence measures:

- **Kullback–Leibler (KL) divergence**

- **Jensen–Shannon (JS) divergence**
- **Chernoff entropy**

These will appear in several quantecon lectures.

Let's start by importing the necessary Python tools.

```
import matplotlib.pyplot as plt
import numpy as np
from numba import vectorize, jit
from math import gamma
from scipy.integrate import quad
from scipy.optimize import minimize_scalar
from scipy.special import betaln, digamma
from scipy.stats import beta as beta_dist
import pandas as pd
from IPython.display import display, Math
```

21.2 Primer on entropy, cross-entropy, KL divergence

Before diving in, we'll introduce some useful concepts in a simple setting.

We'll temporarily assume that f and g are two probability mass functions for discrete random variables on state space $I = \{1, 2, \ldots, n\}$ that satisfy $f_i \geq 0, \sum_i f_i = 1, g_i \geq 0, \sum_i g_i = 1$.

We follow some statisticians and information theorists who define the **surprise** or **surprisal** associated with having observed a single draw $x = i$ from distribution f as

$$\log\left(\frac{1}{f_i}\right)$$

They then define the **information** that you can anticipate to gather from observing a single realization as the expected surprisal

$$H(f) = \sum_i f_i \log\left(\frac{1}{f_i}\right).$$

Claude Shannon [Shannon, 1948] called $H(f)$ the **entropy** of distribution f.

Note

By maximizing $H(f)$ with respect to $\{f_1, f_2, \ldots, f_n\}$ subject to $\sum_i f_i = 1$, we can verify that the distribution that maximizes entropy is the uniform distribution $f_i = \frac{1}{n}$. Entropy $H(f)$ for the uniform distribution equals $\log n$.

Kullback and Leibler [1951] define the amount of information that a single draw $x = i$ provides for distinguishing f from g as the log likelihood ratio

$$\log\frac{f_i}{g_i}$$

The following two concepts are widely used to compare two distributions f and g.

Cross-Entropy:

$$H(f, g) = -\sum_i f_i \log g_i \tag{21.1}$$

Kullback-Leibler (KL) Divergence:

$$D_{KL}(f \parallel g) = \sum_i f_i \log \left[\frac{f_i}{g_i} \right] \tag{21.2}$$

These concepts are related by the following equality.

$$D_{KL}(f \parallel g) = H(f,g) - H(f) \tag{21.3}$$

To prove (21.3), note that

$$\begin{aligned} D_{KL}(f \parallel g) &= \sum_i f_i \log \left[\frac{f_i}{g_i} \right] && (21.4)\\ &= \sum_i f_i \left[\log f_i - \log g_i \right] && (21.5)\\ &= \sum_i f_i \log f_i - \sum_i f_i \log g_i && (21.6)\\ &= -H(f) + H(f,g) && (21.7)\\ &= H(f,g) - H(f) && (21.8) \end{aligned}$$

Remember that $H(f)$ is the anticipated surprisal from drawing x from f.

The cross-entropy $H(f,g) = \sum_i f_i \log(1/g_i)$ is the anticipated surprisal when x is drawn from f but surprisal is computed with g.

Then the above equation tells us that the KL divergence is an anticipated "excess surprise" that comes from anticipating that x is drawn from g when it is actually drawn from f.

21.3 Two Beta distributions: running example

We'll use Beta distributions extensively to illustrate concepts.

The Beta distribution is particularly convenient as it's defined on $[0,1]$ and exhibits diverse shapes by appropriately choosing its two parameters.

The density of a Beta distribution with parameters a and b is given by

$$f(z;a,b) = \frac{\Gamma(a+b) z^{a-1}(1-z)^{b-1}}{\Gamma(a)\Gamma(b)} \quad \text{where} \quad \Gamma(p) := \int_0^\infty x^{p-1} e^{-x} dx$$

We introduce two Beta distributions $f(x)$ and $g(x)$, which we will use to illustrate the different divergence measures.

Let's define parameters and density functions in Python

```
# Parameters in the two Beta distributions
F_a, F_b = 1, 1
G_a, G_b = 3, 1.2

@vectorize
def p(x, a, b):
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x** (a-1) * (1 - x) ** (b-1)

# The two density functions
f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))
```

(continues on next page)

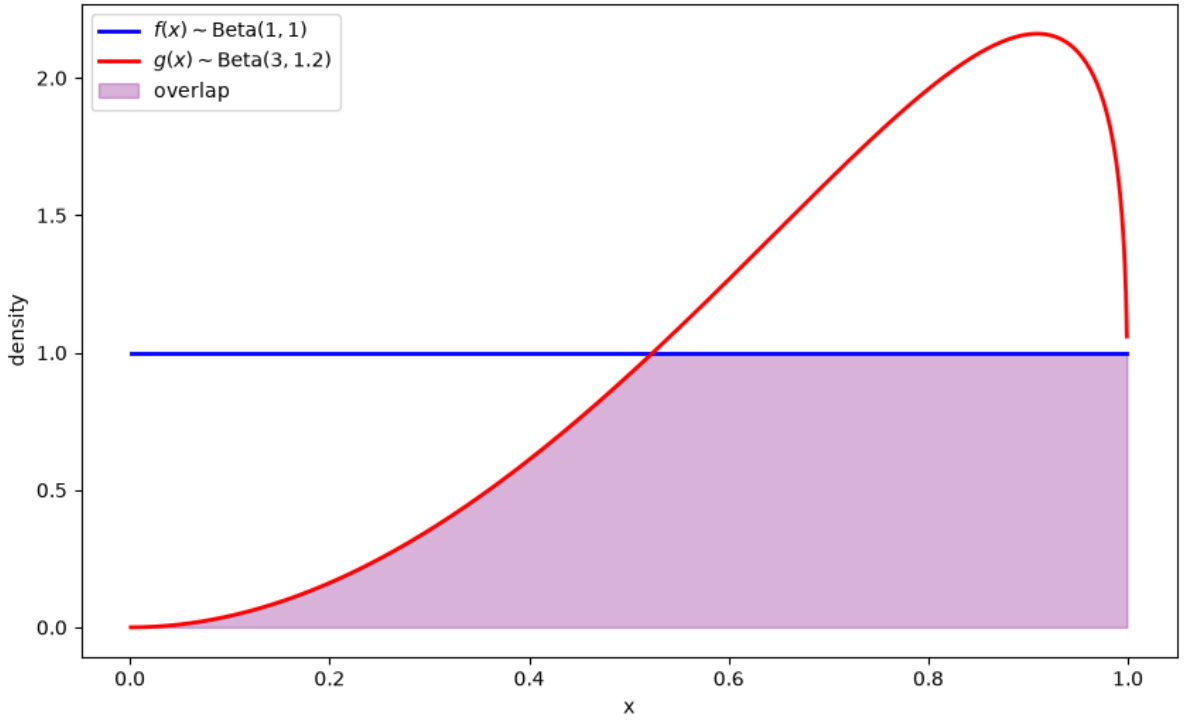
(continued from previous page)

```
# Plot the distributions
x_range = np.linspace(0.001, 0.999, 1000)
f_vals = [f(x) for x in x_range]
g_vals = [g(x) for x in x_range]

plt.figure(figsize=(10, 6))
plt.plot(x_range, f_vals, 'b-', linewidth=2, label=r'$f(x) \sim \text{Beta}(1,1)$')
plt.plot(x_range, g_vals, 'r-', linewidth=2, label=r'$g(x) \sim \text{Beta}(3,1.2)$')

# Fill overlap region
overlap = np.minimum(f_vals, g_vals)
plt.fill_between(x_range, 0, overlap, alpha=0.3, color='purple', label='overlap')

plt.xlabel('x')
plt.ylabel('density')
plt.legend()
plt.show()
```



21.4 Kullback–Leibler divergence

Our first divergence function is the **Kullback–Leibler (KL) divergence**.

For probability densities (or pmfs) f and g it is defined by

$$D_{KL}(f\|g) = KL(f,g) = \int f(x) \log \frac{f(x)}{g(x)} \, dx.$$

We can interpret $D_{KL}(f\|g)$ as the expected excess log loss (expected excess surprisal) incurred when we use g while the data are generated by f.

$D_{KL}(f\|g) = +\infty$ unless f is absolutely continuous with respect to g ($f \ll g$), that is, unless $g(x) > 0$ wherever $f(x) > 0$.

It has several important properties:

- Non-negativity (Gibbs' inequality): $D_{KL}(f\|g) \geq 0$ with equality if and only if $f = g$ almost everywhere.
- Asymmetry: $D_{KL}(f\|g) \neq D_{KL}(g\|f)$ in general (hence it is not a metric)
- Information decomposition: $D_{KL}(f\|g) = H(f,g) - H(f)$, where $H(f,g)$ is the cross entropy and $H(f)$ is the Shannon entropy of f.
- Chain rule: For joint distributions $f(x,y)$ and $g(x,y)$, $D_{KL}(f(x,y)\|g(x,y)) = D_{KL}(f(x)\|g(x)) + E_f\left[D_{KL}(f(y|x)\|g(y|x))\right]$
- Data processing inequality: if a Markov kernel κ maps x into a coarser signal y, and $f\kappa$ and $g\kappa$ denote the induced distributions of y, then $D_{KL}(f\kappa\|g\kappa) \leq D_{KL}(f\|g)$.

The data processing inequality says that garbling data cannot make f and g easier to tell apart.

The section on *the data processing inequality and coarse-graining* in *Blackwell's Theorem on Comparing Experiments* connects it to Blackwell's ordering of experiments.

JS divergence and Chernoff entropy, defined below, also satisfy the data processing inequality.

KL divergence plays a central role in statistical inference, including model selection and hypothesis testing.

Likelihood Ratio Processes describes a link between KL divergence and the expected log likelihood ratio.

A Problem that Stumped Milton Friedman connects KL divergence to the expected stopping time of a sequential probability ratio test.

Incorrect Models uses two KL divergences to determine the limit of a Bayesian's posterior probability when the data come from a mixture of f and g that the Bayesian's model excludes.

Let's compute the KL divergence between our example distributions f and g.

A U-shaped Beta density diverges at 0 or 1.

We integrate over all of $(0,1)$ and split the interval at $1/2$ so that `quad` handles one endpoint singularity at a time.

```
def integrate(F):
    """Integrate F over (0, 1), splitting the interval at 1/2."""
    left, _ = quad(F, 0, 0.5, limit=200)
    right, _ = quad(F, 0.5, 1, limit=200)
    return left + right

def compute_KL(f, g):
    """
    Compute KL divergence KL(f, g) via numerical integration
    """
    def integrand(w):
        fw = f(w)
        gw = g(w)
        return fw * np.log(fw / gw)
    return integrate(integrand)

# Compute KL divergences between our example distributions
kl_fg = compute_KL(f, g)
kl_gf = compute_KL(g, f)

print(f"KL(f, g) = {kl_fg:.4f}")
print(f"KL(g, f) = {kl_gf:.4f}")
```

```
KL(f, g) = 0.7592
KL(g, f) = 0.3436
```

The asymmetry of KL divergence has important practical implications.

$D_{KL}(f\|g)$ heavily penalizes regions where $f > 0$ but g is close to zero, because f puts weight on the large values of $\log(f/g)$ there.

$D_{KL}(f\|g)$ barely penalizes regions where $g > 0$ but f is close to zero, because f puts almost no weight there.

$D_{KL}(g\|f)$ reverses these roles.

In our example, g vanishes at $x = 0$ while $f = 1$ there, so $D_{KL}(f\|g)$ exceeds $D_{KL}(g\|f)$.

21.5 Jensen-Shannon divergence

Sometimes we want a symmetric measure of divergence that captures the difference between two distributions without favoring one over the other.

This often arises in applications like clustering, where we want to compare distributions without assuming one is the true model.

The **Jensen-Shannon (JS) divergence** symmetrizes KL divergence by comparing both distributions to their mixture:

$$JS(f, g) = \frac{1}{2} D_{KL}(f\|m) + \frac{1}{2} D_{KL}(g\|m), \quad m = \frac{1}{2}(f + g).$$

where m is a mixture distribution that averages f and g

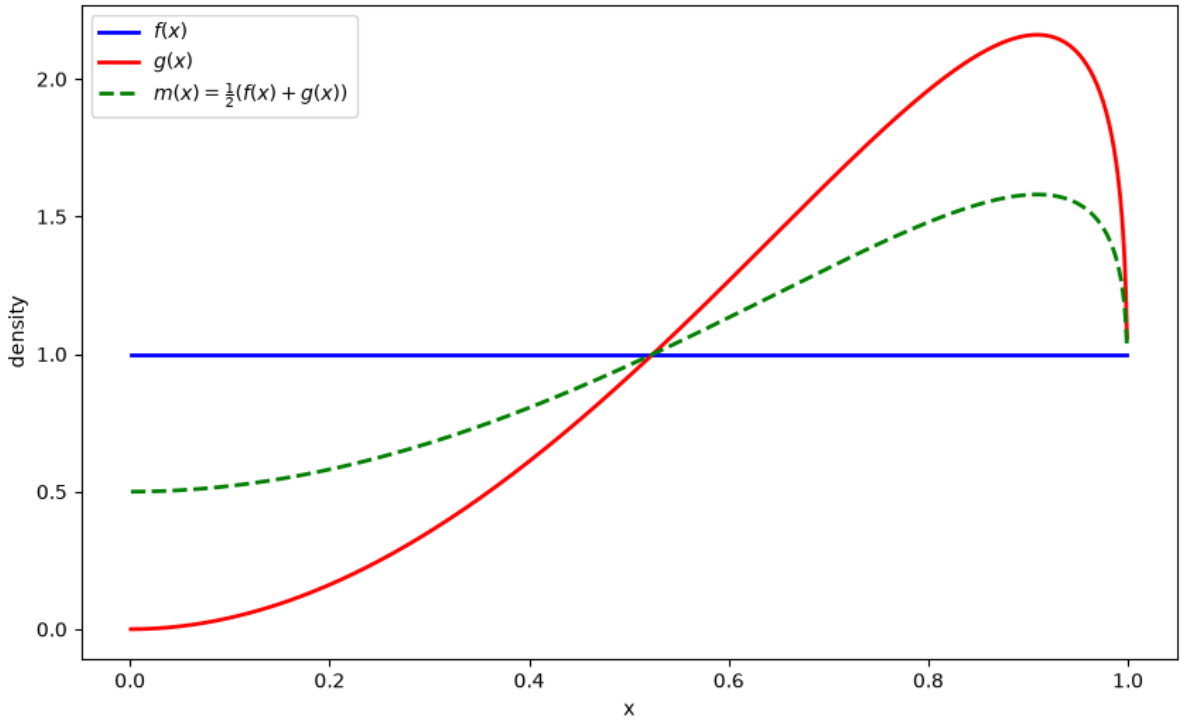
Let's also visualize the mixture distribution m:

```
def m(x):
    return 0.5 * (f(x) + g(x))

m_vals = [m(x) for x in x_range]

plt.figure(figsize=(10, 6))
plt.plot(x_range, f_vals, 'b-', linewidth=2, label=r'$f(x)$')
plt.plot(x_range, g_vals, 'r-', linewidth=2, label=r'$g(x)$')
plt.plot(x_range, m_vals, 'g--', linewidth=2, label=r'$m(x) = \frac{1}{2}(f(x) + g(x))
↪$')

plt.xlabel('x')
plt.ylabel('density')
plt.legend()
plt.show()
```



The JS divergence has several useful properties:

- Symmetry: $JS(f, g) = JS(g, f)$.
- Boundedness: $0 \leq JS(f, g) \leq \log 2$.
- Its square root $\sqrt{JS}$ is a metric (Jensen–Shannon distance) on the space of probability distributions.
- JS divergence equals the mutual information between a binary random variable $Z \sim \text{Bernoulli}(1/2)$ indicating the source and a sample X drawn from f if $Z = 0$ or from g if $Z = 1$.

The Jensen–Shannon divergence appears in the original formulation of generative adversarial networks: when the discriminator is optimal, the generator minimizes $2\,JS(p_{\text{data}}, p_{\text{model}}) - \log 4$ [Goodfellow *et al.*, 2014].

Boundedness has a cost.

When two distributions have nearly disjoint supports, JS sits near its ceiling $\log 2$ and is nearly flat in the parameters of the model, so gradient-based training stalls [Arjovsky and Bottou, 2017].

Let's compute the JS divergence between our example distributions f and g

```
def compute_JS(f, g):
    """Compute Jensen-Shannon divergence."""
    def integrand(w):
        fw, gw = f(w), g(w)
        mw = 0.5 * (fw + gw)
        return 0.5 * fw * np.log(fw / mw) + 0.5 * gw * np.log(gw / mw)
    return integrate(integrand)

js_div = compute_JS(f, g)
print(f"Jensen-Shannon divergence JS(f,g) = {js_div:.4f}")
```

```
Jensen-Shannon divergence JS(f,g) = 0.0984
```

We can easily generalize to more than two distributions using the generalized Jensen-Shannon divergence with weights $\alpha = (\alpha_i)_{i=1}^n$:

$$JS_\alpha(f_1, \ldots, f_n) = H\left(\sum_{i=1}^n \alpha_i f_i\right) - \sum_{i=1}^n \alpha_i H(f_i)$$

where:

- $\alpha_i \geq 0$ and $\sum_{i=1}^n \alpha_i = 1$, and
- $H(f) = -\int f(x) \log f(x) dx$ is the **differential entropy** of density f

Unlike the entropy of a pmf, differential entropy can be negative: the uniform density on $[0, 1/2]$ has differential entropy $-\log 2$.

JS_α is nonnegative nonetheless, because H is concave.

21.6 Chernoff entropy

Chernoff entropy originates from early applications of the theory of large deviations, which refines central limit approximations by providing exponential decay rates for rare events.

For densities f and g the Chernoff entropy, which Chernoff [1952] introduced and information theorists call the Chernoff information, is

$$C(f, g) = -\log \min_{\phi \in (0,1)} \rho(\phi), \qquad \rho(\phi) = \int f^\phi(x) g^{1-\phi}(x)\, dx.$$

Remarks:

- The integral $\rho(\phi)$ is the **Chernoff coefficient**.
- At $\phi = 1/2$ it becomes the **Bhattacharyya coefficient** $\int \sqrt{fg}$.
- In binary hypothesis testing with T IID observations and equal prior probabilities on f and g, the smallest attainable probability of error $P_e(T)$ decays at exponential rate $C(f, g)$: $-\frac{1}{T} \log P_e(T) \to C(f, g)$.

Merging of Opinions: The Blackwell–Dubins Theorem calls the Bhattacharyya coefficient the Hellinger affinity.

Kakutani's theorem there implies that when $\int \sqrt{fg} < 1$, the distributions of infinite IID sequences drawn from f and from g are mutually singular.

We will see an example of the third point in the lecture *Likelihood Ratio Processes*, where we study the Chernoff entropy in the context of model selection.

Chernoff entropy lies between the Bhattacharyya distance $-\log \int \sqrt{fg}$ and a combination of the two KL divergences:

$$-\log \int \sqrt{fg}\, dx \;\leq\; C(f, g) \;\leq\; \frac{D_{KL}(f\|g)\, D_{KL}(g\|f)}{D_{KL}(f\|g) + D_{KL}(g\|f)} \;\leq\; \min\{D_{KL}(f\|g),\, D_{KL}(g\|f)\}. \tag{21.9}$$

When f and g are positive on the same set, $\log \rho(\phi)$ is a convex function of ϕ that equals 0 at $\phi = 0$ and $\phi = 1$.

Its slope is $-D_{KL}(g\|f)$ at $\phi = 0$ and $D_{KL}(f\|g)$ at $\phi = 1$.

The lower bound in (21.9) evaluates $\log \rho$ at $\phi = 1/2$; the upper bound uses the two tangent lines at the endpoints, which lie below $\log \rho$.

Exercise 21.10.2 asks you to prove these statements.

Let's compute the Chernoff entropy between our example distributions f and g and check the bounds (21.9).

```
def chernoff_integrand(ϕ, f, g):
    """Integral entering Chernoff entropy for a given ϕ."""
    def integrand(w):
        return f(w)**ϕ * g(w)**(1-ϕ)
    return integrate(integrand)

def compute_chernoff_entropy(f, g):
    """Compute Chernoff entropy C(f,g)."""
    def objective(ϕ):
        return chernoff_integrand(ϕ, f, g)
    result = minimize_scalar(objective, bounds=(1e-5, 1-1e-5), method='bounded')
    min_value = result.fun
    ϕ_optimal = result.x
    chernoff_entropy = -np.log(min_value)
    return chernoff_entropy, ϕ_optimal

C_fg, ϕ_optimal = compute_chernoff_entropy(f, g)
print(f"Chernoff entropy C(f,g) = {C_fg:.4f}")
print(f"Optimal ϕ = {ϕ_optimal:.4f}")

bhattacharyya = -np.log(chernoff_integrand(0.5, f, g))
kl_bound = kl_fg * kl_gf / (kl_fg + kl_gf)
print(f"{bhattacharyya:.4f} ≤ {C_fg:.4f} ≤ {kl_bound:.4f} "
      f"≤ {min(kl_fg, kl_gf):.4f}")
```

```
Chernoff entropy C(f,g) = 0.1212
Optimal ϕ = 0.5969
0.1169 ≤ 0.1212 ≤ 0.2365 ≤ 0.3436
```

21.7 Comparing divergence measures

We now compare these measures across several pairs of Beta distributions.

For Beta densities, the KL divergence and the Chernoff coefficient have closed forms.

Let $B(a,b)$ be the Beta function and $\psi = \Gamma'/\Gamma$ the digamma function.

Then

$$D_{KL}(\text{Beta}(a_1,b_1) \,\|\, \text{Beta}(a_2,b_2)) = \log\frac{B(a_2,b_2)}{B(a_1,b_1)} + (a_1-a_2)\psi(a_1) + (b_1-b_2)\psi(b_1) + (a_2-a_1+b_2-b_1)\psi(a_1+b_1)$$

and

$$\rho(\phi) = \frac{B(\phi a_1 + (1-\phi)a_2,\ \phi b_1 + (1-\phi)b_2)}{B(a_1,b_1)^{\phi}\, B(a_2,b_2)^{1-\phi}}.$$

Exercise 21.10.1 asks you to derive these formulas.

JS divergence has no closed form, so we compute it by numerical integration over $(0,1)$.

The integration must cover the whole interval: $\text{Beta}(0.1,0.2)$ puts a quarter of its probability mass outside $[10^{-5}, 1-10^{-5}]$.

```
def KL_beta(a1, b1, a2, b2):
    """KL(Beta(a1, b1) ‖ Beta(a2, b2)) in closed form."""
    return (betaln(a2, b2) - betaln(a1, b1)
```

(continues on next page)

(continued from previous page)

```
            + (a1 - a2) * digamma(a1) + (b1 - b2) * digamma(b1)
            + (a2 - a1 + b2 - b1) * digamma(a1 + b1))

def log_chernoff_coef_beta(ϕ, a1, b1, a2, b2):
    """log ρ(ϕ) for Beta(a1, b1) and Beta(a2, b2) in closed form."""
    return (betaln(ϕ * a1 + (1 - ϕ) * a2, ϕ * b1 + (1 - ϕ) * b2)
            - ϕ * betaln(a1, b1) - (1 - ϕ) * betaln(a2, b2))

def chernoff_beta(a1, b1, a2, b2):
    """Chernoff entropy and minimizing ϕ for two Beta densities."""
    res = minimize_scalar(log_chernoff_coef_beta, bounds=(0, 1),
                          args=(a1, b1, a2, b2), method='bounded')
    return -res.fun, res.x

mass_out = beta_dist.cdf(1e-5, 0.1, 0.2) + beta_dist.sf(1 - 1e-5, 0.1, 0.2)
print(f"Mass of Beta(0.1, 0.2) outside [1e-5, 1-1e-5] = {mass_out:.3f}")
```

```
Mass of Beta(0.1, 0.2) outside [1e-5, 1-1e-5] = 0.251
```

Pair (f, g)	KL(f, g)	KL(g, f)	JS	C
$\text{Beta}(1,1), \text{Beta}(1.1, 1.05)$	0.0028	0.0026	0.0007	0.0007
$\text{Beta}(1,1), \text{Beta}(1.2, 1.1)$	0.0106	0.0092	0.0024	0.0025
$\text{Beta}(1,1), \text{Beta}(0.9, 0.8)$	0.0142	0.0168	0.0038	0.0039
$\text{Beta}(1,1), \text{Beta}(1.5, 1.2)$	0.0590	0.0437	0.0121	0.0126
$\text{Beta}(1,1), \text{Beta}(0.7, 0.6)$	0.0673	0.0965	0.0189	0.0199
$\text{Beta}(1,1), \text{Beta}(2, 1.5)$	0.1782	0.1081	0.0309	0.0339
$\text{Beta}(1,1), \text{Beta}(0.5, 0.5)$	0.1447	0.2416	0.0413	0.0457
$\text{Beta}(1,1), \text{Beta}(2.5, 1.8)$	0.3326	0.1731	0.0502	0.0577
$\text{Beta}(1,1), \text{Beta}(0.3, 0.4)$	0.3316	0.7710	0.0957	0.1189
$\text{Beta}(1,1), \text{Beta}(3, 1.2)$	0.7592	0.3436	0.0984	0.1212
$\text{Beta}(1,1), \text{Beta}(0.3, 0.3)$	0.3934	0.9533	0.1129	0.1437
$\text{Beta}(1,1), \text{Beta}(4, 1)$	1.6137	0.6363	0.1733	0.2341
$\text{Beta}(1,1), \text{Beta}(5, 1)$	2.3906	0.8094	0.2162	0.3128
$\text{Beta}(1,1), \text{Beta}(0.1, 0.2)$	0.9810	4.9773	0.2652	0.4431

JS divergence and Chernoff entropy rank the fourteen pairs in the same order.

The two KL divergences rank them differently.

For $g = \text{Beta}(0.1, 0.2)$, $KL(g, f)$ is by far the largest in the table, because g puts most of its mass near the endpoints, where g greatly exceeds $f = 1$.

For $g = \text{Beta}(4, 1)$ and $\text{Beta}(5, 1)$, $KL(f, g)$ exceeds its value for $\text{Beta}(0.1, 0.2)$, because these g's vanish at 0, where $f = 1$.

Next we visualize relationships among KL, JS, and Chernoff entropy.

```
kl_fg_values = [result['KL(f, g)'] for result in results]
js_values = [result['JS'] for result in results]
chernoff_values = [result['C'] for result in results]

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

axes[0].scatter(kl_fg_values, js_values, alpha=0.7, s=60)
```

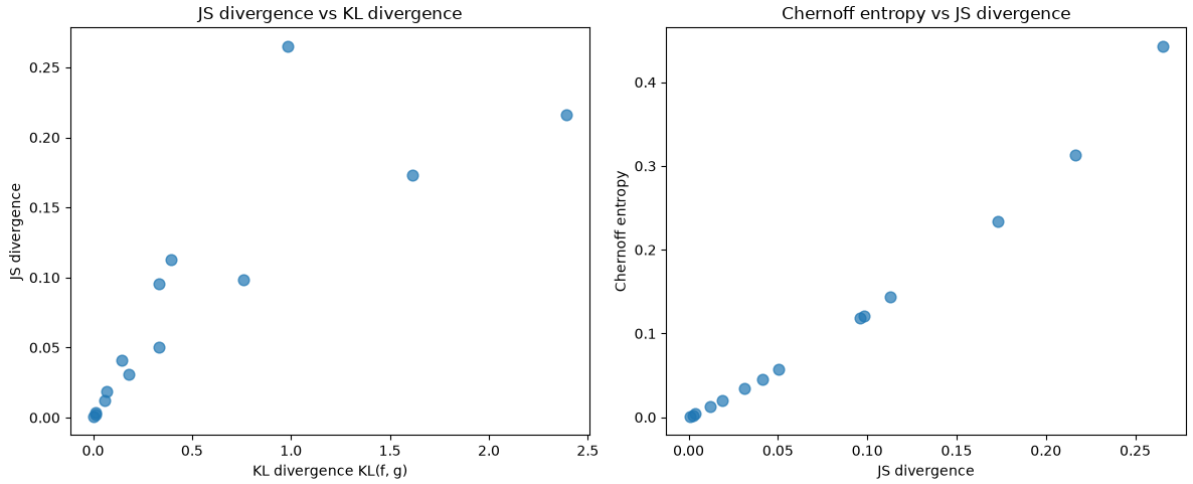
(continues on next page)

(continued from previous page)

```
axes[0].set_xlabel('KL divergence KL(f, g)')
axes[0].set_ylabel('JS divergence')
axes[0].set_title('JS divergence vs KL divergence')

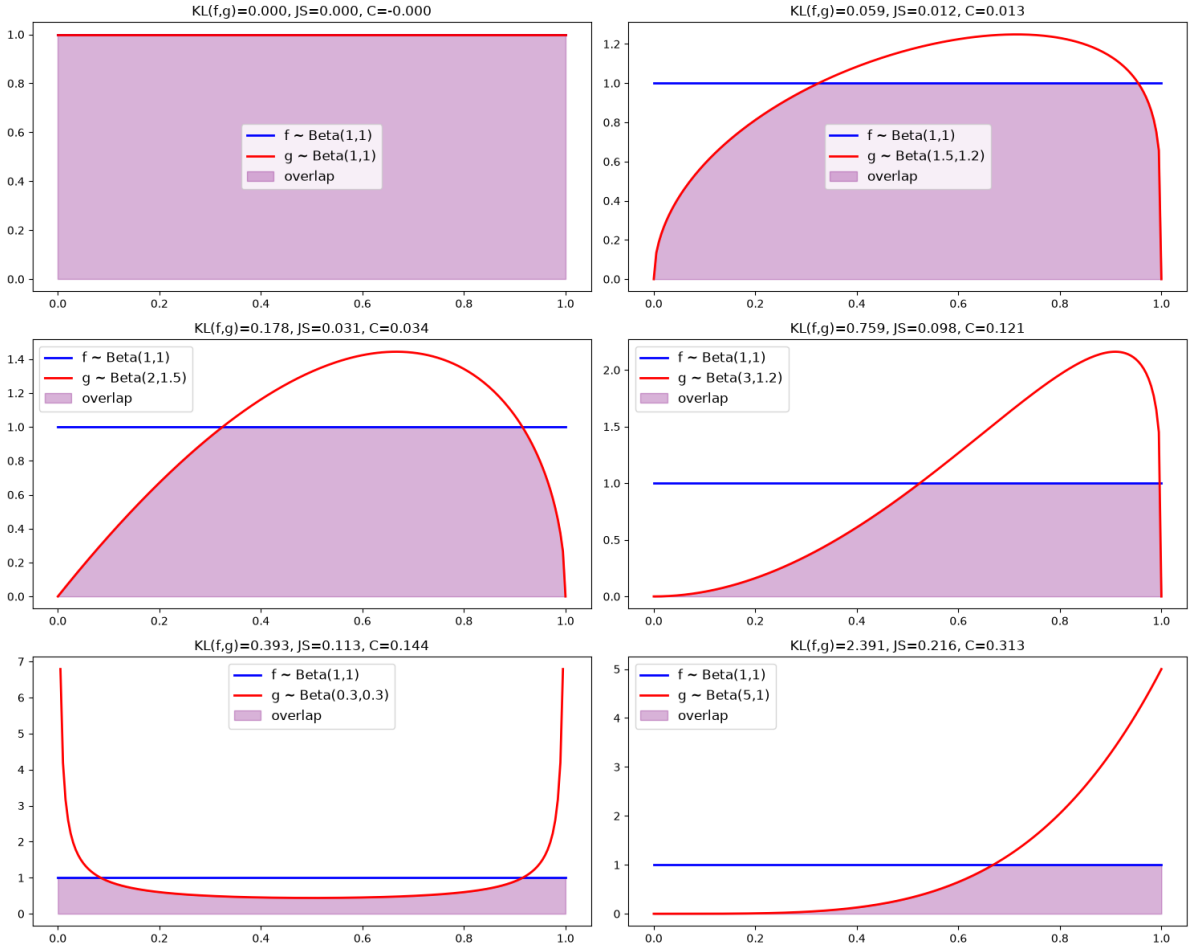
axes[1].scatter(js_values, chernoff_values, alpha=0.7, s=60)
axes[1].set_xlabel('JS divergence')
axes[1].set_ylabel('Chernoff entropy')
axes[1].set_title('Chernoff entropy vs JS divergence')

plt.tight_layout()
plt.show()
```



We now generate plots illustrating how overlap visually diminishes as divergence measures increase.

```
param_grid = [
    ((1, 1), (1, 1)),
    ((1, 1), (1.5, 1.2)),
    ((1, 1), (2, 1.5)),
    ((1, 1), (3, 1.2)),
    ((1, 1), (0.3, 0.3)),
    ((1, 1), (5, 1))
]
```



21.8 KL divergence and maximum-likelihood estimation

Let $x_1, x_2, \ldots, x_N$ be IID draws from a density p_0.

A statistician fits a parametric family of densities $\{p_\theta\}$ by maximum likelihood.

The **empirical distribution** p_e puts probability $1/N$ on each observation:

$$p_e(x) = \frac{1}{N} \sum_{i=1}^{N} \delta(x - x_i),$$

where $\delta(x - x_i)$ is the Dirac delta function centered at x_i; it satisfies $\delta(x - x_i) = 0$ for $x \neq x_i$ and $\int \delta(x - x_i)\, dx = 1$.

The log-likelihood function for parameter θ is

$$\ell_N(\theta) = \sum_{i=1}^{N} \log p_\theta(x_i).$$

The average log likelihood is an expectation under p_e:

$$\frac{1}{N} \ell_N(\theta) = \int p_e(x) \log p_\theta(x)\, dx = -H(p_e, p_\theta),$$

where $H(p_e, p_\theta)$ is the cross-entropy defined in the primer above.

Maximum likelihood chooses θ to minimize the cross-entropy $H(p_e, p_\theta)$.

The KL divergence $D_{KL}(p_e \| p_\theta)$ is of no use here.

When p_θ is a density, p_e puts probability one on the finite set $\{x_1, \ldots, x_N\}$, to which p_θ assigns probability zero.

So p_e is not absolutely continuous with respect to p_θ, and $D_{KL}(p_e \| p_\theta) = +\infty$ for every θ.

KL divergence enters through the population.

By the strong law of large numbers, for each θ

$$\frac{1}{N}\ell_N(\theta) \to E_{p_0} \log p_\theta(X) = -H(p_0) - D_{KL}(p_0 \| p_\theta) \quad \text{almost surely,}$$

where the equality is (21.3) for densities, and we assume that the entropy $H(p_0)$ is finite.

The entropy $H(p_0)$ does not depend on θ.

Maximum likelihood is the sample analogue of minimizing $D_{KL}(p_0 \| p_\theta)$ over θ.

If $p_0 = p_{\theta_0}$ for some θ_0, Gibbs' inequality implies that θ_0 minimizes $D_{KL}(p_0 \| p_\theta)$.

If no member of the family equals p_0, then under regularity conditions the maximum likelihood estimator converges to the **pseudo-true value**

$$\theta^* = \arg\min_\theta D_{KL}(p_0 \| p_\theta),$$

the member of the family closest to p_0 in KL divergence [White, 1982].

Differentiating under the integral sign gives

$$\nabla_\theta D_{KL}(p_0 \| p_\theta) = -E_{p_0} \nabla_\theta \log p_\theta(X).$$

The expected score is the negative gradient of the KL divergence, and it vanishes at θ^*.

Which Moments to Match? builds an estimator on this fact.

Maximum Likelihood Estimation describes how to compute maximum likelihood estimates.

Exercise 21.10.3 fits a misspecified Beta model to data drawn from a mixture.

21.9 Related lectures

This lecture has introduced tools that we'll encounter elsewhere.

- Other quantecon lectures that apply connections between divergence measures and statistical inference include *Likelihood Ratio Processes*, *A Problem that Stumped Milton Friedman*, and *Incorrect Models*.
- Statistical divergence functions also take center stage in *Heterogeneous Beliefs and Financial Markets* that studies Lawrence Blume and David Easley's model of heterogeneous beliefs and financial markets.
- *Mean of a Likelihood Ratio Process* uses the χ^2 divergence $\chi^2(f\|q) = \int (f-q)^2/q\, dx = \int f^2/q\, dx - 1$, which equals the variance under q of the importance weight f/q.
- *Merging of Opinions: The Blackwell–Dubins Theorem* uses the Bhattacharyya coefficient, under the name Hellinger affinity, to decide when two Bayesians' opinions merge.

21.10 Exercises

Exercise 21.10.1

Let $X \sim \text{Beta}(a, b)$, with density $x^{a-1}(1-x)^{b-1}/B(a,b)$ on $(0,1)$.

1. Differentiate $B(a,b) = \int_0^1 x^{a-1}(1-x)^{b-1}dx$ with respect to a and b to show that $E\log X = \psi(a) - \psi(a+b)$ and $E\log(1-X) = \psi(b) - \psi(a+b)$.
2. Use part 1 to derive the closed form for $D_{KL}(\text{Beta}(a_1,b_1) \,\|\, \text{Beta}(a_2,b_2))$ stated above.
3. Show that when f and g are Beta densities, $f^\phi g^{1-\phi}$ is proportional to a Beta density, and use this fact to derive the closed form for $\rho(\phi)$.
4. Let $f = \text{Beta}(1,1)$ and $g = \text{Beta}(0.1, 0.2)$. Compute $KL(g,f)$, $JS(f,g)$, and $C(f,g)$ by applying `quad` to the truncated interval $[10^{-5}, 1-10^{-5}]$. Compare the results with the values in the table above and explain the differences.

Solution

Part 1. Differentiating under the integral sign gives

$$\frac{\partial B(a,b)}{\partial a} = \int_0^1 \log x \; x^{a-1}(1-x)^{b-1}dx = B(a,b)\, E\log X.$$

Since $B(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$,

$$E\log X = \frac{\partial \log B(a,b)}{\partial a} = \psi(a) - \psi(a+b).$$

The same argument with b in place of a gives $E\log(1-X) = \psi(b) - \psi(a+b)$.

Part 2. Let $f = \text{Beta}(a_1,b_1)$ and $g = \text{Beta}(a_2,b_2)$.

Then

$$\log f(x) - \log g(x) = (a_1 - a_2)\log x + (b_1 - b_2)\log(1-x) + \log B(a_2,b_2) - \log B(a_1,b_1).$$

Taking expectations under f and using part 1 gives

$$D_{KL}(f\|g) = \log\frac{B(a_2,b_2)}{B(a_1,b_1)} + (a_1-a_2)[\psi(a_1) - \psi(a_1+b_1)] + (b_1-b_2)[\psi(b_1) - \psi(a_1+b_1)],$$

which rearranges to the formula in the text.

Part 3. Let $a_\phi = \phi a_1 + (1-\phi)a_2$ and $b_\phi = \phi b_1 + (1-\phi)b_2$.

Then

$$f(x)^\phi g(x)^{1-\phi} = \frac{x^{a_\phi - 1}(1-x)^{b_\phi - 1}}{B(a_1,b_1)^\phi B(a_2,b_2)^{1-\phi}},$$

which is proportional to the $\text{Beta}(a_\phi, b_\phi)$ density.

Integrating over $(0,1)$ gives

$$\rho(\phi) = \frac{B(a_\phi, b_\phi)}{B(a_1,b_1)^\phi B(a_2,b_2)^{1-\phi}}.$$

Part 4. Here is code that repeats the computations on the truncated interval.

```
def KL_truncated(f, g):
    """KL(f, g) by quad over [1e-5, 1-1e-5]."""
    integrand = lambda w: f(w) * np.log(f(w) / g(w))
    return quad(integrand, 1e-5, 1 - 1e-5)[0]

def chernoff_truncated(f, g):
    """Chernoff entropy by quad over [1e-5, 1-1e-5]."""
    ρ = lambda ϕ: quad(lambda w: f(w)**ϕ * g(w)**(1-ϕ),
                       1e-5, 1 - 1e-5)[0]
    res = minimize_scalar(ρ, bounds=(0, 1), method='bounded')
    return -np.log(res.fun)

f_ex = jit(lambda x: p(x, 1, 1))
g_ex = jit(lambda x: p(x, 0.1, 0.2))
m_ex = lambda x: 0.5 * (f_ex(x) + g_ex(x))

truncated = [KL_truncated(g_ex, f_ex),
             0.5 * KL_truncated(f_ex, m_ex) + 0.5 * KL_truncated(g_ex, m_ex),
             chernoff_truncated(f_ex, g_ex)]
accurate = [KL_beta(0.1, 0.2, 1, 1),
            compute_JS(f_ex, g_ex),
            chernoff_beta(1, 1, 0.1, 0.2)[0]]

for name, t, a in zip(['KL(g, f)', 'JS(f, g)', 'C(f, g)'], truncated, accurate):
    print(f"{name}: truncated = {t:.4f}, accurate = {a:.4f}")
print(f"Mass of g outside the truncated interval = {mass_out:.4f}")
```

```
KL(g, f): truncated = 1.0036, accurate = 4.9773
JS(f, g): truncated = 0.1783, accurate = 0.2652
C(f, g): truncated = 0.4556, accurate = 0.4431
Mass of g outside the truncated interval = 0.2509
```

The density $g = \text{Beta}(0.1, 0.2)$ puts a quarter of its mass outside $[10^{-5}, 1 - 10^{-5}]$, where g is enormous and $f = 1$.

The omitted part of $KL(g, f) = \int g \log g$ is large, because $\log g$ is large there.

The omitted part of $JS(f, g)$ is close to $\frac{1}{2} \times 0.25 \times \log 2 \approx 0.087$, because $g \log(2g/(f+g)) \approx g \log 2$ where $g \gg f$.

Truncation also omits part of $\rho(\phi)$, which lowers ρ and raises the computed Chernoff entropy.

Exercise 21.10.2

Let f and g be densities that are positive on the same set, and let $\rho(\phi) = \int f^\phi g^{1-\phi} dx$.

1. Use Hölder's inequality to show that $\log \rho(\phi)$ is convex on $[0, 1]$.
2. Show that $\log \rho(0) = \log \rho(1) = 0$, that the derivative of $\log \rho$ at $\phi = 0$ is $-D_{KL}(g \| f)$, and that its derivative at $\phi = 1$ is $D_{KL}(f \| g)$.
3. Deduce the bounds (21.9).
4. Verify the bounds numerically for every pair in the table above, and plot $-\log \rho(\phi)$ together with the two tangent lines for our example f and g.

Solution

Part 1. Take $\phi_1, \phi_2 \in [0,1]$, $\lambda \in (0,1)$, and $\phi = \lambda\phi_1 + (1-\lambda)\phi_2$.

Write

$$f^{\phi} g^{1-\phi} = \left(f^{\phi_1} g^{1-\phi_1}\right)^{\lambda} \left(f^{\phi_2} g^{1-\phi_2}\right)^{1-\lambda}.$$

Hölder's inequality with exponents $1/\lambda$ and $1/(1-\lambda)$ gives

$$\rho(\phi) \leq \rho(\phi_1)^{\lambda} \rho(\phi_2)^{1-\lambda}.$$

Taking logarithms gives $\log \rho(\phi) \leq \lambda \log \rho(\phi_1) + (1-\lambda) \log \rho(\phi_2)$.

Part 2. Because f and g are positive on the same set, $\rho(0) = \int g = 1$ and $\rho(1) = \int f = 1$.

Differentiating under the integral sign gives

$$\frac{d \log \rho(\phi)}{d\phi} = \frac{1}{\rho(\phi)} \int f^{\phi} g^{1-\phi} \log \frac{f}{g} \, dx.$$

At $\phi = 0$ the derivative is $\int g \log(f/g) = -D_{KL}(g\|f)$.

At $\phi = 1$ the derivative is $\int f \log(f/g) = D_{KL}(f\|g)$.

Part 3. The lower bound holds because $C(f,g) = -\min_{\phi} \log \rho(\phi) \geq -\log \rho(1/2)$.

A convex function lies above its tangent lines, so

$$\log \rho(\phi) \geq -\phi \, D_{KL}(g\|f) \quad \text{and} \quad \log \rho(\phi) \geq -(1-\phi) \, D_{KL}(f\|g).$$

Hence

$$-\log \rho(\phi) \leq \min\{\phi \, D_{KL}(g\|f), \; (1-\phi) \, D_{KL}(f\|g)\}.$$

The right side is largest where the two lines cross, at $\phi = D_{KL}(f\|g)/(D_{KL}(f\|g) + D_{KL}(g\|f))$, where it equals $D_{KL}(f\|g) \, D_{KL}(g\|f)/(D_{KL}(f\|g) + D_{KL}(g\|f))$.

Maximizing the left side over ϕ gives the middle inequality.

The last inequality holds because $xy/(x+y) \leq \min\{x, y\}$ for positive x and y.

Part 4. We check the bounds for every pair in the table.

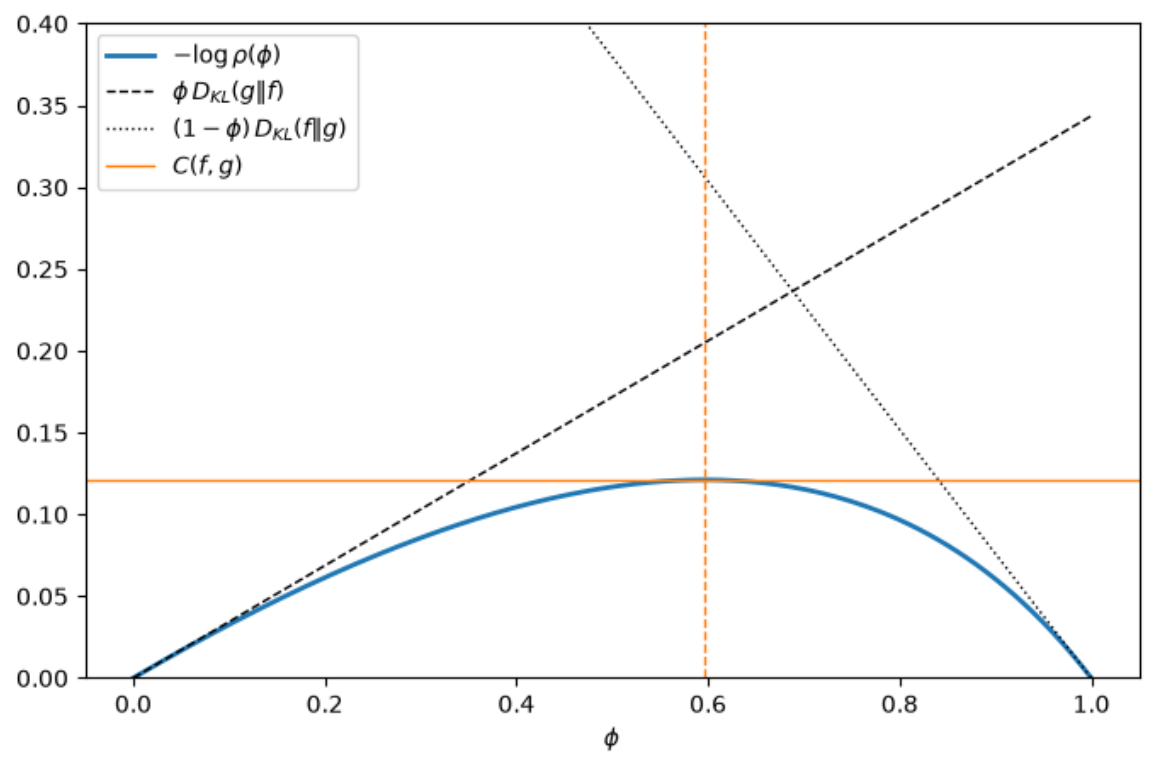
```
for (f_a, f_b), (g_a, g_b) in distribution_pairs:
    K_fg = KL_beta(f_a, f_b, g_a, g_b)
    K_gf = KL_beta(g_a, g_b, f_a, f_b)
    lower = -log_chernoff_coef_beta(0.5, f_a, f_b, g_a, g_b)
    C, _ = chernoff_beta(f_a, f_b, g_a, g_b)
    upper = K_fg * K_gf / (K_fg + K_gf)
    holds = lower <= C <= upper <= min(K_fg, K_gf)
    print(f"Beta({g_a},{g_b}): {lower:.4f} ≤ {C:.4f} ≤ {upper:.4f} "
          f"≤ {min(K_fg, K_gf):.4f}  {holds}")
```

```
Beta(0.1,0.2): 0.3929 ≤ 0.4431 ≤ 0.8195 ≤ 0.9810  True
Beta(0.3,0.3): 0.1374 ≤ 0.1437 ≤ 0.2785 ≤ 0.3934  True
Beta(0.3,0.4): 0.1141 ≤ 0.1189 ≤ 0.2319 ≤ 0.3316  True
Beta(0.5,0.5): 0.0450 ≤ 0.0457 ≤ 0.0905 ≤ 0.1447  True
Beta(0.7,0.6): 0.0198 ≤ 0.0199 ≤ 0.0396 ≤ 0.0673  True
Beta(0.9,0.8): 0.0039 ≤ 0.0039 ≤ 0.0077 ≤ 0.0142  True
Beta(1.1,1.05): 0.0007 ≤ 0.0007 ≤ 0.0014 ≤ 0.0026  True
Beta(1.2,1.1): 0.0025 ≤ 0.0025 ≤ 0.0049 ≤ 0.0092  True
Beta(1.5,1.2): 0.0125 ≤ 0.0126 ≤ 0.0251 ≤ 0.0437  True
Beta(2,1.5): 0.0334 ≤ 0.0339 ≤ 0.0673 ≤ 0.1081  True
Beta(2.5,1.8): 0.0562 ≤ 0.0577 ≤ 0.1138 ≤ 0.1731  True
Beta(3,1.2): 0.1169 ≤ 0.1212 ≤ 0.2365 ≤ 0.3436  True
Beta(4,1): 0.2231 ≤ 0.2341 ≤ 0.4564 ≤ 0.6363  True
Beta(5,1): 0.2939 ≤ 0.3128 ≤ 0.6047 ≤ 0.8094  True
```

The next figure plots $-\log \rho(\phi)$ and the two tangent lines for $f = \text{Beta}(1, 1)$ and $g = \text{Beta}(3, 1.2)$.

```
ϕ_grid = np.linspace(0, 1, 200)
K_fg = KL_beta(F_a, F_b, G_a, G_b)
K_gf = KL_beta(G_a, G_b, F_a, F_b)
C_ex, ϕ_star = chernoff_beta(F_a, F_b, G_a, G_b)

fig, ax = plt.subplots(figsize=(8, 5))
ax.plot(ϕ_grid, -log_chernoff_coef_beta(ϕ_grid, F_a, F_b, G_a, G_b),
        lw=2, label=r'$-\log \rho(\phi)$')
ax.plot(ϕ_grid, ϕ_grid * K_gf, 'k--', lw=1,
        label=r'$\phi \, D_{KL}(g\|f)$')
ax.plot(ϕ_grid, (1 - ϕ_grid) * K_fg, 'k:', lw=1,
        label=r'$(1-\phi) \, D_{KL}(f\|g)$')
ax.axhline(C_ex, color='C1', lw=1, label=r'$C(f,g)$')
ax.axvline(ϕ_star, color='C1', lw=1, ls='--')
ax.set_ylim(0, 0.4)
ax.set_xlabel(r'$\phi$')
ax.legend()
plt.show()
```



Exercise 21.10.3

Let the data $x_1, \ldots, x_N$ be IID draws from the mixture

$$h(x) = \frac{1}{2} f(x) + \frac{1}{2} g(x)$$

of our example densities $f = \text{Beta}(1, 1)$ and $g = \text{Beta}(3, 1.2)$.

A statistician fits the Beta(a, b) family by maximum likelihood.

The family excludes h.

1. Show that the pseudo-true value $(a^*, b^*) = \arg\min_{a,b} D_{KL}(h \,\|\, \text{Beta}(a, b))$ solves $\psi(a) - \psi(a+b) = E_h \log X$ and $\psi(b) - \psi(a+b) = E_h \log(1 - X)$.
2. Show that the maximum likelihood estimator $(\hat{a}_N, \hat{b}_N)$ solves the same equations with the expectations replaced by the sample means of $\log x_i$ and $\log(1 - x_i)$.
3. Compute (a^*, b^*) and plot h against the Beta(a^*, b^*) density.
4. Simulate samples of size $N = 10^2, 10^3, \ldots, 10^6$ and verify that $(\hat{a}_N, \hat{b}_N)$ approaches (a^*, b^*).

Solution

Parts 1 and 2. The log density of Beta(a, b) is

$$\log p_{a,b}(x) = (a-1)\log x + (b-1)\log(1-x) - \log B(a, b).$$

Hence

$$D_{KL}(h \| p_{a,b}) = E_h \log h(X) + Q(a, b; s_1, s_2), \qquad Q(a, b; s_1, s_2) = \log B(a, b) - (a-1)s_1 - (b-1)s_2,$$

with $s_1 = E_h \log X$ and $s_2 = E_h \log(1 - X)$.

The average negative log likelihood is $-\ell_N(a, b)/N = Q(a, b; \hat{s}_1, \hat{s}_2)$, where $\hat{s}_1$ and $\hat{s}_2$ are the sample means of $\log x_i$ and $\log(1 - x_i)$.

The function $\log B(a, b)$ is convex, so Q is convex in (a, b).

By part 1 of Exercise 21.10.1, $\partial \log B(a, b)/\partial a = \psi(a) - \psi(a + b)$ and $\partial \log B(a, b)/\partial b = \psi(b) - \psi(a + b)$.

Setting the gradient of Q to zero gives the stated equations, with (s_1, s_2) for the pseudo-true value and $(\hat{s}_1, \hat{s}_2)$ for the maximum likelihood estimator.

Part 3. Part 1 of Exercise 21.10.1 and linearity of expectations give s_1 and s_2 in closed form.

We minimize Q over $(\log a, \log b)$ to keep a and b positive.

```
from scipy.optimize import minimize

def Q(z, s1, s2):
    a, b = np.exp(z)
    return betaln(a, b) - (a - 1) * s1 - (b - 1) * s2

def fit_beta(s1, s2):
    res = minimize(Q, x0=np.zeros(2), args=(s1, s2), method='BFGS')
    return np.exp(res.x)

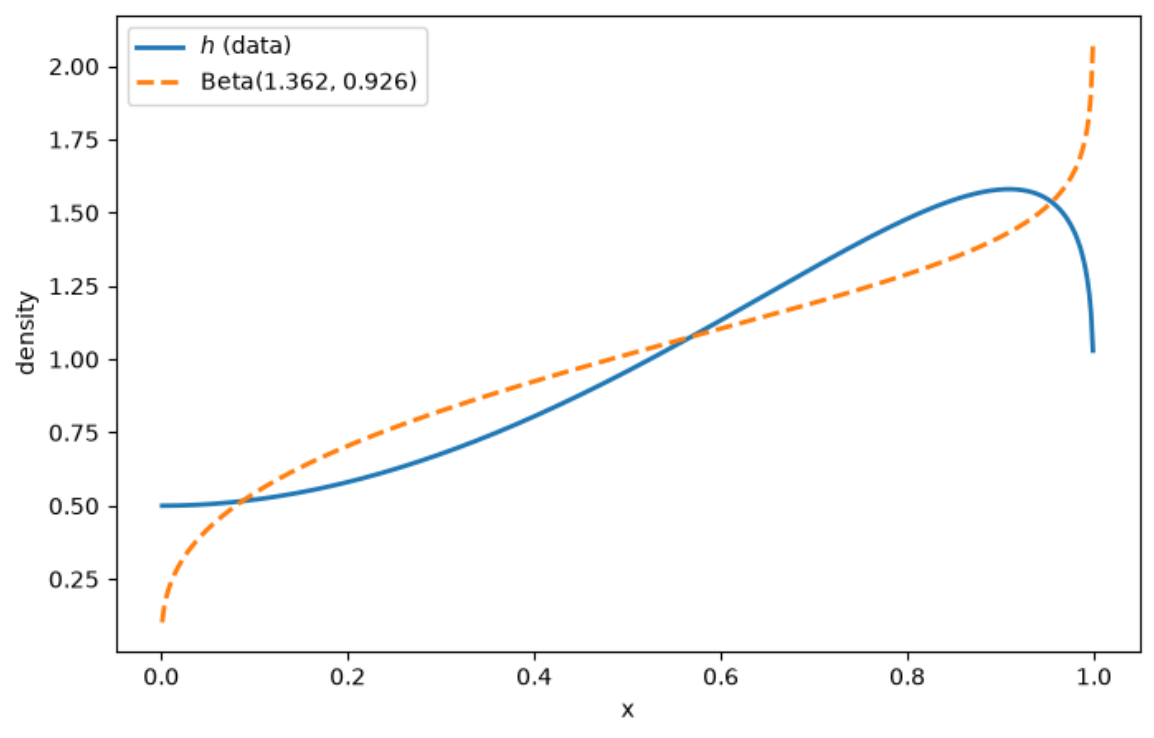
E_log_x = lambda a, b: digamma(a) - digamma(a + b)
s1 = 0.5 * E_log_x(F_a, F_b) + 0.5 * E_log_x(G_a, G_b)
s2 = 0.5 * E_log_x(F_b, F_a) + 0.5 * E_log_x(G_b, G_a)

a_star, b_star = fit_beta(s1, s2)
print(f"pseudo-true value: a* = {a_star:.4f}, b* = {b_star:.4f}")

h = lambda x: 0.5 * p(x, F_a, F_b) + 0.5 * p(x, G_a, G_b)
p_star = lambda x: p(x, a_star, b_star)
print(f"KL(h, Beta(a*, b*)) = {compute_KL(h, p_star):.4f}")

x_grid = np.linspace(0.001, 0.999, 500)
plt.figure(figsize=(8, 5))
plt.plot(x_grid, h(x_grid), lw=2, label='$h$ (data)')
plt.plot(x_grid, p_star(x_grid), lw=2, ls='--',
         label=f'Beta({a_star:.3f}, {b_star:.3f})')
plt.xlabel('x')
plt.ylabel('density')
```

```
pseudo-true value: a* = 1.3621, b* = 0.9255
KL(h, Beta(a*, b*)) = 0.0124
```



Part 4. We draw samples from h and compute $(\hat{a}_N, \hat{b}_N)$ from the sample moments.

```
rng = np.random.default_rng(1234)

for N in [10**2, 10**3, 10**4, 10**5, 10**6]:
    from_f = rng.random(N) < 0.5
    x = np.where(from_f, rng.beta(F_a, F_b, N), rng.beta(G_a, G_b, N))
    a_hat, b_hat = fit_beta(np.mean(np.log(x)), np.mean(np.log1p(-x)))
    print(f"N = {N:>7}: a_hat = {a_hat:.4f}, b_hat = {b_hat:.4f}")
```

```
N =     100: a_hat = 1.2524, b_hat = 0.9221
N =    1000: a_hat = 1.5154, b_hat = 0.9324
N =   10000: a_hat = 1.3437, b_hat = 0.9186
N =  100000: a_hat = 1.3472, b_hat = 0.9216
N = 1000000: a_hat = 1.3601, b_hat = 0.9258
```

The estimates approach (a^*, b^*), the member of the Beta family closest to h in KL divergence.

Incorrect Models studies a Bayesian whose model contains only f and g and who sees data from h; KL divergences from h to f and to g determine where that Bayesian's beliefs settle.

Which Moments to Match? uses the pseudo-true value of an auxiliary model as the target of an estimator.

CHAPTER

TWENTYTWO

LIKELIHOOD RATIO PROCESSES

Contents

- *Likelihood Ratio Processes*
  - *Overview*
  - *Likelihood Ratio Process*
  - *Nature permanently draws from density g*
  - *Peculiar property*
  - *Nature permanently draws from density f*
  - *Likelihood ratio test*
  - *Hypothesis testing and classification*
  - *Markov chains*
  - *Related lectures*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install --upgrade quantecon
```

22.1 Overview

This lecture describes likelihood ratio processes and some of their uses.

We'll study the same setting that is also used in *this lecture on exchangeability*.

Among the things that we'll learn are

- How a likelihood ratio process is a key ingredient in frequentist hypothesis testing
- How a **receiver operating characteristic curve** summarizes information about a false alarm probability and power in frequentist hypothesis testing
- How a statistician can combine frequentist probabilities of type I and type II errors to form posterior probabilities of mistakes in a model selection or in an individual-classification problem

- How to use a Kullback-Leibler divergence to quantify the difference between two probability distributions with the same support
- How during World War II the United States Navy devised a decision rule for doing quality control on lots of ammunition, a topic that sets the stage for *this lecture*
- A peculiar property of likelihood ratio processes

Let's start by importing some Python tools.

```
import matplotlib.pyplot as plt
import numpy as np
from numba import vectorize, jit
from math import gamma
from scipy.integrate import quad
from scipy.optimize import brentq, minimize_scalar
from scipy.stats import beta as beta_dist
from scipy.special import betaln, digamma
import pandas as pd
from IPython.display import display, Math
import quantecon as qe

rng = np.random.default_rng()
```

22.2 Likelihood Ratio Process

A nonnegative random variable W has one of two probability density functions, either f or g.

Before the beginning of time, nature once and for all decides whether she will draw a sequence of IID draws from either f or g.

We let q be the density that nature chose once and for all, so that q is either f or g, permanently.

Nature knows which density it permanently draws from, but we the observers do not.

We know both f and g but we don't know which density nature chose.

But we want to know.

To do that, we use observations.

We observe a sequence $\{w_t\}_{t=1}^T$ of T IID draws that we know came from either f or g.

We want to use these observations to infer whether nature chose f or g.

A **likelihood ratio process** is a useful tool for this task.

To begin, we define a key component of a likelihood ratio process, namely, the time t likelihood ratio as the random variable

$$\ell(w_t) = \frac{f(w_t)}{g(w_t)}, \quad t \geq 1.$$

We assume that f and g both put positive probabilities on the same intervals of possible realizations of the random variable W.

That means that under the g density, $\ell(w_t) = \frac{f(w_t)}{g(w_t)}$ is a nonnegative random variable with mean 1.

A **likelihood ratio process** for sequence $\{w_t\}_{t=1}^\infty$ is defined as

$$L(w^t) = \prod_{i=1}^t \ell(w_i),$$

where $w^t = \{w_1, \ldots, w_t\}$ is a history of observations up to and including time t.

Sometimes for shorthand we'll write $L_t = L(w^t)$.

Notice that the likelihood process satisfies the *recursion*

$$L(w^t) = \ell(w_t) L(w^{t-1}).$$

The likelihood ratio and its logarithm are key tools for making inferences using a classic frequentist approach due to Neyman and Pearson [Neyman and Pearson, 1933].

To help us appreciate how things work, the following Python code evaluates f and g as two different Beta distributions, then computes and simulates an associated likelihood ratio process by generating a sequence w^t from one of the two probability distributions, for example, a sequence of IID draws from g.

```
# Parameters for the two Beta distributions
F_a, F_b = 1, 1
G_a, G_b = 3, 1.2

@vectorize
def p(x, a, b):
    """Beta distribution density function."""
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x** (a-1) * (1 - x) ** (b-1)

f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))

def create_beta_density(a, b):
    """Create a beta density function with specified parameters."""
    return jit(lambda x: p(x, a, b))

def likelihood_ratio(w, f_func, g_func):
    """Compute likelihood ratio for observation(s) w."""
    return f_func(w) / g_func(w)

@jit
def simulate_likelihood_ratios(a, b, f_func, g_func, rng, T=50, N=500):
    """
    Generate N sets of T observations of the likelihood ratio.
    """
    l_arr = np.empty((N, T))
    for i in range(N):
        for j in range(T):
            w = rng.beta(a, b)
            l_arr[i, j] = f_func(w) / g_func(w)
    return l_arr

def simulate_sequences(distribution, f_func, g_func, rng,
        F_params=(1, 1), G_params=(3, 1.2), T=50, N=500):
    """
    Generate N sequences of T observations from specified distribution.
    """
    if distribution == 'f':
        a, b = F_params
    elif distribution == 'g':
        a, b = G_params
    else:
        raise ValueError("distribution must be 'f' or 'g'")
```

(continues on next page)

(continued from previous page)

```
    l_arr = simulate_likelihood_ratios(a, b, f_func, g_func, rng, T, N)
    l_seq = np.cumprod(l_arr, axis=1)
    return l_arr, l_seq

def plot_likelihood_paths(l_seq, title="Likelihood ratio paths",
                          ylim=None, n_paths=None):
    """Plot likelihood ratio paths."""
    N, T = l_seq.shape
    n_show = n_paths or min(N, 100)

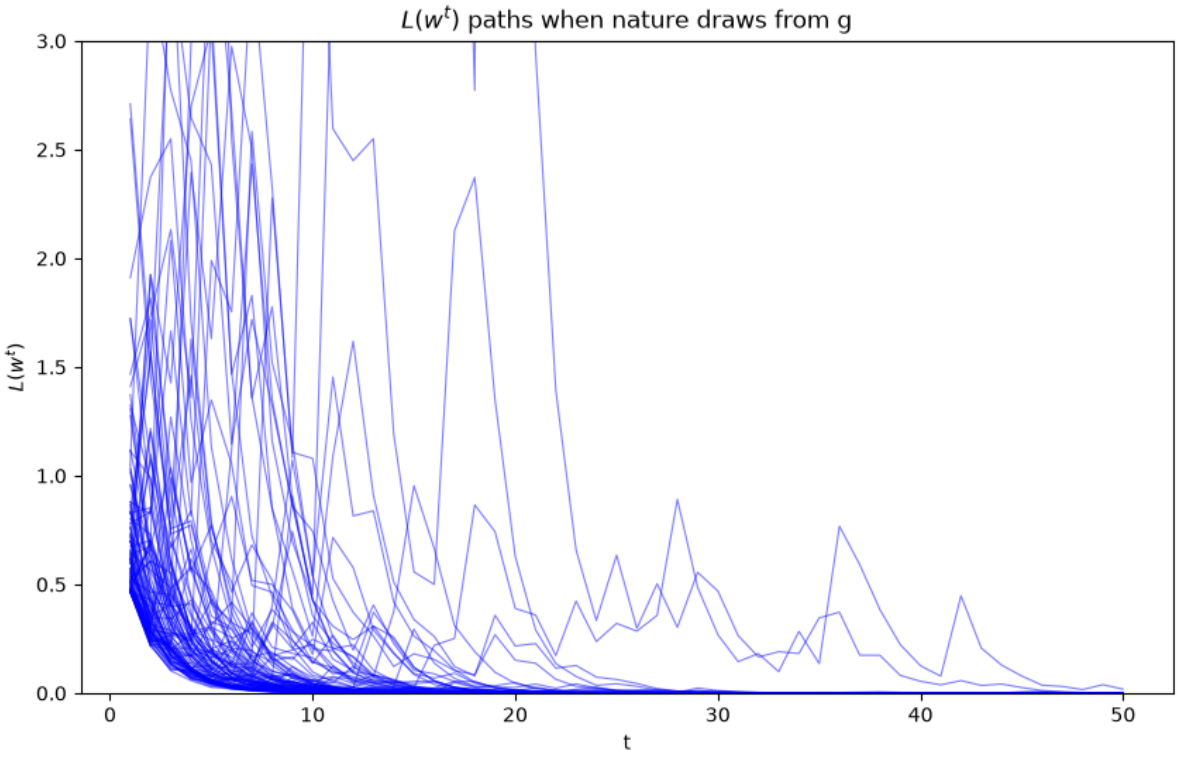
    plt.figure(figsize=(10, 6))
    for i in range(n_show):
        plt.plot(np.arange(1, T+1), l_seq[i, :], color='b', lw=0.8, alpha=0.5)

    if ylim:
        plt.ylim(ylim)
    plt.title(title)
    plt.xlabel('t')
    plt.ylabel('$L(w^t)$')
    plt.show()
```

22.3 Nature permanently draws from density g

We first simulate the likelihood ratio process when nature permanently draws from g.

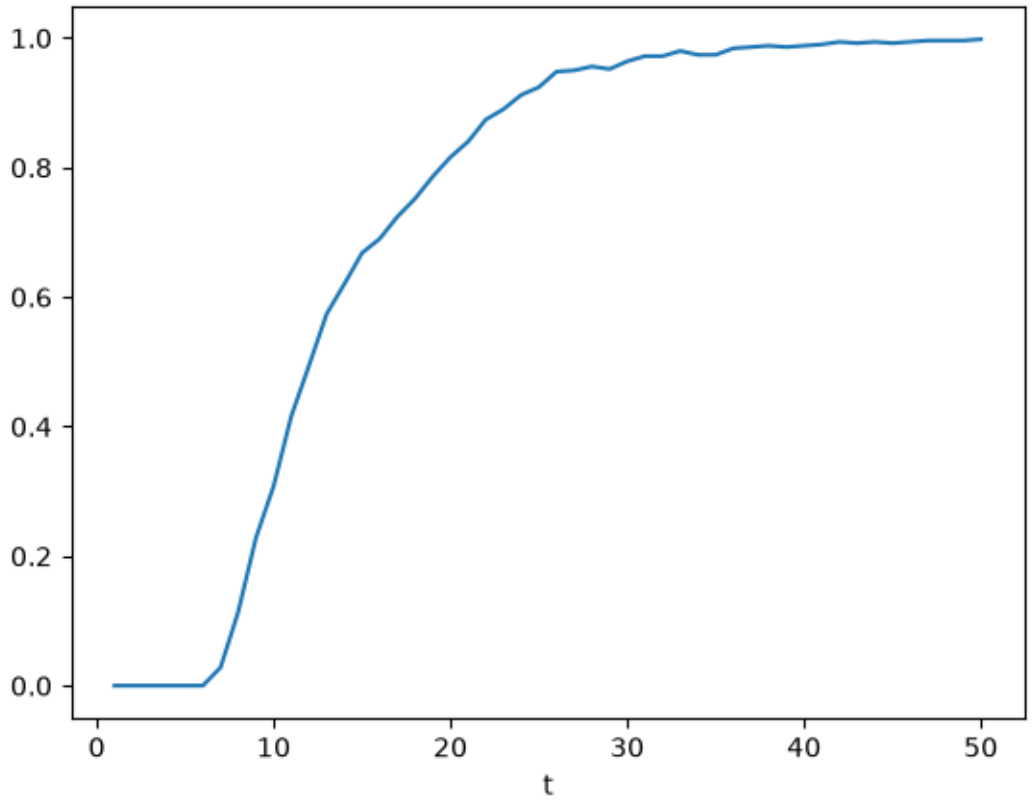
```
# Simulate when nature draws from g
l_arr_g, l_seq_g = simulate_sequences('g', f, g, rng, (F_a, F_b), (G_a, G_b))
plot_likelihood_paths(l_seq_g,
                      title="$L(w^{t})$ paths when nature draws from g",
                      ylim=[0, 3])
```



Evidently, as sample length T grows, most probability mass shifts toward zero.

To see this more clearly, we plot over time the fraction of paths $L\left(w^t\right)$ that fall in the interval $[0, 0.01]$.

```
N, T = l_arr_g.shape
plt.plot(np.arange(1, T+1), np.sum(l_seq_g <= 0.01, axis=0) / N)
plt.xlabel('t')
plt.show()
```



Despite the evident convergence of most probability mass to a very small interval near 0, the unconditional mean of $L(w^t)$ under probability density g is identically 1 for all t.

To verify this assertion, first notice that as mentioned earlier the unconditional mean $E[\ell(w_t) \mid q = g]$ is 1 for all t:

$$\begin{aligned} E[\ell(w_t) \mid q = g] &= \int \frac{f(w_t)}{g(w_t)} g(w_t)\, dw_t \\ &= \int f(w_t)\, dw_t \\ &= 1, \end{aligned}$$

which immediately implies

$$\begin{aligned} E[L(w^1) \mid q = g] &= E[\ell(w_1) \mid q = g] \\ &= 1. \end{aligned}$$

Because $L(w^t) = \ell(w_t) L(w^{t-1})$ and $\{w_t\}_{t=1}^{\infty}$ is an IID sequence, we have

$$\begin{aligned} E[L(w^t) \mid q = g] &= E[L(w^{t-1})\, \ell(w_t) \mid q = g] \\ &= E[L(w^{t-1})\, E[\ell(w_t) \mid q = g, w^{t-1}] \mid q = g] \\ &= E[L(w^{t-1})\, E[\ell(w_t) \mid q = g] \mid q = g] \\ &= E[L(w^{t-1}) \mid q = g] \end{aligned}$$

for any $t \geq 1$.

Mathematical induction implies $E[L(w^t) \mid q = g] = 1$ for all $t \geq 1$.

22.4 Peculiar property

How can $E\left[L\left(w^t\right) \mid q = g\right] = 1$ possibly be true when most probability mass of the likelihood ratio process is piling up near 0 as $t \to +\infty$?

The answer is that as $t \to +\infty$, the distribution of L_t becomes more and more fat-tailed: enough mass shifts to larger and larger values of L_t to make the mean of L_t continue to be one despite most of the probability mass piling up near 0.

Under g, the process $\{L_t\}$ is a nonnegative martingale with mean 1.

By the strong law of large numbers, $\frac{1}{t}\log L_t \to E_g \log \ell(w) = -D_{KL}(g\|f) < 0$ almost surely, so $L_t \to 0$ almost surely.

Thus $E[L_t \mid q = g] = 1$ for every t while $E[\lim_{t\to\infty} L_t \mid q = g] = 0$.

The martingale $\{L_t\}$ converges almost surely but not in mean, because it is not uniformly integrable.

Merging of Opinions: The Blackwell–Dubins Theorem studies this likelihood-ratio martingale.

Kakutani's theorem there states that when $\int \sqrt{fg}\, dw < 1$, the distributions of infinite sequences of IID draws from f and from g are mutually singular, which is another way of saying that $L_t \to 0$ almost surely under g.

To illustrate this peculiar property, we simulate many paths and calculate the unconditional mean of $L\left(w^t\right)$ by averaging across these many paths at each t.

```
l_arr_g, l_seq_g = simulate_sequences('g',
                f, g, rng, (F_a, F_b), (G_a, G_b), N=50000)
```

It would be useful to use simulations to verify that unconditional means $E\left[L\left(w^t\right)\right]$ equal unity by averaging across sample paths.

But it would be too computer-time-consuming for us to do that here simply by applying a standard Monte Carlo simulation approach.

The reason is that the distribution of $L\left(w^t\right)$ is extremely skewed for large values of t.

Because the probability density in the right tail is close to 0, it just takes too much computer time to sample enough points from the right tail.

We explain the problem in more detail in *this lecture*.

There we describe an alternative way to compute the mean of a likelihood ratio by computing the mean of a *different* random variable by sampling from a *different* probability distribution.

22.5 Nature permanently draws from density f

Now suppose that before time 0 nature permanently decided to draw repeatedly from density f.

While the mean of the likelihood ratio $\ell\left(w_t\right)$ under density g is 1, its mean under the density f exceeds one.

To see this, we compute

$$\begin{aligned} E\left[\ell\left(w_t\right) \mid q=f\right] &= \int \frac{f\left(w_t\right)}{g\left(w_t\right)} f\left(w_t\right) dw_t \\ &= \int \frac{f\left(w_t\right)}{g\left(w_t\right)} \frac{f\left(w_t\right)}{g\left(w_t\right)} g\left(w_t\right) dw_t \\ &= \int \ell\left(w_t\right)^2 g\left(w_t\right) dw_t \\ &= E\left[\ell\left(w_t\right)^2 \mid q=g\right] \\ &= E\left[\ell\left(w_t\right) \mid q=g\right]^2 + Var\left(\ell\left(w_t\right) \mid q=g\right) \\ &> E\left[\ell\left(w_t\right) \mid q=g\right]^2 = 1 \end{aligned}$$

For the f and g in our example, $E\left[\ell\left(w_t\right) \mid q=f\right]$ is infinite.

Because f is the uniform density, $E\left[\ell\left(w_t\right) \mid q=f\right] = \int_0^1 1/g(w)\, dw$.

As $w \to 0$, $g(w)$ vanishes like $w^{G_a - 1} = w^2$, so the integral diverges.

Exercise 22.10.4 shows that $E\left[\ell\left(w_t\right) \mid q=f\right] = B(G_a, G_b)B(2-G_a, 2-G_b)$ when $G_a < 2$ and $G_b < 2$, where B is the Beta function, and that it is infinite otherwise.

Because the w_t's are IID, $E\left[L\left(w^t\right) \mid q=f\right] = +\infty$ for every $t \geq 1$.

A sample average of $L\left(w^t\right)$ across simulated paths therefore does not settle down as the number of paths grows.

Mean of a Likelihood Ratio Process encounters the same divergent integral as the infinite variance of $\ell(w_t)$ under g, since $E\left[\ell\left(w_t\right)^2 \mid q=g\right] = E\left[\ell\left(w_t\right) \mid q=f\right]$.

The logarithm of $L\left(w^t\right)$ behaves well.

Under f, $E\left[\log L\left(w^t\right) \mid q=f\right] = t\, D_{KL}(f\|g)$, and by the strong law of large numbers $\frac{1}{t}\log L\left(w^t\right) \to D_{KL}(f\|g) > 0$ almost surely, so $L\left(w^t\right) \to +\infty$ almost surely.

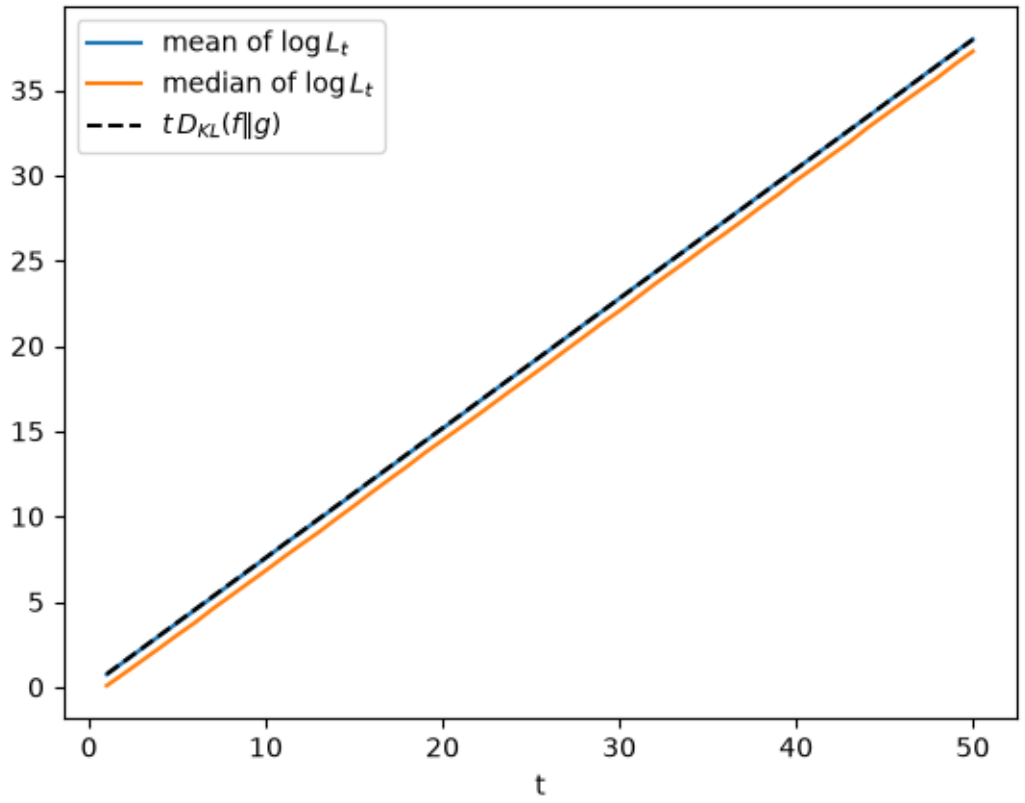
```
# Simulate when nature draws from f
l_arr_f, l_seq_f = simulate_sequences('f', f, g, rng,
                        (F_a, F_b), (G_a, G_b), N=50000)
```

The next figure plots the mean and the median across paths of $\log L\left(w^t\right)$, together with the line $t\, D_{KL}(f\|g)$.

```
N, T = l_arr_f.shape
t_grid = np.arange(1, T+1)
KL_fg, _ = quad(lambda w: f(w) * np.log(f(w) / g(w)), 0, 1)
log_l_seq_f = np.log(l_seq_f)

plt.plot(t_grid, np.mean(log_l_seq_f, axis=0), label=r'mean of $\log L_t$')
plt.plot(t_grid, np.median(log_l_seq_f, axis=0), label=r'median of $\log L_t$')
plt.plot(t_grid, KL_fg * t_grid, 'k--', label=r'$t \, D_{KL}(f\|g)$')
plt.xlabel('t')
plt.legend()
plt.show()

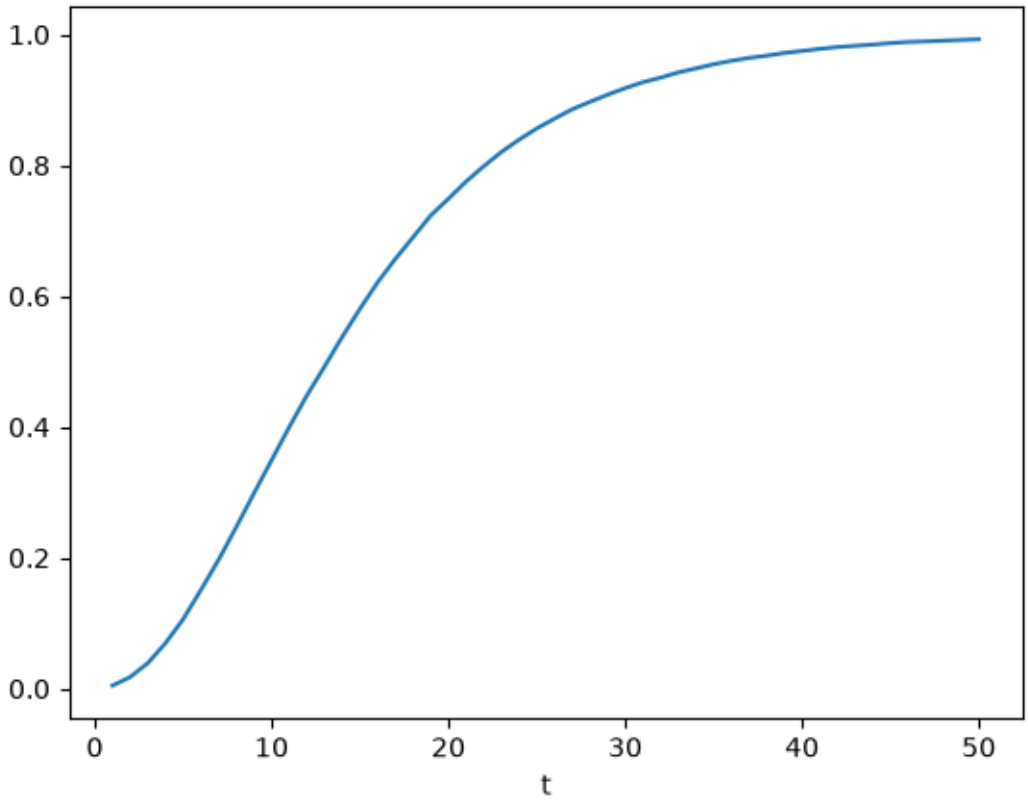
print(f"D_KL(f||g) = {KL_fg:.4f}")
```




```
D_KL(f||g) = 0.7592
```

We also plot the probability that $L\left(w^t\right)$ falls into the interval $[10000, \infty)$ as a function of time and watch how fast probability mass diverges to $+\infty$.

```
plt.plot(t_grid, np.sum(l_seq_f > 10000, axis=0) / N)
plt.xlabel('t')
plt.show()
```



22.6 Likelihood ratio test

We now describe how to employ the machinery of Neyman and Pearson [Neyman and Pearson, 1933] to test the hypothesis that history w^t is generated by repeated IID draws from density f.

Denote q as the data generating process, so that $q = f$ or g.

Upon observing a sample $\{W_i\}_{i=1}^t$, we want to decide whether nature is drawing from g or from f by performing a (frequentist) hypothesis test.

We specify

- Null hypothesis H_0: $q = f$,
- Alternative hypothesis H_1: $q = g$.

Neyman and Pearson proved that the best way to test this hypothesis is to use a **likelihood ratio test** that takes the form:

- accept H_0 if $L(W^t) > c$,
- reject H_0 if $L(W^t) < c$,

where c is a given discrimination threshold.

Setting $c = 1$ is a common choice.

When f and g have equal prior probabilities, $c = 1$ selects the model with the larger posterior probability, which minimizes the probability of selecting the wrong model.

We'll discuss consequences of other choices of c below.

This test is *best* in the sense that it is **most powerful**.

To understand what this means, we have to define probabilities of two important events that allow us to characterize a test associated with a given threshold c.

The two probabilities are:

- Probability of a Type I error in which we reject H_0 when it is true:

$$\alpha \equiv \Pr\{L(w^t) < c \mid q = f\}$$

- Probability of a Type II error in which we accept H_0 when it is false:

$$\beta \equiv \Pr\{L(w^t) > c \mid q = g\}$$

The probability α is also called the probability of false alarm or the significance level.

The probability $1 - \beta = \Pr\{L(w^t) < c \mid q = g\}$ is called the probability of detection or the power of the test.

The Neyman-Pearson Lemma states that among all possible tests, a likelihood ratio test maximizes the probability of detection for a given probability of false alarm.

Another way to say the same thing is that among all possible tests, a likelihood ratio test maximizes **power** for a given **significance level**.

We want a small probability of false alarm and a large probability of detection.

With sample size t fixed, we can change our two probabilities by adjusting c.

A troublesome "that's life" fact is that these two probabilities move in the same direction as we vary the critical value c.

Without specifying quantitative losses from making Type I and Type II errors, there is little that we can say about how we *should* trade off probabilities of the two types of mistakes.

We do know that increasing sample size t improves statistical inference.

The experiment that reveals w^t is more informative in Blackwell's sense than the experiment that reveals w^s for $t > s$, because a statistician can always discard the last $t - s$ observations.

When there are only two hypotheses, one experiment is more informative than another in Blackwell's sense if and only if its receiver operating characteristic curve, defined below, lies everywhere above the other's.

Blackwell's Theorem on Comparing Experiments describes Blackwell's ordering of experiments; its section on *Bayesian likelihood-ratio learning* applies it to the setting of this lecture.

Below we plot some informative figures that illustrate this.

We also present a classical frequentist method for choosing a sample size t.

Let's start with a case in which we fix the threshold c at 1.

```
c = 1
```

Below we plot empirical distributions of logarithms of the cumulative likelihood ratios simulated above, which are generated by either f or g.

Taking logarithms has no effect on calculating the probabilities because the log is a monotonic transformation.

As t increases, the probabilities of making Type I and Type II errors both decrease, which is good.

This is because most of the probability mass of $\log(L(w^t))$ moves toward $-\infty$ when g is the data generating process, while $\log(L(w^t))$ goes to ∞ when data are generated by f.

That disparate behavior of $\log(L(w^t))$ under f and g is what makes it possible eventually to distinguish $q = f$ from $q = g$.

```
def plot_log_histograms(l_seq_f, l_seq_g, c=1, time_points=[1, 7, 14, 21]):
    """Plot log likelihood ratio histograms."""
    fig, axs = plt.subplots(2, 2, figsize=(12, 8))

    for i, t in enumerate(time_points):
        nr, nc = i // 2, i % 2

        axs[nr, nc].axvline(np.log(c), color="k", ls="--")

        # Column t-1 of l_seq holds L(w^t)
        hist_f, x_f = np.histogram(np.log(l_seq_f[:, t-1]), 200, density=True)
        hist_g, x_g = np.histogram(np.log(l_seq_g[:, t-1]), 200, density=True)

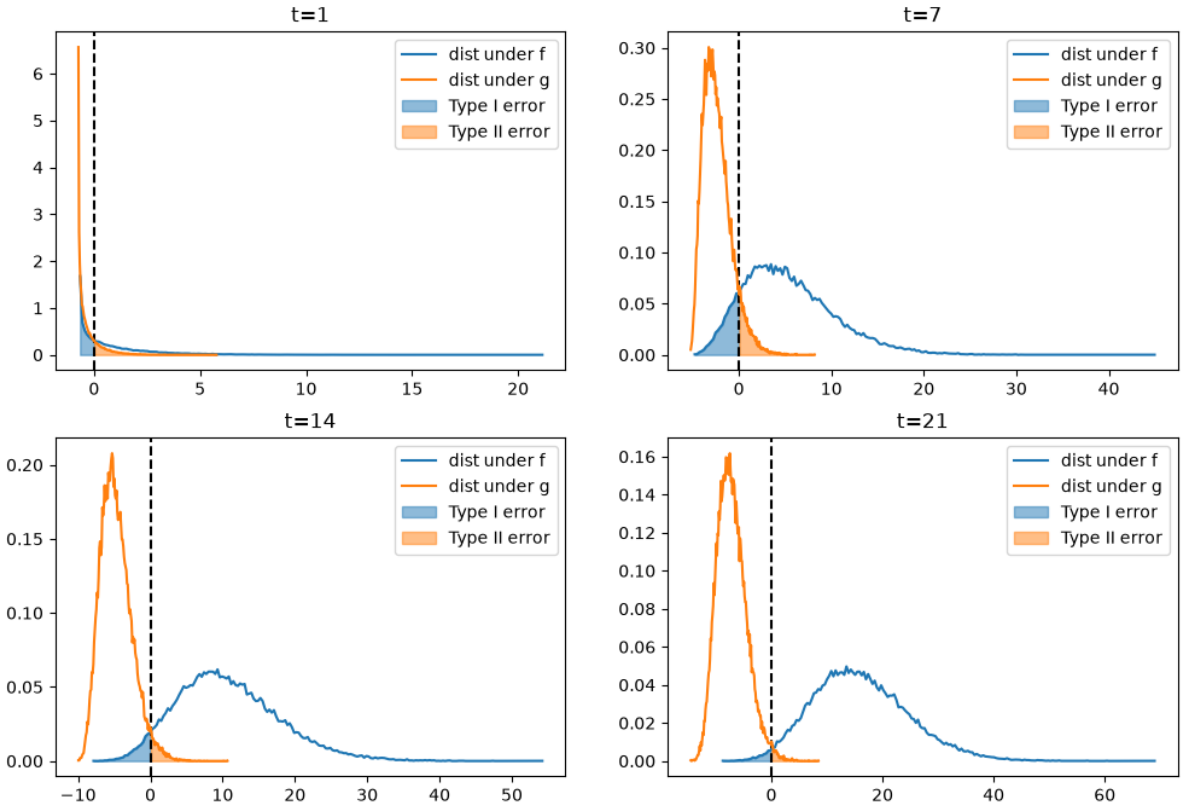
        axs[nr, nc].plot(x_f[1:], hist_f, color="C0", label="dist under f")
        axs[nr, nc].plot(x_g[1:], hist_g, color="C1", label="dist under g")

        # Fill error regions
        for j, (x, hist, label) in enumerate(
            zip([x_f, x_g], [hist_f, hist_g],
            ["Type I error", "Type II error"])):
            ind = x[1:] <= np.log(c) if j == 0 else x[1:] > np.log(c)
            axs[nr, nc].fill_between(x[1:][ind], hist[ind], color=f"C{j}",
                                     alpha=0.5, label=label)

        axs[nr, nc].legend()
        axs[nr, nc].set_title(f"t={t}")

    plt.show()

plot_log_histograms(l_seq_f, l_seq_g, c=c)
```



In the above graphs,

- the blue areas equal the probabilities α of a type I error, up to the binning of the histograms
- the orange areas equal the probabilities β of a type II error, up to the binning of the histograms

Because the logarithm is monotone, the event $\{L_t < 1\}$ is the event $\{\log L_t < 0\}$, so the area under the density of $\log L_t$ to the left of 0 is $\Pr\{L_t < 1\}$.

When we hold c fixed at $c = 1$, the following graph shows that

- the probability of detection monotonically increases with increases in t
- the probability of a false alarm monotonically decreases with increases in t.

```
def compute_error_probabilities(l_seq_f, l_seq_g, c=1):
    """
    Compute Type I and Type II error probabilities.
    """
    N, T = l_seq_f.shape

    # Type I error (false alarm) - reject H0 when true
    PFA = np.array([np.sum(l_seq_f[:, t] < c) / N for t in range(T)])

    # Type II error - accept H0 when false
    beta = np.array([np.sum(l_seq_g[:, t] >= c) / N for t in range(T)])

    # Probability of detection (power)
    PD = np.array([np.sum(l_seq_g[:, t] < c) / N for t in range(T)])

    return {
```

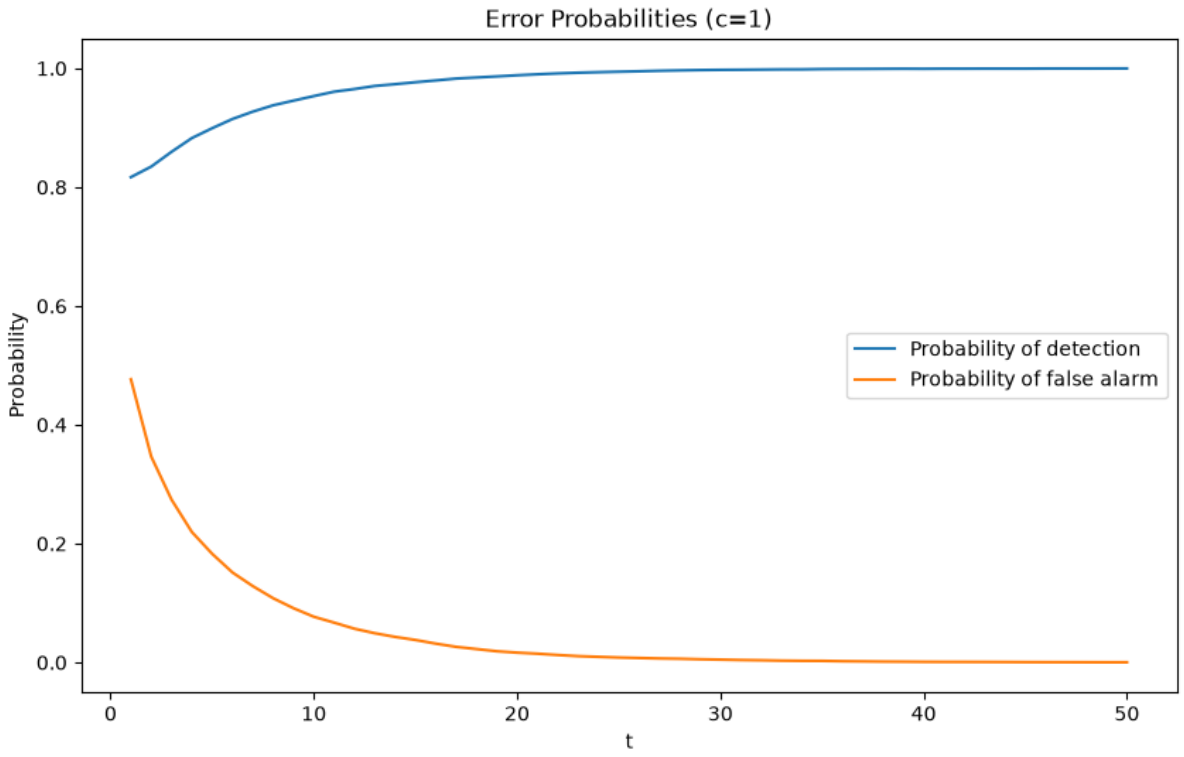
(continues on next page)

(continued from previous page)

```
        'alpha': PFA,
        'beta': beta,
        'PD': PD,
        'PFA': PFA
    }

def plot_error_probabilities(error_dict, T, c=1, title_suffix=""):
    """Plot error probabilities over time."""
    t_grid = np.arange(1, T+1)
    plt.figure(figsize=(10, 6))
    plt.plot(t_grid, error_dict['PD'], label="Probability of detection")
    plt.plot(t_grid, error_dict['PFA'], label="Probability of false alarm")
    plt.xlabel("t")
    plt.ylabel("Probability")
    plt.title(f"Error Probabilities (c={c}){title_suffix}")
    plt.legend()
    plt.show()

error_probs = compute_error_probabilities(l_seq_f, l_seq_g, c=c)
N, T = l_seq_f.shape
plot_error_probabilities(error_probs, T, c)
```



For a given sample size t, the threshold c uniquely pins down probabilities of both types of error.

If for a fixed t we now free up and move c, we will sweep out the probability of detection as a function of the probability of false alarm.

This produces a receiver operating characteristic curve (ROC curve).

Below, we plot receiver operating characteristic curves for different sample sizes t.

```
def plot_roc_curves(l_seq_f, l_seq_g, t_values=[1, 5, 9, 13], N=None):
    """Plot ROC curves for different sample sizes."""
    if N is None:
        N = l_seq_f.shape[0]

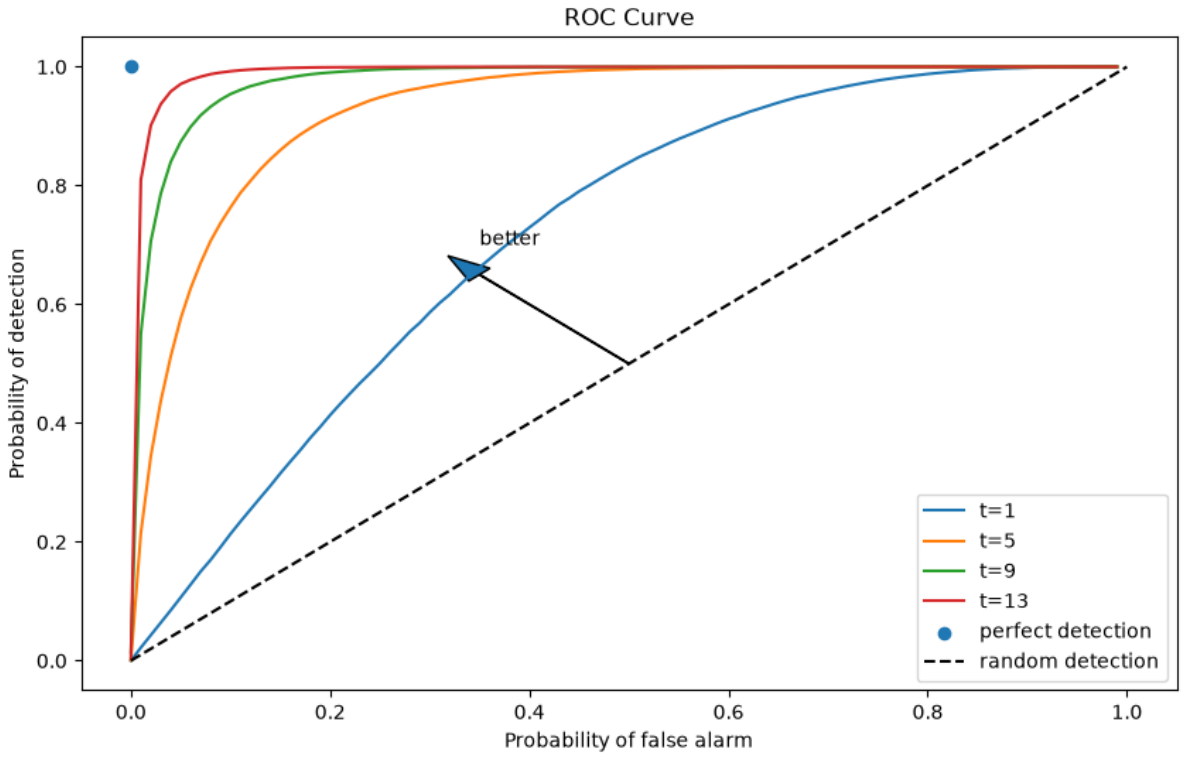
    PFA = np.arange(0, 100, 1)

    plt.figure(figsize=(10, 6))
    for t in t_values:
        # Column t-1 of l_seq holds L(w^t)
        percentile = np.percentile(l_seq_f[:, t-1], PFA)
        PD = [np.sum(l_seq_g[:, t-1] < p) / N for p in percentile]
        plt.plot(PFA / 100, PD, label=f"t={t}")

    plt.scatter(0, 1, label="perfect detection")
    plt.plot([0, 1], [0, 1], color='k', ls='--', label="random detection")

    plt.arrow(0.5, 0.5, -0.15, 0.15, head_width=0.03)
    plt.text(0.35, 0.7, "better")
    plt.xlabel("Probability of false alarm")
    plt.ylabel("Probability of detection")
    plt.legend()
    plt.title("ROC Curve")
    plt.show()


plot_roc_curves(l_seq_f, l_seq_g, t_values=range(1, 15, 4), N=N)
```



Notice that as t increases, we are assured a larger probability of detection and a smaller probability of false alarm associated with a given discrimination threshold c.

For a given sample size t, both α and β change as we vary c.

As we increase c

- $\alpha \equiv \Pr\{L(w^t) < c \mid q = f\}$ increases
- $\beta \equiv \Pr\{L(w^t) > c \mid q = g\}$ decreases

As $t \to +\infty$, we approach the perfect detection curve that is indicated by a right angle hinging on the blue dot.

For a given sample size t, the discrimination threshold c determines a point on the receiver operating characteristic curve.

It is up to the test designer to trade off probabilities of making the two types of errors.

But we know how to choose the smallest sample size to achieve given targets for the probabilities.

Typically, frequentists aim for a high probability of detection that respects an upper bound on the probability of false alarm.

Below we show an example in which we fix the probability of false alarm at 0.05.

The required sample size for making a decision is then determined by a target probability of detection, for example, 0.9, as depicted in the following graph.

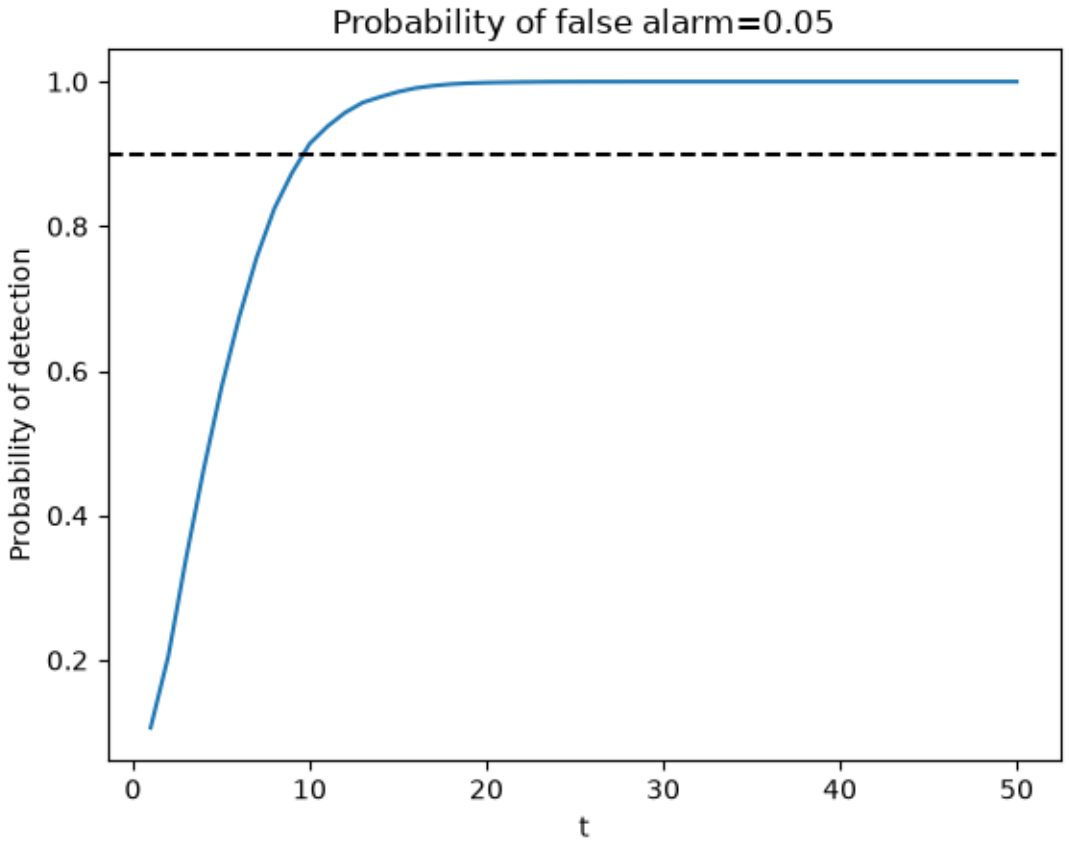
```
PFA = 0.05
PD = np.empty(T)

for t in range(T):

    c = np.percentile(l_seq_f[:, t], PFA * 100)
    PD[t] = np.sum(l_seq_g[:, t] < c) / N

# PD[t] refers to sample size t+1
plt.plot(np.arange(1, T+1), PD)
plt.axhline(0.9, color="k", ls="--")

plt.xlabel("t")
plt.ylabel("Probability of detection")
plt.title(f"Probability of false alarm={PFA}")
plt.show()
```



The United States Navy evidently used a procedure like this to select a sample size t for doing quality control tests during World War II.

A Navy Captain who had been ordered to perform tests of this kind had doubts about it that he presented to Milton Friedman, as we describe in *this lecture*.

A Bayesian Formulation of Friedman and Wald's Problem formulates the Captain's problem as a Bayesian sequential decision problem, and *Bayesian versus Frequentist Decision Rules* compares the Navy's frequentist rule with the Bayesian decision rule.

22.6.1 A third distribution h

Now let's consider a case in which neither g nor f generates the data.

Instead, a third distribution h does.

Let's study how accumulated likelihood ratios L behave when h governs the data.

A key tool here is called **Kullback–Leibler divergence** we studied in *Statistical Divergence Measures*.

In our application, we want to measure how much f or g diverges from h

Two Kullback–Leibler divergences pertinent for us are K_f and K_g defined as

$$\begin{aligned} K_f = D_{KL}(h\|f) = KL(h,f) &= E_h\left[\log\frac{h(w)}{f(w)}\right] \\ &= \int \log\left(\frac{h(w)}{f(w)}\right) h(w)dw. \end{aligned}$$

$$K_g = D_{KL}(h\|g) = KL(h,g) = E_h\left[\log\frac{h(w)}{g(w)}\right]$$
$$= \int \log\left(\frac{h(w)}{g(w)}\right) h(w)dw.$$

Let's compute the Kullback–Leibler discrepancies using the closed form for Beta densities from *Statistical Divergence Measures*.

```
def KL_beta(a1, b1, a2, b2):
    """KL(Beta(a1, b1) ‖ Beta(a2, b2)) in closed form."""
    return (betaln(a2, b2) - betaln(a1, b1)
            + (a1 - a2) * digamma(a1) + (b1 - b2) * digamma(b1)
            + (a2 - a1 + b2 - b1) * digamma(a1 + b1))

def compute_KL_h(h_params, F_params=(F_a, F_b), G_params=(G_a, G_b)):
    """
    Compute K_f = KL(h, f) and K_g = KL(h, g) for Beta densities h, f, g
    """
    Kf = KL_beta(*h_params, *F_params)
    Kg = KL_beta(*h_params, *G_params)
    return Kf, Kg
```

22.6.2 A helpful formula

There is a mathematical relationship between likelihood ratios and KL divergence.

When data is generated by distribution h, the expected log likelihood ratio is:

$$\frac{1}{t}E_h[\log L_t] = K_g - K_f \tag{22.1}$$

where $L_t = \prod_{j=1}^{t}\frac{f(w_j)}{g(w_j)}$ is the likelihood ratio process.

Equation (22.1) tells us that:

- When $K_g < K_f$ (i.e., g is closer to h than f is), the expected log likelihood ratio is negative, so $L\left(w^t\right) \to 0$.
- When $K_g > K_f$ (i.e., f is closer to h than g is), the expected log likelihood ratio is positive, so $L\left(w^t\right) \to +\infty$.

Let's verify this using simulation.

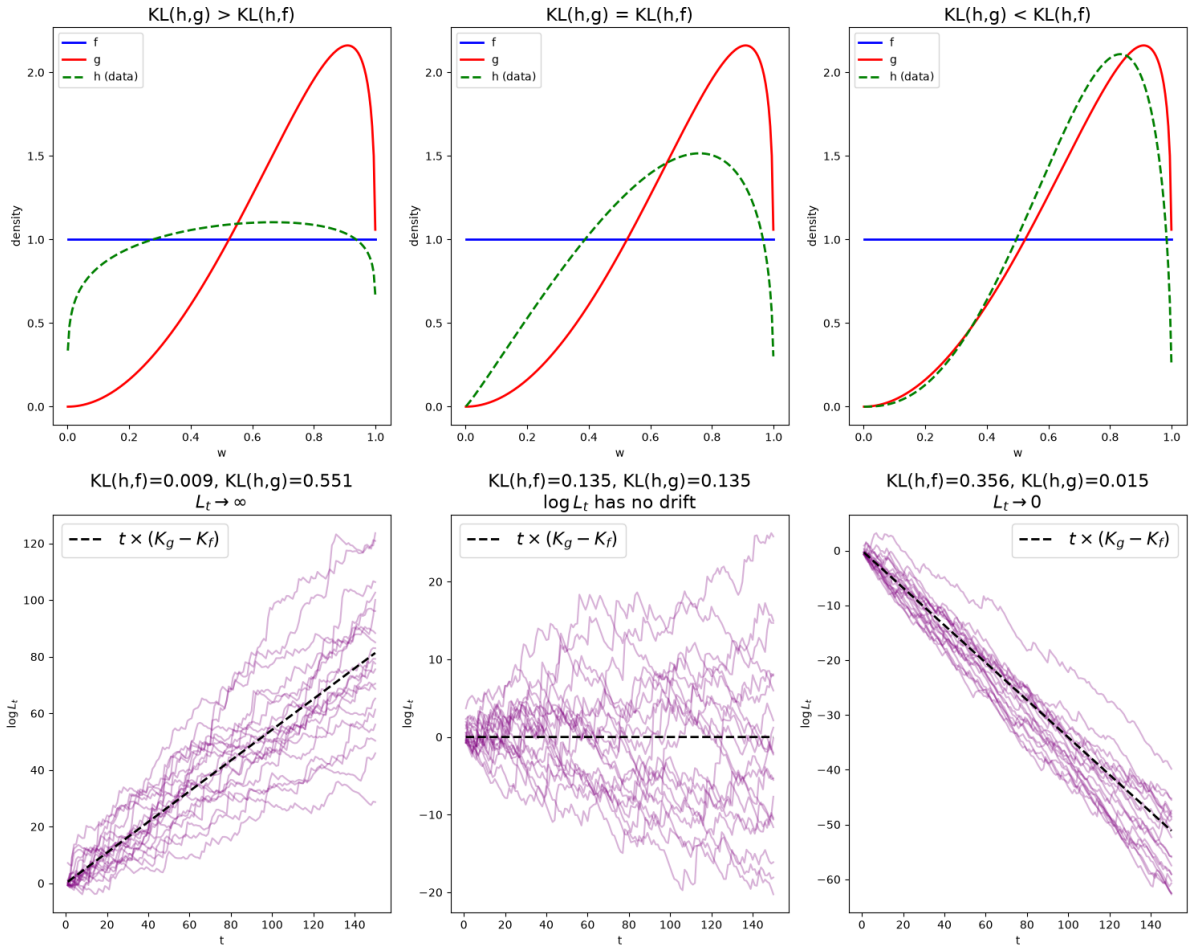
In the simulation, we generate multiple paths using Beta distributions f, g, and h, and compute the paths of $\log(L(w^t))$.

First, we write a function to compute the likelihood ratio process

```
def compute_likelihood_ratios(sequences, f, g):
    """Compute likelihood ratios and cumulative products."""
    l_ratios = f(sequences) / g(sequences)
    L_cumulative = np.cumprod(l_ratios, axis=1)
    return l_ratios, L_cumulative
```

We consider three cases: (1) h is closer to f, (2) f and g are equidistant from h, and (3) h is closer to g.

In case (2) we set $h = \text{Beta}(2.1035, 1.35)$, for which K_f and K_g agree to four decimal places.



Note that

- In the first figure, $\log L(w^t)$ diverges to ∞ because $K_g > K_f$.
- In the second figure, $K_g = K_f$, so $\log L(w^t)$ is a random walk with IID increments of mean zero; its paths wander away from zero in both directions, with a spread that grows like $\sqrt{t}$.
- In the last figure, $\log L(w^t)$ diverges to $-\infty$ because $K_g < K_f$.
- The black dashed line, $t\left(D_{KL}(h\|g) - D_{KL}(h\|f)\right)$, passes through the center of the paths, as (22.1) predicts.

These observations align with the theory.

In *Heterogeneous Beliefs and Financial Markets*, we will see an application of these ideas.

22.7 Hypothesis testing and classification

This section discusses another application of likelihood ratio processes.

We describe how a statistician can combine frequentist probabilities of type I and type II errors in order to

- compute an anticipated frequency of selecting a wrong model based on a sample length T
- compute an anticipated error rate in a classification problem

We consider a situation in which nature generates data by mixing known densities f and g with known mixing parameter $\pi_{-1} \in (0, 1)$ so that the random variable w is drawn from the density

$$h(w) = \pi_{-1} f(w) + (1 - \pi_{-1}) g(w)$$

We assume that the statistician knows the densities f and g and also the mixing parameter π_{-1}.

Below, we'll set $\pi_{-1} = .5$, although much of the analysis would follow through with other settings of $\pi_{-1} \in (0, 1)$.

We assume that f and g both put positive probabilities on the same intervals of possible realizations of the random variable W.

In the simulations below, we specify that f is a Beta$(1, 1)$ distribution and that g is Beta$(3, 1.2)$ distribution.

We consider two alternative timing protocols.

- Timing protocol 1 is for the model selection problem
- Timing protocol 2 is for the individual classification problem

Timing Protocol 1: Nature flips a coin only **once** at time $t = -1$ and with probability π_{-1} generates a sequence $\{w_t\}_{t=1}^T$ of IID draws from f and with probability $1 - \pi_{-1}$ generates a sequence $\{w_t\}_{t=1}^T$ of IID draws from g.

Timing Protocol 2: Nature flips a coin **often**. At each time $t \geq 1$, nature flips a coin and with probability π_{-1} draws w_t from f and with probability $1 - \pi_{-1}$ draws w_t from g.

Here is Python code that we'll use to implement timing protocol 1 and 2

```
def protocol_1(π_minus_1, T, rng, N=1000, F_params=(1, 1), G_params=(3, 1.2)):
    """
    Simulate Protocol 1: Nature decides once at t=-1 which model to use.
    """
    F_a, F_b = F_params
    G_a, G_b = G_params

    # Single coin flip for the true model
    true_models_F = rng.random(N) < π_minus_1
    sequences = np.empty((N, T))

    n_f = np.sum(true_models_F)
    n_g = N - n_f

    if n_f > 0:
        sequences[true_models_F, :] = rng.beta(F_a, F_b, (n_f, T))
    if n_g > 0:
        sequences[~true_models_F, :] = rng.beta(G_a, G_b, (n_g, T))

    return sequences, true_models_F

def protocol_2(π_minus_1, T, rng, N=1000, F_params=(1, 1), G_params=(3, 1.2)):
    """
    Simulate Protocol 2: Nature decides at each time step which model to use.
    """
    F_a, F_b = F_params
    G_a, G_b = G_params

    # Coin flips for each time step
    true_models_F = rng.random((N, T)) < π_minus_1
    sequences = np.empty((N, T))

    n_f = np.sum(true_models_F)
```

(continues on next page)

(continued from previous page)

```
    n_g = N * T - n_f

    if n_f > 0:
        sequences[true_models_F] = rng.beta(F_a, F_b, n_f)
    if n_g > 0:
        sequences[~true_models_F] = rng.beta(G_a, G_b, n_g)

    return sequences, true_models_F
```

Remark: Under timing protocol 2, the $\{w_t\}_{t=1}^T$ is a sequence of IID draws from $h(w)$. Under timing protocol 1, the $\{w_t\}_{t=1}^T$ is not IID. It is **conditionally IID** – meaning that with probability π_{-1} it is a sequence of IID draws from $f(w)$ and with probability $1-\pi_{-1}$ it is a sequence of IID draws from $g(w)$. For more about this, see *this lecture about exchangeability*.

Likelihood Ratio Processes and Bayesian Learning studies a Bayesian who knows that timing protocol 1 governs the data and learns which density nature chose.

Incorrect Models studies a Bayesian who wrongly believes timing protocol 1 when timing protocol 2 generates the data.

We again deploy a **likelihood ratio process** with time t component being the likelihood ratio

$$\ell(w_t) = \frac{f(w_t)}{g(w_t)}, \quad t \geq 1.$$

The **likelihood ratio process** for sequence $\{w_t\}_{t=1}^\infty$ is

$$L(w^t) = \prod_{i=1}^t \ell(w_i),$$

For shorthand we'll write $L_t = L(w^t)$.

22.7.1 Model selection mistake probability

We first study a problem that assumes timing protocol 1.

Consider a decision maker who wants to know whether model f or model g governs a data set of length T observations.

The decision maker has observed a sequence $\{w_t\}_{t=1}^T$.

On the basis of that observed sequence, a likelihood ratio test selects model f when $L_T \geq 1$ and model g when $L_T < 1$.

When model f generates the data, the probability that the likelihood ratio test selects the wrong model is

$$p_f = \text{Prob}\left(L_T < 1 \middle| f\right) = \alpha_T.$$

When model g generates the data, the probability that the likelihood ratio test selects the wrong model is

$$p_g = \text{Prob}\left(L_T \geq 1 \middle| g\right) = \beta_T.$$

We can construct a probability that the likelihood ratio selects the wrong model by assigning a Bayesian prior probability of $\pi_{-1} = .5$ that nature selects model f then averaging p_f and p_g to form the prior (ex ante) probability of a detection error equal to

$$p(\text{wrong decision}) = \frac{1}{2}(\alpha_T + \beta_T). \tag{22.2}$$

Now let's simulate timing protocol 1 and compute the error probabilities

```
def compute_protocol_1_errors(π_minus_1, T_max, N_simulations, f_func, g_func,
                              rng, F_params=(1, 1), G_params=(3, 1.2)):
    """
    Compute error probabilities for Protocol 1.
    """
    sequences, true_models = protocol_1(
        π_minus_1, T_max, rng, N_simulations, F_params, G_params)
    l_ratios, L_cumulative = compute_likelihood_ratios(sequences,
                                    f_func, g_func)

    T_range = np.arange(1, T_max + 1)

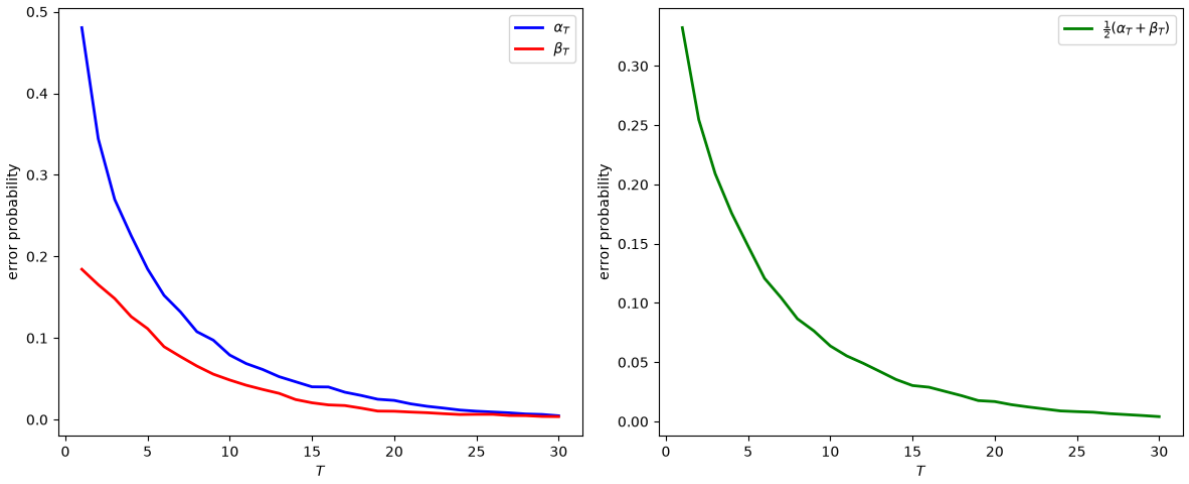
    mask_f = true_models
    mask_g = ~true_models

    L_f = L_cumulative[mask_f, :]
    L_g = L_cumulative[mask_g, :]

    α_T = np.mean(L_f < 1, axis=0)
    β_T = np.mean(L_g >= 1, axis=0)
    error_prob = 0.5 * (α_T + β_T)

    return {
        'T_range': T_range,
        'alpha': α_T,
        'beta': β_T,
        'error_prob': error_prob,
        'L_cumulative': L_cumulative,
        'true_models': true_models
    }
```

The following code visualizes the error probabilities for timing protocol 1




```
At T=30:
α_30 = 0.0046
β_30 = 0.0036
Model selection error probability = 0.0041
```

Notice how the model selection error probability approaches zero as T grows.

22.7.2 Classification

We now consider a problem that assumes timing protocol 2.

A decision maker wants to classify components of an observed sequence $\{w_t\}_{t=1}^T$ as having been drawn from either f or g.

The decision maker uses the following classification rule:

$$w_t \text{ is from } f \text{ if } l_t \geq 1$$
$$w_t \text{ is from } g \text{ if } l_t < 1.$$

Under this rule, the expected misclassification rate is

$$p(\text{misclassification}) = \frac{1}{2}(\tilde{\alpha}_t + \tilde{\beta}_t) \tag{22.3}$$

where $\tilde{\alpha}_t = \text{Prob}(l_t < 1 \mid f)$ and $\tilde{\beta}_t = \text{Prob}(l_t \geq 1 \mid g)$.

Now let's write some code to simulate it

```
def compute_protocol_2_errors(π_minus_1, T_max, N_simulations, f_func, g_func,
                              rng, F_params=(1, 1), G_params=(3, 1.2)):
    """
    Compute error probabilities for Protocol 2.
    """
    sequences, true_models = protocol_2(π_minus_1,
                        T_max, rng, N_simulations, F_params, G_params)
    l_ratios, _ = compute_likelihood_ratios(sequences, f_func, g_func)

    T_range = np.arange(1, T_max + 1)

    accuracy = np.empty(T_max)
    for t in range(T_max):
        predictions = (l_ratios[:, t] >= 1)
        actual = true_models[:, t]
        accuracy[t] = np.mean(predictions == actual)

    return {
        'T_range': T_range,
        'accuracy': accuracy,
        'l_ratios': l_ratios,
        'true_models': true_models
    }
```

Since for each t, the decision boundary is the same, the decision boundary can be computed as

```
root = brentq(lambda w: f(w) / g(w) - 1, 0.001, 0.99)
root_2 = brentq(lambda w: f(w) / g(w) - 1, 0.99, 1 - 1e-12)
print(f"f/g = 1 at w = {root:.4f} and at w = {root_2:.5f}")
```

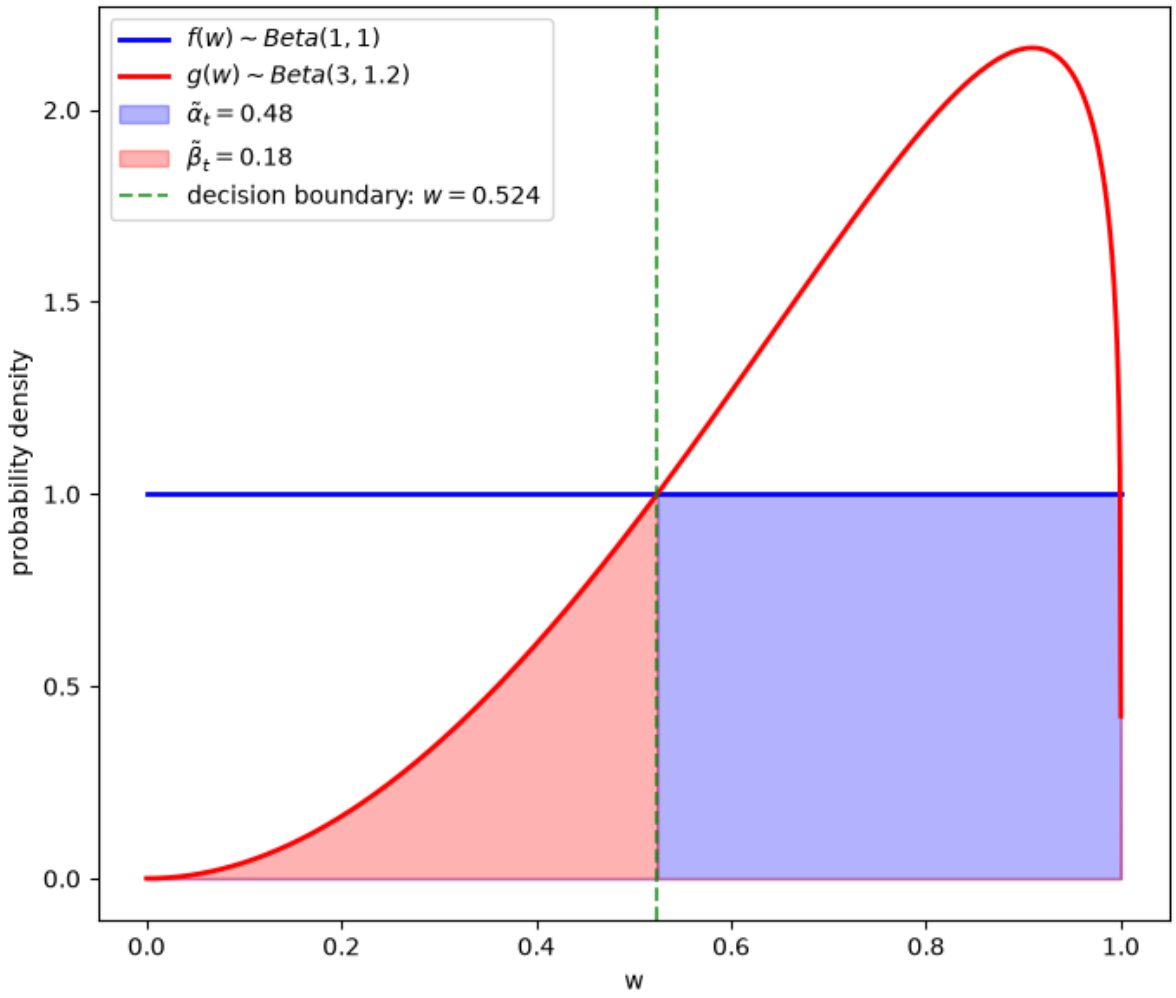
```
f/g = 1 at w = 0.5241 and at w = 0.99925
```

The likelihood ratio f/g equals 1 at two points, because g vanishes at both ends of $(0, 1)$.

The rule classifies w_t as coming from g when w_t lies between the two roots and as coming from f otherwise.

The second root lies so close to 1 that the interval to its right has probability below 0.001 under both f and g.

We can plot the distributions of f and g and the decision boundary



To the left of the green vertical line $g < f$, so $l_t > 1$; therefore a w_t that falls to the left of the green line is classified as a type f individual.

- The shaded red area equals $\tilde{\beta}_t$ – the probability of classifying someone as a type f individual when it is really a type g individual.

To the right of the green vertical line and to the left of the second root $w = 0.99925$, $g > f$, so $l_t < 1$; therefore a w_t that falls in this interval is classified as a type g individual.

- The shaded blue area equals $\tilde{\alpha}_t$ – the probability of classifying someone as a type g individual when it is really a type f individual.

This gives us clues about how to compute the theoretical classification error probability

```
# tilde α_t = P(l_t < 1 | f) = P(root < w < root_2 | f)
α_theory = beta_dist.cdf(root_2, F_a, F_b) - beta_dist.cdf(root, F_a, F_b)

# tilde β_t = P(l_t >= 1 | g) = P(w <= root or w >= root_2 | g)
β_theory = beta_dist.cdf(root, G_a, G_b) + beta_dist.sf(root_2, G_a, G_b)

theory_error = 0.5 * (α_theory + β_theory)

print(f"theoretical tilde α_t = {α_theory:.4f}")
```

(continues on next page)

(continued from previous page)

```
print(f"theoretical tilde β_t = {β_theory:.4f}")
print(f"theoretical classification error probability = {theory_error:.4f}")
```

```
theoretical tilde α_t = 0.4752
theoretical tilde β_t = 0.1836
theoretical classification error probability = 0.3294
```

Now we simulate timing protocol 2 and compute the classification error probability.

In the next cell, we also compare the theoretical classification accuracy to the empirical classification accuracy

```
def analyze_protocol_2(π_minus_1, T_max, N_simulations, f_func, g_func, rng,
                       theory_error=None, F_params=(1, 1), G_params=(3, 1.2)):
    """Analyze Protocol 2."""
    result = compute_protocol_2_errors(π_minus_1, T_max, N_simulations,
                                       f_func, g_func, rng, F_params, G_params)

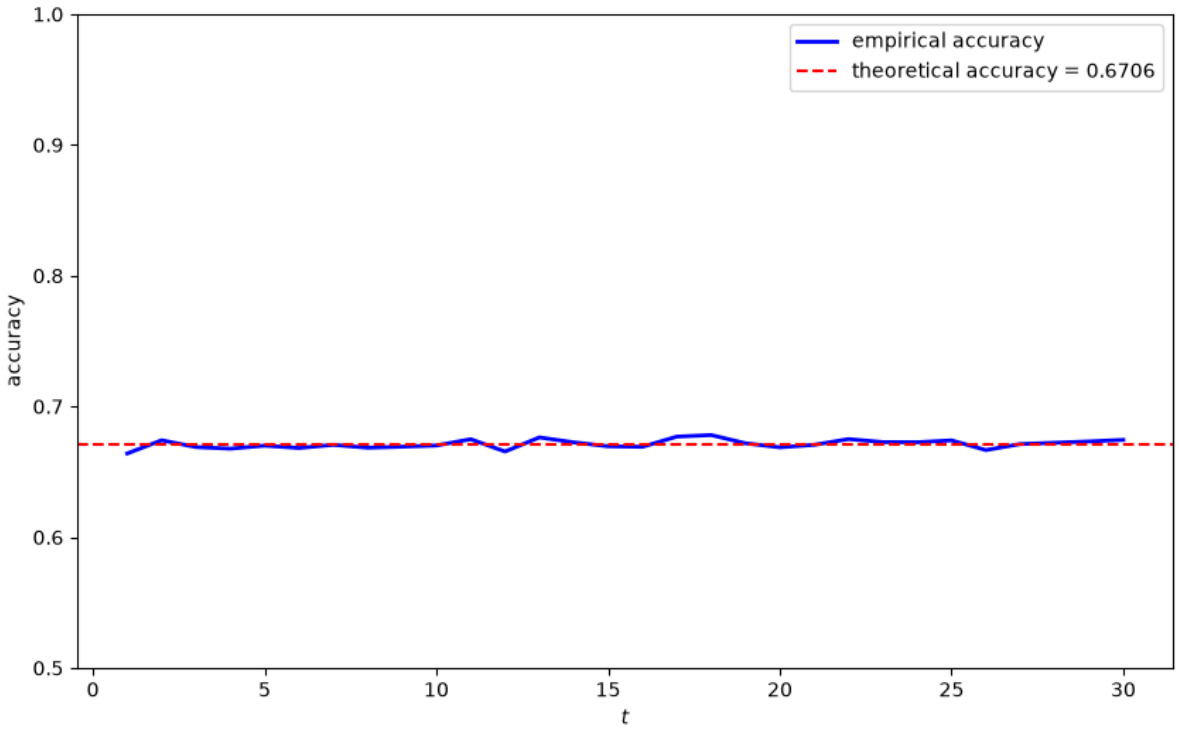
    # Plot results
    plt.figure(figsize=(10, 6))
    plt.plot(result['T_range'], result['accuracy'],
             'b-', linewidth=2, label='empirical accuracy')

    if theory_error is not None:
        plt.axhline(1 - theory_error, color='r', linestyle='--',
                    label=f'theoretical accuracy = {1 - theory_error:.4f}')

    plt.xlabel('$t$')
    plt.ylabel('accuracy')
    plt.legend()
    plt.ylim(0.5, 1.0)
    plt.show()

    return result

# Analyze Protocol 2
result_p2 = analyze_protocol_2(π_minus_1, T_max, N_simulations, f, g, rng,
                               theory_error, (F_a, F_b), (G_a, G_b))
```



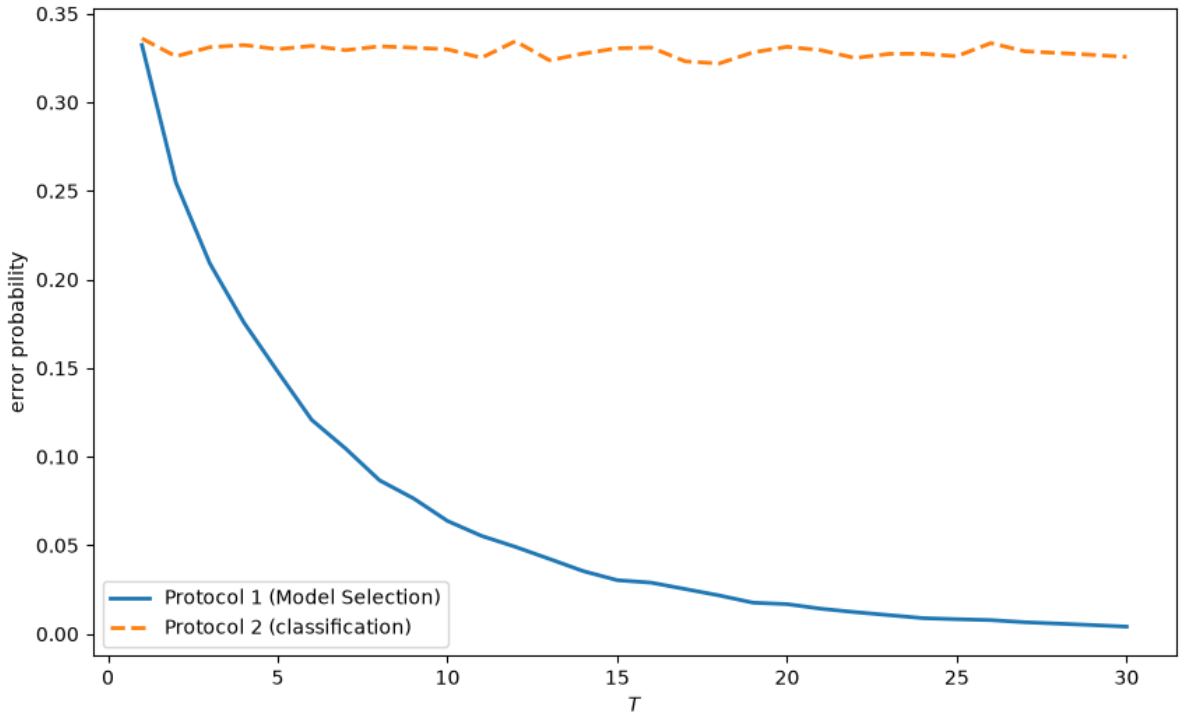
Let's watch decisions made by the two timing protocols as more and more observations accrue.

```
def compare_protocols(result1, result2):
    """Compare results from both protocols."""
    plt.figure(figsize=(10, 6))

    plt.plot(result1['T_range'], result1['error_prob'], linewidth=2,
             label='Protocol 1 (Model Selection)')
    plt.plot(result2['T_range'], 1 - result2['accuracy'],
             linestyle='--', linewidth=2,
             label='Protocol 2 (classification)')

    plt.xlabel('$T$')
    plt.ylabel('error probability')
    plt.legend()
    plt.show()

compare_protocols(result_p1, result_p2)
```



From the figure above, we can see:

- For both timing protocols, the error probability starts at the same level, subject to a little randomness.
- For timing protocol 1, the error probability decreases as the sample size increases because we are making just **one** decision – i.e., selecting whether f or g governs **all** individuals. More data provides better evidence.
- For timing protocol 2, the error probability remains constant because we are making **many** decisions – one classification decision for each observation.

Remark: Think about how laws of large numbers are applied to compute error probabilities for the model selection problem and the classification problem.

22.7.3 Error probability and divergence measures

A plausible guess is that the ability of a likelihood ratio to distinguish distributions f and g depends on how "different" they are.

We have learnt some measures of "difference" between distributions in *Statistical Divergence Measures*.

Let's now study two more measures of "difference" between distributions that are useful in the context of model selection and classification.

Recall that Chernoff entropy between probability densities f and g is defined as:

$$C(f,g) = -\log \min_{\phi \in (0,1)} \rho(\phi), \qquad \rho(\phi) = \int f^{\phi}(x) g^{1-\phi}(x) dx$$

Let $f_T(w^T) = \prod_{t=1}^{T} f(w_t)$ and $g_T(w^T) = \prod_{t=1}^{T} g(w_t)$ be the joint densities of w^T under the two models.

The likelihood ratio test selects f when $f_T \geq g_T$ and g when $f_T < g_T$.

Its model selection error probability (22.2) is therefore

$$\frac{1}{2}(\alpha_T + \beta_T) = \frac{1}{2}\int_{f_T<g_T} f_T\, dw^T + \frac{1}{2}\int_{f_T\geq g_T} g_T\, dw^T = \frac{1}{2}\int \min\{f_T, g_T\}\, dw^T.$$

For $x, y \geq 0$ and $\phi \in [0,1]$, $\min\{x,y\} \leq x^\phi y^{1-\phi}$.

Because the w_t's are IID,

$$\frac{1}{2}(\alpha_T + \beta_T) \leq \frac{1}{2}\int f_T^\phi\, g_T^{1-\phi}\, dw^T = \frac{1}{2}\rho(\phi)^T$$

for every $\phi \in [0,1]$.

Minimizing the right side over ϕ gives an upper bound on model selection error probability:

$$\frac{1}{2}(\alpha_T + \beta_T) \leq \frac{1}{2}e^{-C(f,g)T}. \tag{22.4}$$

The weaker bound $e^{-C(f,g)T}$ also holds.

Chernoff [1952] showed that the bound (22.4) delivers the correct exponential rate:

$$-\frac{1}{T}\log\frac{1}{2}(\alpha_T + \beta_T) \to C(f,g).$$

The approach to the limit is slow, because the error probability contains a factor that decays polynomially in T, which the bound omits.

For Beta densities, the Chernoff coefficient $\rho(\phi)$ has a closed form that we derived in *Statistical Divergence Measures*.

Let's compute Chernoff entropy with some Python code

```
def log_chernoff_coef_beta(ϕ, a1, b1, a2, b2):
    """log ρ(ϕ) for Beta(a1, b1) and Beta(a2, b2) in closed form."""
    return (betaln(ϕ * a1 + (1 - ϕ) * a2, ϕ * b1 + (1 - ϕ) * b2)
            - ϕ * betaln(a1, b1) - (1 - ϕ) * betaln(a2, b2))

def chernoff_beta(a1, b1, a2, b2):
    """Chernoff entropy C(f, g) and minimizing ϕ for two Beta densities."""
    res = minimize_scalar(log_chernoff_coef_beta, bounds=(0, 1),
                          args=(a1, b1, a2, b2), method='bounded')
    return -res.fun, res.x

C_fg, ϕ_optimal = chernoff_beta(F_a, F_b, G_a, G_b)
print(f"Chernoff entropy C(f,g) = {C_fg:.4f}")
print(f"Optimal ϕ = {ϕ_optimal:.4f}")
```

```
Chernoff entropy C(f,g) = 0.1212
Optimal ϕ = 0.5969
```

Now let's compare the model selection error probability that we simulated under timing protocol 1 with the bounds $\frac{1}{2}e^{-C(f,g)T}$ and $e^{-C(f,g)T}$.

The right panel plots $-\frac{1}{T}$ log of the simulated error probability against T.

```
T_range = np.arange(1, T_max+1)
chernoff_bound = 0.5 * np.exp(-C_fg * T_range)

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(14, 5))
```

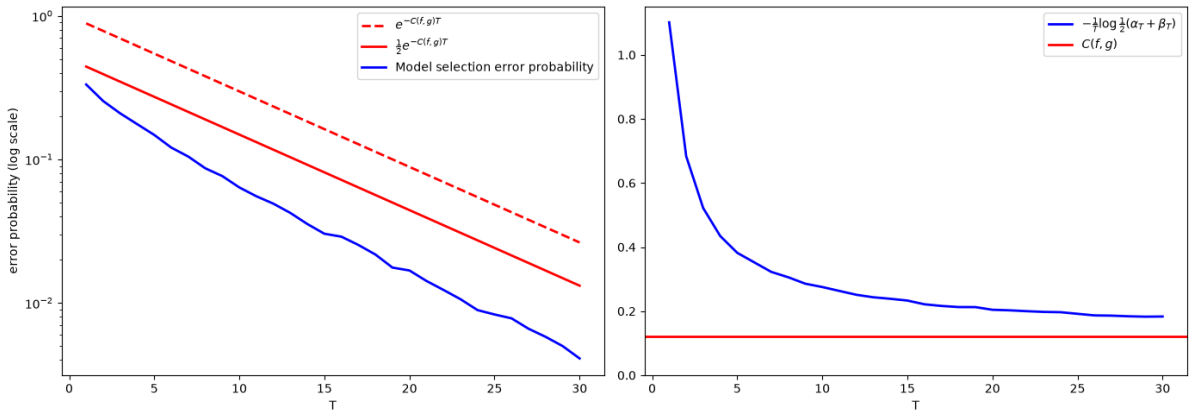
(continues on next page)

(continued from previous page)

```
ax1.semilogy(T_range, 2 * chernoff_bound, 'r--', linewidth=2,
             label=f'$e^{{-C(f,g)T}}$')
ax1.semilogy(T_range, chernoff_bound, 'r-', linewidth=2,
             label=r'$\frac{1}{2} e^{-C(f,g)T}$')
ax1.semilogy(T_range, result_p1['error_prob'], 'b-', linewidth=2,
             label='Model selection error probability')
ax1.set_xlabel('T')
ax1.set_ylabel('error probability (log scale)')
ax1.legend()

ax2.plot(T_range, -np.log(result_p1['error_prob']) / T_range, 'b-',
         linewidth=2, label=r'$-\frac{1}{T}\log \frac{1}{2}(\alpha_T + \beta_T)$')
ax2.axhline(C_fg, color='r', linestyle='-', linewidth=2, label='$C(f,g)$')
ax2.set_xlabel('T')
ax2.set_ylim(0, None)
ax2.legend()

plt.tight_layout()
plt.show()
```



The simulated error probability lies below $\frac{1}{2}e^{-C(f,g)T}$, as (22.4) requires.

Its negative log divided by T declines toward $C(f,g)$ slowly.

In *Statistical Divergence Measures*, we also studied **Jensen-Shannon divergence** as a symmetric measure of distance between distributions.

We can use Jensen-Shannon divergence to measure the distance between distributions f and g and compute how it covaries with the model selection error probability.

We also compute Jensen-Shannon divergence numerically with some Python code.

We integrate over all of $(0,1)$, splitting the interval at $1/2$, because U-shaped Beta densities diverge at the endpoints.

```
def compute_JS(f, g):
    """
    Compute Jensen-Shannon divergence
    """
    def integrand(w):
        fw, gw = f(w), g(w)
        mw = 0.5 * (fw + gw)
        return 0.5 * fw * np.log(fw / mw) + 0.5 * gw * np.log(gw / mw)
```

(continues on next page)

(continued from previous page)

```
    left, _ = quad(integrand, 0, 0.5, limit=200)
    right, _ = quad(integrand, 0.5, 1, limit=200)
    return left + right
```

Now let's return to our guess that the error probability at large sample sizes is related to the Chernoff entropy between two distributions.

We verify this by computing the correlation between the log of the error probability at $T = 50$ under Timing Protocol 1 and the divergence measures.

We use the following pairs of Beta distributions for f and g as test cases

```
distribution_pairs = [
    # (f_params, g_params)
    ((1, 1), (0.1, 0.2)),
    ((1, 1), (0.3, 0.3)),
    ((1, 1), (0.3, 0.4)),
    ((1, 1), (0.5, 0.5)),
    ((1, 1), (0.7, 0.6)),
    ((1, 1), (0.9, 0.8)),
    ((1, 1), (1.1, 1.05)),
    ((1, 1), (1.2, 1.1)),
    ((1, 1), (1.5, 1.2)),
    ((1, 1), (2, 1.5)),
    ((1, 1), (2.5, 1.8)),
    ((1, 1), (3, 1.2)),
    ((1, 1), (4, 1)),
    ((1, 1), (5, 1))
]
```

For the most distinct pair, the Chernoff bound (22.4) puts the error probability at $T = 50$ below $\frac{1}{2}e^{-0.443\times 50} \approx 1.2 \times 10^{-10}$.

A simulation of a few thousand sequences from f and from g typically records no errors at all for several of these pairs.

We compute the error probabilities instead by importance sampling, a technique that *Mean of a Likelihood Ratio Process* describes.

We draw w^T from the **tilted density**

$$q(w) = \frac{f(w)^{\phi^*} g(w)^{1-\phi^*}}{\rho(\phi^*)},$$

where ϕ^* attains the minimum in the definition of $C(f, g)$.

For Beta densities f and g, q is again a Beta density, as *Statistical Divergence Measures* shows.

Since $f/q = \rho(\phi^*)\,\ell^{1-\phi^*}$ and $g/q = \rho(\phi^*)\,\ell^{-\phi^*}$,

$$\frac{1}{2}(\alpha_T + \beta_T) = \frac{1}{2}\int \min\left\{\frac{f_T}{q_T}, \frac{g_T}{q_T}\right\} q_T\, dw^T = \frac{1}{2} e^{-C(f,g)T}\, E_q\left[\min\left\{L_T^{1-\phi^*}, L_T^{-\phi^*}\right\}\right],$$

where $q_T(w^T) = \prod_{t=1}^T q(w_t)$.

The random variable inside the expectation lies between 0 and 1.

That fact delivers the bound (22.4) once more and gives the sample average of $\min\{L_T^{1-\phi^*}, L_T^{-\phi^*}\}$ a small variance.

```
def log_beta_pdf(w, a, b):
    """Log density of Beta(a, b)."""
    return (a - 1) * np.log(w) + (b - 1) * np.log1p(-w) - betaln(a, b)

def error_prob_tilted(f_params, g_params, T, N, rng):
    """
    Model selection error probability (α_T + β_T) / 2 by
    importance sampling from the Chernoff-tilted density.
    """
    (f_a, f_b), (g_a, g_b) = f_params, g_params
    C, ϕ = chernoff_beta(f_a, f_b, g_a, g_b)
    q_a, q_b = ϕ * f_a + (1 - ϕ) * g_a, ϕ * f_b + (1 - ϕ) * g_b

    # Keep draws off the boundary, where the log densities are infinite
    w = np.clip(rng.beta(q_a, q_b, (N, T)), 1e-300, 1 - 1e-16)
    log_L = np.sum(log_beta_pdf(w, f_a, f_b)
                   - log_beta_pdf(w, g_a, g_b), axis=1)
    weights = np.exp(np.minimum((1 - ϕ) * log_L, -ϕ * log_L))
    return 0.5 * np.exp(-C * T) * np.mean(weights)
```

As a check, we apply the importance sampling estimator to our f and g at $T = 30$ and compare it with the simulation of timing protocol 1 above.

```
is_estimate = error_prob_tilted((F_a, F_b), (G_a, G_b), T_max, 5000, rng)
print(f"Importance sampling estimate at T={T_max}: {is_estimate:.5f}")
print(f"Simulated protocol 1 error at T={T_max}:   "
      f"{result_p1['error_prob'][-1]:.5f}")
```

```
Importance sampling estimate at T=30: 0.00355
Simulated protocol 1 error at T=30:   0.00410
```

Now we compute the divergence measures and the error probability at $T = 50$ for every pair

```
# Parameters for simulation
T_large = 50
N_sims = 5000

# Initialize arrays
n_pairs = len(distribution_pairs)
kl_fg_vals = np.zeros(n_pairs)
kl_gf_vals = np.zeros(n_pairs)
js_vals = np.zeros(n_pairs)
chernoff_vals = np.zeros(n_pairs)
error_probs = np.zeros(n_pairs)
pair_names = []

for i, ((f_a, f_b), (g_a, g_b)) in enumerate(distribution_pairs):
    # Create density functions
    f_pair = jit(lambda x, a=f_a, b=f_b: p(x, a, b))
    g_pair = jit(lambda x, a=g_a, b=g_b: p(x, a, b))

    # Compute divergence measures
    kl_fg_vals[i] = KL_beta(f_a, f_b, g_a, g_b)
    kl_gf_vals[i] = KL_beta(g_a, g_b, f_a, f_b)
    js_vals[i] = compute_JS(f_pair, g_pair)
    chernoff_vals[i], _ = chernoff_beta(f_a, f_b, g_a, g_b)
```

(continues on next page)

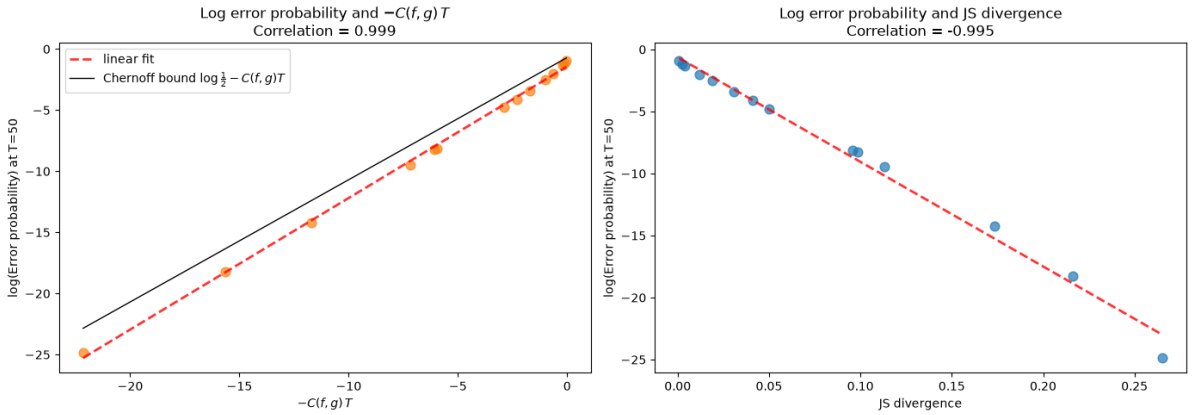
(continued from previous page)

```
    # Error probability by importance sampling
    error_probs[i] = error_prob_tilted((f_a, f_b), (g_a, g_b),
                                       T_large, N_sims, rng)
    pair_names.append(f"Beta({f_a},{f_b}) and Beta({g_a},{g_b})")

cor_data = {
    'kl_fg': kl_fg_vals,
    'kl_gf': kl_gf_vals,
    'js': js_vals,
    'chernoff': chernoff_vals,
    'error_prob': error_probs,
    'names': pair_names,
    'T': T_large}
```

Now let's visualize the correlations.

The left panel plots the log error probability against $-C(f,g)T$, together with the Chernoff bound $\log \frac{1}{2} - C(f,g)T$.



All fourteen points lie below the Chernoff bound.

Evidently, Chernoff entropy and Jensen-Shannon divergence each covary tightly with the model selection error probability.

The vertical distance between the points and the Chernoff bound is the log of $E_q\left[\min\left\{L_T^{1-\phi^*}, L_T^{-\phi^*}\right\}\right]$, which varies across pairs much less than $C(f,g)T$ does.

We'll encounter related ideas in *A Problem that Stumped Milton Friedman* very soon.

22.8 Markov chains

Let's now look at a likelihood ratio process for a sequence of random variables that is not independently and identically distributed.

Here we assume that the sequence is generated by a Markov chain on a finite state space.

We consider two n-state irreducible and aperiodic Markov chain models on the same state space $\{1, 2, \ldots, n\}$ with positive transition matrices $P^{(f)}$, $P^{(g)}$ and initial distributions $\pi_0^{(f)}$, $\pi_0^{(g)}$.

We assume that nature samples from chain f.

For a sample path $(x_0, x_1, \ldots, x_T)$, let N_{ij} count transitions from state i to j.

The likelihood process under model $m \in \{f, g\}$ is

$$L_T^{(m)} = \pi_{0,x_0}^{(m)} \prod_{i=1}^{n} \prod_{j=1}^{n} \left(P_{ij}^{(m)}\right)^{N_{ij}}$$

Hence,

$$\log L_T^{(m)} = \log \pi_{0,x_0}^{(m)} + \sum_{i,j} N_{ij} \log P_{ij}^{(m)}$$

The log-likelihood ratio is

$$\log \frac{L_T^{(f)}}{L_T^{(g)}} = \log \frac{\pi_{0,x_0}^{(f)}}{\pi_{0,x_0}^{(g)}} + \sum_{i,j} N_{ij} \log \frac{P_{ij}^{(f)}}{P_{ij}^{(g)}} \tag{22.5}$$

22.8.1 KL divergence rate

By the ergodic theorem for irreducible Markov chains, we have

$$\frac{N_{ij}}{T} \xrightarrow{a.s.} \pi_i^{(f)} P_{ij}^{(f)} \quad \text{as } T \to \infty$$

where $\pi^{(f)}$ is the stationary distribution satisfying $\pi^{(f)} = \pi^{(f)} P^{(f)}$.

Therefore,

$$\frac{1}{T} \log \frac{L_T^{(f)}}{L_T^{(g)}} = \frac{1}{T} \log \frac{\pi_{0,x_0}^{(f)}}{\pi_{0,x_0}^{(g)}} + \frac{1}{T} \sum_{i,j} N_{ij} \log \frac{P_{ij}^{(f)}}{P_{ij}^{(g)}}$$

Taking the limit as $T \to \infty$, we have:

- The first term: $\frac{1}{T} \log \frac{\pi_{0,x_0}^{(f)}}{\pi_{0,x_0}^{(g)}} \to 0$
- The second term: $\frac{1}{T} \sum_{i,j} N_{ij} \log \frac{P_{ij}^{(f)}}{P_{ij}^{(g)}} \xrightarrow{a.s.} \sum_{i,j} \pi_i^{(f)} P_{ij}^{(f)} \log \frac{P_{ij}^{(f)}}{P_{ij}^{(g)}}$

Define the **KL divergence rate** as

$$h_{KL}(f, g) = \sum_{i=1}^{n} \pi_i^{(f)} \underbrace{\sum_{j=1}^{n} P_{ij}^{(f)} \log \frac{P_{ij}^{(f)}}{P_{ij}^{(g)}}}_{=: KL(P_{i\cdot}^{(f)}, P_{i\cdot}^{(g)})}$$

where $KL(P_{i\cdot}^{(f)}, P_{i\cdot}^{(g)})$ is the row-wise KL divergence.

By the ergodic theorem, we have

$$\frac{1}{T} \log \frac{L_T^{(f)}}{L_T^{(g)}} \xrightarrow{a.s.} h_{KL}(f, g) \quad \text{as } T \to \infty$$

Taking expectations and using the dominated convergence theorem, we obtain

$$\frac{1}{T} E_f \left[\log \frac{L_T^{(f)}}{L_T^{(g)}} \right] \to h_{KL}(f, g) \quad \text{as } T \to \infty$$

Here we invite readers to pause and compare this result with (22.1).

Likelihood Processes For VAR Models extends these calculations to vector autoregressions.

Let's confirm this in the simulation below.

22.8.2 Simulations

Let's implement simulations to illustrate these concepts with a three-state Markov chain.

We start by writing functions to compute the stationary distribution and the KL divergence rate for Markov chain models.

Now let's create an example with two different 3-state Markov chains.

We are now ready to simulate paths and visualize how likelihood ratios evolve.

We draw x_0 from the stationary distribution $\pi^{(f)}$ and set $\pi_0^{(f)} = \pi^{(f)}$ and $\pi_0^{(g)} = \pi^{(g)}$.

The chain is then stationary under f, so $E_f N_{ij} = T\pi_i^{(f)} P_{ij}^{(f)}$ and (22.5) implies

$$E_f\left[\log \frac{L_T^{(f)}}{L_T^{(g)}}\right] = D_{KL}(\pi^{(f)} \| \pi^{(g)}) + T\, h_{KL}(f, g)$$

for every T.

We verify this formula by plotting both the empirical average and the line predicted by the theory

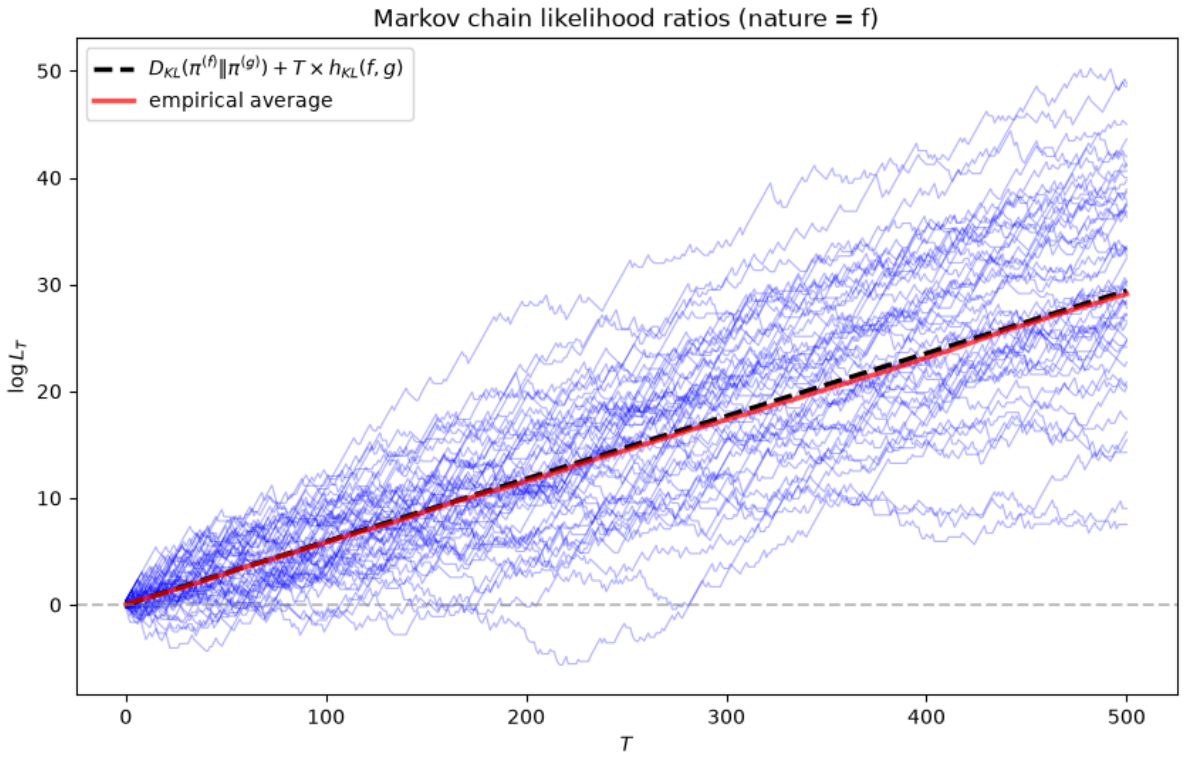
```
# Define example Markov chain transition matrices
P_f = np.array([[0.7, 0.2, 0.1],
                [0.3, 0.5, 0.2],
                [0.1, 0.3, 0.6]])

P_g = np.array([[0.5, 0.3, 0.2],
                [0.2, 0.6, 0.2],
                [0.2, 0.2, 0.6]])

markov_results = analyze_markov_chains(P_f, P_g, rng)
```

```
Stationary distribution (f): [0.41176471 0.32352941 0.26470588]
Stationary distribution (g): [0.28571429 0.38095238 0.33333333]

KL divergence rate h(f, g): 0.0588
KL divergence rate h(g, f): 0.0563
KL(π_f, π_g): 0.0366
```



22.9 Related lectures

Likelihood processes play an important role in Bayesian learning, as described in *Likelihood Ratio Processes and Bayesian Learning* and as applied in *Job Search IX: Search with Learning*.

Likelihood ratio processes are central to Lawrence Blume and David Easley's answer to their question "If you're so smart, why aren't you rich?" [Blume and Easley, 2006], the subject of the lecture *Heterogeneous Beliefs and Financial Markets*.

Likelihood ratio processes also appear in Additive and Multiplicative Functionals, which contains another illustration of the **peculiar property** of likelihood ratio processes described above.

22.10 Exercises

Exercise 22.10.1

Consider the setting where nature generates data from a third density h.

Let $\{w_t\}_{t=1}^T$ be IID draws from h, and let $L_t = L(w^t)$ be the likelihood ratio process defined as in the lecture.

Assume that K_f and K_g are finite and that $E_h|\log f(W)| < \infty$ and $E_h|\log g(W)| < \infty$.

Show that:

$$\frac{1}{t} E_h[\log L_t] = K_g - K_f$$

Hint: Start by expressing $\log L_t$ as a sum of $\log \ell(w_i)$ terms and compare with the definition of K_f and K_g.

Solution

Since $w_1, \ldots, w_t$ are IID draws from h, we can write

$$\log L_t = \log \prod_{i=1}^{t} \ell(w_i) = \sum_{i=1}^{t} \log \ell(w_i) = \sum_{i=1}^{t} \log \frac{f(w_i)}{g(w_i)}$$

Taking the expectation under h

$$E_h[\log L_t] = E_h\left[\sum_{i=1}^{t} \log \frac{f(w_i)}{g(w_i)}\right] = \sum_{i=1}^{t} E_h\left[\log \frac{f(w_i)}{g(w_i)}\right]$$

Since the w_i are identically distributed

$$E_h[\log L_t] = t \cdot E_h\left[\log \frac{f(w)}{g(w)}\right]$$

where $w \sim h$.

Therefore

$$\frac{1}{t} E_h[\log L_t] = E_h\left[\log \frac{f(w)}{g(w)}\right] = E_h[\log f(w)] - E_h[\log g(w)]$$

Now, from the definition of Kullback-Leibler divergence

$$K_f = \int h(w) \log \frac{h(w)}{f(w)} dw = E_h[\log h(w)] - E_h[\log f(w)]$$

This gives us

$$E_h[\log f(w)] = E_h[\log h(w)] - K_f$$

Similarly

$$E_h[\log g(w)] = E_h[\log h(w)] - K_g$$

Substituting back

$$\begin{aligned} \frac{1}{t} E_h[\log L_t] &= E_h[\log f(w)] - E_h[\log g(w)] \\ &= [E_h[\log h(w)] - K_f] - [E_h[\log h(w)] - K_g] \\ &= K_g - K_f \end{aligned}$$

Exercise 22.10.2

Building on Exercise 22.10.1, use the result to explain what happens to L_t as $t \to \infty$ in the following cases:

1. When $K_g > K_f$ (i.e., f is "closer" to h than g is)
2. When $K_g < K_f$ (i.e., g is "closer" to h than f is)

Relate your answer to the simulation results shown in *this section*.

Solution

From Exercise 22.10.1, we know that:

$$\frac{1}{t}E_h[\log L_t] = K_g - K_f$$

Case 1: When $K_g > K_f$

Here, f is "closer" to h than g is. Since $K_g - K_f > 0$

$$E_h[\log L_t] = t \cdot (K_g - K_f) \to +\infty \text{ as } t \to \infty$$

By the strong law of large numbers, $\frac{1}{t}\log L_t \to K_g - K_f > 0$ almost surely.

Therefore $L_t \to +\infty$ almost surely.

Case 2: When $K_g < K_f$

Here, g is "closer" to h than f is. Since $K_g - K_f < 0$

$$E_h[\log L_t] = t \cdot (K_g - K_f) \to -\infty \text{ as } t \to \infty$$

Therefore by similar reasoning $L_t \to 0$ almost surely.

Exercise 22.10.3

This exercise extends the Chernoff bound (22.4) to the Markov chains of *this section*.

Let x_0 be drawn from $\pi_0^{(f)}$ under model f and from $\pi_0^{(g)}$ under model g.

The likelihood ratio test selects f when $L_T^{(f)}/L_T^{(g)} \geq 1$ and g otherwise; let α_T and β_T be its two error probabilities.

For $\phi \in [0,1]$, let Q_ϕ be the $n \times n$ matrix with entries $Q_\phi(i,j) = (P_{ij}^{(f)})^{\phi}(P_{ij}^{(g)})^{1-\phi}$, and let $\lambda(\phi)$ be its largest eigenvalue (the Perron root).

1. Show that

$$\frac{1}{2}(\alpha_T + \beta_T) \leq \frac{1}{2}v_\phi^\top Q_\phi^T \mathbf{1},$$

   where $v_\phi(i) = (\pi_{0,i}^{(f)})^{\phi}(\pi_{0,i}^{(g)})^{1-\phi}$ and $\mathbf{1}$ is a vector of ones.

2. Use the Perron–Frobenius theorem to conclude that the error probability decays at least at the exponential rate $C_M = -\log \min_{\phi\in[0,1]} \lambda(\phi)$.

3. Compute C_M and the minimizing ϕ for the matrices $P^{(f)}$ and $P^{(g)}$ in the lecture, and compare C_M with the KL divergence rates $h_{KL}(f,g)$ and $h_{KL}(g,f)$.

4. Set $\pi_0^{(f)} = \pi^{(f)}$ and $\pi_0^{(g)} = \pi^{(g)}$. Use `compute_markov_selection_error` to simulate the error probability for several values of T, and plot it together with the bound in part 1.

Solution

Part 1. Let $p_f(x^T) = \pi_{0,x_0}^{(f)} \prod_{t=1}^{T} P_{x_{t-1}x_t}^{(f)}$ and define $p_g(x^T)$ in the same way.

As in the IID case, the error probability equals $\frac{1}{2}\sum_{x^T}\min\{p_f(x^T), p_g(x^T)\}$.

Since $\min\{x, y\} \leq x^\phi y^{1-\phi}$,

$$\frac{1}{2}(\alpha_T + \beta_T) \leq \frac{1}{2}\sum_{x_0,\ldots,x_T} v_\phi(x_0)\prod_{t=1}^{T} Q_\phi(x_{t-1}, x_t) = \frac{1}{2} v_\phi^\top Q_\phi^T \mathbf{1}.$$

Part 2. The matrix Q_ϕ has positive entries.

By the Perron–Frobenius theorem, $Q_\phi^T = \lambda(\phi)^T(r_\phi \ell_\phi^\top + o(1))$ as $T \to \infty$, where r_ϕ and ℓ_ϕ are positive right and left eigenvectors normalized so that $\ell_\phi^\top r_\phi = 1$.

Hence $\frac{1}{T}\log(v_\phi^\top Q_\phi^T \mathbf{1}) \to \log\lambda(\phi)$, and the error probability decays at least at rate $-\log\lambda(\phi)$ for every ϕ.

The best rate comes from the ϕ that minimizes $\lambda(\phi)$.

Because $P^{(f)}$ and $P^{(g)}$ are stochastic matrices, $\lambda(0) = \lambda(1) = 1$.

Large deviation theory for Markov chains implies that C_M is the exact rate [Dembo and Zeitouni, 1998].

Part 3.

```
def perron_root(ϕ, P_f, P_g):
    """Largest eigenvalue of the matrix with entries P_f^ϕ P_g^(1-ϕ)."""
    Q = P_f**ϕ * P_g**(1 - ϕ)
    return np.max(np.abs(np.linalg.eigvals(Q)))

res = minimize_scalar(perron_root, bounds=(0, 1), args=(P_f, P_g),
                      method='bounded')
C_M, ϕ_M = -np.log(res.fun), res.x

print(f"C_M = {C_M:.5f} at ϕ = {ϕ_M:.4f}")
print(f"h_KL(f, g) = {markov_results['kl_rate_fg']:.4f}, "
      f"h_KL(g, f) = {markov_results['kl_rate_gf']:.4f}")
```

```
C_M = 0.01427 at ϕ = 0.5055
h_KL(f, g) = 0.0588, h_KL(g, f) = 0.0563
```

The Chernoff rate $C_M = 0.0143$ is about a quarter of either KL divergence rate.

The error probability therefore falls much more slowly than $\log L_T^{(f)}/L_T^{(g)}$ drifts.

Part 4.

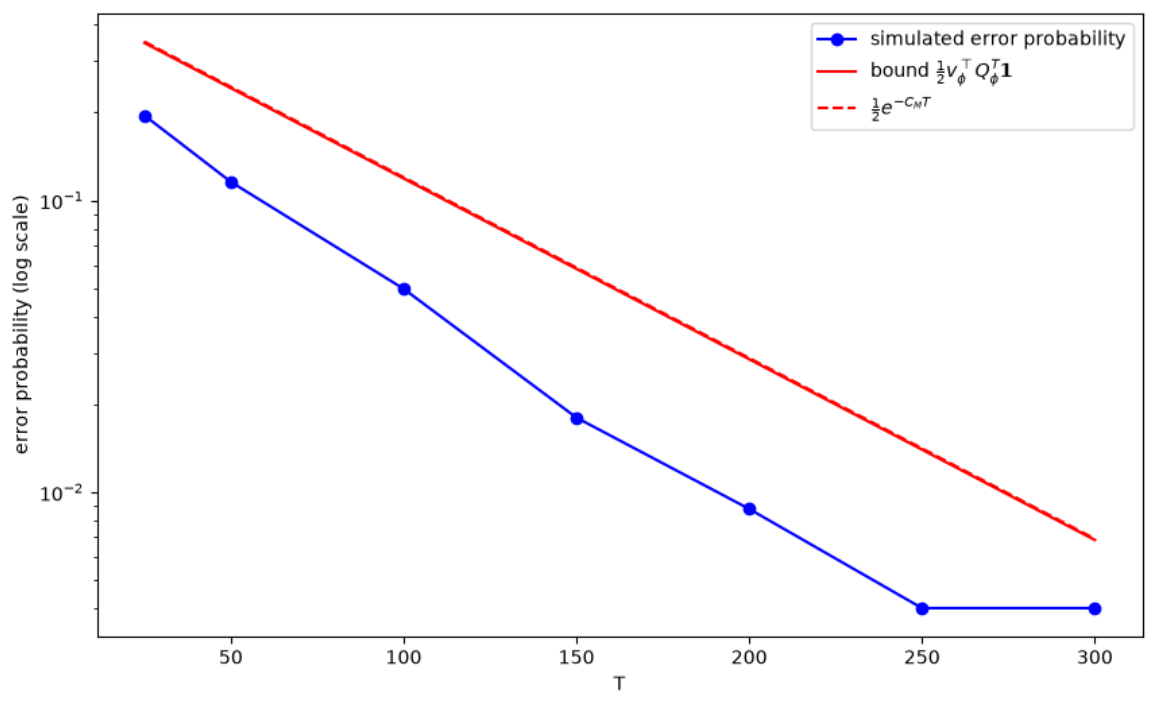
```
π_f = markov_results['stationary_f']
π_g = markov_results['stationary_g']
T_values = np.array([25, 50, 100, 150, 200, 250, 300])

errors = compute_markov_selection_error(T_values, P_f, P_g, π_f, π_g,
                                        rng, N_sim=4000)

Q_M = P_f**ϕ_M * P_g**(1 - ϕ_M)
v_M = π_f**ϕ_M * π_g**(1 - ϕ_M)
bound = [0.5 * v_M @ np.linalg.matrix_power(Q_M, T) @ np.ones(3)
         for T in T_values]

plt.figure(figsize=(10, 6))
plt.semilogy(T_values, errors, 'bo-', label='simulated error probability')
plt.semilogy(T_values, bound, 'r-',
             label=r'bound $\frac{1}{2} v_\phi^\top Q_\phi^T \mathbf{1}$')
plt.semilogy(T_values, 0.5 * np.exp(-C_M * T_values), 'r--',
             label=r'$\frac{1}{2} e^{-C_M T}$')
```

```
plt.xlabel('T')
plt.ylabel('error probability (log scale)')
plt.legend()
plt.show()
```



The bound from part 1 nearly coincides with $\frac{1}{2}e^{-C_M T}$.

The simulated error probabilities lie below the bound.

On the log scale they decline roughly in parallel with the bound, a little faster because of a polynomial factor that the bound omits.

Exercise 22.10.4

Let $f = \text{Beta}(1,1)$ and $g = \text{Beta}(G_a, G_b)$.

1. Show that

$$E\left[\ell(w_t) \mid q = f\right] = \int_0^1 \frac{f(w)^2}{g(w)} dw = B(G_a, G_b)\, B(2 - G_a, 2 - G_b)$$

   when $G_a < 2$ and $G_b < 2$, and that $E\left[\ell(w_t) \mid q = f\right] = +\infty$ otherwise.

2. Compute $E\left[\ell(w_t) \mid q = f\right]$ for $g = \text{Beta}(1.5, 1.2)$ and for the lecture's $g = \text{Beta}(3, 1.2)$.

3. For each g, draw $w_1, w_2, \ldots, w_N$ from f and plot the running sample average $\frac{1}{n}\sum_{i=1}^n \ell(w_i)$ against n. Explain the difference between the two plots.

Solution

Part 1. Since $f = 1$ on $(0, 1)$,

$$\int_0^1 \frac{f(w)^2}{g(w)} dw = B(G_a, G_b) \int_0^1 w^{1-G_a}(1-w)^{1-G_b} dw.$$

The integral converges if and only if $1 - G_a > -1$ and $1 - G_b > -1$, that is, if and only if $G_a < 2$ and $G_b < 2$. When it converges, it equals $B(2 - G_a, 2 - G_b)$.

Part 2.

```
def mean_ell_under_f(G_a, G_b):
    """E[ℓ(w) | q=f] for f = Beta(1, 1) and g = Beta(G_a, G_b)."""
    if G_a < 2 and G_b < 2:
        return np.exp(betaln(G_a, G_b) + betaln(2 - G_a, 2 - G_b))
    return np.inf

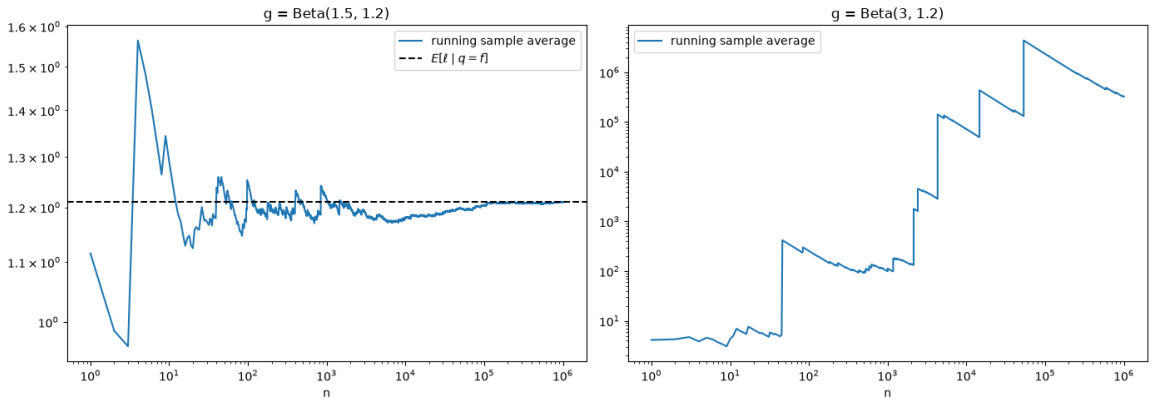
for G_params in [(1.5, 1.2), (3, 1.2)]:
    print(f"g = Beta{G_params}: E[ℓ | f] = {mean_ell_under_f(*G_params):.4f}")
```

```
g = Beta(1.5, 1.2): E[ℓ | f] = 1.2112
g = Beta(3, 1.2): E[ℓ | f] = inf
```

Part 3.

```
N = 1_000_000
n_grid = np.arange(1, N+1)

fig, axes = plt.subplots(1, 2, figsize=(14, 5))
for ax, (g_a, g_b) in zip(axes, [(1.5, 1.2), (3, 1.2)]):
    w = rng.uniform(size=N)
    ℓ = np.exp(-log_beta_pdf(w, g_a, g_b))
    ax.plot(n_grid, np.cumsum(ℓ) / n_grid, lw=1.5,
            label='running sample average')
    mean = mean_ell_under_f(g_a, g_b)
    if np.isfinite(mean):
        ax.axhline(mean, color='k', ls='--', label=r'$E[\ell \mid q=f]$')
    ax.set_xscale('log')
    ax.set_yscale('log')
    ax.set_xlabel('n')
    ax.set_title(f'g = Beta({g_a}, {g_b})')
    ax.legend()
plt.tight_layout()
plt.show()
```



For $g = \text{Beta}(1.5, 1.2)$, the mean of ℓ under f is finite and the strong law of large numbers drives the running average to 1.2112.

For $g = \text{Beta}(3, 1.2)$, the mean is infinite.

Rare draws near $w = 0$ produce enormous values of $\ell = 1/g$, and the running average keeps jumping upward without settling.

For this reason we plotted $\log L_t$ in *the section in which nature draws from f*.

CHAPTER

TWENTYTHREE

HETEROGENEOUS BELIEFS AND FINANCIAL MARKETS

Contents

- *Heterogeneous Beliefs and Financial Markets*
  - *Overview*
  - *Review: likelihood ratio processes*
  - *Blume and Easley's setting*
  - *Nature and agents' beliefs*
  - *A socialist risk-sharing arrangement*
  - *Social planner's allocation problem*
  - *If you're so smart, …*
  - *Competitive equilibrium prices*
  - *Simulations*
  - *Related lectures*
  - *Exercises*

23.1 Overview

A likelihood ratio process lies behind Lawrence Blume and David Easley's answer to their question "If you're so smart, why aren't you rich?" [Blume and Easley, 2006].

Blume and Easley constructed formal models to study how differences of opinions about probabilities governing risky income processes would influence outcomes and be reflected in prices of stocks, bonds, and insurance policies that individuals use to share and hedge risks.

Note

[Alchian, 1950] and [Friedman, 1953] conjectured that, by rewarding traders with more realistic probability models, competitive markets in financial securities put wealth in the hands of better informed traders and help make prices of risky assets reflect realistic probability assessments.

Here we'll provide an example that illustrates basic components of Blume and Easley's analysis.

Our main tools are the likelihood ratio processes and Kullback–Leibler divergences studied in *Likelihood Ratio Processes*.

We'll focus only on their analysis of an environment with complete markets in which trades in all conceivable risky securities are possible.

We'll study two alternative arrangements:

- perfect socialism in which individuals surrender their endowments of consumption goods each period to a central planner who then dictatorially allocates those goods
- a decentralized system of competitive markets in which selfish price-taking individuals voluntarily trade with each other in competitive markets

The fundamental theorems of welfare economics will apply and assure us that these two arrangements end up producing exactly the same allocation of consumption goods to individuals **provided** that the social planner assigns an appropriate set of **Pareto weights**.

Note

You can learn about how the two welfare theorems are applied in modern macroeconomic models in *this lecture on a planning problem* and *this lecture on a related competitive equilibrium*. *This quantecon lecture* presents a recursive formulation of complete markets models with homogeneous beliefs.

Let's start by importing some Python tools.

```
import matplotlib.pyplot as plt
import numpy as np
from numba import vectorize, jit
from math import gamma
from scipy.integrate import quad
from scipy.special import expit

rng = np.random.default_rng()
```

23.2 Review: likelihood ratio processes

Likelihood Ratio Processes studies likelihood ratio processes in detail.

We recall the facts that we use here.

A nonnegative random variable W has one of two probability density functions, either f or g.

We assume that f and g both put positive probabilities on the same intervals of possible realizations of the random variable W.

For a sequence $\{w_t\}_{t=1}^{\infty}$ of IID draws, the time t likelihood ratio is

$$\ell(w_t) = \frac{f(w_t)}{g(w_t)}, \quad t \geq 1,$$

and the **likelihood ratio process** is

$$L(w^t) = \prod_{i=1}^{t} \ell(w_i),$$

where $w^t = \{w_1, \ldots, w_t\}$ is a history of observations up to and including time t.

We write $L_t = L(w^t)$.

The likelihood ratio process satisfies the recursion $L(w^t) = \ell(w_t) L(w^{t-1})$.

Suppose that a density h generates the data.

Let $K_f = KL(h, f)$ and $K_g = KL(h, g)$ be the Kullback–Leibler divergences of f and g from h.

The strong law of large numbers applied to $\log L_t = \sum_{i=1}^{t} \log \ell(w_i)$ implies

$$\frac{1}{t} \log L_t \to K_g - K_f \quad \text{almost surely under } h.$$

A helpful formula in *Likelihood Ratio Processes* derives the corresponding formula for $\frac{1}{t} E_h[\log L_t]$.

When $K_f < K_g$, $L_t \to +\infty$; when $K_f > K_g$, $L_t \to 0$.

When $h = f$, the drift of $\log L_t$ is $KL(f, g) > 0$; when $h = g$, it is $-KL(g, f) < 0$.

The following Python code defines f and g as two Beta densities.

```
# Parameters in the two Beta distributions.
F_a, F_b = 1, 1
G_a, G_b = 3, 1.2

@vectorize
def p(x, a, b):
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x** (a-1) * (1 - x) ** (b-1)

# The two density functions.
f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))
```

23.3 Blume and Easley's setting

Let the random variable $s_t \in (0, 1)$ at time $t = 0, 1, 2, \ldots$ be distributed according to the same Beta distribution with parameters $\theta = \{\theta_1, \theta_2\}$.

We'll denote this probability density as

$$\pi(s_t | \theta)$$

Below, we'll often just write $\pi(s_t)$ instead of $\pi(s_t|\theta)$ to save space.

Let $s_t \equiv y_t^1$ be the endowment of a nonstorable consumption good that a person we'll call "agent 1" receives at time t.

Let a history $s^t = [s_t, s_{t-1}, \ldots, s_0]$ be a sequence of i.i.d. random variables with joint distribution

$$\pi_t(s^t) = \pi(s_t)\pi(s_{t-1}) \cdots \pi(s_0)$$

So in our example, the history s^t is a comprehensive record of agent 1's endowments of the consumption good from time 0 up to time t.

If agent 1 were to live on an island by himself, agent 1's consumption $c^1(s_t)$ at time t is

$$c^1(s_t) = y_t^1 = s_t.$$

But in our model, agent 1 is not alone.

23.4 Nature and agents' beliefs

Nature draws i.i.d. sequences $\{s_t\}_{t=0}^{\infty}$ from $\pi_t(s^t)$.

- so π without a superscript is nature's model
- but in addition to nature, there are other entities inside our model – artificial people that we call "agents"
- each agent has a sequence of probability distributions over s^t for $t = 0, \ldots$
- agent i thinks that nature draws i.i.d. sequences $\{s_t\}_{t=0}^{\infty}$ from $\{\pi_t^i(s^t)\}_{t=0}^{\infty}$
  - agent i is mistaken unless $\pi_t^i(s^t) = \pi_t(s^t)$

Note

A **rational expectations** model would set $\pi_t^i(s^t) = \pi_t(s^t)$ for all agents i.

There are two agents named $i = 1$ and $i = 2$.

At time t, agent 1 receives an endowment

$$y_t^1 = s_t$$

of a nonstorable consumption good, while agent 2 receives an endowment of

$$y_t^2 = 1 - s_t$$

The aggregate endowment of the consumption good is

$$y_t^1 + y_t^2 = 1$$

at each date $t \geq 0$.

At date t agent i consumes $c_t^i(s^t)$ of the good.

A (non wasteful) feasible allocation of the aggregate endowment of 1 each period satisfies

$$c_t^1 + c_t^2 = 1.$$

23.5 A socialist risk-sharing arrangement

In order to share risks, a benevolent social planner dictates a history-dependent consumption allocation that takes the form of a sequence of functions

$$c_t^i = c_t^i(s^t)$$

that satisfy

$$c_t^1(s^t) + c_t^2(s^t) = 1 \tag{23.1}$$

for all s^t for all $t \geq 0$.

To design a socially optimal allocation, the social planner wants to know what each agent i believes about the endowment sequence and how they feel about bearing risks.

As for the endowment sequences, agent i believes that nature draws i.i.d. sequences from joint densities

$$\pi_t^i(s^t) = \pi^i(s_t)\pi^i(s_{t-1})\cdots\pi^i(s_0)$$

As for attitudes toward bearing risks, agent i has a one-period utility function

$$u(c_t^i) = \ln(c_t^i)$$

with marginal utility of consumption in period t

$$u'(c_t^i) = \frac{1}{c_t^i}$$

Putting its beliefs about its random endowment sequence and its attitudes toward bearing risks together, agent i has intertemporal utility function

$$V^i = \sum_{t=0}^{\infty}\sum_{s^t}\delta^t u(c_t^i(s^t))\pi_t^i(s^t), \tag{23.2}$$

where $\delta \in (0,1)$ is an intertemporal discount factor, and $u(\cdot)$ is a strictly increasing, concave one-period utility function.

Because s_t has a density, $\sum_{s^t}$ stands for an integral $\int\cdots ds^t$ over histories s^t, here and below.

23.6 Social planner's allocation problem

The benevolent dictator has all the information it requires to choose a consumption allocation that maximizes the social welfare criterion

$$W = \lambda V^1 + (1-\lambda)V^2 \tag{23.3}$$

where $\lambda \in [0,1]$ is a Pareto weight that tells how much the planner likes agent 1 and $1-\lambda$ is a Pareto weight that tells how much the social planner likes agent 2.

Setting $\lambda = .5$ expresses "egalitarian" social preferences.

Notice how social welfare criterion (23.3) takes into account both agents' preferences as represented by formula (23.2).

This means that the social planner knows and respects

- each agent's one period utility function $u(\cdot) = \ln(\cdot)$
- each agent i's probability model $\{\pi_t^i(s^t)\}_{t=0}^{\infty}$

Consequently, it is natural to anticipate that these objects will appear in the social planner's rule for allocating the aggregate endowment each period.

Let $\theta_t(s^t)$ be a Lagrange multiplier on feasibility constraint (23.1) at history s^t.

First-order necessary conditions for maximizing welfare criterion (23.3) subject to the feasibility constraint (23.1) are

$$\lambda\delta^t\frac{\pi_t^1(s^t)}{c_t^1(s^t)} = \theta_t(s^t), \qquad (1-\lambda)\delta^t\frac{\pi_t^2(s^t)}{c_t^2(s^t)} = \theta_t(s^t). \tag{23.4}$$

Dividing the second condition by the first gives

$$\frac{\pi_t^2(s^t)}{\pi_t^1(s^t)}\frac{(1/c_t^2(s^t))}{(1/c_t^1(s^t))} = \frac{\lambda}{1-\lambda}$$

which can be rearranged to become

$$\frac{c_t^1(s^t)}{c_t^2(s^t)} = \frac{\lambda}{1-\lambda} l_t(s^t) \tag{23.5}$$

where

$$l_t(s^t) = \frac{\pi_t^1(s^t)}{\pi_t^2(s^t)}$$

is the likelihood ratio of agent 1's joint density to agent 2's joint density.

The cumulative likelihood ratio $l_t(s^t)$ is the counterpart of L_t in the review above, with $f = \pi^1$, $g = \pi^2$, and histories that start at date 0.

The one-period likelihood ratio is $\pi^1(s_t)/\pi^2(s_t)$.

Using

$$c_t^1(s^t) + c_t^2(s^t) = 1$$

we can rewrite allocation rule (23.5) as

$$\frac{c_t^1(s^t)}{1 - c_t^1(s^t)} = \frac{\lambda}{1-\lambda} l_t(s^t)$$

or

$$c_t^1(s^t) = \frac{\lambda}{1-\lambda} l_t(s^t)(1 - c_t^1(s^t))$$

which implies that the social planner's allocation rule is

$$c_t^1(s^t) = \frac{\lambda l_t(s^t)}{1 - \lambda + \lambda l_t(s^t)} \tag{23.6}$$

If we define a temporary or **continuation Pareto weight** process as

$$\lambda_t(s^t) = \frac{\lambda l_t(s^t)}{1 - \lambda + \lambda l_t(s^t)},$$

then we can represent the social planner's allocation rule as

$$c_t^1(s^t) = \lambda_t(s^t)$$

and of course

$$c_t^2(s^t) = 1 - \lambda_t(s^t).$$

Substituting $c_t^1(s^t) = \lambda_t(s^t)$ into the first condition in (23.4) gives

$$\theta_t(s^t) = \delta^t \left[\lambda \pi_t^1(s^t) + (1-\lambda)\pi_t^2(s^t)\right]. \tag{23.7}$$

The multiplier $\theta_t(s^t)$ is the planner's **shadow price** of one unit of consumption at date t after history s^t, measured in units of social welfare W.

23.7 If you're so smart, …

Let's compute some values of limiting allocations (23.6) for some interesting possible limiting values of the likelihood ratio process $l_t(s^t)$.

As our first case, let's suppose that the two agents have identical beliefs, $\pi^1 = \pi^2$, so that

$$l_t(s^t) = 1 \text{ for all } t; \quad c_t^1 = \lambda$$

- In this case, the consumption allocation stays put at a $\lambda, 1-\lambda$ split between the two agents.

As our second case, let's suppose that

$$l_\infty(s^\infty) = 0; \quad c_\infty^1 = 0$$

- In this case, agent 2 is "smarter" than agent 1, and agent 1's share of the aggregate endowment converges to zero.

As our third case, let's suppose that

$$l_\infty(s^\infty) = \infty; \quad c_\infty^1 = 1$$

- In this case, agent 1 is smarter than agent 2, and agent 1's share of the aggregate endowment converges to 1.

Kullback–Leibler divergences of the agents' beliefs from nature's density determine which case occurs.

Let

$$KL(\pi, \pi^i) = \int \pi(s) \log \frac{\pi(s)}{\pi^i(s)} \, ds.$$

When nature draws s_t IID from π, $\log l_t(s^t) = \sum_{j=0}^t \log[\pi^1(s_j)/\pi^2(s_j)]$ is a sum of IID random variables with mean

$$E_\pi \left[\log \frac{\pi^1(s)}{\pi^2(s)} \right] = KL(\pi, \pi^2) - KL(\pi, \pi^1).$$

The strong law of large numbers implies

$$\frac{1}{t} \log l_t(s^t) \to KL(\pi, \pi^2) - KL(\pi, \pi^1) \quad \text{almost surely.} \tag{23.8}$$

Formula (23.8) is the formula in *A helpful formula* of *Likelihood Ratio Processes* with $h = \pi$, $f = \pi^1$, and $g = \pi^2$.

It implies:

- if $KL(\pi, \pi^1) < KL(\pi, \pi^2)$, then $l_t \to +\infty$ and $c_t^1 \to 1$
- if $KL(\pi, \pi^1) > KL(\pi, \pi^2)$, then $l_t \to 0$ and $c_t^1 \to 0$
- if $KL(\pi, \pi^1) = KL(\pi, \pi^2)$ and $\pi^1 \neq \pi^2$, then $\log l_t$ is a random walk with zero drift and positive variance; its lim sup is $+\infty$ and its lim inf is $-\infty$, so c_t^1 swings between values near 0 and values near 1 and never settles down

Exercise 23.11.3 studies the third case.

"Smarter" means closer to nature's π in Kullback–Leibler divergence.

An agent with correct beliefs, $\pi^i = \pi$, has $KL(\pi, \pi^i) = 0$ and drives out any agent whose beliefs differ.

Sandroni [2000] and Blume and Easley [2006] prove market selection results of this kind for more general economies.

Note

These three cases are somehow telling us about how relative **wealths** of the agents evolve as time passes.

- when the two agents have identical beliefs and $\lambda \in (0, 1)$, agent 1's wealth share stays at λ perpetually.
- when agent 1 is smarter and $\lambda \in (0, 1)$, agent 1 eventually "owns" the entire continuation endowment and agent 2 eventually "owns" nothing.
- when agent 2 is smarter and $\lambda \in (0, 1)$, agent 2 eventually "owns" the entire continuation endowment and agent 1 eventually "owns" nothing.

We define continuation wealths precisely after we construct a competitive equilibrium **price** system below.

Soon we'll do some simulations that will shed further light on possible outcomes.

But before we do that, let's take a detour and study the "shadow prices" $\theta_t(s^t)$ in (23.7) for the social planning problem.

We shall soon see that these shadow prices can readily be converted to "equilibrium prices" for a competitive equilibrium.

Doing this will allow us to connect our analysis with claims by [Alchian, 1950] and [Friedman, 1953] that transfers of wealth coming from trades in competitive asset markets eventually make prices of risky assets reflect realistic probability assessments.

23.8 Competitive equilibrium prices

Two fundamental welfare theorems for general equilibrium models lead us to anticipate that there is a connection between the allocation that solves the social planning problem we have been studying and the allocation in a **competitive equilibrium** with complete markets in contingent claims to time- and history-dependent consumption goods.

Note

For the two welfare theorems and their history, see https://en.wikipedia.org/wiki/Fundamental_theorems_of_welfare_economics. Again, for applications to a classic macroeconomic growth model, see *this lecture on a planning problem* and *this lecture on a related competitive equilibrium*

Such a connection prevails for our model.

We'll sketch it now.

In a competitive equilibrium, no social planner dictatorially collects everybody's endowments and then reallocates them.

Instead, there is a comprehensive centralized market that meets at one point in time.

There are **prices** at which price-taking agents can buy or sell whatever goods that they want.

Trade is multilateral in the sense that a "Walrasian auctioneer" who lives outside the model finds prices at which supply equals demand for every good.

Each agent faces a budget constraint.

That budget constraint requires that the total value of the agent's endowment stream be at least as large as the total value of its consumption stream.

These values are computed at price vectors that the agents take as given – the agents are "price-takers" who assume that they can buy or sell whatever quantities that they want at those prices.

Suppose that at time -1, before time 0 starts, agent i can purchase one unit $c_t(s^t)$ of consumption at time t after history s^t at price $p_t(s^t)$.

The price **vector** is very long:

- there is one price $p_t(s^t)$ for each history s^t at every date $t = 0, 1, \ldots,$.
- so there are as many prices as there are histories and dates.

The auctioneer sets these prices at time -1, before the economy starts.

The market meets once at time -1.

At times $t = 0, 1, 2, \ldots$ trades made at time -1 are simply executed, i.e., the promised deliveries are made.

- in the background, there is an "enforcement" procedure that forces agents to carry out the exchanges or "deliveries" that they agreed to at time -1.

We want to study how agents' probability models influence equilibrium prices.

Agent i faces a **single** intertemporal budget constraint

$$\sum_{t=0}^{\infty} \sum_{s^t} p_t(s^t) c_t^i(s^t) \leq \sum_{t=0}^{\infty} \sum_{s^t} p_t(s^t) y_t^i(s^t) \tag{23.9}$$

Because s_t has a density, $p_t(s^t)$ is a price density: $p_t(s^t)\, ds^t$ is the time -1 price of one unit of consumption at date t delivered if the history lands in a small set ds^t around s^t.

According to budget constraint (23.9), trade is **multilateral** in the following sense

- imagine that agent i first sells his random endowment stream $\{y_t^i(s^t)\}$ and then uses the proceeds (i.e., his "wealth") to purchase a random consumption stream $\{c_t^i(s^t)\}$.

Agent i attaches a Lagrange multiplier μ_i to budget constraint (23.9) and once-and-for-all chooses a consumption plan $\{c_t^i(s^t)\}_{t=0}^{\infty}$ that maximizes criterion (23.2) subject to budget constraint (23.9).

This means that the agent i chooses many objects, namely, $c_t^i(s^t)$ for all s^t for $t = 0, 1, 2, \ldots$.

For convenience, let's remind ourselves of criterion V^i defined in (23.2):

$$V^i = \sum_{t=0}^{\infty} \sum_{s^t} \delta^t u(c_t^i(s^t)) \pi_t^i(s^t)$$

First-order necessary conditions for maximizing objective V^i defined in (23.2) with respect to $c_t^i(s^t)$ are

$$\delta^t u'(c_t^i(s^t)) \pi_t^i(s^t) = \mu_i p_t(s^t),$$

which we can rearrange to obtain

$$p_t(s^t) = \frac{\delta^t \pi_t^i(s^t)}{\mu_i c_t^i(s^t)} \tag{23.10}$$

for $i = 1, 2$.

If we divide equation (23.10) for agent 1 by the appropriate version of equation (23.10) for agent 2, impose the feasibility condition $c_t^2(s^t) = 1 - c_t^1(s^t)$, and do some algebra, we'll obtain

$$c_t^1(s^t) = \frac{\mu_2 l_t(s^t)}{\mu_1 + \mu_2 l_t(s^t)}. \tag{23.11}$$

We now embark on an extended "guess-and-verify" exercise that involves matching objects in our competitive equilibrium with objects in our social planning problem.

- we'll match consumption allocations in the planning problem with equilibrium consumption allocations in the competitive equilibrium
- we'll match "shadow" prices in the planning problem with competitive equilibrium prices.

Formula (23.11) has the same form as formula (23.6) with $\lambda/(1-\lambda) = \mu_2/\mu_1$.

Setting $\mu_1 = 1-\lambda$ and $\mu_2 = \lambda$ makes the two formulas agree.

Only the ratio μ_2/μ_1 matters for the allocation.

- multiplying μ_1 and μ_2 by a positive constant and dividing all prices $p_t(s^t)$ by the same constant leaves budget constraints and first-order conditions unchanged
- choosing that constant amounts to choosing a **numeraire** or normalization for the price system $\{p_t(s^t)\}_{t=0}^{\infty}$

Note

For information about how a numeraire must be chosen to pin down the absolute price level in a model like ours that determines only relative prices, see https://en.wikipedia.org/wiki/Numéraire.

The ratio μ_2/μ_1 is not a normalization.

Budget constraints (23.9) pin it down, and with it the Pareto weight λ, as we show below.

If we set $\mu_1 = 1-\lambda$ and $\mu_2 = \lambda$, substitute formula (23.11) for $c_t^1(s^t)$ into formula (23.10), and rearrange, we obtain

$$p_t(s^t) = \frac{\delta^t}{\lambda(1-\lambda)}\pi_t^2(s^t)[1-\lambda+\lambda l_t(s^t)]$$

or

$$p_t(s^t) = \frac{\delta^t}{\lambda(1-\lambda)}[(1-\lambda)\pi_t^2(s^t) + \lambda\pi_t^1(s^t)] \tag{23.12}$$

Comparing (23.12) with (23.7) shows that $p_t(s^t) = \theta_t(s^t)/(\lambda(1-\lambda))$.

Competitive equilibrium prices are proportional to the planner's shadow prices.

The normalization $\mu_1 = 1/\lambda$ and $\mu_2 = 1/(1-\lambda)$ makes them equal: $p_t(s^t) = \theta_t(s^t)$.

According to formula (23.12), we have the following possible limiting cases:

- when $l_\infty = 0$, $c_\infty^1 = 0$ and tails of competitive equilibrium prices reflect agent 2's probability model $\pi_t^2(s^t)$ according to $p_t(s^t) \propto \delta^t\pi_t^2(s^t)$
- when $l_\infty = \infty$, $c_\infty^1 = 1$ and tails of competitive equilibrium prices reflect agent 1's probability model $\pi_t^1(s^t)$ according to $p_t(s^t) \propto \delta^t\pi_t^1(s^t)$
- for small t's, competitive equilibrium prices reflect both agents' probability models.

23.8.1 The competitive equilibrium Pareto weight

Let $E^i[s] = \int s\,\pi^i(s)\,ds$ be agent i's expected value of s_t.

Substituting allocation (23.6) and prices (23.12) into agent 1's budget constraint (23.9) with $y_t^1 = s_t$ and imposing equality gives

$$\lambda = \lambda E^1[s] + (1-\lambda)E^2[s],$$

or

$$\lambda = \frac{E^2[s]}{1 - E^1[s] + E^2[s]}. \tag{23.13}$$

Exercise 23.11.1 asks you to verify this.

Agent 2's budget constraint then holds by Walras' law.

Finding Pareto weights at which the planner's allocation satisfies every agent's budget constraint is the method of Negishi [1960].

The left side of the first equation is agent 1's share of total wealth at time -1.

The right side is the market value of agent 1's endowment stream divided by the market value of the aggregate endowment stream.

For $\pi^1 = f = \text{Beta}(1, 1)$ and $\pi^2 = g = \text{Beta}(3, 1.2)$, $E^1[s] = 1/2$, $E^2[s] = 5/7$, and $\lambda = 10/17 \approx 0.588$.

Agent 2 expects agent 1's endowment to be high.

Agent 2's beliefs raise the price of claims on agent 1's endowment and make agent 1 rich.

The simulations below set $\lambda = 0.5$.

By the second welfare theorem, that allocation is a competitive equilibrium after a time -1 transfer of wealth from agent 1 to agent 2.

23.8.2 Wealth

Agent i's **continuation wealth** at history s^t is the value of its continuation consumption plan in units of date t consumption:

$$W_t^i(s^t) = \frac{1}{p_t(s^t)} \sum_{\tau=t}^{\infty} \int p_\tau(s^\tau) c_\tau^i(s^\tau)\, ds_{t+1} \cdots ds_\tau.$$

With logarithmic utility, first-order condition (23.10) implies

$$W_t^i(s^t) = \frac{c_t^i(s^t)}{1 - \delta}. \tag{23.14}$$

Total wealth $W_t^1(s^t) + W_t^2(s^t) = 1/(1 - \delta)$ is the value of a claim to the aggregate endowment.

Agent 1's wealth share equals its consumption share $c_t^1(s^t) = \lambda_t(s^t)$.

The limiting cases above describe wealth shares as well as consumption shares.

23.8.3 Arrow prices and market selection

Let $q(s_{t+1}|s^t) = p_{t+1}(s^{t+1})/p_t(s^t)$ be the Arrow price density at history s^t of claims to consumption at $t+1$ contingent on s_{t+1}, measured in units of date t consumption.

Formula (23.12) and $\pi_{t+1}^i(s^{t+1}) = \pi^i(s_{t+1})\pi_t^i(s^t)$ imply

$$q(s_{t+1}|s^t) = \delta\big[\lambda_t(s^t)\pi^1(s_{t+1}) + (1 - \lambda_t(s^t))\pi^2(s_{t+1})\big]. \tag{23.15}$$

Arrow prices equal δ times a mixture of the two agents' one-period densities.

Wealth shares are the mixing weights.

Integrating (23.15) over s_{t+1} shows that a one-period riskless bond costs δ at every history.

The continuation Pareto weight

$$\lambda_t(s^t) = \frac{\lambda l_t(s^t)}{\lambda l_t(s^t) + 1 - \lambda}$$

is the posterior probability on model π^1 of a Bayesian who puts prior probability λ on model π^1 and $1-\lambda$ on model π^2 and then observes s^t, as in *Likelihood Ratio Processes and Bayesian Learning*.

The bracketed term in (23.15) is that Bayesian's predictive density of s_{t+1}.

Competitive markets price risks as if such a Bayesian set prices.

Formula (23.15) makes the Alchian–Friedman conjecture precise.

When agent 1's beliefs are correct, $\pi^1 = \pi$, formula (23.8) implies $\lambda_t \to 1$.

Arrow prices then converge to $\delta\pi(s_{t+1})$ and reflect true probabilities.

When neither agent's beliefs are correct and $KL(\pi, \pi^1) \neq KL(\pi, \pi^2)$, Arrow prices converge to δ times the one-period density of the agent whose beliefs are closer to π in Kullback–Leibler divergence.

Markets then select the less wrong model, and prices do not converge to the truth.

23.9 Simulations

Now let's implement some simulations when agent 1 believes marginal density

$$\pi^1(s_t) = f(s_t)$$

and agent 2 believes marginal density

$$\pi^2(s_t) = g(s_t)$$

where f and g are Beta distributions like ones that we used in earlier sections of this lecture.

Meanwhile, we'll assume that nature believes a marginal density

$$\pi(s_t) = h(s_t)$$

where $h(s_t)$ is perhaps a mixture of f and g.

First, we write a function to compute the likelihood ratio process

```
def compute_likelihood_ratios(sequences, f, g):
    """Compute likelihood ratios and cumulative products."""
    l_ratios = f(sequences) / g(sequences)
    L_cumulative = np.cumprod(l_ratios, axis=1)
    return l_ratios, L_cumulative
```

Let's compute the Kullback–Leibler discrepancies by quadrature integration.

```
def compute_KL(f, g):
    """
    Compute KL divergence KL(f, g)
    """
    integrand = lambda w: f(w) * np.log(f(w) / g(w))
    val, _ = quad(integrand, 1e-5, 1-1e-5)
    return val
```

We also create a helper function to compute KL divergence with respect to a reference distribution h

```
def compute_KL_h(h, f, g):
    """
    Compute KL divergence with reference distribution h
    """

    Kf = compute_KL(h, f)
    Kg = compute_KL(h, g)

    return Kf, Kg
```

Let's write a Python function that computes agent 1's consumption share

```
def simulate_blume_easley(sequences, f_belief=f, g_belief=g, λ=0.5):
    """Simulate Blume-Easley model consumption shares."""
    l_ratios, l_cumulative = compute_likelihood_ratios(sequences, f_belief, g_belief)
    c1_share = λ * l_cumulative / (1 - λ + λ * l_cumulative)
    return l_cumulative, c1_share
```

Now let's use this function to generate sequences in which

- nature draws from f each period, or
- nature draws from g each period, or
- nature flips a fair coin each period to decide whether to draw from f or g

```
λ = 0.5
T = 100
N = 10000

# Nature follows f, g, or mixture
s_seq_f = rng.beta(F_a, F_b, (N, T))
s_seq_g = rng.beta(G_a, G_b, (N, T))

h = jit(lambda x: 0.5 * f(x) + 0.5 * g(x))
model_choices = rng.random((N, T)) < 0.5
s_seq_h = np.empty((N, T))
s_seq_h[model_choices] = rng.beta(F_a, F_b, size=model_choices.sum())
s_seq_h[~model_choices] = rng.beta(G_a, G_b, size=(~model_choices).sum())

l_cum_f, c1_f = simulate_blume_easley(s_seq_f)
l_cum_g, c1_g = simulate_blume_easley(s_seq_g)
l_cum_h, c1_h = simulate_blume_easley(s_seq_h)
```

Before looking at the figure below, have some fun by guessing whether agent 1 or agent 2 will have a larger and larger consumption share as time passes in our three cases.

To make better guesses, let's visualize instances of the likelihood ratio processes in the three cases.

```
fig, axes = plt.subplots(2, 3, figsize=(15, 10))

titles = ["Nature = f", "Nature = g", "Nature = mixture"]
data_pairs = [(l_cum_f, c1_f), (l_cum_g, c1_g), (l_cum_h, c1_h)]

for i, ((l_cum, c1), title) in enumerate(zip(data_pairs, titles)):
    # Likelihood ratios
    ax = axes[0, i]
```

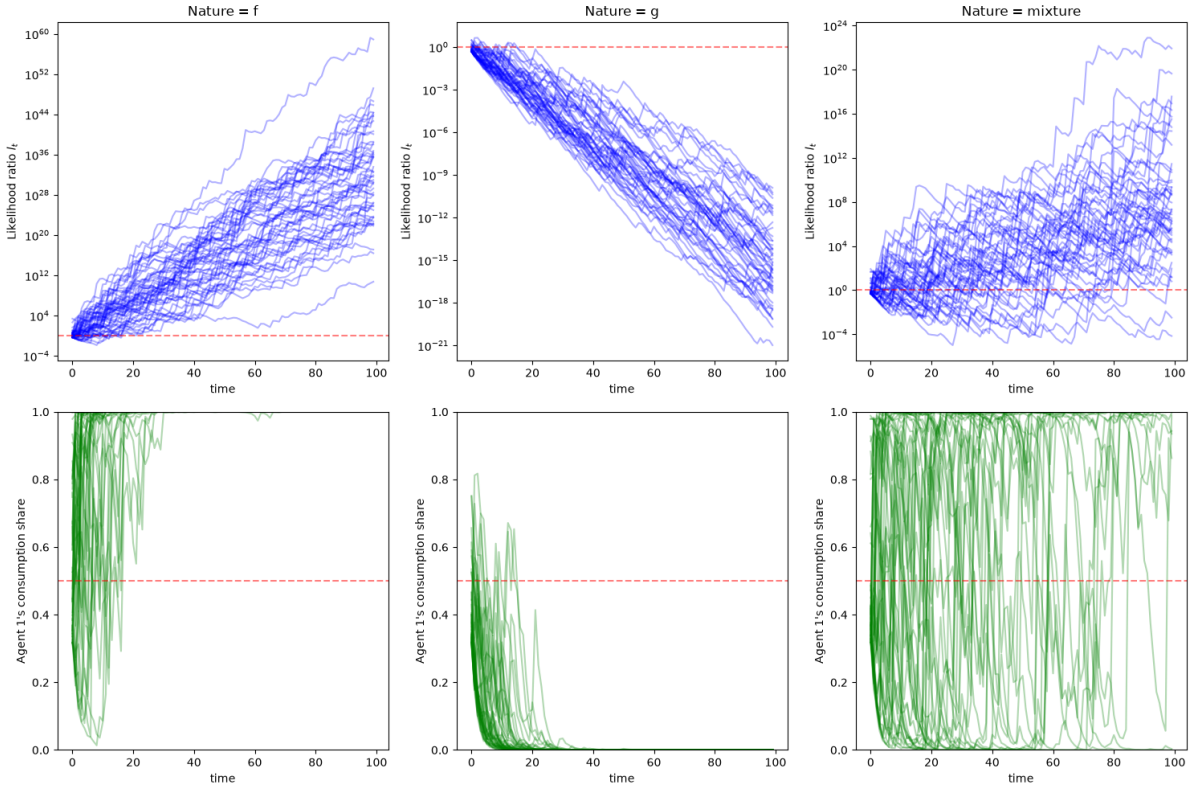
(continues on next page)

(continued from previous page)

```
    for j in range(min(50, l_cum.shape[0])):
        ax.plot(l_cum[j, :], alpha=0.3, color='blue')
    ax.set_yscale('log')
    ax.set_xlabel('time')
    ax.set_ylabel('Likelihood ratio $l_t$')
    ax.set_title(title)
    ax.axhline(y=1, color='red', linestyle='--', alpha=0.5)

    # Consumption shares
    ax = axes[1, i]
    for j in range(min(50, c1.shape[0])):
        ax.plot(c1[j, :], alpha=0.3, color='green')
    ax.set_xlabel('time')
    ax.set_ylabel("Agent 1's consumption share")
    ax.set_ylim([0, 1])
    ax.axhline(y=λ, color='red', linestyle='--', alpha=0.5)

plt.tight_layout()
plt.show()
```



In the left panel, nature chooses f. Agent 1's consumption reaches 1 very quickly.

In the middle panel, nature chooses g. Agent 1's consumption ratio tends to move towards 0 but not as fast as in the first case.

In the right panel, nature flips coins each period. Agent 1's consumption share again moves toward 1, more slowly than in the left panel.

The figures in the top panel remind us of the discussion in *this section* of *Likelihood Ratio Processes*.

We invite readers to revisit *that section* and try to infer the relationships among $D_{KL}(f\|g)$, $D_{KL}(g\|f)$, $D_{KL}(h\|f)$, and

$D_{KL}(h\|g)$.

Let's compute values of KL divergence

```
shares = [np.mean(c1_f[:, -1]), np.mean(c1_g[:, -1]), np.mean(c1_h[:, -1])]
Kf_g, Kg_f = compute_KL(f, g), compute_KL(g, f)
Kf_h, Kg_h = compute_KL_h(h, f, g)

print(f"Final shares: f={shares[0]:.3f}, g={shares[1]:.3f}, mix={shares[2]:.3f}")
print(f"KL divergences: \nKL(f,g)={Kf_g:.3f}, KL(g,f)={Kg_f:.3f}")
print(f"KL(h,f)={Kf_h:.3f}, KL(h,g)={Kg_h:.3f}")
```

```
Final shares: f=1.000, g=0.000, mix=0.923
KL divergences:
KL(f,g)=0.759, KL(g,f)=0.344
KL(h,f)=0.073, KL(h,g)=0.281
```

We find that $KL(f,g) > KL(g,f)$ and $KL(h,g) > KL(h,f)$.

The first inequality tells us that the average "surprise" from having belief g when nature chooses f is greater than the "surprise" from having belief f when nature chooses g.

This explains the difference between the first two panels we noted above.

The second inequality tells us that agent 1's belief distribution f is closer to nature's pick than agent 2's belief g.

By formula (23.8), $\log l_t$ drifts upward at rate $KL(f,g) = 0.759$ per period when nature draws from f and at rate $KL(h,g) - KL(h,f) = 0.208$ per period when nature draws from the mixture h.

The drift has the same sign in the left and right panels and is about 3.6 times smaller in the right panel.

To make this idea more concrete, let's compare two cases:

- agent 1's belief distribution f is close to agent 2's belief distribution g;
- agent 1's belief distribution f is far from agent 2's belief distribution g.

We use the two distributions visualized below

```
def plot_distribution_overlap(ax, x_range, f_vals, g_vals,
                              f_label='f', g_label='g',
                              f_color='blue', g_color='red'):
    """Plot two distributions with their overlap region."""
    ax.plot(x_range, f_vals, color=f_color, linewidth=2, label=f_label)
    ax.plot(x_range, g_vals, color=g_color, linewidth=2, label=g_label)

    overlap = np.minimum(f_vals, g_vals)
    ax.fill_between(x_range, 0, overlap, alpha=0.3, color='purple', label='Overlap')
    ax.set_xlabel('x')
    ax.set_ylabel('Density')
    ax.legend()

# Define close and far belief distributions
f_close = jit(lambda x: p(x, 1, 1))
g_close = jit(lambda x: p(x, 1.1, 1.05))

f_far = jit(lambda x: p(x, 1, 1))
g_far = jit(lambda x: p(x, 3, 1.2))

# Visualize the belief distributions
fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12, 5))
```

(continues on next page)

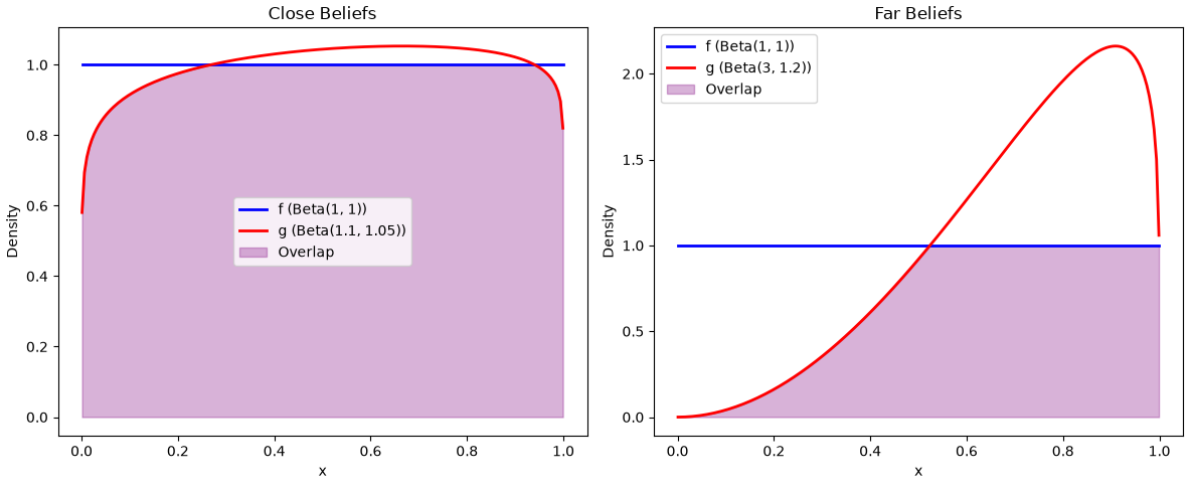
(continued from previous page)

```
x_range = np.linspace(0.001, 0.999, 200)

# Close beliefs
f_close_vals = [f_close(x) for x in x_range]
g_close_vals = [g_close(x) for x in x_range]
plot_distribution_overlap(ax1, x_range, f_close_vals, g_close_vals,
                          f_label='f (Beta(1, 1))', g_label='g (Beta(1.1, 1.05))')
ax1.set_title(f'Close Beliefs')

# Far beliefs
f_far_vals = [f_far(x) for x in x_range]
g_far_vals = [g_far(x) for x in x_range]
plot_distribution_overlap(ax2, x_range, f_far_vals, g_far_vals,
                          f_label='f (Beta(1, 1))', g_label='g (Beta(3, 1.2))')
ax2.set_title(f'Far Beliefs')

plt.tight_layout()
plt.show()
```



Let's draw the same consumption ratio plots as above for agent 1.

We replace the simulation paths with median and percentiles to make the figure cleaner.

In the third column, nature draws from $h = \text{Beta}(2, 1.5)$, a different h from the mixture used above.

Staring at the figure below, can we infer the relation between $KL(f, g)$ and $KL(g, f)$?

From the right panel, can we infer the relation between $KL(h, g)$ and $KL(h, f)$?

```
fig, axes = plt.subplots(2, 3, figsize=(15, 10))
nature_params = {'close': [(1, 1), (1.1, 1.05), (2, 1.5)],
                 'far':   [(1, 1), (3, 1.2),    (2, 1.5)]}
nature_labels = ["Nature = f", "Nature = g", "Nature = h = Beta(2, 1.5)"]
colors = {'close': 'blue', 'far': 'red'}

threshold = 1e-5  # "close to zero" cutoff

for row, (f_belief, g_belief, label) in enumerate([
                        (f_close, g_close, 'close'),
                        (f_far, g_far, 'far')]):
```

(continues on next page)

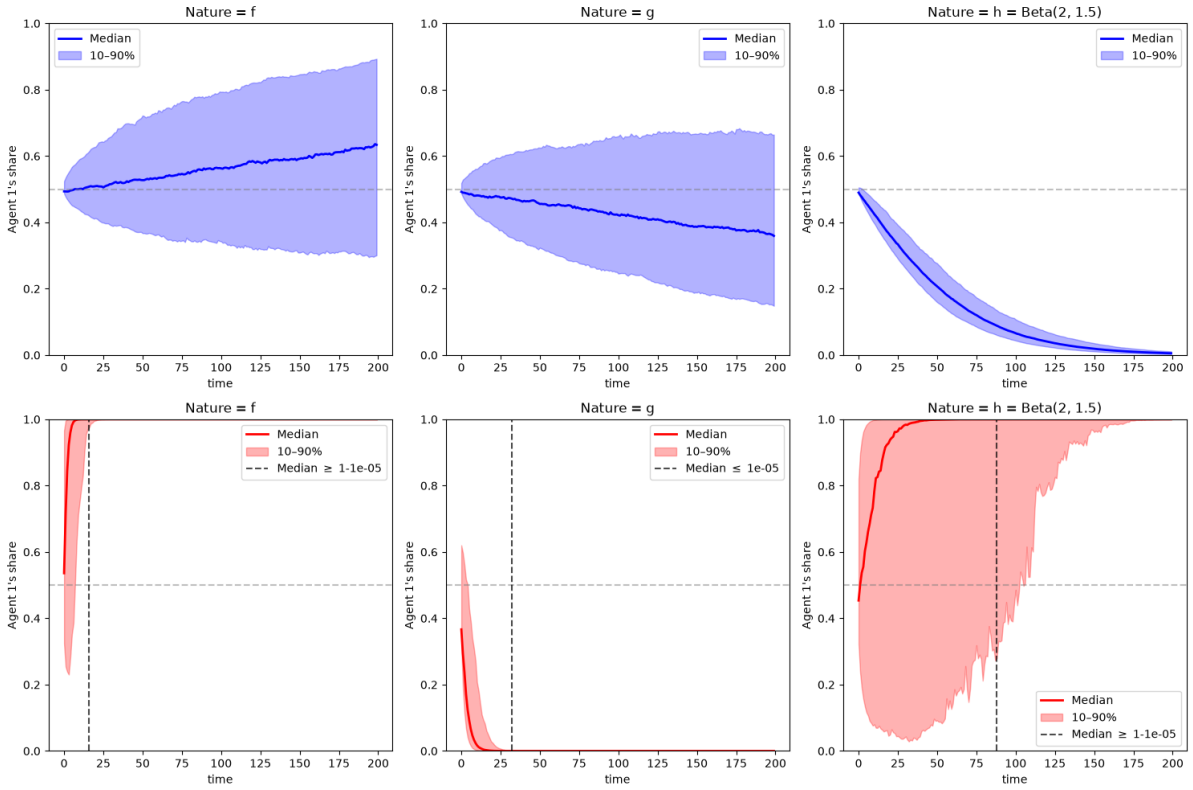
(continued from previous page)

```
    for col, nature_label in enumerate(nature_labels):
        params = nature_params[label][col]
        s_seq = rng.beta(params[0], params[1], (1000, 200))
        _, c1 = simulate_blume_easley(s_seq, f_belief, g_belief, λ)

        median_c1 = np.median(c1, axis=0)
        p10, p90 = np.percentile(c1, [10, 90], axis=0)

        ax = axes[row, col]
        color = colors[label]
        ax.plot(median_c1, color=color, linewidth=2, label='Median')
        ax.fill_between(range(len(median_c1)), p10, p90, alpha=0.3, color=color,␣
↪label='10-90%')
        ax.set_xlabel('time')
        ax.set_ylabel("Agent 1's share")
        ax.set_ylim([0, 1])
        ax.set_title(nature_label)
        ax.axhline(y=λ, color='gray', linestyle='--', alpha=0.5)
        below = np.where(median_c1 < threshold)[0]
        above = np.where(median_c1 > 1-threshold)[0]
        if below.size > 0: first_zero = (below[0], True)
        elif above.size > 0: first_zero = (above[0], False)
        else: first_zero = None
        if first_zero is not None:
            ax.axvline(x=first_zero[0], color='black', linestyle='--',
                       alpha=0.7,
                       label=fr'Median $\leq$ {threshold}' if first_zero[1]
                       else fr'Median $\geq$ 1-{threshold}')
        ax.legend()

plt.tight_layout()
plt.show()
```



Holding to our guesses, let's calculate the four values

```
# Nature's density in the third column
h_beta = jit(lambda x: p(x, 2, 1.5))

# Close case
Kf_g, Kg_f = compute_KL(f_close, g_close), compute_KL(g_close, f_close)
Kf_h, Kg_h = compute_KL_h(h_beta, f_close, g_close)

print(f"KL divergences (close): \nKL(f,g)={Kf_g:.3f}, KL(g,f)={Kg_f:.3f}")
print(f"KL(h,f)={Kf_h:.3f}, KL(h,g)={Kg_h:.3f}")

# Far case
Kf_g, Kg_f = compute_KL(f_far, g_far), compute_KL(g_far, f_far)
Kf_h, Kg_h = compute_KL_h(h_beta, f_far, g_far)

print(f"KL divergences (far): \nKL(f,g)={Kf_g:.3f}, KL(g,f)={Kg_f:.3f}")
print(f"KL(h,f)={Kf_h:.3f}, KL(h,g)={Kg_h:.3f}")
```

```
KL divergences (close):
KL(f,g)=0.003, KL(g,f)=0.003
KL(h,f)=0.108, KL(h,g)=0.082
KL divergences (far):
KL(f,g)=0.759, KL(g,f)=0.344
KL(h,f)=0.108, KL(h,g)=0.241
```

We find that in the first case, $KL(f,g) \approx KL(g,f)$ and both are relatively small, so although either agent 1 or agent 2 will eventually consume everything, convergence displayed in the first two panels on the top is pretty slow.

In the first two panels at the bottom, we see convergence occurring faster (as indicated by the black dashed line) because the divergence gaps $KL(f,g)$ and $KL(g,f)$ are larger.

Since $KL(f,g) > KL(g,f)$, we see faster convergence in the first panel at the bottom when nature chooses f than in the second panel where nature chooses g.

The third column holds nature's h and agent 1's f fixed and changes agent 2's g.

In both rows, $KL(h,f) = 0.108$.

The close g has $KL(h,g) = 0.082 < KL(h,f)$, so agent 1's consumption share drifts toward 0.

The far g has $KL(h,g) = 0.241 > KL(h,f)$, so agent 1's consumption share goes to 1.

Agent 1 loses to the close g and wins against the far g.

In each case the winner is the agent whose beliefs have the smaller KL divergence from h.

These outcomes agree with formula (23.8) and with *A helpful formula* in *Likelihood Ratio Processes*.

23.10 Related lectures

Complete markets models with homogeneous beliefs, a kind often used in macroeconomics and finance, are studied in this quantecon lecture *Competitive Equilibria with Arrow Securities*.

[Blume *et al.*, 2018] discuss a paternalistic case against complete markets. They study the consequences of assuming that a social planner disregards individuals preferences in the sense that it ignores the subjective belief components of their preferences and replaces it with the social planner's beliefs about probabilities.

Survival and Long-Run Dynamics under Recursive Preferences studies Borovička [2020], who shows that market selection can fail when agents have Epstein–Zin recursive preferences: agents with distorted beliefs can survive and can even dominate.

Merging of Opinions: The Blackwell–Dubins Theorem presents the merging-of-opinions theorem of Blackwell and Dubins [1962].

In Exercise 23.11.4 and Exercise 23.11.6, nature's probability law is absolutely continuous with respect to each agent's subjective probability law, each agent's predictions merge with the truth, and both agents survive.

In Exercise 23.11.7, agent 2's dogmatic prior breaks absolute continuity, and agent 2 does not survive.

Incorrect Models studies what a Bayesian learns when nature draws from a mixture h of f and g, the case in the right panels of our first figure.

Statistical Divergence Measures describes Kullback–Leibler divergence and other statistical divergences.

Likelihood processes play an important role in Bayesian learning, as described in *Likelihood Ratio Processes and Bayesian Learning* and as applied in *Job Search IX: Search with Learning*.

Likelihood ratio processes appear again in Additive and Multiplicative Functionals.

23.11 Exercises

Exercise 23.11.1

Agent 1's endowment is $y_t^1 = s_t$ and agent 2's endowment is $y_t^2 = 1 - s_t$.

Let $E^i[s] = \int s\,\pi^i(s)\,ds$.

Take prices (23.12) and allocation (23.6) as given.

1. Show that agent 1's budget constraint (23.9) holds with equality if and only if λ satisfies (23.13).

2. Show that λ is agent 1's share of total wealth at time -1.
3. Show that continuation wealth satisfies (23.14).
4. Derive Arrow-price formula (23.15).
5. Set $\pi^1 = f = \text{Beta}(1,1)$, $\pi^2 = g = \text{Beta}(3, 1.2)$, and $\delta = 0.9$. Compute λ from (23.13). Use Monte Carlo simulation to compute the values of agent 1's consumption and endowment streams when λ satisfies (23.13) and when $\lambda = 0.5$.

Solution

Part 1. Allocation (23.6) is

$$c_t^1(s^t) = \frac{\lambda \pi_t^1(s^t)}{\lambda \pi_t^1(s^t) + (1-\lambda)\pi_t^2(s^t)}.$$

Multiplying by prices (23.12) gives

$$p_t(s^t) c_t^1(s^t) = \frac{\delta^t \pi_t^1(s^t)}{1-\lambda}.$$

Because $\int \pi_t^1(s^t)\, ds^t = 1$, the value of agent 1's consumption stream is

$$\sum_{t=0}^{\infty} \int p_t(s^t) c_t^1(s^t)\, ds^t = \frac{1}{(1-\delta)(1-\lambda)}.$$

Integrating over all dates except t gives $\int \pi_t^i(s^t) s_t\, ds^t = E^i[s]$.

The value of agent 1's endowment stream is

$$\sum_{t=0}^{\infty} \int p_t(s^t) s_t\, ds^t = \frac{\lambda E^1[s] + (1-\lambda) E^2[s]}{(1-\delta)\lambda(1-\lambda)}.$$

Equating the two values gives $\lambda = \lambda E^1[s] + (1-\lambda) E^2[s]$.

Solving for λ gives (23.13).

Part 2. The value of the aggregate endowment stream is

$$\sum_{t=0}^{\infty} \int p_t(s^t) \cdot 1\, ds^t = \frac{1}{(1-\delta)\lambda(1-\lambda)}.$$

Dividing the value of agent 1's consumption stream by this value gives λ.

Part 3. First-order condition (23.10) gives $p_\tau(s^\tau) c_\tau^i(s^\tau) = \delta^\tau \pi_\tau^i(s^\tau)/\mu_i$.

Integrating over $s_{t+1}, \ldots, s_\tau$ gives $\delta^\tau \pi_t^i(s^t)/\mu_i$.

Summing over $\tau \geq t$ gives $\delta^t \pi_t^i(s^t)/(\mu_i(1-\delta))$.

Dividing by $p_t(s^t) = \delta^t \pi_t^i(s^t)/(\mu_i c_t^i(s^t))$ gives $W_t^i(s^t) = c_t^i(s^t)/(1-\delta)$.

Part 4. Using $\pi_{t+1}^i(s^{t+1}) = \pi_t^i(s^t)\pi^i(s_{t+1})$ in (23.12),

$$\begin{aligned} q(s_{t+1}|s^t) &= \frac{p_{t+1}(s^{t+1})}{p_t(s^t)} \\ &= \delta \frac{\lambda \pi_t^1(s^t)\pi^1(s_{t+1}) + (1-\lambda)\pi_t^2(s^t)\pi^2(s_{t+1})}{\lambda \pi_t^1(s^t) + (1-\lambda)\pi_t^2(s^t)} \\ &= \delta\big[\lambda_t(s^t)\pi^1(s_{t+1}) + (1-\lambda_t(s^t))\pi^2(s_{t+1})\big], \end{aligned}$$

where the last line uses $\lambda_t(s^t) = \lambda\pi_t^1(s^t)/[\lambda\pi_t^1(s^t) + (1-\lambda)\pi_t^2(s^t)]$.

Arrow prices from s^t onward have the same form as prices at time -1, with $\lambda_t(s^t)$ in place of λ.

Part 2 applied at s^t gives agent 1's wealth share $\lambda_t(s^t)$, which agrees with part 3.

Part 5. Prices (23.12) turn the value of any stream $x_t(s^t)$ into a mixture of expectations under the two agents' models:

$$\sum_{t=0}^{\infty} \int p_t(s^t) x_t(s^t)\, ds^t = \frac{1}{\lambda(1-\lambda)} \sum_{t=0}^{\infty} \delta^t \left(\lambda E^1[x_t] + (1-\lambda) E^2[x_t]\right).$$

We estimate $E^i[x_t]$ by averaging over histories drawn from π^i.

We report values as shares of the value of the aggregate endowment, which removes the factor $1/(\lambda(1-\lambda))$.

```
f = jit(lambda x: p(x, 1, 1))
g = jit(lambda x: p(x, 3, 1.2))

E1 = quad(lambda s: s * f(s), 0, 1)[0]
E2 = quad(lambda s: s * g(s), 0, 1)[0]
λ_ce = E2 / (1 - E1 + E2)

print(f"E^1[s] = {E1:.4f}, E^2[s] = {E2:.4f}, λ = {λ_ce:.4f}")
```

```
E^1[s] = 0.5000, E^2[s] = 0.7143, λ = 0.5882
```

```
def value_shares(λ, δ=0.9, T=200, N=10_000):
    """
    Monte Carlo values of agent 1's consumption and endowment streams,
    as shares of the value of the aggregate endowment stream.
    """
    discount = δ ** np.arange(T)
    value_c, value_y = 0.0, 0.0

    # Histories drawn from agent 1's model f and agent 2's model g
    for (a, b), weight in [((1, 1), λ), ((3, 1.2), 1 - λ)]:
        s = rng.beta(a, b, (N, T))
        _, l = compute_likelihood_ratios(s, f, g)
        c1 = λ * l / (1 - λ + λ * l)
        value_c += weight * discount @ c1.mean(axis=0)
        value_y += weight * discount @ s.mean(axis=0)

    return value_c / discount.sum(), value_y / discount.sum()


for λ_val in [λ_ce, 0.5]:
    share_c, share_y = value_shares(λ_val)
    exact_y = λ_val * E1 + (1 - λ_val) * E2
    print(f"λ = {λ_val:.3f}: value of c^1 = {share_c:.3f} (exact {λ_val:.3f}), "
          f"value of y^1 = {share_y:.3f} (exact {exact_y:.3f})")
```

```
λ = 0.588: value of c^1 = 0.588 (exact 0.588), value of y^1 = 0.588 (exact 0.
↪588)
λ = 0.500: value of c^1 = 0.500 (exact 0.500), value of y^1 = 0.607 (exact 0.
↪607)
```

At $\lambda = 10/17$, the values of agent 1's consumption and endowment streams agree.

At $\lambda = 0.5$, agent 1's endowment stream is worth more than its consumption stream.

Agent 1 can then afford more consumption, so $\lambda = 0.5$ is not a competitive equilibrium Pareto weight without a transfer.

Exercise 23.11.2

Suppose that agent i has its own discount factor δ_i, so that

$$V^i = \sum_{t=0}^{\infty} \sum_{s^t} \delta_i^t u(c_t^i(s^t)) \pi_t^i(s^t).$$

Keep all other assumptions of the planning problem.

1. Show that the planner's allocation satisfies

$$\frac{c_t^1(s^t)}{c_t^2(s^t)} = \frac{\lambda}{1-\lambda} \left(\frac{\delta_1}{\delta_2}\right)^t l_t(s^t).$$

2. Show that when nature draws s_t IID from π,

$$\frac{1}{t} \log \frac{c_t^1(s^t)}{c_t^2(s^t)} \to [\log \delta_1 - KL(\pi, \pi^1)] - [\log \delta_2 - KL(\pi, \pi^2)] \quad \text{almost surely.}$$

   The quantity $\log \delta_i - KL(\pi, \pi^i)$ is agent i's **survival index**.

3. Let nature draw from the mixture $h = 0.5f + 0.5g$ with $f = \text{Beta}(1,1)$ and $g = \text{Beta}(3, 1.2)$, as in the right panels of our first figure. Agent 1 believes f and agent 2 believes g. For which values of δ_1/δ_2 does agent 1's consumption share converge to 1?

4. Set $\delta_2 = 0.99$ and $\lambda = 0.5$. Simulate agent 1's consumption share for values of δ_1 on both sides of the threshold you found.

Solution

Part 1. The planner's first-order conditions become

$$\lambda \delta_1^t \frac{\pi_t^1(s^t)}{c_t^1(s^t)} = \theta_t(s^t) = (1-\lambda)\delta_2^t \frac{\pi_t^2(s^t)}{c_t^2(s^t)}.$$

Equating the first and third expressions and rearranging gives the result.

Part 2. Taking logarithms gives

$$\log \frac{c_t^1(s^t)}{c_t^2(s^t)} = \log \frac{\lambda}{1-\lambda} + t \log \frac{\delta_1}{\delta_2} + \log l_t(s^t).$$

Dividing by t and applying (23.8) gives the result.

Part 3. Agent 1's consumption share converges to 1 when its survival index exceeds agent 2's:

$$\log \frac{\delta_1}{\delta_2} + KL(h, g) - KL(h, f) > 0,$$

or

$$\frac{\delta_1}{\delta_2} > \exp(-[KL(h, g) - KL(h, f)]).$$

```
f = jit(lambda x: p(x, 1, 1))
g = jit(lambda x: p(x, 3, 1.2))
h = jit(lambda x: 0.5 * f(x) + 0.5 * g(x))

Kh_f, Kh_g = compute_KL_h(h, f, g)
ratio_threshold = np.exp(-(Kh_g - Kh_f))

print(f"KL(h,g) - KL(h,f) = {Kh_g - Kh_f:.3f}")
print(f"agent 1's share converges to 1 if δ_1/δ_2 > {ratio_threshold:.3f}")
```

```
KL(h,g) - KL(h,f) = 0.208
agent 1's share converges to 1 if δ_1/δ_2 > 0.812
```

Agent 1 has better beliefs.

A more patient agent 2 drives agent 1 out when $\delta_1/\delta_2 < 0.81$.

An agent can make up for worse beliefs by saving more.

Part 4. With $\delta_2 = 0.99$, the threshold is $\delta_1 = 0.99 \times 0.81 \approx 0.80$.

We simulate $\delta_1 = 0.95$ and $\delta_1 = 0.85$ above the threshold and $\delta_1 = 0.75$ below it.

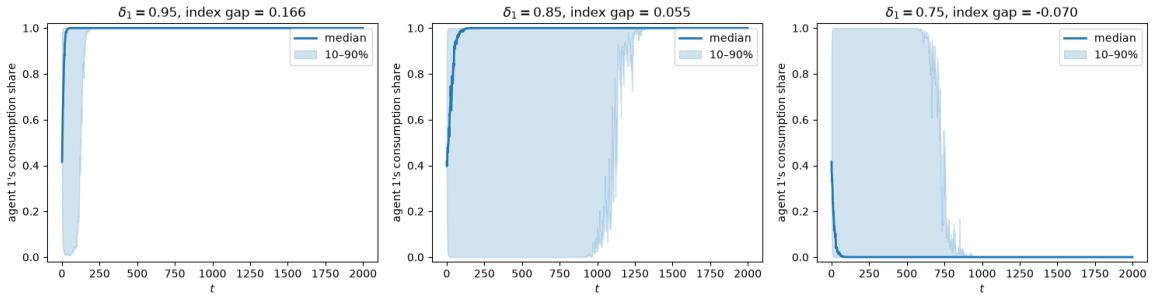
```
def simulate_patience(sequences, δ_1, δ_2, λ=0.5):
    """Agent 1's consumption share when agents discount at δ_1 and δ_2."""
    log_l = np.cumsum(np.log(f(sequences) / g(sequences)), axis=1)
    t = np.arange(sequences.shape[1])
    return expit(np.log(λ / (1 - λ)) + t * np.log(δ_1 / δ_2) + log_l)


T, N = 2000, 1000
δ_2 = 0.99

# Nature draws from the mixture h
draw_f = rng.random((N, T)) < 0.5
s_seq_h = np.where(draw_f, rng.beta(1, 1, (N, T)), rng.beta(3, 1.2, (N, T)))

fig, axes = plt.subplots(1, 3, figsize=(15, 4))
for ax, δ_1 in zip(axes, [0.95, 0.85, 0.75]):
    c1 = simulate_patience(s_seq_h, δ_1, δ_2)
    gap = np.log(δ_1 / δ_2) + Kh_g - Kh_f
    p10, p50, p90 = np.percentile(c1, [10, 50, 90], axis=0)
    ax.plot(p50, color='C0', linewidth=2, label='median')
    ax.fill_between(range(T), p10, p90, color='C0', alpha=0.2, label='10-90%')
    ax.set_title(rf'$\delta_1 = {δ_1}$, index gap = {gap:.3f}')
    ax.set_xlabel('$t$')
    ax.set_ylabel("agent 1's consumption share")
    ax.set_ylim([-0.02, 1.02])
    ax.legend()

plt.tight_layout()
plt.show()
```



The panel titles report the difference between agent 1's and agent 2's survival indexes.

When the difference is positive, agent 1's consumption share goes to 1.

When it is negative, agent 1's consumption share goes to 0 even though agent 1 has better beliefs.

Exercise 23.11.3

Let nature draw s_t IID from $\pi = \text{Beta}(2, 2)$, let agent 1 believe $\pi^1 = \text{Beta}(2, 3)$, and let agent 2 believe $\pi^2 = \text{Beta}(3, 2)$.

Set $\lambda = 0.5$.

1. Show that $KL(\pi, \pi^1) = KL(\pi, \pi^2)$.
2. Show that agent 1's consumption share c_t^1 does not converge.
3. Simulate many paths and plot histograms of c_t^1 at $t = 10, 100, 2000$.

Solution

Part 1. The densities satisfy $\pi^1(s) = \pi^2(1-s)$ and $\pi(s) = \pi(1-s)$.

The change of variables $s \mapsto 1 - s$ maps $KL(\pi, \pi^1)$ into $KL(\pi, \pi^2)$.

Part 2. Both densities have normalizing constant 12, so

$$\log \frac{\pi^1(s)}{\pi^2(s)} = \log \frac{s(1-s)^2}{s^2(1-s)} = \log \frac{1-s}{s}.$$

Under π, this random variable has mean 0 by symmetry and variance $\sigma^2 > 0$.

The log likelihood ratio $\log l_t$ is a random walk with zero drift and increment variance σ^2.

The law of the iterated logarithm implies $\limsup_t \log l_t = +\infty$ and $\liminf_t \log l_t = -\infty$ almost surely.

With $\lambda = 0.5$, $c_t^1 = l_t/(1 + l_t)$, so $\limsup_t c_t^1 = 1$ and $\liminf_t c_t^1 = 0$.

By the central limit theorem, $\log l_t$ is approximately normal with mean 0 and variance $(t+1)\sigma^2$.

The probability that c_t^1 lies near 0.5 shrinks at rate $1/\sqrt{t}$, and the distribution of c_t^1 piles up near 0 and 1.

Part 3.

```
rng = np.random.default_rng(1234)

nature = jit(lambda x: p(x, 2, 2))
f = jit(lambda x: p(x, 2, 3))
g = jit(lambda x: p(x, 3, 2))

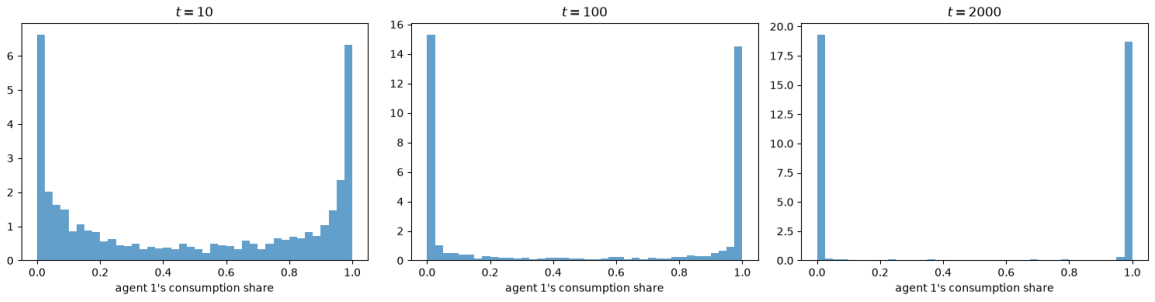
K1, K2 = compute_KL_h(nature, f, g)
σ2 = quad(lambda s: nature(s) * np.log((1 - s) / s)**2, 0, 1)[0]
print(f"KL(π, π^1) = {K1:.4f}, KL(π, π^2) = {K2:.4f}, σ^2 = {σ2:.3f}")
```

```
KL(π, π^1) = 0.1402, KL(π, π^2) = 0.1402, σ^2 = 1.290
```

```
N, T = 2000, 2001
s_seq = rng.beta(2, 2, (N, T))
log_l = np.cumsum(np.log(f(s_seq) / g(s_seq)), axis=1)
c1 = expit(log_l)      # λ = 0.5

fig, axes = plt.subplots(1, 3, figsize=(15, 4))
for ax, t in zip(axes, [10, 100, 2000]):
    ax.hist(c1[:, t], bins=40, range=(0, 1), density=True, alpha=0.7)
    ax.set_title(f'$t = {t}$')
    ax.set_xlabel("agent 1's consumption share")
plt.tight_layout()
plt.show()

for t in [10, 100, 2000]:
    near_half = np.mean(np.abs(c1[:, t] - 0.5) < 0.05)
    extreme = np.mean((c1[:, t] < 0.05) | (c1[:, t] > 0.95))
    print(f"t = {t:4d}: P(|c1 - 0.5| < 0.05) = {near_half:.3f}, "
          f"P(c1 < 0.05 or c1 > 0.95) = {extreme:.3f}")
```

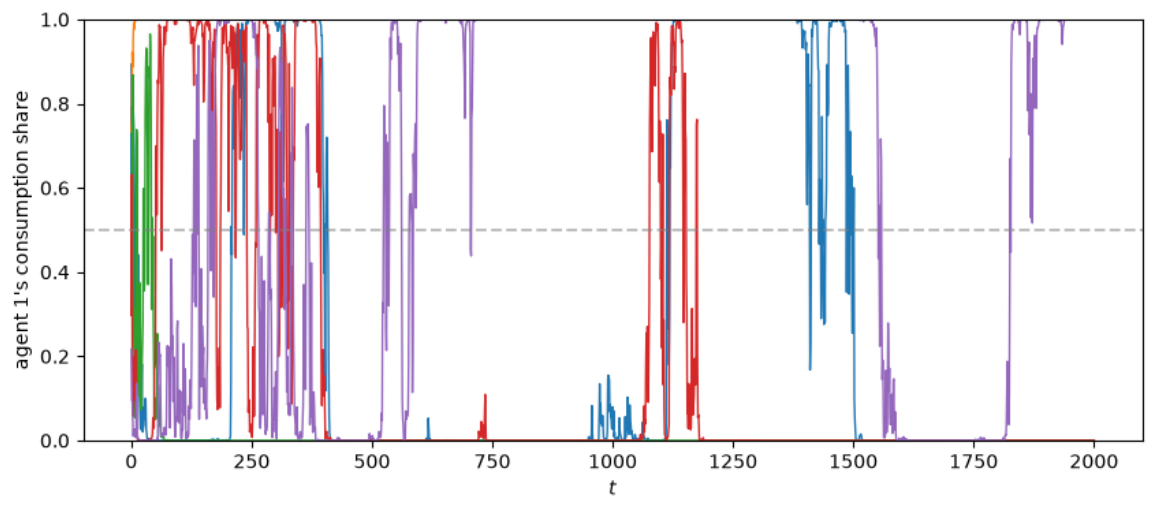



```
t =   10: P(|c1 - 0.5| < 0.05) = 0.035, P(c1 < 0.05 or c1 > 0.95) = 0.433
t =  100: P(|c1 - 0.5| < 0.05) = 0.011, P(c1 < 0.05 or c1 > 0.95) = 0.794
t = 2000: P(|c1 - 0.5| < 0.05) = 0.003, P(c1 < 0.05 or c1 > 0.95) = 0.961
```

As t grows, the histograms put more mass near 0 and 1 and less near 0.5.

The next figure shows five sample paths.

```
fig, ax = plt.subplots(figsize=(10, 4))
for j in range(5):
    ax.plot(c1[j, :], linewidth=1)
ax.axhline(y=0.5, color='gray', linestyle='--', alpha=0.5)
ax.set_xlabel('$t$')
ax.set_ylabel("agent 1's consumption share")
ax.set_ylim([0, 1])
plt.show()
```



Each path keeps crossing 0.5 and spends long spells near 0 or near 1.

Neither agent is driven out, and neither agent's share settles down.

Exercise 23.11.4

In this exercise, each of our two agents is a Bayesian who does not know whether nature draws from f or from g.

In this exercise and the following ones, dates start at $t = 1$: the history is $s^t = (s_1, \ldots, s_t)$.

The following two models are on the table

$$f(s^t) = f(s_1)f(s_2)\cdots f(s_t)$$

and

$$g(s^t) = g(s_1)g(s_2)\cdots g(s_t)$$

as is an associated likelihood ratio process

$$L(s^t) = \frac{f(s^t)}{g(s^t)}.$$

The two agents share f and g but have different prior probabilities $\pi_0^1, \pi_0^2 \in (0, 1)$ on model f.

Agent i's joint density of s^t is the prior-weighted mixture

$$m^i(s^t) = \pi_0^i f(s^t) + (1 - \pi_0^i)g(s^t). \tag{23.16}$$

As in *Likelihood Ratio Processes and Bayesian Learning*, Bayes' law implies that agent i's posterior probability on model f after observing s^t is

$$\pi_t^i = \frac{\pi_0^i L(s^t)}{\pi_0^i L(s^t) + 1 - \pi_0^i}.$$

Joint density (23.16) is a product of one-step-ahead predictive densities:

$$m^i(s^t) = \prod_{j=1}^{t} \left[\pi_{j-1}^i f(s_j) + (1 - \pi_{j-1}^i)g(s_j)\right].$$

The predictive density of s_j uses the posterior π^i_{j-1} that agent i forms before observing s_j.

We now hand probability models (23.16) for $i = 1, 2$ to the social planner.

1. Show that the likelihood ratio $l_t(s^t) = m^1(s^t)/m^2(s^t)$ satisfies

$$l_t(s^t) = \frac{\pi^1_0 L(s^t) + 1 - \pi^1_0}{\pi^2_0 L(s^t) + 1 - \pi^2_0},$$

   and find the limits of agent 1's consumption share when nature draws from f and when nature draws from g.

2. Set $f \sim \text{Beta}(1.5, 1)$, $g \sim \text{Beta}(1, 1.5)$, and $\lambda = 0.5$. Write Python code that simulates posteriors and consumption shares for prior pairs $(\pi^1_0, \pi^2_0) = (0.3, 0.7), (0.7, 0.3), (0.1, 0.9)$ when nature draws from f and when nature draws from g.

3. Does either agent's consumption share converge to zero? How do the priors π^1_0 and π^2_0 affect outcomes?

Solution

Part 1. Dividing the numerator and the denominator of $m^1(s^t)/m^2(s^t)$ by $g(s^t)$ gives the formula for $l_t(s^t)$.

When nature draws from f, $L(s^t) \to +\infty$ and $l_t \to \pi^1_0/\pi^2_0$, so

$$c^1_t \to \frac{\lambda\pi^1_0}{\lambda\pi^1_0 + (1-\lambda)\pi^2_0}.$$

When nature draws from g, $L(s^t) \to 0$ and $l_t \to (1-\pi^1_0)/(1-\pi^2_0)$, so

$$c^1_t \to \frac{\lambda(1-\pi^1_0)}{\lambda(1-\pi^1_0) + (1-\lambda)(1-\pi^2_0)}.$$

Both limits lie strictly between 0 and 1.

Both agents survive.

Part 2. The following function computes posteriors, the likelihood ratio l_t, and agent 1's consumption share.

It builds l_t from one-step-ahead predictive densities that use posteriors formed before s_t is observed.

```
def simulate_learning_blume_easley(sequences, f_belief, g_belief,
                                   π_0_1, π_0_2, λ=0.5):
    """
    Simulate the Blume-Easley model with two Bayesian agents.

    Column t-1 of `sequences` holds s_t.
    """
    N, T = sequences.shape
    f_vals = f_belief(sequences)
    g_vals = g_belief(sequences)

    π_1_seq = np.empty((N, T))
    π_2_seq = np.empty((N, T))
    l_agents_seq = np.empty((N, T))

    # Priors and l_0 = 1
    π_1 = np.full(N, π_0_1)
    π_2 = np.full(N, π_0_2)
    l_agents = np.ones(N)
```

```
    for t in range(T):
        # Predictive densities of s_{t+1} use posteriors given s^t
        m1 = π_1 * f_vals[:, t] + (1 - π_1) * g_vals[:, t]
        m2 = π_2 * f_vals[:, t] + (1 - π_2) * g_vals[:, t]
        l_agents = l_agents * m1 / m2

        # Bayes' law after observing s_{t+1}
        π_1 = π_1 * f_vals[:, t] / m1
        π_2 = π_2 * f_vals[:, t] / m2

        π_1_seq[:, t] = π_1
        π_2_seq[:, t] = π_2
        l_agents_seq[:, t] = l_agents

    c1_share = λ * l_agents_seq / (1 - λ + λ * l_agents_seq)

    return {
        'π_1': π_1_seq,
        'π_2': π_2_seq,
        'c1_share': c1_share,
        'l_agents': l_agents_seq
    }
```

The next function computes the limits of agent 1's consumption share from part 1.

```
def c1_limit(π_0_1, π_0_2, nature_type, λ=0.5):
    """Limit of agent 1's consumption share."""
    if nature_type == 'f':
        w_1, w_2 = π_0_1, π_0_2
    else:
        w_1, w_2 = 1 - π_0_1, 1 - π_0_2
    return λ * w_1 / (λ * w_1 + (1 - λ) * w_2)
```

Let's run simulations for different scenarios.

We use $\lambda = 0.5$, $T = 40$, and $N = 1000$.

```
λ = 0.5
T = 40
N = 1000

F_a, F_b = 1.5, 1
G_a, G_b = 1, 1.5

f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))
```

We'll start with different initial priors $\pi_0^i \in (0, 1)$ and widen the gap between them.

```
# Different initial priors
π_0_scenarios = [
    (0.3, 0.7),
    (0.7, 0.3),
    (0.1, 0.9),
]
```

Now we can run simulations for different scenarios

```
rng = np.random.default_rng()

# Nature follows f
s_seq_f = rng.beta(F_a, F_b, (N, T))

# Nature follows g
s_seq_g = rng.beta(G_a, G_b, (N, T))

results_f = {}
results_g = {}

for i, (π_0_1, π_0_2) in enumerate(π_0_scenarios):
    # When nature follows f
    results_f[i] = simulate_learning_blume_easley(
            s_seq_f, f, g, π_0_1, π_0_2, λ)
    # When nature follows g
    results_g[i] = simulate_learning_blume_easley(
            s_seq_g, f, g, π_0_1, π_0_2, λ)
```

The recursion reproduces the closed form for l_t in part 1

```
_, L = compute_likelihood_ratios(s_seq_f, f, g)
π_0_1, π_0_2 = π_0_scenarios[0]
l_closed = (π_0_1 * L + 1 - π_0_1) / (π_0_2 * L + 1 - π_0_2)
print(np.max(np.abs(results_f[0]['l_agents'] / l_closed - 1)))
```

```
6.873390745454344e-13
```

Let's visualize the results

```
def plot_learning_results(results, π_0_scenarios, nature_type, λ=0.5):
    """
    Plot median posteriors and consumption shares for learning agents.
    """
    fig, axes = plt.subplots(3, 2, figsize=(10, 15))
    truth = 1.0 if nature_type == 'f' else 0.0

    for row, (π_0_1, π_0_2) in enumerate(π_0_scenarios):
        res = results[row]
        t_grid = np.arange(1, res['c1_share'].shape[1] + 1)
        scenario_label = rf'$\pi_0^1 = {π_0_1}, \pi_0^2 = {π_0_2}$'

        # Posterior probabilities on f
        ax = axes[row, 0]
        ax.plot(t_grid, np.median(res['π_1'], axis=0), 'C0',
                label='agent 1', linewidth=2)
        ax.plot(t_grid, np.median(res['π_2'], axis=0), 'C1',
                label='agent 2', linewidth=2)
        ax.axhline(y=truth, color='gray', linestyle='--',
                   alpha=0.5, label=f'truth ({nature_type})')
        ax.set_title(f'Beliefs when nature = {nature_type}\n{scenario_label}')
        ax.set_ylabel(r'median $\pi^i_t$')
        ax.set_ylim([-0.05, 1.05])
        ax.legend()

        # Consumption shares and their limit
        ax = axes[row, 1]
        ax.plot(t_grid, np.median(res['c1_share'], axis=0), 'g-',
                linewidth=2, label='median')
```

```
        ax.axhline(y=c1_limit(π_0_1, π_0_2, nature_type, λ),
                   color='k', linestyle='--', alpha=0.7, label='limit')
        ax.set_title(f'Agent 1 consumption share (Nature = {nature_type})')
        ax.set_ylabel('consumption share')
        ax.set_ylim([0, 1])
        ax.legend()

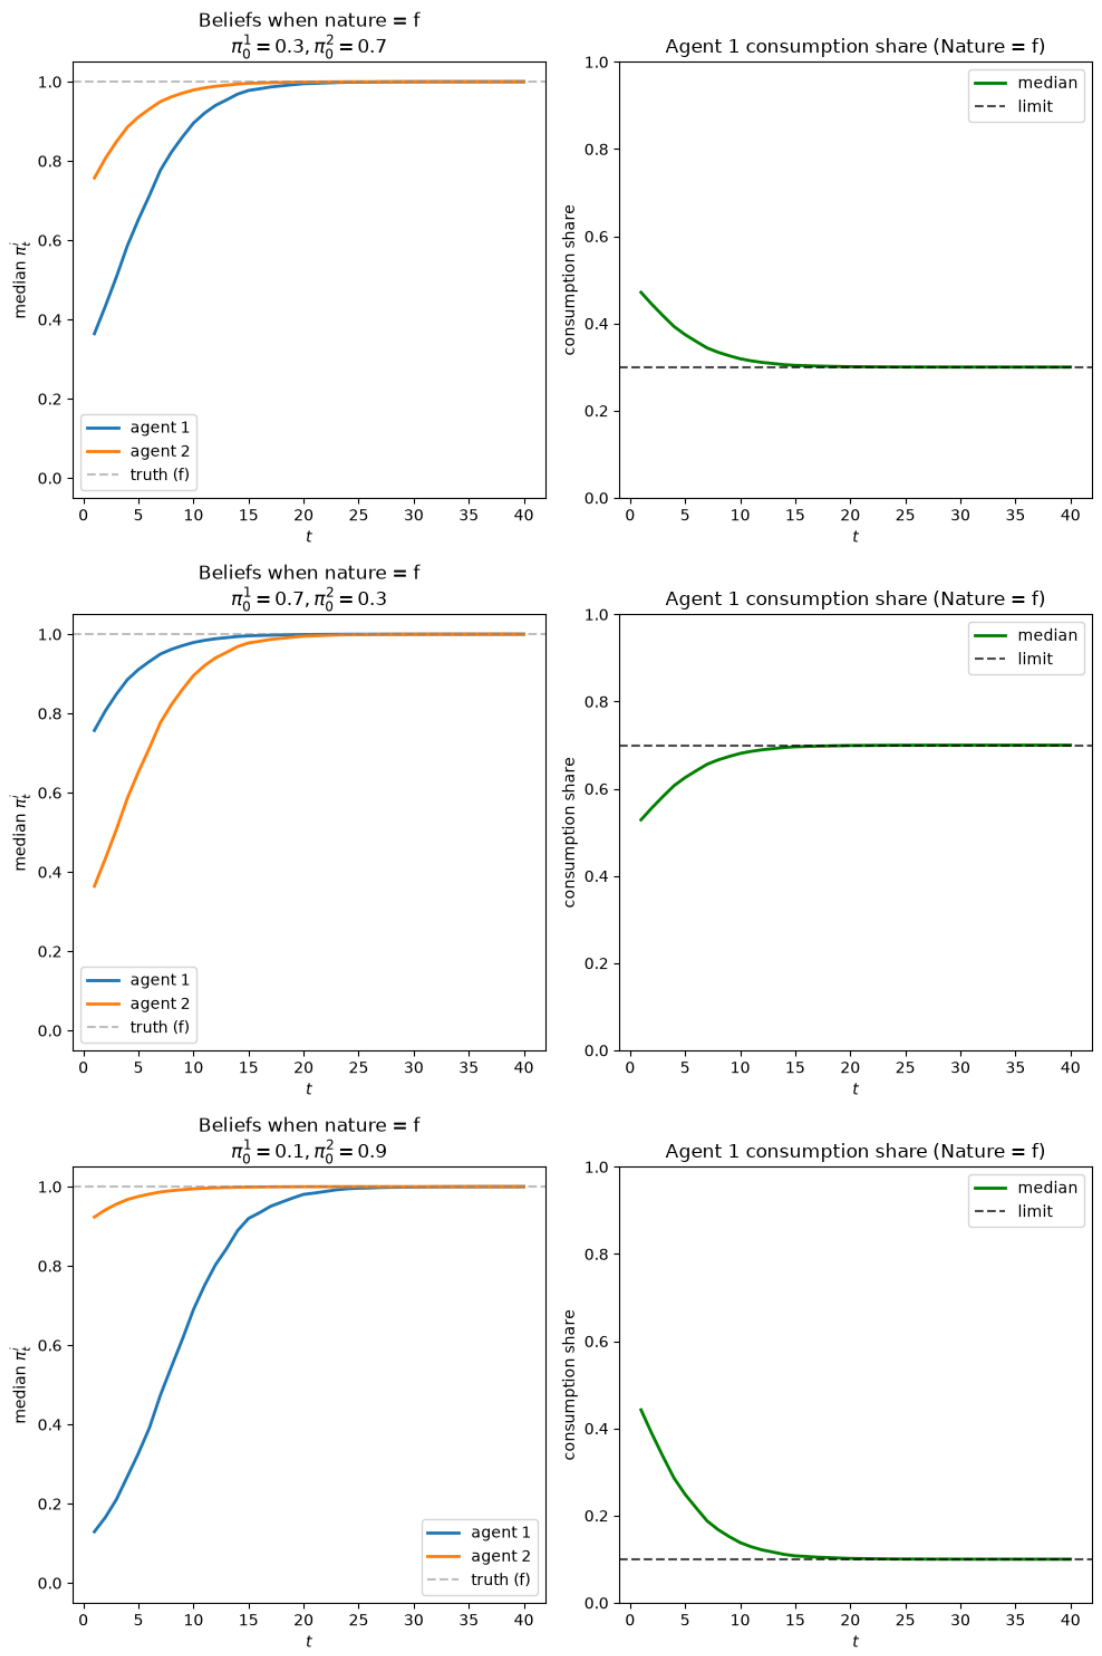
        for col in range(2):
            axes[row, col].set_xlabel('$t$')

    plt.tight_layout()
    return fig, axes


def print_limits(results, π_0_scenarios, nature_type, λ=0.5):
    """Compare median final consumption shares with their limits."""
    for i, (π_0_1, π_0_2) in enumerate(π_0_scenarios):
        c1_T = np.median(results[i]['c1_share'][:, -1])
        limit = c1_limit(π_0_1, π_0_2, nature_type, λ)
        print(f"nature = {nature_type}, priors = ({π_0_1}, {π_0_2}): "
              f"median c1 at T = {c1_T:.3f}, limit = {limit:.3f}")
```

Now we'll plot outcomes when nature follows f:

```
fig_f, axes_f = plot_learning_results(results_f, π_0_scenarios, 'f', λ)
plt.show()
print_limits(results_f, π_0_scenarios, 'f', λ)
```


Beliefs when nature = f
$\pi_0^1 = 0.3, \pi_0^2 = 0.7$
median π_t^i
agent 1
agent 2
truth (f)
t
Agent 1 consumption share (Nature = f)
consumption share
median
limit
Beliefs when nature = f
$\pi_0^1 = 0.7, \pi_0^2 = 0.3$
Agent 1 consumption share (Nature = f)
Beliefs when nature = f
$\pi_0^1 = 0.1, \pi_0^2 = 0.9$
Agent 1 consumption share (Nature = f)

```
nature = f, priors = (0.3, 0.7): median c1 at T = 0.300, limit = 0.300
nature = f, priors = (0.7, 0.3): median c1 at T = 0.700, limit = 0.700
nature = f, priors = (0.1, 0.9): median c1 at T = 0.100, limit = 0.100
```

Agent 1's median consumption share settles at the limit from part 1 in each scenario.

Neither agent's consumption share goes to zero.

Both agents learn that nature draws from f.

Their posterior log odds $\log[\pi_t^i/(1-\pi_t^i)] = \log[\pi_0^i/(1-\pi_0^i)] + \log L(s^t)$ move in parallel.

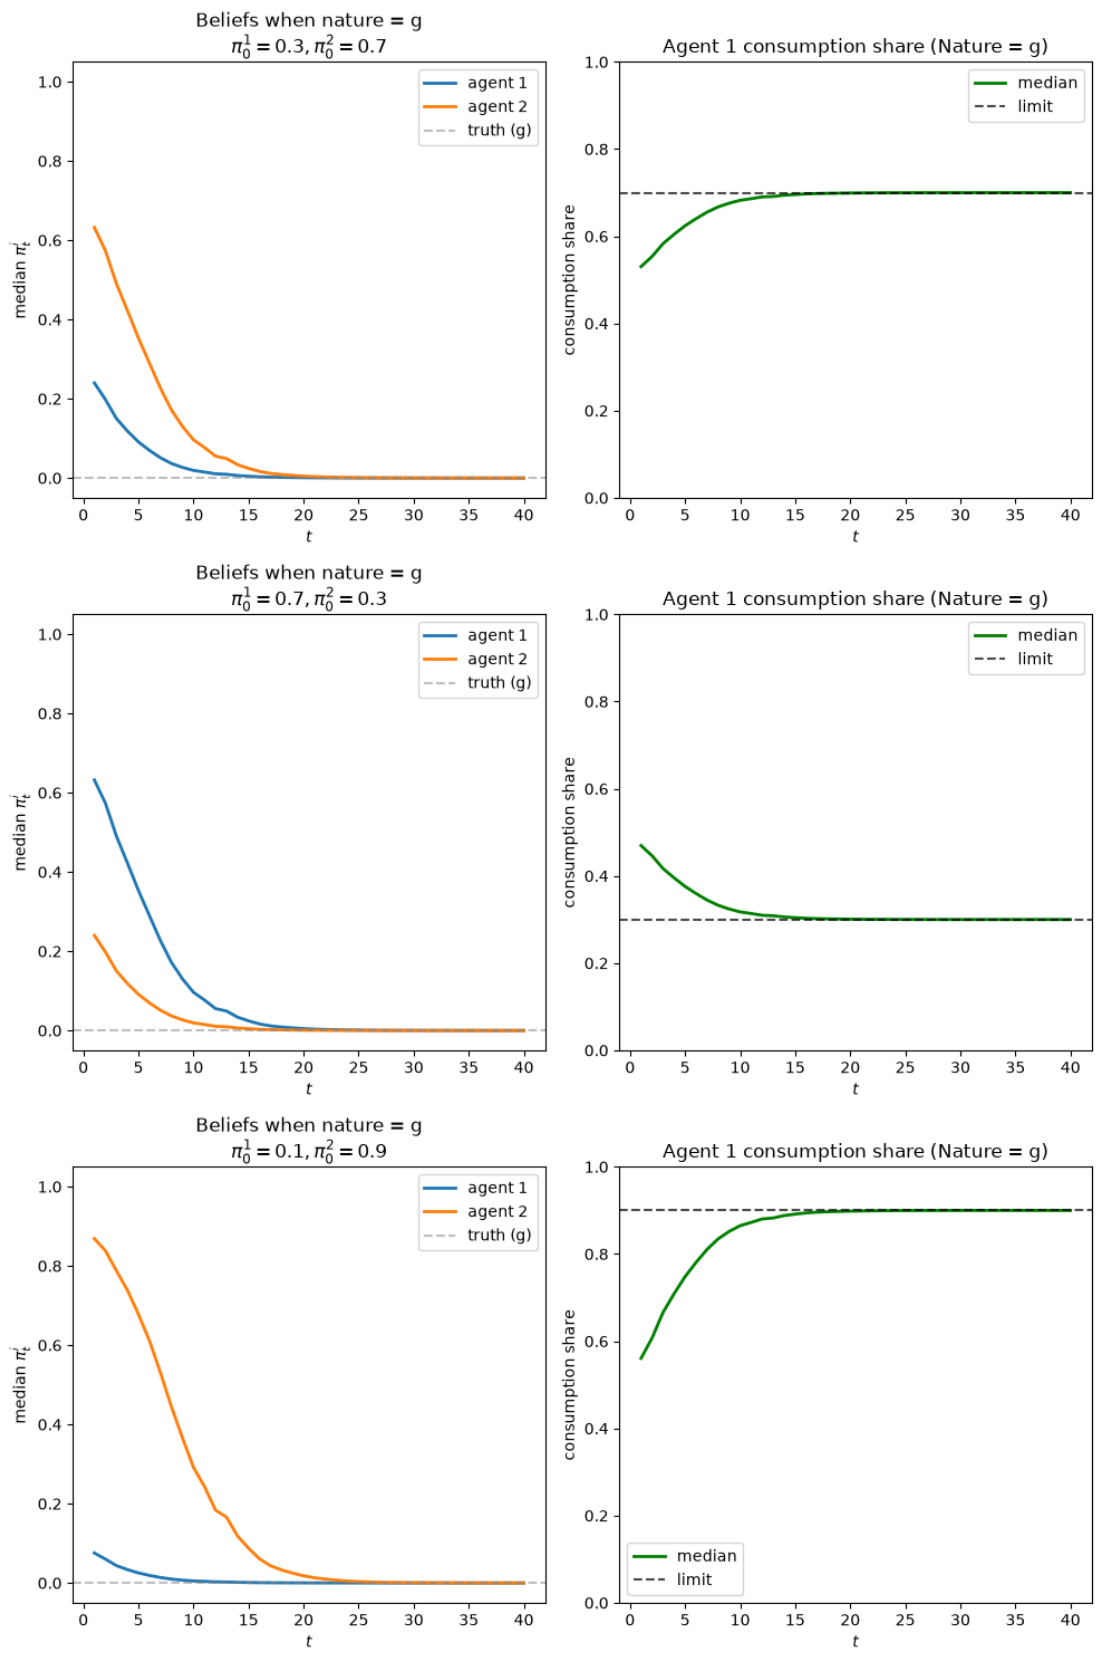
The agent with the larger prior on f has the larger posterior on f at every date.

Once both agents have learned, their predictive densities agree and l_t stops moving.

The agent who puts more prior probability on the true model keeps the larger consumption share.

Now let's plot outcomes when nature follows g:

```
fig_g, axes_g = plot_learning_results(results_g, π_0_scenarios, 'g', λ)
plt.show()
print_limits(results_g, π_0_scenarios, 'g', λ)
```


Beliefs when nature = g
$\pi_0^1 = 0.3, \pi_0^2 = 0.7$
agent 1
agent 2
truth (g)
median π_t^i
t
Agent 1 consumption share (Nature = g)
median
limit
consumption share
t
Beliefs when nature = g
$\pi_0^1 = 0.7, \pi_0^2 = 0.3$
agent 1
agent 2
truth (g)
median π_t^i
t
Agent 1 consumption share (Nature = g)
median
limit
consumption share
t
Beliefs when nature = g
$\pi_0^1 = 0.1, \pi_0^2 = 0.9$
agent 1
agent 2
truth (g)
median π_t^i
t
Agent 1 consumption share (Nature = g)
median
limit
consumption share
t

```
nature = g, priors = (0.3, 0.7): median c1 at T = 0.700, limit = 0.700
nature = g, priors = (0.7, 0.3): median c1 at T = 0.300, limit = 0.300
nature = g, priors = (0.1, 0.9): median c1 at T = 0.900, limit = 0.900
```

We observe symmetrical outcomes.

Bayesian learners who put positive prior probability on the true model all survive.

Nature's probability law $f(s^t)$ or $g(s^t)$ is absolutely continuous with respect to each agent's joint density $m^i(s^t)$.

As *Merging of Opinions: The Blackwell–Dubins Theorem* explains, each agent's predictions then merge with the truth.

Exercise 23.11.5

In the previous exercise, we purposefully set the two beta distributions to be relatively close to each other.

That made it challenging to distinguish the distributions.

Now let's study outcomes when the distributions are further apart.

Let's set $f \sim \text{Beta}(2, 5)$ and $g \sim \text{Beta}(5, 2)$.

Please use the Python code you have written to study outcomes.

Solution

Here is one solution

```
λ = 0.5
T = 40
N = 1000

rng = np.random.default_rng()

F_a, F_b = 2, 5
G_a, G_b = 5, 2

f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))

π_0_scenarios = [
    (0.3, 0.7),
    (0.7, 0.3),
    (0.1, 0.9),
]

s_seq_f = rng.beta(F_a, F_b, (N, T))
s_seq_g = rng.beta(G_a, G_b, (N, T))

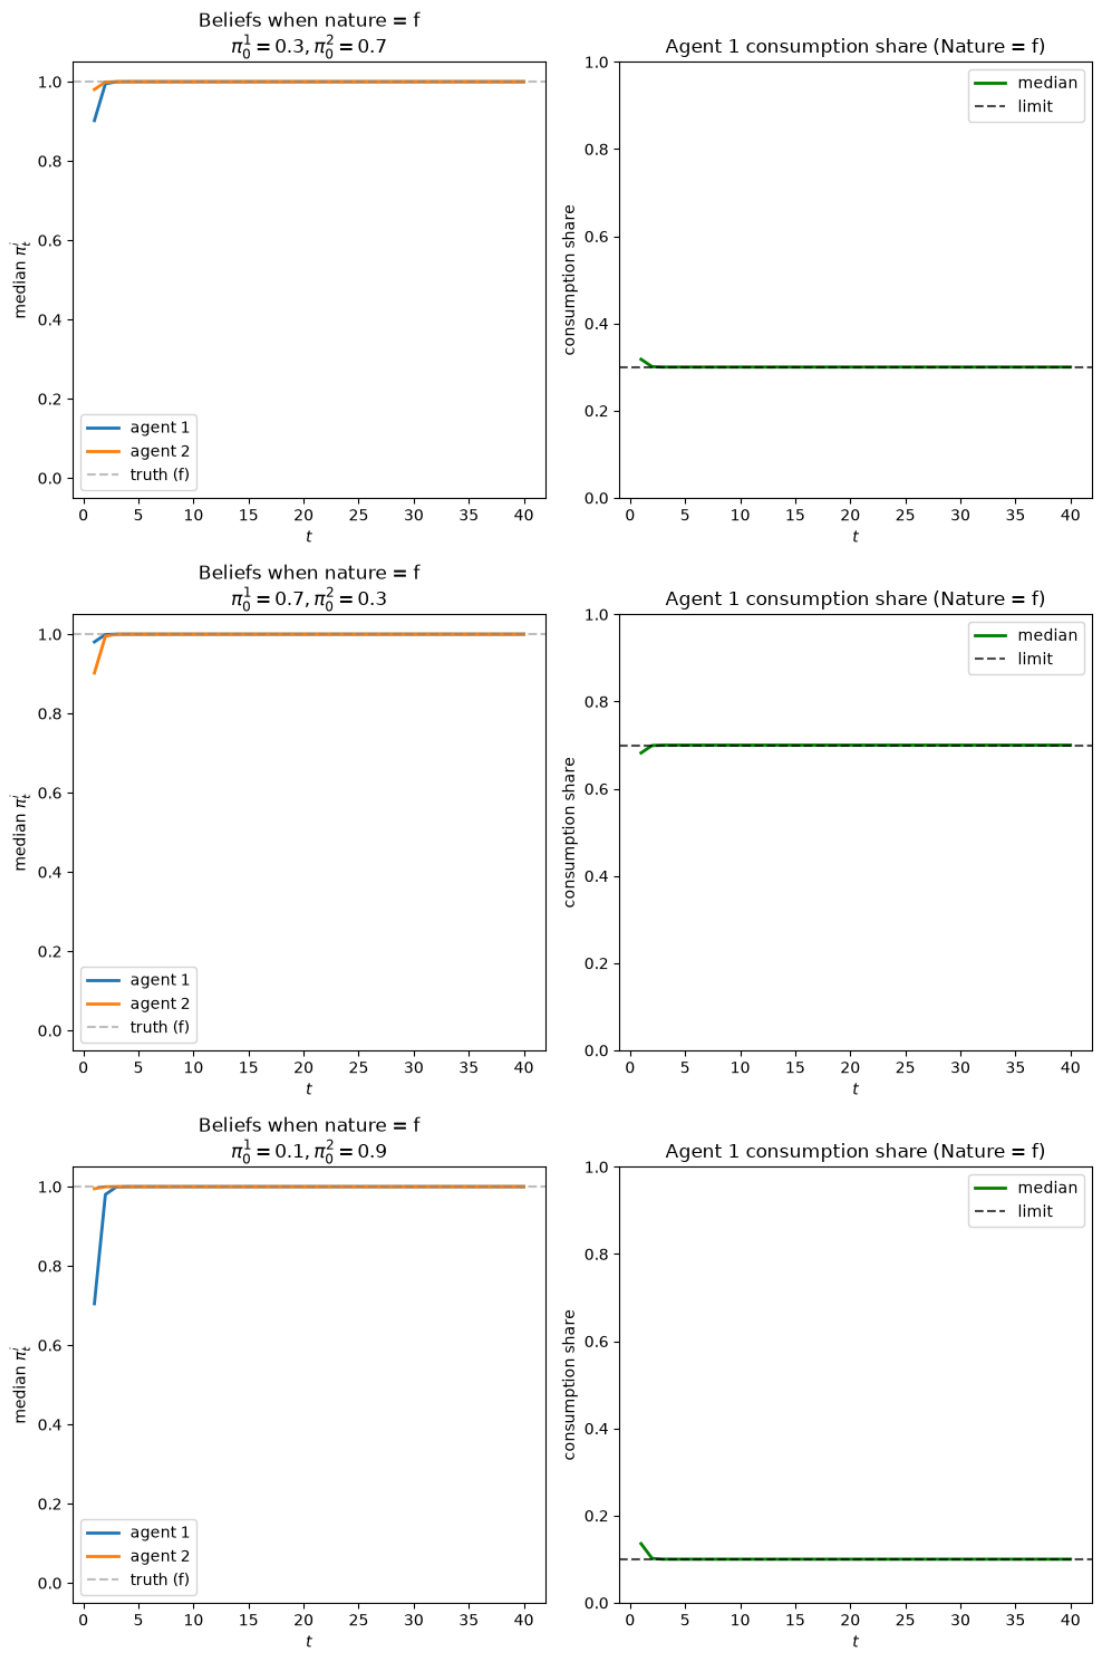
results_f = {}
results_g = {}

for i, (π_0_1, π_0_2) in enumerate(π_0_scenarios):
    # When nature follows f
    results_f[i] = simulate_learning_blume_easley(
            s_seq_f, f, g, π_0_1, π_0_2, λ)
    # When nature follows g
```

```
    results_g[i] = simulate_learning_blume_easley(
            s_seq_g, f, g, π_0_1, π_0_2, λ)
```

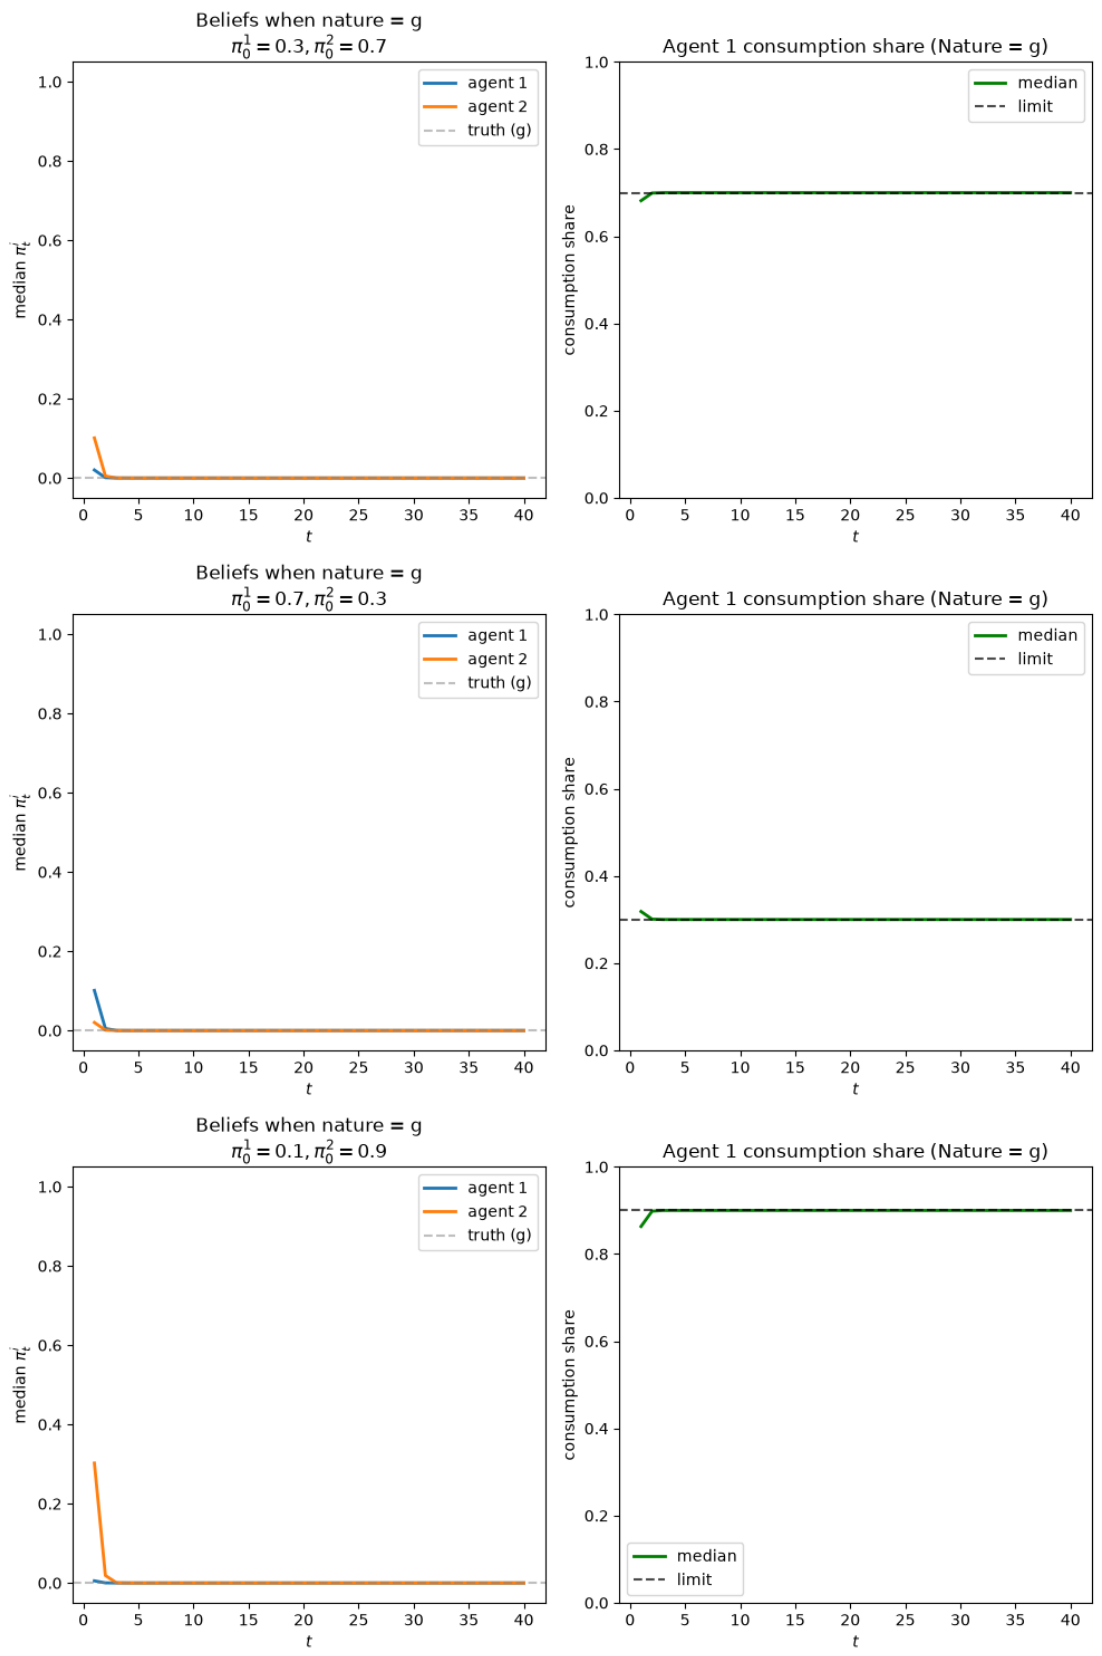
Now let's visualize the results

```
fig_f, axes_f = plot_learning_results(results_f, π_0_scenarios, 'f', λ)
plt.show()
print_limits(results_f, π_0_scenarios, 'f', λ)
```


Beliefs when nature = f
$\pi_0^1 = 0.3, \pi_0^2 = 0.7$
median π_t^i
agent 1
agent 2
truth (f)
t
Agent 1 consumption share (Nature = f)
consumption share
median
limit
Beliefs when nature = f
$\pi_0^1 = 0.7, \pi_0^2 = 0.3$
Agent 1 consumption share (Nature = f)
Beliefs when nature = f
$\pi_0^1 = 0.1, \pi_0^2 = 0.9$
Agent 1 consumption share (Nature = f)

```
nature = f, priors = (0.3, 0.7): median c1 at T = 0.300, limit = 0.300
nature = f, priors = (0.7, 0.3): median c1 at T = 0.700, limit = 0.700
nature = f, priors = (0.1, 0.9): median c1 at T = 0.100, limit = 0.100
```

```
fig_g, axes_g = plot_learning_results(results_g, π_0_scenarios, 'g', λ)
plt.show()
print_limits(results_g, π_0_scenarios, 'g', λ)
```


Beliefs when nature = g
$\pi_0^1 = 0.3, \pi_0^2 = 0.7$
agent 1
agent 2
truth (g)
median π_t^i
t
Agent 1 consumption share (Nature = g)
median
limit
consumption share
t
Beliefs when nature = g
$\pi_0^1 = 0.7, \pi_0^2 = 0.3$
agent 1
agent 2
truth (g)
median π_t^i
t
Agent 1 consumption share (Nature = g)
median
limit
consumption share
t
Beliefs when nature = g
$\pi_0^1 = 0.1, \pi_0^2 = 0.9$
agent 1
agent 2
truth (g)
median π_t^i
t
Agent 1 consumption share (Nature = g)
median
limit
consumption share
t

```
nature = g, priors = (0.3, 0.7): median c1 at T = 0.700, limit = 0.700
nature = g, priors = (0.7, 0.3): median c1 at T = 0.300, limit = 0.300
nature = g, priors = (0.1, 0.9): median c1 at T = 0.900, limit = 0.900
```

Evidently, because the two distributions are further apart, it is easier to distinguish them.

So learning occurs more quickly.

So does convergence of consumption shares.

The limits are the same as in Exercise 23.11.4 because they depend only on λ and the priors.

Exercise 23.11.6

Two agents have different beliefs about three possible models.

Assume $f(x) \geq 0$, $g(x) \geq 0$, and $h(x) \geq 0$ for $x \in X$ with:

- $\int_X f(x)dx = 1$
- $\int_X g(x)dx = 1$
- $\int_X h(x)dx = 1$

Each model asserts that s_t is IID, so that $f(s^t) = f(s_1) \cdots f(s_t)$, and similarly for g and h.

Let $\pi_0^{f,i}$, $\pi_0^{g,i}$, and $\pi_0^{h,i}$ be agent i's prior probabilities on models f, g, and h.

We'll consider two agents:

- Agent 1: $\pi_0^{f,1} = \pi_0^{g,1} = 1/2$, $\pi_0^{h,1} = 0$ (attaches positive probability only to models f and g)
- Agent 2: $\pi_0^{f,2} = \pi_0^{g,2} = \pi_0^{h,2} = 1/3$ (attaches equal weights to all three models)

Let f and g be two beta distributions with $f \sim \text{Beta}(3, 2)$ and $g \sim \text{Beta}(2, 3)$, and set $h = xf + (1 - x)g$ with $x = 0.5$.

Agent i's joint density of s^t is

$$m^i(s^t) = \pi_0^{f,i} f(s^t) + \pi_0^{g,i} g(s^t) + \pi_0^{h,i} h(s^t).$$

Bayes' Law tells us that agent i's posterior probabilities on models f and g are

$$\pi^{f,i}(s^t) = \frac{\pi_0^{f,i} f(s^t)}{m^i(s^t)}$$

and

$$\pi^{g,i}(s^t) = \frac{\pi_0^{g,i} g(s^t)}{m^i(s^t)}.$$

Please simulate and visualize evolutions of posterior probabilities and consumption allocations when:

- Nature permanently draws from f
- Nature permanently draws from g

Solution

Let's implement this three-model case with two agents having different beliefs.

Let's define f and g far apart, with h being a mixture of f and g.

```
F_a, F_b = 3, 2
G_a, G_b = 2, 3
λ = 0.5
x = 0.5     # weight on f in the mixture h

f = jit(lambda s: p(s, F_a, F_b))
g = jit(lambda s: p(s, G_a, G_b))
h = jit(lambda s: x * f(s) + (1 - x) * g(s))
```

The following function computes each agent's joint density $m^i(s^t)$, posterior probabilities, and agent 1's consumption share.

It works with logarithms of joint densities to avoid overflow and underflow over long histories.

```
def simulate_three_model_allocation(sequences, f_func, g_func, h_func,
                                    priors_1, priors_2, λ=0.5):
    """
    Simulate the Blume-Easley model with two Bayesian agents
    whose prior probabilities on models (f, g, h) are priors_1 and priors_2.

    Column t-1 of `sequences` holds s_t.
    """
    # log f(s^t), log g(s^t), log h(s^t), with shape (3, N, T)
    log_dens = np.stack([np.cumsum(np.log(dens(sequences)), axis=1)
                         for dens in (f_func, g_func, h_func)])

    results = {}
    log_m = []
    for i, priors in enumerate((priors_1, priors_2), start=1):
        # log of prior × joint density; a zero prior gives -inf
        with np.errstate(divide='ignore'):
            log_w = np.log(np.asarray(priors))[:, None, None] + log_dens
        w_max = log_w.max(axis=0)
        log_m_i = w_max + np.log(np.exp(log_w - w_max).sum(axis=0))

        # Posterior probabilities after observing s^t
        posteriors = np.exp(log_w - log_m_i)
        for k, model in enumerate(['f', 'g', 'h']):
            results[f'π_{model}_{i}'] = posteriors[k]
        log_m.append(log_m_i)

    # log l_t = log m^1(s^t) - log m^2(s^t)
    log_l = log_m[0] - log_m[1]
    results['c1_share'] = expit(np.log(λ / (1 - λ)) + log_l)
    return results
```

The following code cell defines a plotting function to show evolutions of beliefs and consumption ratios

Now let's run the simulation.

In the simulation below, agent 1 assigns positive probabilities only to f and g, while agent 2 puts equal weights on all three models.

```
T = 100
N = 1000

rng = np.random.default_rng()

# Generate sequences for nature f and g
s_seq_f = rng.beta(F_a, F_b, (N, T))
s_seq_g = rng.beta(G_a, G_b, (N, T))

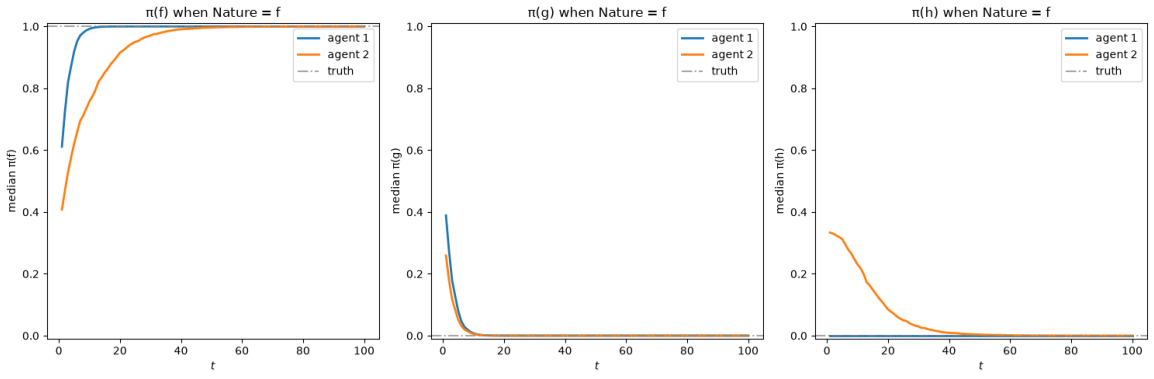
priors_1 = (0.5, 0.5, 0.0)
priors_2 = (1/3, 1/3, 1/3)

# Run simulations
results_f = simulate_three_model_allocation(s_seq_f, f, g, h,
                                            priors_1, priors_2, λ)
results_g = simulate_three_model_allocation(s_seq_g, f, g, h,
                                            priors_1, priors_2, λ)
```

Plots below show the evolution of beliefs for each model (f, g, h) separately.

First we show the figure when nature chooses f

```
plot_belief_evolution(results_f, nature='f', figsize=(15, 5))
plt.show()
```



Agent 1's posterior beliefs are depicted in blue and agent 2's posterior beliefs are depicted in orange.

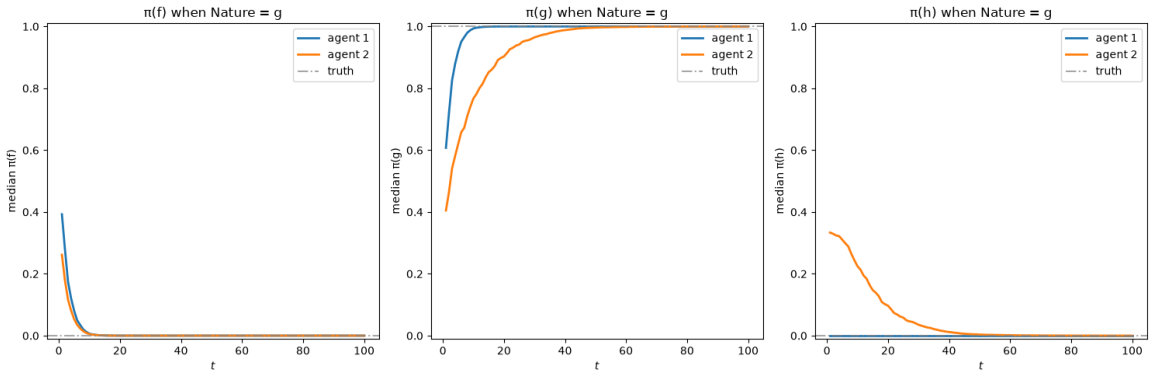
Evidently, when nature draws from f, agent 1 learns faster than agent 2, who, unlike agent 1, attaches a positive prior probability to model h:

- In the leftmost panel, both agents' beliefs for $\pi(f)$ converge toward 1 (the truth)
- Agent 2's belief in model h (rightmost panel) equals $1/3$ after the first draw and then falls toward 0

Agent 2's posterior on h after one draw is $h(s_1)/[f(s_1) + g(s_1) + h(s_1)] = 1/3$ because $f + g = 2h$.

Now let's plot the belief evolution when nature chooses g:

```
plot_belief_evolution(results_g, nature='g', figsize=(15, 5))
plt.show()
```



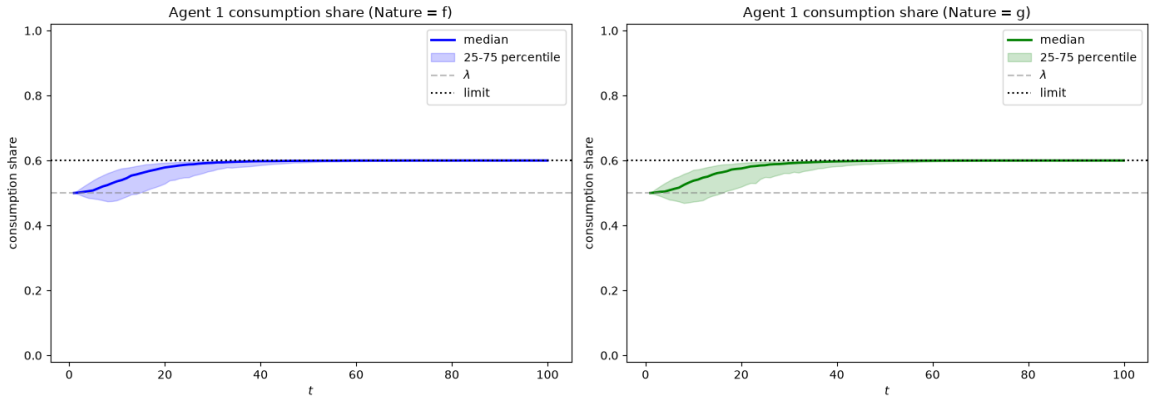
Again, agent 1 learns faster than agent 2.

Before reading the next figure, please guess how consumption shares evolve.

Remember that agent 1 reaches the correct model faster than agent 2

```
c1_lim = λ * priors_1[0] / (λ * priors_1[0] + (1 - λ) * priors_2[0])
plot_consumption_dynamics(results_f, results_g, λ=0.5,
                          c1_limits=(c1_lim, c1_lim), figsize=(14, 5))
plt.show()

for name, res in [('f', results_f), ('g', results_g)]:
    print(f"nature = {name}: median c1 at T = "
          f"{np.median(res['c1_share'][:, -1]):.3f}, limit = {c1_lim:.3f}")
```




```
nature = f: median c1 at T = 0.600, limit = 0.600
nature = g: median c1 at T = 0.600, limit = 0.600
```

As we expected, agent 1 has a higher consumption share compared to agent 2.

In this exercise, the "truth" is among possible outcomes according to both agents.

Agent 2's model is "more general" because it allows a possibility – that nature is drawing from h – that agent 1's model does not include.

When nature draws from f, $g(s^t)/f(s^t) \to 0$ and $h(s^t)/f(s^t) \to 0$, so $m^i(s^t)/f(s^t) \to \pi_0^{f,i}$.

The likelihood ratio l_t converges to $\pi_0^{f,1}/\pi_0^{f,2} = (1/2)/(1/3) = 1.5$, and agent 1's consumption share converges to $0.5 \times 1.5/(0.5 \times 1.5 + 0.5) = 0.6$.

The same limit prevails when nature draws from g.

Both agents survive.

Agent 1's limiting share exceeds agent 2's because agent 1 puts prior probability $1/2$ on the true model while agent 2 puts $1/3$.

Agent 2 pays for model h with prior probability taken from the true model.

Exercise 23.11.7

Now consider two agents with extreme priors about three models.

Consider the same setup as the previous exercise, but now:

- Agent 1: $\pi_0^{f,1} = \pi_0^{g,1} = \frac{\epsilon}{2} > 0$ and $\pi_0^{h,1} = 1 - \epsilon$, where ϵ is close to 0 (e.g., $\epsilon = 0.01$)
- Agent 2: $\pi_0^{f,2} = \pi_0^{g,2} = 0$ and $\pi_0^{h,2} = 1$ (rigid belief in model h)

Set $f \sim \text{Beta}(1, 1)$, $g \sim \text{Beta}(3, 1.2)$, and $h \sim \text{Beta}(1.2, 1.1)$.

Model h is close to f and far from g as measured by KL divergence.

Please simulate and visualize evolutions of posterior probabilities and consumption allocations when:

- Nature permanently draws from f
- Nature permanently draws from g

Solution

To explore this exercise, we increase T to 1000.

Let's specify f, g, and h and verify that h and f are closer than h and g

```
F_a, F_b = 1, 1
G_a, G_b = 3, 1.2
H_a, H_b = 1.2, 1.1

f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))
h = jit(lambda x: p(x, H_a, H_b))

Kh_f = compute_KL(h, f)
Kh_g = compute_KL(h, g)
Kf_h = compute_KL(f, h)
Kg_h = compute_KL(g, h)

print(f"KL divergences:")
print(f"KL(h,f) = {Kh_f:.4f}, KL(h,g) = {Kh_g:.4f}")
print(f"KL(f,h) = {Kf_h:.4f}, KL(g,h) = {Kg_h:.4f}")
```

```
KL divergences:
KL(h,f) = 0.0092, KL(h,g) = 0.5514
KL(f,h) = 0.0105, KL(g,h) = 0.2919
```

Now we can set the belief models for the two agents

```
ε = 0.01
λ = 0.5

# Agent 1: π_f = ε/2, π_g = ε/2, π_h = 1-ε
# (almost rigid about h)
priors_1 = (ε/2, ε/2, 1 - ε)

# Agent 2: π_f = 0, π_g = 0, π_h = 1
# (fully rigid about h)
priors_2 = (0.0, 0.0, 1.0)
```

Now we can run the simulation

```
T = 1000
N = 1000

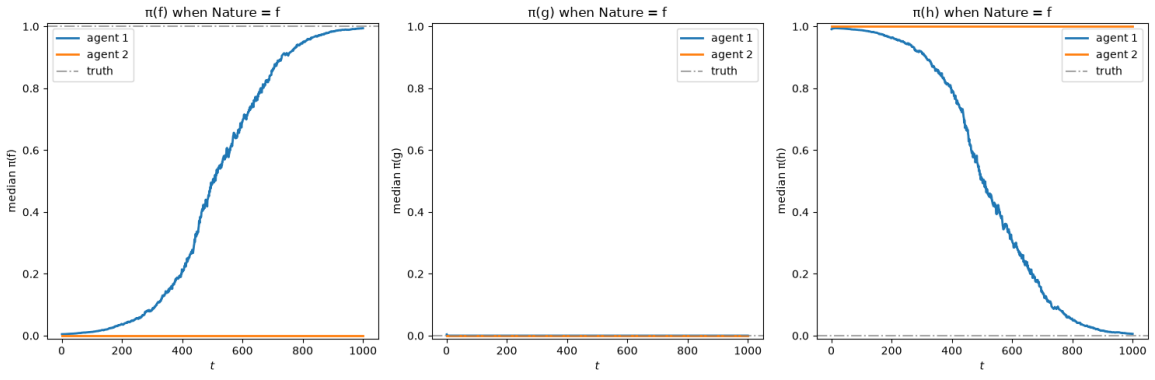
rng = np.random.default_rng()

# Generate sequences for different nature scenarios
s_seq_f = rng.beta(F_a, F_b, (N, T))
s_seq_g = rng.beta(G_a, G_b, (N, T))

# Run simulations for both scenarios
results_f = simulate_three_model_allocation(s_seq_f, f, g, h,
                                            priors_1, priors_2, λ)
results_g = simulate_three_model_allocation(s_seq_g, f, g, h,
                                            priors_1, priors_2, λ)
```

Let's plot the belief evolution when nature chooses f

```
plot_belief_evolution(results_f, nature='f', figsize=(15, 5))
plt.show()
```



Agent 1 learns the truth slowly in the leftmost panel showing $\pi(f)$.

Agent 2 does not update at all.

Bayes' law cannot revise a prior probability of zero, so agent 2's posterior probability on h stays at 1 whatever the data.

Agent 2's dogmatic prior keeps it from learning.

The closeness of h to f sets only how long agent 1 takes to overcome its own small prior on f.

Agent 1's posterior on f reaches $1/2$ roughly when $\log[f(s^t)/h(s^t)]$ reaches $\log[2(1-\epsilon)/\epsilon]$.

When nature draws from f, $\log[f(s^t)/h(s^t)]$ grows at rate $KL(f,h)$ per period.

The following code computes the implied waiting times, along with the waiting times for a nearly dogmatic agent who puts prior probability 10^{-10} on each of f and g.

```
def first_crossing(post, level=0.5):
    """First date at which the median posterior exceeds level."""
    above = np.where(np.median(post, axis=0) > level)[0]
    return above[0] + 1 if above.size > 0 else None

for nature, K in [('f', Kf_h), ('g', Kg_h)]:
    print(f"nature = {nature}: log(2(1-ε)/ε)/KL = {np.log(2*(1-ε)/ε)/K:.0f}, "
          f"log(1e10)/KL = {np.log(1e10)/K:.0f}")

print("first t with median π(f) of agent 1 > 1/2 when nature = f:",
      first_crossing(results_f['π_f_1']))
print("first t with median π(g) of agent 1 > 1/2 when nature = g:",
      first_crossing(results_g['π_g_1']))
```

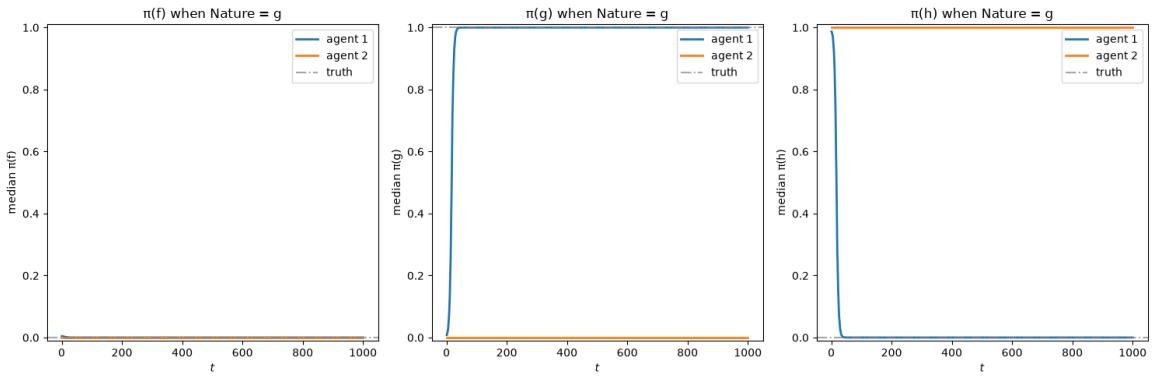
```
nature = f: log(2(1-ε)/ε)/KL = 502, log(1e10)/KL = 2186
nature = g: log(2(1-ε)/ε)/KL = 18, log(1e10)/KL = 79
first t with median π(f) of agent 1 > 1/2 when nature = f: 502
first t with median π(g) of agent 1 > 1/2 when nature = g: 18
```

Agent 1 needs about 500 periods to put probability $1/2$ on f when nature draws from f.

A nearly dogmatic agent 2 would need about 2200 periods.

Now let's plot the belief evolution when nature chooses g

```
plot_belief_evolution(results_g, nature='g', figsize=(15, 5))
plt.show()
```



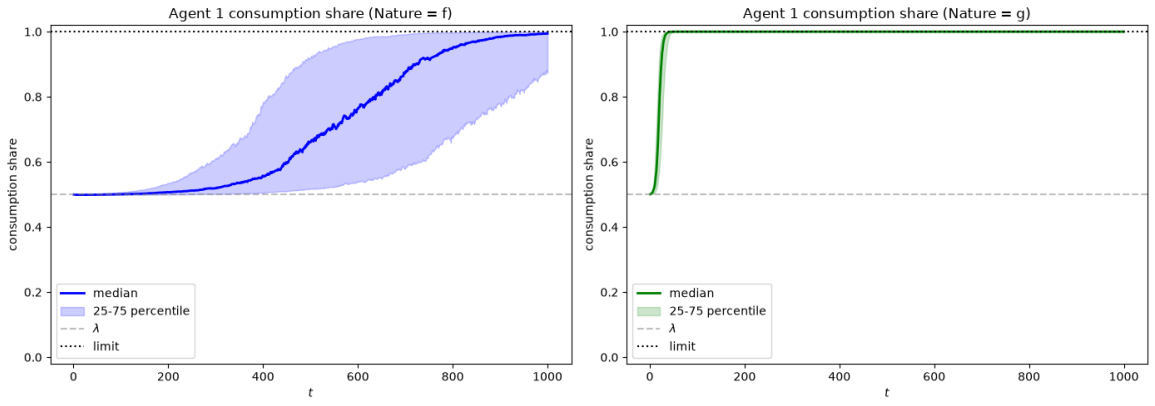
When nature draws from g, $KL(g,h) = 0.2919$ is much larger than $KL(f,h) = 0.0105$.

Agent 1 puts probability $1/2$ on g after about 18 periods.

Agent 2 again does not update.

A nearly dogmatic agent 2 would need about 79 periods.

```
plot_consumption_dynamics(results_f, results_g, λ=0.5,
                          c1_limits=(1.0, 1.0), figsize=(14, 5))
plt.show()
```



Agent 1's consumption share converges to 1 both when nature permanently draws from f and when nature permanently draws from g.

Agent 2's probability law over infinite histories puts probability zero on a set of histories that has probability one when nature draws from f, and likewise for g.

Nature's probability law is not absolutely continuous with respect to agent 2's, and agent 2's predictions never merge with the truth (see *Merging of Opinions: The Blackwell–Dubins Theorem*).

The likelihood ratio $l_t = m^1(s^t)/h(s^t)$ eventually grows at rate $KL(f,h)$ when nature draws from f and at rate $KL(g,h)$ when nature draws from g.

A nearly dogmatic agent 2 with prior probability 10^{-10} on each of f and g would eventually learn the truth and survive.

Its limiting consumption share would be $(1-\lambda)10^{-10}/[\lambda\epsilon/2 + (1-\lambda)10^{-10}] \approx 2 \times 10^{-8}$.

CHAPTER

TWENTYFOUR

LIKELIHOOD PROCESSES FOR VAR MODELS

Contents

- *Likelihood Processes For VAR Models*
  - *Overview*
  - *VAR model setup*
  - *Likelihood ratio process*
  - *Example 1: two AR(1) processes*
  - *Example 2: bivariate VAR models*
  - *KL divergence rate*
  - *Application: Samuelson multiplier-accelerator*
  - *Exercises*

24.1 Overview

This lecture extends our analysis of likelihood ratio processes to Vector Autoregressions (VARs).

We'll

- Construct likelihood functions for VAR models
- Form likelihood ratio processes for comparing two VAR models
- Visualize the evolution of likelihood ratios over time
- Compute the Kullback–Leibler divergence rate that governs the drift of a log likelihood ratio
- Connect VAR likelihood ratios to the Samuelson multiplier-accelerator model

Our analysis builds on concepts from:

- *Likelihood Ratio Processes*
- *Statistical Divergence Measures*
- *Linear State Space Models*
- *Samuelson Multiplier-Accelerator*

Let's start by importing helpful libraries:

```
import numpy as np
import matplotlib.pyplot as plt
from scipy import linalg
from scipy.stats import multivariate_normal as mvn
from collections import namedtuple

rng = np.random.default_rng(1234)
```

24.2 VAR model setup

Consider a VAR model of the form:

$$x_{t+1} = Ax_t + Cw_{t+1}$$
$$x_0 \sim \mathcal{N}(\mu_0, \Sigma_0)$$

where:

- x_t is an $n \times 1$ state vector
- $w_{t+1} \sim \mathcal{N}(0, I)$ is an $m \times 1$ vector of shocks
- A is an $n \times n$ transition matrix
- C is an $n \times m$ volatility matrix

Let's define the necessary data structures for the VAR model

```
VARModel = namedtuple('VARModel', ['A', 'C', 'μ_0', 'Σ_0',
                                   'CC', 'CC_inv', 'log_det_CC',
                                   'Σ_0_inv', 'log_det_Σ_0'])
def compute_stationary_var(A, C):
    """
    Compute stationary mean and covariance for VAR model
    """
    n = A.shape[0]

    # Check stability
    eigenvalues = np.linalg.eigvals(A)
    if np.max(np.abs(eigenvalues)) >= 1:
        raise ValueError("VAR is not stationary")

    μ_0 = np.zeros(n)

    # Stationary covariance: solve discrete Lyapunov equation
    # Σ_0 = A @ Σ_0 @ A.T + C @ C.T
    CC = C @ C.T
    Σ_0 = linalg.solve_discrete_lyapunov(A, CC)

    return μ_0, Σ_0

def create_var_model(A, C, μ_0=None, Σ_0=None, stationary=True):
    """
    Create a VAR model with parameters and precomputed matrices
    """
    A = np.asarray(A)
```

(continues on next page)

(continued from previous page)

```
    C = np.asarray(C)
    n = A.shape[0]
    CC = C @ C.T

    if stationary:
        μ_0_comp, Σ_0_comp = compute_stationary_var(A, C)
    else:
        μ_0_comp = μ_0 if μ_0 is not None else np.zeros(n)
        Σ_0_comp = Σ_0 if Σ_0 is not None else np.eye(n)

    # Check if CC is singular
    det_CC = np.linalg.det(CC)
    if np.abs(det_CC) < 1e-10:
        # Use pseudo-inverse for singular case
        CC_inv = np.linalg.pinv(CC)
        CC_reg = CC + 1e-10 * np.eye(CC.shape[0])
        log_det_CC = np.log(np.linalg.det(CC_reg))
    else:
        CC_inv = np.linalg.inv(CC)
        log_det_CC = np.log(det_CC)

    # Same check for Σ_0
    det_Σ_0 = np.linalg.det(Σ_0_comp)
    if np.abs(det_Σ_0) < 1e-10:
        Σ_0_inv = np.linalg.pinv(Σ_0_comp)
        Σ_0_reg = Σ_0_comp + 1e-10 * np.eye(Σ_0_comp.shape[0])
        log_det_Σ_0 = np.log(np.linalg.det(Σ_0_reg))
    else:
        Σ_0_inv = np.linalg.inv(Σ_0_comp)
        log_det_Σ_0 = np.log(det_Σ_0)

    return VARModel(A=A, C=C, μ_0=μ_0_comp, Σ_0=Σ_0_comp,
                    CC=CC, CC_inv=CC_inv, log_det_CC=log_det_CC,
                    Σ_0_inv=Σ_0_inv, log_det_Σ_0=log_det_Σ_0)
```

24.2.1 Joint distribution

The joint probability density $p(x_T, x_{T-1}, \dots, x_0)$ can be factored as:

$$p(x_T, \dots, x_0) = p(x_T|x_{T-1})p(x_{T-1}|x_{T-2}) \cdots p(x_1|x_0)p(x_0)$$

Since the VAR is Markovian, $p(x_{t+1}|x_t, \dots, x_0) = p(x_{t+1}|x_t)$.

We write p for a generic density and reserve f and g for the names of two models.

24.2.2 Conditional densities

Given the Gaussian structure, the conditional distribution $p(x_{t+1}|x_t)$ is Gaussian with:

- Mean: Ax_t
- Covariance: $CC^\top$

The log conditional density is

$$\log p(x_{t+1}|x_t) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det(CC^\top) - \frac{1}{2}(x_{t+1} - Ax_t)^\top (CC^\top)^{-1}(x_{t+1} - Ax_t) \tag{24.1}$$

```
def log_likelihood_transition(x_next, x_curr, model):
    """
    Compute log likelihood of transition from x_curr to x_next
    """
    x_next = np.atleast_1d(x_next)
    x_curr = np.atleast_1d(x_curr)
    n = len(x_next)
    diff = x_next - model.A @ x_curr
    return -0.5 * (n * np.log(2 * np.pi) + model.log_det_CC +
                   diff @ model.CC_inv @ diff)
```

The log density of the initial state is:

$$\log p(x_0) = -\frac{n}{2}\log(2\pi) - \frac{1}{2}\log\det(\Sigma_0) - \frac{1}{2}(x_0 - \mu_0)^\top \Sigma_0^{-1}(x_0 - \mu_0) \tag{24.2}$$

```
def log_likelihood_initial(x_0, model):
    """
    Compute log likelihood of initial state
    """
    x_0 = np.atleast_1d(x_0)
    n = len(x_0)
    diff = x_0 - model.μ_0
    return -0.5 * (n * np.log(2 * np.pi) + model.log_det_Σ_0 +
                   diff @ model.Σ_0_inv @ diff)
```

The following function simulates paths from a VAR model

```
def simulate_var(model, T, rng, N_paths=1):
    """
    Simulate paths from the VAR model
    """
    n = model.A.shape[0]
    m = model.C.shape[1]
    paths = np.zeros((N_paths, T+1, n))

    for i in range(N_paths):
        # Draw initial state
        x = mvn.rvs(mean=model.μ_0, cov=model.Σ_0, random_state=rng)
        x = np.atleast_1d(x)
        paths[i, 0] = x

        # Simulate forward
        for t in range(T):
            w = rng.standard_normal(m)
            x = model.A @ x + model.C @ w
```

(continues on next page)

(continued from previous page)

```
            paths[i, t+1] = x

    return paths if N_paths > 1 else paths[0]
```

24.3 Likelihood ratio process

Now let's compute likelihood ratio processes for comparing two VAR models f and g.

Let p_f and p_g denote densities under models f and g respectively.

As in *Likelihood Ratio Processes*, L_t denotes the likelihood ratio of the sample $x_0, x_1, \ldots, x_t$:

$$L_t = \frac{p_f(x_t, \ldots, x_0)}{p_g(x_t, \ldots, x_0)} = \frac{p_f(x_0)}{p_g(x_0)} \prod_{s=1}^{t} \frac{p_f(x_s|x_{s-1})}{p_g(x_s|x_{s-1})}$$

where the conditional densities are given by (24.1) and the densities of x_0 by (24.2).

The log likelihood ratio process is

$$\log L_t = \log \frac{p_f(x_0)}{p_g(x_0)} + \sum_{s=1}^{t} \ell_s, \qquad \ell_s = \log \frac{p_f(x_s|x_{s-1})}{p_g(x_s|x_{s-1})} \tag{24.3}$$

It obeys the recursion $\log L_t = \log L_{t-1} + \ell_t$ with initial condition $\log L_0 = \log p_f(x_0) - \log p_g(x_0)$.

Let's write those equations in Python

```
def compute_likelihood_ratio_var(paths, model_f, model_g):
    """
    Compute likelihood ratio process for VAR models
    """
    if paths.ndim == 2:
        paths = paths[np.newaxis, :]

    N_paths, T_plus_1, n = paths.shape
    T = T_plus_1 - 1
    log_L_ratios = np.zeros((N_paths, T+1))

    for i in range(N_paths):
        X = paths[i]

        # Initial log likelihood ratio
        log_L_f_0 = log_likelihood_initial(X[0], model_f)
        log_L_g_0 = log_likelihood_initial(X[0], model_g)
        log_L_ratios[i, 0] = log_L_f_0 - log_L_g_0

        # Recursive computation
        for t in range(1, T+1):
            log_L_f_t = log_likelihood_transition(X[t], X[t-1], model_f)
            log_L_g_t = log_likelihood_transition(X[t], X[t-1], model_g)

            # Update log likelihood ratio
            log_diff = log_L_f_t - log_L_g_t

            log_L_prev = log_L_ratios[i, t-1]
            log_L_new = log_L_prev + log_diff
```

(continues on next page)

(continued from previous page)

```
            log_L_ratios[i, t] = log_L_new

    return log_L_ratios if N_paths > 1 else log_L_ratios[0]
```

24.4 Example 1: two AR(1) processes

Let's start with a simple example comparing two univariate AR(1) processes with $A_f = 0.8$, $A_g = 0.5$, and $C_f = 0.3$, $C_g = 0.4$

```
# Model f: AR(1) with persistence ρ = 0.8
A_f = np.array([[0.8]])
C_f = np.array([[0.3]])

# Model g: AR(1) with persistence ρ = 0.5
A_g = np.array([[0.5]])
C_g = np.array([[0.4]])

# Create VAR models
model_f = create_var_model(A_f, C_f)
model_g = create_var_model(A_g, C_g)
```

Let's generate 100 paths of length 200 from model f and compute the likelihood ratio processes

```
# Simulate from model f
T = 200
N_paths = 100
paths_from_f = simulate_var(model_f, T, rng, N_paths)

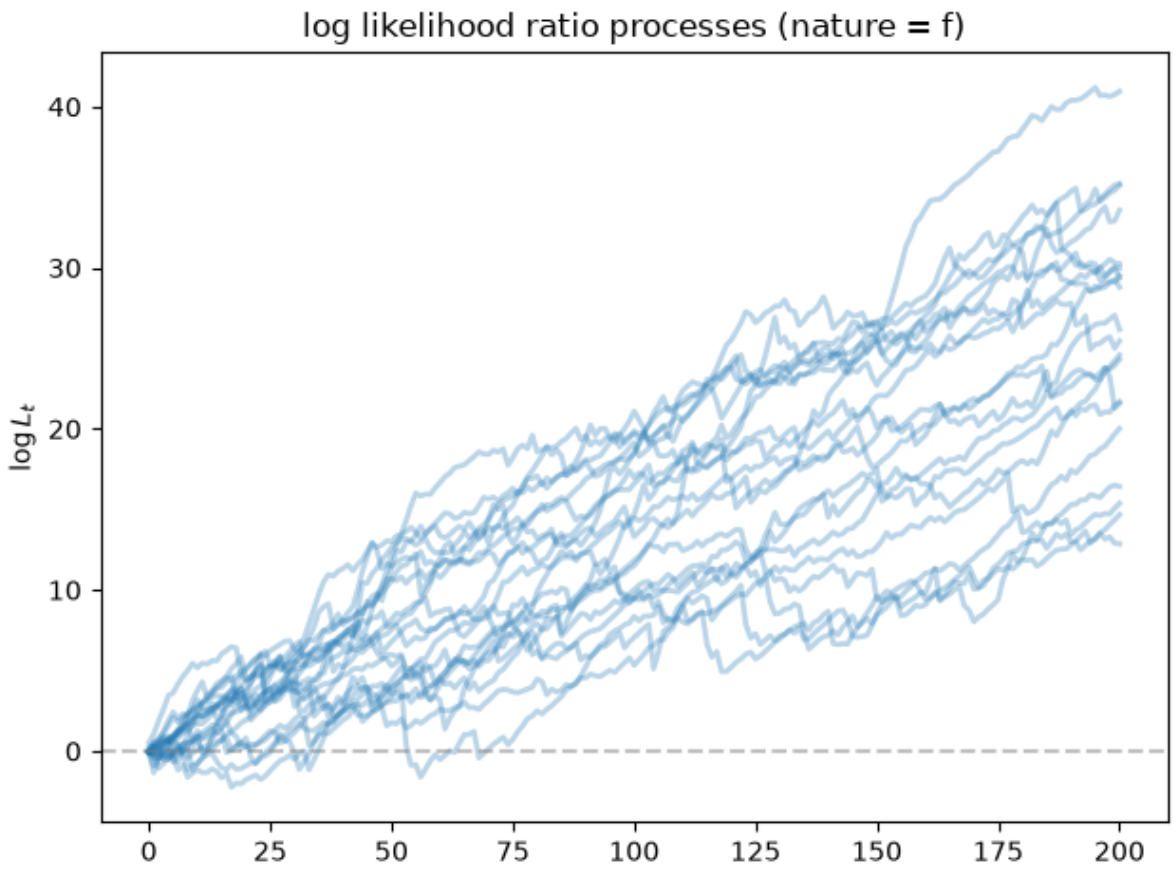
L_ratios_f = compute_likelihood_ratio_var(paths_from_f, model_f, model_g)

fig, ax = plt.subplots()

for i in range(min(20, N_paths)):
    ax.plot(L_ratios_f[i], alpha=0.3, color='C0', lw=2)

ax.axhline(y=0, color='gray', linestyle='--', alpha=0.5)
ax.set_ylabel(r'$\log L_t$')
ax.set_title('log likelihood ratio processes (nature = f)')

plt.tight_layout()
plt.show()
```



As we expected, the log likelihood ratio processes go to $+\infty$ as t increases, indicating that model f is chosen correctly by our algorithm.

24.5 Example 2: bivariate VAR models

Now let's consider an example with bivariate VAR models with

$$A_f = \begin{bmatrix} 0.7 & 0.2 \\ 0.1 & 0.6 \end{bmatrix}, \quad C_f = \begin{bmatrix} 0.3 & 0.1 \\ 0.1 & 0.3 \end{bmatrix}$$

and

$$A_g = \begin{bmatrix} 0.5 & 0.3 \\ 0.2 & 0.5 \end{bmatrix}, \quad C_g = \begin{bmatrix} 0.4 & 0.0 \\ 0.0 & 0.4 \end{bmatrix}$$

```
A_f = np.array([[0.7, 0.2],
                [0.1, 0.6]])

C_f = np.array([[0.3, 0.1],
                [0.1, 0.3]])

A_g = np.array([[0.5, 0.3],
                [0.2, 0.5]])

C_g = np.array([[0.4, 0.0],
```

(continues on next page)

(continued from previous page)

```
                [0.0, 0.4]])

# Create VAR models
model2_f = create_var_model(A_f, C_f)
model2_g = create_var_model(A_g, C_g)

# Check stationarity
print("model f eigenvalues:", np.linalg.eigvals(A_f))
print("model g eigenvalues:", np.linalg.eigvals(A_g))
```

```
model f eigenvalues: [0.8 0.5]
model g eigenvalues: [0.74494897 0.25505103]
```

Let's generate 50 paths of length 50 from both models and compute the likelihood ratio processes

```
# Simulate from both models
T = 50
N_paths = 50

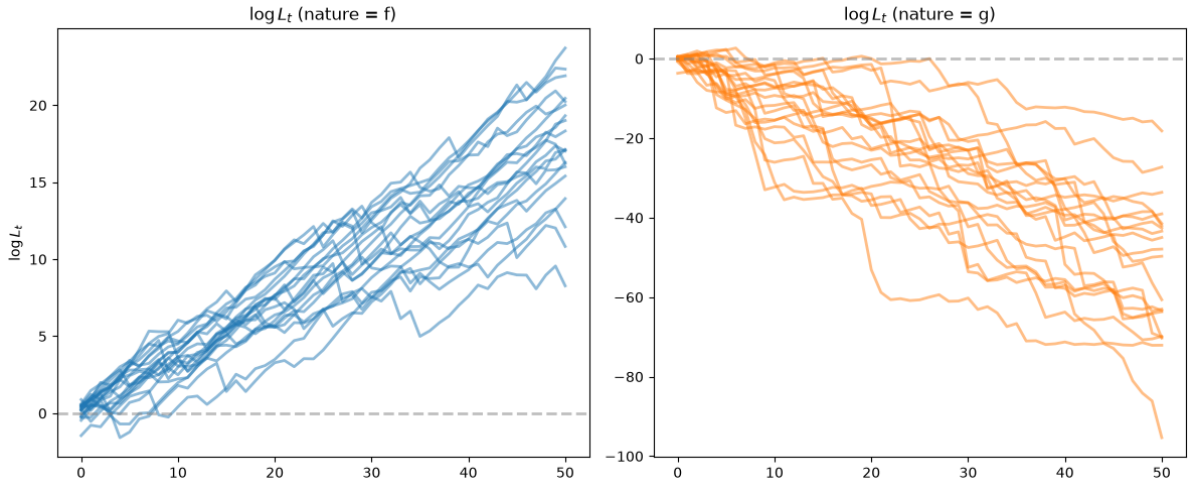
paths_from_f = simulate_var(model2_f, T, rng, N_paths)
paths_from_g = simulate_var(model2_g, T, rng, N_paths)

# Compute likelihood ratios
L_ratios_ff = compute_likelihood_ratio_var(paths_from_f, model2_f, model2_g)
L_ratios_gf = compute_likelihood_ratio_var(paths_from_g, model2_f, model2_g)
```

```
# Visualize the results
fig, axes = plt.subplots(1, 2, figsize=(12, 5))

ax = axes[0]
for i in range(min(20, N_paths)):
    ax.plot(L_ratios_ff[i], alpha=0.5, color='C0', lw=2)
ax.axhline(y=0, color='gray', linestyle='--', alpha=0.5, lw=2)
ax.set_title(r'$\log L_t$ (nature = f)')
ax.set_ylabel(r'$\log L_t$')

ax = axes[1]
for i in range(min(20, N_paths)):
    ax.plot(L_ratios_gf[i], alpha=0.5, color='C1', lw=2)
ax.axhline(y=0, color='gray', linestyle='--', alpha=0.5, lw=2)
ax.set_title(r'$\log L_t$ (nature = g)')
plt.tight_layout()
plt.show()
```



For paths generated from model f, the log likelihood ratio processes tend to go to $+\infty$, while for paths from model g, they tend to go to $-\infty$.

The paths fall faster under g than they rise under f; the section *KL divergence rate* explains why.

Let's apply the decision rule described in the section on *model selection mistake probability* in *Likelihood Ratio Processes*: select model f when $\log L_T \geq 0$ and model g when $\log L_T < 0$.

This rule is a likelihood ratio test with threshold $L_T = 1$.

Suppose that a Bayesian puts prior probability $1/2$ on each model.

Then the posterior probability of model f is $L_T/(1 + L_T)$, as in *Likelihood Ratio Processes and Bayesian Learning*.

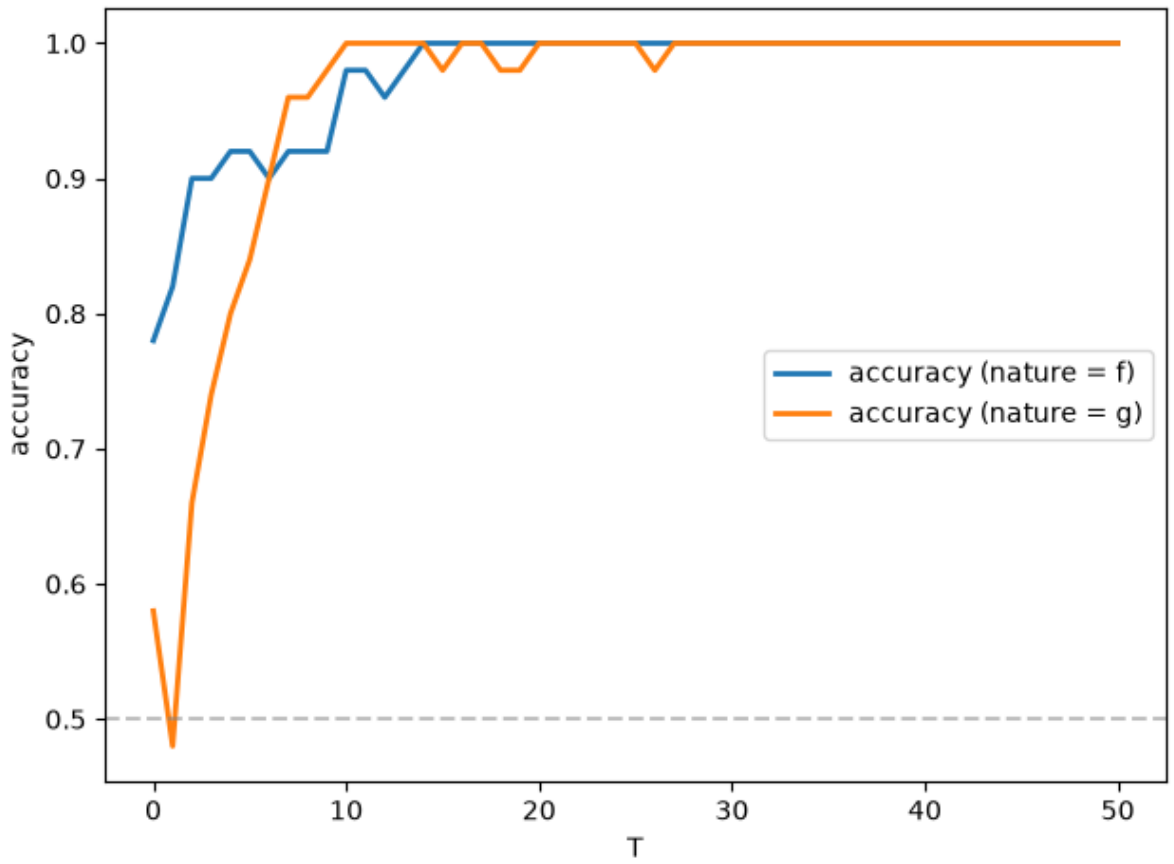
The rule selects the model with the higher posterior probability, so it minimizes the probability $\frac{1}{2}(\alpha_T + \beta_T)$ of selecting the wrong model, where $\alpha_T = \text{Prob}(\log L_T < 0|f)$ and $\beta_T = \text{Prob}(\log L_T \geq 0|g)$.

```
fig, ax = plt.subplots()
T_values = np.arange(0, T+1)
accuracy_f = np.zeros(len(T_values))
accuracy_g = np.zeros(len(T_values))

for i, t in enumerate(T_values):
    # Correct selection when data from f
    accuracy_f[i] = np.mean(L_ratios_ff[:, t] > 0)
    # Correct selection when data from g
    accuracy_g[i] = np.mean(L_ratios_gf[:, t] < 0)

ax.plot(T_values, accuracy_f, 'C0', linewidth=2, label='accuracy (nature = f)')
ax.plot(T_values, accuracy_g, 'C1', linewidth=2, label='accuracy (nature = g)')
ax.axhline(y=0.5, color='gray', linestyle='--', alpha=0.5)
ax.set_xlabel('T')
ax.set_ylabel('accuracy')
ax.legend()

plt.tight_layout()
plt.show()
```



Evidently, the accuracy approaches 1 as T increases, and it does so very quickly.

Let's also check the type I and type II errors α_T and β_T as functions of T

```
def model_selection_analysis(T_values, model_f, model_g, N_sim=500):
    """
    Analyze model selection performance for different sample sizes
    """
    errors_f = []  # Type I errors
    errors_g = []  # Type II errors

    for T in T_values:
        # Simulate from model f
        paths_f = simulate_var(model_f, T, rng, N_sim//2)
        L_ratios_f = compute_likelihood_ratio_var(paths_f, model_f, model_g)

        # Simulate from model g
        paths_g = simulate_var(model_g, T, rng, N_sim//2)
        L_ratios_g = compute_likelihood_ratio_var(paths_g, model_f, model_g)

        # Decision rule: choose f if log L_T >= 0
        errors_f.append(np.mean(L_ratios_f[:, -1] < 0))
        errors_g.append(np.mean(L_ratios_g[:, -1] >= 0))

    return np.array(errors_f), np.array(errors_g)

T_values = np.arange(1, 50, 1)
```

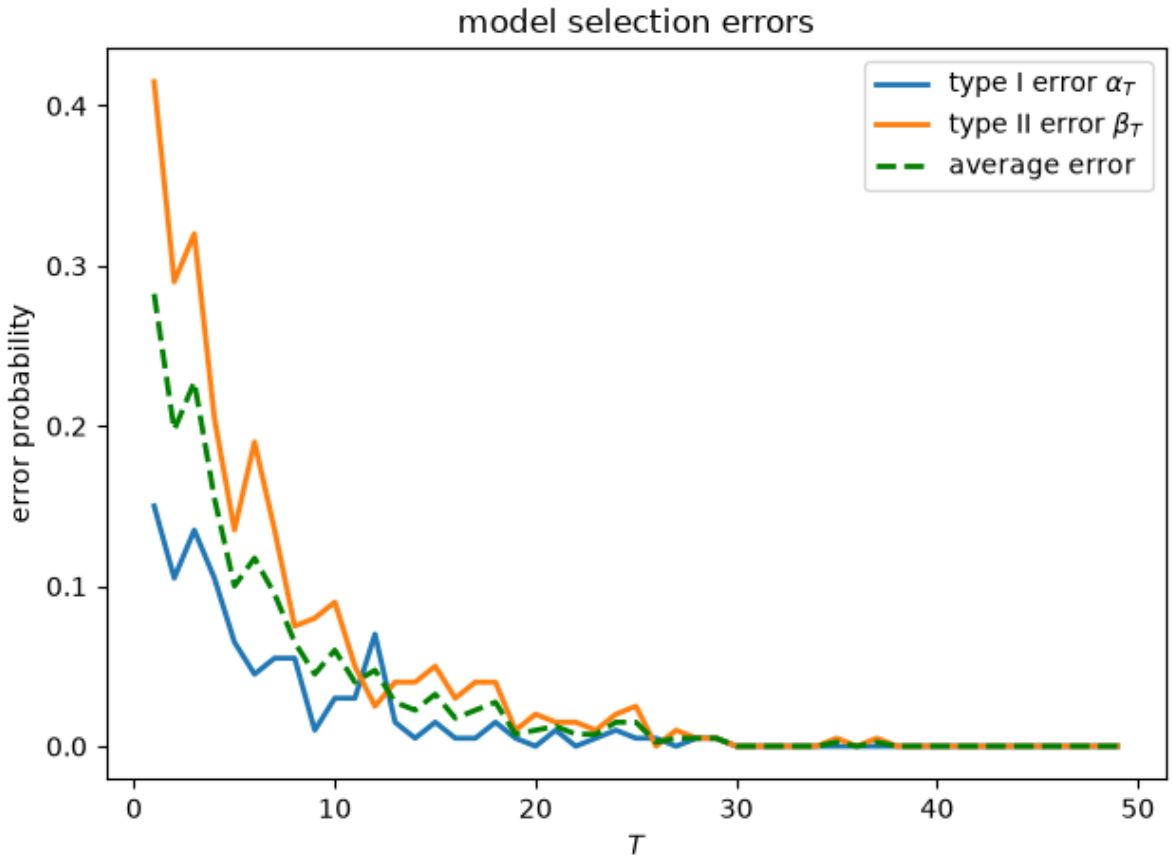
(continues on next page)

(continued from previous page)

```
errors_f, errors_g = model_selection_analysis(T_values, model2_f, model2_g, N_sim=400)

fig, ax = plt.subplots()

ax.plot(T_values, errors_f, 'C0', linewidth=2, label=r'type I error $\alpha_T$')
ax.plot(T_values, errors_g, 'C1', linewidth=2, label=r'type II error $\beta_T$')
ax.plot(T_values, 0.5 * (errors_f + errors_g), 'g--',
linewidth=2, label='average error')
ax.set_xlabel('$T$')
ax.set_ylabel('error probability')
ax.set_title('model selection errors')
ax.legend()
plt.tight_layout()
plt.show()
```



The decision rule above fixes the sample size T in advance.

Wald's sequential probability ratio test, described in *A Problem that Stumped Milton Friedman*, instead stops sampling the first time $\log L_t$ leaves an interval.

Exercise Exercise **??** of that lecture applies it to two VAR models.

24.6 KL divergence rate

The average slopes of the log likelihood ratio paths in the figures above are Kullback–Leibler divergence rates.

They are VAR counterparts of the *KL divergence rate for Markov chains* in *Likelihood Ratio Processes*.

Statistical Divergence Measures describes Kullback–Leibler divergence.

Let $\Sigma_f = C_f C_f^\top$ and $\Sigma_g = C_g C_g^\top$, and let V_f solve the discrete Lyapunov equation

$$V_f = A_f V_f A_f^\top + \Sigma_f$$

so that $\mathcal{N}(0, V_f)$ is the stationary distribution of x_t under model f.

Suppose that model f generates the data and that x_0 is drawn from $\mathcal{N}(0, V_f)$, as in Examples 1 and 2.

Then every increment ℓ_t in (24.3) has the same expected value

$$\begin{aligned} h_{KL}(f,g) &= E_f\left[\log \frac{p_f(x_t|x_{t-1})}{p_g(x_t|x_{t-1})}\right] \\ &= \frac{1}{2}\left[\operatorname{tr}(\Sigma_g^{-1}\Sigma_f) - n + \log\frac{\det\Sigma_g}{\det\Sigma_f} + \operatorname{tr}\left(\Sigma_g^{-1}(A_f - A_g)V_f(A_f - A_g)^\top\right)\right] \end{aligned} \tag{24.4}$$

The first three terms in the brackets measure the discrepancy between the shock covariance matrices Σ_f and Σ_g.

The last term measures the discrepancy between the conditional means $A_f x_{t-1}$ and $A_g x_{t-1}$, averaged over the stationary distribution of x_{t-1}.

Because $\{x_t\}$ is ergodic under f,

$$\frac{1}{t}\log L_t \to h_{KL}(f,g) \quad \text{almost surely as } t \to \infty$$

When model g generates the data, $\frac{1}{t}\log L_t \to -h_{KL}(g,f)$, where $h_{KL}(g,f)$ is (24.4) with the roles of f and g reversed.

Exercise Exercise 24.8.1 asks you to derive (24.4) and to verify these limits by simulation.

The function below computes the two parts of (24.4).

It takes V_f from `model_f.Σ_0`, so `model_f` must be created with `stationary=True`.

```
def kl_rate_terms(model_f, model_g):
    """
    Covariance and conditional-mean terms of the KL divergence rate
    h_KL(f, g), with x_t drawn from the stationary distribution of model f
    """
    n = model_f.A.shape[0]
    D = model_f.A - model_g.A
    V_f = model_f.Σ_0
    cov_term = 0.5 * (np.trace(model_g.CC_inv @ model_f.CC) - n
                      + model_g.log_det_CC - model_f.log_det_CC)
    mean_term = 0.5 * np.trace(model_g.CC_inv @ D @ V_f @ D.T)
    return cov_term, mean_term

def kl_rate_var(model_f, model_g):
    """
    KL divergence rate h_KL(f, g) between two stationary Gaussian VARs
    """
    return sum(kl_rate_terms(model_f, model_g))
```

(continues on next page)

(continued from previous page)

```
examples = [("Example 1", model_f, model_g),
            ("Example 2", model2_f, model2_g)]

for name, m_f, m_g in examples:
    for label, m1, m2 in [("h_KL(f, g)", m_f, m_g), ("h_KL(g, f)", m_g, m_f)]:
        cov_term, mean_term = kl_rate_terms(m1, m2)
        print(f"{name}: {label} = {cov_term + mean_term:.4f}  "
              f"(covariance terms {cov_term:.4f}, mean term {mean_term:.4f})")
```

```
Example 1: h_KL(f, g) = 0.1392  (covariance terms 0.0689, mean term 0.0703)
Example 1: h_KL(g, f) = 0.2079  (covariance terms 0.1012, mean term 0.1067)
Example 2: h_KL(f, g) = 0.3443  (covariance terms 0.3181, mean term 0.0261)
Example 2: h_KL(g, f) = 0.9614  (covariance terms 0.8069, mean term 0.1546)
```

For Example 2, $h_{KL}(g, f) = 0.961$ is almost three times $h_{KL}(f, g) = 0.344$.

That is why $\log L_t$ falls faster when g generates the data than it rises when f generates the data.

The covariance terms account for most of both rates.

Both models assign variance 0.16 to the shock to $(x_{1t} + x_{2t})/\sqrt{2}$.

Model g assigns variance 0.16 to the shock to $(x_{1t} - x_{2t})/\sqrt{2}$; model f assigns it variance 0.04.

Let r be the variance of the shock to $(x_{1t} - x_{2t})/\sqrt{2}$ under the model that generates the data divided by its variance under the other model.

Because the two models agree about the other direction, the covariance terms equal $\frac{1}{2}(r - 1 - \log r)$.

This expression equals 0.807 for $r = 4$ (nature $= g$) and 0.318 for $r = 1/4$ (nature $= f$).

Kullback–Leibler divergence penalizes a model that understates a variance more heavily than a model that overstates it.

In Example 1, the two parts contribute about equally, and $h_{KL}(g, f) = 0.208$ exceeds $h_{KL}(f, g) = 0.139$.

Let's compare the theoretical rates for Example 2 with sample averages of $\frac{1}{T} \log L_T$ across the paths simulated above

```
T2 = L_ratios_ff.shape[1] - 1

print(f"nature = f: mean of log L_T / T = {L_ratios_ff[:, -1].mean() / T2:.4f}, "
      f"h_KL(f, g) = {kl_rate_var(model2_f, model2_g):.4f}")
print(f"nature = g: mean of log L_T / T = {L_ratios_gf[:, -1].mean() / T2:.4f}, "
      f"-h_KL(g, f) = {-kl_rate_var(model2_g, model2_f):.4f}")
```

```
nature = f: mean of log L_T / T = 0.3472, h_KL(f, g) = 0.3443
nature = g: mean of log L_T / T = -0.9924, -h_KL(g, f) = -0.9614
```

The sample averages include the initial-condition term $\log L_0$ divided by $T = 50$, so they need not equal the rates exactly.

The KL rates describe the means of the increments ℓ_t.

Error probabilities for small T depend on the entire distribution of $\log L_T$.

```
for nature, L_ratios in [('f', L_ratios_ff), ('g', L_ratios_gf)]:
    ℓ = np.diff(L_ratios, axis=1).ravel()
    print(f"nature = {nature}: increments ℓ_t have mean {ℓ.mean():.2f}, "
          f"median {np.median(ℓ):.2f}, standard deviation {ℓ.std():.2f}")
```

```
nature = f: increments ℓ_t have mean 0.34, median 0.53, standard deviation 0.54
nature = g: increments ℓ_t have mean -0.98, median -0.07, standard deviation 2.35
```

Under g, the increments ℓ_t have a mean near -1, a median near 0, and a large standard deviation.

Under f, their median exceeds their mean and their standard deviation is small.

That is why β_T exceeds α_T for small T in the error-probability figure above, although $\log L_t$ falls faster under g than it rises under f.

24.7 Application: Samuelson multiplier-accelerator

Now let's connect to the Samuelson multiplier-accelerator model.

The model consists of:

- Consumption: $C_t = \gamma + aY_{t-1}$ where $a \in (0, 1)$ is the marginal propensity to consume
- Investment: $I_t = b(Y_{t-1} - Y_{t-2}) + \sigma\epsilon_t$ where $b > 0$ is the accelerator coefficient and $\{\epsilon_t\}$ is an IID sequence of standard normal random variables
- Government spending: $G_t = G$ (constant)

We have the national income identity

$$Y_t = C_t + I_t + G_t$$

Samuelson Multiplier-Accelerator writes α and β for a and b.

That lecture adds the shock $\sigma\epsilon_t$ directly to the second-order difference equation for Y_t.

We put the shock in the investment equation.

This leads to the same difference equation for Y_t and also determines how C_t and I_t respond to ϵ_t.

Substituting the consumption and investment equations into the national income identity yields the second-order difference equation:

$$Y_t = (\gamma + G) + (a + b)Y_{t-1} - bY_{t-2} + \sigma\epsilon_t$$

With $\rho_1 = a + b$ and $\rho_2 = -b$, we have:

$$Y_t = (\gamma + G) + \rho_1 Y_{t-1} + \rho_2 Y_{t-2} + \sigma\epsilon_t$$

When the roots of $z^2 - \rho_1 z - \rho_2$ lie inside the unit circle, Y_t has a stationary distribution with mean

$$Y^* = \frac{\gamma + G}{1 - \rho_1 - \rho_2} = \frac{\gamma + G}{1 - a}$$

When $\sigma = 0$, Y_t converges to the steady state Y^*, where investment is zero and $Y^* = \gamma + aY^* + G$.

To fit into our discussion, we write it into state-space representation.

To handle the constant term properly, we use an augmented state vector $x_t = [1, Y_t, Y_{t-1}]^\top$:

$$x_{t+1} = \begin{bmatrix} 1 \\ Y_{t+1} \\ Y_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ \gamma + G & \rho_1 & \rho_2 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} 1 \\ Y_t \\ Y_{t-1} \end{bmatrix} + \begin{bmatrix} 0 \\ \sigma \\ 0 \end{bmatrix} \epsilon_{t+1}$$

The observation equation extracts the economic variables:

$$y_t = \begin{bmatrix} Y_t \\ C_t \\ I_t \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ \gamma & 0 & a \\ -(\gamma + G) & 1 & -a \end{bmatrix} \begin{bmatrix} 1 \\ Y_t \\ Y_{t-1} \end{bmatrix}$$

This gives us:

- Y_t (total output)
- $C_t = \gamma \cdot 1 + aY_{t-1}$ (consumption)
- $I_t = Y_t - C_t - G = b(Y_{t-1} - Y_{t-2}) + \sigma\epsilon_t$ (investment)

The third row uses the national income identity to recover I_t from Y_t and Y_{t-1}, so the state vector need not include Y_{t-2}.

The matrix $CC^\top$ has rank one: given x_t, the first and third components of x_{t+1} are known and only Y_{t+1} is random.

Consequently (24.1) is not a density with respect to Lebesgue measure on $\mathbb{R}^3$.

`create_var_model` replaces $(CC^\top)^{-1}$ with a pseudo-inverse and adds $10^{-10}I$ to $CC^\top$ before taking the log determinant.

The regularization terms cancel in ℓ_{t+1}, which equals the log ratio of the conditional densities of Y_{t+1} given x_t under the two models.

The same device handles the singular covariance matrix Σ_0 of x_0, whose first component is always 1.

An econometrician who observes Y_t and Y_{t-1} observes the state x_t.

When an econometrician observes only some components of a state vector, or observes them with error, a Kalman filter produces innovations from which to compute the likelihood, as in *A First Look at the Kalman Filter* and *The Kalman Filter and Vector Autoregressions*.

```
def samuelson_to_var(a, b, γ, G, σ):
    """
    Convert Samuelson model parameters to VAR form with augmented state

    Samuelson model:
    - Y_t = C_t + I_t + G
    - C_t = γ + a*Y_{t-1}
    - I_t = b*(Y_{t-1} - Y_{t-2}) + σ*ε_t

    Reduced form: Y_t = (γ+G) + (a+b)*Y_{t-1} - b*Y_{t-2} + σ*ε_t

    State vector is [1, Y_t, Y_{t-1}]'
    """
    ρ_1 = a + b
    ρ_2 = -b

    # State transition matrix for augmented state
    A = np.array([[1,       0,      0],
                  [γ + G,   ρ_1,    ρ_2],
                  [0,       1,      0]])

    # Shock loading matrix
    C = np.array([[0],
                  [σ],
                  [0]])

    # Observation matrix (extracts Y_t, C_t, I_t from [1, Y_t, Y_{t-1}]')
    G_obs = np.array([[0,          1,   0],     # Y_t
                      [γ,          0,   a],     # C_t = γ + a Y_{t-1}
                      [-(γ + G),   1,  -a]])    # I_t = Y_t - C_t - G

    return A, C, G_obs
```

We define functions in the code cell below to get the initial conditions and check stability

With `stationary_init=True`, the initial state $[1, Y_0, Y_{-1}]^\top$ has the stationary distribution: (Y_0, Y_{-1}) has mean (Y^*, Y^*) and a covariance matrix that solves a discrete Lyapunov equation.

Let's implement it and inspect the likelihood ratio processes induced by two Samuelson models with different parameters.

```
def create_samuelson_var_model(a, b, γ, G, σ, stationary_init=False,
                               y_0=None, y_m1=None):
    """
    Create a VAR model from Samuelson parameters
    """
    A, C, G_obs = samuelson_to_var(a, b, γ, G, σ)

    μ_0, Σ_0 = get_samuelson_initial_conditions(
        a, b, γ, G, σ, y_0, y_m1, stationary_init
    )

    # Create VAR model
    model = create_var_model(A, C, μ_0, Σ_0, stationary=False)
    is_stable, roots, max_root, dynamics = check_samuelson_stability(a, b)
    info = {
        'a': a, 'b': b, 'γ': γ, 'G': G, 'σ': σ,
        'ρ_1': a + b, 'ρ_2': -b,
        'steady_state': (γ + G) / (1 - a),
        'is_stable': is_stable,
        'roots': roots,
        'max_abs_root': max_root,
        'dynamics': dynamics
    }

    return model, G_obs, info

def simulate_samuelson(model, G_obs, T, rng, N_paths=1):
    """
    Simulate Samuelson model
    """
    # Simulate state paths
    states = simulate_var(model, T, rng, N_paths)

    # Extract observables using G matrix
    if N_paths == 1:
        # Single path: states is (T+1, 3)
        observables = (G_obs @ states.T).T
    else:
        # Multiple paths: states is (N_paths, T+1, 3)
        observables = np.zeros((N_paths, T+1, 3))
        for i in range(N_paths):
            observables[i] = (G_obs @ states[i].T).T

    return states, observables
```

Now let's simulate two Samuelson models with different accelerator coefficients and plot their sample paths

Both models have the same steady state $Y^* = 20/(1 - 0.98) = 1000$.

Both start from $Y_0 \sim \mathcal{N}(100, 25)$ and $Y_{-1} \sim \mathcal{N}(95, 25)$, far below Y^*.

```
# Model f: Higher accelerator coefficient
a_f, b_f = 0.98, 0.9
γ_f, G_f, σ_f = 10, 10, 0.5
```

(continues on next page)

(continued from previous page)

```
# Model g: Lower accelerator coefficient
a_g, b_g = 0.98, 0.85
y_g, G_g, σ_g = 10, 10, 0.5


model_sam_f, G_obs_f, info_f = create_samuelson_var_model(
    a_f, b_f, y_f, G_f, σ_f,
    stationary_init=False,
    y_0=100, y_m1=95
)

model_sam_g, G_obs_g, info_g = create_samuelson_var_model(
    a_g, b_g, y_g, G_g, σ_g,
    stationary_init=False,
    y_0=100, y_m1=95
)

T = 50
N_paths = 50

# Get both states and observables
states_f, obs_f = simulate_samuelson(model_sam_f, G_obs_f, T, rng, N_paths)
states_g, obs_g = simulate_samuelson(model_sam_g, G_obs_g, T, rng, N_paths)

output_paths_f = obs_f[:, :, 0]
output_paths_g = obs_g[:, :, 0]

print("model f:")
print(f"  ρ_1 = a + b = {info_f['ρ_1']:.2f}")
print(f"  ρ_2 = -b = {info_f['ρ_2']:.2f}")
print(f"  roots: {info_f['roots']}")
print(f"  dynamics: {info_f['dynamics']}")
print(f"  steady state: {info_f['steady_state']:.1f}")

print("\nmodel g:")
print(f"  ρ_1 = a + b = {info_g['ρ_1']:.2f}")
print(f"  ρ_2 = -b = {info_g['ρ_2']:.2f}")
print(f"  roots: {info_g['roots']}")
print(f"  dynamics: {info_g['dynamics']}")
print(f"  steady state: {info_g['steady_state']:.1f}")


fig, ax = plt.subplots(1, 1)

for i in range(min(20, N_paths)):
    ax.plot(output_paths_f[i], alpha=0.6, color='C0', linewidth=0.8)
    ax.plot(output_paths_g[i], alpha=0.6, color='C1', linewidth=0.8)
ax.axhline(info_f['steady_state'], color='gray', linestyle='--', alpha=0.5)
ax.set_xlabel('$t$')
ax.set_ylabel('$Y_t$')
ax.legend(['model f', 'model g'], loc='lower right')
plt.tight_layout()
plt.show()
```

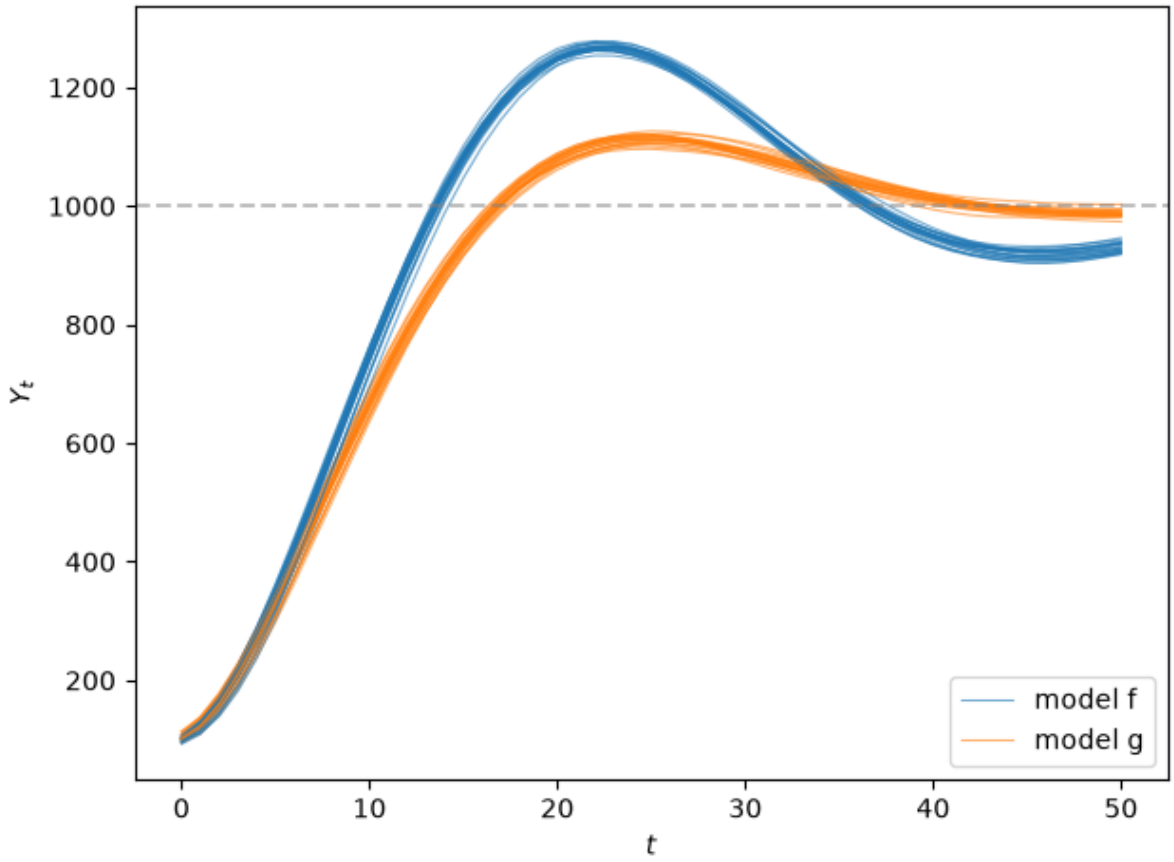
```
model f:
```

(continues on next page)

(continued from previous page)

```
  ρ_1 = a + b = 1.88
  ρ_2 = -b = -0.90
  roots: [0.94+0.12806248j 0.94-0.12806248j]
  dynamics: Damped oscillations
  steady state: 1000.0

model g:
  ρ_1 = a + b = 1.83
  ρ_2 = -b = -0.85
  roots: [0.915+0.11302655j 0.915-0.11302655j]
  dynamics: Damped oscillations
  steady state: 1000.0
```




```
# Compute likelihood ratios
L_ratios_ff = compute_likelihood_ratio_var(states_f, model_sam_f, model_sam_g)
L_ratios_gf = compute_likelihood_ratio_var(states_g, model_sam_f, model_sam_g)

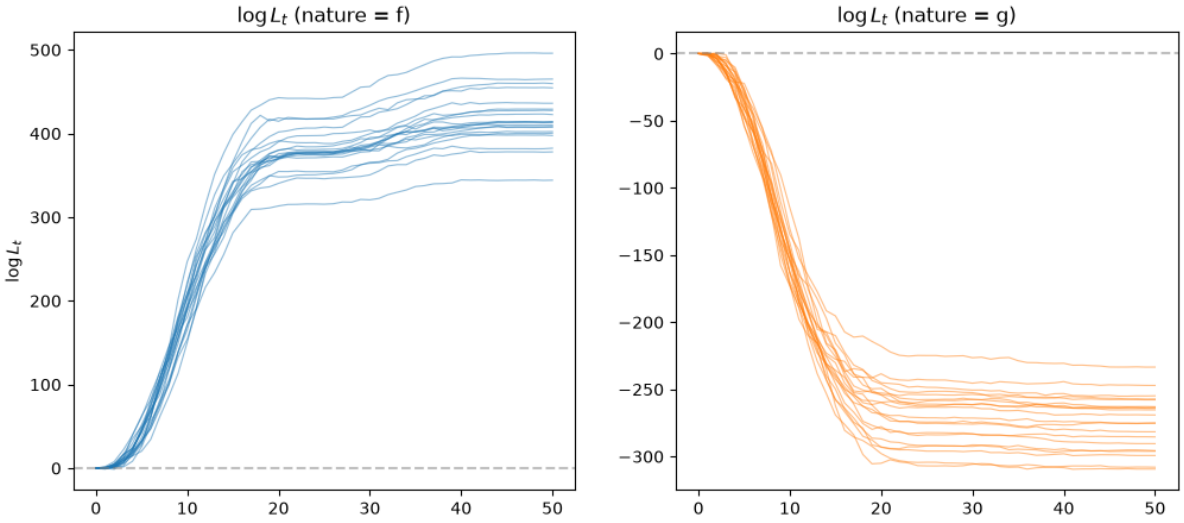
fig, axes = plt.subplots(1, 2, figsize=(12, 5))

ax = axes[0]
for i in range(min(20, N_paths)):
    ax.plot(L_ratios_ff[i], alpha=0.5, color='C0', lw=0.8)
ax.axhline(y=0, color='gray', linestyle='--', alpha=0.5)
ax.set_title(r'$\log L_t$ (nature = f)')
ax.set_ylabel(r'$\log L_t$')
```

(continues on next page)

(continued from previous page)

```
ax = axes[1]
for i in range(min(20, N_paths)):
    ax.plot(L_ratios_gf[i], alpha=0.5, color='C1', lw=0.8)
ax.axhline(y=0, color='gray', linestyle='--', alpha=0.5)
ax.set_title(r'$\log L_t$ (nature = g)')
plt.show()
```



In the figure on the left, data are generated by f and the log likelihood ratio diverges to plus infinity.

In the figure on the right, data are generated by g and the log likelihood ratio diverges to negative infinity.

In both cases, the likelihood ratio processes eventually lead us to select the correct model.

Almost all of the information that discriminates between the two models comes from the transient that carries Y_t from $Y_0 \approx 100$ toward $Y^* = 1000$.

The two models share a, γ, G, σ, and Y^*; they differ only in the accelerator coefficient b.

Their forecasts of Y_{t+1} differ by $(b_f - b_g)(Y_t - Y_{t-1}) = 0.05\,\Delta Y_t$.

Because the two models share σ, the expected increment of $\log L_t$ under f is

$$E_f[\ell_{t+1} | x_t] = \frac{(b_f - b_g)^2 (\Delta Y_t)^2}{2\sigma^2}$$

and $E_g[-\ell_{t+1} | x_t]$ is given by the same formula.

Along the transient, $|\Delta Y_t|$ is large relative to $\sigma = 0.5$.

In the stationary distribution, ΔY_t is small, and the KL divergence rate is $h_{KL}(f, g) = (b_f - b_g)^2 E_f[(\Delta Y_t)^2]/(2\sigma^2)$.

The function below computes this rate from the stationary distribution of x_t for a model created with `stationary_init=True`.

```
def samuelson_kl_rate(model_f, model_g):
    """
    KL divergence rate h_KL(f, g) for two Samuelson models created with
    stationary_init=True; only the Y component of x_{t+1} is random
    """
    σ2_f, σ2_g = model_f.CC[1, 1], model_g.CC[1, 1]
```

(continues on next page)

(continued from previous page)

```
    d = model_f.A[1] - model_g.A[1]      # difference in forecasts of Y_{t+1}
    M_f = model_f.Σ_0 + np.outer(model_f.μ_0, model_f.μ_0)   # E_f[x_t x_t']
    return 0.5 * (σ2_f / σ2_g - 1 + np.log(σ2_g / σ2_f)
                  + d @ M_f @ d / σ2_g)

model_stat_f, _, _ = create_samuelson_var_model(
    a_f, b_f, γ_f, G_f, σ_f, stationary_init=True)
model_stat_g, _, _ = create_samuelson_var_model(
    a_g, b_g, γ_g, G_g, σ_g, stationary_init=True)

h_fg = samuelson_kl_rate(model_stat_f, model_stat_g)
h_gf = samuelson_kl_rate(model_stat_g, model_stat_f)

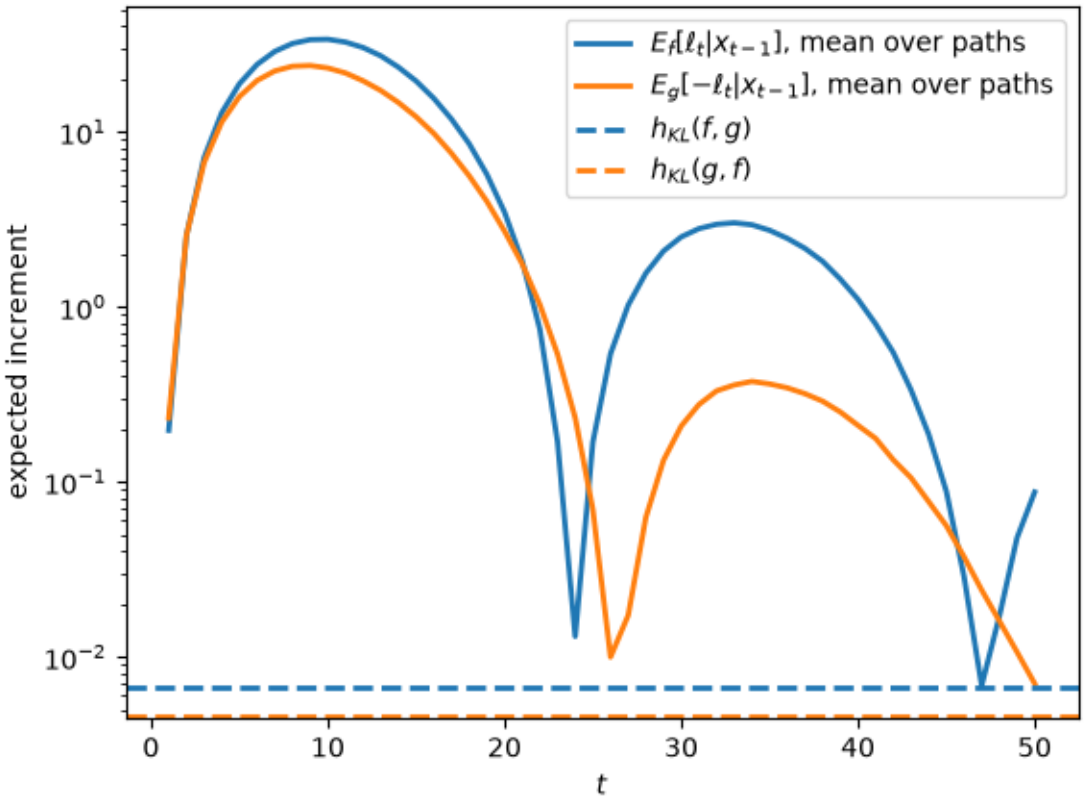
print(f"median log L_{T}: nature = f: {np.median(L_ratios_ff[:, -1]):.1f}, "
      f"nature = g: {np.median(L_ratios_gf[:, -1]):.1f}")
print(f"stationary sd of ΔY_t under f: "
      f"{np.sqrt(2 * (model_stat_f.Σ_0[1, 1] - model_stat_f.Σ_0[1, 2])):.2f}")
print(f"h_KL(f, g) = {h_fg:.4f}, h_KL(g, f) = {h_gf:.4f}")
print(f"{T} × h_KL(f, g) = {T * h_fg:.2f}, {T} × h_KL(g, f) = {T * h_gf:.2f}")

# Expected increments E[ℓ_t | x_{t-1}] along the simulated paths
d = model_sam_f.A[1] - model_sam_g.A[1]
σ2 = σ_f**2
inc_f = ((states_f[:, :-1, :] @ d)**2 / (2 * σ2)).mean(axis=0)
inc_g = ((states_g[:, :-1, :] @ d)**2 / (2 * σ2)).mean(axis=0)
t_grid = np.arange(1, T + 1)

print(f"largest mean expected increment: "
      f"nature = f: {inc_f.max():.1f} at t = {t_grid[inc_f.argmax()]}, "
      f"nature = g: {inc_g.max():.1f} at t = {t_grid[inc_g.argmax()]}")

fig, ax = plt.subplots()
ax.semilogy(t_grid, inc_f, color='C0', lw=2,
            label=r'$E_f[\ell_t | x_{t-1}]$, mean over paths')
ax.semilogy(t_grid, inc_g, color='C1', lw=2,
            label=r'$E_g[-\ell_t | x_{t-1}]$, mean over paths')
ax.axhline(h_fg, color='C0', ls='--', lw=2, label=r'$h_{KL}(f, g)$')
ax.axhline(h_gf, color='C1', ls='--', lw=2, label=r'$h_{KL}(g, f)$')
ax.set_xlabel('$t$')
ax.set_ylabel('expected increment')
ax.legend()
plt.show()
```

```
median log L_50: nature = f: 413.5, nature = g: -279.2
stationary sd of ΔY_t under f: 1.15
h_KL(f, g) = 0.0066, h_KL(g, f) = 0.0045
50 × h_KL(f, g) = 0.33, 50 × h_KL(g, f) = 0.23
largest mean expected increment: nature = f: 33.8 at t = 10, nature = g: 23.9 at t␣
↪= 9
```



The expected increments reach about 34 under f and 24 under g near $t = 10$, when Y_t rises fastest.

They approach the stationary rates $h_{KL}(f, g)$ and $h_{KL}(g, f)$ only near the turning points of the oscillation, where ΔY_t is close to zero.

The stationary standard deviation of ΔY_t under f is 1.15.

From the stationary distribution, 50 periods raise $E_f[\log L_t]$ by only about 0.33, while the transient raises the median of $\log L_{50}$ to about 400.

Exercise Exercise 24.8.2 starts both models from their stationary distributions and studies how much harder the two models are to tell apart.

24.8 Exercises

Exercise 24.8.1

Derive formula (24.4) for the KL divergence rate $h_{KL}(f, g)$.

Use the formula for the Kullback–Leibler divergence of two multivariate normal distributions:

$$KL(\mathcal{N}(\mu_1, \Sigma_1), \mathcal{N}(\mu_2, \Sigma_2)) = \frac{1}{2}\left[\text{tr}(\Sigma_2^{-1}\Sigma_1) - n + (\mu_1 - \mu_2)^\top \Sigma_2^{-1}(\mu_1 - \mu_2) + \log\frac{\det\Sigma_2}{\det\Sigma_1}\right]$$

Then verify by simulation that $\frac{1}{t}\log L_t \to h_{KL}(f, g)$ when f generates the data and $\frac{1}{t}\log L_t \to -h_{KL}(g, f)$ when g generates the data.

Do this for Examples 1 and 2 with $T = 2000$.

Solution

Under model f, $x_t|x_{t-1} \sim \mathcal{N}(A_f x_{t-1}, \Sigma_f)$; under model g, $x_t|x_{t-1} \sim \mathcal{N}(A_g x_{t-1}, \Sigma_g)$.

So $E_f[\ell_t|x_{t-1}] = KL(\mathcal{N}(A_f x_{t-1}, \Sigma_f), \mathcal{N}(A_g x_{t-1}, \Sigma_g))$.

Setting $\mu_1 - \mu_2 = (A_f - A_g)x_{t-1}$ in the formula above gives

$$E_f[\ell_t|x_{t-1}] = \frac{1}{2}\left[\text{tr}(\Sigma_g^{-1}\Sigma_f) - n + \log\frac{\det\Sigma_g}{\det\Sigma_f} + x_{t-1}^\top (A_f - A_g)^\top \Sigma_g^{-1}(A_f - A_g)x_{t-1}\right]$$

Under f, $x_{t-1} \sim \mathcal{N}(0, V_f)$.

For any matrix M, $E[x_{t-1}^\top M x_{t-1}] = \text{tr}(MV_f)$.

With $M = (A_f - A_g)^\top \Sigma_g^{-1}(A_f - A_g)$ and the cyclic property of the trace,

$$E_f[x_{t-1}^\top M x_{t-1}] = \text{tr}\left(\Sigma_g^{-1}(A_f - A_g)V_f(A_f - A_g)^\top\right)$$

Taking expectations of $E_f[\ell_t|x_{t-1}]$ over x_{t-1} gives (24.4).

Because x_{t-1} has the same distribution $\mathcal{N}(0, V_f)$ for every t, every increment ℓ_t has mean $h_{KL}(f, g)$ and

$$E_f[\log L_t] = E_f\left[\log\frac{p_f(x_0)}{p_g(x_0)}\right] + t\, h_{KL}(f, g)$$

The ergodic theorem implies $\frac{1}{t}\sum_{s=1}^t \ell_s \to h_{KL}(f, g)$ almost surely, and $\frac{1}{t}\log\frac{p_f(x_0)}{p_g(x_0)} \to 0$.

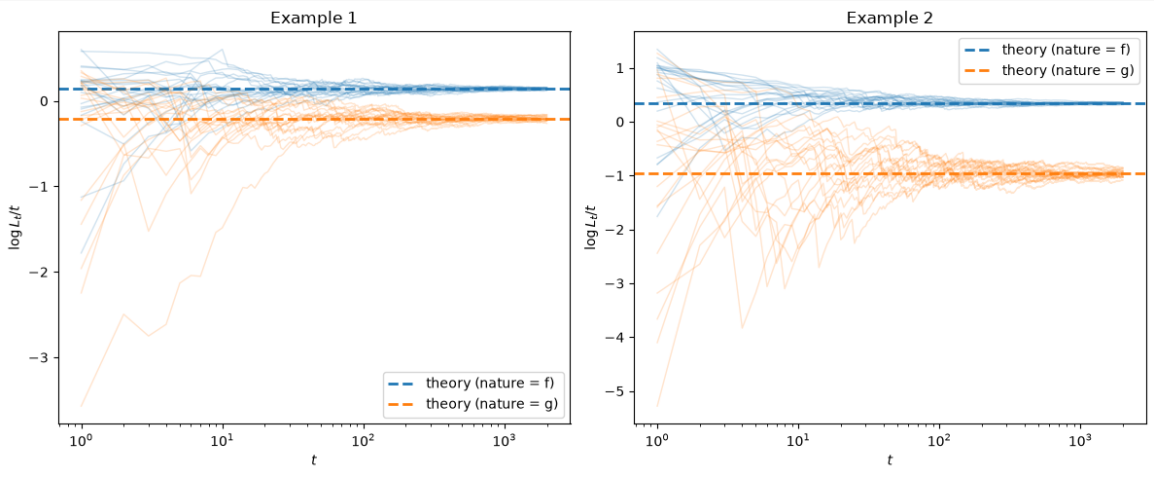
Here is the simulation.

```
T_long = 2000
N_long = 20
t_grid = np.arange(1, T_long + 1)

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

for ax, (name, m_f, m_g) in zip(axes, examples):
    for nature, m_true, rate, color in [
            ('f', m_f, kl_rate_var(m_f, m_g), 'C0'),
            ('g', m_g, -kl_rate_var(m_g, m_f), 'C1')]:
        paths = simulate_var(m_true, T_long, rng, N_long)
        log_L = compute_likelihood_ratio_var(paths, m_f, m_g)
        avg = log_L[:, 1:] / t_grid
        ax.plot(t_grid, avg.T, color=color, alpha=0.2, lw=1)
        ax.axhline(rate, color=color, ls='--', lw=2,
                   label=f'theory (nature = {nature})')
        print(f"{name}, nature = {nature}: mean of log L_T / T = "
              f"{avg[:, -1].mean():.4f}, theory = {rate:.4f}")
    ax.set_xscale('log')
    ax.set_xlabel('$t$')
    ax.set_ylabel(r'$\log L_t / t$')
    ax.set_title(name)
    ax.legend()
```

```
Example 1, nature = f: mean of log L_T / T = 0.1420, theory = 0.1392
Example 1, nature = g: mean of log L_T / T = -0.2068, theory = -0.2079
Example 2, nature = f: mean of log L_T / T = 0.3423, theory = 0.3443
Example 2, nature = g: mean of log L_T / T = -0.9629, theory = -0.9614
```



In all four cases, $\frac{1}{t}\log L_t$ settles down near the theoretical rate.

Exercise 24.8.2

In the Samuelson example, both models start from $Y_0 \approx 100$, far below the common steady state $Y^* = 1000$.

Start both models from their stationary distributions instead by using the models `model_stat_f` and `model_stat_g` created with `stationary_init=True`.

1. Simulate 500 paths of length $T = 50$ from each model. Report the median of $\log L_{50}$ and the error probabilities α_{50}, β_{50}, and $\frac{1}{2}(\alpha_{50} + \beta_{50})$. Compare them with those for the paths that start from $Y_0 \approx 100$.
2. Use `samuelson_kl_rate` and the Kullback–Leibler divergence of the two stationary distributions of (Y_0, Y_{-1}) to compute the number of periods t at which $E_f[\log L_t] = 5$.
3. Simulate paths of that length and plot α_t and β_t.

Solution

Here are the error probabilities at $T = 50$ for the two initial conditions.

```
T_s, N_s = 50, 500

paths_stat_f = simulate_var(model_stat_f, T_s, rng, N_s)
paths_stat_g = simulate_var(model_stat_g, T_s, rng, N_s)
log_L_stat_f = compute_likelihood_ratio_var(paths_stat_f,
                                            model_stat_f, model_stat_g)
log_L_stat_g = compute_likelihood_ratio_var(paths_stat_g,
                                            model_stat_f, model_stat_g)

for label, log_L_f, log_L_g in [
        ("stationary start", log_L_stat_f, log_L_stat_g),
        ("Y_0 ≈ 100", L_ratios_ff, L_ratios_gf)]:
```

```
    α = np.mean(log_L_f[:, -1] < 0)
    β = np.mean(log_L_g[:, -1] >= 0)
    print(f"{label}:")
    print(f"  median log L_50: nature = f: {np.median(log_L_f[:, -1]):.3f}, "
          f"nature = g: {np.median(log_L_g[:, -1]):.3f}")
    print(f"  α_50 = {α:.3f}, β_50 = {β:.3f}, "
          f"average error = {0.5 * (α + β):.3f}")
```

```
stationary start:
  median log L_50: nature = f: 0.132, nature = g: -0.452
  α_50 = 0.426, β_50 = 0.244, average error = 0.335
Y_0 ≈ 100:
  median log L_50: nature = f: 413.526, nature = g: -279.175
  α_50 = 0.000, β_50 = 0.000, average error = 0.000
```

From the stationary distributions, the medians of $\log L_{50}$ are within 0.5 of zero, α_{50} is about 0.4, and the likelihood ratio test selects the wrong model in about a third of the samples.

From $Y_0 \approx 100$, it selects the correct model in all 50 samples from each model.

Under f, $E_f[\log L_t] = KL_0 + t\, h_{KL}(f,g)$, where KL_0 is the Kullback–Leibler divergence of the stationary distribution of (Y_0, Y_{-1}) under g from that under f.

```
def kl_gauss(μ_1, Σ_1, μ_2, Σ_2):
    """
    KL divergence of N(μ_2, Σ_2) from N(μ_1, Σ_1)
    """
    n = len(μ_1)
    Σ_2_inv = np.linalg.inv(Σ_2)
    d = μ_1 - μ_2
    return 0.5 * (np.trace(Σ_2_inv @ Σ_1) - n + d @ Σ_2_inv @ d
                  + np.log(np.linalg.det(Σ_2) / np.linalg.det(Σ_1)))

KL_0 = kl_gauss(model_stat_f.μ_0[1:], model_stat_f.Σ_0[1:, 1:],
                model_stat_g.μ_0[1:], model_stat_g.Σ_0[1:, 1:])
t_star = int(np.ceil((5 - KL_0) / h_fg))

print(f"KL_0 = {KL_0:.4f}, h_KL(f, g) = {h_fg:.5f}")
print(f"E_f[log L_50]: theory = {KL_0 + 50 * h_fg:.3f}, "
      f"simulation = {log_L_stat_f[:, -1].mean():.3f}")
print(f"stationary start: E_f[log L_t] = 5 at t = {t_star}")

# Transient from Y_0 ≈ 100: cumulative mean expected increments
t_transient = np.argmax(np.cumsum(inc_f) >= 5) + 1
print(f"Y_0 ≈ 100: cumulative expected increments reach 5 at t = {t_transient}")
```

```
KL_0 = 0.0881, h_KL(f, g) = 0.00661
E_f[log L_50]: theory = 0.419, simulation = 0.391
stationary start: E_f[log L_t] = 5 at t = 743
Y_0 ≈ 100: cumulative expected increments reach 5 at t = 3
```

About 750 periods of stationary data contain as much information about b, on average, as the first three periods of the transient from $Y_0 \approx 100$.

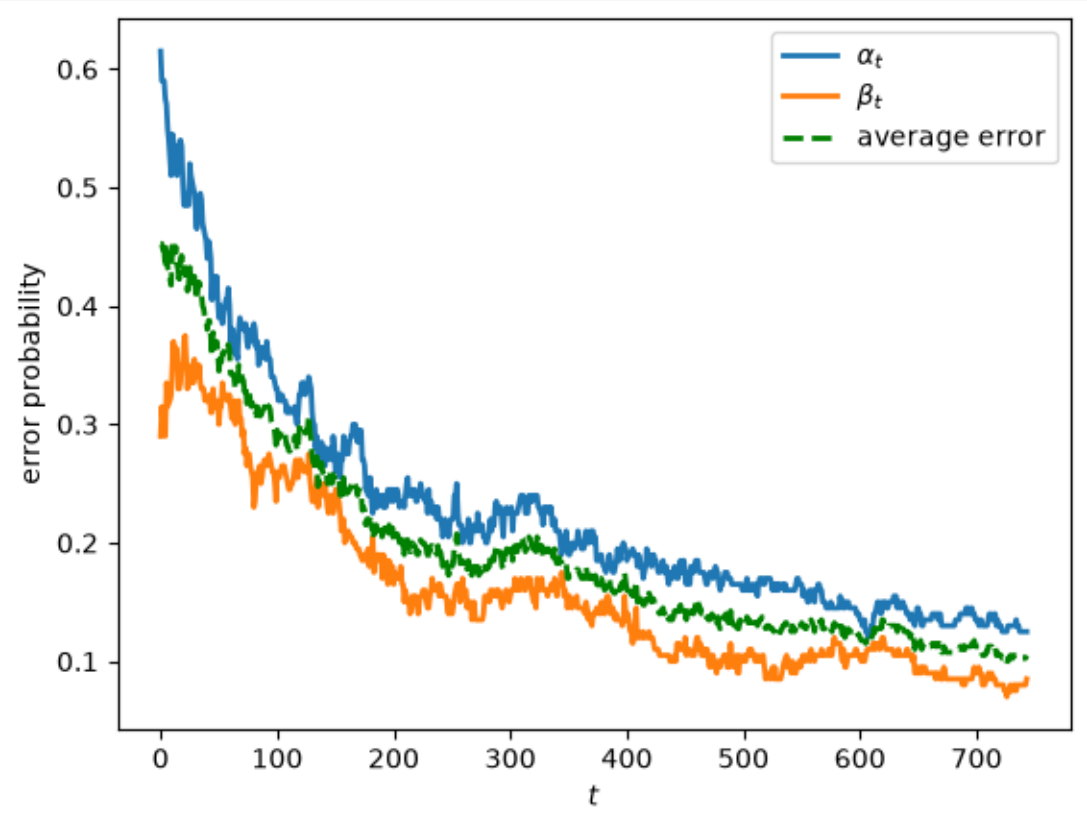
```
N_star = 200

paths_f = simulate_var(model_stat_f, t_star, rng, N_star)
paths_g = simulate_var(model_stat_g, t_star, rng, N_star)
log_L_f = compute_likelihood_ratio_var(paths_f, model_stat_f, model_stat_g)
log_L_g = compute_likelihood_ratio_var(paths_g, model_stat_f, model_stat_g)
```

```
α_t = np.mean(log_L_f < 0, axis=0)
β_t = np.mean(log_L_g >= 0, axis=0)

fig, ax = plt.subplots()
ax.plot(α_t, color='C0', lw=2, label=r'$\alpha_t$')
ax.plot(β_t, color='C1', lw=2, label=r'$\beta_t$')
ax.plot(0.5 * (α_t + β_t), 'g--', lw=2, label='average error')
ax.set_xlabel('$t$')
ax.set_ylabel('error probability')
ax.legend()
plt.show()

print(f"at t = {t_star}: α = {α_t[-1]:.3f}, β = {β_t[-1]:.3f}")
```




```
at t = 743: α = 0.125, β = 0.085
```

Even after about 750 periods, both error probabilities remain roughly 0.1.

Exercise 24.8.3

Suppose that neither f nor g of Example 2 generates the data.

Instead a third VAR h with matrices (A_h, C_h) generates the data.

Show that

$$\frac{1}{t}\log L_t \to h_{KL}(h,g) - h_{KL}(h,f)$$

almost surely, where $h_{KL}(h, \cdot)$ is (24.4) with h in place of f.

This is the VAR counterpart of the formula in *A helpful formula* of *Likelihood Ratio Processes*.

Verify the limit by simulation for $(A_h, C_h) = (A_f, C_g)$ and for $(A_h, C_h) = (A_g, C_f)$.

Which model does the likelihood ratio select in each case?

Solution

Write

$$\ell_t = \log \frac{p_h(x_t|x_{t-1})}{p_g(x_t|x_{t-1})} - \log \frac{p_h(x_t|x_{t-1})}{p_f(x_t|x_{t-1})}$$

Taking expectations under h, with x_{t-1} drawn from the stationary distribution of h, gives $E_h[\ell_t] = h_{KL}(h, g) - h_{KL}(h, f)$.

The ergodic theorem under h then gives the almost sure limit.

```
model_h1 = create_var_model(model2_f.A, model2_g.C)
model_h2 = create_var_model(model2_g.A, model2_f.C)

T_h, N_h = 1000, 20
t_grid = np.arange(1, T_h + 1)

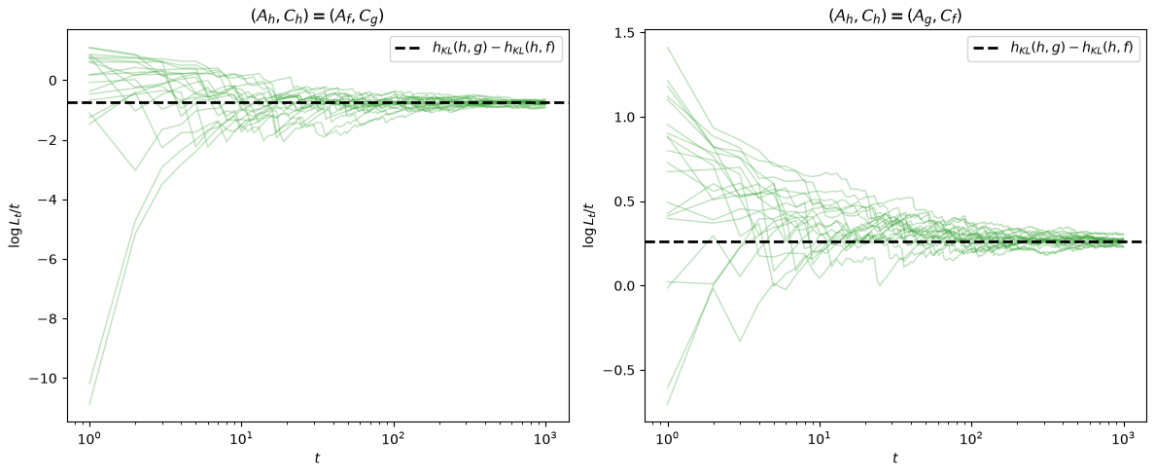
fig, axes = plt.subplots(1, 2, figsize=(12, 5))

cases = [("(A_f, C_g)", r'$(A_h, C_h) = (A_f, C_g)$', model_h1),
         ("(A_g, C_f)", r'$(A_h, C_h) = (A_g, C_f)$', model_h2)]

for ax, (name, title, m_h) in zip(axes, cases):
    K_f = kl_rate_var(m_h, model2_f)
    K_g = kl_rate_var(m_h, model2_g)
    paths = simulate_var(m_h, T_h, rng, N_h)
    log_L = compute_likelihood_ratio_var(paths, model2_f, model2_g)
    avg = log_L[:, 1:] / t_grid
    ax.plot(t_grid, avg.T, color='C2', alpha=0.3, lw=1)
    ax.axhline(K_g - K_f, color='k', ls='--', lw=2,
               label=r'$h_{KL}(h, g) - h_{KL}(h, f)$')
    ax.set_xscale('log')
    ax.set_xlabel('$t$')
    ax.set_ylabel(r'$\log L_t / t$')
    ax.set_title(title)
    ax.legend()
    print(f"(A_h, C_h) = {name}: h_KL(h, f) = {K_f:.4f}, "
          f"h_KL(h, g) = {K_g:.4f}, difference = {K_g - K_f:.4f}, "
          f"mean of log L_T / T = {avg[:, -1].mean():.4f}")

plt.tight_layout()
plt.show()
```

```
(A_h, C_h) = (A_f, C_g): h_KL(h, f) = 0.8069, h_KL(h, g) = 0.0586, difference =␣
↪-0.7482, mean of log L_T / T = -0.7711
(A_h, C_h) = (A_g, C_f): h_KL(h, f) = 0.0542, h_KL(h, g) = 0.3181, difference =␣
↪0.2639, mean of log L_T / T = 0.2643
```



When $(A_h, C_h) = (A_f, C_g)$, $\log L_t \to -\infty$ and the likelihood ratio selects g, although f has the correct transition matrix.

When $(A_h, C_h) = (A_g, C_f)$, $\log L_t \to +\infty$ and the likelihood ratio selects f, although g has the correct transition matrix.

In both cases the likelihood ratio selects the model with the correct shock covariance matrix.

When $(A_h, C_h) = (A_f, C_g)$, $h_{KL}(h, f)$ consists only of covariance terms and $h_{KL}(h, g)$ only of a conditional-mean term.

When $(A_h, C_h) = (A_g, C_f)$, the roles of f and g are reversed.

In both cases the covariance terms are larger, as in the decomposition for Example 2 above.

CHAPTER

TWENTYFIVE

MEAN OF A LIKELIHOOD RATIO PROCESS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Mean of a Likelihood Ratio Process*
  - *Overview*
  - *Mathematical expectation of likelihood ratio*
  - *Importance sampling*
  - *Approximating a cumulative likelihood ratio*
  - *Distribution of sample mean*
  - *Choosing a sampling distribution*
  - *Exercises*

25.1 Overview

In *Likelihood Ratio Processes* we described a peculiar property of a likelihood ratio process, namely, that its mean equals one for all $t \geq 0$ despite its converging to zero almost surely.

While it is easy to verify that peculiar property analytically (i.e., in population), it is challenging to use a computer simulation to verify it via an application of a law of large numbers that entails studying sample averages of repeated simulations.

To confront this challenge, this lecture puts **importance sampling** to work to accelerate convergence of sample averages to population means.

We use importance sampling to estimate the mean of a cumulative likelihood ratio $L\left(\omega^t\right) = \prod_{i=1}^t \ell\left(\omega_i\right)$.

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install jax
```

We start by importing some Python packages.

```
import jax
import jax.numpy as jnp
import numpy as np
import matplotlib.pyplot as plt
from jax.scipy.stats import beta
from typing import NamedTuple
from functools import partial

# Set JAX to use 64-bit floats
jax.config.update("jax_enable_x64", True)
```

25.2 Mathematical expectation of likelihood ratio

In *Likelihood Ratio Processes*, we studied a likelihood ratio $\ell\left(\omega_t\right)$

$$\ell\left(\omega_t\right) = \frac{f\left(\omega_t\right)}{g\left(\omega_t\right)}$$

where f is the Beta$(F_a, F_b) =$ Beta$(1, 1)$ density, the uniform density on $[0, 1]$, and g is the Beta$(G_a, G_b) =$ Beta$(3, 1.2)$ density.

Assume that an i.i.d. random variable $\omega_t \in \Omega$ is generated by g.

The **cumulative likelihood ratio** $L\left(\omega^t\right)$ is

$$L\left(\omega^t\right) = \prod_{i=1}^t \ell\left(\omega_i\right)$$

Our goal is to approximate the mathematical expectation $E\left[L\left(\omega^t\right)\right]$ well.

In *this section* of *Likelihood Ratio Processes*, we showed that $E\left[L\left(\omega^t\right)\right]$ equals 1 for all t.

We want to check out how well this holds if we replace E with sample averages from simulations.

This turns out to be easier said than done because for Beta distributions assumed above, $L\left(\omega^t\right)$ has a very skewed distribution with a very long tail as $t \to \infty$.

This property makes it difficult efficiently and accurately to estimate the mean by standard Monte Carlo simulation methods.

In this lecture we explore how a standard Monte Carlo method fails.

We also show how **importance sampling** provides a more computationally efficient way to approximate the mean of the cumulative likelihood ratio.

We first take a look at the density functions `f` and `g` .

```
# Parameters for the model
class ImpSampleParams(NamedTuple):
    F_a: float = 1.0    # Beta parameters for f
```

(continues on next page)

(continued from previous page)

```
    F_b: float = 1.0
    G_a: float = 3.0    # Beta parameters for g
    G_b: float = 1.2

params = ImpSampleParams()

def f(w):
    return beta.pdf(w, params.F_a, params.F_b)

def g(w):
    return beta.pdf(w, params.G_a, params.G_b)
```

```
w_range = np.linspace(1e-2, 1-1e-5, 1000)

plt.plot(w_range, g(w_range), lw=2, label='g')
plt.plot(w_range, f(w_range), lw=2, label='f')
plt.xlabel(r'$\omega$')
plt.legend()
plt.show()
```

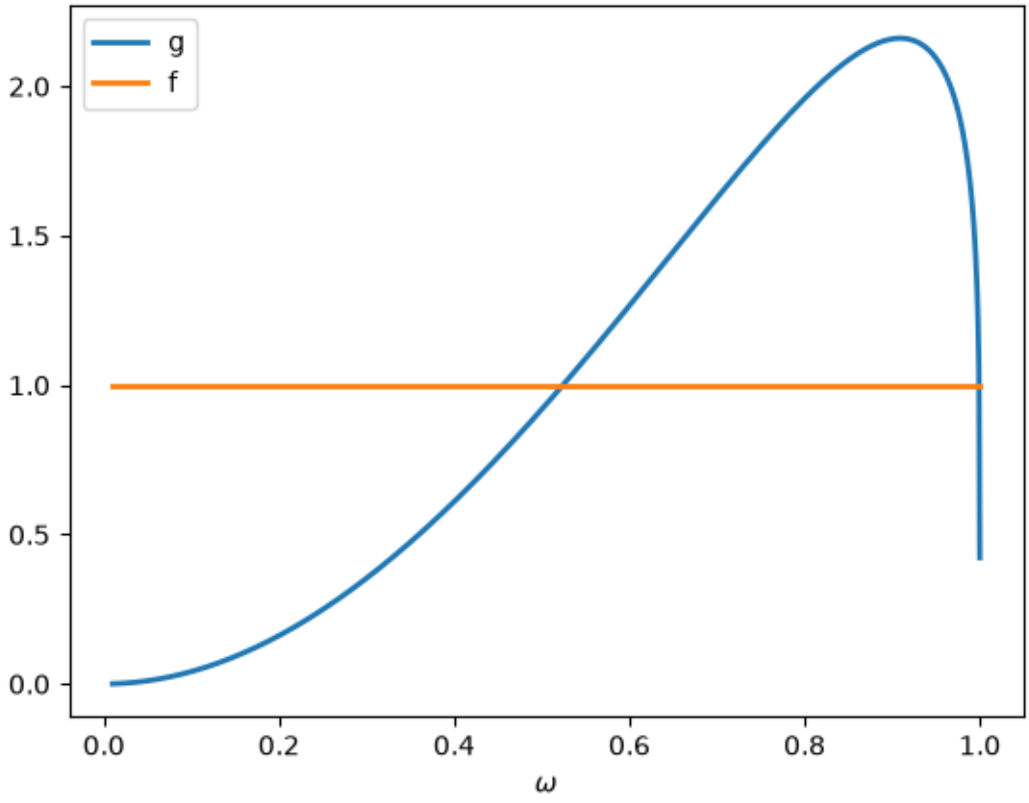


Fig. 25.1: Beta density functions f and g

The likelihood ratio is `l(w)=f(w)/g(w)`.

```
def l(w):
    return f(w) / g(w)
```

```
plt.plot(w_range, l(w_range), lw=2)
plt.yscale('log')
plt.xlabel(r'$\omega$')
plt.ylabel(r'$\ell(\omega)$')
plt.show()
```

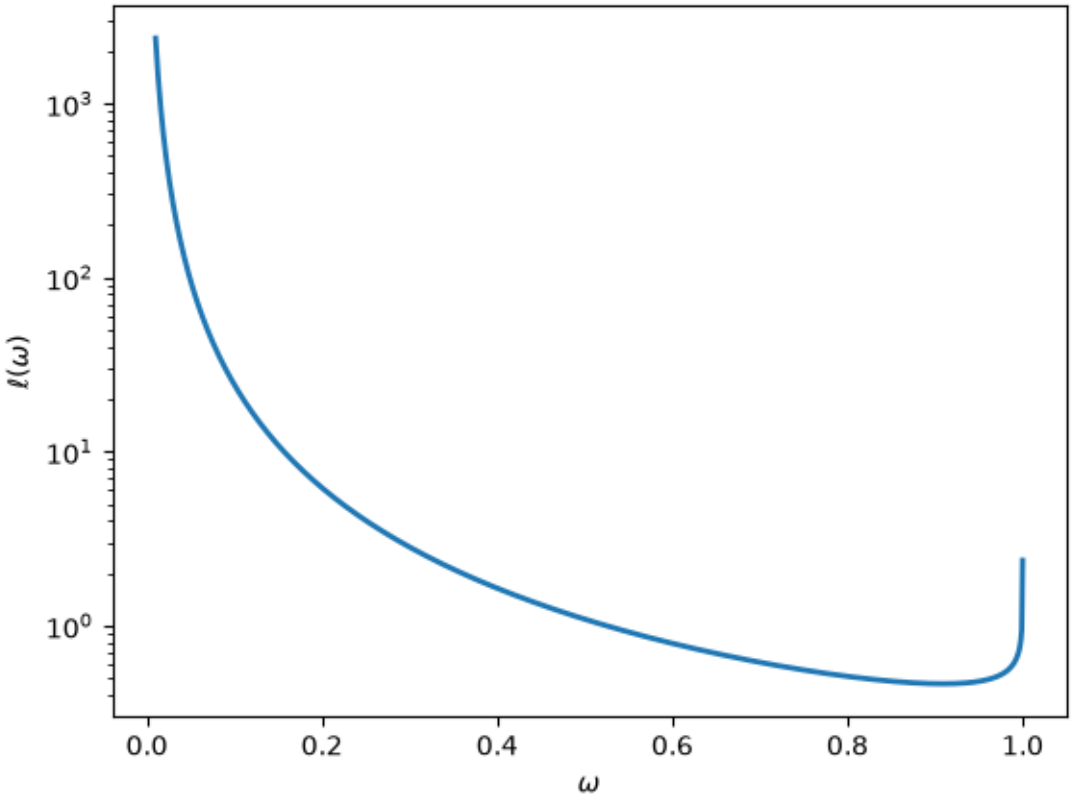


Fig. 25.2: Likelihood ratio $\ell(\omega)$, log scale

Note the log scale on the vertical axis: on a linear scale the entire curve would be flattened against zero by the spike at the left-hand edge.

Fig. 25.1 shows that as $\omega \to 0$, $f(\omega)$ is unchanged while $g(\omega) \to 0$.

Hence the likelihood ratio in Fig. 25.2 diverges to infinity, at rate ω^{-2}.

The same thing happens as $\omega \to 1$, since $G_b > 1$ also forces $g(\omega) \to 0$ there, which explains the upturn at the right-hand edge.

But that divergence is far slower — at rate $(1-\omega)^{-1/5}$ — and it is the behavior near $\omega = 0$ that causes the trouble below.

A Monte Carlo approximation of $E[L(\omega^t)]$ would repeatedly draw a sequence $\omega^t = (\omega_1, \ldots, \omega_t)$ of t independent observations from g, form the product $L(\omega^t) = \prod_{i=1}^{t} \ell(\omega_i)$ for each such sequence, then average these products across independently drawn sequences.

Because $g(\omega) \to 0$ as $\omega \to 0$, such a simulation procedure undersamples a part of the sample space $[0, 1]$ that it is important to visit often in order to do a good job of approximating the mathematical expectation of $\ell(\omega)$.

Every one of the t factors in $L(\omega^t)$ is distorted in the same way, so the problem compounds as t grows.

We illustrate this numerically below.

25.3 Importance sampling

We circumvent the issue by using a *change of distribution* called **importance sampling**.

Instead of drawing from g to generate data during the simulation, we use an alternative distribution h to generate draws of ω.

In *Likelihood Ratio Processes*, h denotes a third density that generates the data; here h denotes an importance distribution that we choose.

The idea is to design h so that it oversamples the region of Ω where $\ell(\omega_t)$ has large values but low density under g.

After we construct a sample in this way, we must then weight each realization by the likelihood ratio of g and h when we compute the empirical mean of the likelihood ratio.

By doing this, we properly account for the fact that we are using h and not g to simulate data.

To illustrate, suppose we were interested in $E[\ell(\omega)]$.

We could simply compute:

$$\hat{E}^g[\ell(\omega)] = \frac{1}{N}\sum_{i=1}^{N} \ell(\omega_i^g)$$

where ω_i^g indicates that ω_i is drawn from g.

But using our insight from importance sampling, we could instead calculate the object:

$$\hat{E}^h\left[\ell(\omega)\frac{g(\omega)}{h(\omega)}\right] = \frac{1}{N}\sum_{i=1}^{N} \ell(\omega_i^h)\frac{g(\omega_i^h)}{h(\omega_i^h)}$$

where ω_i is now drawn from importance distribution h.

Notice that the above two are exactly the same population objects:

$$E^g[\ell(\omega)] = \int_\Omega \ell(\omega)g(\omega)d\omega = \int_\Omega \ell(\omega)\frac{g(\omega)}{h(\omega)}h(\omega)d\omega = E^h\left[\ell(\omega)\frac{g(\omega)}{h(\omega)}\right]$$

The second equality requires $h(\omega) > 0$ wherever $\ell(\omega)g(\omega) = f(\omega) > 0$.

Since we must use an h that has larger mass in parts of the distribution to which g puts low mass, we use $h = Beta(0.5, 0.5)$ as our importance distribution.

The plots compare g and h.

```
g_a, g_b = params.G_a, params.G_b
h_a, h_b = 0.5, 0.5
```

```
w_range = np.linspace(1e-5, 1-1e-5, 1000)

plt.plot(w_range, g(w_range),
         lw=2, label=f'g=Beta({g_a}, {g_b})')
```

(continues on next page)

(continued from previous page)

```
plt.plot(w_range, beta.pdf(w_range, 0.5, 0.5),
         lw=2, label=f'h=Beta({h_a}, {h_b})')
plt.legend()
plt.ylim([0., 3.])
plt.show()
```

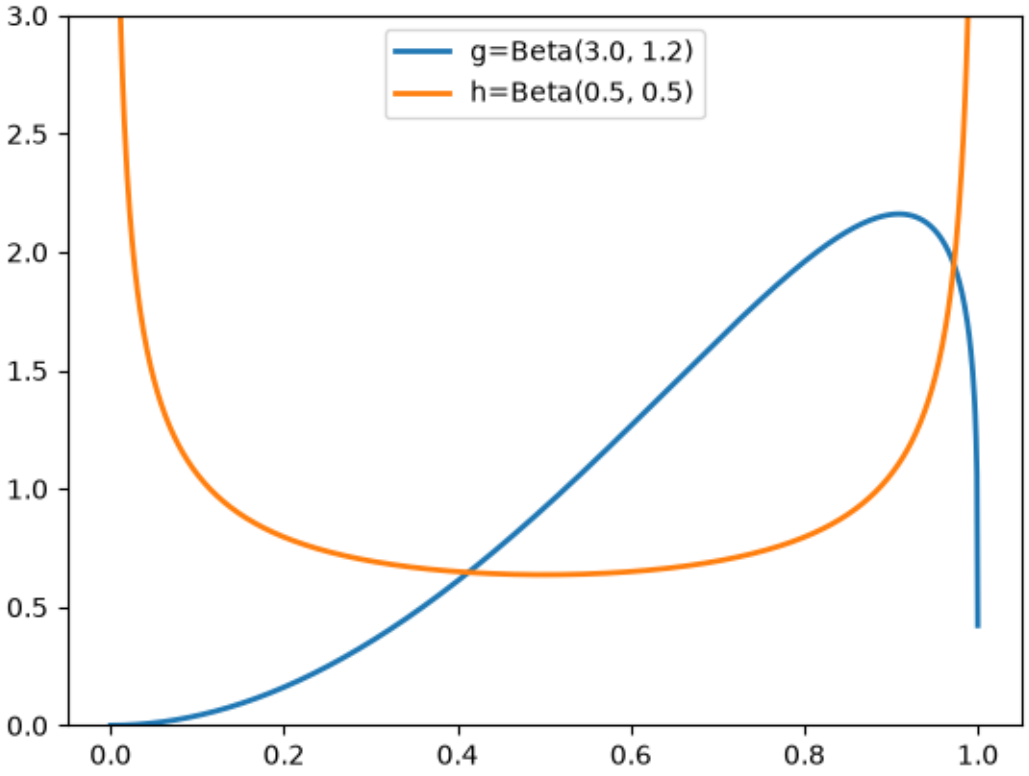


Fig. 25.3: Data and importance sampling distributions

25.4 Approximating a cumulative likelihood ratio

We now study how to use importance sampling to approximate $E\left[L\left(\omega^t\right)\right] = E\left[\prod_{i=1}^t \ell\left(\omega_i\right)\right]$.

To write code for general Beta densities, we let p denote the density under which we want the expectation and q the density from which we draw.

In our application $p = g$ and q is an importance distribution such as h.

In *Likelihood Ratio Processes*, q denotes the density that nature chose; here q denotes the importance distribution.

As above, our plan is to draw sequences ω^t from q and then re-weight the likelihood ratio appropriately.

The change of distribution holds in population, exactly as it did for a single draw:

$$E^p\left[L\left(\omega^t\right)\right] = E^p\left[\prod_{i=1}^t \ell\left(\omega_i\right)\right] = E^q\left[\prod_{i=1}^t \ell\left(\omega_i\right)\frac{p\left(\omega_i\right)}{q\left(\omega_i\right)}\right]$$

This suggests the estimator

$$\hat{E}^q\left[L\left(\omega^t\right)\right] = \frac{1}{N}\sum_{n=1}^{N}\left(\prod_{i=1}^{t}\ell\left(\omega_{n,i}^q\right)\frac{p\left(\omega_{n,i}^q\right)}{q\left(\omega_{n,i}^q\right)}\right)$$

where $\omega_{n,i}^q$ is the i-th observation of the n-th sequence, drawn from the importance distribution q.

Here $\frac{p\left(\omega_{n,i}^q\right)}{q\left(\omega_{n,i}^q\right)}$ is the weight we assign to each data point $\omega_{n,i}^q$.

Below we prepare a Python function for computing the importance sampling estimates given any beta distributions p, q.

```
def estimate_single_path(key, p_a, p_b, q_a, q_b, T):
    """
    Estimation for a single sample path.
    """

    def loop_body(i, carry):
        L, weight, key_state = carry
        key_state, subkey = jax.random.split(key_state)
        w = jax.random.beta(subkey, q_a, q_b)

        # Keep draws off the boundary, where the log densities
        # below produce the undefined form 0 * inf
        w = jnp.clip(w, 1e-12, 1 - 1e-12)

        # Compute likelihood ratio using f/g functions
        likelihood_ratio = f(w) / g(w)
        L = L * likelihood_ratio

        # Importance sampling weight
        p_w = beta.pdf(w, p_a, p_b)
        q_w = beta.pdf(w, q_a, q_b)
        weight = weight * (p_w / q_w)

        return (L, weight, key_state)

    # Use fori_loop for dynamic T values
    final_L, final_weight, _ = jax.lax.fori_loop(
        0, T, loop_body, (1.0, 1.0, key)
    )
    return final_L * final_weight

@partial(jax.jit, static_argnames=['N'])
def estimate(key, p_a, p_b, q_a, q_b, T=1, N=10000):
    """Estimation of a batch of sample paths."""
    keys = jax.random.split(key, N)

    # Vectorize over keys, holding the parameters fixed
    estimates = jax.vmap(
        estimate_single_path,
        in_axes=(0, None, None, None, None, None)
    )(keys, p_a, p_b, q_a, q_b, T)

    return jnp.mean(estimates)
```

Consider the case when $T = 1$, which amounts to approximating $E\left[\ell\left(\omega\right)\right]$.

For the standard Monte Carlo estimate, we can set $p = g$ and $q = g$.

```
estimate(jax.random.key(0), g_a, g_b, g_a, g_b,
         T=1, N=10000)
```

```
Array(0.98300826, dtype=float64)
```

For our importance sampling estimate, we set $q = h$.

```
estimate(jax.random.key(1), g_a, g_b, h_a, h_b,
         T=1, N=10000)
```

```
Array(1.00722995, dtype=float64)
```

Evidently, even at $T = 1$, our importance sampling estimate is closer to 1 than is the Monte Carlo estimate.

Bigger differences arise when computing expectations over longer sequences, $E\left[L\left(\omega^t\right)\right]$.

Setting $T = 10$, we find that the Monte Carlo method severely underestimates the mean while importance sampling still produces an estimate close to its theoretical value of unity.

```
estimate(jax.random.key(2), g_a, g_b, g_a, g_b,
         T=10, N=10000)
```

```
Array(0.60388316, dtype=float64)
```

```
estimate(jax.random.key(3), g_a, g_b, h_a, h_b,
         T=10, N=10000)
```

```
Array(0.99185478, dtype=float64)
```

The Monte Carlo method underestimates because the likelihood ratio $L(\omega^T) = \prod_{t=1}^{T} \frac{f(\omega_t)}{g(\omega_t)}$ has a highly skewed distribution under g.

Most samples from g produce small likelihood ratios, while the true mean requires occasional very large values that are rarely sampled.

In our case, since $g(\omega) \to 0$ as $\omega \to 0$ while $f(\omega)$ remains constant, the Monte Carlo procedure undersamples precisely where the likelihood ratio $\frac{f(\omega)}{g(\omega)}$ is largest.

In fact the situation is worse than skewness — the Monte Carlo estimator has **infinite variance**.

To see this, note that the second moment of a single likelihood ratio is

$$E^g\left[\ell\left(\omega\right)^2\right] = \int_0^1 \left(\frac{f\left(\omega\right)}{g\left(\omega\right)}\right)^2 g\left(\omega\right) d\omega = \int_0^1 \frac{f\left(\omega\right)^2}{g\left(\omega\right)} d\omega$$

Here f is the uniform density, while $g\left(\omega\right)$ vanishes like $\omega^{G_a - 1} = \omega^2$ as $\omega \to 0$.

The integrand therefore behaves like ω^{-2} near the origin, and the integral diverges.

So $\ell\left(\omega\right)$ has no finite variance, and since the draws are independent, neither does $L\left(\omega^t\right)$ for any t.

Likelihood Ratio Processes meets the same divergent integral as the mean of $\ell(\omega)$ under f, since $E^g\left[\ell(\omega)^2\right] = E^f\left[\ell(\omega)\right]$.

The mean $E^g\left[\ell(\omega)\right] = 1$ is finite, so the strong law of large numbers still applies: the Monte Carlo average converges to 1 almost surely as the number of draws N grows.

The ordinary central limit theorem does not apply, and its $N^{-1/2}$ rate of convergence fails.

Near the origin, $g(\omega) \approx k\omega^2$ with $k = 1/B(3, 1.2)$, where B is the Beta function.

For large x, the event $\ell(\omega) > x$ occurs mainly when $\omega < (kx)^{-1/2}$, an event with probability about $\frac{1}{3}k^{-1/2}x^{-3/2}$ under g.

A tail that decays like $x^{-3/2}$ puts $\ell(\omega)$ in the domain of attraction of a stable law with index $3/2$.

By the generalized central limit theorem, the error of a sample mean of N draws shrinks like $N^{-1/3}$, and its limiting distribution is skewed, with a long right tail.

Exercise 25.7.2 confirms the tail index and the $N^{-1/3}$ rate numerically.

As T increases, the problem worsens, making standard Monte Carlo increasingly unreliable.

Importance sampling with $q = h$ fixes the infinite variance by sampling more often near $\omega = 0$, where ℓ is large.

25.5 Distribution of sample mean

We next study the bias and efficiency of the Monte Carlo and importance sampling approaches.

The code below repeats the estimate `N_simu` times, so that we can look at the distribution of the estimates that each method produces.

```
@partial(jax.jit, static_argnames=['N_simu', 'N_samples'])
def simulate(key, p_a, p_b, q_a, q_b, N_simu, T=1,
             N_samples=10000):
    """Repeat the estimate N_simu times, drawing from q."""
    keys = jax.random.split(key, N_simu)
    return jax.vmap(
        lambda k: estimate(k, p_a, p_b, q_a, q_b, T,
                           N_samples)
    )(keys)
```

Setting $q = p$ recovers standard Monte Carlo, since every importance weight is then identically one.

Again, we first consider estimating $E[\ell(\omega)]$ by setting T=1.

We simulate 1000 times for each method.

```
N_simu = 1000
μ_L_g = simulate(jax.random.key(4), g_a, g_b,
                 g_a, g_b, N_simu)
μ_L_h = simulate(jax.random.key(5), g_a, g_b,
                 h_a, h_b, N_simu)
```

```
# standard Monte Carlo (mean and variance)
jnp.mean(μ_L_g), jnp.var(μ_L_g)
```

```
(Array(0.99558545, dtype=float64), Array(0.00448641, dtype=float64))
```

```
# importance sampling (mean and variance)
jnp.mean(μ_L_h), jnp.var(μ_L_h)
```

```
(Array(0.99985196, dtype=float64), Array(2.46227759e-05, dtype=float64))
```

Although both methods tend to provide a mean estimate of $E[\ell(\omega)]$ close to 1, the importance sampling estimates have smaller variance.

Next, we present distributions of estimates for $\hat{E}[L(\omega^t)]$, in cases for $T = 1, 5, 10, 20$.

```
def simulate_multiple_T(key, p_a, p_b, q_a, q_b, T_values,
                        N_simu, N_samples=10000):
    """Run simulate once per T, returning a dict keyed by T."""
    keys = jax.random.split(key, len(T_values))
    return {T: simulate(keys[i], p_a, p_b, q_a, q_b,
                        N_simu, T, N_samples)
            for i, T in enumerate(T_values)}
```

The next function draws the histograms that we use to compare the two methods.

```
def plot_estimates(T_values, mc, imp, imp_label, n_rows=1):
    """Compare Monte Carlo and importance sampling estimates."""
    n_cols = len(T_values) // n_rows
    fig, axs = plt.subplots(n_rows, n_cols,
                            figsize=(14, 5 * n_rows))
    μ_range = np.linspace(0, 2, 100)

    for ax, T in zip(np.ravel(axs), T_values):
        μ_L_g, μ_L_h = np.asarray(mc[T]), np.asarray(imp[T])

        ax.set_xlabel('$μ_L$')
        ax.set_ylabel('frequency')
        ax.set_title(f'$T$={T}')
        ax.hist(μ_L_g, bins=μ_range,
                color='r', alpha=0.5, label='$g$ generating')
        ax.hist(μ_L_h, bins=μ_range,
                color='b', alpha=0.5, label=imp_label)
        ax.legend(loc=4)

        # Summarize each distribution in an upper corner
        for μ_L, color, x, ha in ((μ_L_g, 'r', 0.02, 'left'),
                                  (μ_L_h, 'b', 0.98, 'right')):
            ax.text(x, 0.98, transform=ax.transAxes, ha=ha,
                    va='top', color=color, fontsize=9,
                    s=r'$\hat{μ}$=' + f'{np.mean(μ_L):.3g}' +
                      '\n' + 'med=' + f'{np.median(μ_L):.3g}' +
                      '\n' + r'$\hat{σ}^2$=' + f'{np.var(μ_L):.3g}')

    plt.show()
```

The Monte Carlo estimates do not depend on the importance distribution, so we compute them once and reuse them in each of the figures below.

```
T_values = [1, 5, 10, 20]

mc = simulate_multiple_T(jax.random.key(6), g_a, g_b,
                         g_a, g_b, T_values, N_simu)
```

```
imp_h1 = simulate_multiple_T(jax.random.key(7), g_a, g_b,
                             h_a, h_b, T_values, N_simu)

plot_estimates(T_values, mc, imp_h1, '$h$ generating',
               n_rows=2)
```

The simulation exercises above show that the importance sampling estimates stay centered on 1 for every T, while the distribution of the standard Monte Carlo estimates drifts steadily to the left as T increases.

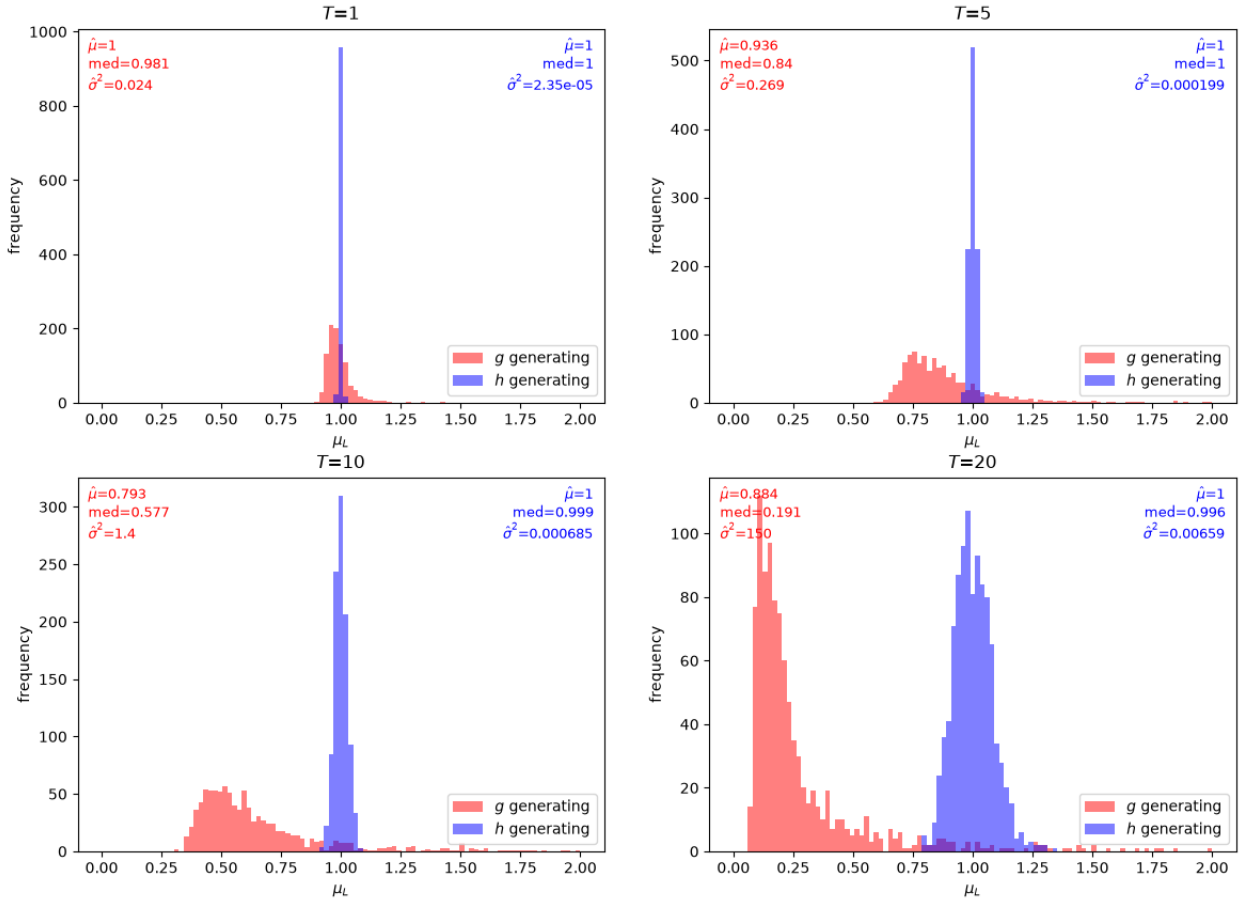


Fig. 25.4: Monte Carlo and importance sampling estimates

It is worth being careful about what is going wrong here.

The Monte Carlo estimator is unbiased in population — its expectation is exactly 1 for every T.

What deteriorates as T grows is the *shape* of its sampling distribution: almost all of the mass collapses towards zero, while the expectation is sustained by rare, very large outliers.

This is why each panel reports a median alongside the mean.

The median falls steadily in T, whereas the reported mean $\hat{\mu}$ is an unreliable statistic — recall that the underlying variance is infinite, so the $\hat{\sigma}^2$ printed in each panel estimates a quantity that does not exist, and a different seed can move $\hat{\mu}$ a long way.

Importance sampling with h removes the infinite variance, but it does not stop the variance from growing with T.

Drawing a path from h and reweighting gives the path the value $\prod_{i=1}^{T} \ell(\omega_i) g(\omega_i)/h(\omega_i) = \prod_{i=1}^{T} f(\omega_i)/h(\omega_i)$.

Its second moment is $\left(\int_0^1 f^2/h \, d\omega\right)^T = \left(1 + \chi^2(f\|h)\right)^T$, where $\chi^2(f\|h) = \int_0^1 (f-h)^2/h \, d\omega$ is the χ^2 divergence of f from h (see *Statistical Divergence Measures*).

For $h = \text{Beta}(0.5, 0.5)$, each weight $f(\omega)/h(\omega) = \pi\sqrt{\omega(1-\omega)}$ is at most $\pi/2$, and $\int_0^1 f^2/h \, d\omega = \pi^2/8 \approx 1.234$.

The variance of a single path's contribution, $(\pi^2/8)^T - 1$, is finite but grows geometrically: it equals 66 at $T = 20$ and about 3.6×10^4 at $T = 50$.

25.6 Choosing a sampling distribution

Above, we arbitrarily chose $h = Beta(0.5, 0.5)$ as the importance distribution.

Is there an optimal importance distribution?

An importance sampling estimator has zero variance when the importance density is proportional to the integrand, here ℓg.

Since $\ell g = f$, the zero-variance importance density for a single draw is f itself; for a path of length T it is the product density $f(\omega_1) f(\omega_2) \cdots f(\omega_T)$.

To use it we need only be able to draw from f; we do not need to know in advance that $E\left[L\left(\omega^t\right)\right] = 1$.

Thus, suppose that we simply use $h = f$.

When estimating the mean of the likelihood ratio (T=1), we get:

$$\hat{E}^f\left[\ell(\omega)\frac{g(\omega)}{f(\omega)}\right] = \hat{E}^f\left[\frac{f(\omega)}{g(\omega)}\frac{g(\omega)}{f(\omega)}\right] = \frac{1}{N}\sum_{i=1}^{N} \ell\left(\omega_i^f\right)\frac{g\left(\omega_i^f\right)}{f\left(\omega_i^f\right)} = 1$$

```
μ_L_f = simulate(jax.random.key(8), g_a, g_b,
                 params.F_a, params.F_b, N_simu)
```

```
# importance sampling (mean and variance)
jnp.mean(μ_L_f), jnp.var(μ_L_f)
```

```
(Array(1., dtype=float64), Array(0., dtype=float64))
```

We could also use other distributions as our importance distribution.

Below we choose just a few and compare their sampling properties.

```
a_list = [0.5, 1., 2.]
b_list = [0.5, 1.2, 5.]
```

```
w_range = np.linspace(1e-5, 1-1e-5, 1000)

plt.plot(w_range, g(w_range),
         lw=2, label=f'g=Beta({g_a}, {g_b})')
plt.plot(w_range, beta.pdf(w_range, a_list[0], b_list[0]),
         lw=2, label=f'$h_1$=Beta({a_list[0]},{b_list[0]})')
plt.plot(w_range, beta.pdf(w_range, a_list[1], b_list[1]),
         lw=2, label=f'$h_2$=Beta({a_list[1]},{b_list[1]})')
plt.plot(w_range, beta.pdf(w_range, a_list[2], b_list[2]),
         lw=2, label=f'$h_3$=Beta({a_list[2]},{b_list[2]})')
plt.legend()
plt.ylim([0., 3.])
plt.show()
```

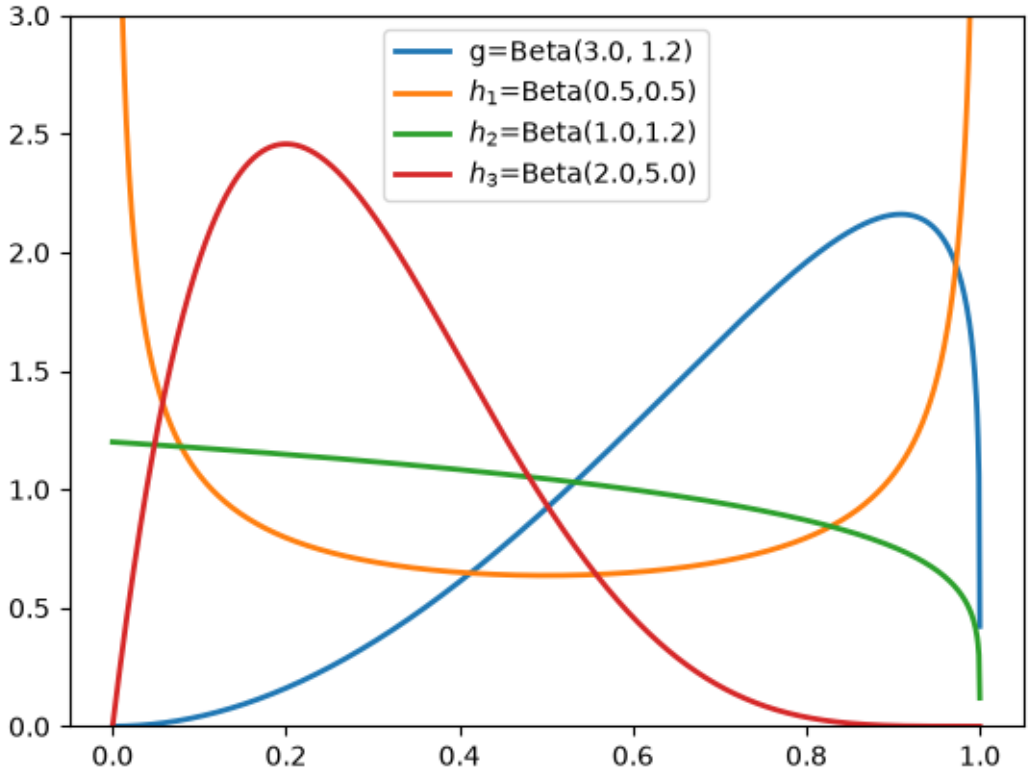


Fig. 25.5: Comparison of importance sampling distributions

We consider two additional distributions.

As a reminder h_1 is the original $Beta(0.5, 0.5)$ distribution that we used above.

h_2 is the $Beta(1, 1.2)$ distribution.

Note how h_2 has a similar shape to g at higher values of ω but more mass at lower values.

Our hunch is that h_2 should be a good importance sampling distribution.

h_3 is the $Beta(2,5)$ distribution.

Note how h_3 has almost no mass at values very close to 0 and at values close to 1.

Our hunch is that h_3 will be a poor importance sampling distribution.

The variance calculation above lets us make these hunches precise.

Drawing from h and reweighting gives each observation the value $\ell(\omega) g(\omega) / h(\omega) = f(\omega) / h(\omega)$, so the estimator has finite variance exactly when

$$\int_0^1 \frac{f(\omega)^2}{h(\omega)} d\omega < \infty$$

For a Beta(a,b) importance density h and the uniform f,

$$\int_0^1 \frac{f(\omega)^2}{h(\omega)} d\omega = B(a,b) \int_0^1 \omega^{1-a}(1-\omega)^{1-b} d\omega = B(a,b)\, B(2-a, 2-b)$$

when $a < 2$ and $b < 2$, and the integral is infinite otherwise.

The integral equals $\pi^2/8 \approx 1.234$ for h_1, $1/0.96 \approx 1.042$ for h_2, and 1 for $h = f$.

The proposal h_2 is good because it is close to f.

For h_3 it diverges at **both** endpoints, since h_3 vanishes like ω at the origin and like $(1-\omega)^4$ at one.

So importance sampling with h_3 has infinite variance, just as standard Monte Carlo does — which is what makes it a poor choice.

We first simulate and plot the distribution of estimates for $\hat{E}[L(\omega^t)]$ using h_2 as the importance sampling distribution.

```
T_values_h2 = [1, 20]
imp_h2 = simulate_multiple_T(jax.random.key(9), g_a, g_b,
                             a_list[1], b_list[1],
                             T_values_h2, N_simu)

plot_estimates(T_values_h2, mc, imp_h2, '$h_2$ generating')
```

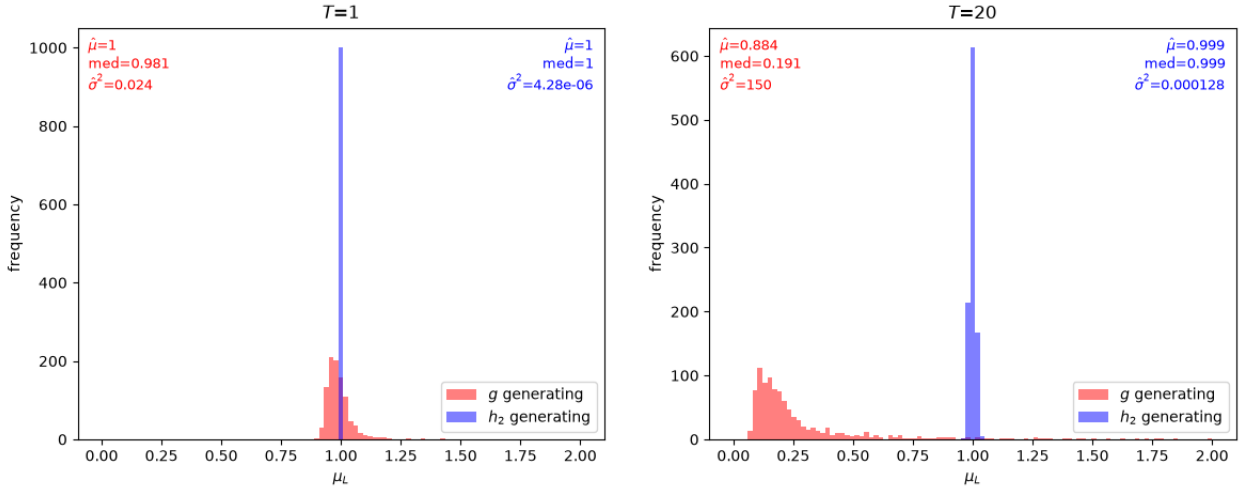


Fig. 25.6: Estimates using importance distribution h_2

Our simulations suggest that indeed h_2 is a quite good importance sampling distribution for our problem.

Even at $T = 20$, the mean is very close to 1 and the variance is small.

```
T_values_h3 = [1, 20]
imp_h3 = simulate_multiple_T(jax.random.key(10), g_a, g_b,
                             a_list[2], b_list[2],
                             T_values_h3, N_simu)

plot_estimates(T_values_h3, mc, imp_h3, '$h_3$ generating')
```

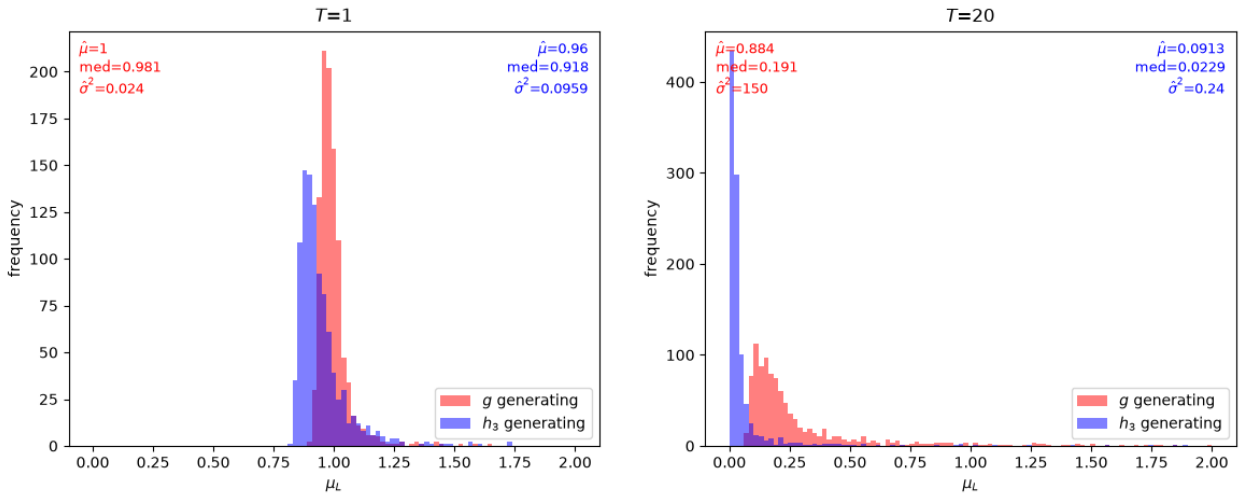


Fig. 25.7: Estimates using importance distribution h_3

However, h_3 is evidently a poor importance sampling distribution for our problem, with a mean estimate far away from 1 for $T = 20$.

Because $h_3 > 0$ on $(0, 1)$, the importance sampling estimator with h_3 is unbiased for every T, just as the Monte Carlo estimator under g is.

The trouble lies in the shape of its sampling distribution.

Even at $T = 1$, the estimates under h_3 are more dispersed than the Monte Carlo estimates under g, and the median estimate lies farther below 1.

Near $\omega = 1$, h_3 vanishes like $(1 - \omega)^4$ while $f = 1$, so the weight f/h_3 has tail index $5/4$ under h_3, smaller than the tail index $3/2$ of f/g under g.

The error of a sample mean of N draws therefore shrinks like $N^{-1/5}$ under h_3 and like $N^{-1/3}$ under g (Exercise 25.7.2).

Thus, our simulations suggest that for our problem we would be better off simply using Monte Carlo approximations under g than using h_3 as an importance sampling distribution.

25.7 Exercises

Exercise 25.7.1

Consider estimating $E^g\left[L\left(\omega^T\right)\right] = 1$ by drawing paths of length T from a Beta(a, b) importance density q.

1. Show that each path contributes $\prod_{i=1}^T f(\omega_i)/q(\omega_i)$ to the estimator and that the second moment of this contribution is $M(a,b)^T$, where $M(a,b) = \int_0^1 f^2/q\, d\omega = B(a,b)B(2-a, 2-b)$ when $a, b < 2$ and $M(a,b) = +\infty$ otherwise.
2. Use the Cauchy–Schwarz inequality to show that $M(a,b) \geq 1$, with equality only at $(a,b) = (1,1)$. Plot contours of $\log M(a,b)$ on $(0,2) \times (0,2)$.
3. The effective sample size of N weighted paths is approximately $N/M(a,b)^T$. Plot the ratio of the effective sample size to N against T for h_1 and h_2. How many paths does h_1 need for an effective sample size of 1000 at $T = 20$ and at $T = 50$?

Solution

Part 1. Each path contributes $\prod_{i=1}^T \ell(\omega_i)g(\omega_i)/q(\omega_i) = \prod_{i=1}^T f(\omega_i)/q(\omega_i)$.

Because the ω_i's are IID draws from q, the second moment is $\left(E^q\left[(f/q)^2\right]\right)^T = M(a,b)^T$ with $M(a,b) = \int_0^1 f^2/q\, d\omega$.

Since $f = 1$, $M(a,b) = B(a,b)\int_0^1 \omega^{1-a}(1-\omega)^{1-b}d\omega$, which equals $B(a,b)B(2-a,2-b)$ when $a < 2$ and $b < 2$ and is infinite otherwise.

Part 2. By the Cauchy–Schwarz inequality,

$$1 = \left(\int_0^1 f\, d\omega\right)^2 = \left(\int_0^1 \frac{f}{\sqrt{q}}\sqrt{q}\, d\omega\right)^2 \leq \int_0^1 \frac{f^2}{q} d\omega \int_0^1 q\, d\omega = M(a,b),$$

with equality if and only if $f/\sqrt{q}$ is proportional to $\sqrt{q}$, that is, $q = f$.

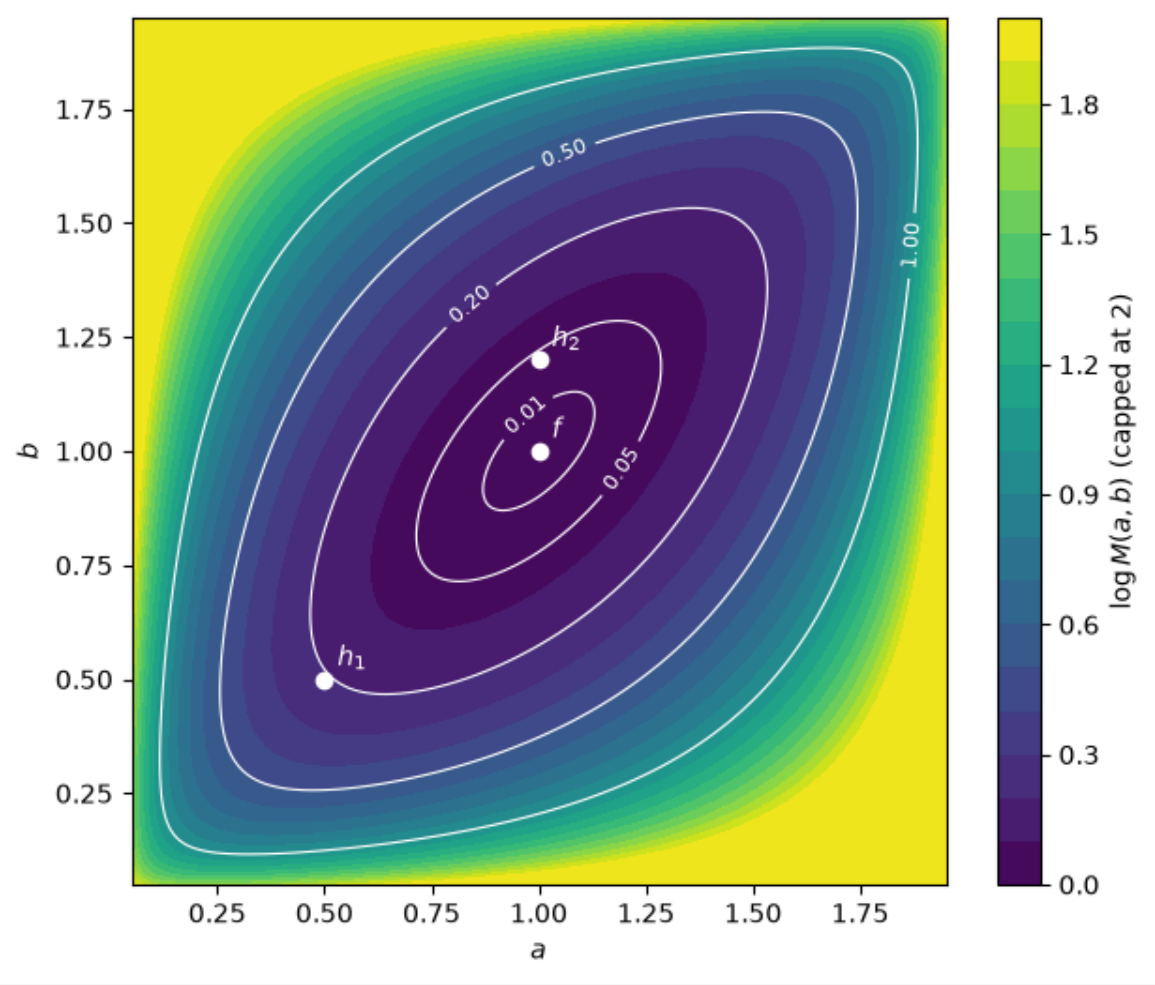
The variance of a single weight is $M(a,b) - 1 = \chi^2(f\|q)$.

```
from scipy.special import betaln

def M(a, b):
    """Second moment of the importance weight f/q for q = Beta(a, b)."""
    if a < 2 and b < 2:
        return np.exp(betaln(a, b) + betaln(2 - a, 2 - b))
    return np.inf

grid = np.linspace(0.05, 1.95, 191)
log_M = np.array([[np.log(M(a, b)) for a in grid] for b in grid])

fig, ax = plt.subplots(figsize=(7, 6))
cs = ax.contourf(grid, grid, np.minimum(log_M, 2),
                 levels=np.linspace(0, 2, 21), cmap='viridis')
fig.colorbar(cs, label=r'$\log M(a, b)$ (capped at 2)')
lines = ax.contour(grid, grid, log_M, levels=[0.01, 0.05, 0.2, 0.5, 1],
                   colors='w', linewidths=0.8)
ax.clabel(lines, fmt='%.2f', fontsize=8)
for (a, b), name in zip([(0.5, 0.5), (1.0, 1.2), (1.0, 1.0)],
                        ['$h_1$', '$h_2$', '$f$']):
    ax.plot(a, b, 'wo')
    ax.annotate(name, (a, b), xytext=(5, 5), textcoords='offset points',
                color='w')
ax.set_xlabel('$a$')
ax.set_ylabel('$b$')
plt.show()
```

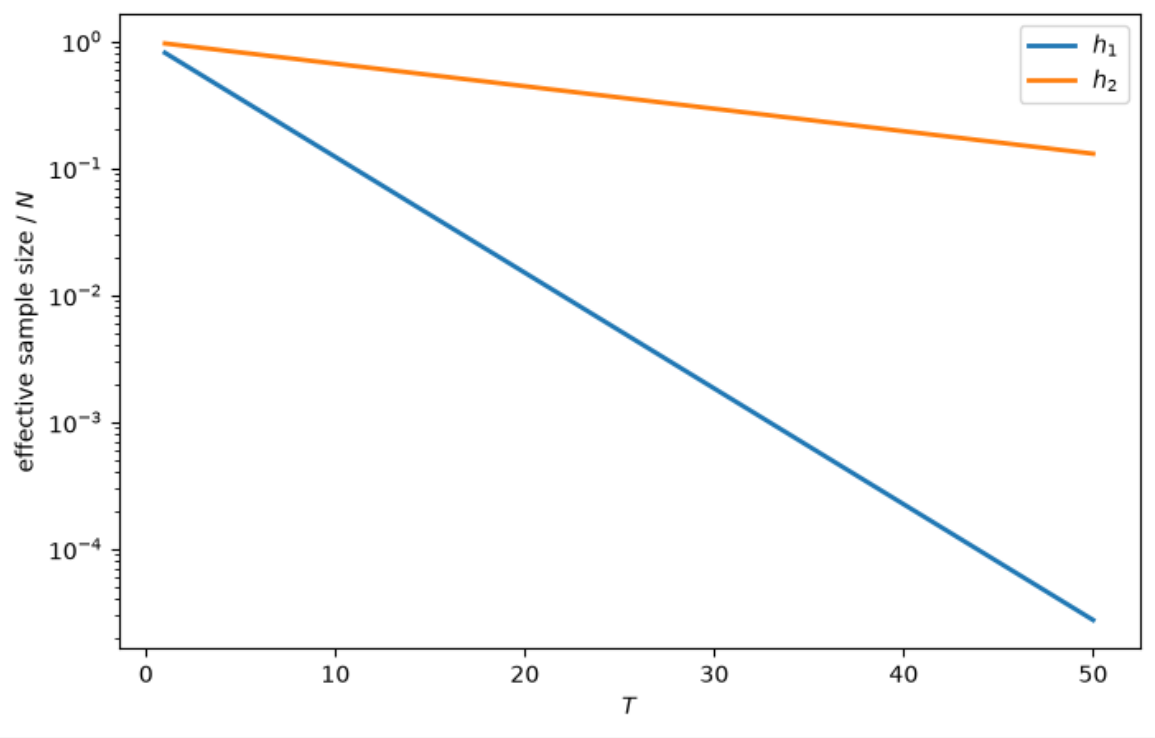



```
Minimum of M on the grid at (a, b) = (1.00, 1.00)
M(h_1) = 1.2337, M(h_2) = 1.0417
```

Part 3.

```
T_grid = np.arange(1, 51)
plt.figure(figsize=(8, 5))
for (a, b), name in zip([(0.5, 0.5), (1.0, 1.2)], ['$h_1$', '$h_2$']):
    plt.semilogy(T_grid, M(a, b)**(-T_grid), lw=2, label=name)
plt.xlabel('$T$')
plt.ylabel('effective sample size / $N$')
plt.legend()
plt.show()

for T in [20, 50]:
    print(f"T = {T}: h_1 needs N = {1000 * M(0.5, 0.5)**T:,.0f} paths, "
          f"h_2 needs N = {1000 * M(1, 1.2)**T:,.0f} paths")
```




```
T = 20: h_1 needs N = 66,711 paths, h_2 needs N = 2,262 paths
T = 50: h_1 needs N = 36,348,619 paths, h_2 needs N = 7,699 paths
```

The effective sample size under h_1 shrinks by the factor $8/\pi^2$ with each additional period.

Exercise 25.7.2

1. Let $\omega \sim g$. Show that $\Pr\{\ell(\omega) > x\}$ is approximately proportional to $x^{-3/2}$ for large x, so that $\ell(\omega)$ has tail index $3/2$ under g.
2. Let $\omega \sim h_3 = \text{Beta}(2, 5)$. Show that the weight $f(\omega)/h_3(\omega)$ has tail index $5/4$ under h_3. Which endpoint of $(0, 1)$ determines the tail index?
3. Confirm both tail indices with the Hill estimator [Hill, 1975], which uses the k largest of n draws $X_{(1)} \geq X_{(2)} \geq \cdots \geq X_{(k+1)}$ to estimate the tail index by

$$\hat{\alpha}_k = \left(\frac{1}{k} \sum_{i=1}^{k} \log \frac{X_{(i)}}{X_{(k+1)}} \right)^{-1}.$$

4. Set $T = 1$. For $N = 10^2, 10^3, 10^4, 10^5$, compute 500 replicate sample means of f/q for N draws from q when q is g (standard Monte Carlo), h_2, and h_3. Plot the interquartile range of the replicate means against N on log-log scales, and compare the slopes with $-1/3$, $-1/2$, and $-1/5$.

Solution

Part 1. Near the origin, $g(\omega) \approx k\omega^2$ with $k = 1/B(3, 1.2)$.

For large x, $\ell(\omega) = 1/g(\omega) > x$ occurs mainly when $\omega < (kx)^{-1/2}$.

Since $\Pr_g\{\omega < u\} \approx ku^3/3$ for small u,

$$\Pr_g\{\ell(\omega) > x\} \approx \frac{k}{3}(kx)^{-3/2} = \frac{1}{3}k^{-1/2}x^{-3/2}.$$

Near $\omega = 1$, g vanishes like $(1-\omega)^{0.2}$, which contributes a tail proportional to x^{-6}, too thin to matter.

Part 2. The density $h_3(\omega) = 30\,\omega(1-\omega)^4$ vanishes at both endpoints.

Near $\omega = 0$, $1/h_3 > x$ occurs when $\omega < 1/(30x)$, which has probability about $15/(30x)^2$, proportional to x^{-2}.

Near $\omega = 1$, $1/h_3 > x$ occurs when $1 - \omega < (30x)^{-1/4}$, which has probability about $6(30x)^{-5/4}$, proportional to $x^{-5/4}$.

The heavier tail, with index $5/4$, comes from the right endpoint.

By the generalized central limit theorem, a tail index $\alpha \in (1, 2)$ makes the error of a sample mean shrink like $N^{1/\alpha - 1}$: $N^{-1/3}$ for $\alpha = 3/2$ and $N^{-1/5}$ for $\alpha = 5/4$.

Part 3.

```
def log_beta_pdf(w, a, b):
    """Log density of Beta(a, b)."""
    return (a - 1) * np.log(w) + (b - 1) * np.log1p(-w) - betaln(a, b)

def hill(x, k):
    """Hill estimator of the tail index from the k largest values of x."""
    x_sorted = np.sort(x)[::-1]
    return 1 / np.mean(np.log(x_sorted[:k] / x_sorted[k]))

rng = np.random.default_rng(0)
n = 1_000_000
w_g = rng.beta(g_a, g_b, n)
w_h3 = rng.beta(2, 5, n)
weights = {'f/g under g': np.exp(-log_beta_pdf(w_g, g_a, g_b)),
           'f/h_3 under h_3': np.exp(-log_beta_pdf(w_h3, 2, 5))}

for name, x in weights.items():
    estimates = [f"{hill(x, k):.3f}" for k in (500, 1000, 2000, 5000)]
    print(f"{name}: Hill estimates for k = 500, 1000, 2000, 5000: {estimates}")
```

```
f/g under g: Hill estimates for k = 500, 1000, 2000, 5000: ['1.483', '1.476',
 ↪'1.497', '1.514']
f/h_3 under h_3: Hill estimates for k = 500, 1000, 2000, 5000: ['1.248', '1.265
 ↪', '1.267', '1.271']
```

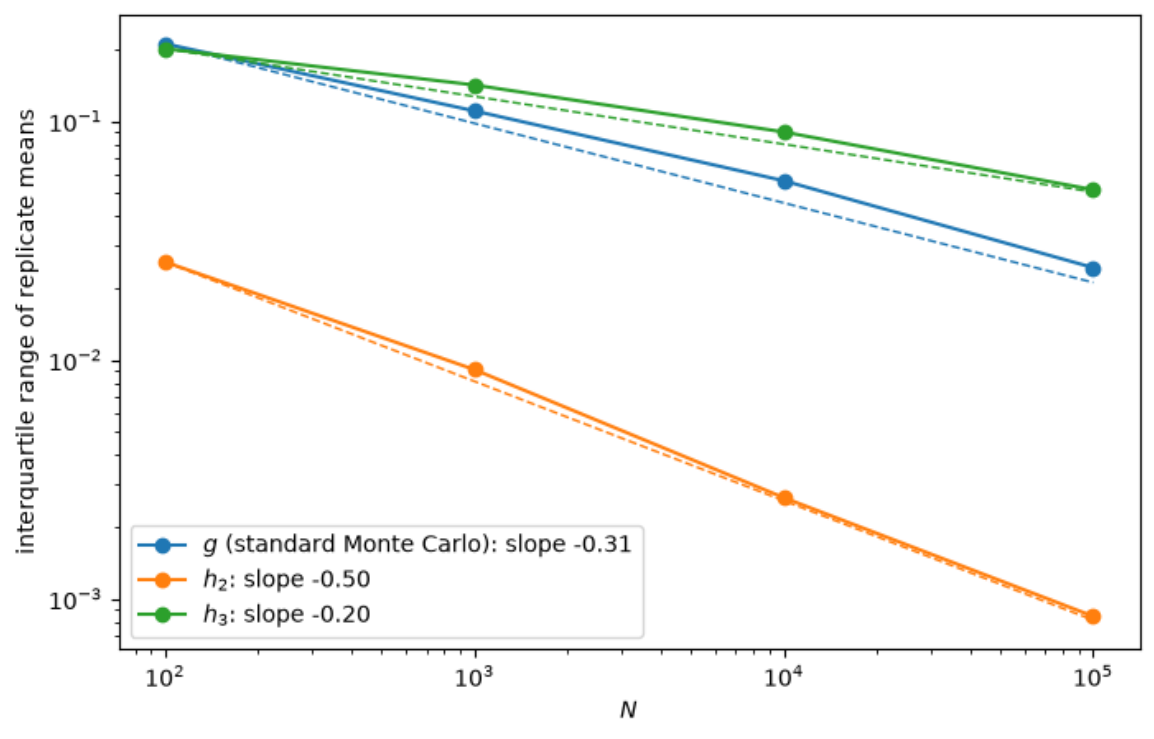
The Hill estimates are close to $3/2$ and $5/4$.

Part 4.

```
N_values = np.array([10**2, 10**3, 10**4, 10**5])
R = 500
proposals = {'$g$ (standard Monte Carlo)': (g_a, g_b, 1/3),
             '$h_2$': (1.0, 1.2, 1/2),
             '$h_3$': (2.0, 5.0, 1/5)}

plt.figure(figsize=(8, 5))
for name, (a, b, rate) in proposals.items():
    iqr = []
    for N in N_values:
        means = np.empty(R)
```

```
        for r in range(R):
            w = np.clip(rng.beta(a, b, N), 1e-16, 1 - 1e-16)
            means[r] = np.mean(np.exp(-log_beta_pdf(w, a, b)))
        q75, q25 = np.quantile(means, [0.75, 0.25])
        iqr.append(q75 - q25)
    slope = np.polyfit(np.log(N_values), np.log(iqr), 1)[0]
    line, = plt.loglog(N_values, iqr, 'o-',
                       label=f'{name}: slope {slope:.2f}')
    plt.loglog(N_values, iqr[0] * (N_values / N_values[0])**(-rate),
               '--', color=line.get_color(), lw=1)
plt.xlabel('$N$')
plt.ylabel('interquartile range of replicate means')
plt.legend()
plt.show()
```



The dashed lines have slopes $-1/3$, $-1/2$, and $-1/5$.

The fitted slopes are close to them.

Exercise 25.7.3

In *this section* of *Likelihood Ratio Processes*, the probability $\alpha_T = \Pr_f\{L(\omega^T) < 1\}$ of a type I error decays exponentially in T, so standard Monte Carlo cannot estimate it for large T.

Let ϕ^* minimize the Chernoff coefficient $\rho(\phi) = \int_0^1 f^\phi g^{1-\phi} d\omega$, and let $q = f^{\phi^*} g^{1-\phi^*}/\rho(\phi^*)$ be the tilted density.

1. Show that q is the Beta$(\phi^* F_a + (1-\phi^*)G_a,\ \phi^* F_b + (1-\phi^*)G_b)$ density, and compute ϕ^* and the parameters of q.

2. Show that

$$\alpha_T = \rho(\phi^*)^T \, E^q \left[\mathbf{1}\{L(\omega^T) < 1\} \, L(\omega^T)^{1-\phi^*} \right].$$

3. For $T = 30, 60, 100$, estimate α_T by standard Monte Carlo with $N = 20{,}000$ paths drawn from f and by importance sampling with $N = 20{,}000$ paths drawn from q. Report the relative standard error of the importance sampling estimate.

Solution

Part 1. For Beta densities,

$$f(\omega)^\phi g(\omega)^{1-\phi} \propto \omega^{\phi F_a + (1-\phi) G_a - 1} (1-\omega)^{\phi F_b + (1-\phi) G_b - 1},$$

which is proportional to the stated Beta density.

Part 2. Since $f/q = \rho(\phi^*)(f/g)^{1-\phi^*}$,

$$\alpha_T = E^q \left[\mathbf{1}\{L(\omega^T) < 1\} \prod_{i=1}^T \frac{f(\omega_i)}{q(\omega_i)} \right] = \rho(\phi^*)^T \, E^q \left[\mathbf{1}\{L(\omega^T) < 1\} \, L(\omega^T)^{1-\phi^*} \right].$$

On the event $L(\omega^T) < 1$, the weight $\rho(\phi^*)^T L(\omega^T)^{1-\phi^*}$ is at most $\rho(\phi^*)^T = e^{-C(f,g)T}$, so the estimator has a small variance.

At ϕ^*, the derivative of $\log \rho(\phi)$, which equals $E^q \log \ell(\omega)$, is zero.

Thus $\log L(\omega^T)$ has mean zero under q, and the rare event $L(\omega^T) < 1$ under f occurs about half the time under q.

Choosing q in this way is called exponential tilting [Bucklew, 2004].

Part 3.

```
from scipy.optimize import minimize_scalar

F_a, F_b = params.F_a, params.F_b

def log_ρ(ϕ):
    """Log Chernoff coefficient for f = Beta(F_a, F_b) and g = Beta(g_a, g_b)."""
    return (betaln(ϕ * F_a + (1 - ϕ) * g_a, ϕ * F_b + (1 - ϕ) * g_b)
            - ϕ * betaln(F_a, F_b) - (1 - ϕ) * betaln(g_a, g_b))

res = minimize_scalar(log_ρ, bounds=(0, 1), method='bounded')
ϕ_star, log_ρ_star = res.x, res.fun
q_a = ϕ_star * F_a + (1 - ϕ_star) * g_a
q_b = ϕ_star * F_b + (1 - ϕ_star) * g_b
print(f"ϕ* = {ϕ_star:.4f}, C(f, g) = {-log_ρ_star:.4f}, "
      f"q = Beta({q_a:.4f}, {q_b:.4f})")

rng = np.random.default_rng(1)
N = 20_000

for T in [30, 60, 100]:
    # Standard Monte Carlo under f
    w = rng.uniform(size=(N, T))
    log_L = np.sum(log_beta_pdf(w, F_a, F_b) - log_beta_pdf(w, g_a, g_b), axis=1)
    α_mc = np.mean(log_L < 0)

    # Importance sampling under the tilted density q
    w = np.clip(rng.beta(q_a, q_b, (N, T)), 1e-300, 1 - 1e-16)
    log_L = np.sum(log_beta_pdf(w, F_a, F_b) - log_beta_pdf(w, g_a, g_b), axis=1)
    X = np.where(log_L < 0, np.exp(T * log_ρ_star + (1 - ϕ_star) * log_L), 0.0)
    α_is = np.mean(X)
    rel_se = np.std(X) / np.sqrt(N) / α_is

    print(f"T = {T:>3}: standard MC = {α_mc:.2e}, "
```

```
ϕ* = 0.5969, C(f, g) = 0.1212, q = Beta(1.8063, 1.0806)
T =  30: standard MC = 4.15e-03, importance sampling = 4.401e-03 (relative s.e.␣
↪0.011)
T =  60: standard MC = 1.00e-04, importance sampling = 8.448e-05 (relative s.e.␣
↪0.013)
T = 100: standard MC = 0.00e+00, importance sampling = 5.143e-07 (relative s.e.␣
↪0.015)
```

At $T = 30$ the two estimates agree.

At $T = 60$, standard Monte Carlo rests on a handful of paths with $L(\omega^T) < 1$, while importance sampling has a relative standard error near 1 percent.

At $T = 100$, standard Monte Carlo records no paths with $L(\omega^T) < 1$, while importance sampling still delivers an estimate with a relative standard error near 1.5 percent.

CHAPTER

TWENTYSIX

A PROBLEM THAT STUMPED MILTON FRIEDMAN

(and that Abraham Wald solved by inventing sequential analysis)

Contents

- *A Problem that Stumped Milton Friedman*
  - *Overview*
  - *Source of the Problem*
  - *Neyman-Pearson formulation*
  - *Wald's sequential formulation*
  - *Links between A, B and α, β*
  - *Simulations*
  - *Related lectures*
  - *Exercises*

26.1 Overview

This is the first of three lectures about a statistical decision problem that a US Navy Captain presented to Milton Friedman and W. Allen Wallis during World War II when they were analysts at the U.S. Government's Statistical Research Group at Columbia University.

This problem led Abraham Wald [Wald, 1947] to formulate **sequential analysis**, an approach to statistical decision problems that is intimately related to dynamic programming.

In the spirit of *this earlier lecture*, the present lecture and its *sequel* approach the problem from two distinct points of view, one frequentist, the other Bayesian.

In this lecture, we describe Wald's formulation of the problem from the perspective of a frequentist statistician working within the Neyman-Pearson tradition who thinks about testing hypotheses and consequently uses laws of large numbers to investigate limiting properties of particular statistics under a given **hypothesis**, i.e., a vector of **parameters** that pins down a particular member of a manifold of statistical models that interest the statistician.

- From *this lecture on frequentist and bayesian statistics*, please remember that a frequentist statistician routinely calculates functions of sequences of random variables, conditioning on a vector of parameters.

In *this related lecture* we'll discuss another formulation that adopts the perspective of a **Bayesian statistician** who views parameters as random variables that are jointly distributed with observable variables that he is concerned about.

Because we are taking a frequentist perspective that is concerned about relative frequencies conditioned on alternative parameter values, i.e., alternative **hypotheses**, key ideas in this lecture

- Type I and type II statistical errors
  - a type I error occurs when you reject a null hypothesis that is true
  - a type II error occurs when you accept a null hypothesis that is false
- The **power** of a frequentist statistical test
- The **size** of a frequentist statistical test
- The **critical region** of a statistical test
- A **most powerful test**
- The role of a Law of Large Numbers (LLN) in interpreting **power** and **size** of a frequentist statistical test
- Abraham Wald's **sequential probability ratio test**

We'll begin with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from numba import njit, prange, vectorize, jit
from numba.experimental import jitclass
from math import gamma
from scipy.special import betaln, digamma
from scipy.stats import beta
from collections import namedtuple
import pandas as pd
import scipy as sp
```

This lecture uses ideas studied in *the lecture on likelihood ratio processes* and *the lecture on Bayesian learning*.

26.2 Source of the Problem

On pages 137-139 of his 1998 book *Two Lucky People* with Rose Friedman [Friedman and Friedman, 1998], Milton Friedman described a problem presented to him and Allen Wallis during World War II, when they worked at the US Government's Statistical Research Group at Columbia University.

Note

See pages 25 and 26 of Allen Wallis's 1980 article [Wallis, 1980] about the Statistical Research Group at Columbia University during World War II for his account of the episode and for important contributions that Harold Hotelling made to formulating the problem. Also see chapter 5 of Jennifer Burns' book about Milton Friedman [Burns, 2023].

Let's listen to Milton Friedman tell us what happened

> In order to understand the story, it is necessary to have an idea of a simple statistical problem, and of the standard procedure for dealing with it. The actual problem out of which sequential analysis grew will serve. The Navy has two alternative designs (say A and B) for a projectile. It wants to determine which is superior. To do so it undertakes a series of paired firings. On each round, it assigns the value 1 or 0 to A accordingly as

> its performance is superior or inferior to that of B and conversely 0 or 1 to B. The Navy asks the statistician how to conduct the test and how to analyze the results.
>
> The standard statistical answer was to specify a number of firings (say 1,000) and a pair of percentages (e.g., 53% and 47%) and tell the client that if A receives a 1 in more than 53% of the firings, it can be regarded as superior; if it receives a 1 in fewer than 47%, B can be regarded as superior; if the percentage is between 47% and 53%, neither can be so regarded.
>
> When Allen Wallis was discussing such a problem with (Navy) Captain Garret L. Schuyler, the captain objected that such a test, to quote from Allen's account, may prove wasteful. If a wise and seasoned ordnance officer like Schuyler were on the premises, he would see after the first few thousand or even few hundred [rounds] that the experiment need not be completed either because the new method is obviously inferior or because it is obviously superior beyond what was hoped for ….

Friedman and Wallis worked on the problem for a while but didn't completely solve it.

Realizing that, they told Abraham Wald about the problem.

That set Wald on a path that led him to create *Sequential Analysis* [Wald, 1947].

26.3 Neyman-Pearson formulation

It is useful to begin by describing the theory underlying the test that the U.S. Navy told Captain Garret L. Schuyler to use.

Captain Schuyler's doubts motivated him to tell Milton Friedman and Allen Wallis his conjecture that superior practical procedures existed.

Evidently, the Navy had told Captain Schuyler to use what was then a state-of-the-art Neyman-Pearson hypothesis test.

We'll rely on Abraham Wald's [Wald, 1947] elegant summary of Neyman-Pearson theory.

Watch for these features of the setup:

- the assumption of a *fixed* sample size n
- the application of laws of large numbers, conditioned on alternative probability models, to interpret probabilities α and β of the type I and type II errors defined in the Neyman-Pearson theory

In chapter 1 of **Sequential Analysis** [Wald, 1947] Abraham Wald summarizes the Neyman-Pearson approach to hypothesis testing.

Wald frames the problem as making a decision about a probability distribution that is partially known.

(You have to assume that *something* is already known in order to state a well-posed problem – usually, *something* means *a lot*)

By limiting what is unknown, Wald uses the following simple structure to illustrate the main ideas:

- A decision-maker wants to decide which of two distributions f_0, f_1 govern an IID random variable z.
- The null hypothesis H_0 is the statement that f_0 governs the data.
- The alternative hypothesis H_1 is the statement that f_1 governs the data.
- The problem is to devise and analyze a test of hypothesis H_0 against the alternative hypothesis H_1 on the basis of a sample of a fixed number n independent observations $z_1, z_2, \ldots, z_n$ of the random variable z.

To quote Abraham Wald,

> A test procedure leading to the acceptance or rejection of the [null] hypothesis in question is simply a rule specifying, for each possible sample of size n, whether the [null] hypothesis should be accepted or rejected on the basis of the sample. This may also be expressed as follows: A test procedure is simply a subdivision of the totality of all possible samples of size n into two mutually exclusive parts, say part 1 and part 2, together

> with the application of the rule that the [null] hypothesis be accepted if the observed sample is contained in part 2. Part 1 is also called the critical region. Since part 2 is the totality of all samples of size n which are not included in part 1, part 2 is uniquely determined by part 1. Thus, choosing a test procedure is equivalent to determining a critical region.

Let's listen to Wald longer:

> As a basis for choosing among critical regions the following considerations have been advanced by Neyman and Pearson: In accepting or rejecting H_0 we may commit errors of two kinds. We commit an error of the first kind if we reject H_0 when it is true; we commit an error of the second kind if we accept H_0 when H_1 is true. After a particular critical region W has been chosen, the probability of committing an error of the first kind, as well as the probability of committing an error of the second kind is uniquely determined. The probability of committing an error of the first kind is equal to the probability, determined by the assumption that H_0 is true, that the observed sample will be included in the critical region W. The probability of committing an error of the second kind is equal to the probability, determined on the assumption that H_1 is true, that the observed sample will fall outside the critical region W. For any given critical region W we shall denote the probability of an error of the first kind by α and the probability of an error of the second kind by β.

Let's listen carefully to how Wald applies law of large numbers to interpret α and β:

> The probabilities α and β have the following important practical interpretation: Suppose that we draw a large number of samples of size n. Let M be the number of such samples drawn. Suppose that for each of these M samples we reject H_0 if the sample is included in W and accept H_0 if the sample lies outside W. In this way we make M statements of rejection or acceptance. Some of these statements will in general be wrong. If H_0 is true and if M is large, the probability is nearly 1 (i.e., it is practically certain) that the proportion of wrong statements (i.e., the number of wrong statements divided by M) will be approximately α. If H_1 is true, the probability is nearly 1 that the proportion of wrong statements will be approximately β. Thus, we can say that in the long run [here Wald applies law of large numbers by driving $M \to \infty$ (our comment, not Wald's)] the proportion of wrong statements will be α if H_0 is true and β if H_1 is true.

The quantity α is called the *size* of the critical region, and the quantity $1 - \beta$ is called the *power* of the critical region.

Wald notes that

> one critical region W is more desirable than another if it has smaller values of α and β. Although either α or β can be made arbitrarily small by a proper choice of the critical region W, it is impossible to make both α and β arbitrarily small for a fixed value of n, i.e., a fixed sample size.

Wald summarizes Neyman and Pearson's setup as follows:

> Neyman and Pearson show that a region consisting of all samples $(z_1, z_2, \ldots, z_n)$ which satisfy the inequality
>
> $$\frac{f_1(z_1) \cdots f_1(z_n)}{f_0(z_1) \cdots f_0(z_n)} \geq k$$
>
> is a most powerful critical region for testing the hypothesis H_0 against the alternative hypothesis H_1. The term k on the right side is a constant chosen so that the region will have the required size α.

Wald goes on to discuss Neyman and Pearson's concept of *uniformly most powerful* test.

Our H_0 and H_1 each specify a single distribution, so the most powerful test above is all that we need.

Here is how Wald introduces the notion of a sequential test

> A rule is given for making one of the following three decisions at any stage of the experiment (at the m th trial for each integral value of m): (1) to accept the hypothesis H, (2) to reject the hypothesis H, (3) to continue the experiment by making an additional observation. Thus, such a test procedure is carried out sequentially. On the basis of the first observation, one of the aforementioned decision is made. If the first or second decision is made, the process is terminated. If the third decision is made, a second trial is performed. Again, on the basis of the first two observations, one of the three decision is made. If the third decision

is made, a third trial is performed, and so on. The process is continued until either the first or the second decisions is made. The number n of observations required by such a test procedure is a random variable, since the value of n depends on the outcome of the observations.

26.4 Wald's sequential formulation

By way of contrast to Neyman and Pearson's formulation of the problem, in Wald's formulation

- The sample size n is not fixed but rather a random variable.
- Two parameters A and B that are related to but distinct from Neyman and Pearson's α and β; A and B characterize cut-off rules that Wald uses to determine the random variable n as a function of random outcomes.

Here is how Wald sets up the problem.

A decision-maker can observe a sequence of draws of a random variable z.

He (or she) wants to know which of two probability distributions f_0 or f_1 governs z.

We use beta distributions as examples.

We will also work with the Kullback–Leibler divergence introduced in *Statistical Divergence Measures*.

For two beta densities, Kullback–Leibler divergence has a closed form.

```
@vectorize
def p(x, a, b):
    """Beta distribution density function."""
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x** (a-1) * (1 - x) ** (b-1)

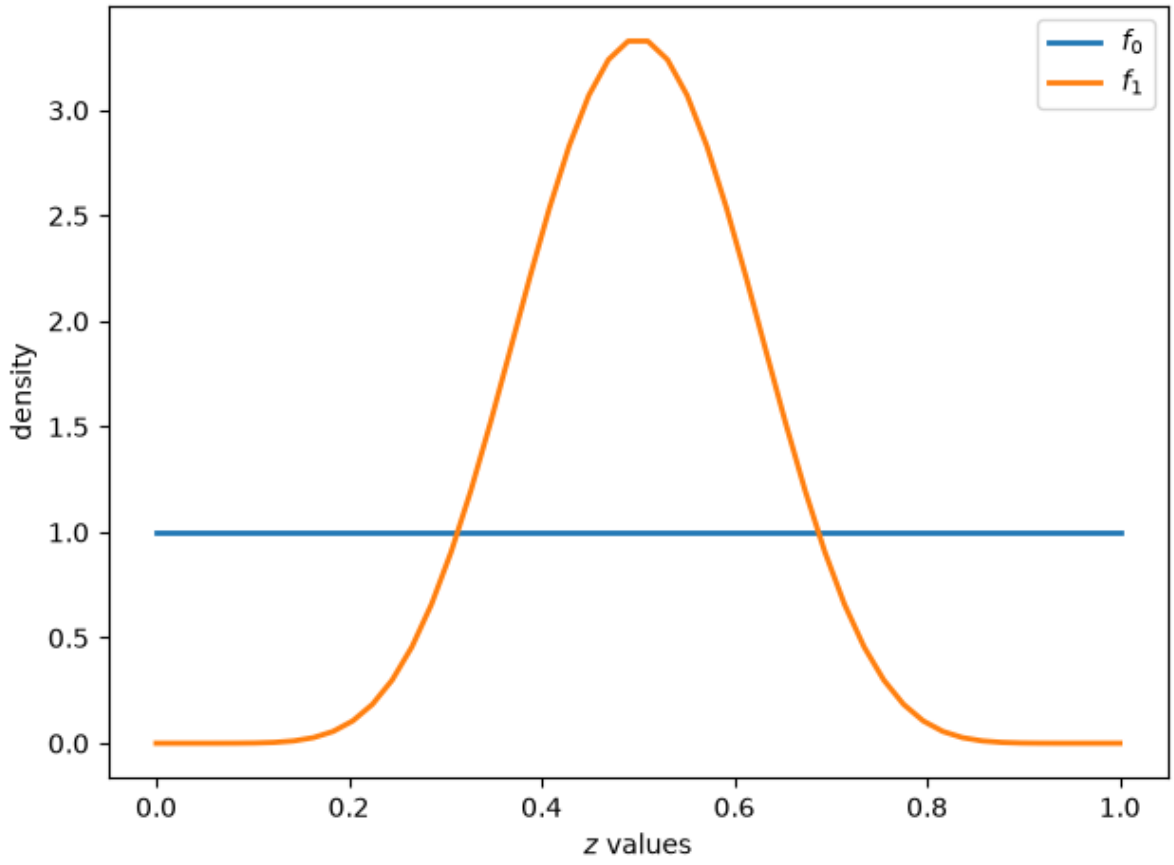
def create_beta_density(a, b):
    """Create a beta density function with specified parameters."""
    return jit(lambda x: p(x, a, b))

def compute_KL(a1, b1, a2, b2):
    """KL divergence KL(f, g) for f = Beta(a1, b1) and g = Beta(a2, b2)."""
    return (betaln(a2, b2) - betaln(a1, b1)
            + (a1 - a2) * digamma(a1) + (b1 - b2) * digamma(b1)
            + (a2 - a1 + b2 - b1) * digamma(a1 + b1))
```

The next figure shows two beta distributions

```
f0 = create_beta_density(1, 1)
f1 = create_beta_density(9, 9)
grid = np.linspace(0, 1, 50)

fig, ax = plt.subplots()
ax.plot(grid, f0(grid), lw=2, label="$f_0$")
ax.plot(grid, f1(grid), lw=2, label="$f_1$")
ax.legend()
ax.set(xlabel="$z$ values", ylabel="density")
plt.tight_layout()
plt.show()
```



Conditional on knowing that successive observations are drawn from distribution f_0, the sequence of random variables is independently and identically distributed (IID).

Conditional on knowing that successive observations are drawn from distribution f_1, the sequence of random variables is also independently and identically distributed (IID).

But the observer does not know which of the two distributions generated the sequence.

For reasons explained in *Exchangeability and Bayesian Updating*, this means that the observer thinks that the sequence is not IID.

Consequently, the observer has something to learn, namely, whether the observations are drawn from f_0 or from f_1.

The decision maker wants to decide which of the two distributions is generating outcomes.

26.4.1 Type I and type II errors

If we regard $f = f_0$ as a null hypothesis and $f = f_1$ as an alternative hypothesis, then

- a type I error is an incorrect rejection of a true null hypothesis (a "false positive")
- a type II error is a failure to reject a false null hypothesis (a "false negative")

To repeat ourselves

- α is the probability of a type I error
- β is the probability of a type II error

26.4.2 Choices

After observing $z_k, z_{k-1}, \ldots, z_1$, the decision-maker chooses among three distinct actions:

- He decides that $f = f_0$ and draws no more z's
- He decides that $f = f_1$ and draws no more z's
- He postpones deciding and instead chooses to draw z_{k+1}

Wald defines

- $p_{0m} = f_0(z_1) \cdots f_0(z_m)$
- $p_{1m} = f_1(z_1) \cdots f_1(z_m)$
- $L_m = \frac{p_{1m}}{p_{0m}}$

for $m \geq 1$ and sets $L_0 = 1$.

Here $\{L_m\}_{m=0}^{\infty}$ is a **likelihood ratio process**.

Wald's sequential decision rule is parameterized by real numbers $0 < B < 1 < A$.

For a given pair A, B, the decision rule is

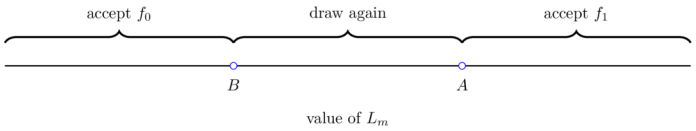
$$\begin{aligned} \text{accept } f = f_1 \text{ if } L_m &\geq A \\ \text{accept } f = f_0 \text{ if } L_m &\leq B \\ \text{draw another } z \text{ if } B < L_m &< A \end{aligned}$$

Note

The sequel *A Bayesian Formulation of Friedman and Wald's Problem* uses A and B for thresholds on the posterior probability π_k that $f = f_1$, and L_0 and L_1 for losses from wrong decisions.

There we write Λ_k for the likelihood ratio and A_Λ, B_Λ for the thresholds on it that correspond to Wald's A and B here.

The following figure illustrates aspects of Wald's procedure.



26.5 Links between A, B and α, β

In chapter 3 of **Sequential Analysis** [Wald, 1947] Wald establishes the inequalities

$$\begin{aligned} \frac{\alpha}{1-\beta} &\leq \frac{1}{A} \\ \frac{\beta}{1-\alpha} &\leq B \end{aligned}$$

These inequalities follow from a change of measure.

Let P_i and E_i denote probabilities and mathematical expectations computed under the assumption that f_i generates the data.

Let N be the random number of draws at which the test stops, and assume that $P_0\{N < \infty\} = P_1\{N < \infty\} = 1$.

On the event $\{N = n\}$, the factor $L_n^{-1} = p_{0n}/p_{1n}$ converts the joint density p_{1n} of $z_1, \ldots, z_n$ into the joint density p_{0n}.

The test rejects H_0 when $L_N \geq A$, so

$$\alpha = P_0\{L_N \geq A\} = E_1\left[L_N^{-1}\mathbf{1}\{L_N \geq A\}\right] \leq \frac{1}{A}P_1\{L_N \geq A\} = \frac{1-\beta}{A}$$

The test accepts H_0 when $L_N \leq B$, so

$$\beta = P_1\{L_N \leq B\} = E_0\left[L_N\mathbf{1}\{L_N \leq B\}\right] \leq B\,P_0\{L_N \leq B\} = B(1-\alpha)$$

The same change of measure underlies importance sampling in *Mean of a Likelihood Ratio Process*.

Both inequalities hold with equality when L_N lands exactly on A or on B.

They are strict when L_N can jump past A or B at the stopping time.

His analysis of these inequalities leads Wald to recommend the following approximations as rules for setting A and B that come close to attaining a decision maker's target values for probabilities α of a type I and β of a type II error:

$$\begin{aligned} A &\approx a(\alpha,\beta) \equiv \frac{1-\beta}{\alpha} \\ B &\approx b(\alpha,\beta) \equiv \frac{\beta}{1-\alpha} \end{aligned} \tag{26.1}$$

For small values of α and β, Wald shows that approximation (26.1) provides a good way to set A and B.

In particular, Wald constructs a mathematical argument that leads him to conclude that the use of approximation (26.1) rather than the true functions $A(\alpha,\beta), B(\alpha,\beta)$ for setting A and B

> ... cannot result in any appreciable increase in the value of either α or β. In other words, for all practical purposes the test corresponding to $A = a(\alpha,\beta), B = b(\alpha,\beta)$ provides at least the same protection against wrong decisions as the test corresponding to $A = A(\alpha,\beta)$ and $B = B(\alpha,\beta)$.
>
> Thus, the only disadvantage that may arise from using $a(\alpha,\beta), b(\alpha,\beta)$ instead of $A(\alpha,\beta), B(\alpha,\beta)$, respectively, is that it may result in an appreciable increase in the number of observations required by the test.

Wald's argument runs as follows.

Let α' and β' be the error probabilities of the test that sets $A = a(\alpha,\beta)$ and $B = b(\alpha,\beta)$.

Applying the inequalities above to that test gives

$$\frac{\alpha'}{1-\beta'} \leq \frac{\alpha}{1-\beta}, \quad \frac{\beta'}{1-\alpha'} \leq \frac{\beta}{1-\alpha}$$

These imply

$$\alpha' \leq \frac{\alpha}{1-\beta}, \quad \beta' \leq \frac{\beta}{1-\alpha}, \quad \alpha'+\beta' \leq \alpha+\beta \tag{26.2}$$

The third inequality comes from cross-multiplying the two inequalities above and adding them.

When α and β are small, $\alpha/(1-\beta)$ and $\beta/(1-\alpha)$ are close to α and β.

The test attains $\alpha' = \alpha$ and $\beta' = \beta$ only when L_N never jumps past A or B.

We'll write some Python code to help us illustrate Wald's claims about how α and β are related to the parameters A and B that characterize his sequential probability ratio test.

26.6 Simulations

We experiment with different distributions f_0 and f_1 to examine how Wald's test performs under various conditions.

Our goal in conducting these simulations is to understand trade-offs between decision speed and accuracy associated with Wald's **sequential probability ratio test**.

Specifically, we will watch how:

- The decision thresholds A and B (or equivalently the target error rates α and β) affect the average stopping time
- The discrepancy between distributions f_0 and f_1 affects average stopping times

We will focus on the case where f_0 and f_1 are beta distributions since it is easy to control the overlapping regions of the two densities by adjusting their shape parameters.

First, we define a namedtuple to store all the parameters we need for our simulation studies.

We also compute Wald's recommended thresholds A and B based on the target type I and type II errors α and β

```
SPRTParams = namedtuple('SPRTParams',
                ['α', 'β',  # Target type I and type II errors
                'a0', 'b0', # Shape parameters for f_0
                'a1', 'b1', # Shape parameters for f_1
                'N',        # Number of simulations
                'seed'])

@njit
def compute_wald_thresholds(α, β):
    """Compute Wald's recommended thresholds."""
    A = (1 - β) / α
    B = β / (1 - α)
    return A, B, np.log(A), np.log(B)
```

Now we can run the simulation following Wald's recommendation.

We'll compare the log-likelihood ratio to logarithms of the thresholds $\log(A)$ and $\log(B)$.

The following algorithm underlies our simulations.

1. Compute thresholds $A = \frac{1-\beta}{\alpha}$, $B = \frac{\beta}{1-\alpha}$ and work with $\log A$, $\log B$.
2. Given true distribution (either f_0 or f_1):
    - Initialize log-likelihood ratio $\log L_0 = 0$
    - Repeat:
        - Draw observation z from the true distribution
        - Update: $\log L_{n+1} \leftarrow \log L_n + (\log f_1(z) - \log f_0(z))$
        - If $\log L_{n+1} \geq \log A$: stop, reject H_0
        - If $\log L_{n+1} \leq \log B$: stop, accept H_0
3. Repeat step 2 for N replications with $N/2$ replications for each distribution, compute the empirical type I error $\hat{\alpha}$ and type II error $\hat{\beta}$ with

$$\hat{\alpha} = \frac{\#\text{ of times reject } H_0 \text{ when } f_0 \text{ is true}}{\#\text{ of replications with } f_0 \text{ true}}$$

$$\hat{\beta} = \frac{\#\text{ of times accept } H_0 \text{ when } f_1 \text{ is true}}{\#\text{ of replications with } f_1 \text{ true}}$$

```
@njit
def sprt_single_run(a0, b0, a1, b1, logA, logB, true_f0, seed):
    """Run a single SPRT until a decision is reached."""
    log_L = 0.0
    n = 0
    np.random.seed(seed)

    while True:
        z = np.random.beta(a0, b0) if true_f0 else np.random.beta(a1, b1)
        n += 1

        # Update log-likelihood ratio
        log_L += np.log(p(z, a1, b1)) - np.log(p(z, a0, b0))

        # Check stopping conditions
        if log_L >= logA:
            return n, False  # Reject H0
        elif log_L <= logB:
            return n, True   # Accept H0

@njit(parallel=True)
def run_sprt_simulation(a0, b0, a1, b1, α, β, N, seed):
    """SPRT simulation."""
    A, B, logA, logB = compute_wald_thresholds(α, β)

    stopping_times = np.zeros(N, dtype=np.int64)
    decisions_h0 = np.zeros(N, dtype=np.bool_)
    truth_h0 = np.zeros(N, dtype=np.bool_)

    for i in prange(N):
        true_f0 = (i % 2 == 0)
        truth_h0[i] = true_f0

        n, accept_f0 = sprt_single_run(
                        a0, b0, a1, b1,
                        logA, logB,
                        true_f0, seed + i)
        stopping_times[i] = n
        decisions_h0[i] = accept_f0

    return stopping_times, decisions_h0, truth_h0

def run_sprt(params):
    """Run SPRT simulations with given parameters."""
    stopping_times, decisions_h0, truth_h0 = run_sprt_simulation(
        params.a0, params.b0, params.a1, params.b1,
        params.α, params.β, params.N, params.seed
    )

    # Calculate error rates
    truth_h0_bool = truth_h0.astype(bool)
    decisions_h0_bool = decisions_h0.astype(bool)

    type_I = np.sum(truth_h0_bool & ~decisions_h0_bool) \
            / np.sum(truth_h0_bool)
    type_II = np.sum(~truth_h0_bool & decisions_h0_bool) \
            / np.sum(~truth_h0_bool)
```

(continues on next page)

(continued from previous page)

```
    return {
        'stopping_times': stopping_times,
        'decisions_h0': decisions_h0_bool,
        'truth_h0': truth_h0_bool,
        'type_I': type_I,
        'type_II': type_II
    }

# Run simulation
params = SPRTParams(α=0.05, β=0.10, a0=2, b0=5, a1=5, b1=2, N=20000, seed=1)
results = run_sprt(params)

print(f"Average stopping time: {results['stopping_times'].mean():.2f}")
print(f"Empirical type I  error: {results['type_I']:.3f} (target = {params.α})")
print(f"Empirical type II error: {results['type_II']:.3f} (target = {params.β})")
```

```
Average stopping time: 1.59
Empirical type I  error: 0.012 (target = 0.05)
Empirical type II error: 0.022 (target = 0.1)
```

As anticipated in the passage above in which Wald discussed the quality of $a(\alpha, \beta), b(\alpha, \beta)$ given in approximation (26.1), we find that the algorithm actually gives **lower** type I and type II error rates than the target values.

The empirical error rates satisfy Wald's bounds (26.2), which here are $\alpha' \leq 0.05/0.90 \approx 0.056$, $\beta' \leq 0.10/0.95 \approx 0.105$, and $\alpha' + \beta' \leq 0.15$.

The bounds are far from tight because $f_0 = \text{Beta}(2, 5)$ and $f_1 = \text{Beta}(5, 2)$ are far apart.

A single draw often carries $\log L_m$ far past $\log A$ or $\log B$, and that overshoot pushes α' and β' below α and β.

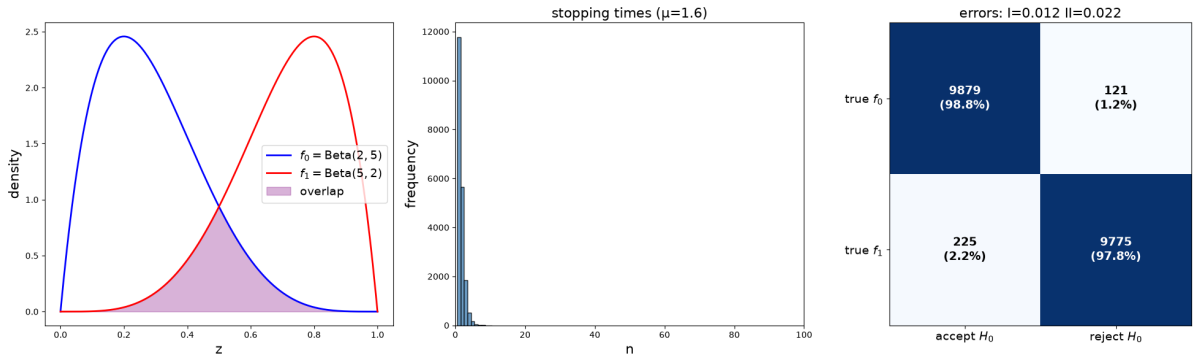
Note

For recent work on the quality of approximation (26.1), see, e.g., [Fischer and Ramdas, 2024].

The following code creates a few graphs that illustrate the results of our simulation.

Let's plot the results of our simulation

```
plot_sprt_results(results, params)
```



In this example, the stopping time never exceeds 10, and the mean stopping time is 1.59.

We can construct a 2×2 "confusion matrix" whose diagonal elements count the number of times that Wald's decision rule correctly accepts and rejects the null hypothesis.

```
print("Confusion Matrix data:")
print(f"Type I error: {results['type_I']:.3f}")
print(f"Type II error: {results['type_II']:.3f}")
```

```
Confusion Matrix data:
Type I error: 0.012
Type II error: 0.022
```

Next we use our code to study three different f_0, f_1 pairs having different discrepancies between distributions.

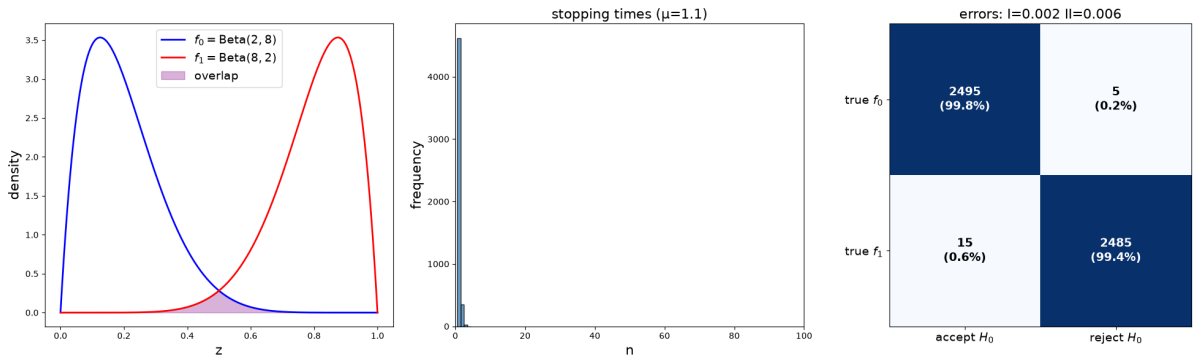
We plot the same three graphs we used above for each pair of distributions

```
params_1 = SPRTParams(α=0.05, β=0.10, a0=2, b0=8, a1=8, b1=2, N=5000, seed=42)
results_1 = run_sprt(params_1)

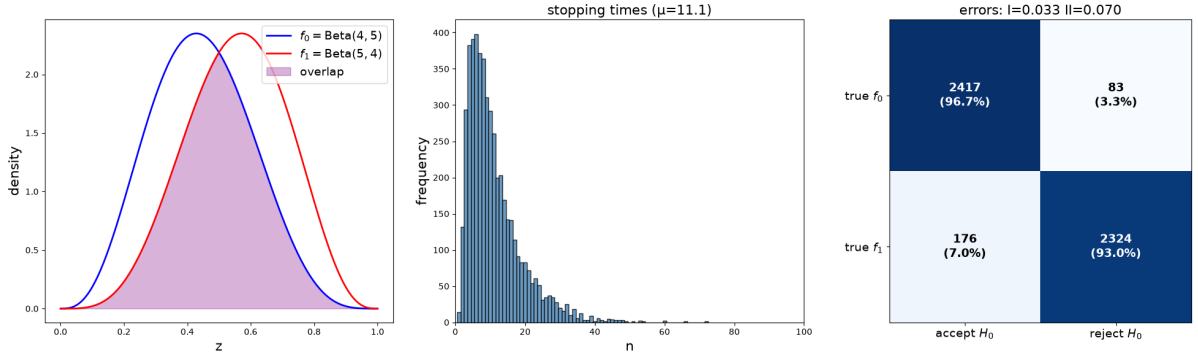
params_2 = SPRTParams(α=0.05, β=0.10, a0=4, b0=5, a1=5, b1=4, N=5000, seed=42)
results_2 = run_sprt(params_2)

params_3 = SPRTParams(α=0.05, β=0.10, a0=0.5, b0=0.4, a1=0.4,
                      b1=0.5, N=5000, seed=42)
results_3 = run_sprt(params_3)
```

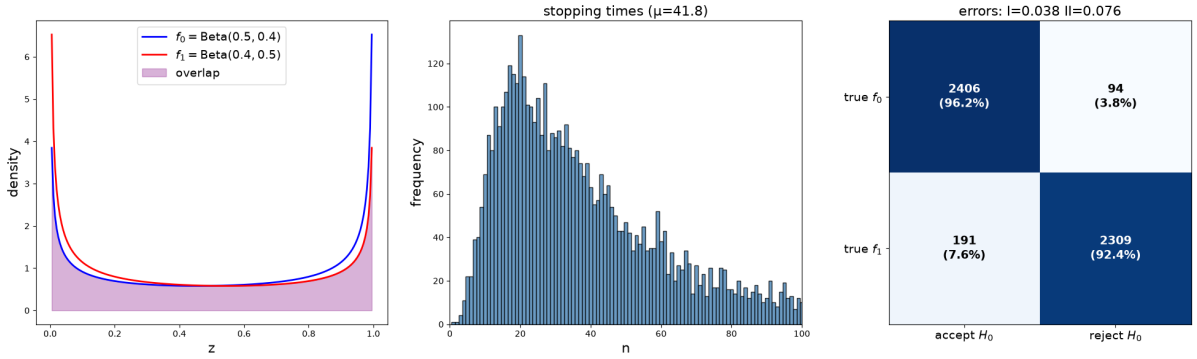
```
plot_sprt_results(results_1, params_1)
```




```
plot_sprt_results(results_2, params_2)
```




```
plot_sprt_results(results_3, params_3)
```



Notice that the stopping times are less when the two distributions are farther apart.

This makes sense.

When two distributions are "far apart", it should not take too long to decide which one is generating the data.

When two distributions are "close", it should take longer to decide which one is generating the data.

The *Kullback–Leibler divergence* described in *Statistical Divergence Measures* measures how far apart two distributions are.

Wald linked it to the expected number of draws that his test requires.

Under H_0, $\log L_N = \sum_{i=1}^{N} \log \frac{f_1(z_i)}{f_0(z_i)}$ is a sum of a random number N of IID terms, each with mean $E_0 \log \frac{f_1(z)}{f_0(z)} = -KL(f_0, f_1)$.

Under H_1, each term has mean $E_1 \log \frac{f_1(z)}{f_0(z)} = KL(f_1, f_0)$.

Because N is a stopping time with finite mean, **Wald's identity** gives

$$E_0[\log L_N] = -E_0[N]\, KL(f_0, f_1), \qquad E_1[\log L_N] = E_1[N]\, KL(f_1, f_0) \tag{26.3}$$

This section of *Likelihood Ratio Processes* describes the link between KL divergence and the expected log likelihood ratio for a fixed number of draws.

Wald ignored overshoot: he approximated $\log L_N$ by $\log A$ when the test rejects H_0 and by $\log B$ when it accepts H_0.

Setting $P_0\{L_N \geq A\} \approx \alpha$ and $P_1\{L_N \leq B\} \approx \beta$ in (26.3) then gives Wald's approximations to the **average sample numbers**

$$E_0[N] \approx \frac{(1-\alpha)\log(1/B) - \alpha \log A}{KL(f_0, f_1)}, \qquad E_1[N] \approx \frac{(1-\beta)\log A - \beta \log(1/B)}{KL(f_1, f_0)} \tag{26.4}$$

KL divergence is not symmetric, so (26.4) divides by $KL(f_0, f_1)$ under H_0 and by $KL(f_1, f_0)$ under H_1.

In the experiments below, $f_0 = \text{Beta}(a, b)$ and $f_1 = \text{Beta}(b, a)$, so $f_1(z) = f_0(1-z)$ and $KL(f_0, f_1) = KL(f_1, f_0)$.

Half of our replications draw from f_0 and half from f_1, so (26.4) makes the mean stopping time proportional to $1/KL(f_0, f_1)$.

We compute $KL(f_0, f_1)$ and the mean stopping time for 100 such pairs and compare the mean stopping time with (26.4).

```
def generate_β_pairs(N=100, T=10.0, d_min=0.5, d_max=9.5):
    ds = np.linspace(d_min, d_max, N)
    a0 = (T - ds) / 2
    b0 = (T + ds) / 2
    return list(zip(a0, b0, b0, a0))

def wald_asn(α, β, KL_01, KL_10):
```

(continues on next page)

(continued from previous page)

```
    """Wald's approximations to E_0[N] and E_1[N]."""
    A, B, logA, logB = compute_wald_thresholds(α, β)
    E0_N = ((1 - α) * (-logB) - α * logA) / KL_01
    E1_N = ((1 - β) * logA - β * (-logB)) / KL_10
    return E0_N, E1_N

param_comb = generate_β_pairs()

KL_divs = []
mean_stopping_times = []
asn_approx = []
param_list = []

for a0, b0, a1, b1 in param_comb:
    KL_01 = compute_KL(a0, b0, a1, b1)
    KL_10 = compute_KL(a1, b1, a0, b0)

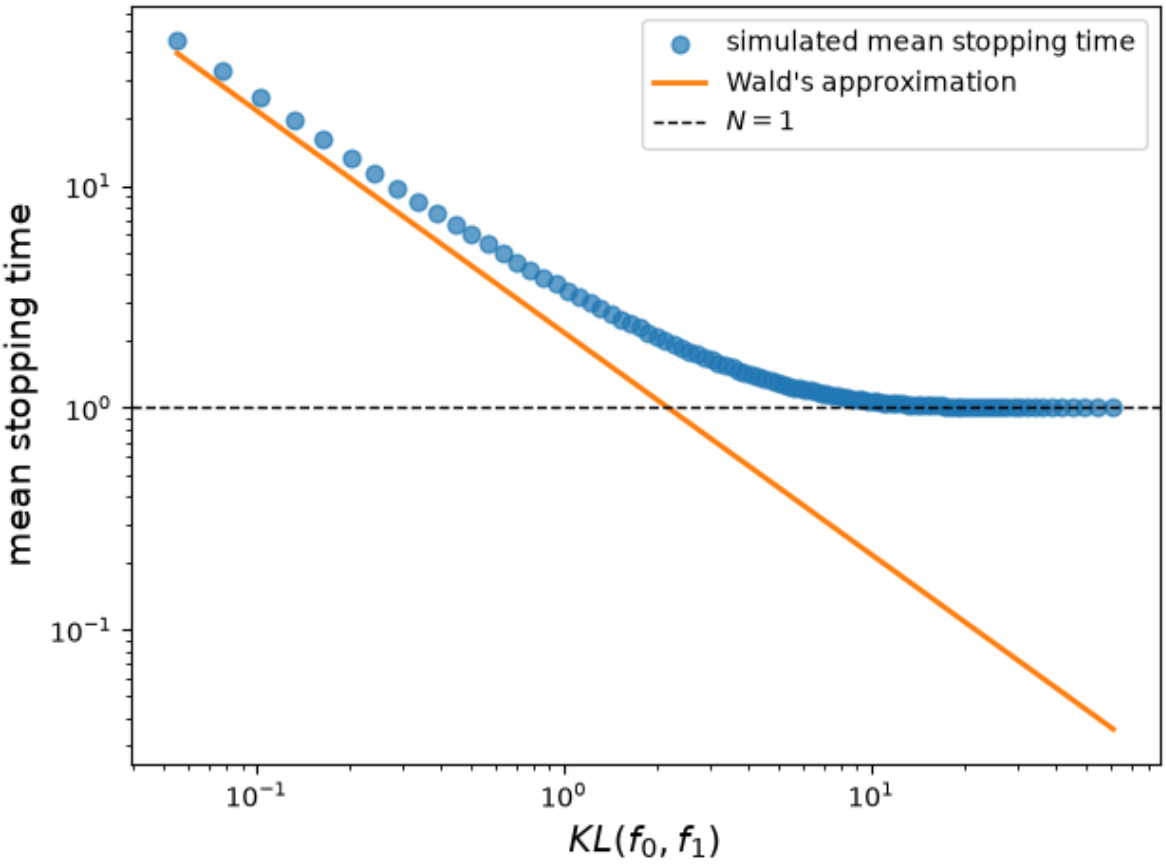
    params_d = SPRTParams(α=0.05, β=0.10, a0=a0, b0=b0,
                          a1=a1, b1=b1, N=5000, seed=42)
    results_d = run_sprt(params_d)
    E0_N, E1_N = wald_asn(params_d.α, params_d.β, KL_01, KL_10)

    KL_divs.append(KL_01)
    mean_stopping_times.append(results_d['stopping_times'].mean())
    asn_approx.append(0.5 * (E0_N + E1_N))
    param_list.append((a0, b0, a1, b1))

KL_divs = np.array(KL_divs)
mean_stopping_times = np.array(mean_stopping_times)
asn_approx = np.array(asn_approx)

fig, ax = plt.subplots()
ax.scatter(KL_divs, mean_stopping_times, s=40, alpha=0.7,
           label='simulated mean stopping time')
ax.plot(KL_divs, asn_approx, color='C1', lw=2,
        label="Wald's approximation")
ax.axhline(1, color='black', linestyle='--', lw=1, label='$N = 1$')
ax.set_xscale('log')
ax.set_yscale('log')
ax.set_xlabel('$KL(f_0, f_1)$', fontsize=14)
ax.set_ylabel('mean stopping time', fontsize=14)
ax.legend()
plt.tight_layout()
plt.show()

ratio = mean_stopping_times / asn_approx
print(f"corr(mean stopping time, 1/KL) = "
      f"{np.corrcoef(mean_stopping_times, 1 / KL_divs)[0, 1]:.4f}")
print(f"simulated / approximate mean stopping time: "
      f"{ratio.min():.2f} (closest pair) to {ratio.max():.1f} (farthest pair)")
```




```
corr(mean stopping time, 1/KL) = 0.9999
simulated / approximate mean stopping time: 1.14 (closest pair) to 28.1 (farthest␣
↪pair)
```

On log scales, Wald's approximation is a straight line with slope -1.

Across the 100 pairs, the correlation between the mean stopping time and $1/KL(f_0, f_1)$ is 0.9999.

The mean stopping time is close to a linear function of $1/KL(f_0, f_1)$.

The simulated mean stopping times lie above Wald's approximation for two reasons.

Overshoot makes $|\log L_N|$ larger than $\log A$ or $\log(1/B)$, so by (26.3) it makes $E_0[N]$ and $E_1[N]$ larger than (26.4).

The test always takes at least one draw, so the mean stopping time cannot fall below 1, while (26.4) goes to zero as $KL(f_0, f_1)$ grows.

Exercise 26.8.4 asks you to study these two sources of discrepancy.

Below are sampled examples from the experiments we have above

```
def plot_beta_distributions_grid(param_list, KL_divs, mean_stopping_times,
                                 selected_indices=None):
    """Plot grid of beta distributions with KL divergence and stopping times."""
    if selected_indices is None:
        selected_indices = [0, len(param_list)//6, len(param_list)//3,
                            len(param_list)//2, 2*len(param_list)//3, -1]

    fig, axes = plt.subplots(2, 3, figsize=(15, 8))
```

(continues on next page)

(continued from previous page)

```
    z_grid = np.linspace(0, 1, 200)

    for i, idx in enumerate(selected_indices):
        row, col = i // 3, i % 3
        a0, b0, a1, b1 = param_list[idx]

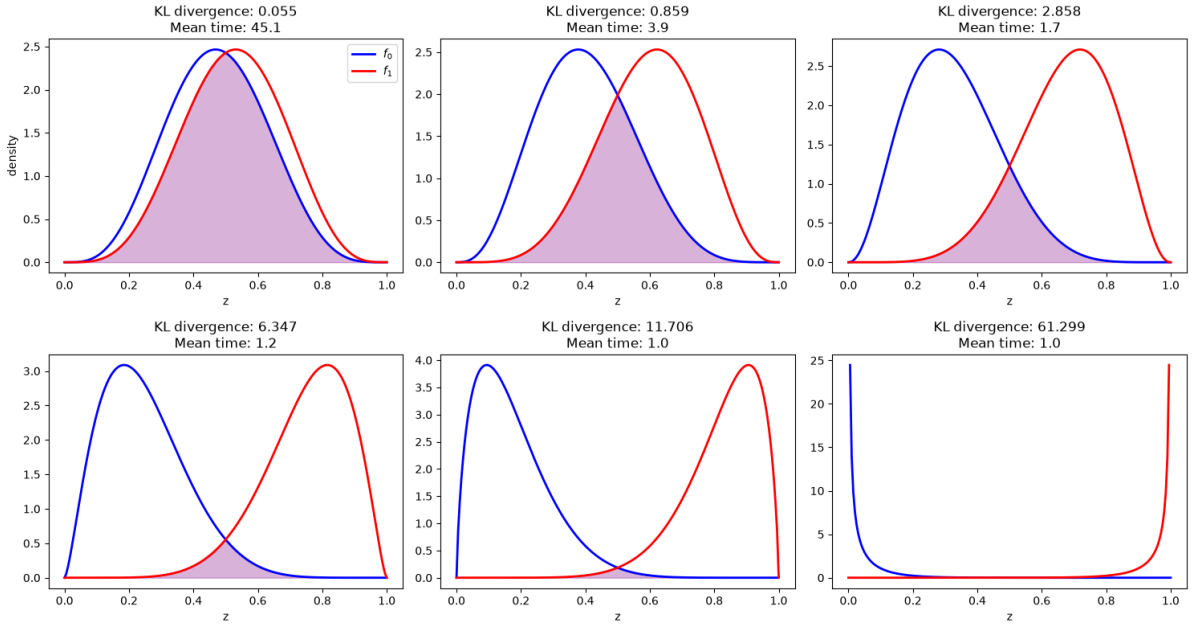
        f0 = create_beta_density(a0, b0)
        f1 = create_beta_density(a1, b1)

        axes[row, col].plot(z_grid, f0(z_grid), 'b-', lw=2, label='$f_0$')
        axes[row, col].plot(z_grid, f1(z_grid), 'r-', lw=2, label='$f_1$')
        axes[row, col].fill_between(z_grid, 0,
                                    np.minimum(f0(z_grid), f1(z_grid)),
                                    alpha=0.3, color='purple')

        axes[row, col].set_title(f'KL divergence: {KL_divs[idx]:.3f}'
                                 f'\nMean time: {mean_stopping_times[idx]:.1f}',
                                 fontsize=12)
        axes[row, col].set_xlabel('z', fontsize=10)
        if i == 0:
            axes[row, col].set_ylabel('density', fontsize=10)
            axes[row, col].legend(fontsize=10)

    plt.tight_layout()
    plt.show()

plot_beta_distributions_grid(param_list, KL_divs, mean_stopping_times)
```



Again, we find that the stopping time is shorter when the distributions are more separated, as measured by KL divergence.

Let's visualize individual likelihood ratio processes to see how they evolve toward the decision boundaries.

```
def plot_likelihood_paths(params, n_highlight=10, n_background=200):
    """visualize likelihood ratio paths."""
    A, B, logA, logB = compute_wald_thresholds(params.α, params.β)
```

(continues on next page)

(continued from previous page)

```
    f0, f1 = map(lambda ab: create_beta_density(*ab),
                 [(params.a0, params.b0),
                  (params.a1, params.b1)])

    fig, axes = plt.subplots(1, 2, figsize=(14, 7))

    for dist_idx, (true_f0, ax, title) in enumerate([
        (True, axes[0], 'true distribution: $f_0$'),
        (False, axes[1], 'true distribution: $f_1$')
    ]):
        rng = np.random.default_rng(seed=42 + dist_idx)
        paths_data = []

        # Generate paths
        for path in range(n_background + n_highlight):
            log_L_path, log_L, n = [0.0], 0.0, 0

            while True:
                z = rng.beta(params.a0, params.b0) if true_f0 \
                    else rng.beta(params.a1, params.b1)
                n += 1
                log_L += np.log(f1(z)) - np.log(f0(z))
                log_L_path.append(log_L)

                if log_L >= logA or log_L <= logB:
                    paths_data.append((log_L_path, n, log_L >= logA))
                    break

        # Plot background paths
        for path, _, decision in paths_data[:n_background]:
            ax.plot(range(len(path)), path, color='C1' if decision else 'C0',
                    alpha=0.2, linewidth=0.5)

        # Plot highlighted paths with labels
        for i, (path, _, decision) in enumerate(paths_data[n_background:]):
            ax.plot(range(len(path)), path, color='C1' if decision else 'C0',
                    alpha=0.8, linewidth=1.5,
                    label='reject $H_0$' if decision and i == 0 else (
                          'accept $H_0$' if not decision and i == 0 else ''))

        # Add threshold lines and formatting
        ax.axhline(y=logA, color='C1', linestyle='--', linewidth=2,
                   label=f'$\\log A = {logA:.2f}$')
        ax.axhline(y=logB, color='C0', linestyle='--', linewidth=2,
                   label=f'$\\log B = {logB:.2f}$')
        ax.axhline(y=0, color='black', linestyle='-', alpha=0.5, linewidth=1)

        ax.set_xlabel(r'$n$')
        ax.set_ylabel(r'$\log(L_n)$')
        ax.set_title(title, fontsize=20)
        ax.legend(fontsize=18, loc='center right')

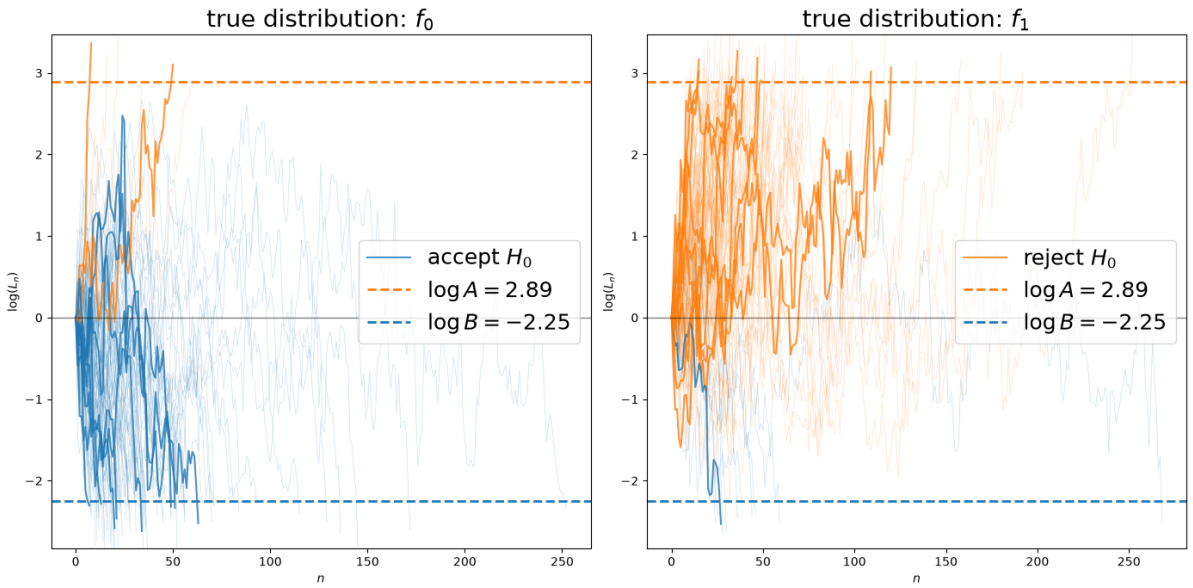
        y_margin = max(abs(logA), abs(logB)) * 0.2
        ax.set_ylim(logB - y_margin, logA + y_margin)

    plt.tight_layout()
    plt.show()
```

(continues on next page)

(continued from previous page)

```
plot_likelihood_paths(params_3, n_highlight=10, n_background=100)
```



Next, let's adjust the decision thresholds A and B and examine how the mean stopping time and the type I and type II error rates change.

In the code below, we adjust Wald's rule by adjusting the thresholds A and B using factors A_f and B_f.

```
@njit(parallel=True)
def run_adjusted_thresholds(a0, b0, a1, b1, α, β, N, seed, A_f, B_f):
    """SPRT simulation with adjusted thresholds."""

    # Calculate original thresholds
    A_original = (1 - β) / α
    B_original = β / (1 - α)

    # Apply adjustment factors
    A_adj = A_original * A_f
    B_adj = B_original * B_f
    logA = np.log(A_adj)
    logB = np.log(B_adj)

    # Pre-allocate arrays
    stopping_times = np.zeros(N, dtype=np.int64)
    decisions_h0 = np.zeros(N, dtype=np.bool_)
    truth_h0 = np.zeros(N, dtype=np.bool_)

    # Run simulations in parallel
    for i in prange(N):
        true_f0 = (i % 2 == 0)
        truth_h0[i] = true_f0

        n, accept_f0 = sprt_single_run(a0, b0, a1, b1,
                        logA, logB, true_f0, seed + i)
        stopping_times[i] = n
        decisions_h0[i] = accept_f0
```

(continues on next page)

(continued from previous page)

```
    return stopping_times, decisions_h0, truth_h0, A_adj, B_adj

def run_adjusted(params, A_f=1.0, B_f=1.0):
    """Wrapper to run SPRT with adjusted A and B thresholds."""

    stopping_times, decisions_h0, truth_h0, A_adj, B_adj = run_adjusted_thresholds(
        params.a0, params.b0, params.a1, params.b1,
        params.α, params.β, params.N, params.seed, A_f, B_f
    )
    truth_h0_bool = truth_h0.astype(bool)
    decisions_h0_bool = decisions_h0.astype(bool)

    # Calculate error rates
    type_I = np.sum(truth_h0_bool
                    & ~decisions_h0_bool) / np.sum(truth_h0_bool)
    type_II = np.sum(~truth_h0_bool
                    & decisions_h0_bool) / np.sum(~truth_h0_bool)

    return {
        'stopping_times': stopping_times,
        'type_I': type_I,
        'type_II': type_II,
        'A_used': A_adj,
        'B_used': B_adj
    }

adjustments = [
    (5.0, 0.5),
    (1.0, 1.0),
    (0.3, 3.0),
    (0.2, 5.0),
    (0.15, 7.0),
]

results_table = []
for A_f, B_f in adjustments:
    result = run_adjusted(params_2, A_f, B_f)
    results_table.append([
        A_f, B_f,
        f"{result['stopping_times'].mean():.1f}",
        f"{result['type_I']:.3f}",
        f"{result['type_II']:.3f}"
    ])

df = pd.DataFrame(results_table,
                  columns=["A_f", "B_f", "mean stop time",
                           "Type I error", "Type II error"])
df = df.set_index(["A_f", "B_f"])
df
```

A_f	B_f	mean stop time	Type I error	Type II error
5.00	0.5	16.1	0.006	0.036
1.00	1.0	11.1	0.033	0.070
0.30	3.0	5.5	0.086	0.195

(continues on next page)

(continued from previous page)

```
0.20 5.0            3.4         0.120         0.304
0.15 7.0            2.2         0.146         0.410
```

Let's pause and think about the table more carefully by referring back to (26.1).

Recall that $A = \frac{1-\beta}{\alpha}$ and $B = \frac{\beta}{1-\alpha}$.

When we multiply A by a factor less than 1 (making A smaller), we are effectively making it easier to reject the null hypothesis H_0.

This increases the probability of Type I errors.

When we multiply B by a factor greater than 1 (making B larger), we are making it easier to accept the null hypothesis H_0.

This increases the probability of Type II errors.

The table confirms this intuition: as A decreases and B increases from Wald's approximate values (26.1), both Type I and Type II error rates increase, while the mean stopping time decreases.

26.7 Related lectures

We'll dig deeper into some of the ideas used here in the following earlier and later lectures:

- In *this sequel*, we reformulate the problem from the perspective of a **Bayesian statistician** who views parameters as vectors of random variables that are jointly distributed with the observables they are concerned about.
- The concept of **exchangeability**, which underlies much of statistical learning, is explored in depth in our *lecture on exchangeable random variables*.
- For a deeper understanding of likelihood ratio processes and their role in frequentist and Bayesian statistical theories, see *Likelihood Ratio Processes*.
- Building on that foundation, *Likelihood Ratio Processes and Bayesian Learning* examines the role of likelihood ratio processes in **Bayesian learning**.
- Finally, *this later lecture* revisits the subject discussed here and examines whether the frequentist decision rule that the Navy ordered the captain to use would perform better or worse than Abraham Wald's sequential decision rule.

26.8 Exercises

In the first two exercises below, please try to rewrite the entire SPRT suite in this lecture.

Exercise 26.8.1

In the first exercise, we apply the sequential probability ratio test to distinguish two models generated by 3-state Markov chains

(For a review on likelihood ratio processes for Markov chains, see *this section*.)

Consider distinguishing between two 3-state Markov chain models using Wald's sequential probability ratio test.

You have competing hypotheses about the transition probabilities:

- H_0: The chain follows transition matrix $P^{(0)}$

- H_1: The chain follows transition matrix $P^{(1)}$

Given transition matrices:

$$P^{(0)} = \begin{bmatrix} 0.7 & 0.2 & 0.1 \\ 0.3 & 0.5 & 0.2 \\ 0.1 & 0.3 & 0.6 \end{bmatrix}, \quad P^{(1)} = \begin{bmatrix} 0.5 & 0.3 & 0.2 \\ 0.2 & 0.6 & 0.2 \\ 0.2 & 0.2 & 0.6 \end{bmatrix}$$

For a sequence of observations $(x_0, x_1, \ldots, x_t)$, the likelihood ratio is:

$$\Lambda_t = \frac{\pi^{(1)}_{x_0}}{\pi^{(0)}_{x_0}} \prod_{s=1}^{t} \frac{P^{(1)}_{x_{s-1},x_s}}{P^{(0)}_{x_{s-1},x_s}}$$

where $\pi^{(i)}$ is the stationary distribution under hypothesis i.

Tasks:

1. Implement the likelihood ratio computation for Markov chains
2. Implement Wald's sequential test with Type I error $\alpha = 0.05$ and Type II error $\beta = 0.10$
3. Run 1000 simulations under each hypothesis and compute empirical error rates
4. Analyze the distribution of stopping times

The test stops when:

- $\Lambda_t \geq A = \frac{1-\beta}{\alpha} = 18$: Reject H_0
- $\Lambda_t \leq B = \frac{\beta}{1-\alpha} = 0.105$: Accept H_0

Solution

Below is one solution to the exercise.

In the lecture, we write the code more verbosely to illustrate the concepts clearly.

In the code below, we simplified some of the code structure for a shorter presentation.

First we define the parameters for the Markov chain SPRT

```
MarkovSPRTParams = namedtuple('MarkovSPRTParams',
            ['α', 'β', 'P_0', 'P_1', 'N', 'seed'])

def compute_stationary_distribution(P):
    """Compute stationary distribution of transition matrix P."""
    eigenvalues, eigenvectors = np.linalg.eig(P.T)
    idx = np.argmin(np.abs(eigenvalues - 1))
    pi = np.real(eigenvectors[:, idx])
    return pi / pi.sum()

@njit
def simulate_markov_chain(P, pi_0, T, seed):
    """Simulate a Markov chain path."""
    np.random.seed(seed)
    path = np.zeros(T, dtype=np.int32)

    cumsum_pi = np.cumsum(pi_0)
    path[0] = np.searchsorted(cumsum_pi, np.random.uniform())

    for t in range(1, T):
        cumsum_row = np.cumsum(P[path[t-1]])
        path[t] = np.searchsorted(cumsum_row, np.random.uniform())

    return path
```

Here we define the function that runs SPRT for Markov chains

```
@njit
def markov_sprt_single_run(P_0, P_1, π_0, π_1,
                           logA, logB, true_P, true_π, seed):
    """Run single SPRT for Markov chains."""
    max_n = 10000
    path = simulate_markov_chain(true_P, true_π, max_n, seed)

    log_L = np.log(π_1[path[0]] / π_0[path[0]])
    if log_L >= logA: return 1, False
    if log_L <= logB: return 1, True

    for t in range(1, max_n):
        prev_state, curr_state = path[t-1], path[t]
        p_1, p_0 = P_1[prev_state, curr_state], P_0[prev_state, curr_state]

        if p_0 > 0:
            log_L += np.log(p_1 / p_0)
        elif p_1 > 0:
            log_L = np.inf

        if log_L >= logA: return t+1, False
        if log_L <= logB: return t+1, True

    return max_n, log_L < 0

def run_markov_sprt(params):
    """Run SPRT for Markov chains."""
    π_0 = compute_stationary_distribution(params.P_0)
    π_1 = compute_stationary_distribution(params.P_1)
    A, B, logA, logB = compute_wald_thresholds(params.α, params.β)

    stopping_times = np.zeros(params.N, dtype=np.int64)
    decisions_h0 = np.zeros(params.N, dtype=bool)
    truth_h0 = np.zeros(params.N, dtype=bool)

    for i in range(params.N):
        true_P, true_π = (params.P_0, π_0) if i % 2 == 0 else (params.P_1, π_1)
        truth_h0[i] = i % 2 == 0

        n, accept_h0 = markov_sprt_single_run(
            params.P_0, params.P_1, π_0, π_1, logA, logB,
            true_P, true_π, params.seed + i)

        stopping_times[i] = n
        decisions_h0[i] = accept_h0

    type_I = np.sum(truth_h0 & ~decisions_h0) / np.sum(truth_h0)
    type_II = np.sum(~truth_h0 & decisions_h0) / np.sum(~truth_h0)

    return {
        'stopping_times': stopping_times, 'decisions_h0': decisions_h0,
        'truth_h0': truth_h0, 'type_I': type_I, 'type_II': type_II
    }
```

Now we can run the SPRT for the Markov chain models and visualize the results

```
# Run Markov chain SPRT
P_0 = np.array([[0.7, 0.2, 0.1],
```

```
                [0.3, 0.5, 0.2],
                [0.1, 0.3, 0.6]])

P_1 = np.array([[0.5, 0.3, 0.2],
                [0.2, 0.6, 0.2],
                [0.2, 0.2, 0.6]])

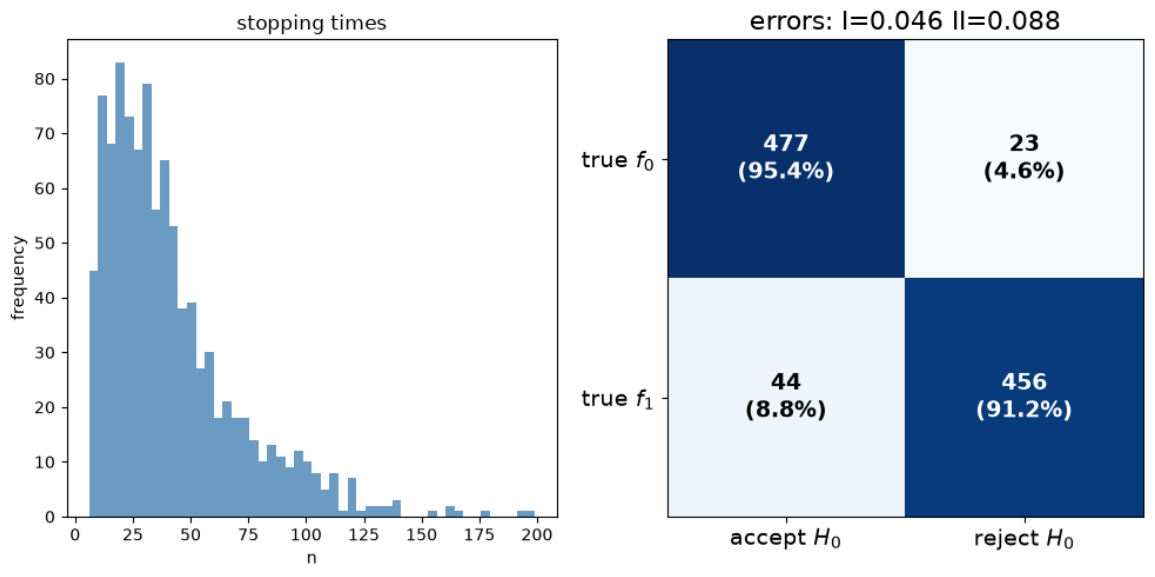
params_markov = MarkovSPRTParams(α=0.05, β=0.10,
                         P_0=P_0, P_1=P_1, N=1000, seed=42)
results_markov = run_markov_sprt(params_markov)


fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(10, 5))

ax1.hist(results_markov['stopping_times'],
            bins=50, color="steelblue", alpha=0.8)
ax1.set_title("stopping times")
ax1.set_xlabel("n")
ax1.set_ylabel("frequency")

plot_confusion_matrix(results_markov, ax2)

plt.tight_layout()
plt.show()
```



Exercise 26.8.2

In this exercise, apply Wald's sequential test to distinguish between two VAR(1) models with different dynamics and noise structures.

For a review of the likelihood ratio process with VAR models, see *Likelihood Processes For VAR Models*.

Given VAR models under each hypothesis:

- H_0: $x_{t+1} = A^{(0)} x_t + C^{(0)} w_{t+1}$

- H_1: $x_{t+1} = A^{(1)} x_t + C^{(1)} w_{t+1}$

where $w_t \sim \mathcal{N}(0, I)$ and:

$$A^{(0)} = \begin{bmatrix} 0.8 & 0.1 \\ 0.2 & 0.7 \end{bmatrix}, \quad C^{(0)} = \begin{bmatrix} 0.3 & 0.1 \\ 0.1 & 0.3 \end{bmatrix}$$

$$A^{(1)} = \begin{bmatrix} 0.6 & 0.2 \\ 0.3 & 0.5 \end{bmatrix}, \quad C^{(1)} = \begin{bmatrix} 0.4 & 0 \\ 0 & 0.4 \end{bmatrix}$$

Tasks:

1. Implement the VAR likelihood ratio using the functions from the VAR lecture
2. Implement Wald's sequential test with $\alpha = 0.05$ and $\beta = 0.10$
3. Analyze performance under both hypotheses and with model misspecification, when a third VAR(1) model generates the data
4. Compare with the Markov chain case in terms of stopping times and accuracy

Solution

Below is one solution to the exercise.

First we define the parameters for the VAR models and simulator

```
VARSPRTParams = namedtuple('VARSPRTParams',
            ['α', 'β', 'A_0', 'C_0', 'A_1', 'C_1', 'N', 'seed'])

def create_var_model(A, C):
    """Create VAR model."""
    μ_0 = np.zeros(A.shape[0])
    CC = C @ C.T
    Σ_0 = sp.linalg.solve_discrete_lyapunov(A, CC)

    CC_inv = np.linalg.inv(CC + 1e-10 * np.eye(CC.shape[0]))
    Σ_0_inv = np.linalg.inv(Σ_0 + 1e-10 * np.eye(Σ_0.shape[0]))

    return {
        'A': A, 'C': C, 'μ_0': μ_0, 'Σ_0': Σ_0,
        'CC_inv': CC_inv, 'Σ_0_inv': Σ_0_inv,
        'log_det_CC': np.log(
            np.linalg.det(CC + 1e-10 * np.eye(CC.shape[0]))),
        'log_det_Σ_0': np.log(
            np.linalg.det(Σ_0 + 1e-10 * np.eye(Σ_0.shape[0])))
    }
```

Now we define the likelihood ratio for the VAR models and the SPRT function similar to the Markov chain case

```
def var_log_likelihood(x_curr, x_prev, model, initial=False):
    """Compute VAR log-likelihood."""
    n = len(x_curr)
    if initial:
        diff = x_curr - model['μ_0']
        return -0.5 * (n * np.log(2 * np.pi) + model['log_det_Σ_0'] +
                       diff @ model['Σ_0_inv'] @ diff)
    else:
        diff = x_curr - model['A'] @ x_prev
```

```
        return -0.5 * (n * np.log(2 * np.pi) + model['log_det_CC'] +
                       diff @ model['CC_inv'] @ diff)

def var_sprt_single_run(model_0, model_1, model_true,
                        logA, logB, seed):
    """Single VAR SPRT run."""
    np.random.seed(seed)
    max_T = 500

    # Generate VAR path
    Σ_chol = np.linalg.cholesky(model_true['Σ_0'])
    x = model_true['μ_0'] + Σ_chol @ np.random.randn(
                len(model_true['μ_0']))

    # Initial likelihood ratio
    log_L = (var_log_likelihood(x, None, model_1, True) -
             var_log_likelihood(x, None, model_0, True))

    if log_L >= logA: return 1, False
    if log_L <= logB: return 1, True

    # Sequential updates
    for t in range(1, max_T):
        x_prev = x.copy()
        w = np.random.randn(model_true['C'].shape[1])
        x = model_true['A'] @ x + model_true['C'] @ w

        log_L += (var_log_likelihood(x, x_prev, model_1) -
                  var_log_likelihood(x, x_prev, model_0))

        if log_L >= logA: return t+1, False
        if log_L <= logB: return t+1, True

    return max_T, log_L < 0

def run_var_sprt(params):
    """Run VAR SPRT."""

    model_0 = create_var_model(params.A_0, params.C_0)
    model_1 = create_var_model(params.A_1, params.C_1)
    A, B, logA, logB = compute_wald_thresholds(params.α, params.β)

    stopping_times = np.zeros(params.N)
    decisions_h0 = np.zeros(params.N, dtype=bool)
    truth_h0 = np.zeros(params.N, dtype=bool)

    for i in range(params.N):
        model_true = model_0 if i % 2 == 0 else model_1
        truth_h0[i] = i % 2 == 0

        n, accept_h0 = var_sprt_single_run(model_0, model_1, model_true,
                                           logA, logB, params.seed + i)
        stopping_times[i] = n
        decisions_h0[i] = accept_h0

    type_I = np.sum(truth_h0 & ~decisions_h0) / np.sum(truth_h0)
    type_II = np.sum(~truth_h0 & decisions_h0) / np.sum(~truth_h0)
```

```
    return {'stopping_times': stopping_times,
            'decisions_h0': decisions_h0,
            'truth_h0': truth_h0,
            'type_I': type_I, 'type_II': type_II}
```

Let's run SPRT and visualize the results

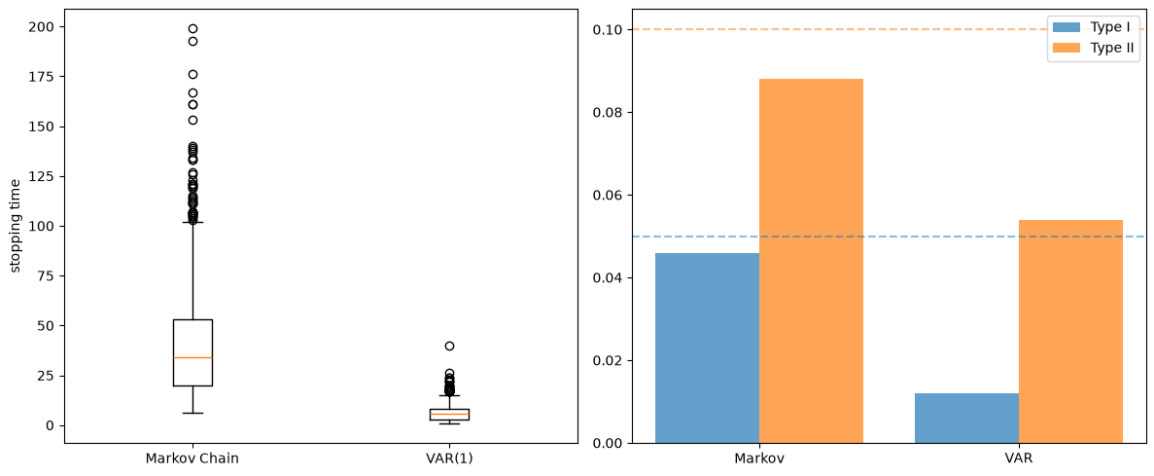
```
# Run VAR SPRT
A_0 = np.array([[0.8, 0.1],
                [0.2, 0.7]])
C_0 = np.array([[0.3, 0.1],
                [0.1, 0.3]])
A_1 = np.array([[0.6, 0.2],
                [0.3, 0.5]])
C_1 = np.array([[0.4, 0.0],
                [0.0, 0.4]])

params_var = VARSPRTParams(α=0.05, β=0.10,
                A_0=A_0, C_0=C_0, A_1=A_1, C_1=C_1,
                N=1000, seed=42)
results_var = run_var_sprt(params_var)

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(12, 5))

ax1.boxplot([results_markov['stopping_times'],
             results_var['stopping_times']],
            tick_labels=['Markov Chain', 'VAR(1)'])
ax1.set_ylabel('stopping time')

x = np.arange(2)
ax2.bar(x - 0.2, [results_markov['type_I'], results_var['type_I']],
        0.4, label='Type I', alpha=0.7)
ax2.bar(x + 0.2, [results_markov['type_II'], results_var['type_II']],
        0.4, label='Type II', alpha=0.7)
ax2.axhline(y=0.05, linestyle='--', alpha=0.5, color='C0')
ax2.axhline(y=0.10, linestyle='--', alpha=0.5, color='C1')
ax2.set_xticks(x), ax2.set_xticklabels(['Markov', 'VAR'])
ax2.legend()
plt.tight_layout()
plt.show()
```



Both tests keep their empirical error rates below the targets, and the VAR test stops much sooner than the Markov chain test.

Now we turn to misspecification.

Suppose that a third VAR(1) model h with

$$A^{(h)} = (1-\lambda)A^{(0)} + \lambda A^{(1)}, \quad C^{(h)} = (1-\lambda)C^{(0)} + \lambda C^{(1)}, \quad \lambda \in [0,1]$$

generates the data.

Let $f_i(x_{t+1}|x_t)$ be the conditional density of model i.

Under h, the expected increment of $\log L_t$ per period is

$$E_h\left[\log \frac{f_1(x_{t+1}|x_t)}{f_0(x_{t+1}|x_t)}\right] = KL(h, f_0) - KL(h, f_1)$$

where $KL(h, f_i)$ is the KL divergence of $f_i(\cdot|x_t)$ from the conditional density of h, averaged over the stationary distribution of x_t under h.

This is the formula in *this section* of *Likelihood Ratio Processes*, applied to conditional densities.

When $KL(h, f_0) > KL(h, f_1)$, $\log L_t$ drifts up and the test tends to reject H_0.

When $KL(h, f_0) < KL(h, f_1)$, $\log L_t$ drifts down and the test tends to accept H_0.

Either way, the test tends to select the model that is closer to h in KL divergence.

Let $\Sigma_i = C^{(i)}C^{(i)\top}$ and let V_h be the stationary covariance matrix of x_t under h.

For our Gaussian models,

$$E_h[\log f_i(x_{t+1}|x_t)] = -\frac{1}{2}\left[2\log 2\pi + \log\det\Sigma_i + \operatorname{tr}(\Sigma_i^{-1}\Sigma_h) + \operatorname{tr}\left((A^{(h)} - A^{(i)})^\top \Sigma_i^{-1}(A^{(h)} - A^{(i)})V_h\right)\right]$$

The code below computes $KL(h, f_0) - KL(h, f_1)$ with this formula and runs the test on 1000 paths from h for several values of λ.

```
def var_drift(A_h, C_h, model_0, model_1):
    """KL(h, f_0) - KL(h, f_1) per period under the stationary law of h."""
    n = A_h.shape[0]
    Σ_h = C_h @ C_h.T
    V_h = sp.linalg.solve_discrete_lyapunov(A_h, Σ_h)

    def expected_loglik(model):
        D = A_h - model['A']
        quad_form = (np.trace(model['CC_inv'] @ Σ_h)
                     + np.trace(D.T @ model['CC_inv'] @ D @ V_h))
        return -0.5 * (n * np.log(2 * np.pi) + model['log_det_CC']
                       + quad_form)

    return expected_loglik(model_1) - expected_loglik(model_0)

model_0 = create_var_model(A_0, C_0)
model_1 = create_var_model(A_1, C_1)
_, _, logA, logB = compute_wald_thresholds(0.05, 0.10)

rows = []
for λ in [0.0, 0.25, 0.4, 0.5, 0.75, 1.0]:
    A_h = (1 - λ) * A_0 + λ * A_1
    C_h = (1 - λ) * C_0 + λ * C_1
```

```
    model_h = create_var_model(A_h, C_h)
    runs = [var_sprt_single_run(model_0, model_1, model_h,
                                logA, logB, 42 + i)
            for i in range(1000)]
    n_stop = np.array([run[0] for run in runs])
    accept_h0 = np.array([run[1] for run in runs])
    rows.append([λ, var_drift(A_h, C_h, model_0, model_1),
                 1 - accept_h0.mean(), n_stop.mean()])

df_misspec = pd.DataFrame(rows, columns=['λ', 'KL(h,f0) - KL(h,f1)',
                                         'share rejecting H0',
                                         'mean stop time'])
df_misspec.set_index('λ').round(3)
```

λ	KL(h,f0) - KL(h,f1)	share rejecting H0	mean stop time
0.00	-0.362	0.015	7.210
0.25	-0.126	0.216	10.084
0.40	0.046	0.444	10.522
0.50	0.174	0.599	9.609
0.75	0.551	0.842	7.579
1.00	1.019	0.940	5.507

At $\lambda = 0$ and $\lambda = 1$, h coincides with one of the two models.

As λ rises, $KL(h, f_0) - KL(h, f_1)$ rises and changes sign between $\lambda = 0.25$ and $\lambda = 0.4$; the share of runs that reject H_0 rises with it.

Near the sign change, $\log L_t$ has little drift, and the test takes the most draws.

A driftless random walk with small steps hits $\log A$ before $\log B$ with probability close to $\log(1/B)/(\log A - \log B) = 0.44$, close to the share of runs that reject H_0 at $\lambda = 0.4$.

Exercise 26.8.3

Wald and Wolfowitz [1948] proved that among all tests, sequential or not, whose type I and type II error probabilities do not exceed those of an SPRT, the SPRT minimizes both $E_0[N]$ and $E_1[N]$.

This exercise compares the SPRT with the fixed-sample Neyman–Pearson test that the Navy told Captain Schuyler to use.

Use `results_2`, which we computed above for $f_0 = \text{Beta}(4, 5)$ and $f_1 = \text{Beta}(5, 4)$.

1. Compute the empirical error probabilities $\hat{\alpha}$ and $\hat{\beta}$ of the SPRT and its mean stopping times under f_0 and under f_1.
2. For each $n = 1, \ldots, 40$, construct the most powerful test of size $\hat{\alpha}$ based on a fixed sample of n draws, and compute its type II error probability $\beta(n)$.
3. Find the smallest n with $\beta(n) \leq \hat{\beta}$ and compare it with the SPRT's mean stopping times.

Hint: for these two densities, $\log\left(f_1(z)/f_0(z)\right) = \log\left(z/(1-z)\right)$.

Solution

By the Neyman–Pearson lemma, the most powerful test based on n draws rejects H_0 when $S_n = \sum_{i=1}^{n} \log\left(f_1(z_i)/f_0(z_i)\right) \geq k_n$, where k_n sets $P_0\{S_n \geq k_n\} = \hat{\alpha}$.

We compute the distributions of S_n under f_0 and f_1 by Monte Carlo.

```
truth_h0 = results_2['truth_h0']
stops = results_2['stopping_times']
α_hat, β_hat = results_2['type_I'], results_2['type_II']
E0_N_sprt = stops[truth_h0].mean()
E1_N_sprt = stops[~truth_h0].mean()

print(f"SPRT: α̂ = {α_hat:.3f}, β = {β_hat:.3f}, "
      f"E_0[N] = {E0_N_sprt:.1f}, E_1[N] = {E1_N_sprt:.1f}")

rng = np.random.default_rng(1234)
M, n_max = 100_000, 40

def llr_partial_sums(a, b):
    """Paths of S_n = sum of log(f_1/f_0), n = 1, ..., n_max."""
    z = rng.beta(a, b, size=(M, n_max))
    return np.cumsum(np.log(z / (1 - z)), axis=1)

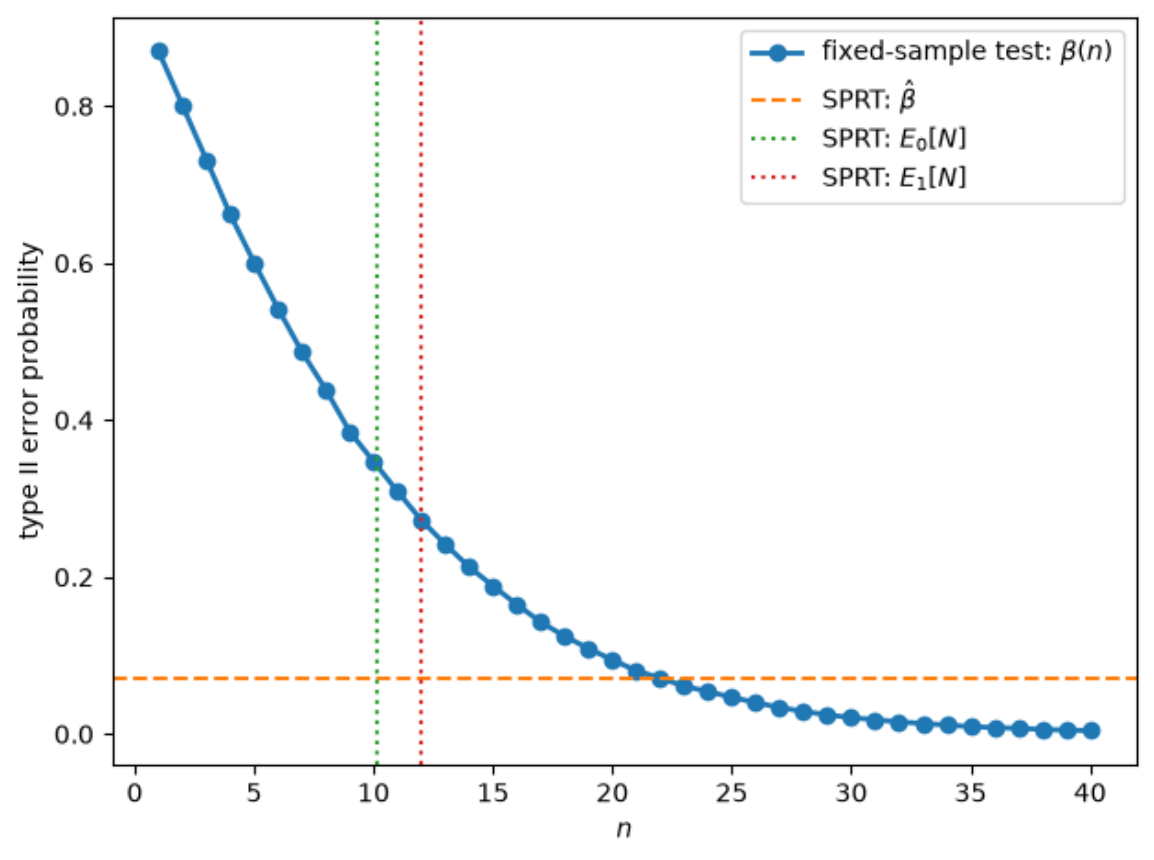
S0 = llr_partial_sums(params_2.a0, params_2.b0)
S1 = llr_partial_sums(params_2.a1, params_2.b1)

ns = np.arange(1, n_max + 1)
k_n = np.quantile(S0, 1 - α_hat, axis=0)   # critical values
β_n = (S1 < k_n).mean(axis=0)              # type II error probabilities
n_star = ns[np.argmax(β_n <= β_hat)]

print(f"smallest fixed n with β(n) <= β: {n_star}")

fig, ax = plt.subplots()
ax.plot(ns, β_n, 'o-', lw=2, label=r'fixed-sample test: $\beta(n)$')
ax.axhline(β_hat, color='C1', linestyle='--', label=r'SPRT: $\hat\beta$')
ax.axvline(E0_N_sprt, color='C2', linestyle=':', label='SPRT: $E_0[N]$')
ax.axvline(E1_N_sprt, color='C3', linestyle=':', label='SPRT: $E_1[N]$')
ax.set_xlabel('$n$')
ax.set_ylabel('type II error probability')
ax.legend()
plt.tight_layout()
plt.show()
```

```
SPRT: α̂ = 0.033, β = 0.070, E_0[N] = 10.1, E_1[N] = 12.0
smallest fixed n with β(n) <= β: 23
```



To match the SPRT's error probabilities, the fixed-sample test needs 23 draws.

The SPRT needs 10.1 draws on average under f_0 and 12.0 under f_1.

On average, the SPRT uses 44 percent as many draws as the fixed-sample test under f_0 and 52 percent as many under f_1.

The theorem of Wald and Wolfowitz implies that no test with error probabilities at most $\hat{\alpha}$ and $\hat{\beta}$ does better on either count.

Exercise 26.8.4

Wald's approximation (26.4) ignores overshoot and the fact that $N \geq 1$.

1. Modify `sprt_single_run` so that it also returns $\log L_N$.
2. For every fifth beta pair generated by `generate_β_pairs`, simulate 50,000 replications under f_0 and 50,000 under f_1 and estimate $E_0[N]$, $E_1[N]$, $E_0[\log L_N]$, and $E_1[\log L_N]$.
3. Verify Wald's identity (26.3).
4. Plot the ratios of simulated $E_0[N]$ and $E_1[N]$ to Wald's approximations (26.4) against $KL(f_0, f_1)$ and explain the pattern.

Solution

We work with log densities to avoid underflow when a shape parameter is small.

```
from math import lgamma

@njit
def log_p(z, a, b):
    """Log of the beta density."""
    return (lgamma(a + b) - lgamma(a) - lgamma(b)
            + (a - 1) * np.log(z) + (b - 1) * np.log(1 - z))

@njit
def sprt_single_run_log_L(a0, b0, a1, b1, logA, logB, true_f0, seed):
    """Run a single SPRT and return N and log L_N."""
    log_L = 0.0
    n = 0
    np.random.seed(seed)

    while True:
        z = np.random.beta(a0, b0) if true_f0 else np.random.beta(a1, b1)
        n += 1
        log_L += log_p(z, a1, b1) - log_p(z, a0, b0)
        if log_L >= logA or log_L <= logB:
            return n, log_L

@njit(parallel=True)
def simulate_N_log_L(a0, b0, a1, b1, α, β, M, true_f0, seed):
    """Draws of N and log L_N from M replications under one hypothesis."""
    A, B, logA, logB = compute_wald_thresholds(α, β)
    N_draws = np.empty(M)
    log_L_draws = np.empty(M)
    for i in prange(M):
        n, log_L = sprt_single_run_log_L(a0, b0, a1, b1,
                                         logA, logB, true_f0, seed + i)
        N_draws[i] = n
        log_L_draws[i] = log_L
    return N_draws, log_L_draws

def t_stat(d):
    """Sample mean of d divided by its standard error."""
    return d.mean() / (d.std() / np.sqrt(len(d)))

α, β, M = 0.05, 0.10, 50_000
rows, t_stats = [], []
for a0, b0, a1, b1 in param_comb[::5]:
    KL_01 = compute_KL(a0, b0, a1, b1)
    KL_10 = compute_KL(a1, b1, a0, b0)
    N0, log_L0 = simulate_N_log_L(a0, b0, a1, b1, α, β, M, True, 0)
    N1, log_L1 = simulate_N_log_L(a0, b0, a1, b1, α, β, M, False, M)
    E0_N_asn, E1_N_asn = wald_asn(α, β, KL_01, KL_10)
    rows.append([KL_01, N0.mean(), -log_L0.mean() / KL_01, E0_N_asn,
                 N1.mean(), log_L1.mean() / KL_10, E1_N_asn])
    # Wald's identity says that these differences have mean zero
    t_stats += [t_stat(N0 + log_L0 / KL_01), t_stat(N1 - log_L1 / KL_10)]

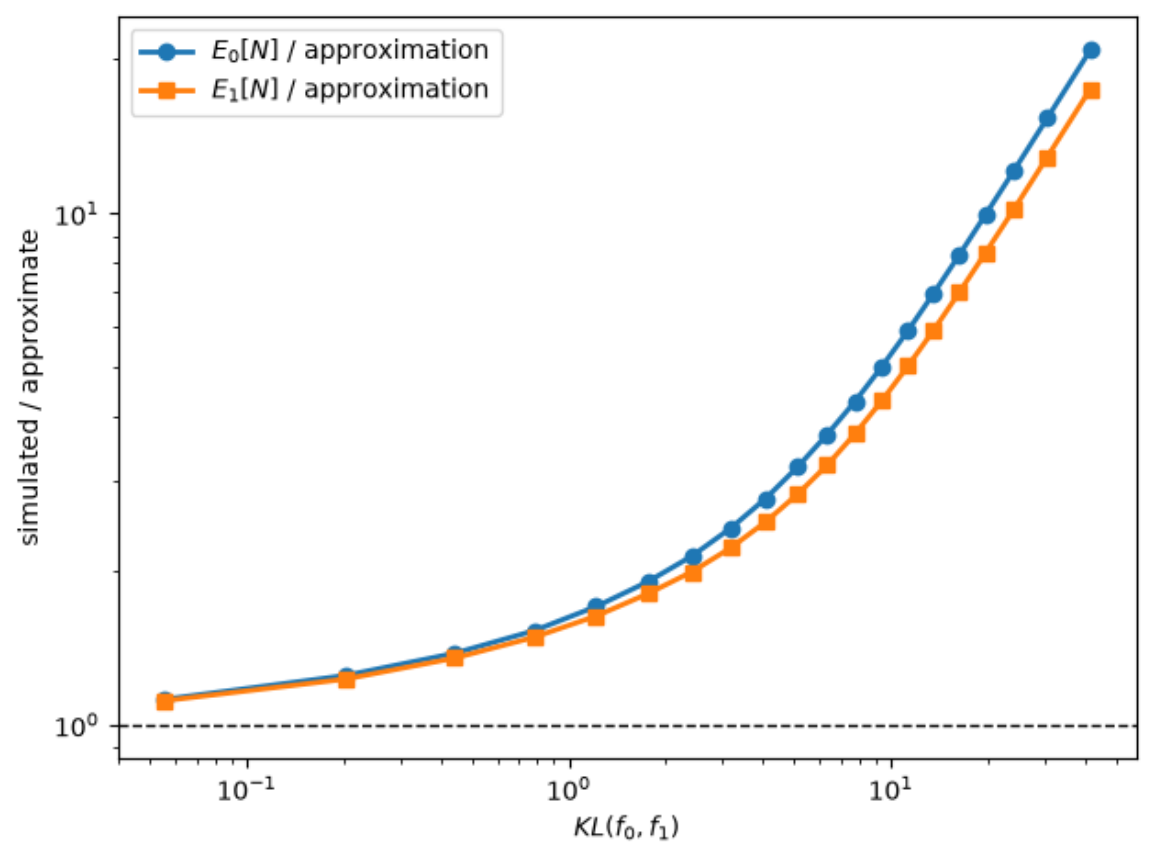
cols = ['KL', 'E0_N', 'E0_N_identity', 'E0_N_asn',
        'E1_N', 'E1_N_identity', 'E1_N_asn']
df_asn = pd.DataFrame(rows, columns=cols)
```

```
print(f"largest |t|-statistic for Wald's identity "
      f"across {len(t_stats)} checks: {np.max(np.abs(t_stats)):.2f}")

fig, ax = plt.subplots()
ax.plot(df_asn['KL'], df_asn['E0_N'] / df_asn['E0_N_asn'], 'o-', lw=2,
        label=r'$E_0[N]$ / approximation')
ax.plot(df_asn['KL'], df_asn['E1_N'] / df_asn['E1_N_asn'], 's-', lw=2,
        label=r'$E_1[N]$ / approximation')
ax.axhline(1, color='black', linestyle='--', lw=1)
ax.set_xscale('log')
ax.set_yscale('log')
ax.set_xlabel('$KL(f_0, f_1)$')
ax.set_ylabel('simulated / approximate')
ax.legend()
plt.tight_layout()
plt.show()

df_asn.iloc[[0, 5, 10, 15, 19]].round(3)
```

```
largest |t|-statistic for Wald's identity across 40 checks: 1.61
```



	KL	E0_N	E0_N_identity	E0_N_asn	E1_N	E1_N_identity	E1_N_asn
0	0.055	40.322	40.275	36.005	47.754	47.620	42.902
5	1.757	2.158	2.163	1.135	2.437	2.437	1.352
10	6.347	1.162	1.160	0.314	1.206	1.210	0.374
15	16.272	1.015	1.015	0.123	1.022	1.023	0.146
19	41.573	1.001	1.000	0.048	1.001	1.002	0.057

The columns `E0_N_identity` and `E1_N_identity` compute $E_0[N]$ and $E_1[N]$ from simulated $E_i[\log L_N]$ and (26.3).

They agree with the direct estimates up to Monte Carlo error: the largest of the 40 t-statistics is 1.61 in absolute value.

All of the gap between the simulated $E_i[N]$ and (26.4) therefore comes from Wald's approximation of $E_i[\log L_N]$, which ignores overshoot.

At the stopping time, $\log L_N$ lies beyond the threshold that it crosses.

Overshoot also pushes the error probabilities below α and β.

Both effects make $|E_i[\log L_N]|$ larger than the value that Wald's approximation assigns to it, so (26.4) understates $E_i[N]$.

For the closest pair, each draw moves $\log L_m$ by a small amount, the overshoot is small relative to $\log A$ and $\log(1/B)$, and the ratio is about 1.1.

As $KL(f_0, f_1)$ grows, a single draw often moves $\log L_m$ by more than the distance between $\log B$ and $\log A$.

Then the test stops after one or two draws, the overshoot dominates, and the ratio grows because (26.4) goes to zero while $N \geq 1$.

CHAPTER

TWENTYSEVEN

A BAYESIAN FORMULATION OF FRIEDMAN AND WALD'S PROBLEM

Contents

- *A Bayesian Formulation of Friedman and Wald's Problem*
  - *Overview*
  - *A Dynamic Programming Approach*
  - *Implementation*
  - *Analysis*
  - *Exercises*

27.1 Overview

This lecture revisits the statistical decision problem presented to Milton Friedman and W. Allen Wallis during World War II when they were analysts at the U.S. Government's Statistical Research Group at Columbia University.

In *an earlier lecture*, we described how Abraham Wald [Wald, 1947] solved the problem by extending frequentist hypothesis testing techniques and formulating the problem sequentially.

Note

Wald's idea of formulating the problem sequentially created links to the **dynamic programming** that Richard Bellman developed in the 1950s.

Arrow *et al.* [1949] computed Bayes solutions of sequential decision problems by backward induction, anticipating Bellman's methods.

Wald and Wolfowitz [1948] used Bayes solutions to prove that Wald's sequential probability ratio test is optimal.

A section below shows that the Bayesian decision rule studied in this lecture is a sequential probability ratio test.

As we learned in *Elementary Probability with Matrices* and *Two Meanings of Probability*, a frequentist statistician views a probability distribution as measuring relative frequencies of a statistic that he anticipates constructing from a very long sequence of i.i.d. draws from a known probability distribution.

That known probability distribution is his 'hypothesis'.

A frequentist statistician studies the distribution of that statistic under that known probability distribution

- when the distribution is a member of a set of parameterized probability distributions, his hypothesis takes the form of a particular parameter vector.
- this is what we mean when we say that the frequentist statistician 'conditions on the parameters'
- he regards the parameters as fixed numbers that are known to nature, but not to him.
- the statistician copes with his ignorance of those parameters by constructing type I and type II errors associated with frequentist hypothesis testing.

In this lecture, we reformulate Friedman and Wald's problem by transforming our point of view from the 'objective' frequentist perspective of *the lecture on Wald's sequential analysis* to an explicitly 'subjective' perspective taken by a Bayesian decision maker who regards parameters not as fixed numbers but as (hidden) random variables that are jointly distributed with the random variables that he can observe by sampling from that joint distribution.

To form that joint distribution, the Bayesian statistician supplements the conditional distributions used by the frequentist statistician with a prior probability distribution over the parameters that represents his personal, subjective opinion about them.

That lets the Bayesian statistician calculate the joint distribution that he requires to calculate the conditional distributions that he wants.

To proceed in this way, we endow our decision maker with

- an initial prior subjective probability $\pi_{-1} \in (0, 1)$ that nature uses to generate $\{z_k\}$ as a sequence of i.i.d. draws from f_1 rather than f_0.
- faith in Bayes' law as a way to revise his subjective beliefs as observations on $\{z_k\}$ sequence arrive.
- a loss function that tells how the decision maker values type I and type II errors.

In our *previous frequentist version*, key ideas in play were:

- Type I and type II statistical errors
  - a type I error occurs when you reject a null hypothesis that is true
  - a type II error occurs when you accept a null hypothesis that is false
- Abraham Wald's **sequential probability ratio test**
- The **power** of a statistical test
- The **critical region** of a statistical test
- A **most powerful test**

In this lecture about a Bayesian reformulation of the problem, additional ideas at work are

- an initial prior probability π_{-1} that model f_1 generates the data
- Bayes' Law
- a sequence of posterior probabilities that model f_1 is generating the data
- dynamic programming

This lecture uses ideas studied in the lectures on *likelihood ratio processes*, *their roles in Bayesian learning*, and *this lecture on exchangeability*.

Blackwell's Theorem on Comparing Experiments studies a complementary question: how to rank experiments by their informativeness.

Here the experiment is fixed, and the decision maker chooses when to stop sampling from it.

We'll begin with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from numba import jit, prange, float64, int64
from numba.experimental import jitclass
from math import gamma
```

27.2 A Dynamic Programming Approach

The following presentation of the problem closely follows Dimitri Bertsekas's treatment in **Dynamic Programming and Stochastic Control** [Bertsekas, 1976].

A decision-maker can observe a sequence of draws of a random variable z.

He (or she) wants to know which of two probability distributions f_0 or f_1 governs z.

Conditional on knowing that successive observations are drawn from distribution f_0, the sequence of random variables is independently and identically distributed (IID).

Conditional on knowing that successive observations are drawn from distribution f_1, the sequence of random variables is also independently and identically distributed (IID).

But the observer does not know which of the two distributions generated the sequence.

For reasons explained in *Exchangeability and Bayesian Updating*, this means that the sequence is not IID.

The observer has something to learn, namely, whether the observations are drawn from f_0 or from f_1.

The decision maker wants to decide which of the two distributions is generating outcomes.

We adopt a Bayesian formulation.

The decision maker begins with a prior probability

$$\pi_{-1} = \mathbb{P}\{f = f_1 \mid \text{ no observations}\} \in (0, 1)$$

Note

In Bertsekas [1976], the belief is associated with the distribution f_0, but here we associate the belief with the distribution f_1 to match the discussions in *the lecture on Wald's sequential analysis*.

After observing $k+1$ observations $z_k, z_{k-1}, \ldots, z_0$, he updates his personal probability that the observations are described by distribution f_1 to

$$\pi_k = \mathbb{P}\{f = f_1 \mid z_k, z_{k-1}, \ldots, z_0\}$$

which is calculated recursively by applying Bayes' law:

$$\pi_{k+1} = \frac{\pi_k f_1(z_{k+1})}{(1-\pi_k) f_0(z_{k+1}) + \pi_k f_1(z_{k+1})}, \quad k = -1, 0, 1, \ldots$$

After observing $z_k, z_{k-1}, \ldots, z_0$, the decision-maker believes that z_{k+1} has probability distribution

$$f_{\pi_k}(v) = (1-\pi_k) f_0(v) + \pi_k f_1(v),$$

which is a mixture of distributions f_0 and f_1, with the weight on f_1 being the posterior probability that $f = f_1$[1].

[1] The decision maker acts as if he believes that the sequence of random variables $[z_0, z_1, \ldots]$ is *exchangeable*. See *Exchangeability and Bayesian Updating* and [Kreps, 1988] chapter 11, for discussions of exchangeability.

To illustrate such a distribution, let's inspect some mixtures of beta distributions.

The density of a beta probability distribution with parameters a and b is

$$f(z;a,b) = \frac{\Gamma(a+b)z^{a-1}(1-z)^{b-1}}{\Gamma(a)\Gamma(b)} \quad \text{where} \quad \Gamma(t) := \int_0^\infty x^{t-1}e^{-x}dx$$

The next figure shows two beta distributions in the top panel.

The bottom panel presents mixtures of these distributions, with various mixing probabilities π_k

```
@jit
def p(x, a, b):
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x**(a-1) * (1 - x)**(b-1)

f0 = lambda x: p(x, 1, 1)
f1 = lambda x: p(x, 9, 9)
grid = np.linspace(0, 1, 50)

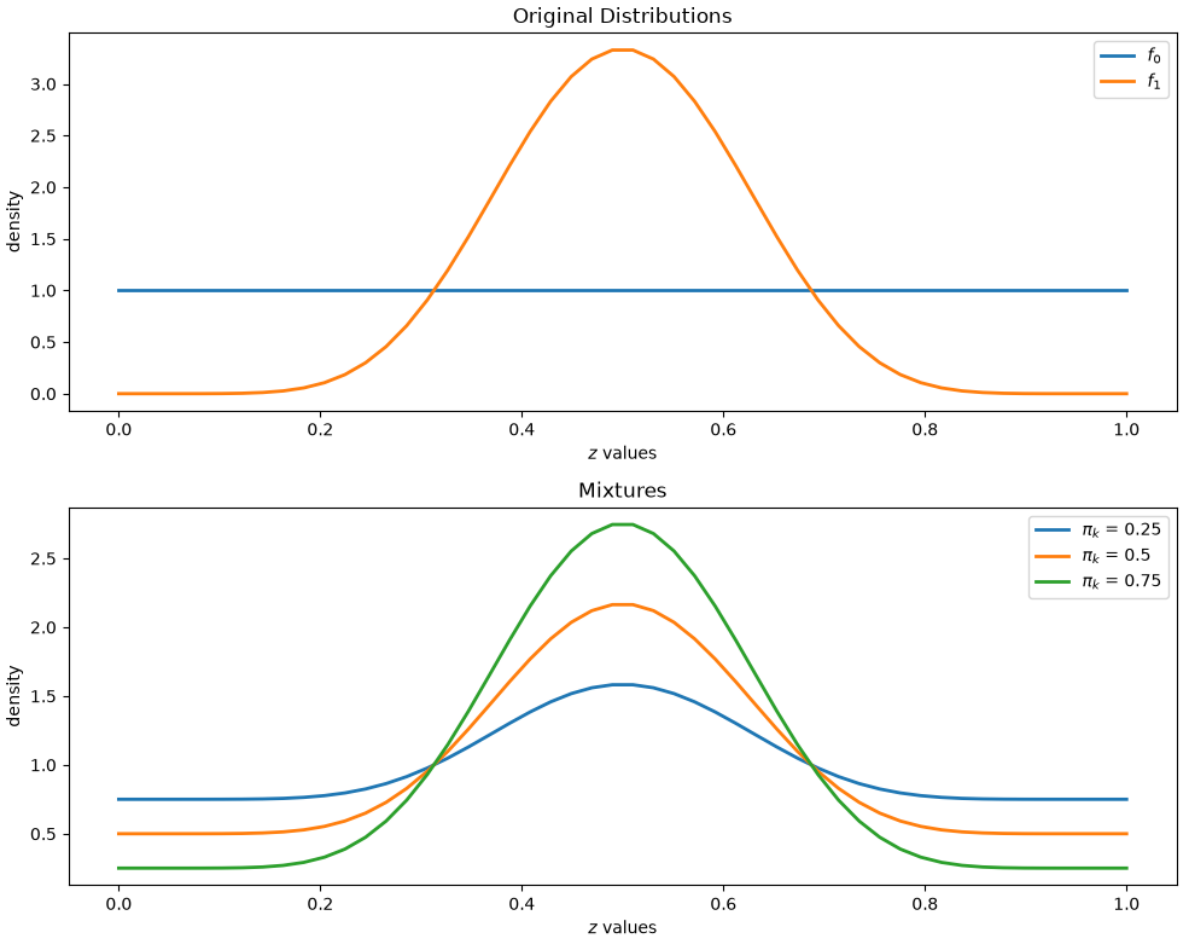
fig, axes = plt.subplots(2, figsize=(10, 8))

axes[0].set_title("Original Distributions")
axes[0].plot(grid, f0(grid), lw=2, label="$f_0$")
axes[0].plot(grid, f1(grid), lw=2, label="$f_1$")

axes[1].set_title("Mixtures")
for π in 0.25, 0.5, 0.75:
    y = (1 - π) * f0(grid) + π * f1(grid)
    axes[1].plot(grid, y, lw=2, label=fr"$\pi_k$ = {π}")

for ax in axes:
    ax.legend()
    ax.set(xlabel="$z$ values", ylabel="density")

plt.tight_layout()
plt.show()
```



27.2.1 Losses and Costs

After observing $z_k, z_{k-1}, \ldots, z_0$, the decision-maker chooses among three distinct actions:

- He decides that $f = f_0$ and draws no more z's
- He decides that $f = f_1$ and draws no more z's
- He postpones deciding now and instead chooses to draw a z_{k+1}

Associated with these three actions, the decision-maker can suffer three kinds of losses:

- A loss L_0 if he decides $f = f_0$ when actually $f = f_1$
- A loss L_1 if he decides $f = f_1$ when actually $f = f_0$
- A cost c if he postpones deciding and chooses instead to draw another z

27.2.2 Digression on Type I and Type II Errors

If we regard $f = f_0$ as a null hypothesis and $f = f_1$ as an alternative hypothesis, then L_1 and L_0 are losses associated with two types of statistical errors

- a type I error is an incorrect rejection of a true null hypothesis (a "false positive")
- a type II error is a failure to reject a false null hypothesis (a "false negative")

So when we treat $f = f_0$ as the null hypothesis

- We can think of L_1 as the loss associated with a type I error.
- We can think of L_0 as the loss associated with a type II error.

27.2.3 Intuition

Before proceeding, let's try to guess what an optimal decision rule might look like.

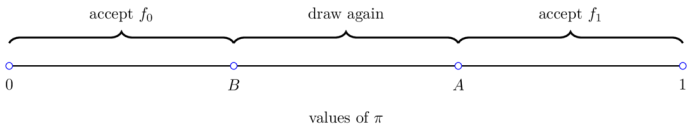
Suppose at some given point in time that π is close to 1.

Then our prior beliefs and the evidence so far point strongly to $f = f_1$.

If, on the other hand, π is close to 0, then $f = f_0$ is strongly favored.

Finally, if π is in the middle of the interval $[0, 1]$, then we are confronted with more uncertainty.

This reasoning suggests a sequential decision rule that we illustrate in the following figure:



As we'll see, this is indeed the correct form of the decision rule.

Our problem is to determine threshold values A, B that somehow depend on the parameters described above.

You might like to pause at this point and try to predict the impact of a parameter such as c or L_0 on A or B.

27.2.4 A Bellman Equation

Let $J(\pi)$ be the total loss for a decision-maker with current belief π who chooses optimally.

Principles of **dynamic programming** teach us that an optimal loss function J satisfies the following Bellman functional equation

$$J(\pi) = \min \left\{ \underbrace{\pi L_0}_{\text{accept } f_0} , \underbrace{(1-\pi) L_1}_{\text{accept } f_1} , \underbrace{c + \mathbb{E}[J(\pi')]}_{\text{draw again}} \right\} \tag{27.1}$$

where π' is the random variable defined by Bayes' Law

$$\pi' = \kappa(z', \pi) = \frac{\pi f_1(z')}{(1-\pi) f_0(z') + \pi f_1(z')}$$

when π is fixed and z' is drawn from the current best guess, which is the distribution f defined by

$$f_\pi(v) = (1-\pi) f_0(v) + \pi f_1(v)$$

In the Bellman equation, minimization is over three actions:

1. Accept the hypothesis that $f = f_0$
2. Accept the hypothesis that $f = f_1$
3. Postpone deciding and draw again

We can represent the Bellman equation as

$$J(\pi) = \min\{\pi L_0,\ (1-\pi)L_1,\ h(\pi)\} \tag{27.2}$$

where $\pi \in [0, 1]$ and

- πL_0 is the expected loss associated with accepting f_0 (i.e., the cost of making a type II error).
- $(1-\pi)L_1$ is the expected loss associated with accepting f_1 (i.e., the cost of making a type I error).
- $h(\pi) := c + \mathbb{E}[J(\pi')]$; this is the continuation value; i.e., the expected cost associated with drawing one more z.

The optimal decision rule is characterized by two numbers $0 < B \leq A < 1$ that satisfy

$$\pi L_0 \leq \min\{(1-\pi)L_1, c + \mathbb{E}[J(\pi')]\} \text{ if } \pi \leq B$$

and

$$(1-\pi)L_1 \leq \min\{\pi L_0, c + \mathbb{E}[J(\pi')]\} \text{ if } \pi \geq A$$

The optimal decision rule is then

$$\begin{aligned} \text{accept } f = f_1 \text{ if } \pi \geq A \\ \text{accept } f = f_0 \text{ if } \pi \leq B \\ \text{draw another } z \text{ if } B < \pi < A \end{aligned}$$

The two-threshold form of this rule follows from three properties of h and of the stopping loss $S(\pi) = \min\{\pi L_0, (1-\pi)L_1\}$.

First, $h(0) = h(1) = c$.

When $\pi \in \{0, 1\}$, Bayes' law sets $\pi' = \pi$, and $J(0) = J(1) = 0$.

Second, h is concave.

Write $\mathbb{E}[J(\pi')] = \int f_\pi(z') J\left(\pi f_1(z')/f_\pi(z')\right) dz'$.

For each z', the integrand is the perspective $tJ(x/t)$ of J evaluated at $x = \pi f_1(z')$ and $t = f_\pi(z')$, both linear in π.

The perspective of a concave function is concave, so $\mathbb{E}[J(\pi')]$ is concave in π whenever J is.

J is the minimum of h and two linear functions, so J is concave whenever h is.

Therefore the operator Q defined below maps concave functions into concave functions, and its fixed point h, a limit of iterates that start from $h = 0$, is concave.

Third, S is linear on each side of $\hat{\pi} = L_1/(L_0 + L_1)$.

On $[0, \hat{\pi}]$, accepting f_0 is optimal where $\pi L_0 - h(\pi) \leq 0$.

The function $\pi L_0 - h(\pi)$ is convex and equals $-c < 0$ at $\pi = 0$, so this set is an interval $[0, B]$ with $B > 0$.

On $[\hat{\pi}, 1]$, the same argument applied to $(1-\pi)L_1 - h(\pi)$ gives an interval $[A, 1]$ with $A < 1$ on which accepting f_1 is optimal.

The decision maker draws again on the interval (B, A) between them.

That interval is empty, and $B = A = \hat{\pi}$, when $h(\hat{\pi}) \geq S(\hat{\pi})$.

Because $h \geq c$, this happens whenever $c \geq L_0 L_1/(L_0 + L_1)$.

Our aim is to compute the cost function J as well as the associated cutoffs A and B.

To help make our computations more manageable, we can use (27.2) to write the continuation cost $h(\pi)$ as

$$\begin{aligned} h(\pi) &= c + \mathbb{E}[J(\pi')] \\ &= c + \mathbb{E}_{\pi'} \min\{\pi' L_0, (1-\pi') L_1, h(\pi')\} \\ &= c + \int \min\{\kappa(z', \pi) L_0, (1 - \kappa(z', \pi)) L_1, h(\kappa(z', \pi))\} f_\pi(z') dz' \end{aligned} \tag{27.3}$$

The equality

$$h(\pi) = c + \int \min\{\kappa(z', \pi) L_0, (1 - \kappa(z', \pi)) L_1, h(\kappa(z', \pi))\} f_\pi(z') dz' \tag{27.4}$$

is an equation in an unknown function h.

Note

Such an equation is called a **functional equation**.

Using the functional equation, (27.4), for the continuation cost, we can back out optimal choices using the right side of (27.2).

This functional equation can be solved by taking an initial guess and iterating to find a fixed point.

Thus, we iterate with an operator Q, where

$$Qh(\pi) = c + \int \min\{\kappa(z', \pi) L_0, (1 - \kappa(z', \pi)) L_1, h(\kappa(z', \pi))\} f_\pi(z') dz'$$

The operator Q is monotone: $h \leq g$ implies $Qh \leq Qg$.

The decision maker does not discount future losses, so Q is not a contraction in the sup norm.

We start from $h = 0$.

Then $Qh \geq h$, and monotonicity makes the iterates $Q^n h$ increase with n.

The iterates are bounded above by $c + L_0 L_1 / (L_0 + L_1)$, so they converge, and the limit solves (27.4).

27.3 Implementation

First, we will construct a `jitclass` to store the parameters of the model

```
wf_data = [('a0', float64),          # Parameters of beta distributions
           ('b0', float64),
           ('a1', float64),
           ('b1', float64),
           ('c', float64),           # Cost of another draw
           ('π_grid_size', int64),
           ('L0', float64),          # Cost of selecting f0 when f1 is true
           ('L1', float64),          # Cost of selecting f1 when f0 is true
           ('π_grid', float64[:]),
           ('mc_size', int64),
           ('z0', float64[:]),
           ('z1', float64[:])]
```

```
@jitclass(wf_data)
class WaldFriedman:

    def __init__(self,
                 c=1.25,
                 a0=1,
                 b0=1,
                 a1=3,
                 b1=1.2,
                 L0=25,
                 L1=25,
                 π_grid_size=200,
                 mc_size=1000,
                 seed=1234):

        self.a0, self.b0 = a0, b0
        self.a1, self.b1 = a1, b1
        self.c, self.π_grid_size = c, π_grid_size
        self.L0, self.L1 = L0, L1
        self.π_grid = np.linspace(0, 1, π_grid_size)
        self.mc_size = mc_size

        np.random.seed(seed)   # Monte Carlo draws used to compute E[J(π')]
        self.z0 = np.random.beta(a0, b0, mc_size)
        self.z1 = np.random.beta(a1, b1, mc_size)

    def f0(self, x):

        return p(x, self.a0, self.b0)

    def f1(self, x):

        return p(x, self.a1, self.b1)

    def f0_rvs(self):
        return np.random.beta(self.a0, self.b0)

    def f1_rvs(self):
        return np.random.beta(self.a1, self.b1)

    def κ(self, z, π):
        """
        Updates π using Bayes' rule and the current observation z
        """

        f0, f1 = self.f0, self.f1

        π_f0, π_f1 = (1 - π) * f0(z), π * f1(z)
        π_new = π_f1 / (π_f0 + π_f1)

        return π_new
```

As in *Optimal Savings III: Stochastic Returns*, to approximate a continuous value function

- We iterate at a finite grid of possible values of π.
- When we evaluate $\mathbb{E}[J(\pi')]$ between grid points, we use linear interpolation.

We define the operator function `Q` below.

```
@jit(nopython=True, parallel=True)
def Q(h, wf):

    c, π_grid = wf.c, wf.π_grid
    L0, L1 = wf.L0, wf.L1
    z0, z1 = wf.z0, wf.z1
    mc_size = wf.mc_size

    κ = wf.κ

    h_new = np.empty_like(π_grid)
    h_func = lambda p: np.interp(p, π_grid, h)

    for i in prange(len(π_grid)):
        π = π_grid[i]

        # Find the expected value of J by integrating over z
        integral_f0, integral_f1 = 0, 0
        for m in range(mc_size):
            π_0 = κ(z0[m], π)  # Draw z from f0 and update π
            integral_f0 += min(π_0 * L0, (1 - π_0) * L1, h_func(π_0))

            π_1 = κ(z1[m], π)  # Draw z from f1 and update π
            integral_f1 += min(π_1 * L0, (1 - π_1) * L1, h_func(π_1))

        integral = ((1 - π) * integral_f0 + π * integral_f1) / mc_size

        h_new[i] = c + integral

    return h_new
```

To solve the key functional equation, we will iterate using `Q` to find the fixed point

```
@jit
def solve_model(wf, tol=1e-4, max_iter=1000):
    """
    Compute the continuation cost function

    * wf is an instance of WaldFriedman
    """

    # Set up loop
    h = np.zeros(len(wf.π_grid))
    i = 0
    error = tol + 1

    while i < max_iter and error > tol:
        h_new = Q(h, wf)
        error = np.max(np.abs(h - h_new))
        i += 1
        h = h_new

    if error > tol:
        print("Failed to converge!")

    return h_new
```

27.4 Analysis

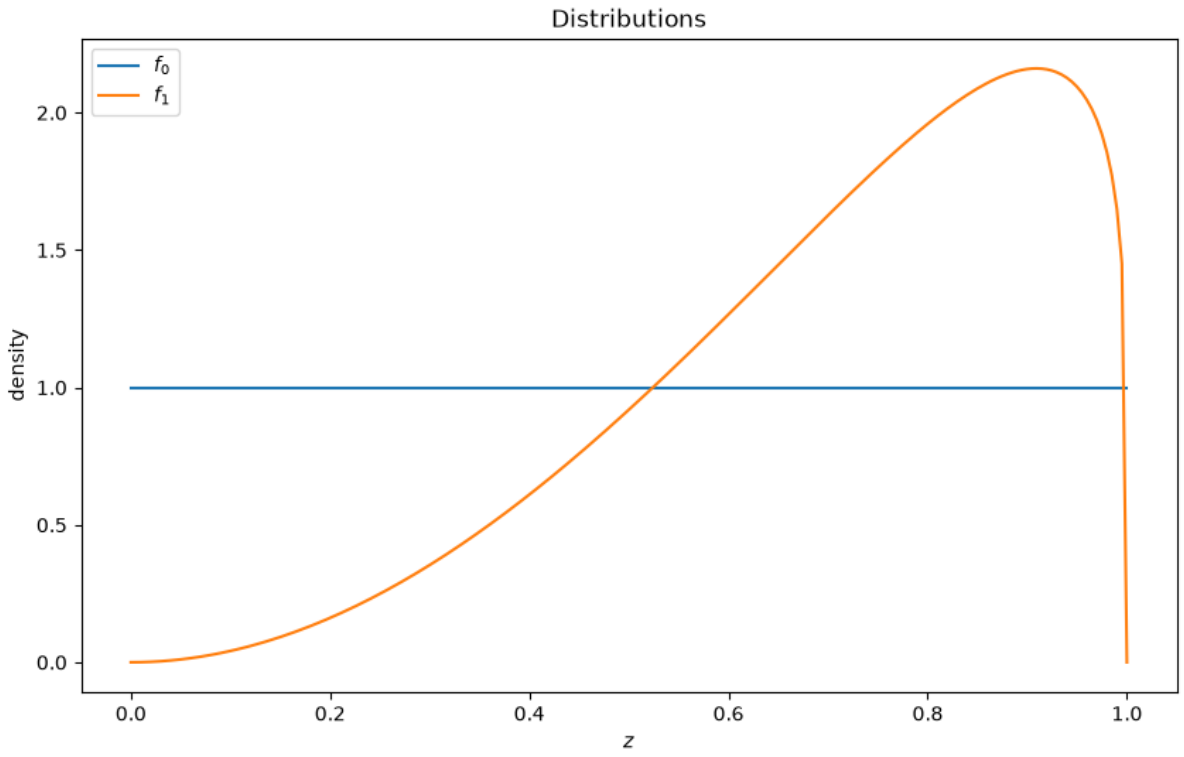
Let's inspect outcomes.

We will be using the default parameterization with distributions like so

```
wf = WaldFriedman()

z_grid = np.linspace(0, 1, 200)
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(z_grid, wf.f0(z_grid), label="$f_0$")
ax.plot(z_grid, wf.f1(z_grid), label="$f_1$")
ax.set(ylabel="density", xlabel="$z$", title="Distributions")
ax.legend()

plt.show()
```



27.4.1 Cost Function

To solve the model, we will call our `solve_model` function

```
h_star = solve_model(wf)     # Solve the model
```

We will also set up a function to compute the cutoffs A and B and plot these on our cost function plot

```
@jit
def find_cutoff_rule(wf, h):

    """
```

(continues on next page)

(continued from previous page)

```
    This function takes a continuation cost function and returns the
    corresponding cutoffs of where you transition between continuing and
    choosing a specific model
    """

    π_grid = wf.π_grid
    L0, L1 = wf.L0, wf.L1

    # Evaluate cost at all points on grid for choosing a model
    cost_f0 = π_grid * L0
    cost_f1 = (1 - π_grid) * L1

    # Find B: largest π where cost_f0 <= min(cost_f1, h)
    optimal_cost = np.minimum(np.minimum(cost_f0, cost_f1), h)
    choose_f0 = (cost_f0 <= cost_f1) & (cost_f0 <= h)

    if np.any(choose_f0):
        B = π_grid[choose_f0][-1]  # Last point where we choose f0
    else:
        assert False, "No point where we choose f0"

    # Find A: smallest π where cost_f1 <= min(cost_f0, h)
    choose_f1 = (cost_f1 <= cost_f0) & (cost_f1 <= h)

    if np.any(choose_f1):
        A = π_grid[choose_f1][0]  # First point where we choose f1
    else:
        assert False, "No point where we choose f1"

    return (B, A)

B, A = find_cutoff_rule(wf, h_star)
cost_L0 = wf.π_grid * wf.L0
cost_L1 = (1 - wf.π_grid) * wf.L1

fig, ax = plt.subplots(figsize=(10, 6))

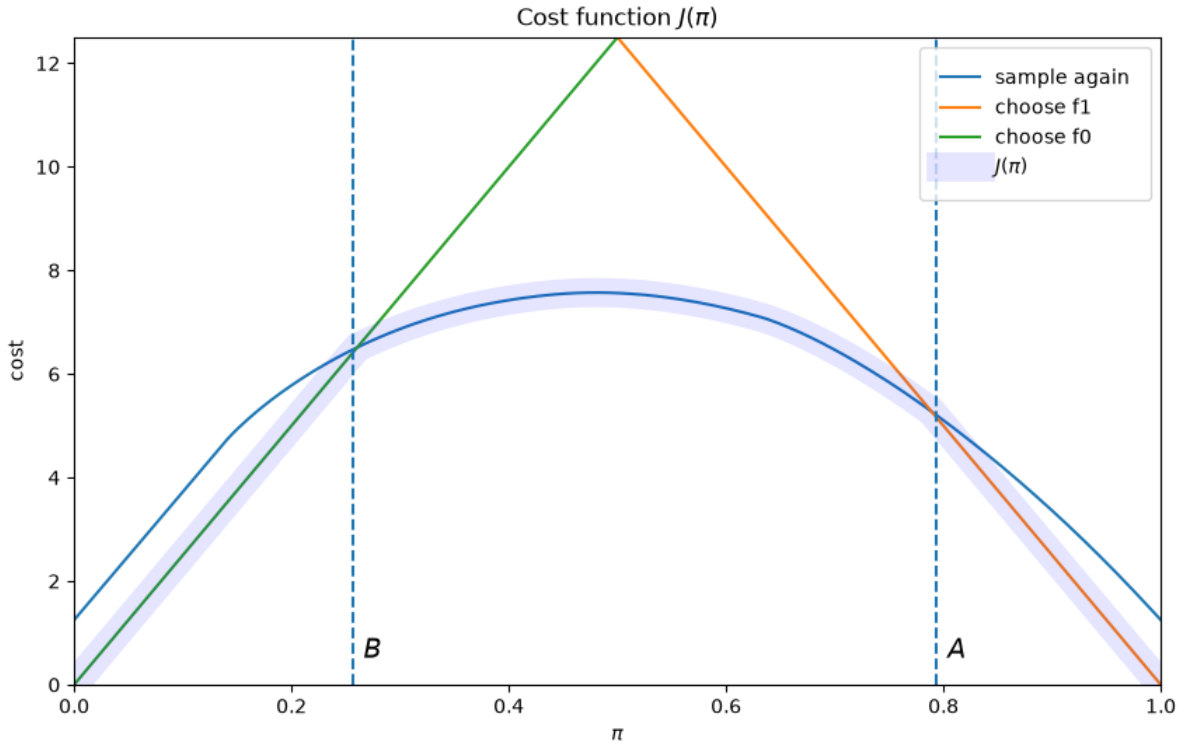
ax.plot(wf.π_grid, h_star, label='sample again')
ax.plot(wf.π_grid, cost_L1, label='choose f1')
ax.plot(wf.π_grid, cost_L0, label='choose f0')
ax.plot(wf.π_grid,
        np.amin(np.column_stack([h_star, cost_L0, cost_L1]),axis=1),
        lw=15, alpha=0.1, color='b', label=r'$J(\pi)$')

ax.annotate(r"$B$", xy=(B + 0.01, 0.5), fontsize=14)
ax.annotate(r"$A$", xy=(A + 0.01, 0.5), fontsize=14)

plt.vlines(B, 0, (1 - B) * wf.L1, linestyle="--")
plt.vlines(A, 0, A * wf.L0, linestyle="--")

ax.set(xlim=(0, 1), ylim=(0, 0.5 * max(wf.L0, wf.L1)), ylabel="cost",
       xlabel=r"$\pi$", title=r"Cost function $J(\pi)$")

plt.legend(borderpad=1.1)
plt.show()
```



The cost function J equals πL_0 for $\pi \leq B$, and $(1-\pi)L_1$ for $\pi \geq A$.

The slopes of the two linear pieces of the cost function $J(\pi)$ are determined by L_0 and $-L_1$.

The cost function J is smooth in the interior region, where the posterior probability assigned to f_1 is in the indecisive region $\pi \in (B, A)$.

The decision-maker continues to sample until the probability that he attaches to model f_1 falls below B or above A.

27.4.2 The Bayesian rule is a sequential probability ratio test

Let

$$\Lambda_k = \prod_{j=0}^{k} \frac{f_1(z_j)}{f_0(z_j)}$$

be the likelihood ratio after the $k+1$ observations $z_0, \ldots, z_k$.

We write Λ_k rather than L_k because L_0 and L_1 denote losses in this lecture.

Iterating on Bayes' law, as in *Likelihood Ratio Processes and Bayesian Learning*, gives

$$\pi_k = \frac{\pi_{-1}\Lambda_k}{\pi_{-1}\Lambda_k + 1 - \pi_{-1}}$$

The posterior π_k is a strictly increasing function of Λ_k.

Therefore $\pi_k \geq A$ if and only if $\Lambda_k \geq A_\Lambda$, and $\pi_k \leq B$ if and only if $\Lambda_k \leq B_\Lambda$, where

$$A_\Lambda = \frac{1-\pi_{-1}}{\pi_{-1}} \frac{A}{1-A}, \qquad B_\Lambda = \frac{1-\pi_{-1}}{\pi_{-1}} \frac{B}{1-B} \tag{27.5}$$

The Bayesian decision rule is Wald's sequential probability ratio test from *A Problem that Stumped Milton Friedman* with thresholds A_Λ and B_Λ on the likelihood ratio.

The thresholds do not depend on k because J and h do not depend on time.

Wald and Wolfowitz [1948] used this connection in the other direction.

They showed that every sequential probability ratio test is a Bayes rule for a problem of this kind with suitably chosen losses and cost per draw.

They used that fact to prove that no other test with the same or smaller error probabilities has a smaller expected number of draws under f_0 or under f_1.

Arrow *et al.* [1949] characterized Bayes and minimax solutions of more general sequential decision problems.

For our default parameters and $\pi_{-1} = 0.5$, the thresholds are

```
def LR_thresholds(B, A, π_minus_1=0.5):
    """Map thresholds on π into thresholds on the likelihood ratio."""
    prior_odds = (1 - π_minus_1) / π_minus_1
    return prior_odds * B / (1 - B), prior_odds * A / (1 - A)

B_Λ, A_Λ = LR_thresholds(B, A)
print(f"B = {B:.3f}, A = {A:.3f}")
print(f"B_Λ = {B_Λ:.2f}, A_Λ = {A_Λ:.2f}")
```

```
B = 0.256, A = 0.794
B_Λ = 0.34, A_Λ = 3.85
```

Exercise 27.5.1 compares the error probabilities of this test with the ones that Wald's formulas in *A Problem that Stumped Milton Friedman* imply.

27.4.3 Simulations

The next figure shows the outcomes of 500 simulations of the decision process.

In all 500 simulations, f_0 generates the data, and the decision maker starts from the prior $\pi_{-1} = 0.5$.

On the left is a histogram of **stopping times**, i.e., the number of draws of z_k required to make a decision.

The average number of draws is 2.82.

On the right is the fraction of correct decisions at the stopping time.

In this case, the decision-maker is correct 80% of the time.

```
@jit
def seed_numba(seed):
    """Seed the random number generator used by numba-compiled code."""
    np.random.seed(seed)

def simulate(wf, true_dist, h_star, π_0=0.5):

    """
    This function takes an initial condition and simulates until it
    stops (when a decision is made)
    """

    f0, f1 = wf.f0, wf.f1
    f0_rvs, f1_rvs = wf.f0_rvs, wf.f1_rvs
    π_grid = wf.π_grid
    κ = wf.κ
```

(continues on next page)

(continued from previous page)

```
    if true_dist == "f0":
        f, f_rvs = wf.f0, wf.f0_rvs
    elif true_dist == "f1":
        f, f_rvs = wf.f1, wf.f1_rvs

    # Find cutoffs
    B, A = find_cutoff_rule(wf, h_star)

    # Initialize a couple of useful variables
    decision_made = False
    π = π_0
    t = 0

    while decision_made is False:
        z = f_rvs()
        t = t + 1
        π = κ(z, π)
        if π <= B:
            decision_made = True
            decision = 0
        elif π >= A:
            decision_made = True
            decision = 1

    if true_dist == "f0":
        if decision == 0:
            correct = True
        else:
            correct = False

    elif true_dist == "f1":
        if decision == 1:
            correct = True
        else:
            correct = False

    return correct, π, t

def stopping_dist(wf, h_star, ndraws=250, true_dist="f0", seed=0):

    """
    Simulates repeatedly to get distributions of time needed to make a
    decision and how often they are correct
    """

    seed_numba(seed)
    tdist = np.empty(ndraws, int)
    cdist = np.empty(ndraws, bool)

    for i in range(ndraws):
        correct, π, t = simulate(wf, true_dist, h_star)
        tdist[i] = t
        cdist[i] = correct

    return cdist, tdist

def simulation_plot(wf):
```

(continues on next page)

(continued from previous page)

```
    h_star = solve_model(wf)
    ndraws = 500
    cdist, tdist = stopping_dist(wf, h_star, ndraws)

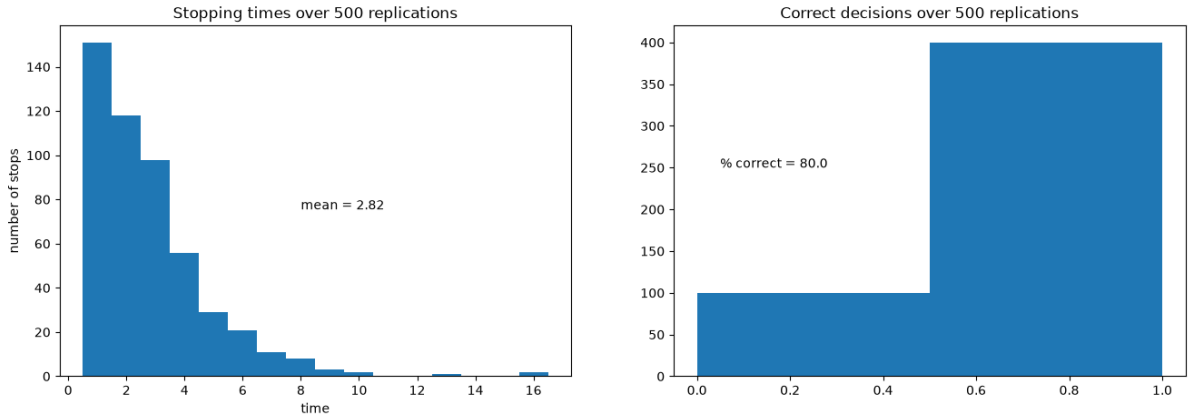
    fig, ax = plt.subplots(1, 2, figsize=(16, 5))

    counts, _, _ = ax[0].hist(tdist, bins=np.arange(1, np.max(tdist) + 2) - 0.5)
    ax[0].set_title(f"Stopping times over {ndraws} replications")
    ax[0].set(xlabel="time", ylabel="number of stops")
    ax[0].annotate(f"mean = {np.mean(tdist):.2f}",
                   xy=(max(tdist) / 2, max(counts) / 2))

    ax[1].hist(cdist.astype(int), bins=2)
    ax[1].set_title(f"Correct decisions over {ndraws} replications")
    ax[1].annotate(f"% correct = {100 * np.mean(cdist):.1f}",
                   xy=(0.05, ndraws / 2))

    plt.show()

simulation_plot(wf)
```



27.4.4 Comparative Statics

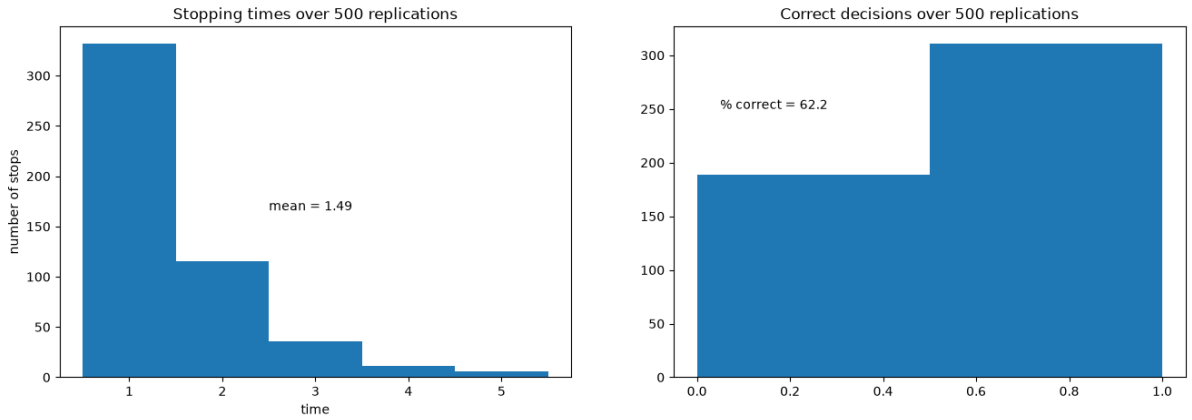
Now let's consider the following exercise.

We double the cost of drawing an additional observation.

Before you look, think about what will happen:

- Will the decision-maker be correct more or less often?
- Will he make decisions sooner or later?

```
wf = WaldFriedman(c=2.5)
simulation_plot(wf)
```



Increased cost per draw has induced the decision-maker to take fewer draws before deciding.

Because he decides with fewer draws, the percentage of time he is correct drops.

This leads to him having a higher expected loss when he puts equal weight on both models.

The code below computes $J(0.5)$ and the mean number of draws $\mathbb{E}[N]$ when $\pi_{-1} = 0.5$ for both values of c.

Because the prior puts equal weight on f_0 and f_1, $\mathbb{E}[N]$ averages simulated mean stopping times under f_0 and under f_1.

```
def J_and_mean_N(c, ndraws=5000):
    """J(0.5) and the mean number of draws when π_{-1} = 0.5."""
    wf_c = WaldFriedman(c=c)
    h_c = solve_model(wf_c)
    S = np.minimum(wf_c.π_grid * wf_c.L0, (1 - wf_c.π_grid) * wf_c.L1)
    J_half = np.interp(0.5, wf_c.π_grid, np.minimum(S, h_c))
    _, t0 = stopping_dist(wf_c, h_c, ndraws, "f0", seed=1)
    _, t1 = stopping_dist(wf_c, h_c, ndraws, "f1", seed=2)
    return J_half, 0.5 * (t0.mean() + t1.mean())

for c in (1.25, 2.5):
    J_half, mean_N = J_and_mean_N(c)
    print(f"c = {c}: J(0.5) = {J_half:.2f}, E[N] = {mean_N:.2f}")
```

```
c = 1.25: J(0.5) = 7.56, E[N] = 3.34
c = 2.5: J(0.5) = 10.42, E[N] = 1.49
```

Doubling c raises $J(0.5)$ from 7.56 to 10.42.

The envelope theorem accounts for the size of the increase.

For a fixed decision rule, the expected loss equals $c\,\mathbb{E}[N]$ plus the expected loss from wrong decisions, a linear function of c with slope $\mathbb{E}[N]$.

J is the minimum of these linear functions over decision rules, so $dJ/dc = \mathbb{E}[N^*]$ at every c where J is differentiable; N^* is the number of draws under the optimal rule.

$\mathbb{E}[N^*]$ falls from 3.34 to 1.49 as c rises from 1.25 to 2.5, so the increase in $J(0.5)$ lies between $1.25 \times 1.49 = 1.86$ and $1.25 \times 3.34 = 4.17$.

The exercises below study how c, L_0, and L_1 shape the thresholds B and A and the error probabilities of the Bayesian decision rule.

27.5 Exercises

Exercise 27.5.1

Use the default parameters, $\pi_{-1} = 0.5$, and the thresholds B_Λ and A_Λ that we computed *above*.

1. Simulate the Bayesian decision rule 20,000 times under f_0 and 20,000 times under f_1, and estimate its type I error probability α (accepting f_1 when f_0 generates the data) and its type II error probability β (accepting f_0 when f_1 generates the data).
2. In *A Problem that Stumped Milton Friedman*, Wald's inequalities $\alpha \leq (1-\beta)/A$ and $\beta \leq B(1-\alpha)$ hold with equality when the likelihood ratio never jumps past its thresholds. Use those equalities with $A = A_\Lambda$ and $B = B_\Lambda$ to compute the error probabilities that the test would have without overshoot, and compare them with your estimates.
3. Record $\log \Lambda_N$ at the stopping time N and plot its distribution under f_0 and under f_1. Use the plot to explain the comparison in step 2.

Solution

Without overshoot, $\alpha = (1-\beta)/A_\Lambda$ and $\beta = B_\Lambda(1-\alpha)$.

Solving these two equations gives

$$\alpha = \frac{1 - B_\Lambda}{A_\Lambda - B_\Lambda}, \qquad \beta = \frac{B_\Lambda(A_\Lambda - 1)}{A_\Lambda - B_\Lambda}$$

With $\pi_{-1} = 0.5$, Λ_N equals the posterior odds $\pi_N/(1-\pi_N)$ at the stopping time.

```
def simulate_many(wf, h_star, ndraws, true_dist, seed):
    """Correctness, terminal posterior, and stopping time of ndraws runs."""
    seed_numba(seed)
    correct = np.empty(ndraws, bool)
    π_N = np.empty(ndraws)
    N = np.empty(ndraws, int)
    for i in range(ndraws):
        correct[i], π_N[i], N[i] = simulate(wf, true_dist, h_star)
    return correct, π_N, N

wf = WaldFriedman()
h_star = solve_model(wf)
B, A = find_cutoff_rule(wf, h_star)
B_Λ, A_Λ = LR_thresholds(B, A)

correct_0, π_N0, N0 = simulate_many(wf, h_star, 20_000, "f0", seed=1)
correct_1, π_N1, N1 = simulate_many(wf, h_star, 20_000, "f1", seed=2)
α_hat, β_hat = 1 - correct_0.mean(), 1 - correct_1.mean()

α_wald = (1 - B_Λ) / (A_Λ - B_Λ)
β_wald = B_Λ * (A_Λ - 1) / (A_Λ - B_Λ)

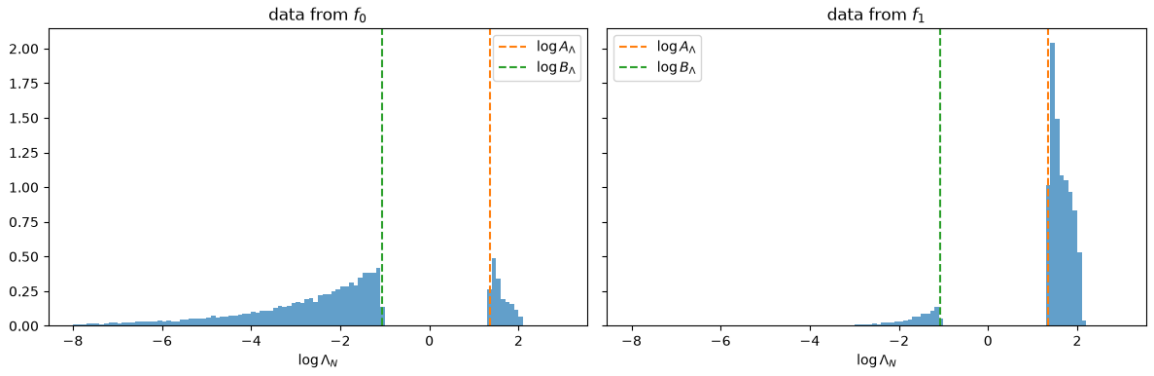
print(f"simulated:           α = {α_hat:.3f}, β = {β_hat:.3f}")
print(f"without overshoot:   α = {α_wald:.3f}, β = {β_wald:.3f}")

log_Λ0 = np.log(π_N0 / (1 - π_N0))
log_Λ1 = np.log(π_N1 / (1 - π_N1))
```

```
log_Λ = np.concatenate((log_Λ0, log_Λ1))
above = log_Λ >= np.log(A_Λ)
print(f"mean of log Λ_N - log A_Λ when the test accepts f1: "
      f"{np.mean(log_Λ[above] - np.log(A_Λ)):.2f}")
print(f"mean of log B_Λ - log Λ_N when the test accepts f0: "
      f"{np.mean(np.log(B_Λ) - log_Λ[~above]):.2f}")
print(f"runs with log Λ_N < -8: {np.sum(log_Λ0 < -8)} under f0, "
      f"{np.sum(log_Λ1 < -8)} under f1")

fig, axes = plt.subplots(1, 2, figsize=(12, 4), sharey=True)
for ax, x, title in zip(axes, (log_Λ0, log_Λ1),
                        ("data from $f_0$", "data from $f_1$")):
    ax.hist(x, bins=np.linspace(-8, 3, 111), density=True, alpha=0.7)
    ax.axvline(np.log(A_Λ), color='C1', linestyle='--',
               label=r'$\log A_\Lambda$')
    ax.axvline(np.log(B_Λ), color='C2', linestyle='--',
               label=r'$\log B_\Lambda$')
    ax.set_title(title)
    ax.set_xlabel(r'$\log \Lambda_N$')
    ax.legend()
plt.tight_layout()
plt.show()
```

```
simulated:           α = 0.176, β = 0.095
without overshoot:   α = 0.187, β = 0.280
mean of log Λ_N - log A_Λ when the test accepts f1: 0.29
mean of log B_Λ - log Λ_N when the test accepts f0: 1.83
runs with log Λ_N < -8: 494 under f0, 0 under f1
```



The histograms omit the 494 runs under f_0 in which $\log \Lambda_N < -8$.

The simulated β is far below the value without overshoot, and the simulated α is somewhat below it.

On average, $\log \Lambda_N$ exceeds $\log A_\Lambda$ by 0.29 when the test accepts f_1 and falls short of $\log B_\Lambda$ by 1.83 when the test accepts f_0.

Because f_0 is uniform, the likelihood ratio of a single draw is $f_1(z)$.

The density $f_1(z)$ never exceeds 2.16, so one draw raises $\log \Lambda$ by at most $\log 2.16 = 0.77$, and $\log \Lambda_N$ lands only a little above $\log A_\Lambda$.

The density $f_1(z)$ goes to zero as $z \to 0$, so one draw can lower $\log \Lambda$ by an arbitrarily large amount, and $\log \Lambda_N$ often lands far below $\log B_\Lambda$.

The change of measure in *A Problem that Stumped Milton Friedman* gives $\beta = E_0[\Lambda_N \mathbf{1}\{\Lambda_N \leq B_\Lambda\}]$, so the further Λ_N falls below B_Λ, the smaller β is relative to $B_\Lambda(1-\alpha)$.

Exercise 27.5.2

Let $S(\pi) = \min\{\pi L_0, (1-\pi)L_1\}$ be the stopping loss.

1. Holding the other parameters at their defaults, solve the model for values of c between 0.5 and 5, and plot $B(c)$ and $A(c)$.
2. Show that the continuation region is empty if and only if $c \geq \bar{c} = \max_\pi\{S(\pi) - \mathbb{E}[S(\pi')]\}$, the largest expected reduction in the stopping loss that one more draw delivers.
3. Compute $\bar{c}$ and compare it with the values of c at which your solutions in step 1 have an empty continuation region.

Solution

Suppose that $c < S(\pi) - \mathbb{E}[S(\pi')]$ for some π.

Since $J \leq S$, $h(\pi) = c + \mathbb{E}[J(\pi')] \leq c + \mathbb{E}[S(\pi')] < S(\pi)$, so drawing again is optimal at π.

Suppose instead that $c \geq S(\pi) - \mathbb{E}[S(\pi')]$ for every π.

Let p_n be the posterior probability of f_1 after n draws, with $p_0 = \pi_{-1}$.

Then $\mathbb{E}[c(n+1) + S(p_{n+1})|p_n] \geq cn + S(p_n)$, so $cn + S(p_n)$ is a submartingale under the decision maker's beliefs.

By the optional stopping theorem, $\mathbb{E}[cN + S(p_N)] \geq S(p_0)$ for every stopping rule N with finite mean.

No rule improves on stopping immediately, so the continuation region is empty.

```
def one_draw_value(wf):
    """S(π) - E[S(π')] on the grid, using the Monte Carlo draws in wf."""
    S = lambda π: np.minimum(π * wf.L0, (1 - π) * wf.L1)
    values = np.empty(len(wf.π_grid))
    for i, π in enumerate(wf.π_grid):
        ES = ((1 - π) * S(wf.κ(wf.z0, π)).mean()
              + π * S(wf.κ(wf.z1, π)).mean())
        values[i] = S(π) - ES
    return values

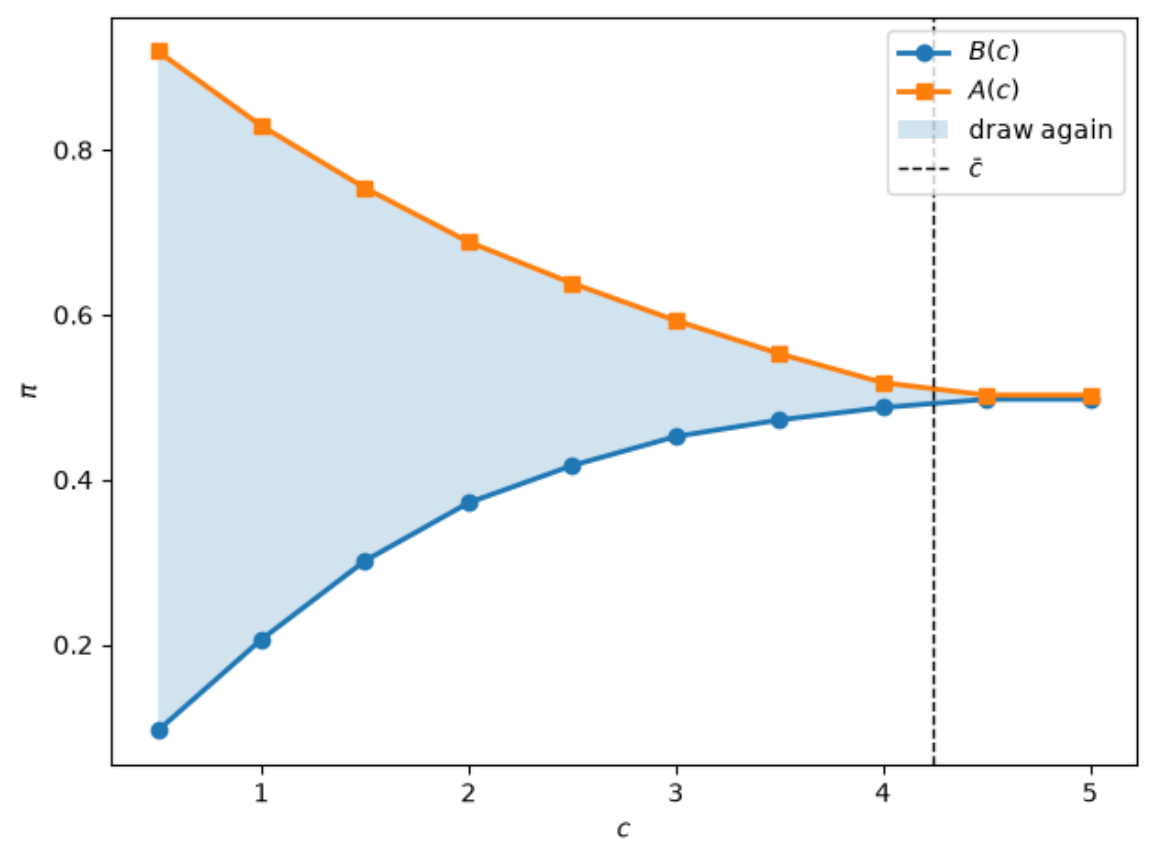
c_bar = one_draw_value(WaldFriedman()).max()
print(f"c̄ = {c_bar:.2f}")

c_grid = np.linspace(0.5, 5, 10)
B_c, A_c, empty = [], [], []
for c in c_grid:
    wf_c = WaldFriedman(c=c)
    h_c = solve_model(wf_c)
    B_val, A_val = find_cutoff_rule(wf_c, h_c)
    S_grid = np.minimum(wf_c.π_grid * wf_c.L0, (1 - wf_c.π_grid) * wf_c.L1)
    B_c.append(B_val)
    A_c.append(A_val)
    empty.append(not np.any(h_c < S_grid))

print("c values with an empty continuation region:",
      np.round(c_grid[np.array(empty)], 2))

fig, ax = plt.subplots()
ax.plot(c_grid, B_c, 'o-', lw=2, label='$B(c)$')
ax.plot(c_grid, A_c, 's-', lw=2, label='$A(c)$')
ax.fill_between(c_grid, B_c, A_c, alpha=0.2, label='draw again')
ax.axvline(c_bar, color='black', linestyle='--', lw=1, label=r'$\bar c$')
ax.set_xlabel('$c$')
ax.set_ylabel(r'$\pi$')
ax.legend(loc='upper right')
plt.tight_layout()
plt.show()
```

```
c̄ = 4.24
c values with an empty continuation region: [4.5 5. ]
```



As c rises, $B(c)$ rises and $A(c)$ falls, and the continuation region shrinks.

The region vanishes at $\bar{c} = 4.24$, far below the bound $L_0 L_1/(L_0 + L_1) = 12.5$ that we derived from $h \geq c$.

When the region is empty, `find_cutoff_rule` returns the two grid points next to $\hat{\pi} = 0.5$.

Exercise 27.5.3

Set (L_0, L_1) to $(50, 25)$, $(25, 25)$, and $(25, 50)$ in turn, holding the other parameters at their defaults.

For each setting, compute B, A, B_Λ, and A_Λ with $\pi_{-1} = 0.5$, and estimate α, β, and the mean stopping times under f_0 and under f_1 by simulation.

Explain how the thresholds and the error probabilities respond to L_0 and L_1.

Solution

We reuse `simulate_many` from the solution to Exercise 27.5.1.

```
for L0, L1 in [(50, 25), (25, 25), (25, 50)]:
    wf_L = WaldFriedman(L0=L0, L1=L1)
    h_L = solve_model(wf_L)
```

```
    B_val, A_val = find_cutoff_rule(wf_L, h_L)
    B_Λ_val, A_Λ_val = LR_thresholds(B_val, A_val)
    correct_0, _, N0 = simulate_many(wf_L, h_L, 20_000, "f0", seed=1)
    correct_1, _, N1 = simulate_many(wf_L, h_L, 20_000, "f1", seed=2)
    print(f"L0 = {L0}, L1 = {L1}: B = {B_val:.3f}, A = {A_val:.3f}, "
          f"B_Λ = {B_Λ_val:.2f}, A_Λ = {A_Λ_val:.2f}, "
          f"α = {1 - correct_0.mean():.3f}, β = {1 - correct_1.mean():.3f}, "
          f"E_0[N] = {N0.mean():.2f}, E_1[N] = {N1.mean():.2f}")
```

```
L0 = 50, L1 = 25: B = 0.116, A = 0.774, B_Λ = 0.13, A_Λ = 3.42, α = 0.214, β =␣
↪0.035, E_0[N] = 3.68, E_1[N] = 3.99
L0 = 25, L1 = 25: B = 0.256, A = 0.794, B_Λ = 0.34, A_Λ = 3.85, α = 0.176, β =␣
↪0.095, E_0[N] = 2.88, E_1[N] = 3.84
L0 = 25, L1 = 50: B = 0.302, A = 0.899, B_Λ = 0.43, A_Λ = 8.95, α = 0.070, β =␣
↪0.136, E_0[N] = 3.16, E_1[N] = 5.73
```

Raising L_1 from 25 to 50 makes accepting f_1 when f_0 is true costlier.

The decision maker requires more evidence before accepting f_1: A rises from 0.794 to 0.899, and A_Λ rises from 3.85 to 8.95.

He also accepts f_0 on weaker evidence: B rises from 0.256 to 0.302, and B_Λ rises from 0.34 to 0.43.

The type I error probability α falls from 0.176 to 0.070, and the type II error probability β rises from 0.095 to 0.136.

Raising L_0 from 25 to 50 works in the opposite direction.

B falls from 0.256 to 0.116 and A falls from 0.794 to 0.774, so β falls to 0.035 and α rises to 0.214.

In both cases the mean stopping times rise because a larger loss from one kind of error makes evidence more valuable.

CHAPTER

TWENTYEIGHT

EXCHANGEABILITY AND BAYESIAN UPDATING

Contents

- *Exchangeability and Bayesian Updating*
  - *Overview*
  - *Independently and identically distributed*
  - *A setting in which past observations are informative*
  - *Relationship between IID and exchangeable*
  - *Exchangeability*
  - *Bayes' Law*
  - *More details about Bayesian updating*
  - *Appendix*
  - *Sequels*
  - *Exercises*

28.1 Overview

This lecture studies learning via Bayes' Law.

We touch foundations of Bayesian statistical inference invented by de Finetti [1937].

The relevance of de Finetti's work for economists is presented forcefully by David Kreps in chapter 11 of [Kreps, 1988].

An example that we study below is a key component of *Job Search IX: Search with Learning*.

That lecture augments the classic job search model of McCall [McCall, 1970], studied in *Job Search I: The McCall Search Model*, by presenting an unemployed worker with a statistical inference problem.

Here we create graphs that illustrate the role that a likelihood ratio plays in Bayes' Law.

We'll use such graphs to provide insights into the mechanics driving outcomes in *Job Search IX: Search with Learning*.

Among other things, this lecture discusses connections between the statistical concepts of sequences of random variables that are

- independently and identically distributed

- exchangeable

Understanding these concepts is essential for appreciating how Bayesian updating works.

You can read about exchangeability here.

De Finetti's theorem states that an infinite exchangeable sequence is **conditionally independently and identically distributed**: it is IID conditional on a random variable whose value the observer does not know.[1]

We want to convey an answer to the question *conditional on what?*

We also tell why an assumption of independence precludes learning while an assumption that a sequence is conditionally IID makes learning possible.

Below, we'll often use

- W to denote a random variable
- w to denote a particular realization of a random variable W

Let's start with some imports:

```
from math import gamma

import matplotlib.pyplot as plt
import numpy as np
import scipy.optimize as op
from scipy.integrate import quad
from scipy.special import expit
```

28.2 Independently and identically distributed

We begin by looking at the notion of an **independently and identically distributed sequence** of random variables.

An independently and identically distributed sequence is often abbreviated as IID.

Two notions are involved

- **independence**
- **identically distributed**

A sequence $W_0, W_1, \ldots$ is **independently distributed** if, for every $n \geq 0$, the joint probability density of $W_0, W_1, \ldots, W_n$ is the **product** of the marginal densities of $W_0, W_1, \ldots, W_n$.

The sequence $W_0, W_1, \ldots$ is **independently and identically distributed** (IID) if in addition the marginal density of W_t is the same for all $t = 0, 1, \ldots$.

For example, let $p(W_0, W_1, \ldots, W_n)$ be the **joint density** of $W_0, W_1, \ldots, W_n$ and let $p(W_t)$ be the common **marginal density** of each W_t.

Then the sequence $W_0, W_1, \ldots$ is IID if

$$p(W_0, W_1, \ldots, W_n) = p(W_0)p(W_1)\cdots p(W_n) \quad \text{for every } n \geq 0$$

so that each joint density is the product of identical marginal densities.

[1] Conditional independence alone does not imply exchangeability. Let $W_t = t\theta + \varepsilon_t$, where θ is a random variable and $\{\varepsilon_t\}$ is an IID $N(0,1)$ sequence independent of θ. Conditional on θ, the W_t are independent, but they are not identically distributed, and the sequence $\{W_t\}$ is not exchangeable.

28.2.1 IID means past observations don't tell us anything about future observations

If a sequence of random variables is IID, past information provides no information about future realizations.

Therefore, there is **nothing to learn** from the past about the future.

To understand these statements, let the joint distribution of a sequence of random variables $\{W_t\}_{t=0}^T$ that is not necessarily IID be

$$p(W_T, W_{T-1}, \dots, W_1, W_0)$$

Using the laws of probability, we can always factor such a joint density into a product of conditional densities:

$$\begin{aligned} p(W_T, W_{T-1}, \dots, W_1, W_0) =& p(W_T|W_{T-1}, \dots, W_0)p(W_{T-1}|W_{T-2}, \dots, W_0)\cdots \\ & \cdots p(W_1|W_0)p(W_0) \end{aligned}$$

In general,

$$p(W_t|W_{t-1}, \dots, W_0) \neq p(W_t)$$

which states that the **conditional density** on the left side does not equal the **marginal density** on the right side.

But in the special IID case,

$$p(W_t|W_{t-1}, \dots, W_0) = p(W_t),$$

so that the partial history $W_{t-1}, \dots, W_0$ contains no information about the probability of W_t.

So in the IID case, there is **nothing to learn** about the densities of future random variables from past random variables.

But when the sequence is not IID, there is something to learn about the future from observations of past random variables.

We turn next to an instance of the general case in which the sequence is not IID.

Please watch for what can be learned from the past and when.

28.3 A setting in which past observations are informative

Let $\{W_t\}_{t=0}^\infty$ be a sequence of nonnegative scalar random variables with a joint probability distribution constructed as follows.

There are two distinct cumulative distribution functions F and G that have densities f and g, respectively, for a nonnegative scalar random variable W.

Before the start of time, say at time $t = -1$, "nature" once and for all selects **either** f **or** g.

Thereafter at each time $t \geq 0$, nature draws a random variable W_t from the selected distribution.

So the data are permanently generated as independently and identically distributed (IID) draws from **either** F **or** G.

We could say that *objectively*, meaning *after* nature has chosen either F or G, the probability that the data are generated as draws from F is either 0 or 1.

We now drop into this setting a partially informed decision maker who

- knows both F and G, but
- does not know whether at $t = -1$ nature had drawn F or whether nature had drawn G once-and-for-all

Thus, although our decision maker knows F and knows G, he does not know which of these two known distributions nature had selected to draw from.

The decision maker describes his ignorance with a **subjective probability** $\tilde{\pi}$ and reasons as if nature had selected F with probability $\tilde{\pi} \in (0,1)$ and G with probability $1-\tilde{\pi}$.

Thus, we assume that the decision maker

- **knows** both F and G
- **doesn't know** which of these two distributions that nature has drawn
- expresses his ignorance by **acting as if** or **thinking that** nature chose distribution F with probability $\tilde{\pi} \in (0,1)$ and distribution G with probability $1-\tilde{\pi}$
- at date $t \geq 0$ knows the partial history $w_t, w_{t-1}, \ldots, w_0$

To proceed, we want to know the decision maker's belief about the joint distribution of the partial history.

We'll discuss that next and in the process describe the concept of **exchangeability**.

28.4 Relationship between IID and exchangeable

Conditional on nature selecting F, the joint density of $W_0, W_1, \ldots, W_n$ is

$$f(W_0)f(W_1)\cdots f(W_n)$$

for every $n \geq 0$.

Conditional on nature selecting G, the joint density of $W_0, W_1, \ldots, W_n$ is

$$g(W_0)g(W_1)\cdots g(W_n)$$

Thus, **conditional on nature having selected** F, the sequence $W_0, W_1, \ldots$ is independently and identically distributed.

Furthermore, **conditional on nature having selected** G, the sequence $W_0, W_1, \ldots$ is also independently and identically distributed.

But what about the **unconditional distribution** of a partial history?

The unconditional joint density of $W_0, W_1, \ldots, W_n$ is evidently

$$h(W_0, W_1, \ldots, W_n) \equiv \tilde{\pi}[f(W_0)f(W_1)\cdots f(W_n)] + (1-\tilde{\pi})[g(W_0)g(W_1)\cdots g(W_n)], \quad n \geq 0 \tag{28.1}$$

Under the unconditional distribution (28.1), the sequence $W_0, W_1, \ldots$ is **not** independently and identically distributed.

To verify this claim, let $h(W_t) = \tilde{\pi}f(W_t) + (1-\tilde{\pi})g(W_t)$ be the marginal density of W_t.

Then

$$h(W_0, W_1) - h(W_0)h(W_1) = \tilde{\pi}(1-\tilde{\pi})(f(W_0) - g(W_0))(f(W_1) - g(W_1)) \tag{28.2}$$

The right side is not zero when $f(W_0) \neq g(W_0)$ and $f(W_1) \neq g(W_1)$, so the joint density $h(W_0, W_1)$ is not the product of the marginal densities $h(W_0)$ and $h(W_1)$.

Thus, the conditional distribution

$$h(W_1|W_0) \equiv \frac{h(W_0, W_1)}{(\tilde{\pi}f(W_0) + (1-\tilde{\pi})g(W_0))} \neq (\tilde{\pi}f(W_1) + (1-\tilde{\pi})g(W_1))$$

This means that random variable W_0 contains information about random variable W_1.

So there is something to learn from the past about the future.

Exercise Exercise 28.10.1 uses (28.2) to show that W_0 and W_1 are positively correlated.

28.5 Exchangeability

While the sequence $W_0, W_1, \ldots$ is not IID, it can be verified that it is **exchangeable**, which means that the joint distributions $h(W_0, W_1)$ and $h(W_1, W_0)$ of the "re-ordered" sequences satisfy

$$h(W_0, W_1) = h(W_1, W_0)$$

and so on.

More generally, a sequence of random variables is said to be **exchangeable** if the joint probability distribution for a sequence does not change when the positions in the sequence in which finitely many of random variables appear are altered.

Equation (28.1) represents our instance of an exchangeable joint density over a sequence of random variables as a **mixture** of two IID joint densities over a sequence of random variables.

A Bayesian statistician interprets the mixing parameter $\tilde{\pi} \in (0, 1)$ as a decision maker's subjective belief – the decision maker's **prior probability** – that nature had selected probability distribution F.

28.5.1 De Finetti's theorem

As de Finetti [1937] proved, every infinite exchangeable sequence is a mixture of IID sequences.

For binary random variables $W_t \in \{0, 1\}$, de Finetti's theorem states that the infinite sequence $W_0, W_1, \ldots$ is exchangeable if and only if there is a probability measure μ on $[0, 1]$ such that

$$\text{Prob}(W_0 = w_0, \ldots, W_{n-1} = w_{n-1}) = \int_0^1 \theta^k (1-\theta)^{n-k} \mu(d\theta), \quad k = \sum_{t=0}^{n-1} w_t \tag{28.3}$$

for every $n \geq 1$ and every $(w_0, \ldots, w_{n-1}) \in \{0, 1\}^n$.

The mixing measure μ is unique.

It is the probability distribution of the limiting relative frequency

$$\theta = \lim_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} W_t,$$

a limit that exists with probability one.

Conditional on θ, the W_t are IID Bernoulli random variables with success probability θ.

This answers the question *conditional on what?*: conditional on the limiting relative frequency θ.

A Bayesian statistician interprets μ as a prior probability distribution over the unknown Bernoulli parameter $\theta \in [0, 1]$.

The measure μ need not have a density; it can put positive probability on finitely many points.

Hewitt and Savage [1955] extended de Finetti's theorem to sequences that take values in more general spaces.

In their version the mixing measure is a probability distribution over probability distributions of W, and the conditioning variable is the limiting empirical distribution of $W_0, W_1, \ldots$.

Equation (28.1) is an instance with a two-point mixing measure that puts probability $\tilde{\pi}$ on F and probability $1 - \tilde{\pi}$ on G.

In that instance the limiting empirical distribution of $W_0, W_1, \ldots$ is F when nature has drawn F and G when nature has drawn G.

De Finetti's theorem requires an infinite sequence.

Exercise Exercise 28.10.1 describes a finite exchangeable sequence that is not a mixture of IID sequences.

The Beta–Bernoulli model in *Merging of Opinions: The Blackwell–Dubins Theorem* is the instance of (28.3) in which μ is a Beta distribution.

Exercise Exercise 28.10.3 describes an urn scheme that generates it.

28.6 Bayes' Law

We noted above that in our example model there is something to learn about the future from past data drawn from our particular instance of a process that is exchangeable but not IID.

But how can we learn?

And about what?

The answer to the *about what* question is: about which distribution, F or G, nature selected at time -1.

The decision maker knows his prior probability $\tilde{\pi}$.

What he learns is expressed in a posterior probability that moves away from $\tilde{\pi}$ as data arrive.

The answer to the *how* question is to use Bayes' Law.

Another way to say *use Bayes' Law* is to say *from a (subjective) joint distribution, compute an appropriate conditional distribution*.

Let's dive into Bayes' Law in this context.

Let q represent the distribution that nature actually draws w from and let

$$\pi = \mathbb{P}\{q = f\}$$

where we regard π as a decision maker's **subjective probability** (also called a **personal probability**).

Suppose that at $t \geq 0$, the decision maker has observed a history $w^t \equiv [w_t, w_{t-1}, \ldots, w_0]$.

We let

$$\pi_t = \mathbb{P}\{q = f | w^t\}$$

where we adopt the convention

$$\pi_{-1} = \tilde{\pi}$$

Conditional on q, the draw w_{t+1} is independent of the history w^t.

The density of w_{t+1} conditional on w^t is then

$$\pi_t f + (1 - \pi_t) g.$$

Bayes' rule for updating π_{t+1} is

$$\pi_{t+1} = \frac{\pi_t f(w_{t+1})}{\pi_t f(w_{t+1}) + (1 - \pi_t) g(w_{t+1})} \tag{28.4}$$

Equation (28.4) follows from Bayes' rule, which tells us that

$$\mathbb{P}\{q = f \,|\, w^{t+1}\} = \frac{p(w_{t+1} \,|\, q = f, w^t)\, \mathbb{P}\{q = f \,|\, w^t\}}{p(w_{t+1} \,|\, w^t)}$$

where p denotes a density and

$$p(w_{t+1} \,|\, w^t) = \sum_{a \in \{f, g\}} p(w_{t+1} \,|\, q = a, w^t)\, \mathbb{P}\{q = a \,|\, w^t\}.$$

Because w_{t+1} is independent of w^t conditional on q, $p(w_{t+1} \,|\, q = f, w^t) = f(w_{t+1})$ and $p(w_{t+1} \,|\, q = g, w^t) = g(w_{t+1})$.

28.7 More details about Bayesian updating

Let's stare at and rearrange Bayes' Law as represented in equation (28.4) with the aim of understanding how the **posterior** probability π_{t+1} is influenced by the **prior** probability π_t and the **likelihood ratio**

$$\ell(w) = \frac{f(w)}{g(w)}$$

It is convenient for us to rewrite the updating rule (28.4) as

$$\pi_{t+1} = \frac{\pi_t f\left(w_{t+1}\right)}{\pi_t f\left(w_{t+1}\right) + \left(1-\pi_t\right) g\left(w_{t+1}\right)} = \frac{\pi_t \frac{f(w_{t+1})}{g(w_{t+1})}}{\pi_t \frac{f(w_{t+1})}{g(w_{t+1})} + \left(1-\pi_t\right)} = \frac{\pi_t \ell\left(w_{t+1}\right)}{\pi_t \ell\left(w_{t+1}\right) + \left(1-\pi_t\right)}$$

This implies that

$$\frac{\pi_{t+1}}{\pi_t} = \frac{\ell\left(w_{t+1}\right)}{\pi_t \ell\left(w_{t+1}\right) + \left(1-\pi_t\right)} \begin{cases} > 1 & \text{if } \ell\left(w_{t+1}\right) > 1 \\ \leq 1 & \text{if } \ell\left(w_{t+1}\right) \leq 1 \end{cases} \tag{28.5}$$

The likelihood ratio alone determines whether an observation w_{t+1} leads the decision maker to increase or decrease the subjective probability he/she attaches to distribution F.

When the likelihood ratio $\ell(w_{t+1})$ exceeds one, the observation w_{t+1} nudges the probability π put on distribution F upward, and when the likelihood ratio $\ell(w_{t+1})$ is less than one, the observation w_{t+1} nudges π downward.

The prior π_t affects the size of the move:

$$\pi_{t+1} - \pi_t = \frac{\pi_t(1-\pi_t)(\ell(w_{t+1}) - 1)}{\pi_t \ell(w_{t+1}) + 1 - \pi_t}$$

The move shrinks to zero as π_t approaches 0 or 1.

Representation (28.5) is the foundation of some graphs that we'll use to display the dynamics of $\{\pi_t\}_{t=0}^{\infty}$ that are induced by Bayes' Law.

We'll plot $\ell\left(w\right)$ as a way to enlighten us about how learning – i.e., Bayesian updating of the probability π that nature has chosen distribution f – works.

We build up the picture in three steps, each of which produces one graph.

All three are built from the same ingredients: the densities f and g, and the values of w at which the likelihood ratio $\ell(w) = f(w)/g(w)$ crosses one.

Both f and g are Beta densities, so we start with the general Beta density.

```
def p(w, a, b):
    "The Beta density with parameters a and b."
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * w**(a - 1) * (1 - w)**(b - 1)
```

The next function assembles the ingredients for a given pair of Beta distributions.

```
def create_model(F_a=1, F_b=1, G_a=3, G_b=1.2):
    """
    Build the densities f and g, along with the two values of w at which
    the likelihood ratio ℓ(w) = f(w) / g(w) equals one.

    The root finder assumes that f is uniform (F_a = F_b = 1) and that
    G_a > 1 and G_b > 1.  Then g vanishes at 0 and 1 and has a single
```

(continues on next page)

(continued from previous page)

```
    interior mode at which g exceeds one, so ℓ(w) = 1 / g(w) equals one
    exactly once on each side of the mode of g.
    """
    f = lambda w: p(w, F_a, F_b)
    g = lambda w: p(w, G_a, G_b)

    # The mode of g divides [0, 1] into two intervals, each holding one root
    G_mode = (G_a - 1) / (G_a + G_b - 2)
    obj = lambda w: f(w) / g(w) - 1
    roots = np.array([op.root_scalar(obj, bracket=[1e-10, G_mode]).root,
                      op.root_scalar(obj, bracket=[G_mode, 1 - 1e-10]).root])
    return f, g, roots
```

28.7.1 The likelihood ratio

Our first graph plots the likelihood ratio $\ell(w)$ on the abscissa axis against w on the ordinate axis.

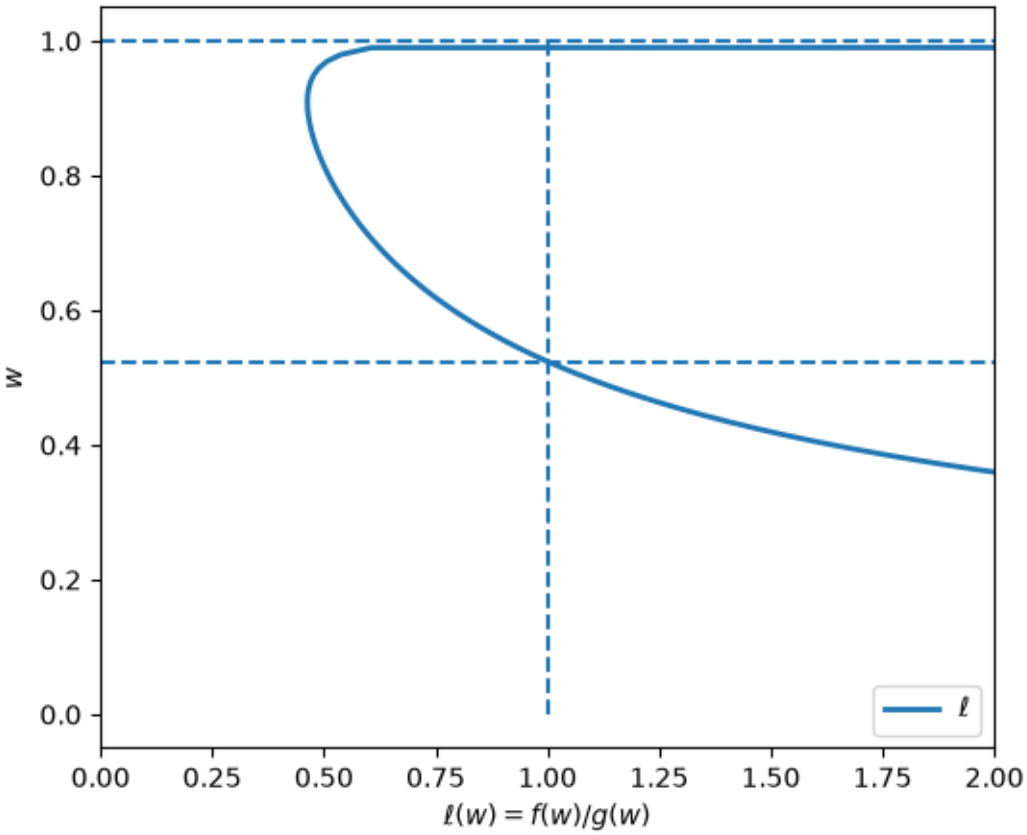
We orient it this way so that w shares an axis with the two graphs that follow.

```
def plot_likelihood_ratio(F_a=1, F_b=1, G_a=3, G_b=1.2):
    f, g, roots = create_model(F_a, F_b, G_a, G_b)
    w_grid = np.linspace(1e-12, 1 - 1e-12, 100)

    fig, ax = plt.subplots(figsize=(6, 5))
    ax.plot(f(w_grid) / g(w_grid), w_grid, label=r'$\ell$', lw=2)
    ax.vlines(1, 0, 1, linestyle='--')
    ax.hlines(roots, 0, 2, linestyle='--')
    ax.set_xlim(0, 2)
    ax.legend(loc=4)
    ax.set(xlabel=r'$\ell(w) = f(w) / g(w)$', ylabel='$w$')
    plt.show()
```

We begin with f uniform on $[0, 1]$ — that is, Beta with $F_a = 1, F_b = 1$ — and g Beta with $G_a = 3, G_b = 1.2$.

```
plot_likelihood_ratio()
```



The two horizontal dashed lines mark the values of w at which $\ell(w) = 1$.

Between them the likelihood ratio is below one, so by (28.5) a draw in that region pushes π down.

Outside them it exceeds one, so a draw there pushes π up.

28.7.2 The densities and the probabilities of moving in each direction

Our second graph plots $f(w)$ and $g(w)$ against w, and shades the regions delineated by those same two values of w.

Shading them lets us attach a probability to each direction of movement, which we compute by integrating the relevant density over the region.

```
def plot_densities(F_a=1, F_b=1, G_a=3, G_b=1.2):
    f, g, roots = create_model(F_a, F_b, G_a, G_b)
    w_grid = np.linspace(0, 1, 100)

    fig, ax = plt.subplots(figsize=(6, 5))
    ax.plot(f(w_grid), w_grid, label='$f$', lw=2)
    ax.plot(g(w_grid), w_grid, label='$g$', lw=2)
    ax.vlines(1, 0, 1, linestyle='--')
    ax.hlines(roots, 0, 2, linestyle='--')
    ax.legend(loc=4)
    ax.set(xlabel='$f(w), g(w)$', ylabel='$w$')

    # Probability of landing in each region, under whichever density is shaded
    area_lower = quad(f, 0, roots[0])[0]
```

(continues on next page)

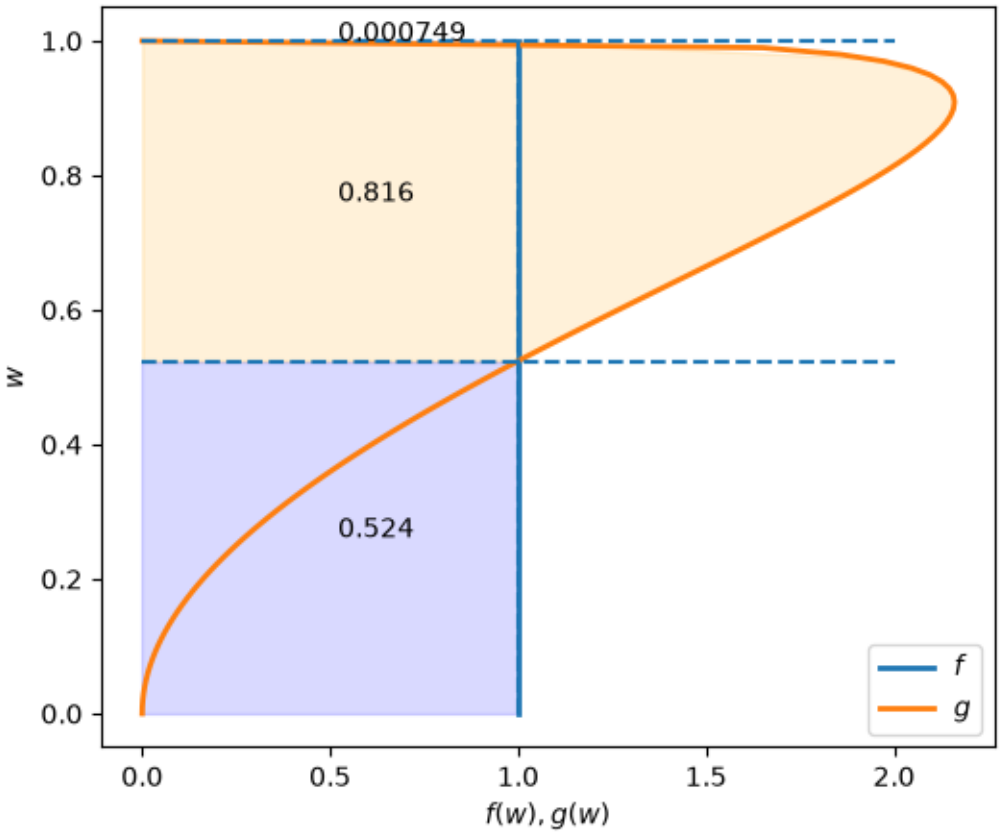
(continued from previous page)

```
    area_middle = quad(g, roots[0], roots[1])[0]
    area_upper = quad(f, roots[1], 1)[0]

    ax.fill_between([0, 1], 0, roots[0], color='blue', alpha=0.15)
    ax.text((f(0) + f(roots[0])) / 4, roots[0] / 2, f"{area_lower: .3g}")
    w_middle = np.linspace(roots[0], roots[1], 20)
    ax.fill_betweenx(w_middle, 0, g(w_middle), color='orange', alpha=0.15)
    ax.text(np.mean(g(roots)) / 2, np.mean(roots), f"{area_middle: .3g}")
    ax.fill_between([0, 1], roots[1], 1, color='blue', alpha=0.15)
    ax.text((f(roots[1]) + f(1)) / 4, (roots[1] + 1) / 2, f"{area_upper: .3g}")
    plt.show()
```

Let's look at the same pair of distributions as before.

```
plot_densities()
```



The fractions in the colored areas are probabilities that a realization of w falls into the region beside them.

The blue regions are integrals of f and the orange region is an integral of g.

For example, under true distribution F, π will be updated toward 0 if w falls into the interval $[0.524, 0.999]$, which occurs with probability $0.999 - 0.524 \approx 0.475$ under F.

But this would occur with probability 0.816 if G were the true distribution.

28.7.3 Dynamics of the belief

Our third graph attaches an arrow to each point in the (π, w) plane showing the change in π that Bayes' Law induces when the current belief is π and the new draw is w.

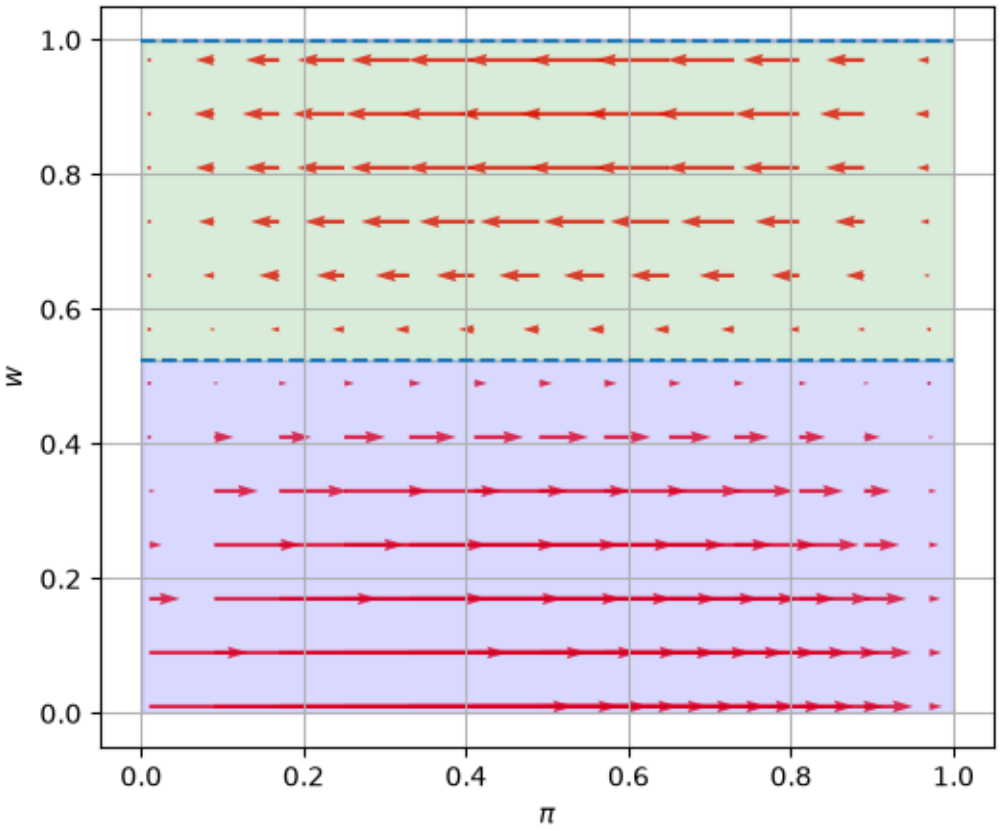
```
def plot_belief_dynamics(F_a=1, F_b=1, G_a=3, G_b=1.2):
    f, g, roots = create_model(F_a, F_b, G_a, G_b)
    π_grid = np.linspace(1e-3, 1 - 1e-3, 100)

    # Change in π at each point of a coarse grid over (π, w)
    W = np.arange(0.01, 0.99, 0.08)
    Π = np.arange(0.01, 0.99, 0.08)
    lw = (f(W) / g(W))[:, None]      # likelihood ratio at each w, as a column
    ΔΠ = Π * (lw / (Π * lw + 1 - Π) - 1)
    ΔW = np.zeros_like(ΔΠ)

    fig, ax = plt.subplots(figsize=(6, 5))
    ax.quiver(Π, W, ΔΠ, ΔW, scale=2, color='r', alpha=0.8)
    ax.fill_between(π_grid, 0, roots[0], color='blue', alpha=0.15)
    ax.fill_between(π_grid, roots[0], roots[1], color='green', alpha=0.15)
    ax.fill_between(π_grid, roots[1], 1, color='blue', alpha=0.15)
    ax.hlines(roots, 0, 1, linestyle='--')
    ax.set(xlabel=r'$\pi$', ylabel='$w$')
    ax.grid()
    plt.show()
```

Again we use the same pair of distributions.

```
plot_belief_dynamics()
```



Arrows pointing right show when Bayes' Law makes π increase and arrows pointing left show when it makes π decrease.

Lengths of the arrows show magnitudes of the force from Bayes' Law impelling π to change.

These lengths depend on both the prior probability π on the abscissa axis and the evidence in the form of the current draw of w on the ordinate axis.

Arrows point right in the blue regions, where $\ell(w) > 1$, and left in the green region between the dashed lines, where $\ell(w) < 1$.

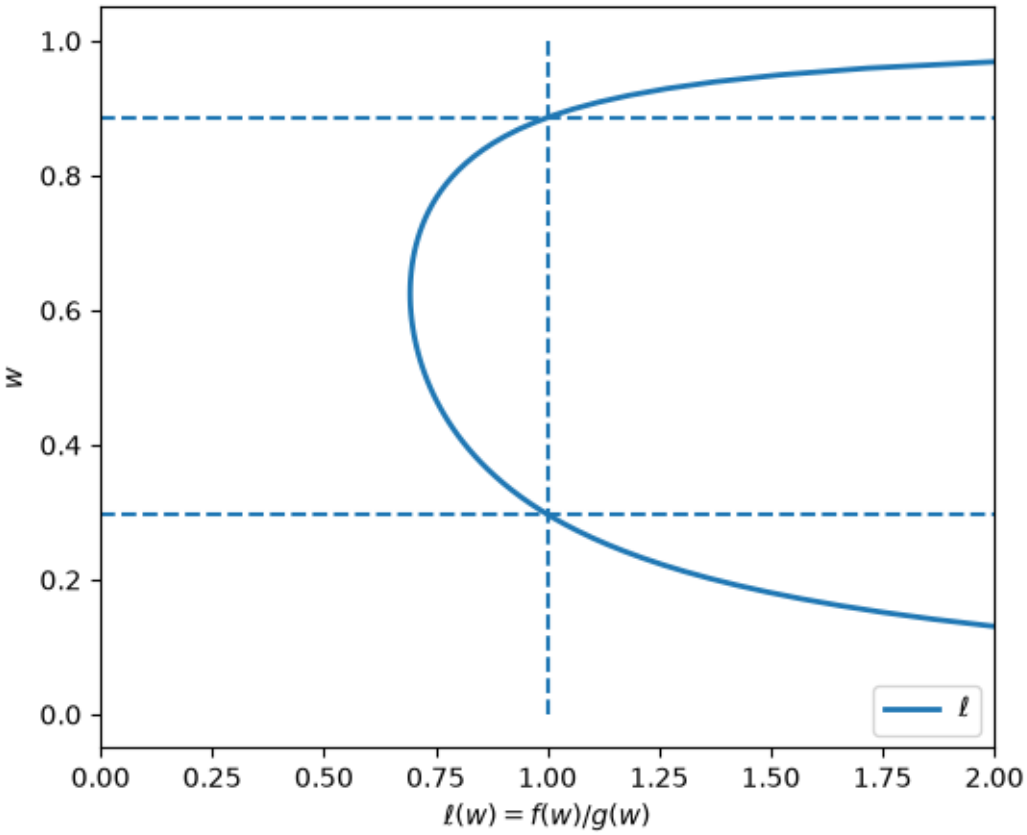
For these parameters the upper blue region is a thin sliver just below $w = 1$, since $\ell(w)$ returns above one only for draws very close to the upper limit.

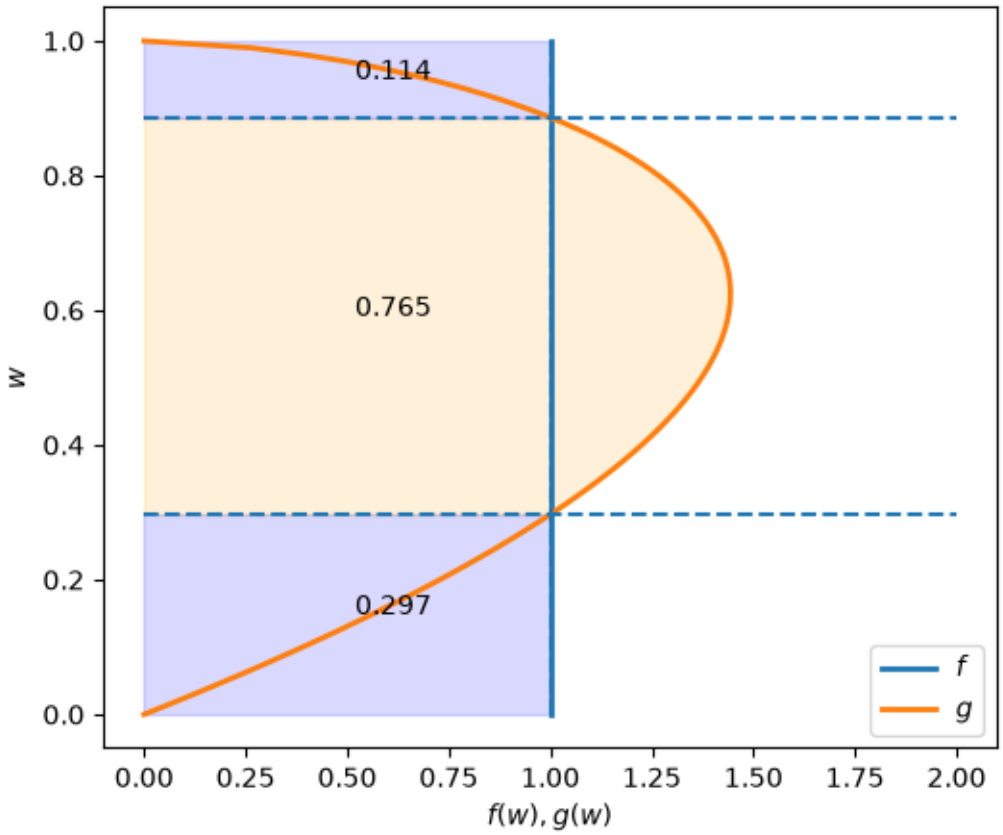
28.7.4 Another instance

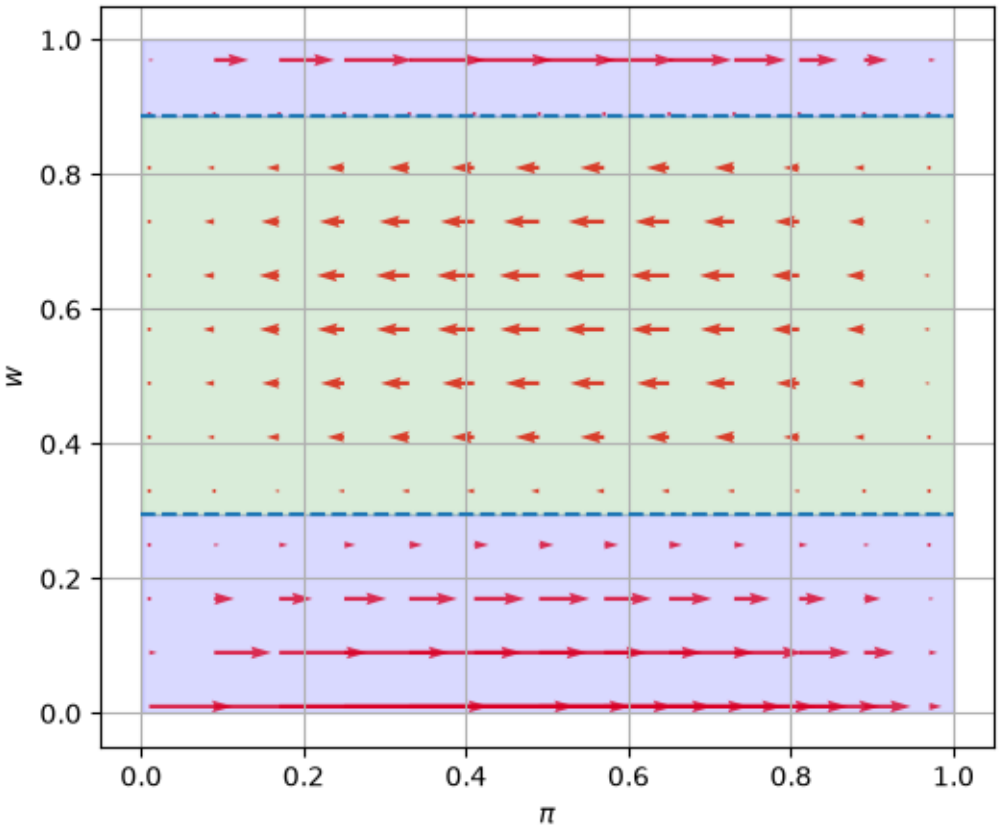
Next we use our code to create graphs for another instance of our model.

We keep F the same as in the preceding instance, namely a uniform distribution, but now assume that G is a Beta distribution with parameters $G_a = 2, G_b = 1.6$.

```
plot_likelihood_ratio(G_a=2, G_b=1.6)
plot_densities(G_a=2, G_b=1.6)
plot_belief_dynamics(G_a=2, G_b=1.6)
```


1.0
0.8
0.6
0.4
0.2
0.0
w
0.00
0.25
0.50
0.75
1.00
1.25
1.50
1.75
2.00
$\ell(w) = f(w)/g(w)$
ℓ


0.114
0.765
0.297
f
g
w
f(w), g(w)
1.0
0.8
0.6
0.4
0.2
0.0
0.00
0.25
0.50
0.75
1.00
1.25
1.50
1.75
2.00



Notice how the likelihood ratio, the densities, and the arrows compare with the previous instance of our example.

28.8 Appendix

28.8.1 Sample paths of π_t

Now we'll have some fun by plotting multiple realizations of sample paths of π_t under two possible assumptions about nature's choice of distribution, namely

- that nature permanently draws from F
- that nature permanently draws from G

Outcomes depend on a peculiar property of likelihood ratio processes discussed in *this section* of *Likelihood Ratio Processes*.

Before simulating, it pays to rewrite Bayes' Law in terms of **odds** rather than probabilities.

In terms of the odds $\pi/(1-\pi)$, the updating rule (28.4) becomes

$$\frac{\pi_{t+1}}{1-\pi_{t+1}} = \ell(w_{t+1})\frac{\pi_t}{1-\pi_t} \tag{28.6}$$

So Bayes' Law is *multiplicative* in the odds, and iterating (28.6) back to the prior π_{-1} yields

$$\frac{\pi_t}{1-\pi_t} = \frac{\pi_{-1}}{1-\pi_{-1}}\prod_{s=0}^{t}\ell(w_s) \tag{28.7}$$

The belief path is thus a running product of likelihood ratios, which is why properties of likelihood ratio processes govern its behavior.

Taking logarithms turns the product into a sum:

$$\log \frac{\pi_t}{1-\pi_t} = \log \frac{\pi_{-1}}{1-\pi_{-1}} + \sum_{s=0}^{t} \log \ell(w_s)$$

We simulate with this sum.

It avoids overflow in the product and loss of precision when π_t is close to 0 or 1.

It also means that we can simulate an entire ensemble of paths with a cumulative sum, with no iteration over dates.

```
def simulate(rng, a, b, T=50, N=1000, π_init=0.5,
             F_a=1, F_b=1, G_a=3, G_b=1.2):
    """
    Simulate N paths of the belief π when nature draws w_0, ..., w_{T-1}
    IID from Beta(a, b).  Returns an array of shape (N, T+1) whose column 0
    holds the prior π_{-1} = π_init and whose column t+1 holds π_t.
    """
    w = rng.beta(a, b, size=(N, T))
    log_l = np.log(p(w, F_a, F_b)) - np.log(p(w, G_a, G_b))
    log_odds = np.log(π_init / (1 - π_init)) + np.cumsum(log_l, axis=1)
    return np.column_stack([np.full(N, π_init), expit(log_odds)])
```

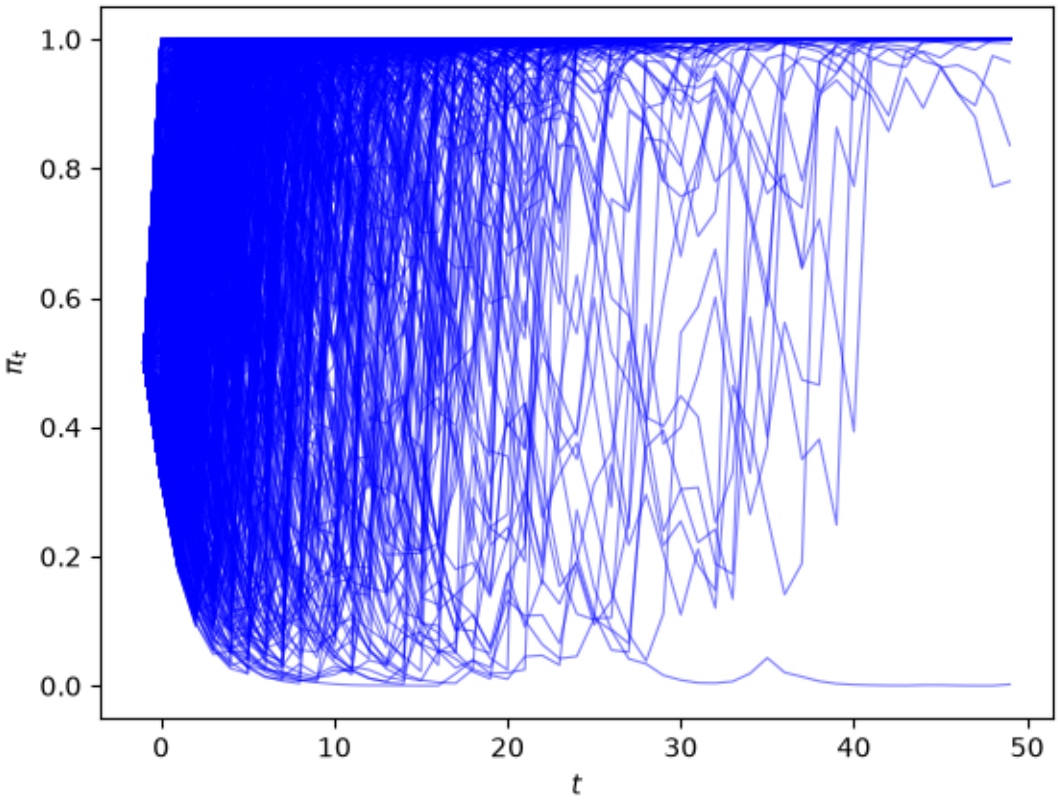
The paths are plotted by the next function.

```
def plot_paths(π_paths):
    "Plot belief paths against the dates t = -1, 0, ..., T-1."
    t = np.arange(π_paths.shape[1]) - 1
    fig, ax = plt.subplots()
    ax.plot(t, π_paths.T, color='b', lw=0.8, alpha=0.5)
    ax.set(xlabel='$t$', ylabel=r'$\pi_t$')
    plt.show()
```

We begin by generating N simulated $\{\pi_t\}$ paths with T periods when the sequence is truly IID draws from F. We set an initial prior $\pi_{-1} = .5$.

```
rng = np.random.default_rng(42)
T = 50
```

```
# when nature selects F
π_paths_F = simulate(rng, a=1, b=1, T=T, N=1000)
plot_paths(π_paths_F)
```

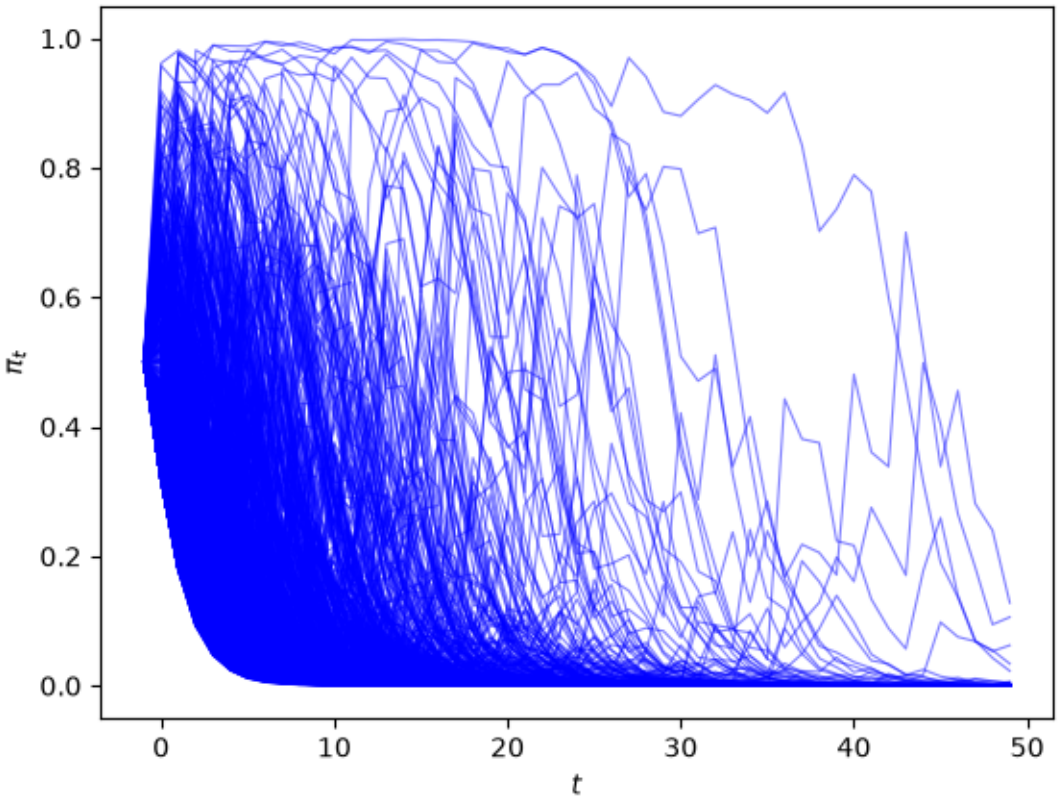


In the above example, for most paths $\pi_t \to 1$.

So Bayes' Law evidently eventually discovers the truth for most of our paths.

Next, we generate paths with T periods when the sequence is truly IID draws from G. Again, we set the initial prior $\pi_{-1} = .5$.

```
# when nature selects G
π_paths_G = simulate(rng, a=3, b=1.2, T=T, N=1000)
plot_paths(π_paths_G)
```



In the above graph we observe that now most paths $\pi_t \to 0$.

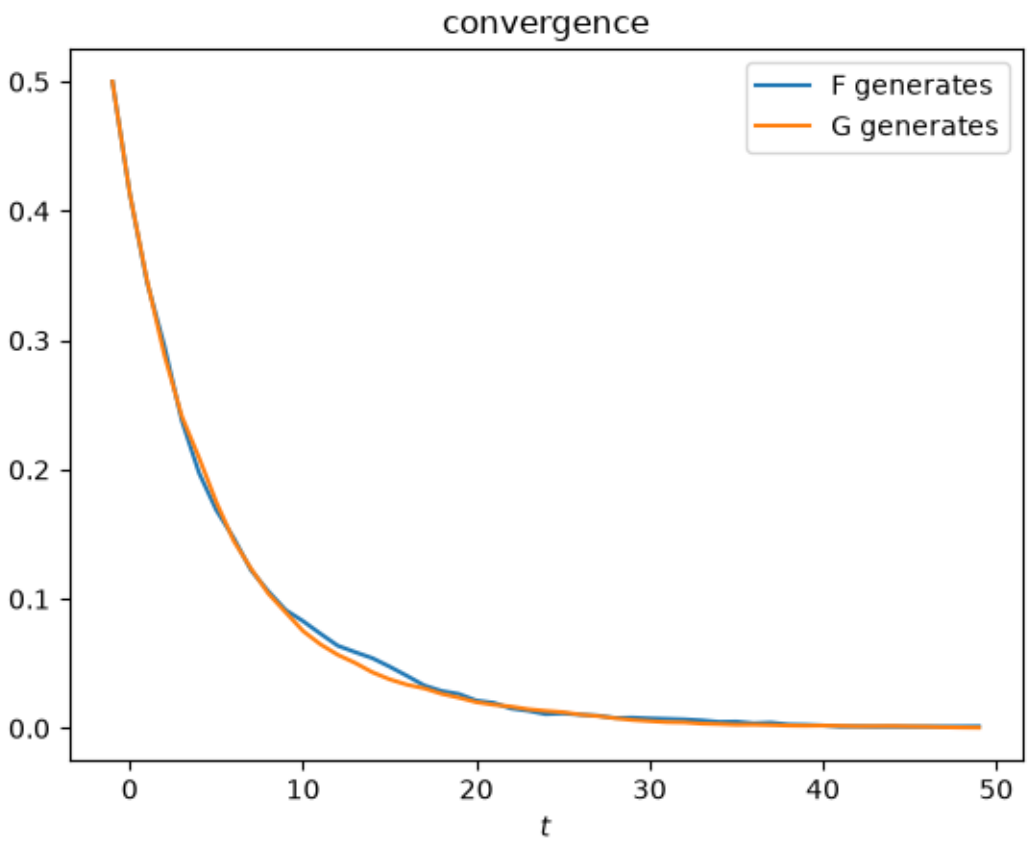
28.8.2 Rates of convergence

We study rates of convergence of π_t to 1 when nature generates the data as IID draws from F and of convergence of π_t to 0 when nature generates IID draws from G.

We do this by averaging across simulated paths of $\{\pi_t\}_{t=-1}^{T-1}$.

Using N simulated π_t paths, we compute $1 - \frac{1}{N}\sum_{i=1}^{N} \pi_{i,t}$ at each t when the data are generated as draws from F and compute $\frac{1}{N}\sum_{i=1}^{N} \pi_{i,t}$ when the data are generated as draws from G.

```
t_grid = np.arange(-1, T)
fig, ax = plt.subplots()
ax.plot(t_grid, 1 - np.mean(π_paths_F, axis=0), label='F generates')
ax.plot(t_grid, np.mean(π_paths_G, axis=0), label='G generates')
ax.set(xlabel='$t$', title='convergence')
ax.legend()
plt.show()
```



The two curves nearly coincide.

An exact identity explains why:

$$\pi_{-1} E_f[1 - \pi_t] = (1 - \pi_{-1}) E_g[\pi_t] = E[\pi_t (1 - \pi_t)] \tag{28.8}$$

Here E_f and E_g denote expectations under IID draws from F and from G, and E denotes the expectation under the decision maker's subjective joint distribution (28.1).

To prove (28.8), let ϕ be a bounded function of the history w^t.

Because $\pi_t = \mathbb{P}\{q = f \mid w^t\}$, the law of iterated expectations implies

$$E[\pi_t \phi(w^t)] = E[\mathbf{1}\{q = f\} \phi(w^t)] = \pi_{-1} E_f[\phi(w^t)].$$

Setting $\phi = 1 - \pi_t$ equates the first and third terms of (28.8).

The same argument applied to $1 - \pi_t = \mathbb{P}\{q = g \mid w^t\}$ with $\phi = \pi_t$ equates the second and third terms.

With $\pi_{-1} = 0.5$, the identity says that $E_f[1 - \pi_t] = E_g[\pi_t]$ for every t, so the two simulated curves differ only by sampling error.

With another prior, the two curves differ by the constant factor $(1 - \pi_{-1})/\pi_{-1}$.

The identity does not say that typical paths converge at the same rate under F and under G.

By (28.7) and the law of large numbers, $\frac{1}{t+1} \log \frac{\pi_t}{1 - \pi_t}$ converges to $E_f[\log \ell(w)] = KL(f, g) = 0.759$ when F generates the data and to $E_g[\log \ell(w)] = -KL(g, f) = -0.344$ when G generates the data, where $KL(f, g) = \int f(w) \log \frac{f(w)}{g(w)} dw$ is the Kullback–Leibler divergence of g from f.

So a typical path approaches 1 under F faster than it approaches 0 under G.

The averages plotted above are dominated by rare paths along which π_t leans the wrong way for a long time.

Both averages decay at the exponential rate

$$C(f,g) = -\log \min_{s\in[0,1]} \int f(w)^s g(w)^{1-s} dw = 0.121,$$

the Chernoff entropy of f and g described in *Statistical Divergence Measures*.

Exercise Exercise 28.10.2 verifies these claims.

28.8.3 Graph of ensemble dynamics of π_t

More insights about the dynamics of $\{\pi_t\}$ can be gleaned by computing conditional expectations of $\frac{\pi_{t+1}}{\pi_t}$ as functions of π_t via integration with respect to the pertinent probability distribution:

$$\begin{aligned} E\left[\frac{\pi_{t+1}}{\pi_t}\,\middle|\, q=a,\pi_t\right] &= E\left[\frac{\ell\left(w_{t+1}\right)}{\pi_t\ell\left(w_{t+1}\right)+\left(1-\pi_t\right)}\,\middle|\, q=a,\pi_t\right], \\ &= \int_0^1 \frac{\ell\left(w_{t+1}\right)}{\pi_t\ell\left(w_{t+1}\right)+\left(1-\pi_t\right)} a\left(w_{t+1}\right) dw_{t+1} \end{aligned}$$

where $a = f, g$.

The following code approximates the integral above:

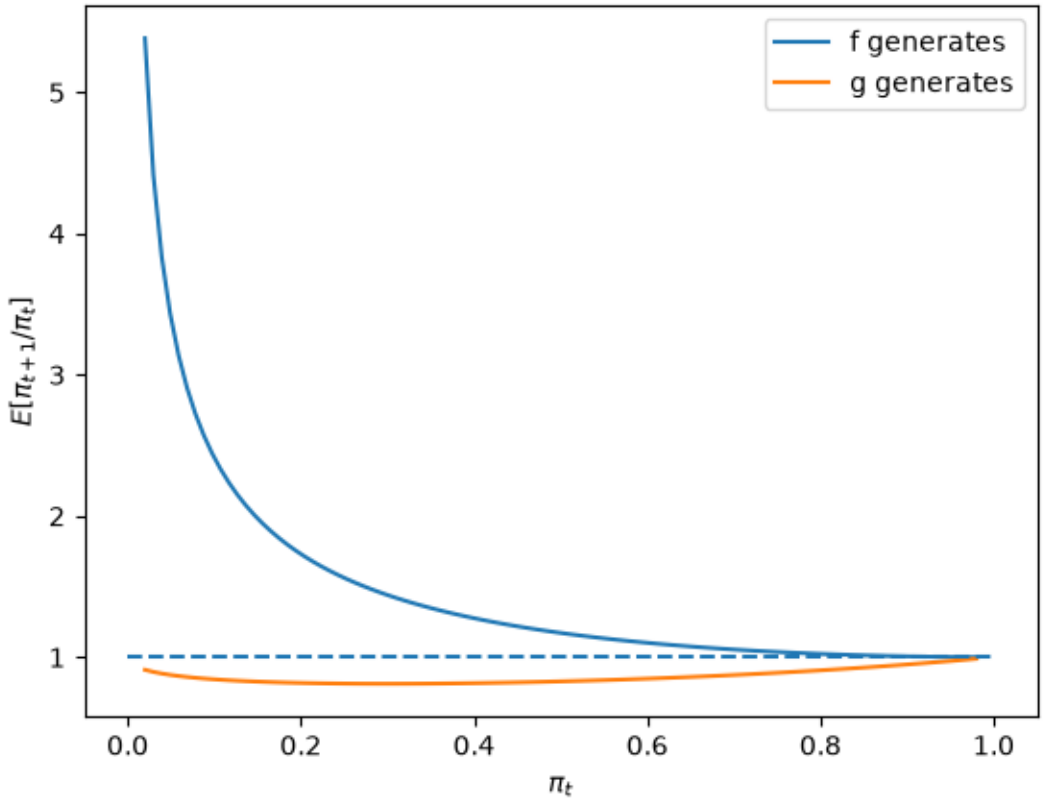
```
def plot_expected_ratio(F_a=1, F_b=1, G_a=3, G_b=1.2):
    # Build f and g directly: here they can coincide, so ℓ(w) = 1 has no root
    f = lambda w: p(w, F_a, F_b)
    g = lambda w: p(w, G_a, G_b)
    l = lambda w: f(w) / g(w)
    π_grid = np.linspace(0.02, 0.98, 100)

    fig, ax = plt.subplots()
    for label, a in [('f', f), ('g', g)]:
        integrand = lambda w, π: a(w) * l(w) / (π * l(w) + 1 - π)
        ratios = [quad(integrand, 0, 1, args=(π,))[0] for π in π_grid]
        ax.plot(π_grid, ratios, label=f'{label} generates')

    ax.hlines(1, 0, 1, linestyle='--')
    ax.set(xlabel=r'$\pi_t$', ylabel=r'$E[\pi_{t+1} / \pi_t]$')
    ax.legend()
    plt.show()
```

First, consider the case where $F_a = F_b = 1$ and $G_a = 3, G_b = 1.2$.

```
plot_expected_ratio()
```



The above graph shows that when F generates the data, π_t on average always heads north, while when G generates the data, π_t heads south.

Two inequalities hold for every pair $f \neq g$ and every $\pi_t = \pi \in (0, 1)$.

Let $r(\ell) = \ell/(\pi\ell + 1 - \pi)$ be the ratio π_{t+1}/π_t as a function of $\ell = \ell(w_{t+1})$.

The function r is strictly concave and $E_g[\ell(w)] = \int f(w)dw = 1$, so Jensen's inequality implies

$$E_g[r(\ell(w))] < r(1) = 1.$$

The π-weighted average of the two expectations equals one:

$$\pi E_f[r(\ell(w))] + (1-\pi)E_g[r(\ell(w))] = \int \frac{f(w)}{\pi f(w) + (1-\pi)g(w)}(\pi f(w) + (1-\pi)g(w))dw = 1.$$

This equality is the martingale property of π_t under the decision maker's subjective distribution, studied in *Likelihood Ratio Processes and Bayesian Learning*.

It and the first inequality imply

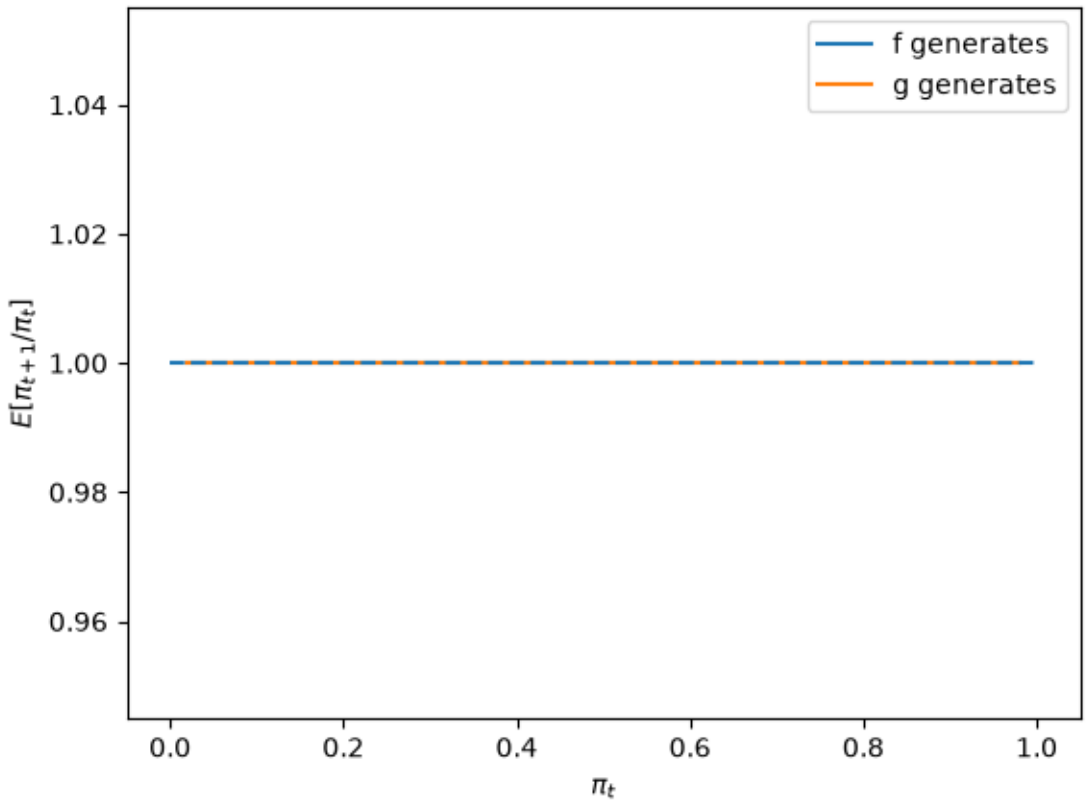
$$E_f[r(\ell(w))] > 1.$$

Thus π_t is a submartingale when F generates the data and a supermartingale when G generates the data.

Next, we'll look at a degenerate case in which f and g are identical beta distributions, and $F_a = G_a = 3, F_b = G_b = 1.2$.

In a sense, here there is nothing to learn.

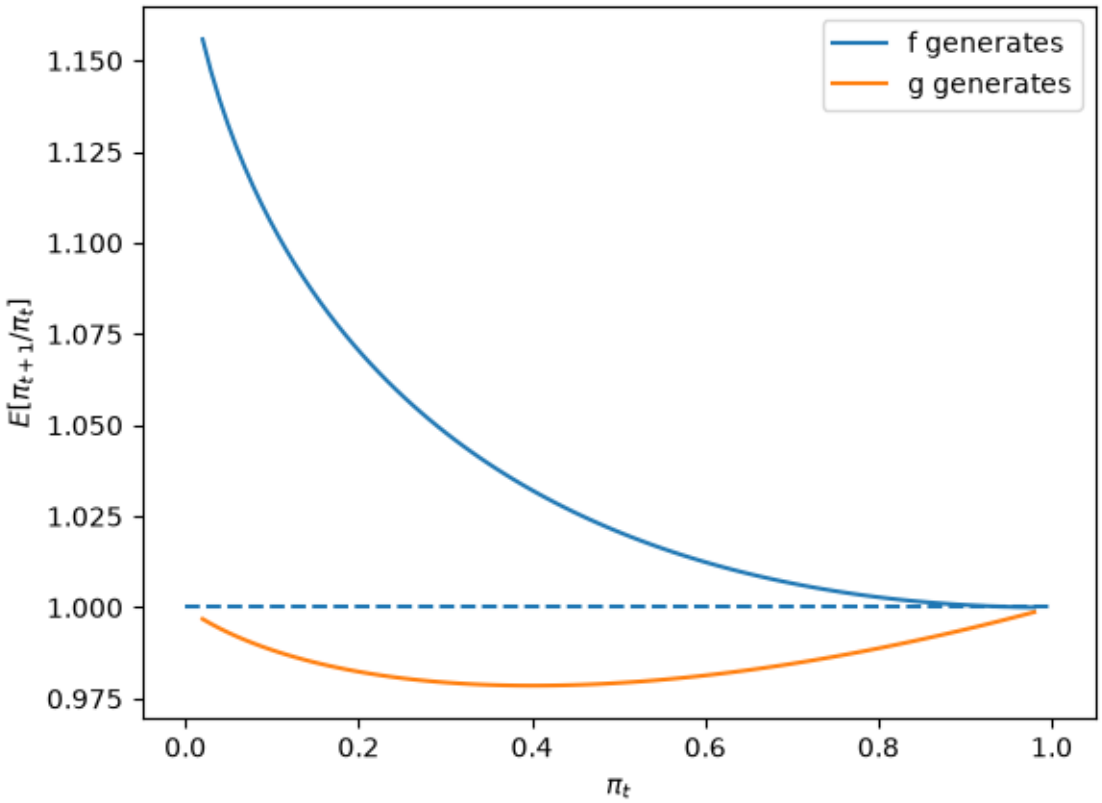
```
plot_expected_ratio(F_a=3, F_b=1.2)
```



The above graph says that π_t is inert and remains at its initial value.

Finally, let's look at a case in which f and g are neither very different nor identical, in particular one in which $F_a = 2, F_b = 1$ and $G_a = 3, G_b = 1.2$.

```
plot_expected_ratio(F_a=2, F_b=1, G_a=3, G_b=1.2)
```



The two inequalities still hold, but both curves now stay close to one.

On the plotted grid, $E_f[\pi_{t+1}/\pi_t]$ never exceeds 1.16 and $E_g[\pi_{t+1}/\pi_t]$ never falls below 0.978.

In the first case the corresponding numbers are 5.38 and 0.811.

The densities Beta$(2, 1)$ and Beta$(3, 1.2)$ are close to each other: $KL(f, g) = 0.052$ and $KL(g, f) = 0.039$, compared with 0.759 and 0.344 in the first case.

A draw moves π_t only a little on average, so learning is slow.

28.9 Sequels

We'll apply and dig deeper into some of the ideas presented in this lecture:

- *Likelihood Ratio Processes* describes **likelihood ratio processes** and their role in frequentist and Bayesian statistical theories
- *Likelihood Ratio Processes and Bayesian Learning* studies how a likelihood ratio process drives the posterior probability π_t and shows that π_t is a martingale under the decision maker's subjective distribution (28.1)
- *Incorrect Models* studies a decision maker who uses this lecture's statistical model when nature actually mixes F and G anew each period
- *Merging of Opinions: The Blackwell–Dubins Theorem* studies the Beta–Bernoulli model, an instance of de Finetti's representation (28.3), and asks when the opinions of two Bayesians with different priors merge
- *Statistical Divergence Measures* describes the Kullback–Leibler divergence and the Chernoff entropy that govern rates of learning

- *A Bayesian Formulation of Friedman and Wald's Problem* studies a decision maker with this lecture's beliefs who decides sequentially when to stop sampling
- *Bayesian versus Frequentist Decision Rules* studies a World War II US Navy Captain's hunch that a (frequentist) decision rule that the Navy had told him to use was inferior to a sequential rule that Abraham Wald had not yet designed

28.10 Exercises

Exercise 28.10.1

This exercise shows that a mixture of IID sequences has nonnegative serial covariances and that a finite exchangeable sequence need not be a mixture of IID sequences.

1. Prove identity (28.2).
2. Let ϕ be a function for which $E_f[\phi(W)]$ and $E_g[\phi(W)]$ are finite. Use (28.2) to show that under the joint density (28.1)

$$\text{Cov}(\phi(W_0), \phi(W_1)) = \tilde{\pi}(1-\tilde{\pi})(E_f[\phi(W)] - E_g[\phi(W)])^2 \geq 0.$$

3. For this lecture's densities $f = \text{Beta}(1,1)$ and $g = \text{Beta}(3, 1.2)$ and $\tilde{\pi} = 0.5$, compute $\text{Cov}(W_0, W_1)$ from the formula in part 2 and compare it with a simulation.
4. An urn contains one red ball and one black ball. We draw both balls without replacement. Let $W_t = 1$ if draw $t \in \{0, 1\}$ is red and $W_t = 0$ if it is black. Show that (W_0, W_1) is exchangeable and that $\text{Cov}(W_0, W_1) < 0$. Explain why (W_0, W_1) is not a mixture of IID pairs and why de Finetti's theorem requires an infinite sequence.

Solution

Part 1.

With $h(W_t) = \tilde{\pi} f(W_t) + (1-\tilde{\pi}) g(W_t)$,

$$h(W_0)h(W_1) = \tilde{\pi}^2 f(W_0)f(W_1) + \tilde{\pi}(1-\tilde{\pi})(f(W_0)g(W_1) + g(W_0)f(W_1)) + (1-\tilde{\pi})^2 g(W_0)g(W_1).$$

Subtracting this from $h(W_0, W_1) = \tilde{\pi} f(W_0)f(W_1) + (1-\tilde{\pi}) g(W_0)g(W_1)$ and using $\tilde{\pi} - \tilde{\pi}^2 = (1-\tilde{\pi}) - (1-\tilde{\pi})^2 = \tilde{\pi}(1-\tilde{\pi})$ gives

$$h(W_0, W_1) - h(W_0)h(W_1) = \tilde{\pi}(1-\tilde{\pi})(f(W_0)f(W_1) + g(W_0)g(W_1) - f(W_0)g(W_1) - g(W_0)f(W_1)),$$

which factors into the right side of (28.2).

Part 2.

Multiply (28.2) by $\phi(W_0)\phi(W_1)$ and integrate over W_0 and W_1:

$$\text{Cov}(\phi(W_0), \phi(W_1)) = \tilde{\pi}(1-\tilde{\pi}) \int \phi(w)(f(w) - g(w))dw \int \phi(w)(f(w) - g(w))dw = \tilde{\pi}(1-\tilde{\pi})(E_f[\phi(W)] - E_g[\phi(W)])^2.$$

The same conclusion holds for a mixture of IID sequences with any mixing measure.

Let Q be the distribution, drawn from the mixing measure, from which the W_t are IID draws.

The law of total covariance gives

$$\text{Cov}(\phi(W_0), \phi(W_1)) = E[\text{Cov}(\phi(W_0), \phi(W_1) \mid Q)] + \text{Cov}(E[\phi(W_0)|Q], E[\phi(W_1)|Q]) = 0 + \text{Var}(E[\phi(W)|Q]) \geq 0.$$

Part 3.

The means of Beta$(1,1)$ and Beta$(3,1.2)$ are $1/2$ and $3/4.2$.

```
π_tilde = 0.5
μ_f, μ_g = 1 / (1 + 1), 3 / (3 + 1.2)
cov_theory = π_tilde * (1 - π_tilde) * (μ_f - μ_g)**2

rng = np.random.default_rng(1234)
N = 1_000_000
nature_draws_f = rng.random(N) < π_tilde
W = np.where(nature_draws_f[:, None],
             rng.beta(1, 1, size=(N, 2)),
             rng.beta(3, 1.2, size=(N, 2)))
cov_sim = np.cov(W[:, 0], W[:, 1])[0, 1]

print(f"Cov(W_0, W_1): formula {cov_theory:.5f}, simulation {cov_sim:.5f}")
```

```
Cov(W_0, W_1): formula 0.01148, simulation 0.01147
```

Part 4.

The first draw is red with probability $1/2$, and then the second draw is black with probability one.

So Prob$(W_0 = 1, W_1 = 0)$ = Prob$(W_0 = 0, W_1 = 1) = 1/2$ and Prob$(W_0 = W_1) = 0$.

The joint distribution is invariant to exchanging W_0 and W_1, so (W_0, W_1) is exchangeable.

$E[W_0] = E[W_1] = 1/2$ and $E[W_0 W_1] = 0$, so Cov$(W_0, W_1) = -1/4$.

By part 2, every mixture of IID pairs has Cov$(W_0, W_1) \geq 0$.

So (W_0, W_1) is not a mixture of IID pairs.

If (W_0, W_1) were the first two members of an infinite exchangeable sequence, de Finetti's theorem would make the pair a mixture of IID pairs.

So the urn sequence cannot be extended to an infinite exchangeable sequence.

More generally, let $W_0, \ldots, W_{n-1}$ be exchangeable with common variance $\sigma^2 > 0$ and common pairwise correlation ρ.

Then $0 \leq$ Var$(\sum_t W_t) = n\sigma^2(1 + (n-1)\rho)$, so $\rho \geq -1/(n-1)$.

The urn attains this bound with $n = 2$ and $\rho = -1$.

The bound approaches 0 as $n \to \infty$.

Exercise 28.10.2

This exercise studies identity (28.8) and rates of convergence of π_t.

1. Prove (28.8) by a second route. Use $\pi_t = \pi_{-1} L_t / (\pi_{-1} L_t + 1 - \pi_{-1})$, where $L_t = \prod_{s=0}^{t} \ell(w_s)$, and $E_f[\phi(w^t)] = E_g[L_t \phi(w^t)]$ for every bounded function ϕ of w^t.
2. Set $\pi_{-1} = 0.3$. Plot $1 - E_f[\pi_t]$ and $E_g[\pi_t]$ against t, then plot $\pi_{-1}(1 - E_f[\pi_t])$ and $(1 - \pi_{-1}) E_g[\pi_t]$ against t.
3. Compute the slopes over t of the medians across paths of $\log(\pi_t/(1-\pi_t))$ under F and under G. Compare them with $KL(f,g)$ and $-KL(g,f)$.
4. Compute the slope over t of $\log(1 - E_f[\pi_t])$ and compare it with $-C(f,g)$, where $C(f,g)$ is the Chernoff

entropy.

Hint for part 4: rare paths dominate $1 - E_f[\pi_t]$, so plain Monte Carlo estimates of it are imprecise for large t.

Let $A(s) = \int f(w)^s g(w)^{1-s} dw$, let s^* minimize $A(s)$ on $[0, 1]$, and let $m(w) = f(w)^{s^*} g(w)^{1-s^*} / A(s^*)$.

Then $f(w)/m(w) = A(s^*)\ell(w)^{1-s^*}$, so

$$E_f[\phi(w^t)] = A(s^*)^{t+1} E_m[\phi(w^t) L_t^{1-s^*}].$$

Draw from m to estimate the expectation on the right.

With f uniform, m is the Beta$(3 - 2s^*, 1.2 - 0.2s^*)$ density.

Solution

Part 1.

Under F, the density of w^t is $f(w_0) \cdots f(w_t) = L_t \, g(w_0) \cdots g(w_t)$, which gives $E_f[\phi(w^t)] = E_g[L_t \phi(w^t)]$.

Since $1 - \pi_t = (1 - \pi_{-1})/(\pi_{-1} L_t + 1 - \pi_{-1})$,

$$\pi_{-1} E_f[1 - \pi_t] = E_g\left[\frac{\pi_{-1} L_t (1 - \pi_{-1})}{\pi_{-1} L_t + 1 - \pi_{-1}}\right] = (1 - \pi_{-1}) E_g[\pi_t].$$

The subjective distribution (28.1) has density $\pi_{-1} f(w^t) + (1 - \pi_{-1}) g(w^t) = (\pi_{-1} L_t + 1 - \pi_{-1}) g(w^t)$, so

$$E[\pi_t(1 - \pi_t)] = E_g[(\pi_{-1} L_t + 1 - \pi_{-1}) \pi_t (1 - \pi_t)] = E_g[\pi_{-1} L_t (1 - \pi_t)] = \pi_{-1} E_f[1 - \pi_t].$$

Part 2.

We write a version of `simulate` that returns log odds, because π_t rounds to exactly 1 once its log odds exceed about 37.

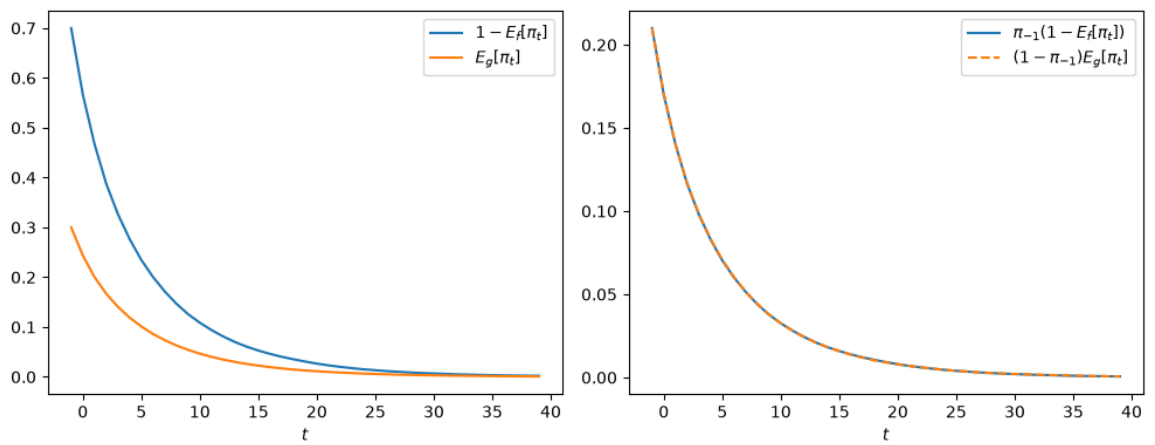
```
def simulate_log_odds(rng, a, b, T, N, π_init):
    "Log odds of π_t for t = -1, 0, ..., T-1 along N paths of IID Beta(a, b) draws.
 ↪"
    w = rng.beta(a, b, size=(N, T))
    log_l = np.log(p(w, 1, 1)) - np.log(p(w, 3, 1.2))
    log_odds_init = np.log(π_init / (1 - π_init))
    log_odds = log_odds_init + np.cumsum(log_l, axis=1)
    return np.column_stack([np.full(N, log_odds_init), log_odds])

π_init, T_ex, N_ex = 0.3, 40, 100_000
rng = np.random.default_rng(2024)
lo_F = simulate_log_odds(rng, 1, 1, T_ex, N_ex, π_init)
lo_G = simulate_log_odds(rng, 3, 1.2, T_ex, N_ex, π_init)
t_ex = np.arange(-1, T_ex)

err_F = np.mean(expit(-lo_F), axis=0)       # 1 - E_f[π_t]
err_G = np.mean(expit(lo_G), axis=0)        # E_g[π_t]

fig, axes = plt.subplots(1, 2, figsize=(10, 4))
axes[0].plot(t_ex, err_F, label=r'$1 - E_f[\pi_t]$')
axes[0].plot(t_ex, err_G, label=r'$E_g[\pi_t]$')
axes[1].plot(t_ex, π_init * err_F, label=r'$\pi_{-1}(1 - E_f[\pi_t])$')
axes[1].plot(t_ex, (1 - π_init) * err_G, '--',
             label=r'$(1 - \pi_{-1}) E_g[\pi_t]$')
for ax in axes:
    ax.set_xlabel('$t$')
    ax.legend()
plt.tight_layout()
plt.show()
```

```
for t in [0, 5, 10, 20]:
    print(f"t = {t:2d}: (1 - E_f[π_t]) / E_g[π_t] = {err_F[t+1] / err_G[t+1]:.3f}")
print(f"(1 - π_init) / π_init = {(1 - π_init) / π_init:.3f}")
```




```
t =  0: (1 - E_f[π_t]) / E_g[π_t] = 2.325
t =  5: (1 - E_f[π_t]) / E_g[π_t] = 2.320
t = 10: (1 - E_f[π_t]) / E_g[π_t] = 2.335
t = 20: (1 - E_f[π_t]) / E_g[π_t] = 2.354
(1 - π_init) / π_init = 2.333
```

The two curves in the left panel differ by the factor $7/3$; the two curves in the right panel coincide up to sampling error.

Part 3.

```
f = lambda w: p(w, 1, 1)
g = lambda w: p(w, 3, 1.2)
KL_fg = quad(lambda w: f(w) * np.log(f(w) / g(w)), 0, 1)[0]
KL_gf = quad(lambda w: g(w) * np.log(g(w) / f(w)), 0, 1)[0]

slope_F = np.polyfit(t_ex, np.median(lo_F, axis=0), 1)[0]
slope_G = np.polyfit(t_ex, np.median(lo_G, axis=0), 1)[0]
print(f"slope of median log odds under F: {slope_F:.3f},  KL(f, g) = {KL_fg:.3f}")
print(f"slope of median log odds under G: {slope_G:.3f}, -KL(g, f) = {-KL_gf:.3f}")
```

```
slope of median log odds under F: 0.754,  KL(f, g) = 0.759
slope of median log odds under G: -0.345, -KL(g, f) = -0.344
```

Typical paths converge faster under F than under G.

Part 4.

We compute the Chernoff entropy and the minimizer s^*.

```
A = lambda s: quad(lambda w: f(w)**s * g(w)**(1 - s), 0, 1)[0]
res = op.minimize_scalar(A, bounds=(0, 1), method='bounded')
s_star, C = res.x, -np.log(res.fun)
print(f"s* = {s_star:.3f}, C(f, g) = {C:.4f}")
```

```
s* = 0.597, C(f, g) = 0.1212
```

Next we estimate $\log(1 - E_f[\pi_t])$ by importance sampling from m for t up to 199.

```
T_is, N_is = 200, 100_000
rng = np.random.default_rng(2025)
w = rng.beta(3 - 2 * s_star, 1.2 - 0.2 * s_star, size=(N_is, T_is))
log_L = np.cumsum(np.log(f(w)) - np.log(g(w)), axis=1)     # log L_t, t = 0, ...,␣
↪199
t_is = np.arange(T_is)

log_odds = np.log(π_init / (1 - π_init)) + log_L
weights = np.exp((1 - s_star) * log_L)                      # L_t^(1 - s*)
log_err_F = (t_is + 1) * np.log(A(s_star)) \
            + np.log(np.mean(expit(-log_odds) * weights, axis=0))

for t in [10, 20, 30]:
    print(f"t = {t}: log(1 - E_f[π_t]) importance sampling {log_err_F[t]:.3f}, "
          f"plain Monte Carlo {np.log(err_F[t+1]):.3f}")

late = t_is >= 100
slope = np.polyfit(t_is[late], log_err_F[late], 1)[0]
slope_adj = np.polyfit(t_is[late], log_err_F[late] + 0.5 * np.log(t_is[late] + 1),␣
↪1)[0]
print(f"slope of log(1 - E_f[π_t]) for t in [100, 199]: {slope:.4f}")
print(f"slope after adding log(t+1)/2:                  {slope_adj:.4f}")
print(f"-C(f, g):                                       {-C:.4f}")
```

```
t = 10: log(1 - E_f[π_t]) importance sampling -2.215, plain Monte Carlo -2.219
t = 20: log(1 - E_f[π_t]) importance sampling -3.638, plain Monte Carlo -3.626
t = 30: log(1 - E_f[π_t]) importance sampling -4.994, plain Monte Carlo -5.000
slope of log(1 - E_f[π_t]) for t in [100, 199]: -0.1243
slope after adding log(t+1)/2:                  -0.1209
-C(f, g):                                       -0.1212
```

The two estimates agree at small t.

Under m, $E_m[\log \ell(w)] = A'(s^*)/A(s^*) = 0$, so $\log L_t$ is a random walk without drift.

The expectation $E_m[(1-\pi_t)L_t^{1-s^*}]$ therefore decays only like $t^{-1/2}$, and

$$1 - E_f[\pi_t] \approx K\, t^{-1/2} e^{-C(f,g)t}$$

for a constant K.

The fitted slope is close to $-C(f,g)$, and adding $\frac{1}{2}\log(t+1)$ to remove the factor $t^{-1/2}$ brings it closer.

By (28.8), $E_g[\pi_t]$ decays at the same rate.

Exercise 28.10.3

A Pólya urn starts with a red balls and b black balls.

At each date $t = 0, 1, \ldots$ we draw a ball at random, set $W_t = 1$ if it is red and $W_t = 0$ if it is black, and return it to the urn together with one more ball of the same color.

1. Show that $\text{Prob}(W_0 = 1, W_1 = 0) = \text{Prob}(W_0 = 0, W_1 = 1)$. More generally, show that the probability of a sequence $(w_0, \ldots, w_{n-1})$ depends only on the number k of red draws, so that $W_0, W_1, \ldots$ is exchangeable. Find the mixing measure μ in de Finetti's representation (28.3).
2. Simulate many urns with $a = 2$ and $b = 3$ for $n = 2000$ draws. Compare the distribution across urns of the

fraction of red draws with the Beta(a, b) density.

3. Show that the fraction of red balls in the urn after n draws, k of them red, equals the probability Prob$(W_n = 1 \,|\, w_0, \ldots, w_{n-1})$ that a Bayesian with a Beta(a, b) prior over the success probability θ of IID Bernoulli draws attaches to the next draw being red.

Solution

Part 1.

Prob$(W_0 = 1, W_1 = 0) = \frac{a}{a+b}\frac{b}{a+b+1}$ and Prob$(W_0 = 0, W_1 = 1) = \frac{b}{a+b}\frac{a}{a+b+1}$.

The two are equal.

For a sequence of n draws with k red draws, the denominators of the successive conditional probabilities are $a + b, a + b + 1, \ldots, a + b + n - 1$ in every order.

The numerators for the red draws are $a, a + 1, \ldots, a + k - 1$, and the numerators for the black draws are $b, b + 1, \ldots, b + n - k - 1$, in every order.

So

$$\text{Prob}(W_0 = w_0, \ldots, W_{n-1} = w_{n-1}) = \frac{\prod_{j=0}^{k-1}(a+j)\prod_{j=0}^{n-k-1}(b+j)}{\prod_{j=0}^{n-1}(a+b+j)} = \frac{B(a+k, b+n-k)}{B(a,b)},$$

where B is the Beta function.

The right side equals $\int_0^1 \theta^k (1-\theta)^{n-k} \mu(d\theta)$ with μ the Beta(a, b) distribution.

That is de Finetti's representation (28.3).

Part 2.

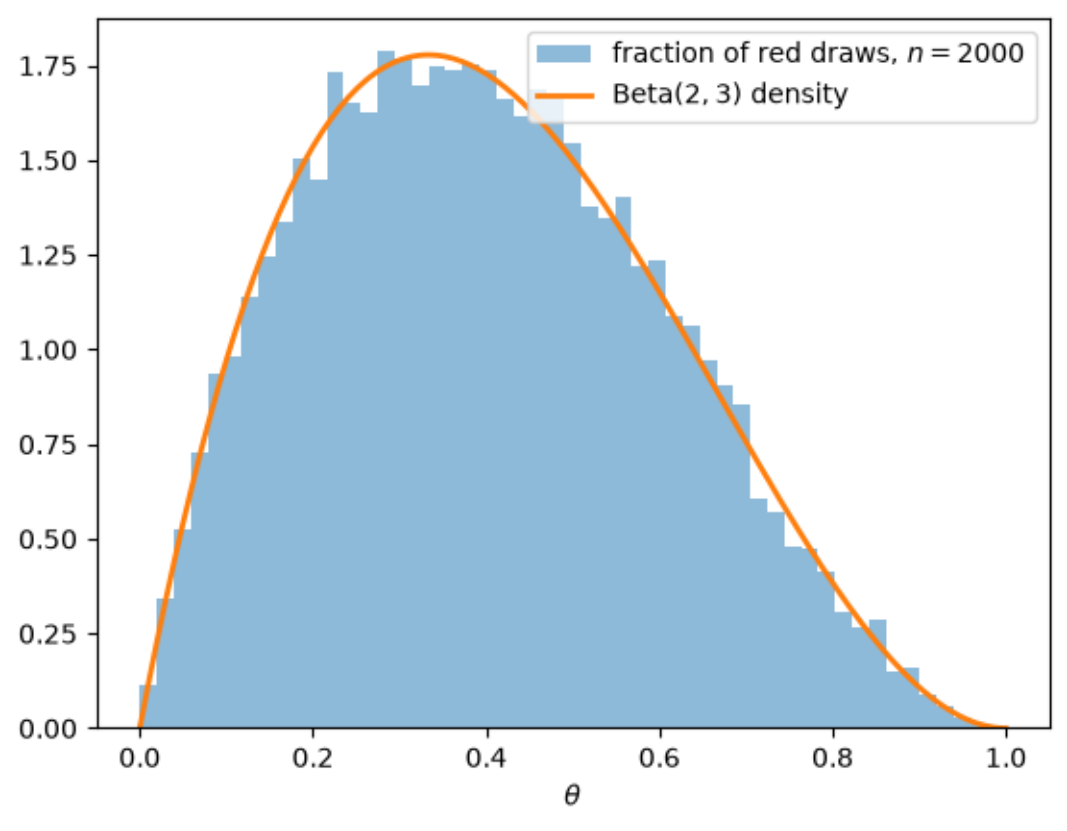
```
from scipy.stats import beta as beta_dist, kstest

def polya_urn(rng, a, b, n, N):
    "Simulate N Pólya urns for n draws; return the N x n array of draws."
    red = np.full(N, float(a))
    total = float(a + b)
    W = np.empty((N, n), dtype=np.int8)
    for t in range(n):
        W[:, t] = rng.random(N) < red / total
        red += W[:, t]
        total += 1
    return W

a, b, n, N = 2, 3, 2000, 20_000
rng = np.random.default_rng(7)
W = polya_urn(rng, a, b, n, N)
freq = W.mean(axis=1)

grid = np.linspace(0, 1, 200)
fig, ax = plt.subplots()
ax.hist(freq, bins=50, density=True, alpha=0.5,
        label=f'fraction of red draws, $n = {n}$')
ax.plot(grid, beta_dist.pdf(grid, a, b), lw=2, label=f'Beta$({a}, {b})$ density')
ax.set_xlabel(r'$\theta$')
ax.legend()
plt.show()

print(f"Prob(W_0=1, W_1=0) = {np.mean((W[:, 0] == 1) & (W[:, 1] == 0)):.4f}, "
      f"Prob(W_0=0, W_1=1) = {np.mean((W[:, 0] == 0) & (W[:, 1] == 1)):.4f}, "
      f"a * b / ((a + b) * (a + b + 1)) = {a * b / ((a + b) * (a + b + 1)):.4f}")
print(f"Kolmogorov-Smirnov test against Beta({a}, {b}): "
      f"p-value = {kstest(freq, beta_dist(a, b).cdf).pvalue:.3f}")
```




```
Prob(W_0=1, W_1=0) = 0.2033, Prob(W_0=0, W_1=1) = 0.2003, exact = 0.2000
Kolmogorov-Smirnov test against Beta(2, 3): p-value = 0.209
```

The fraction of red draws is distributed approximately Beta$(2, 3)$ across urns, and the Kolmogorov–Smirnov test does not reject that distribution.

Part 3.

After n draws with k red draws, the urn holds $a + k$ red balls out of $a + b + n$ balls, so the next draw is red with probability $(a + k)/(a + b + n)$.

A Bayesian with a Beta(a, b) prior over θ has the Beta$(a + k, b + n - k)$ posterior after observing k successes in n IID Bernoulli trials.

His probability that the next draw is a success is the posterior mean of θ, namely $(a + k)/(a + b + n)$.

The urn and the Bayesian agree.

The next cell checks this by grouping the simulated urns according to the number of red balls among the first 10 draws.

Frequencies computed from few urns are less precise.

```
n0 = 10
k = W[:, :n0].sum(axis=1)
for k_val in range(0, n0 + 1, 2):
    sel = k == k_val
    print(f"k = {k_val:2d} ({sel.sum():4d} urns): "
          f"frequency of red at draw {n0} = {W[sel, n0].mean():.3f}, "
          f"(a + k) / (a + b + n) = {(a + k_val) / (a + b + n0):.3f}")
```

```
k =  0 (1378 urns): frequency of red at draw 10 = 0.129, (a + k) / (a + b + n)␣
↪= 0.133
k =  2 (2686 urns): frequency of red at draw 10 = 0.265, (a + k) / (a + b + n)␣
↪= 0.267
k =  4 (2850 urns): frequency of red at draw 10 = 0.403, (a + k) / (a + b + n)␣
↪= 0.400
k =  6 (2077 urns): frequency of red at draw 10 = 0.535, (a + k) / (a + b + n)␣
↪= 0.533
k =  8 (1118 urns): frequency of red at draw 10 = 0.684, (a + k) / (a + b + n)␣
↪= 0.667
k = 10 ( 218 urns): frequency of red at draw 10 = 0.839, (a + k) / (a + b + n)␣
↪= 0.800
```

CHAPTER

TWENTYNINE

LIKELIHOOD RATIO PROCESSES AND BAYESIAN LEARNING

29.1 Overview

This lecture describes the role that *likelihood ratio processes* play in *Bayesian learning*.

As in *Likelihood Ratio Processes*, we'll use a simple statistical setting from *Exchangeability and Bayesian Updating*.

We'll focus on how a likelihood ratio process and a *prior* probability determine a *posterior* probability.

We'll derive a convenient recursion for today's posterior as a function of yesterday's posterior and today's multiplicative increment to a likelihood process.

We'll also present a useful generalization of that formula that represents today's posterior in terms of an initial prior and today's realization of the likelihood ratio process.

We'll study how, at least in our setting, a Bayesian eventually learns the probability distribution that generates the data, an outcome that rests on the asymptotic behavior of likelihood ratio processes studied in *Likelihood Ratio Processes*.

We'll also drill down into the psychology of our Bayesian learner and study dynamics under his subjective beliefs.

This lecture provides technical results that underlie outcomes to be studied in *Job Search IX: Search with Learning*, *A Problem that Stumped Milton Friedman*, and *Bayesian versus Frequentist Decision Rules*.

We'll begin by loading some Python modules.

```
import matplotlib.pyplot as plt
import numpy as np
from numba import vectorize, jit
from math import gamma
import pandas as pd
from scipy.integrate import quad
from scipy.special import expit, logit
from scipy.stats import ks_2samp


import seaborn as sns
colors = sns.color_palette()

rng = np.random.default_rng(142857)
```

29.2 The setting

We begin by reviewing the setting in *Likelihood Ratio Processes*, which we adopt here too.

A nonnegative random variable W has one of two probability density functions, either f or g.

Before the beginning of time, nature once and for all decides whether she will draw a sequence of IID draws from f or from g.

We let q be the density that nature chose once and for all, so that q is either f or g, permanently.

Nature knows which density it permanently draws from, but we the observers do not.

We do know both f and g, but we don't know which density nature chose.

But we want to know.

To do that, we use observations.

We observe a sequence $w_0, w_1, \ldots$ of IID draws from either f or g.

We date the first draw $t = 0$ and the prior $t = -1$, as in *Exchangeability and Bayesian Updating*.

Likelihood Ratio Processes dates the first draw $t = 1$.

We want to use these observations to infer whether nature chose f or g.

A likelihood ratio process is a useful tool for this task.

To begin, we define the key component of a likelihood ratio process, namely, the time t likelihood ratio as the random variable

$$\ell(w_t) = \frac{f(w_t)}{g(w_t)}, \quad t \geq 0.$$

We assume that f and g both put positive probabilities on the same intervals of possible realizations of the random variable W.

That means that under the g density, $\ell(w_t) = \frac{f(w_t)}{g(w_t)}$ is evidently a nonnegative random variable with mean 1.

A **likelihood ratio process** for sequence $\{w_t\}_{t=0}^{\infty}$ is defined as

$$L(w^t) = \prod_{s=0}^{t} \ell(w_s),$$

where $w^t = \{w_0, \ldots, w_t\}$ is a history of observations up to and including time t.

Sometimes for shorthand we'll write

$$L_t = L(w^t) = \frac{f(w^t)}{g(w^t)}$$

where we use the conventions that $f(w^t) = f(w_0)f(w_1)\ldots f(w_t)$ and $g(w^t) = g(w_0)g(w_1)\ldots g(w_t)$.

Notice that the likelihood process satisfies the **recursion** or **multiplicative decomposition**

$$L(w^t) = \ell(w_t)L(w^{t-1}),$$

with $L(w^{-1}) \equiv 1$.

The likelihood ratio and its logarithm are key tools for making inferences using a classic frequentist approach due to Neyman and Pearson [Neyman and Pearson, 1933].

We'll again deploy the following Python code from *Likelihood Ratio Processes* that evaluates f and g as two different beta distributions, then computes and simulates an associated likelihood ratio process by generating a sequence w^t from *some* probability distribution, for example, a sequence of IID draws from g.

```
# Parameters in the two beta distributions.
F_a, F_b = 1, 1
G_a, G_b = 3, 1.2

@vectorize
def p(x, a, b):
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x**(a-1) * (1 - x)**(b-1)

# The two density functions.
f = jit(lambda x: p(x, F_a, F_b))
g = jit(lambda x: p(x, G_a, G_b))
```

```
@jit
def simulate(a, b, rng, T=50, N=500):
    '''
    Generate N sets of T observations of the likelihood ratio,
    return as N x T matrix.  Column s holds ℓ(w_s), s = 0, ..., T-1.
    '''

    l_arr = np.empty((N, T))

    for i in range(N):

        for j in range(T):
            w = rng.beta(a, b)
            l_arr[i, j] = f(w) / g(w)

    return l_arr
```

We'll also use the following Python code to prepare some informative simulations.

Column t of `l_seq_g` and `l_seq_f` holds $L(w^t)$.

```
l_arr_g = simulate(G_a, G_b, rng, N=50000)
l_seq_g = np.cumprod(l_arr_g, axis=1)
```

```
l_arr_f = simulate(F_a, F_b, rng, N=50000)
l_seq_f = np.cumprod(l_arr_f, axis=1)
```

29.3 Likelihood ratio processes and Bayes' law

Let $\pi_{-1} \in [0, 1]$ be a Bayesian statistician's prior probability that nature generates w^t as a sequence of IID draws from distribution f.

- here "probability" is to be interpreted as a way to summarize or express a subjective opinion
- it does *not* mean an anticipated relative frequency as sample size grows without limit

Let π_t be a Bayesian posterior probability defined as

$$\pi_t = \text{Prob}(q = f | w^t) \tag{29.1}$$

The likelihood ratio process is a principal actor in the formula that governs the evolution of the posterior probability π_t, an instance of *Bayes' law*.

Let's derive a couple of formulas for π_t, one in terms of the likelihood ratio $\ell(w_t)$, the other in terms of $L(w^t)$.

To begin, we use the notational conventions

- $f(w^t) \equiv f(w_0)f(w_1)\cdots f(w_t)$
- $g(w^t) \equiv g(w_0)g(w_1)\cdots g(w_t)$
- $\pi_{-1} = \text{Prob}(q = f|\emptyset)$
- $\pi_t = \text{Prob}(q = f|w^t)$

Here the symbol $\emptyset$ means "empty set" or "no data".

With no data in hand, our Bayesian statistician thinks that the probability density of the sequence w^t is

$$\text{Prob}(w^t|\emptyset) = \pi_{-1}f(w^t) + (1-\pi_{-1})g(w^t)$$

Laws of probability say that the joint distribution $\text{Prob}(AB)$ of events A and B are connected to the conditional distributions $\text{Prob}(A|B)$ and $\text{Prob}(B|A)$ by

$$\text{Prob}(AB) = \text{Prob}(A|B)\text{Prob}(B) = \text{Prob}(B|A)\text{Prob}(A). \tag{29.2}$$

We are interested in events

$$A = \{q = f\}, \quad B = \{w^t\},$$

where braces $\{\cdot\}$ are our shorthand for "event".

So in our setting, probability laws (29.2) imply that

$$\text{Prob}(q = f|w^t)\text{Prob}(w^t|\emptyset) = \text{Prob}(w^t|q = f)\text{Prob}(q = f|\emptyset)$$

or

$$\pi_t\left[\pi_{-1}f(w^t) + (1-\pi_{-1})g(w^t)\right] = f(w^t)\pi_{-1}$$

or

$$\pi_t = \frac{f(w^t)\pi_{-1}}{\pi_{-1}f(w^t) + (1-\pi_{-1})g(w^t)}$$

Dividing both the numerator and the denominator on the right side of the above equation by $g(w^t)$ implies

$$\pi_t = \frac{\pi_{-1}L\left(w^t\right)}{\pi_{-1}L\left(w^t\right) + 1 - \pi_{-1}}. \tag{29.3}$$

Formula (29.3) can be regarded as a one step revision of prior probability π_{-1} after seeing the batch of data $\{w_s\}_{s=0}^t$.

Formula (29.3) shows the key role that the likelihood ratio process $L\left(w^t\right)$ plays in determining the posterior probability π_t.

Formula (29.3) is the foundation for the insight that, because of how the likelihood ratio process behaves as $t \to +\infty$, the likelihood ratio process dominates the initial prior π_{-1} in determining the limiting behavior of π_t.

29.3.1 A recursive formula

We can use a similar line of argument to get a recursive version of formula (29.3).

The laws of probability imply

$$\text{Prob}(q = f|w^{t+1}) = \frac{\text{Prob}(q = f|w^t)f(w_{t+1})}{\text{Prob}(q = f|w^t)f(w_{t+1}) + (1 - \text{Prob}(q = f|w^t))g(w_{t+1})}$$

or

$$\pi_{t+1} = \frac{\pi_t f(w_{t+1})}{\pi_t f(w_{t+1}) + (1 - \pi_t)g(w_{t+1})} \tag{29.4}$$

Evidently, the above equation asserts that

$$\text{Prob}(q = f|w^{t+1}) = \frac{\text{Prob}(q = f|w^t)f(w_{t+1})}{\text{Prob}(w_{t+1}|w^t)}$$

where

$$\text{Prob}(w_{t+1}|w^t) = \pi_t f(w_{t+1}) + (1 - \pi_t)g(w_{t+1})$$

is the density of w_{t+1} conditional on the history w^t.

Conditional on q, the draw w_{t+1} is independent of w^t, so $f(w_{t+1})$ is the density of w_{t+1} conditional on $q = f$ and w^t.

Dividing both the numerator and the denominator on the right side of the equation (29.4) by $g(w_{t+1})$ implies the recursion

$$\pi_{t+1} = \frac{\pi_t \ell(w_{t+1})}{\pi_t \ell(w_{t+1}) + 1 - \pi_t} \tag{29.5}$$

with π_{-1} being a Bayesian prior probability that $q = f$, i.e., a personal or subjective belief about q based on our having seen no data.

Formula (29.3) can be deduced by iterating on equation (29.5).

In terms of log odds, recursion (29.5) is

$$\log \frac{\pi_{t+1}}{1 - \pi_{t+1}} = \log \frac{\pi_t}{1 - \pi_t} + \log \ell(w_{t+1}) \tag{29.6}$$

and formula (29.3) is

$$\log \frac{\pi_t}{1 - \pi_t} = \log \frac{\pi_{-1}}{1 - \pi_{-1}} + \log L(w^t).$$

Our code iterates on log odds.

Iterating on (29.5) itself accumulates rounding errors that push π_t slightly above 1 or below 0 once π_t is close to 1 or 0.

Below we define a Python function that updates the log odds of $q = f$ using likelihood ratio ℓ according to recursion (29.6)

```
@jit
def update(log_odds, l):
    "Update the log odds of q = f using likelihood ratio l"
    return log_odds + np.log(l)
```

To illustrate how the likelihood ratio process dominates the initial prior, below we will plot graphs showing *one* simulated path of the likelihood ratio process L_t along with two paths of π_t that are associated with the *same* realization of the likelihood ratio process but *different* initial prior probabilities π_{-1}.

First, we tell Python two values of π_{-1}.

```
π1, π2 = 0.2, 0.8
```

Next we generate paths of the likelihood ratio process L_t and the posterior π_t for a history of IID draws from density f.

Column 0 of `π_seq_f` holds the prior π_{-1} and column $t+1$ holds π_t.

```
T = l_arr_f.shape[1]
log_odds_f = np.empty((2, T+1))
log_odds_f[:, 0] = logit([π1, π2])

for t in range(T):
    for i in range(2):
        log_odds_f[i, t+1] = update(log_odds_f[i, t], l_arr_f[0, t])

π_seq_f = expit(log_odds_f)
```

```
fig, ax1 = plt.subplots()

for i in range(2):
    ax1.plot(range(-1, T), π_seq_f[i, :],
    label=fr"$\pi_{{-1}}$={π_seq_f[i, 0]:.1f}", lw=2)

ax1.set_ylabel(r"$\pi_t$")
ax1.set_xlabel(r"$t$")
ax1.legend()

ax2 = ax1.twinx()
ax2.plot(range(T), np.log(l_seq_f[0, :]), '--', color='b', lw=2)
ax2.set_ylabel(r"$\log(L(w^{t}))$")

plt.show()
```

The dashed line in the graph above records the logarithm of the likelihood ratio process $\log L(w^t)$.

Please note that there are two different scales on the y axis.

Now let's study what happens when the history consists of IID draws from density g

```
T = l_arr_g.shape[1]
log_odds_g = np.empty((2, T+1))
log_odds_g[:, 0] = logit([π1, π2])

for t in range(T):
    for i in range(2):
        log_odds_g[i, t+1] = update(log_odds_g[i, t], l_arr_g[0, t])

π_seq_g = expit(log_odds_g)
```

```
fig, ax1 = plt.subplots()

for i in range(2):
    ax1.plot(range(-1, T), π_seq_g[i, :],
        label=fr"$\pi_{{-1}}$={π_seq_g[i, 0]:.1f}", lw=2)

ax1.set_ylabel(r"$\pi_t$")
ax1.set_xlabel(r"$t$")
```

(continues on next page)

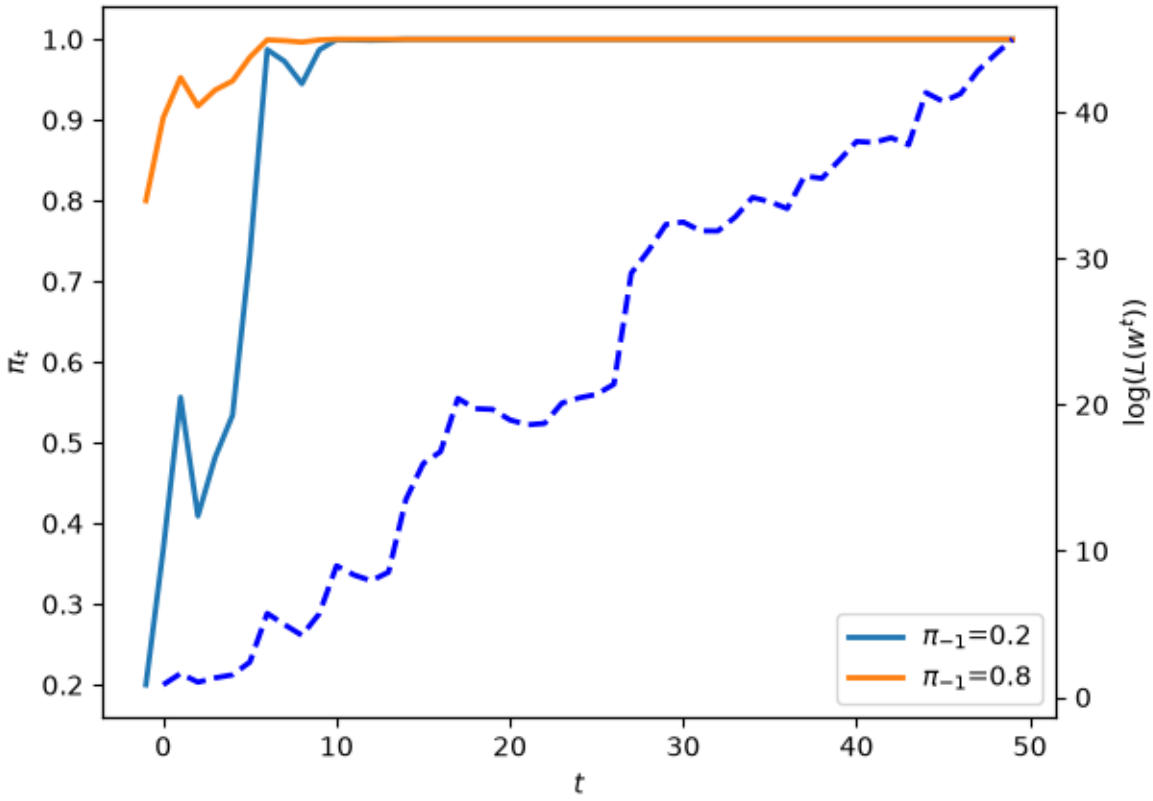


Fig. 29.1: Posterior paths and log likelihood ratio

(continued from previous page)

```
ax1.legend()

ax2 = ax1.twinx()
ax2.plot(range(T), np.log(l_seq_g[0, :]), '--', color='b', lw=2)
ax2.set_ylabel(r"$\log(L(w^{t}))$")

plt.show()
```

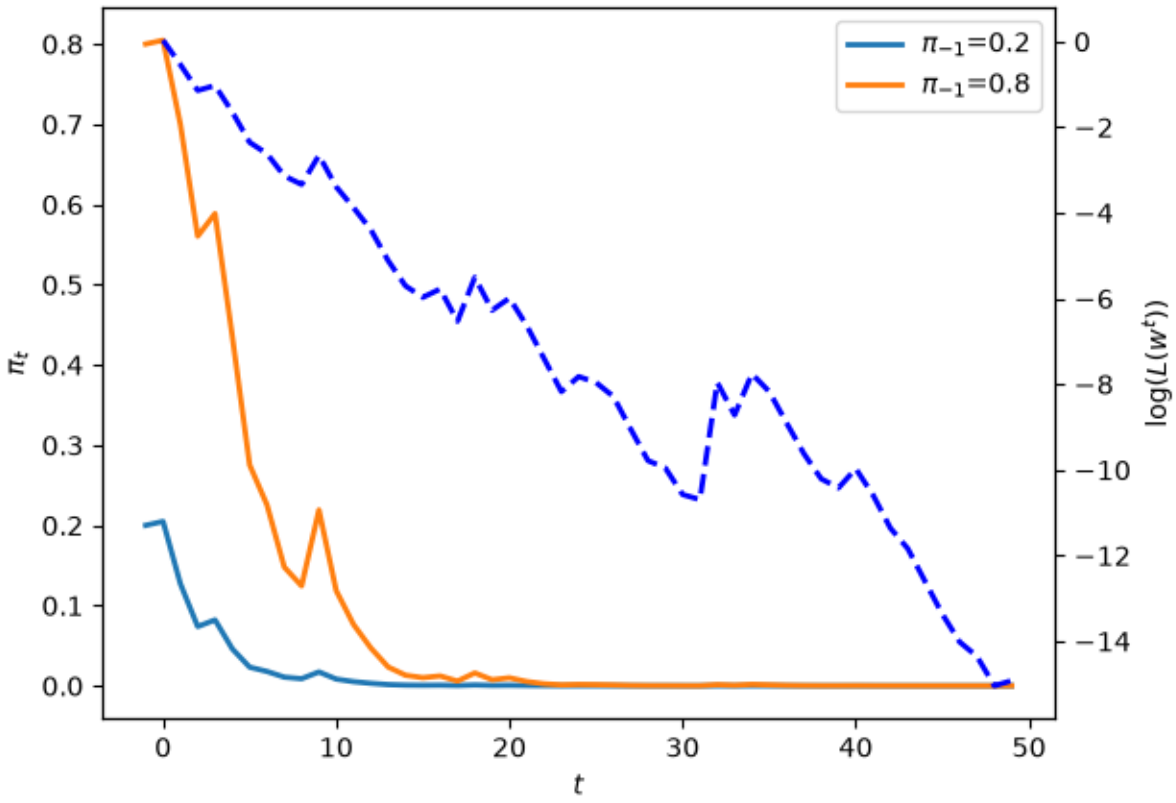


Fig. 29.2: Posterior paths and log likelihood ratio

The next cell checks that the batch formula (29.3) reproduces the paths of π_t that recursion (29.6) generated from the draws from density f.

```
π_seq = np.empty((2, T+1))
π_seq[:, 0] = π1, π2

for i in range(2):
    πL = π_seq[i, 0] * l_seq_f[0, :]
    π_seq[i, 1:] = πL / (πL + 1 - π_seq[i, 0])
```

```
np.abs(π_seq - π_seq_f).max() < 1e-10
```

```
np.True_
```

We thus conclude that the likelihood ratio process is a key ingredient of the formula (29.3) for a Bayesian's posterior

probability that nature has drawn history w^t as repeated draws from density f.

29.4 Another timing protocol

Let's study how the posterior probability $\pi_t = \text{Prob}(q = f|w^t)$ behaves when nature generates the history $w^t = \{w_0, w_1, \ldots, w_t\}$ under a different timing protocol.

Until now we assumed that before time 0 nature somehow chose to draw w^t as an IID sequence from either f or g.

Nature's decision about whether to draw from f or g was thus *permanent*.

We now assume a different timing protocol in which before *each period* $t = 0, 1, \ldots$ nature

- flips an α-weighted coin, then
- draws from f if it has drawn a "head"
- draws from g if it has drawn a "tail".

Under this timing protocol, nature draws permanently from *neither* f *nor* g, so a statistician who thinks that nature is drawing IID draws *permanently* from one of them is mistaken.

- in truth, nature draws IID from the mixture $h(w|\alpha) = \alpha f(w) + (1-\alpha)g(w)$, a density that is neither f nor g when $\alpha \in (0, 1)$

The statistician's posterior π_t remains his probability that $q = f$, as defined in (29.1).

It is not an estimate of the mixing parameter α.

The statistician's model is misspecified, and Kullback–Leibler divergences of f and g from h determine the limit of π_t.

By the law of large numbers, $\frac{1}{t+1} \log L(w^t)$ converges with probability one to

$$E_h[\log \ell(w)] = E_h\left[\log \frac{h(w|\alpha)}{g(w)}\right] - E_h\left[\log \frac{h(w|\alpha)}{f(w)}\right] = KL(h,g) - KL(h,f),$$

where E_h denotes the expectation under $h(\cdot|\alpha)$ and $KL(h,f) = \int h(w|\alpha) \log \frac{h(w|\alpha)}{f(w)} dw$; see *A helpful formula* in *Likelihood Ratio Processes*.

So $\log L(w^t) \to +\infty$ and $\pi_t \to 1$ when $KL(h,f) < KL(h,g)$.

And $\log L(w^t) \to -\infty$ and $\pi_t \to 0$ when $KL(h,f) > KL(h,g)$.

The statistician comes to believe the member of $\{f, g\}$ that is closer to h in Kullback–Leibler divergence.

Let's write some Python code to check this when $\alpha = 0.5$.

```
@jit
def simulate_mixture_path(α, T, rng):
    """
    Simulate draws w_0, ..., w_{T-1} under the mixture timing protocol.
    """
    w = np.empty(T)
    for t in range(T):
        if rng.random() < α:
            w[t] = rng.beta(F_a, F_b)
        else:
            w[t] = rng.beta(G_a, G_b)
    return w
```

(continues on next page)

(continued from previous page)

```
α = 0.5
T_mix = 200
rng = np.random.default_rng(142857)
w_mix = simulate_mixture_path(α, T_mix, rng)
```

We compute paths of π_t for three priors π_{-1}.

```
log_L_mix = np.cumsum(np.log(f(w_mix) / g(w_mix)))

fig, ax = plt.subplots(figsize=(10, 6))
for i, π_init in enumerate([0.25, 0.5, 0.75]):
    π_path = expit(logit(π_init) + log_L_mix)
    ax.plot(range(-1, T_mix), np.concatenate(([π_init], π_path)),
            label=fr'$\pi_{{-1}} = ${π_init:.2f}', color=colors[i], linewidth=2)

ax.set_xlabel('$t$')
ax.set_ylabel(r'$\pi_t$')
ax.legend()
plt.show()
```

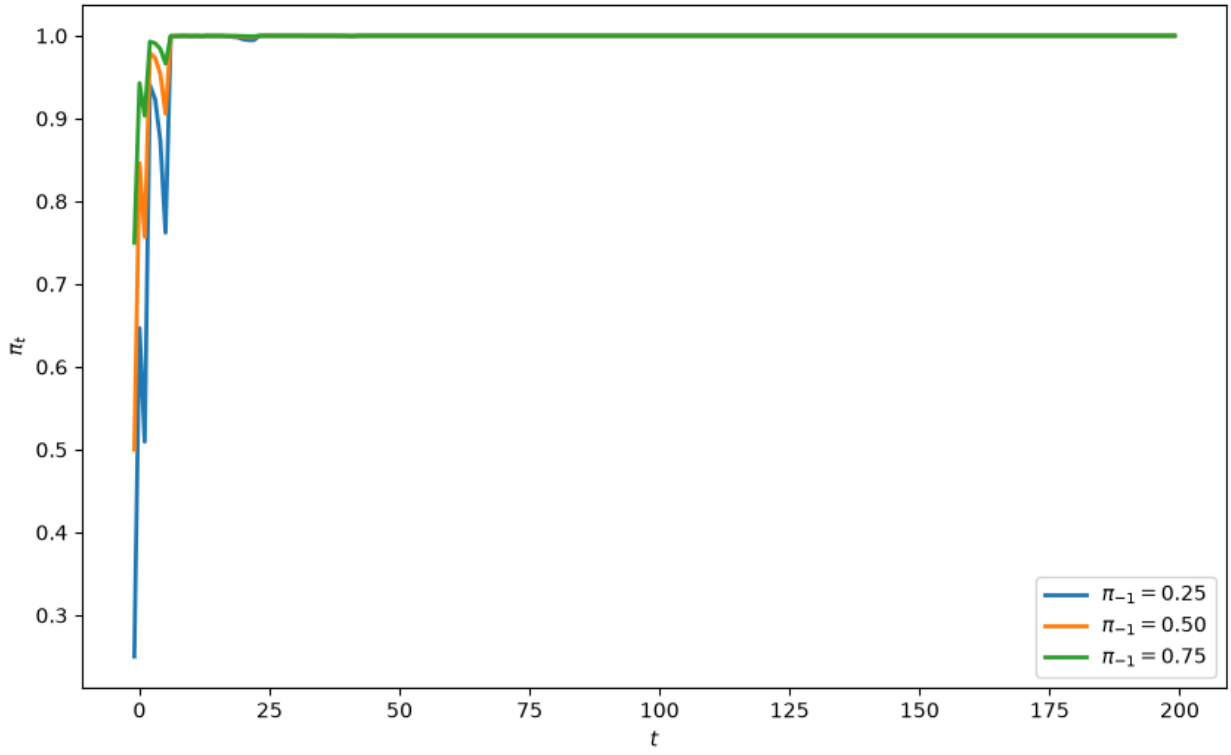


Fig. 29.3: Posterior under mixture model

Evidently, π_t converges to 1.

Let's check the *KL divergences* of f and g from h.

```
def compute_KL(f, g):
    """
    Compute KL divergence KL(f, g)
    """
    integrand = lambda w: f(w) * np.log(f(w) / g(w))
    val, _ = quad(integrand, 1e-5, 1-1e-5)
    return val


h = lambda w: α * f(w) + (1 - α) * g(w)
KL_hf, KL_hg = compute_KL(h, f), compute_KL(h, g)

print(f'KL(h, f) = {KL_hf:.3f}\nKL(h, g) = {KL_hg:.3f}')
```

```
KL(h, f) = 0.073
KL(h, g) = 0.281
```

Since $KL(h,f) < KL(h,g)$, f is closer to h, $\log L(w^t) \to +\infty$, and $\pi_t \to 1$ for every $\pi_{-1} \in (0,1)$.

Incorrect Models studies this misspecified statistician further, including the value of α at which $KL(h,f) = KL(h,g)$ and π_t does not converge.

It also studies a statistician who knows that nature mixes each period and uses Bayes' law to learn α.

29.5 Behavior of posterior probability $\{\pi_t\}$ under subjective probability distribution

We'll end this lecture by briefly studying what our Bayesian learner expects to learn under the subjective beliefs π_t cranked out by Bayes' law.

This will provide us with some perspective on our application of Bayes' law as a theory of learning.

As we shall see, at each time t, the Bayesian learner knows that he will be surprised.

But he expects that new information will not lead him to change his beliefs.

And it won't on average under his subjective beliefs.

We'll continue with our setting in which a McCall worker knows that successive draws of his wage are drawn from either F or G, but does not know which of these two distributions nature has drawn once-and-for-all before time 0.

We'll review and reiterate and rearrange some formulas that we have encountered above and in associated lectures.

The worker's initial beliefs induce a joint probability distribution over a potentially infinite sequence of draws $w_0, w_1, \ldots$.

Bayes' law is simply an application of laws of probability to compute the conditional distribution of the tth draw w_t conditional on $[w_0, \ldots, w_{t-1}]$.

After our worker puts a subjective probability π_{-1} on nature having selected distribution F, we have in effect assumed from the start that the decision maker *knows* the joint distribution for the process $\{w_t\}_{t=0}^{\infty}$.

We assume that the worker also knows the laws of probability theory.

A respectable view is that Bayes' law is less a theory of learning than a statement about the consequences of information inflows for a decision maker who thinks he knows the truth (i.e., a joint probability distribution) from the beginning.

29.5.1 Mechanical details again

At time 0 *before* drawing a wage offer, the worker attaches probability $\pi_{-1} \in (0,1)$ to the distribution being F.

Before drawing a wage at time 0, the worker thus believes that the density of w_0 is

$$h(w_0|\pi_{-1}) = \pi_{-1} f(w_0) + (1-\pi_{-1}) g(w_0).$$

Here and below $h(w|\pi) \equiv \pi f(w) + (1-\pi) g(w)$ denotes the mixture of f and g with weight π on f.

Recall that $q \in \{f, g\}$ indicates whether nature chose permanently to draw from distribution f or from distribution g.

After drawing w_0, the worker uses Bayes' law to deduce that the posterior probability $\pi_0 = \text{Prob}(q = f|w_0)$ that the density is $f(w)$ is

$$\pi_0 = \frac{\pi_{-1} f(w_0)}{\pi_{-1} f(w_0) + (1-\pi_{-1}) g(w_0)}.$$

More generally, after making the tth draw and having observed $w_t, w_{t-1}, \ldots, w_0$, the worker believes that the probability that w_{t+1} is being drawn from distribution F is

$$\pi_t = \pi_t(w_t|\pi_{t-1}) \equiv \frac{\pi_{t-1} f(w_t)/g(w_t)}{\pi_{t-1} f(w_t)/g(w_t) + (1-\pi_{t-1})} \tag{29.7}$$

or

$$\pi_t = \frac{\pi_{t-1} \ell(w_t)}{\pi_{t-1} \ell(w_t) + 1 - \pi_{t-1}}$$

and that the density of w_{t+1} conditional on $w_t, w_{t-1}, \ldots, w_0$ is

$$h(w_{t+1}|\pi_t) = \pi_t f(w_{t+1}) + (1-\pi_t) g(w_{t+1}).$$

Notice that

$$\begin{aligned} E(\pi_t|\pi_{t-1}) &= \int \Big[\frac{\pi_{t-1} f(w)}{\pi_{t-1} f(w) + (1-\pi_{t-1}) g(w)}\Big] \Big[\pi_{t-1} f(w) + (1-\pi_{t-1}) g(w)\Big] dw \\ &= \pi_{t-1} \int f(w) dw \\ &= \pi_{t-1}, \end{aligned}$$

so that the process π_t is a *martingale*.

Indeed, it is a *bounded martingale* because each π_t, being a probability, is between 0 and 1.

In the first line in the above string of equalities, the term in the first set of brackets is just π_t as a function of w_t, while the term in the second set of brackets is the density of w_t conditional on $w_{t-1}, \ldots, w_0$ or equivalently conditional on the *sufficient statistic* π_{t-1} for $w_{t-1}, \ldots, w_0$.

Notice that here we are computing $E(\pi_t|\pi_{t-1})$ under the *subjective* density described in the second term in brackets.

Because $\{\pi_t\}$ is a bounded martingale sequence, it follows from the *martingale convergence theorem* that π_t converges almost surely to a random variable in $[0,1]$.

Practically, this means that probability one is attached to sample paths $\{\pi_t\}_{t=0}^{\infty}$ that converge.

According to the theorem, different sample paths can converge to different limiting values.

Thus, let $\{\pi_t(\omega)\}_{t=0}^{\infty}$ denote a particular sample path indexed by a particular $\omega \in \Omega$.

We can think of nature as drawing an $\omega \in \Omega$ from a probability distribution $\text{Prob}(\Omega)$ and then generating a single realization (or *simulation*) $\{\pi_t(\omega)\}_{t=0}^{\infty}$ of the process.

The limit points of $\{\pi_t(\omega)\}_{t=0}^{\infty}$ as $t \to +\infty$ are realizations of a random variable that is swept out as we sample ω from Ω and construct repeated draws of $\{\pi_t(\omega)\}_{t=0}^{\infty}$.

Let

$$\pi_\infty(\omega) = \lim_{t \to +\infty} \pi_t(\omega).$$

The martingale convergence theorem does not say where $\pi_t(\omega)$ goes.

Two facts settle that.

First, π_t is the conditional expectation $E[\mathbf{1}\{q = f\} \mid w^t]$ of the indicator of the event $q = f$.

Second, an observer of the infinite history $w_0, w_1, \ldots$ can determine q.

By the strong law of large numbers, $\frac{1}{t+1} \log L(w^t)$ converges to $E_f[\log \ell(w)] = KL(f,g) > 0$ when $q = f$ and to $E_g[\log \ell(w)] = -KL(g,f) < 0$ when $q = g$.

Doob's martingale convergence theorem [Doob, 1953] then implies that π_t converges with probability one to $E[\mathbf{1}\{q = f\} \mid w_0, w_1, \ldots] = \mathbf{1}\{q = f\}$.

Thus, for realizations along which nature drew f, $\lim_{t \to +\infty} \pi_t(\omega) = 1$, while for realizations along which nature drew g, $\lim_{t \to +\infty} \pi_t(\omega) = 0$.

The section *Drilling down a little bit* and Exercise Exercise 29.9.2 give a second proof that $\pi_\infty(\omega) \in \{0, 1\}$.

Now let's remember that $\{\pi_t\}_{t=0}^{\infty}$ is a martingale and apply the law of iterated expectations.

The law of iterated expectations implies

$$E_t \pi_{t+j} = \pi_t$$

and in particular

$$E_{-1} \pi_{t+j} = \pi_{-1}.$$

Applying the above formula to π_∞, we obtain

$$E_{-1} \pi_\infty(\omega) = \pi_{-1} \tag{29.8}$$

where the mathematical expectation E_{-1} here is taken with respect to the probability measure $\text{Prob}(\Omega)$.

Since the only two values that $\pi_\infty(\omega)$ can take are 1 and 0, we know that for some $\lambda \in [0, 1]$

$$\text{Prob}\Big(\pi_\infty(\omega) = 1\Big) = \lambda, \quad \text{Prob}\Big(\pi_\infty(\omega) = 0\Big) = 1 - \lambda$$

and consequently that

$$E_{-1} \pi_\infty(\omega) = \lambda \cdot 1 + (1 - \lambda) \cdot 0 = \lambda$$

Combining this equation with equation (29.8), we deduce that the probability that $\text{Prob}(\Omega)$ attaches to $\pi_\infty(\omega)$ being 1 must be π_{-1}.

The same answer follows from $\pi_\infty = \mathbf{1}\{q = f\}$ and $\text{Prob}(q = f) = \pi_{-1}$.

Thus, under the worker's subjective distribution, π_{-1} of the sample paths of $\{\pi_t\}$ will converge pointwise to 1 and $1 - \pi_{-1}$ of the sample paths will converge pointwise to 0.

29.5.2 Some simulations

Let's watch the martingale convergence theorem at work in some simulations of our learning model under the worker's subjective distribution.

Let us simulate $\{\pi_t\}_{t=-1}^{T-1}$, $\{w_t\}_{t=0}^{T-1}$ paths where for each $t \geq 0$, w_t is drawn from the subjective distribution

$$h(w_t|\pi_{t-1}) = \pi_{t-1} f(w_t) + (1 - \pi_{t-1}) g(w_t)$$

This sampling scheme gives the history w^t the same joint distribution as drawing q once and for all with Prob$(q = f) = \pi_{-1}$ and then drawing $w_0, w_1, \ldots$ IID from q:

$$\prod_{s=0}^{t} h(w_s|\pi_{s-1}) = \pi_{-1} f(w^t) + (1 - \pi_{-1}) g(w^t).$$

Exercise Exercise 29.9.1 asks you to prove and check this.

Per-period mixing with a *fixed* weight α gives a different joint density, $\prod_{s=0}^{t} h(w_s|\alpha)$; that is the timing protocol of *Incorrect Models*.

We'll plot a large sample of paths.

Our code tracks log odds of $q = f$ and converts them to probabilities.

```
@jit
def martingale_simulate(π_prior, rng, N=5000, T=200):
    """
    Simulate N paths of w_0, ..., w_{T-1} from the subjective distribution.
    Column 0 of π_path holds the prior π_{-1}; column t+1 holds π_t.
    """
    π_path = np.empty((N,T+1))
    w_path = np.empty((N,T))
    π_path[:,0] = π_prior

    for n in range(N):
        log_odds = np.log(π_prior) - np.log(1 - π_prior)
        π = π_prior
        for t in range(T):
            # draw w from h(w | π)
            if rng.random() <= π:
                w = rng.beta(F_a, F_b)
            else:
                w = rng.beta(G_a, G_b)
            log_odds = update(log_odds, f(w) / g(w))
            π = 1 / (1 + np.exp(-log_odds))
            π_path[n,t+1] = π
            w_path[n,t] = w

    return π_path, w_path

def fraction_0_1(π_prior, rng, N, T, decimals):

    π_path, w_path = martingale_simulate(π_prior, rng, N=N, T=T)
    values, counts = np.unique(np.round(π_path[:,-1], decimals=decimals), return_
↪counts=True)
    return values, counts

def create_table(π_priors, rng, N=10000, T=500, decimals=2):
```

(continues on next page)

(continued from previous page)

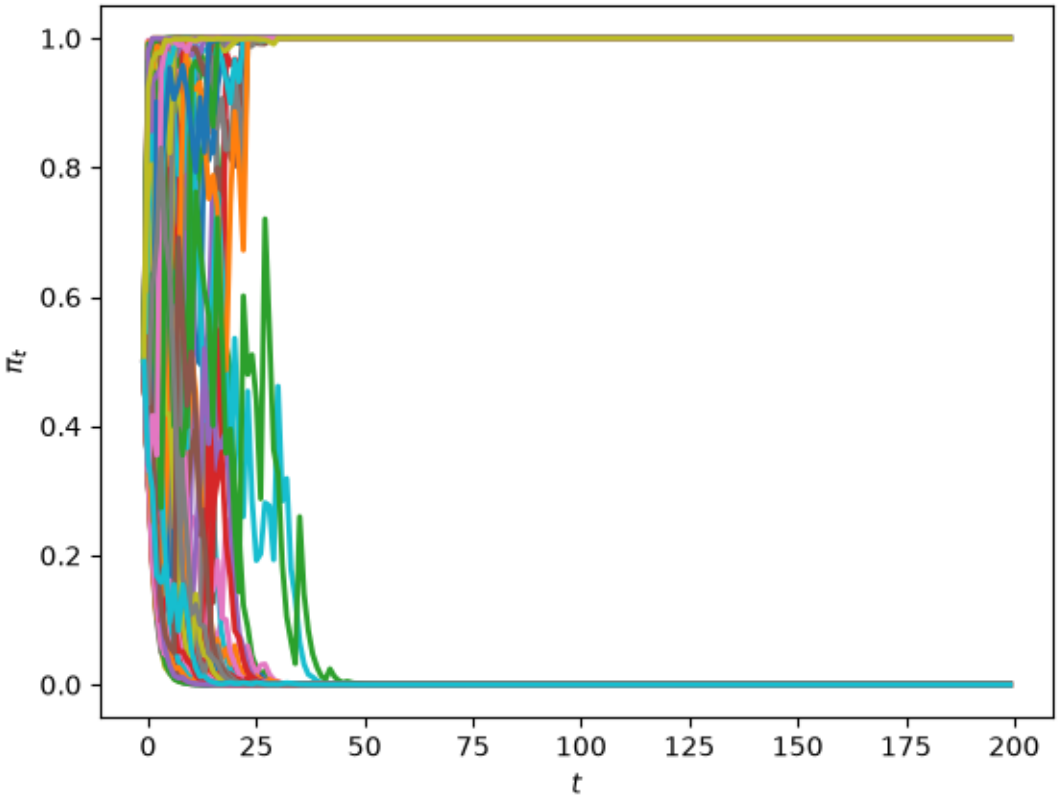
```
    outcomes = []
    for π_prior in π_priors:
        values, counts = fraction_0_1(π_prior, rng, N=N, T=T, decimals=decimals)
        freq = counts/N
        outcomes.append(dict(zip(values, freq)))
    table = pd.DataFrame(outcomes).sort_index(axis=1).fillna(0)
    table.index = π_priors
    return table

# simulate
T = 200
π_prior = .5

π_path, w_path = martingale_simulate(π_prior=π_prior, rng=rng, T=T, N=10000)
```

```
fig, ax = plt.subplots()
for i in range(100):
    ax.plot(range(-1, T), π_path[i, :], lw=2)

ax.set_xlabel('$t$')
ax.set_ylabel(r'$\pi_t$')
plt.show()
```



The above graph indicates that

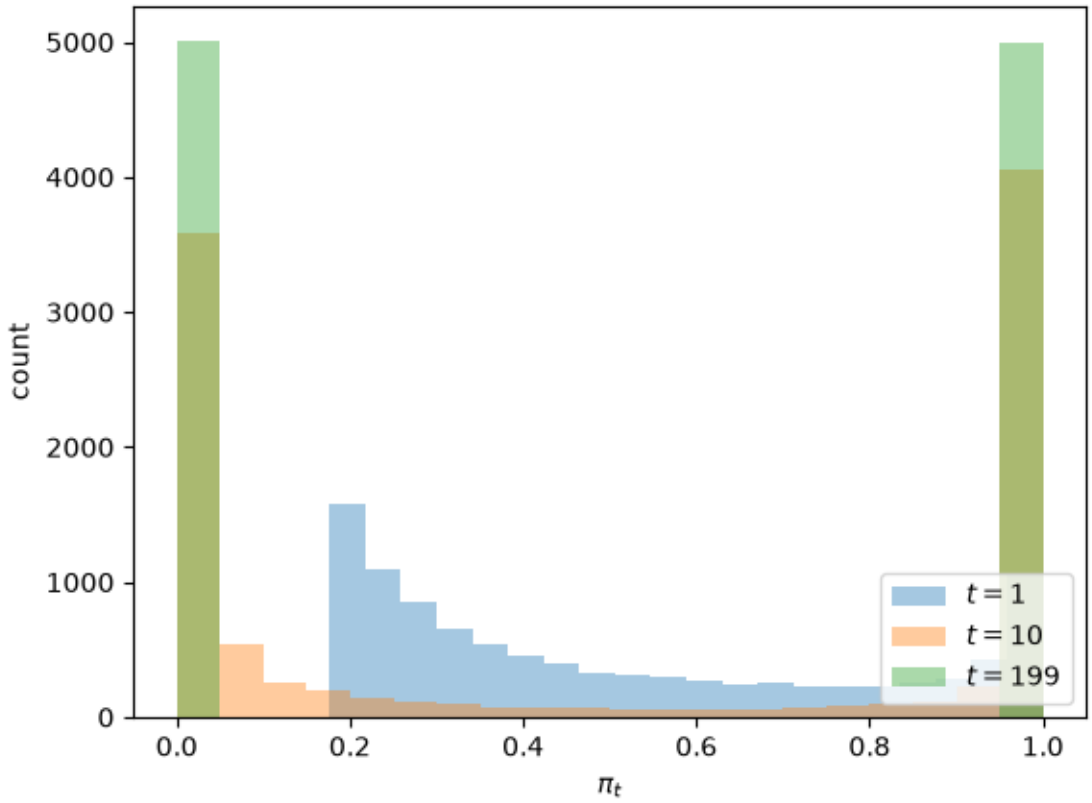
- each of the paths converges
- some of the paths converge to 1

- some of the paths converge to 0
- none of the paths converge to a limit point not equal to 0 or 1

Convergence actually occurs pretty fast, as the following graph of the cross-ensemble distribution of π_t for $t = 1, 10,$ and 199 indicates.

```
fig, ax = plt.subplots()
for t in [1, 10, T-1]:
    ax.hist(π_path[:, t+1], bins=20, alpha=0.4, label=f'$t={t}$')

ax.set_ylabel('count')
ax.set_xlabel(r'$\pi_t$')
ax.legend(loc='lower right')
plt.show()
```



The next cell reports the fractions of paths with π_{199} above 0.99, below 0.01, and in between.

```
def report_fractions(π_path):
    π_T = π_path[:, -1]
    print(f"fraction with π_T > 0.99:         {np.mean(π_T > 0.99):.4f}")
    print(f"fraction with π_T < 0.01:         {np.mean(π_T < 0.01):.4f}")
    print(f"fraction with 0.01 <= π_T <= 0.99: {np.mean((π_T >= 0.01) & (π_T <= 0.
↪99)):.4f}")

report_fractions(π_path)
```

```
fraction with π_T > 0.99:         0.4991
fraction with π_T < 0.01:         0.5009
fraction with 0.01 <= π_T <= 0.99: 0.0000
```

Evidently, by $t = 199$, π_t has converged to either 0 or 1 along almost every path.

The fraction of paths that have converged to 1 is close to .5.

The fraction of paths that have converged to 0 is also close to .5.

Does the fraction .5 ring a bell?

Yes, it does: it equals the value of $\pi_{-1} = .5$ that we used to generate each sequence in the ensemble.

So let's change π_{-1} to .3 and watch what happens to the distribution of the ensemble of π_t's for various t's.

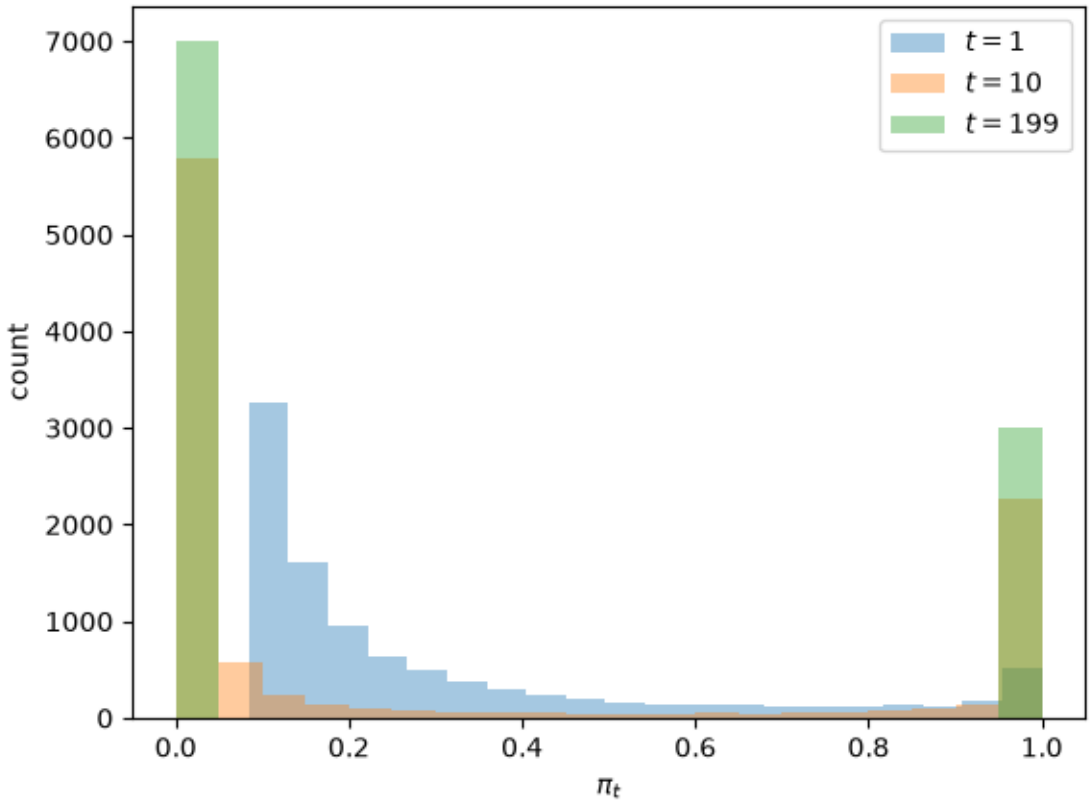
```
# simulate
T = 200
π_prior = .3

π_path3, w_path3 = martingale_simulate(π_prior=π_prior, rng=rng, T=T, N=10000)
```

```
fig, ax = plt.subplots()
for t in [1, 10, T-1]:
    ax.hist(π_path3[:, t+1], bins=20, alpha=0.4, label=f'$t={t}$')

ax.set_ylabel('count')
ax.set_xlabel(r'$\pi_t$')
ax.legend(loc='upper right')
plt.show()

report_fractions(π_path3)
```



```
fraction with π_T > 0.99:         0.2996
fraction with π_T < 0.01:         0.7004
fraction with 0.01 <= π_T <= 0.99: 0.0000
```

For the preceding ensemble that assumed $\pi_{-1} = .5$, the following graph shows two paths of w_t's and the π_t sequences that gave rise to them.

We choose the path with the largest π_{199} and the path with the smallest π_{199}.

Draws $w \in [0.524, 0.999]$ have $\ell(w) < 1$ and push π_t downward (see *Exchangeability and Bayesian Updating*).

The path along which π_t goes to 1 has systematically lower w_t's than the path along which π_t goes to 0.

The luck of the draw early in a simulation pushes the subjective distribution to draw from F more frequently along a sample path, and this pushes π_t toward 1.

Because predictive sampling has the same joint distribution as drawing q once and for all, a path along which $\pi_t \to 1$ looks like a sequence of IID draws from f, and a path along which $\pi_t \to 0$ looks like a sequence of IID draws from g.

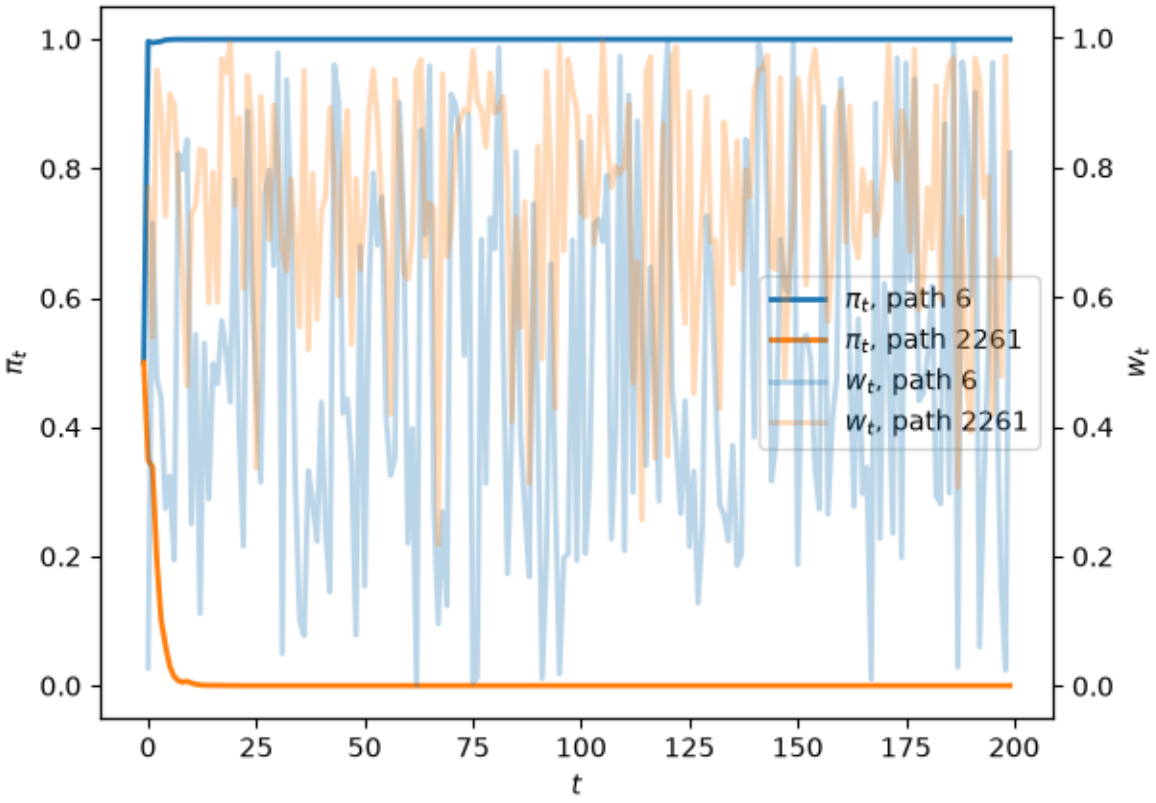
```
j_up, j_down = np.argmax(π_path[:, -1]), np.argmin(π_path[:, -1])

fig, ax = plt.subplots()
ax2 = ax.twinx()
for i, j in enumerate([j_up, j_down]):
    ax.plot(range(-1, T), π_path[j, :], color=colors[i],
            label=fr'$\pi_t$, path {j}', lw=2)
    ax2.plot(range(T), w_path[j, :], color=colors[i],
             label=fr'$w_t$, path {j}', alpha=0.3, lw=2)

ax.set_xlabel('$t$')
ax.set_ylabel(r'$\pi_t$')
ax2.set_ylabel(r'$w_t$')
handles = ax.get_legend_handles_labels()[0] + ax2.get_legend_handles_labels()[0]
ax.legend(handles=handles, loc='center right')
plt.show()

for name, j in [('π_t → 1', j_up), ('π_t → 0', j_down)]:
    print(f"path with {name}: mean of w_t = {w_path[j, :].mean():.3f}")

up = π_path[:, -1] > 0.5
print(f"all paths with π_199 > 0.5: mean of w_t = {w_path[up, :].mean():.3f}")
print(f"all paths with π_199 < 0.5: mean of w_t = {w_path[~up, :].mean():.3f}")
print(f"means of f and g: {F_a / (F_a + F_b):.3f}, {G_a / (G_a + G_b):.3f}")
```




```
path with π_t → 1: mean of w_t = 0.496
path with π_t → 0: mean of w_t = 0.755
all paths with π_199 > 0.5: mean of w_t = 0.500
all paths with π_199 < 0.5: mean of w_t = 0.714
means of f and g: 0.500, 0.714
```

Across all paths, the mean of w_t along paths with $\pi_{199} > 0.5$ is close to the mean of f, and the mean of w_t along paths with $\pi_{199} < 0.5$ is close to the mean of g.

29.6 Initial prior is verified by paths drawn from subjective conditional densities

Now let's use our Python code to generate a table that checks out our earlier claims about the probability distribution of the pointwise limits $\pi_\infty(\omega)$.

We'll use our simulations to generate a histogram of this distribution.

In the following table, the left column in bold face reports an assumed value of π_{-1}.

The second column reports the fraction of $N = 10000$ simulations for which π_t had converged to 0 at the terminal date $T = 500$ for each simulation.

The third column reports the fraction of $N = 10000$ simulations for which π_t had converged to 1 at the terminal date $T = 500$ for each simulation.

```
# create table
table = create_table(list(np.linspace(0,1,11)), rng, N=10000, T=500)
table
```

```
        0.0     1.0
0.0  1.0000  0.0000
0.1  0.9010  0.0990
0.2  0.7991  0.2009
0.3  0.6962  0.3038
0.4  0.6033  0.3967
0.5  0.5034  0.4966
0.6  0.3991  0.6009
0.7  0.2936  0.7064
0.8  0.2047  0.7953
0.9  0.0955  0.9045
1.0  0.0000  1.0000
```

The fraction of simulations for which π_t had converged to 1 is indeed always close to π_{-1}, as anticipated.

29.7 Drilling down a little bit

To understand how the local dynamics of π_t behaves, it is enlightening to consult the variance of π_t conditional on π_{t-1}.

Under the subjective distribution this conditional variance is defined as

$$\sigma^2(\pi_t|\pi_{t-1}) = \int \Big[\frac{\pi_{t-1}f(w)}{\pi_{t-1}f(w) + (1-\pi_{t-1})g(w)} - \pi_{t-1}\Big]^2 \Big[\pi_{t-1}f(w) + (1-\pi_{t-1})g(w)\Big] dw$$

We can use a Monte Carlo simulation to approximate this conditional variance.

We approximate it for a grid of points $\pi_{t-1} \in [0, 1]$.

Then we'll plot it.

```
@jit
def compute_cond_var(π, rng, mc_size=int(1e6)):
    # Create Monte Carlo draws
    mc_draws = np.zeros(mc_size)

    for i in range(mc_size):
        if rng.random() <= π:
            mc_draws[i] = rng.beta(F_a, F_b)
        else:
            mc_draws[i] = rng.beta(G_a, G_b)

    dev = π*f(mc_draws)/(π*f(mc_draws) + (1-π)*g(mc_draws)) - π
    return np.mean(dev**2)

π_array = np.linspace(0, 1, 40)
cond_var_array = []

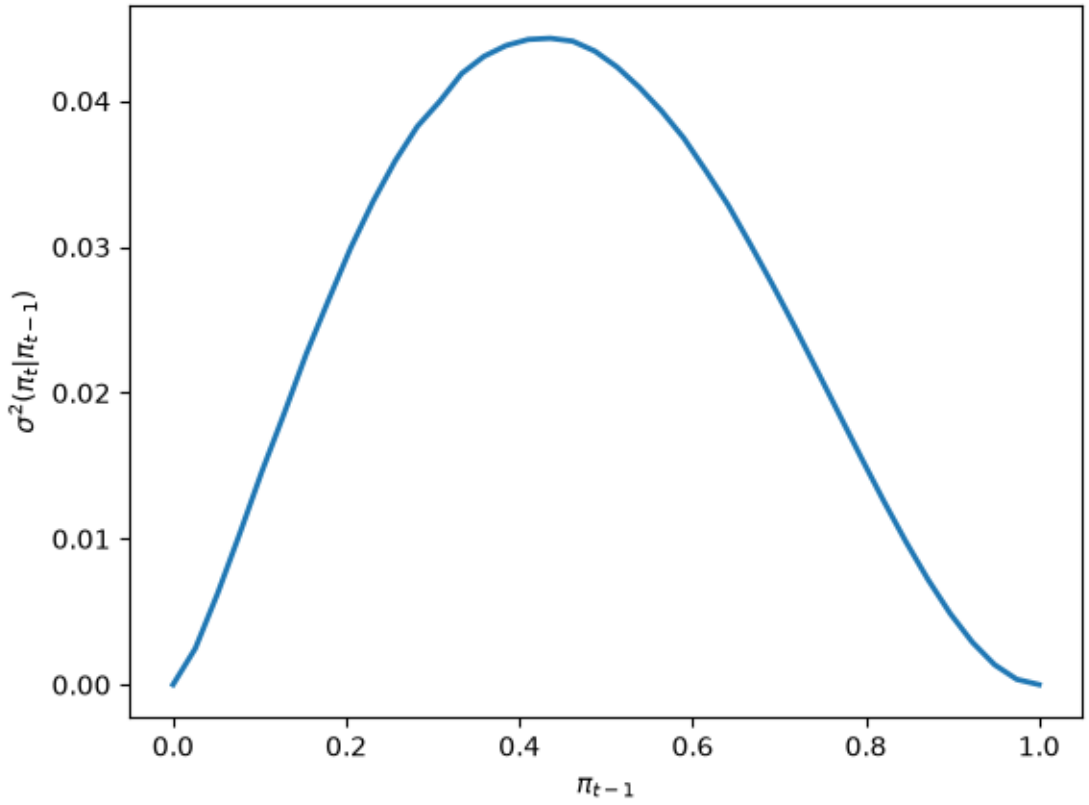
for π in π_array:
    cond_var_array.append(compute_cond_var(π, rng))

fig, ax = plt.subplots()
ax.plot(π_array, cond_var_array, lw=2)
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel(r'$\pi_{t-1}$')
ax.set_ylabel(r'$\sigma^{2}(\pi_{t}\vert \pi_{t-1})$')
plt.show()
```



The shape of the conditional variance as a function of π_{t-1} is informative about the behavior of sample paths of $\{\pi_t\}$.

Notice how the conditional variance approaches 0 for π_{t-1} near either 0 or 1.

The conditional variance is nearly zero only when the agent is almost sure that w_t is drawn from F, or is almost sure it is drawn from G.

The conditional variance is positive at every $\pi_{t-1} \in (0, 1)$.

Because π_t is a martingale, the expected sum of the conditional variances along a path is bounded:

$$E_{-1} \sum_{t=0}^{\infty} \sigma^2(\pi_t|\pi_{t-1}) = E_{-1} \sum_{t=0}^{\infty} (\pi_t - \pi_{t-1})^2 = E_{-1}\pi_\infty^2 - \pi_{-1}^2 \leq \pi_{-1}(1 - \pi_{-1}).$$

So $\sigma^2(\pi_t|\pi_{t-1}) \to 0$ along almost every path, and π_t must approach a point at which the conditional variance is zero, namely 0 or 1.

Exercise Exercise 29.9.2 fills in the details.

29.8 Related lectures

This lecture has been devoted to building some useful infrastructure that will help us understand inferences that are the foundations of results described in *Job Search IX: Search with Learning*, *A Problem that Stumped Milton Friedman*, and *Bayesian versus Frequentist Decision Rules*.

Other lectures that build on it include

- *Exchangeability and Bayesian Updating*, which describes the exchangeable joint distribution of $w_0, w_1, \ldots$ that underlies this lecture's subjective distribution
- *Incorrect Models*, which studies Bayesian learning when nature mixes f and g anew each period
- *Merging of Opinions: The Blackwell–Dubins Theorem*, which studies when the posteriors of two Bayesians with different priors merge
- *Heterogeneous Beliefs and Financial Markets*, which studies likelihood ratio processes and Bayesian learning in further settings
- *A Bayesian Formulation of Friedman and Wald's Problem*, which uses the posterior π_t as the state variable of a sequential decision problem

29.9 Exercises

Exercise 29.9.1

Consider two ways to simulate a history $w_0, w_1, \ldots, w_{T-1}$:

- method (i): draw q once and for all with Prob$(q = f) = \pi_{-1}$, then draw $w_0, w_1, \ldots$ IID from q
- method (ii): draw each w_t from the predictive density $h(w_t|\pi_{t-1})$, as in `martingale_simulate`

1. Prove that $\prod_{s=0}^{t} h(w_s|\pi_{s-1}) = \pi_{-1} f(w^t) + (1-\pi_{-1}) g(w^t)$, so that the two methods give the same joint distribution of w^t for every t.
2. With $\pi_{-1} = 0.3$, simulate 10,000 paths with each method. Compare the distributions of π_t at $t = 1, 10, 100$ with two-sample Kolmogorov–Smirnov tests.
3. Under method (i), check that π_t approaches 1 along paths with $q = f$ and approaches 0 along paths with $q = g$.

Solution

Part 1.

Let $D_s = \pi_{-1} f(w^s) + (1-\pi_{-1}) g(w^s)$ for $s \geq 0$, and let $D_{-1} = 1$.

By (29.3), $\pi_{s-1} = \pi_{-1} f(w^{s-1})/D_{s-1}$ and $1 - \pi_{s-1} = (1-\pi_{-1}) g(w^{s-1})/D_{s-1}$.

Therefore

$$h(w_s|\pi_{s-1}) = \frac{\pi_{-1} f(w^{s-1}) f(w_s) + (1-\pi_{-1}) g(w^{s-1}) g(w_s)}{D_{s-1}} = \frac{D_s}{D_{s-1}}.$$

The product over $s = 0, \ldots, t$ telescopes to $D_t / D_{-1} = D_t$.

Part 2.

Both methods produce paths of log odds, and we compare log odds.

The Kolmogorov–Smirnov statistic is invariant to the increasing transformation that maps log odds into probabilities.

```
@jit
def draw_q_once(π_prior, rng, N, T):
    """
    Method (i). Returns log odds of π_t (column t+1 holds π_t)
    and the indicator of q = f for each path.
    """
    log_odds = np.empty((N, T+1))
    q_is_f = np.empty(N, dtype=np.bool_)
    for n in range(N):
        q_is_f[n] = rng.random() < π_prior
        log_odds[n, 0] = np.log(π_prior) - np.log(1 - π_prior)
        for t in range(T):
            if q_is_f[n]:
                w = rng.beta(F_a, F_b)
            else:
                w = rng.beta(G_a, G_b)
            log_odds[n, t+1] = update(log_odds[n, t], f(w) / g(w))
    return log_odds, q_is_f


@jit
def predictive_sampling(π_prior, rng, N, T):
    "Method (ii). Returns log odds of π_t (column t+1 holds π_t)."
    log_odds = np.empty((N, T+1))
    for n in range(N):
        log_odds[n, 0] = np.log(π_prior) - np.log(1 - π_prior)
        for t in range(T):
            π = 1 / (1 + np.exp(-log_odds[n, t]))
            if rng.random() < π:
                w = rng.beta(F_a, F_b)
            else:
                w = rng.beta(G_a, G_b)
            log_odds[n, t+1] = update(log_odds[n, t], f(w) / g(w))
    return log_odds


π_prior, N_ex, T_ex = 0.3, 10_000, 200
rng_ex = np.random.default_rng(1)
lo_i, q_is_f = draw_q_once(π_prior, rng_ex, N_ex, T_ex)
lo_ii = predictive_sampling(π_prior, rng_ex, N_ex, T_ex)

for t in [1, 10, 100]:
    res = ks_2samp(lo_i[:, t+1], lo_ii[:, t+1])
    print(f"t = {t:3d}: KS statistic = {res.statistic:.4f}, p-value = {res.pvalue:.
↪3f}")
```

```
t =   1: KS statistic = 0.0182, p-value = 0.073
t =  10: KS statistic = 0.0143, p-value = 0.258
t = 100: KS statistic = 0.0125, p-value = 0.415
```

At the 5 percent level, the tests do not reject the hypothesis that the two methods produce the same distribution of π_t.

Part 3.

```
π_T = expit(lo_i[:, -1])
print(f"fraction of paths with q = f: {q_is_f.mean():.4f}")
print(f"among paths with q = f, fraction with π_{T_ex-1} > 0.99: {np.mean(π_T[q_is_
↪f] > 0.99):.4f}")
print(f"among paths with q = g, fraction with π_{T_ex-1} < 0.01: {np.mean(π_T[~q_
↪is_f] < 0.01):.4f}")
```

```
fraction of paths with q = f: 0.3065
among paths with q = f, fraction with π_199 > 0.99: 1.0000
among paths with q = g, fraction with π_199 < 0.01: 1.0000
```

Under method (i), π_t discovers q along every simulated path, as Doob's theorem asserts.

Exercise 29.9.2

Let $\Delta_t = \pi_t - \pi_{t-1}$ for $t \geq 0$.

1. Show that $E[\Delta_t^2 \mid w^{t-1}] = \sigma^2(\pi_t|\pi_{t-1})$ and that
$$\sigma^2(\pi_t|\pi_{t-1}) = \pi_{t-1}^2(1-\pi_{t-1})^2 \int \frac{(f(w)-g(w))^2}{h(w|\pi_{t-1})} dw.$$
2. Use the martingale property to show that $E_{-1}\sum_{s=0}^{t}\Delta_s^2 = E_{-1}\pi_t^2 - \pi_{-1}^2$. Deduce that $E_{-1}\sum_{t=0}^{\infty}\Delta_t^2 = E_{-1}\pi_\infty^2 - \pi_{-1}^2 \leq \pi_{-1}(1-\pi_{-1})$, with equality when $\pi_\infty \in \{0,1\}$ with probability one.
3. Verify the equality by simulation for $\pi_{-1} = 0.3$ and $\pi_{-1} = 0.5$.
4. Use parts 1 and 2 to give a second proof that $\pi_\infty \in \{0,1\}$ with probability one.

Solution

Part 1.

Conditional on w^{t-1}, w_t has density $h(w_t|\pi_{t-1})$ and $E[\pi_t|w^{t-1}] = \pi_{t-1}$, so $E[\Delta_t^2|w^{t-1}]$ is the conditional variance $\sigma^2(\pi_t|\pi_{t-1})$.

Write $\pi = \pi_{t-1}$.

Then $\pi_t - \pi = \pi f/h(w|\pi) - \pi = \pi(1-\pi)(f-g)/h(w|\pi)$, and

$$\sigma^2(\pi_t|\pi_{t-1}) = \int \frac{\pi^2(1-\pi)^2(f(w)-g(w))^2}{h(w|\pi)^2} h(w|\pi)dw = \pi^2(1-\pi)^2 \int \frac{(f(w)-g(w))^2}{h(w|\pi)} dw.$$

The integral is positive when $f \neq g$, so σ^2 is positive on $(0,1)$.

A random variable with values in $[0,1]$ and mean π has variance at most $\pi(1-\pi)$, so $\sigma^2 \to 0$ as $\pi \to 0$ or $\pi \to 1$.

Part 2.

Since $E[\Delta_t|w^{t-1}] = 0$,

$$E_{-1}\pi_t^2 = E_{-1}(\pi_{t-1}+\Delta_t)^2 = E_{-1}\pi_{t-1}^2 + E_{-1}\Delta_t^2 + 2E_{-1}[\pi_{t-1}E[\Delta_t|w^{t-1}]] = E_{-1}\pi_{t-1}^2 + E_{-1}\Delta_t^2.$$

Summing from 0 to t gives $E_{-1}\sum_{s=0}^{t}\Delta_s^2 = E_{-1}\pi_t^2 - \pi_{-1}^2$.

Let $t \to \infty$.

The left side converges by monotone convergence, and the right side converges to $E_{-1}\pi_\infty^2 - \pi_{-1}^2$ by bounded convergence.

Because $0 \leq \pi_\infty \leq 1$, $\pi_\infty^2 \leq \pi_\infty$, so $E_{-1}\pi_\infty^2 \leq E_{-1}\pi_\infty = \pi_{-1}$.

Equality holds if and only if $\pi_\infty^2 = \pi_\infty$, that is, $\pi_\infty \in \{0, 1\}$, with probability one.

Part 3.

```
for π_prior in [0.3, 0.5]:
    π_path_ex, _ = martingale_simulate(π_prior, rng_ex, N=20_000, T=500)
    sum_sq = np.sum(np.diff(π_path_ex, axis=1)**2, axis=1)
    print(f"π_-1 = {π_prior}: mean of sum of squared increments = {sum_sq.mean():.
↪4f}, "
          f"π_-1 (1 - π_-1) = {π_prior * (1 - π_prior):.4f}")
```

```
π_-1 = 0.3: mean of sum of squared increments = 0.2094, π_-1 (1 - π_-1) = 0.2100
π_-1 = 0.5: mean of sum of squared increments = 0.2500, π_-1 (1 - π_-1) = 0.2500
```

Part 4.

By parts 1 and 2,

$$E_{-1}\sum_{t=0}^{\infty}\sigma^2(\pi_t|\pi_{t-1}) = E_{-1}\sum_{t=0}^{\infty}\Delta_t^2 \leq \pi_{-1}(1-\pi_{-1}) < \infty,$$

so $\sum_t \sigma^2(\pi_t|\pi_{t-1}) < \infty$ and $\sigma^2(\pi_t|\pi_{t-1}) \to 0$ with probability one.

The function $\pi \mapsto \sigma^2$ is continuous on $[0, 1]$, positive on $(0, 1)$, and zero at 0 and 1.

Since $\pi_t \to \pi_\infty$, continuity implies that σ^2 evaluated at π_∞ is zero, so $\pi_\infty \in \{0, 1\}$.

Exercise 29.9.3

Two Bayesians observe the same history w^t but start from different priors, $\pi_{-1}^1 = 0.2$ and $\pi_{-1}^2 = 0.8$.

1. Show that the difference between their log odds is the same at every date.
2. Show that $\frac{1}{t+1}\log|\pi_t^1 - \pi_t^2|$ converges to $-KL(f, g)$ with probability one when nature draws from f and to $-KL(g, f)$ with probability one when nature draws from g.
3. Verify part 2 by simulation.

Merging of Opinions: The Blackwell–Dubins Theorem studies when the opinions of Bayesians with different priors merge in general.

Solution

Part 1.

By (29.6), each Bayesian adds the same $\log \ell(w_t)$ to his log odds at each date, so the difference c between their log odds equals its initial value $\log\frac{0.2}{0.8} - \log\frac{0.8}{0.2} = -2\log 4$.

Part 2.

Let x_i be the log odds of Bayesian i, so that $\pi_t^i = 1/(1 + e^{-x_i})$.

Direct computation gives

$$\pi_t^1 - \pi_t^2 = \frac{\sinh(c/2)}{2\cosh(x_1/2)\cosh(x_2/2)}.$$

Since $\log\cosh(x/2) = |x|/2 - \log 2 + o(1)$ as $|x| \to \infty$ and $x_i = x_i(-1) + \log L(w^t)$,

$$\log|\pi_t^1 - \pi_t^2| = -|\log L(w^t)| + O(1).$$

By the strong law of large numbers, $\frac{1}{t+1}\log L(w^t)$ converges to $KL(f,g)$ under f and to $-KL(g,f)$ under g.

Part 3.

We compute $\log|\pi_t^1 - \pi_t^2|$ from the formula in part 2.

Computing $\pi_t^1 - \pi_t^2$ directly would return zero once both probabilities round to 1 or to 0.

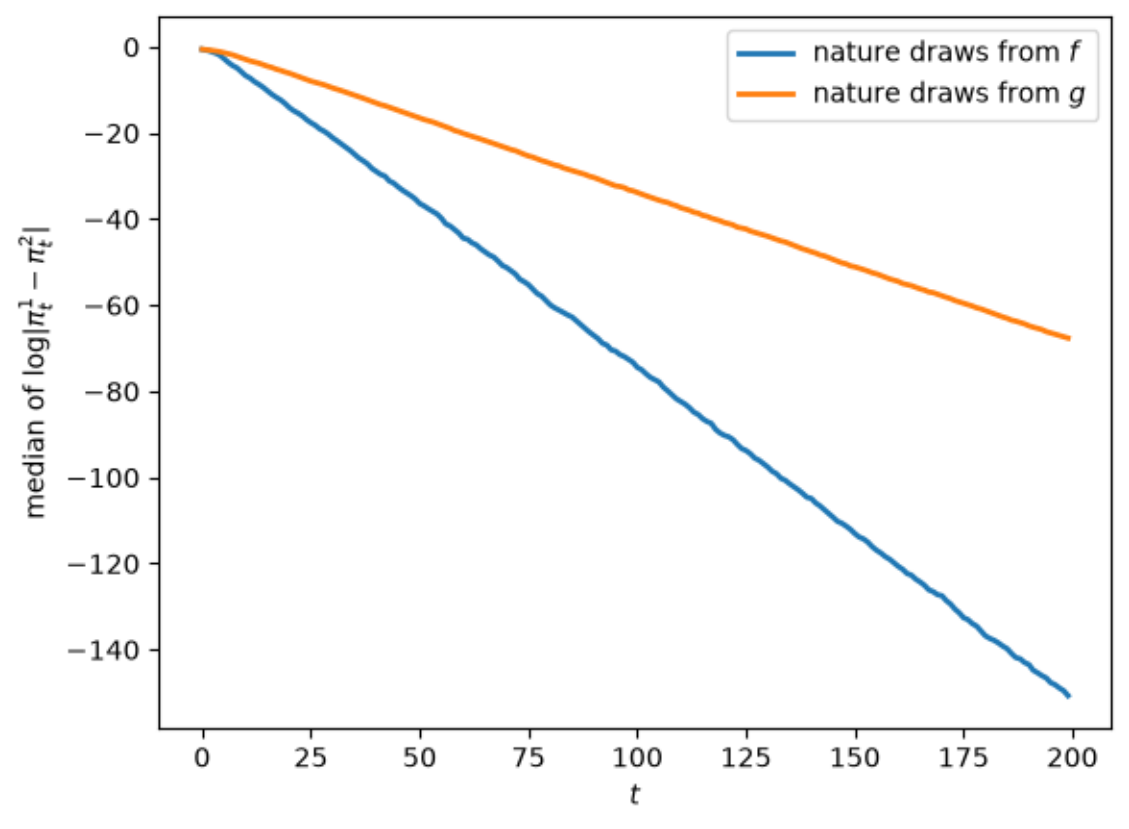
```
def log_gap(x1, x2):
    "log |expit(x1) - expit(x2)|, computed without cancellation."
    log_cosh_half = lambda x: np.logaddexp(x / 2, -x / 2) - np.log(2)
    c = x1 - x2
    return (np.log(np.abs(np.sinh(c / 2))) - np.log(2)
            - log_cosh_half(x1) - log_cosh_half(x2))

KL_fg = compute_KL(f, g)
KL_gf = compute_KL(g, f)

T_m, N_m = 200, 1000
t_m = np.arange(T_m)
fig, ax = plt.subplots()
for name, (a, b), KL in [('f', (F_a, F_b), KL_fg), ('g', (G_a, G_b), KL_gf)]:
    log_L = np.cumsum(np.log(simulate(a, b, rng_ex, T=T_m, N=N_m)), axis=1)
    gap = log_gap(logit(0.2) + log_L, logit(0.8) + log_L)
    med = np.median(gap, axis=0)
    slope = np.polyfit(t_m[50:], med[50:], 1)[0]
    ax.plot(t_m, med, lw=2, label=f'nature draws from ${name}$')
    print(f"nature draws from {name}: slope of median log gap = {slope:.3f}, "
          f"-KL = {-KL:.3f}")

ax.set_xlabel('$t$')
ax.set_ylabel(r'median of $\log|\pi^1_t - \pi^2_t|$')
ax.legend()
plt.show()
```

```
nature draws from f: slope of median log gap = -0.768, -KL = -0.759
nature draws from g: slope of median log gap = -0.344, -KL = -0.344
```



The slopes match $-KL(f,g)$ and $-KL(g,f)$.

The two Bayesians' opinions merge exponentially fast, faster when nature draws from f than when nature draws from g.

CHAPTER

THIRTY

BLACKWELL'S THEOREM ON COMPARING EXPERIMENTS

Contents

- *Blackwell's Theorem on Comparing Experiments*
  - *Overview*
  - *Experiments and stochastic transformations*
  - *Three equivalent criteria*
  - *The main theorem*
  - *Kihlstrom's Bayesian interpretation*
  - *Simulating the Blackwell order with many states*
  - *The DeGroot uncertainty function*
  - *Application 1: demand for information about product quality*
  - *Application 2: sequential experimental design*
  - *The Data Processing Inequality and Coarse-Graining*
  - *Relation to Bayesian likelihood-ratio learning*
  - *Summary*
  - *Exercises*

30.1 Overview

This lecture studies *Blackwell's theorem* [Blackwell, 1951, Blackwell, 1953] on ranking statistical experiments.

Our presentation brings in findings from a Bayesian interpretation of Blackwell's theorem by Kihlstrom [1984].

Blackwell and Kihlstrom study statistical model-selection questions closely related to those encountered in this QuantEcon lecture *Likelihood Ratio Processes and Bayesian Learning*.

To appreciate the connection involved, it is helpful to appreciate how Blackwell's notion of an **experiment** is related to the concept of a "probability distribution" or "parameterized statistical model" appearing in *Likelihood Ratio Processes and Bayesian Learning*

Blackwell studies a situation in which a decision maker wants to know the value of a state s that lives in a space S.

For Blackwell, an **experiment** is a **conditional probability model** $\{\mu(\cdot \mid s) : s \in S\}$, i.e., a family of probability distributions that are conditioned by the same state $s \in S$.

We are free to interpret the "state" as a "parameter" or "parameter vector".

In a two-state case $S = \{s_1, s_2\}$, the two conditional densities $f(\cdot) = \mu(\cdot \mid s_1)$ and $g(\cdot) = \mu(\cdot \mid s_2)$ are the ones used repeatedly in our studies of classical hypothesis testing and Bayesian inference in this QuantEcon lecture *Likelihood Ratio Processes and Bayesian Learning* as well as several other lectures in this suite of QuantEcon lectures.

Kihlstrom [1984] interprets the question *which experiment is more informative?* as asking which conditional probability model allows a Bayesian decision maker with a prior over $\{s_1, s_2\}$ to gather higher expected utility.

We'll use the terms "signal" and "experiment" as synonyms.

Thus, suppose that two signals, $\tilde{x}_\mu$ and $\tilde{x}_\nu$, are both informative about an unknown state $\tilde{s}$.

Signal μ is **at least as informative as** signal ν if every Bayesian decision maker can attain weakly higher expected utility with μ than with ν.

This economic criterion is equivalent to two statistical criteria:

- *Sufficiency* (Blackwell): $\tilde{x}_\nu$ can be generated from $\tilde{x}_\mu$ by an additional randomization.
- *Uncertainty reduction* (DeGroot [1962]): $\tilde{x}_\mu$ lowers expected uncertainty at least as much as $\tilde{x}_\nu$ for every concave uncertainty function.

Kihlstrom's formulation focuses on the *posterior distribution*.

More informative experiments generate posterior distributions that are more dispersed in convex order.

In the two-state case, this becomes a mean-preserving-spread comparison on $[0, 1]$, which can be checked with the integrated-CDF test used for second-order stochastic dominance.

Information and Market Equilibrium and *The Design and Price of Information* build on the results of this lecture.

The lecture proceeds as follows.

1. We define experiments as Markov matrices and garblings as stochastic transformations.
2. We state three equivalent criteria for ranking experiments and Blackwell's theorem.
3. We develop Kihlstrom's Bayesian interpretation through standard experiments, mean-preserving spreads, and convex order.
4. We study DeGroot's uncertainty functions and two applications.
5. We connect Blackwell's order to the data processing inequality and to the likelihood-ratio learning of *Likelihood Ratio Processes and Bayesian Learning*.

Python simulations illustrate each step.

We begin with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import minimize
```

30.2 Experiments and stochastic transformations

30.2.1 The state space and experiments

Let $S = \{s_1, \ldots, s_N\}$ be a finite set of possible states of the world.

An **experiment** is described by the conditional distribution of an observed signal $\tilde{x}$ given the state $\tilde{s}$.

When the signal space is also finite, say $X = \{x_1, \ldots, x_M\}$, an experiment reduces to an $N \times M$ *Markov matrix*

$$\mu = [\mu_{ij}], \qquad \mu_{ij} = \Pr(\tilde{x}_\mu = x_j \mid \tilde{s} = s_i) \geq 0, \quad \sum_{j=1}^{M} \mu_{ij} = 1 \; \forall \, i.$$

Each row i gives the distribution of signals when the true state is s_i.

```
μ = np.array([[0.6, 0.3, 0.1],
              [0.1, 0.3, 0.6]])

Q = np.array([[1.0, 0.0],
              [0.5, 0.5],
              [0.0, 1.0]])

ν = μ @ Q

print("Experiment μ (3 signals, rows sum to 1):")
print(μ)
print("\nStochastic transformation Q (3 × 2):")
print(Q)
print("\nExperiment ν = μ @ Q (2 signals):")
print(ν)
print("\nRow sums μ:", μ.sum(axis=1))
print("Row sums ν:", ν.sum(axis=1))
```

```
Experiment μ (3 signals, rows sum to 1):
[[0.6 0.3 0.1]
 [0.1 0.3 0.6]]

Stochastic transformation Q (3 × 2):
[[1.  0. ]
 [0.5 0.5]
 [0.  1. ]]

Experiment ν = μ @ Q (2 signals):
[[0.75 0.25]
 [0.25 0.75]]

Row sums μ: [1. 1.]
Row sums ν: [1. 1.]
```

30.2.2 Stochastic transformations

A **stochastic transformation** Q maps signals from one experiment to signals from another by further randomization.

In the discrete setting with M input signals and K output signals, Q is an $M \times K$ Markov matrix: $q_{lk} \geq 0$ and $\sum_k q_{lk} = 1$ for every row l.

Definition 30.2.1 (Sufficiency)

Experiment μ is *sufficient for* ν if there exists a stochastic transformation Q (an $M \times K$ Markov matrix) such that

$$\nu = \mu\, Q,$$

meaning that an observer of $\tilde{x}_\mu$ can generate the distribution of $\tilde{x}_\nu$ by passing their signal through Q.

With general signal spaces, Q is a Markov kernel and the condition reads $\nu_i(E) = \int Q(x, E)\, \mu_i(dx)$ for every measurable set E of ν's signals and every state s_i.

We write $\mu \geq \nu$ when μ is sufficient for ν, and we call ν a **garbling** of μ.

If you observe the more informative signal $\tilde{x}_\mu$, then you can always *throw away* information to reproduce a less informative signal.

The reverse generally fails: a garbling cannot restore what it discards.

We can verify this numerically using the two experiments μ and ν defined above.

The function below searches for a stochastic transformation Q that minimizes $\|\nu - \mu\, Q\|$.

If an exact Q exists the residual will be close to zero; otherwise it will be large.

```
def find_stochastic_transform(μ, ν, tol=1e-8):
    """
    Find a row-stochastic matrix Q that minimizes ||ν - μ @ Q||.
    """
    _, M = μ.shape
    _, K = ν.shape

    def unpack(q_flat):
        return q_flat.reshape(M, K)

    def objective(q_flat):
        Q = unpack(q_flat)
        return np.linalg.norm(ν - μ @ Q)**2

    constraints = [
        {"type": "eq", "fun": lambda q_flat,
        row=i: unpack(q_flat)[row].sum() - 1.0}
        for i in range(M)
    ]
    bounds = [(0.0, 1.0)] * (M * K)
    Q0 = np.full((M, K), 1 / K).ravel()

    result = minimize(
        objective,
        Q0,
        method="SLSQP",
        bounds=bounds,
```

(continues on next page)

(continued from previous page)

```
        constraints=constraints,
        options={"ftol": tol, "maxiter": 1_000},
    )

    Q = unpack(result.x)
    residual = np.linalg.norm(ν - μ @ Q)
    return Q, residual

# Forward: find Q such that ν = μ @ Q  (should succeed)
Q_fwd, res_fwd = find_stochastic_transform(μ, ν)
print("Forward (μ to ν):")
print(f"  residual = {res_fwd:.2e}")
print(f"  exact transformation exists: {res_fwd < 1e-6}")

# Reverse: find Q' such that μ = ν @ Q'  (should fail)
Q_rev, res_rev = find_stochastic_transform(ν, μ)
print("\nReverse (ν to μ):")
print(f"  residual = {res_rev:.2e}")
print(f"  exact transformation exists: {res_rev < 1e-6}")
```

```
Forward (μ to ν):
  residual = 5.05e-09
  exact transformation exists: True

Reverse (ν to μ):
  residual = 1.44e-01
  exact transformation exists: False
```

The forward residual is close to zero: a stochastic transformation from μ to ν exists, confirming that μ is sufficient for ν.

The reverse residual is large: no stochastic transformation can recover μ from ν.

Exercise 30.13.1 proves that no stochastic transformation maps ν into μ.

A garbling can only pull posterior beliefs toward the prior, and μ generates posteriors farther from the prior than ν can reach.

30.3 Three equivalent criteria

Blackwell's theorem establishes that three different ways of comparing experiments all turn out to be equivalent.

30.3.1 Criterion 1: the economic criterion

The first criterion compares experiments by their value to decision makers.

Let A be a compact set of actions and $u : A \times S \to \mathbb{R}$ a bounded continuous utility function.

A decision maker observes $x \in X$ and then chooses an action, possibly at random.

A **decision rule** δ assigns to each signal x a probability distribution $\delta(\cdot \mid x)$ over actions.

Let $p = (p_1, \ldots, p_N)$ be the prior over states, and write

$$P = \{(p_1, \ldots, p_N) : p_i \geq 0,\ \textstyle\sum_i p_i = 1\}$$

for the probability simplex.

For fixed A and u, the set of *achievable expected-utility vectors* under experiment μ is

$$B(\mu, A, u) = \left\{ v \in \mathbb{R}^N : v_i = \int_X \int_A u(a, s_i)\, \delta(da \mid x)\, \mu_i(dx) \text{ for some decision rule } \delta \right\}.$$

Each $v \in B(\mu, A, u)$ lists the expected utilities that one decision rule delivers in each state.

Randomized rules make $B(\mu, A, u)$ convex.

They also make the replication argument below exact: a decision maker who sees $\tilde{x}_\mu$ can draw a synthetic $\tilde{x}_\nu$ and then follow any rule designed for ν.

Definition 30.3.1 (Economic criterion)

μ is **at least as informative as** ν in the economic sense if

$$B(\mu, A, u) \supseteq B(\nu, A, u)$$

for every compact action set A and every bounded continuous utility function $u : A \times S \to \mathbb{R}$.

This criterion says that experiment μ is better than experiment ν if anything a decision maker can achieve after seeing ν, they can also achieve after seeing μ.

The reason is that a more informative experiment lets the agent imitate a less informative one by *ignoring* or *garbling* some of the extra information.

But the reverse need not be possible.

So $B(\mu, A, u) \supseteq B(\nu, A, u)$ means that μ gives the decision maker at least as many feasible expected-utility outcomes as ν.

Blackwell's theorem below implies an equivalent comparison of values.

For every prior $p \in P$, every A, and every u, the maximized expected utility $\max_{v \in B(\mu,A,u)} \sum_i p_i v_i$ is at least as large under μ as under ν.

A Bayesian decision maker attains that maximum with a deterministic rule that chooses an optimal action at each posterior; randomization matters only for the comparison of sets.

30.3.2 Criterion 2: the sufficiency criterion

The second criterion is the sufficiency relation of *Definition 30.2.1*: $\mu \geq \nu$ if $\nu = \mu Q$ for some stochastic transformation Q.

30.3.3 Criterion 3: the uncertainty criterion

The third criterion compares experiments by how much they reduce uncertainty about the state.

DeGroot [1962] calls any concave function $U : P \to \mathbb{R}$ an **uncertainty function**.

The prototypical example is Shannon entropy:

$$U(p) = -\sum_{i=1}^{N} p_i \log p_i.$$

Definition 30.3.2 (DeGroot uncertainty criterion)

μ **reduces expected uncertainty at least as much as** ν if, for every prior $p \in P$ and every concave $U : P \to \mathbb{R}$,

$$\int_P U(q)\,\hat{\mu}^p(dq) \;\leq\; \int_P U(q)\,\hat{\nu}^p(dq),$$

where $\hat{\mu}^p$ is the distribution of posterior beliefs induced by experiment μ under prior p.

Every experiment weakly lowers expected uncertainty.

To see why, let $\tilde{\pi} = p^\mu(\tilde{x}_\mu)$ denote the random posterior induced by experiment μ.

Then $\tilde{\pi}$ has distribution $\hat{\mu}^p$, so

$$\mathbb{E}[U(\tilde{\pi})] = \int_P U(q)\,\hat{\mu}^p(dq).$$

The posterior has mean p (*Property 30.5.1* below), so Jensen's inequality gives

$$\mathbb{E}[U(\tilde{\pi})] \leq U(\mathbb{E}[\tilde{\pi}]) = U(p).$$

Hence

$$\int_P U(q)\,\hat{\mu}^p(dq) \leq U(p),$$

so any experiment weakly lowers expected uncertainty.

Criterion 3 quantifies over every prior p.

Blackwell's theorem implies that one prior with full support suffices: if the inequality holds for every concave U at one such prior, it holds at every prior.

Kihlstrom's construction below exploits this fact by working with a fixed prior c, often the uniform prior $(1/N, \ldots, 1/N)$.

30.4 The main theorem

Theorem 30.4.1 (Blackwell's theorem)

The following three conditions are equivalent:

(i) Economic criterion: $B(\mu, A, u) \supseteq B(\nu, A, u)$ for every compact A and every bounded continuous utility function u.

(ii) Sufficiency criterion: There exists a stochastic transformation Q from the signal space of μ to the signal space of ν such that $\nu = \mu Q$.

(iii) Uncertainty criterion: $\int_P U(q)\,\hat{\mu}^p(dq) \leq \int_P U(q)\,\hat{\nu}^p(dq)$ for every concave U, at every prior $p \in P$, or equivalently at one prior with full support.

See also Blackwell [1951], Bohnenblust *et al.* [1949], and DeGroot [1962].

The hard part is the equivalence between the economic and sufficiency criteria.

Sketch (ii $\Rightarrow$ i): If $\nu = \mu Q$, then any decision rule based on $\tilde{x}_\nu$ can be replicated by first observing $\tilde{x}_\mu$, then drawing a synthetic $\tilde{x}_\nu$ from Q, and then applying the same rule.

Sketch (i ⇒ ii): Since $B(\mu, A, u) \supseteq B(\nu, A, u)$ for every A and u, a separating-hyperplane (duality) argument implies that the posterior distribution of μ is a mean-preserving spread of that of ν, and Strassen's theorem [Strassen, 1965] then supplies a posterior-space mean-preserving kernel D sending the standard experiment of ν into that of μ. Passing from these posterior laws back to the original signal spaces then yields the required garbling Q with $\nu = \mu Q$. Thus D is an intermediate randomization on posterior beliefs, not literally the signal-space kernel Q.

Sketch (ii ⇒ iii): Under a garbling, the posterior from the coarser experiment is the conditional expectation of the posterior from the finer experiment, so Jensen's inequality gives the result for every concave U.

Sketch (iii ⇒ ii): The converse, that the inequality for all concave U forces the existence of Q, is proved in [Blackwell, 1953]. Kihlstrom's posterior-based representation makes the geometry transparent.

30.5 Kihlstrom's Bayesian interpretation

30.5.1 Posteriors and standard experiments

The key object in Kihlstrom's analysis is the *posterior belief vector*.

When prior p holds and experiment μ produces signal x, Bayes' rule gives

$$p_i^{\mu}(x) = \Pr(\tilde{s} = s_i \mid \tilde{x}_\mu = x) = \frac{\mu_{ix}\, p_i}{\sum_j \mu_{jx}\, p_j}, \qquad i = 1, \dots, N.$$

The posterior $p^{\mu}(x) \in P$ is a random point in the simplex.

Property 30.5.1 (Mean preservation)

The prior p is the expectation of the posterior:

$$\mathbb{E}[p^{\mu}] = \sum_x \Pr(\tilde{x}_\mu = x)\, p^{\mu}(x) = p.$$

This is sometimes called the *law of iterated expectations for beliefs*.

For a fixed prior c, Kihlstrom's **standard experiment** replaces the raw signals of μ with the posterior beliefs they generate.

Let $\hat{\mu}^c$ denote the distribution over posteriors induced by μ under prior c. Mean preservation implies $\int_P q\, \hat{\mu}^c(dq) = c$.

Two experiments are **informationally equivalent** when they induce the same posterior distribution.

The standard experiment strips away every detail of the signal except its posterior, so it provides a canonical Bayesian representation for comparing experiments.

A stochastic kernel on posterior beliefs lives on the simplex P, whereas a Blackwell garbling Q lives on the original signal space. Kihlstrom's construction uses the former to study convex order and then recovers the latter after passing to standard experiments.

Any two experiments that generate the same distribution over posteriors lead to identical decisions for every Bayesian decision maker, regardless of how different their raw signal spaces may look.

30.5.2 Mean-preserving spreads and Blackwell's order

Kihlstrom's key reformulation is the following.

Theorem 30.5.1 (Kihlstrom's Reformulation)

$\mu \geq \nu$ in Blackwell's sense if and only if $\hat{\mu}^c$ is a **mean-preserving spread** of $\hat{\nu}^c$; that is,

$$\int_P g(p)\,\hat{\mu}^c(dp) \;\geq\; \int_P g(p)\,\hat{\nu}^c(dp)$$

for every convex function $g : P \to \mathbb{R}$.

Equivalently, $\hat{\mu}^c$ is larger than $\hat{\nu}^c$ in convex order.

The prior c can be any prior with full support; the comparison does not depend on which one.

A better experiment spreads posterior beliefs farther from the prior while preserving their mean.

To see this concretely, we define two experiments for the two-state case and compute their posteriors.

```
def compute_posteriors(μ, prior, tol=1e-14):
    """
    Compute the posterior distribution for each signal realisation.
    """
    N, M = μ.shape
    signal_probs = μ.T @ prior
    numerators = μ.T * prior
    posteriors = np.zeros((M, N))
    np.divide(
        numerators,
        signal_probs[:, None],
        out=posteriors,
        where=signal_probs[:, None] > tol,
    )
    return posteriors, signal_probs


def check_mean_preservation(posteriors, signal_probs, prior):
    """Verify E[posterior] == prior."""
    expected_posterior = (posteriors * signal_probs[:, None]).sum(axis=0)
    return expected_posterior, np.allclose(expected_posterior, prior)


N = 2
prior = np.array([0.5, 0.5])

μ_info = np.array([[0.8, 0.2],
                   [0.2, 0.8]])

ν_info = np.array([[0.6, 0.4],
                   [0.4, 0.6]])

post_μ, probs_μ = compute_posteriors(μ_info, prior)
post_ν, probs_ν = compute_posteriors(ν_info, prior)

print("Experiment μ (more informative):\n")
```

(continues on next page)

(continued from previous page)

```
print("Signal probabilities:", probs_μ.round(3))
print("Posteriors (row = signal, col = state):")
print(post_μ.round(3))
mean_μ, ok_μ = check_mean_preservation(post_μ, probs_μ, prior)
print(f"E[posterior] = {mean_μ.round(4)}  (equals prior: {ok_μ})")

print("\n Experiment ν (less informative):\n")
print("Signal probabilities:", probs_ν.round(3))
print("Posteriors:")
print(post_ν.round(3))
mean_ν, ok_ν = check_mean_preservation(post_ν, probs_ν, prior)
print(f"E[posterior] = {mean_ν.round(4)}  (equals prior: {ok_ν})")
```

```
Experiment μ (more informative):

Signal probabilities: [0.5 0.5]
Posteriors (row = signal, col = state):
[[0.8 0.2]
 [0.2 0.8]]
E[posterior] = [0.5 0.5]  (equals prior: True)

 Experiment ν (less informative):

Signal probabilities: [0.5 0.5]
Posteriors:
[[0.6 0.4]
 [0.4 0.6]]
E[posterior] = [0.5 0.5]  (equals prior: True)
```

For $N = 2$ states, the simplex P is the unit interval $[0, 1]$ (the probability of state s_1).

We can directly plot the distribution of posteriors under experiments μ and ν.

```
def plot_posterior_distributions(μ_matrix, ν_matrix, prior,
                                 labels=("μ (more informative)",
                                         "ν (less informative)")):
    """
    For a two-state experiment, plot the distribution of posteriors
    (i.e., the standard experiment distribution) on [0,1].
    """
    posts_μ, probs_μ = compute_posteriors(μ_matrix, prior)
    posts_ν, probs_ν = compute_posteriors(ν_matrix, prior)

    fig, axes = plt.subplots(1, 2, figsize=(11, 4), sharey=False)
    prior_val = prior[0]

    for ax, posts, probs, label in zip(
        axes, [posts_μ, posts_ν], [probs_μ, probs_ν], labels):
        p_s1 = posts[:, 0]
        ax.vlines(p_s1, 0, probs, linewidth=6, color="steelblue", alpha=0.7)
        ax.axvline(prior_val, color="tomato", linestyle="--", linewidth=2,
                   label=f"prior = {prior_val:.2f}")
        ax.set_xlim(0, 1)
        ax.set_xlabel(r"posterior $p(s_1 \mid x)$", fontsize=12)
        ax.set_ylabel("probability mass", fontsize=12)
        mean_post = (p_s1 * probs).sum()
        ax.axvline(mean_post, color="green", linestyle=":", linewidth=2,
```

(continues on next page)

(continued from previous page)

```
                    label=f"E[post] = {mean_post:.2f}")
        ax.text(0.03, 0.94, label, transform=ax.transAxes, va="top")
        ax.legend()

    plt.tight_layout()
    plt.show()

plot_posterior_distributions(μ_info, ν_info, prior)
```

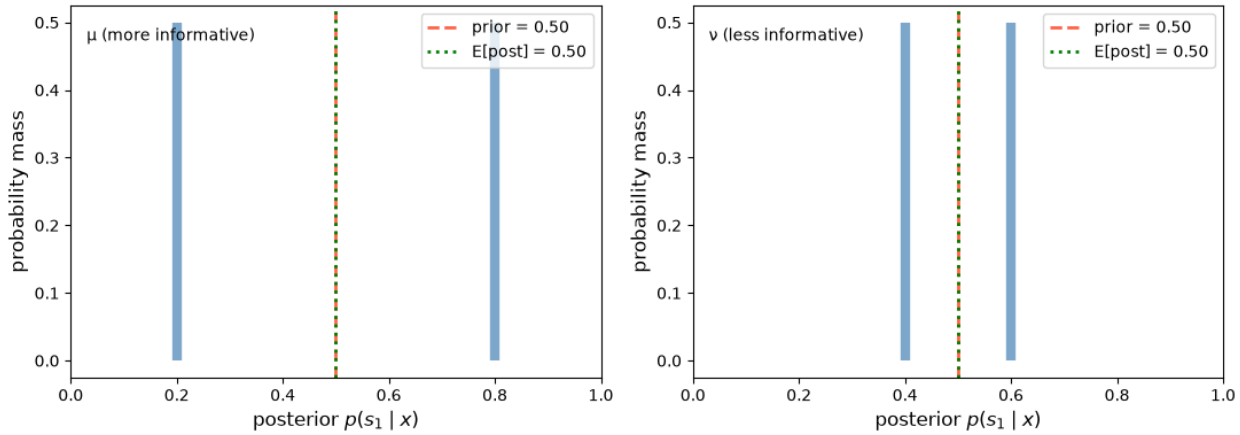


Fig. 30.1: Posterior distributions in the two-state case

This is the mean-preserving spread in action: both distributions have the same mean (equal to the prior), but the more informative experiment μ spreads its posteriors farther apart.

We can verify the mean-preserving spread condition numerically.

The key fact is that, up to an affine term, any convex function can be represented as a mixture of "call option" payoffs $g_t(p) = \max(p - t, 0)$.

Because the two posterior distributions being compared have the same mean, that affine term cancels in the comparison.

So it suffices to check $E[g_t(p^\mu)] \geq E[g_t(p^\nu)]$ for all thresholds $t \in [0, 1]$.

```
def check_mps_convex_functions(μ_matrix, ν_matrix, prior, n_functions=200):
    """
    Verify the mean-preserving spread condition using
    convex functions g(p) = max(p - t, 0).
    """
    posts_μ, probs_μ = compute_posteriors(μ_matrix, prior)
    posts_ν, probs_ν = compute_posteriors(ν_matrix, prior)

    p_μ = posts_μ[:, 0]
    p_ν = posts_ν[:, 0]

    thresholds = np.linspace(0, 1, n_functions)
    diffs = []
    for t in thresholds:
        Eg_μ = (np.maximum(p_μ - t, 0) * probs_μ).sum()
        Eg_ν = (np.maximum(p_ν - t, 0) * probs_ν).sum()
```

(continues on next page)

(continued from previous page)

```
        diffs.append(Eg_μ - Eg_ν)

    fig, ax = plt.subplots(figsize=(8, 4))
    ax.plot(thresholds, diffs, color="steelblue", linewidth=2)
    ax.axhline(0, color="tomato", linestyle="--", linewidth=2)
    ax.fill_between(thresholds, diffs, 0,
                    where=np.array(diffs) >= 0,
                    alpha=0.25, color="steelblue",
                    label="$E[g(p^μ)] - E[g(p^ν)] \\geq 0$")
    ax.set_xlabel("threshold $t$", fontsize=12)
    ax.set_ylabel(r"$E[\max(p-t,0)]$ difference", fontsize=12)
    ax.legend(fontsize=11)
    plt.tight_layout()
    plt.show()

    all_non_negative = all(d >= -1e-10 for d in diffs)
    print(f"μ is a mean-preserving spread of ν: {all_non_negative}")
    return diffs

_ = check_mps_convex_functions(μ_info, ν_info, prior)
```

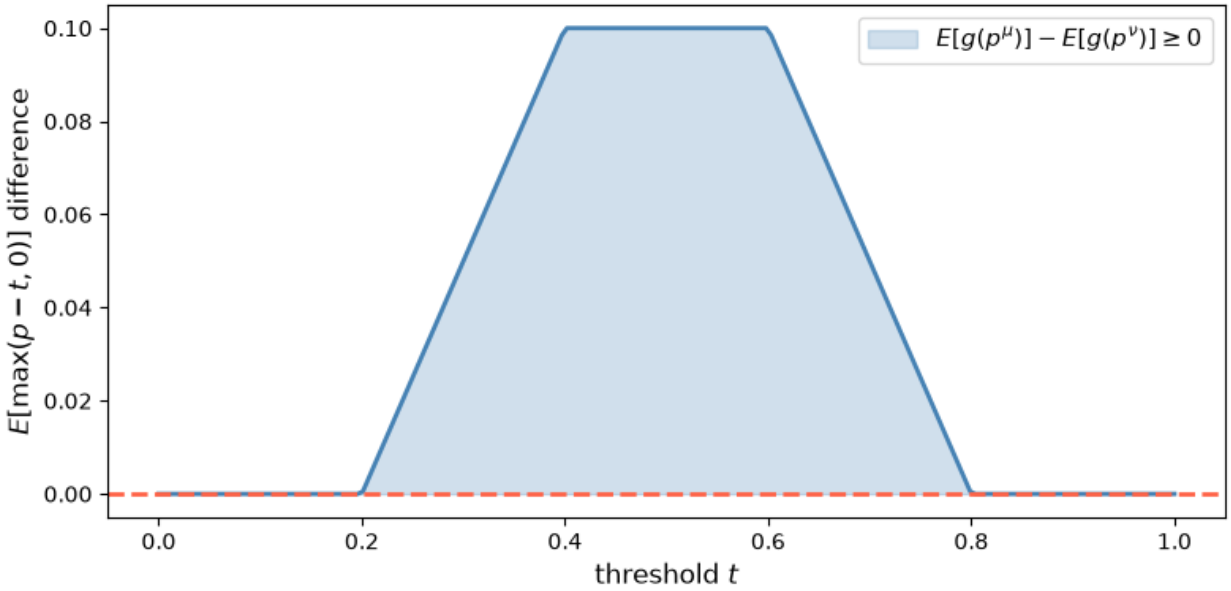


Fig. 30.2: Convex-order check in the two-state case

```
μ is a mean-preserving spread of ν: True
```

The difference $E[g_t(p^\mu)] - E[g_t(p^\nu)]$ is non-negative for every threshold t, confirming that $\hat{\mu}^c$ is a mean-preserving spread of $\hat{\nu}^c$ and therefore $\mu \succeq \nu$ in the Blackwell order.

30.5.3 Convex order and second-order stochastic dominance

A random variable X dominates Y in the **concave order** if $E[u(X)] \geq E[u(Y)]$ for every concave function u.

Taking $u(x) = x$ and $u(x) = -x$ shows that the two means must be equal.

X dominates Y in the concave order if and only if Y is a mean-preserving spread of X.

Second-order stochastic dominance uses only increasing concave u and allows the means to differ; when the means are equal it coincides with the concave order.

Because uncertainty functions are concave, the condition

$$\mathbb{E}[U(p^\mu)] \leq \mathbb{E}[U(p^\nu)] \quad \text{for all concave } U$$

says that the posterior under ν dominates the posterior under μ in the concave order, so that $\hat{\mu}^c$ dominates $\hat{\nu}^c$ in convex order on P.

When $N = 2$, posterior beliefs are scalars in $[0, 1]$, and the comparison reduces to the classical integrated-CDF test.

Given equal means, $\hat{\mu}^c$ is a mean-preserving spread of $\hat{\nu}^c$ if and only if $\int_0^t F_\mu(s)\, ds \geq \int_0^t F_\nu(s)\, ds$ for all $t \in [0, 1]$, where F_μ and F_ν are the CDFs of the posterior on s_1 under each experiment.

Both posterior distributions have mean equal to the prior, so the test applies.

We verify it for the two-state example above.

```
def cdf_data_1d(weights, values):
    """Sort support points and cumulative masses for a discrete distribution."""
    idx = np.argsort(values)
    sorted_vals = values[idx]
    sorted_wts = weights[idx]
    cum_mass = np.cumsum(sorted_wts)
    return sorted_vals, cum_mass


def plot_sosd_posteriors(μ_matrix, ν_matrix, prior):
    """Plot CDFs and integrated CDFs for the posterior-on-s1 distributions."""
    posts_μ, probs_μ = compute_posteriors(μ_matrix, prior)
    posts_ν, probs_ν = compute_posteriors(ν_matrix, prior)

    p_μ = posts_μ[:, 0]
    p_ν = posts_ν[:, 0]

    sv_μ, cm_μ = cdf_data_1d(probs_μ, p_μ)
    sv_ν, cm_ν = cdf_data_1d(probs_ν, p_ν)

    fig, axes = plt.subplots(1, 2, figsize=(11, 4))

    ax = axes[0]
    for sv, cm, lbl, c in [(sv_μ, cm_μ, "μ", "steelblue"),
                           (sv_ν, cm_ν, "ν", "darkorange")]:
        xs = np.concatenate([[0], sv, [1]])
        ys = np.concatenate([[0], cm, [1]])
        ax.step(xs, ys, where="post", label=lbl, color=c, linewidth=2)
    ax.axvline(prior[0], linestyle="--", color="gray", alpha=0.6, linewidth=2,
               label="prior")
    ax.set_xlabel(r"posterior $p(s_1 \mid x)$", fontsize=12)
    ax.set_ylabel("cumulative probability", fontsize=12)
    ax.text(0.03, 0.94, "CDFs", transform=ax.transAxes, va="top")
```

(continues on next page)

(continued from previous page)

```
    ax.legend(fontsize=11)

    ax2 = axes[1]
    grid = np.linspace(0, 1, 200)

    def integrated_cdf(sorted_vals, cum_mass, grid):
        cdf = np.array([cum_mass[sorted_vals <= t].max()
                        if np.any(sorted_vals <= t) else 0.0
                        for t in grid])
        return np.cumsum(cdf) * (grid[1] - grid[0])

    int_μ = integrated_cdf(sv_μ, cm_μ, grid)
    int_ν = integrated_cdf(sv_ν, cm_ν, grid)

    ax2.plot(grid, int_μ, label=r"$\int F_\mu$", color="steelblue", linewidth=2)
    ax2.plot(grid, int_ν, color="darkorange",
             label=r"$\int F_\nu$", linewidth=2)
    ax2.fill_between(grid, int_ν, int_μ,
                     where=int_μ >= int_ν,
                     alpha=0.2, color="steelblue",
                     label=(r"$\int F_\mu \geq \int F_\nu$"
                            r" ($\mu$ is an MPS of $\nu$)"))
    ax2.set_xlabel(r"$t$", fontsize=12)
    ax2.set_ylabel("integrated CDF", fontsize=12)
    ax2.text(0.03, 0.94, "integrated CDFs", transform=ax2.transAxes, va="top")
    ax2.legend(fontsize=10)

    plt.tight_layout()
    plt.show()

plot_sosd_posteriors(μ_info, ν_info, prior)
```

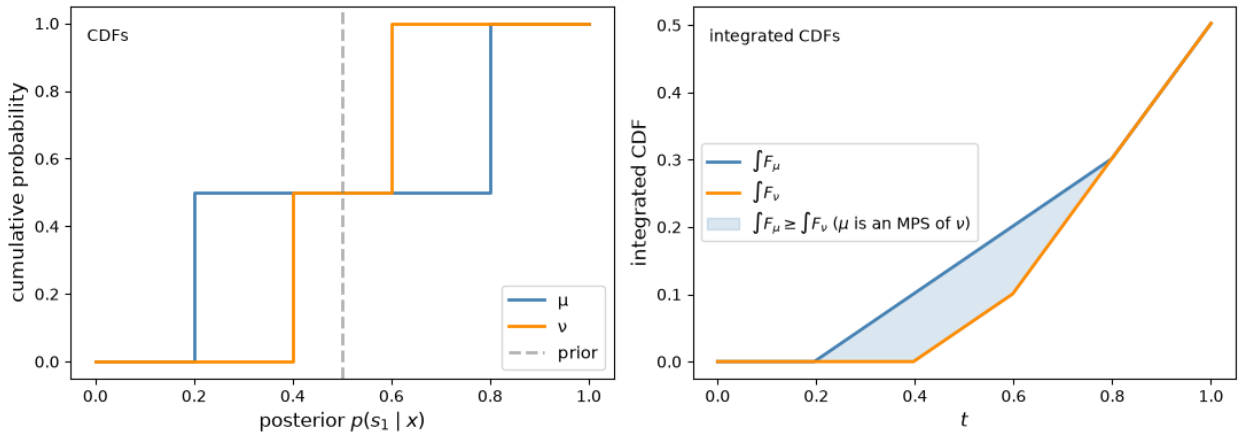


Fig. 30.3: Integrated-CDF check in the two-state case

30.6 Simulating the Blackwell order with many states

We now move to a three-state example.

Experiment μ is strongly correlated with the state, and experiment ν is a garbling of μ.

```
N3 = 3
prior3 = np.array([1/3, 1/3, 1/3])

μ3 = np.array([[0.7, 0.2, 0.1],
               [0.1, 0.7, 0.2],
               [0.2, 0.1, 0.7]])

Q3 = np.array([[0.9, 0.05, 0.05],
               [0.05, 0.8, 0.15],
               [0.05, 0.15, 0.8]])

ν3 = μ3 @ Q3

print("μ (3×3):")
print(np.round(μ3, 2))
print("\nQ (garbling):")
print(np.round(Q3, 2))
print("\nν = μ @ Q:")
print(np.round(ν3, 3))
```

```
μ (3×3):
[[0.7 0.2 0.1]
 [0.1 0.7 0.2]
 [0.2 0.1 0.7]]

Q (garbling):
[[0.9  0.05 0.05]
 [0.05 0.8  0.15]
 [0.05 0.15 0.8 ]]

ν = μ @ Q:
[[0.645 0.21  0.145]
 [0.135 0.595 0.27 ]
 [0.22  0.195 0.585]]
```

For three states, posterior beliefs live in a 2-simplex.

Let's plot the posterior atoms under μ and ν, with marker area proportional to probability

```
def simplex_to_cart(pts):
    """Convert barycentric coordinates on the 2-simplex to 2-D Cartesian coordinates."
 ↪""
    corners = np.array([[0.0, 0.0],
                        [1.0, 0.0],
                        [0.5, np.sqrt(3)/2]])
    return pts @ corners


def plot_simplex_posteriors(μ_matrix, ν_matrix, prior3):
    """Plot posterior atoms on the 2-simplex, with area proportional to probability.""
 ↪"
```

(continues on next page)

(continued from previous page)

```
    corners = simplex_to_cart(np.eye(3))
    prior_cart = simplex_to_cart(prior3[None, :])[0]

    fig, axes = plt.subplots(1, 2, figsize=(12, 5))
    panel_labels = ["μ (more informative)", "ν (garbled)"]
    colors = ["steelblue", "darkorange"]
    labels = ["$s_1$", "$s_2$", "$s_3$"]
    offsets = [(-0.07, -0.05), (0.02, -0.05), (-0.02, 0.03)]

    for ax, m, c, panel_label in zip(axes, [μ_matrix, ν_matrix],
                                     colors, panel_labels):
        posts, probs = compute_posteriors(m, prior3)
        cart = simplex_to_cart(posts)
        tri = plt.Polygon(corners, fill=False, edgecolor="black", linewidth=2)
        ax.add_patch(tri)
        ax.scatter(cart[:, 0], cart[:, 1], s=3000 * probs, alpha=0.6, color=c)
        ax.scatter(*prior_cart, s=120, color="red", zorder=5,
                   label="prior", marker="*")
        for i, (lbl, off) in enumerate(zip(labels, offsets)):
            ax.text(corners[i][0] + off[0], corners[i][1] + off[1],
                    lbl, fontsize=13)
        ax.set_xlim(-0.15, 1.15)
        ax.set_ylim(-0.1, np.sqrt(3)/2 + 0.1)
        ax.set_aspect("equal")
        ax.set_xticks([])
        ax.set_yticks([])
        ax.text(0.03, 0.94, panel_label, transform=ax.transAxes, va="top")
        ax.legend(fontsize=11, loc="upper right")

    plt.tight_layout()
    plt.show()

plot_simplex_posteriors(μ3, ν3, prior3)
```

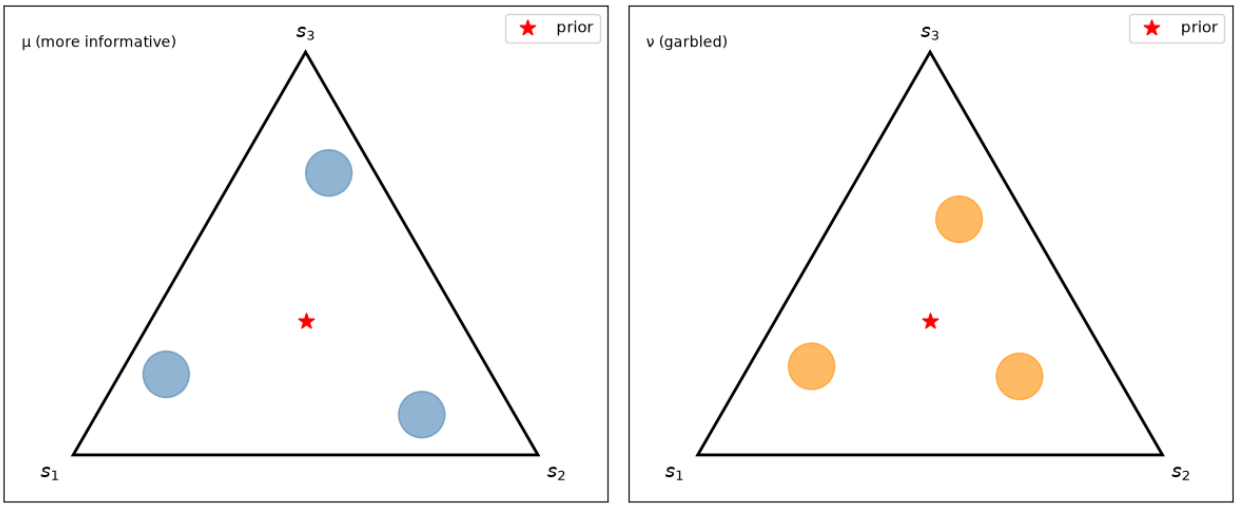


Fig. 30.4: Posterior atoms on the 2-simplex

Under μ, the posterior atoms lie farther toward the vertices.

Under the garbled experiment ν, they lie closer to the prior.

The call-option test does not extend to the simplex.

The garbling Q_3 itself certifies that $\mu_3 \succeq \nu_3$, and the next section checks several concave uncertainty functions in this example.

30.7 The DeGroot uncertainty function

30.7.1 Concave uncertainty functions and the value of information

[DeGroot, 1962] formalizes the value of information through an **uncertainty function** $U : P \to \mathbb{R}$.

DeGroot's definition requires only concavity, so that observing any signal weakly reduces expected uncertainty.

Common examples, such as Shannon entropy, are also symmetric in the components of p and maximized at the uniform prior.

Blackwell's equivalence uses every concave U, symmetric or not; Exercise 30.13.2 shows why asymmetric functions matter.

The *value of experiment μ given prior p* is

$$I(\tilde{x}_\mu;\, \tilde{s};\, U) = U(p) - \mathbb{E}[U(p^\mu)],$$

This quantity is the expected reduction in uncertainty.

Blackwell's order is equivalent to the statement that $I(\tilde{x}_\mu; \tilde{s}; U) \geq I(\tilde{x}_\nu; \tilde{s}; U)$ for *every* concave U.

30.7.2 Shannon entropy as a special case

The canonical uncertainty function is Shannon entropy

$$U_H(p) = -\sum_{i=1}^{N} p_i \log p_i.$$

For any prior p, DeGroot's value formula becomes

$$I(\tilde{x}_\mu, p;\, U_H) = H(\tilde{s}) - H(\tilde{s} \mid \tilde{x}_\mu),$$

where $H(\tilde{s}) = U_H(p)$ is the entropy of the prior and $H(\tilde{s} \mid \tilde{x}_\mu)$ is the conditional entropy of the state given the signal.

Under the uniform prior, $H(\tilde{s}) = \log N$.

To see why, write $H(\tilde{s} \mid \tilde{x}_\mu) = \sum_x \Pr(\tilde{x}_\mu = x)\, H(\tilde{s} \mid \tilde{x}_\mu = x)$, where each conditional entropy term equals $-\sum_i p_i^\mu(x) \log p_i^\mu(x) = U_H(p^\mu(x))$.

Substituting into DeGroot's formula gives $I = U_H(p) - \mathbb{E}[U_H(p^\mu)] = H(\tilde{s}) - H(\tilde{s} \mid \tilde{x}_\mu)$, which is exactly the *mutual information* between $\tilde{x}_\mu$ and $\tilde{s}$.

> **Note**
>
> Blackwell's order implies the entropy inequality.
>
> Because entropy is strictly concave, a strict Blackwell improvement raises mutual information strictly.
>
> The converse fails: a ranking by mutual information does not imply a ranking by Blackwell's order.

Two experiments can be incomparable in Blackwell's order while one has higher mutual information, or while both have the same mutual information, as in Exercise 30.13.2.

The full Blackwell ordering requires the inequality to hold for *every* concave U, not just Shannon entropy.

```
def entropy(p, ε=1e-12):
    """Shannon entropy of a probability vector."""
    p = np.asarray(p, dtype=float)
    p = np.clip(p, ε, 1.0)
    return -np.sum(p * np.log(p))


def degroot_value(μ_matrix, prior, U_func):
    """
    Compute DeGroot's value of information I = U(prior) - E[U(posterior)].
    """
    posts, probs = compute_posteriors(μ_matrix, prior)
    prior_uncertainty = U_func(prior)
    expected_post_uncertainty = sum(
        probs[j] * U_func(posts[j]) for j in range(len(probs)))
    return prior_uncertainty - expected_post_uncertainty


def gini_impurity(p):
    """Gini impurity: 1 - sum(p_i^2)."""
    return 1.0 - np.sum(np.asarray(p)**2)


def tsallis_entropy(p, q=2):
    """Tsallis entropy of order q (concave for q>1)."""
    p = np.clip(p, 1e-12, 1.0)
    return (1 - np.sum(p**q)) / (q - 1)


def tsallis_q15(p):
    """Tsallis entropy with q=1.5 for an independent concavity check."""
    return tsallis_entropy(p, q=1.5)


def sqrt_index(p):
    """Concave uncertainty index based on sum(sqrt(p_i))."""
    p = np.clip(np.asarray(p), 0.0, 1.0)
    return np.sum(np.sqrt(p)) - 1.0

uncertainty_functions = {
    "Shannon entropy": entropy,
    "Gini impurity": gini_impurity,
    "Tsallis (q=1.5)": tsallis_q15,
    "Square-root index": sqrt_index,
}

header = (f"{'Uncertainty function':<22}  "
          f"{'I(μ)':<10}  {'I(ν)':<10}  "
          f"{'I(μ)>=I(ν)?'}")
print(header)
print("-" * 58)
for name, U in uncertainty_functions.items():
```

(continues on next page)

(continued from previous page)

```
    I_μ = degroot_value(μ_info, prior, U)
    I_ν = degroot_value(ν_info, prior, U)
    print(f"{name:<22}  {I_μ:<10.4f}  {I_ν:<10.4f}  {I_μ >= I_ν - 1e-10}")

print("\nThree-state example (μ3, ν3, uniform prior):")
print(header)
print("-" * 58)
for name, U in uncertainty_functions.items():
    I_μ = degroot_value(μ3, prior3, U)
    I_ν = degroot_value(ν3, prior3, U)
    print(f"{name:<22}  {I_μ:<10.4f}  {I_ν:<10.4f}  {I_μ >= I_ν - 1e-10}")
```

```
Uncertainty function    I(μ)        I(ν)        I(μ)>=I(ν)?
----------------------------------------------------------
Shannon entropy         0.1927      0.0201      True
Gini impurity           0.1800      0.0200      True
Tsallis (q=1.5)         0.1958      0.0213      True
Square-root index       0.0726      0.0072      True

Three-state example (μ3, ν3, uniform prior):
Uncertainty function    I(μ)        I(ν)        I(μ)>=I(ν)?
----------------------------------------------------------
Shannon entropy         0.2968      0.1690      True
Gini impurity           0.2067      0.1183      True
Tsallis (q=1.5)         0.2588      0.1487      True
Square-root index       0.1319      0.0734      True
```

As predicted by the theorem, $I(\mu) \geq I(\nu)$ for every concave uncertainty function in both examples, because $\mu \geq \nu$ in the Blackwell sense.

30.7.3 Value of information as a function of experiment quality

We now parameterize a continuum of experiments between the uninformative and perfectly informative cases.

For $N = 2$ states, a natural family is

$$\mu(\theta) = (1-\theta) \cdot \tfrac{1}{2}\mathbf{1}\mathbf{1}^\top + \theta \cdot I_2, \quad \theta \in [0,1],$$

The first term is the completely mixed matrix and I_2 is the identity.

The family is ordered by Blackwell's criterion: for $\theta' < \theta$, $\mu(\theta') = \mu(\theta)\,\mu(\theta'/\theta)$, so a lower-quality experiment is a garbling of a higher-quality one.

```
def make_experiment(θ, N=2):
    """Parameterized experiment: θ=0 is uninformative, θ=1 is perfect."""
    return (1 - θ) * np.ones((N, N)) / N + θ * np.eye(N)


θs = np.linspace(0, 1, 100)
prior2 = np.array([0.5, 0.5])

fig, ax = plt.subplots(figsize=(9, 4))
for name, U in uncertainty_functions.items():
    values = [degroot_value(make_experiment(θ), prior2, U) for θ in θs]
    vmin, vmax = values[0], values[-1]
```

(continues on next page)

(continued from previous page)

```
    normed = (np.array(values) - vmin) / (vmax - vmin + 1e-15)
    ax.plot(θs, normed, label=name, linewidth=2)

ax.set_xlabel("experiment quality θ  (0 = uninformative, 1 = perfect)",
              fontsize=11)
ax.set_ylabel("normalized value of information I(μ(θ))", fontsize=11)
ax.legend(fontsize=10)
plt.tight_layout()
plt.show()
```

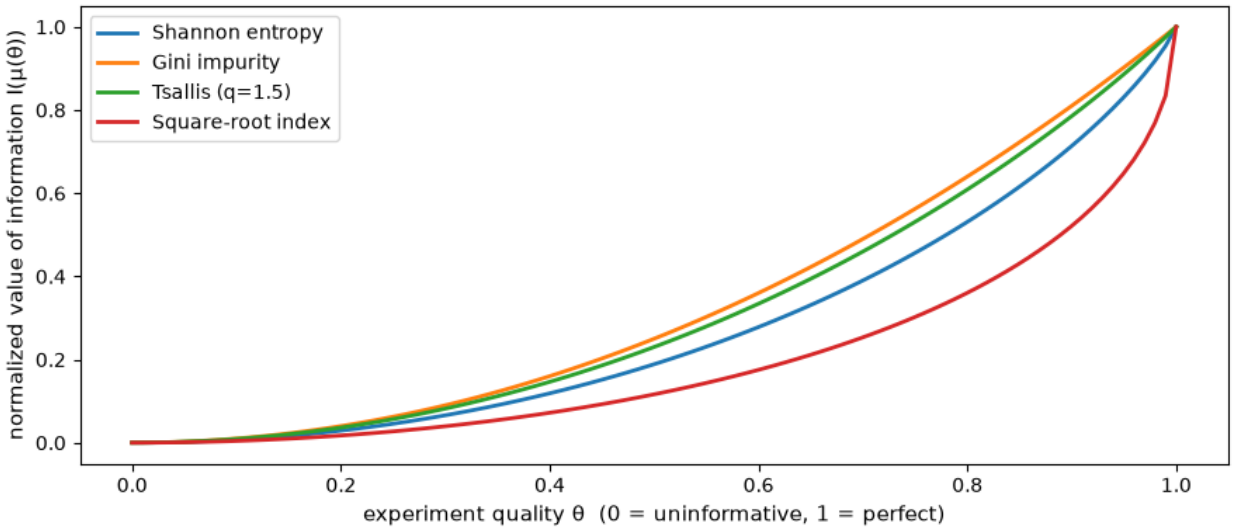


Fig. 30.5: Value of information and experiment quality

Every concave uncertainty function assigns weakly higher value to a more informative experiment.

30.8 Application 1: demand for information about product quality

Kihlstrom [1974] applies Blackwell's theorem to a consumer's demand for information about product quality.

- The unknown state $\tilde{s}$ is the quality of a product.
- A consumer can purchase λ units of information at cost $c\lambda^2$.
- As λ rises, the experiment becomes more informative in Blackwell's sense.

Absent costs, the Blackwell order says that more information is better for every expected-utility maximizer.

With costs, the consumer chooses λ to maximize the net value of information.

We let the accuracy of the experiment rise with diminishing returns, $\phi(\lambda) = 1 - e^{-a\lambda}$, and feed it into the family $\mu(\phi)$ above.

We measure the gross value of information by DeGroot's entropy value, a stand-in for the consumer's expected-utility gain in Kihlstrom [1974].

The Design and Price of Information studies the sale of information to buyers whose values for it differ.

```
def gross_value(λ, prior2, U=entropy, rate=2):
    """Gross value of buying λ units of information."""
    accuracy = 1 - np.exp(-rate * λ)
    return degroot_value(make_experiment(accuracy), prior2, U)


λ_grid = np.linspace(0, 1, 200)
c = 0.6

gross_vals = np.array([gross_value(λ, prior2) for λ in λ_grid])
cost_vals = c * λ_grid**2
net_vals = gross_vals - cost_vals
marginal_vals = np.gradient(gross_vals, λ_grid)
marginal_cost = 2 * c * λ_grid
opt_idx = int(np.argmax(net_vals))

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

ax = axes[0]
ax.plot(λ_grid, gross_vals, label="Gross value I(λ)",
        color="steelblue", linewidth=2)
ax.plot(λ_grid, cost_vals, label=r"Cost $c\lambda^2$",
        color="tomato", linestyle="--", linewidth=2)
ax.plot(λ_grid, net_vals, label="Net value", color="green", linewidth=2)
ax.axvline(λ_grid[opt_idx], color="green", linestyle=":", linewidth=2,
           label=f"λ* ≈ {λ_grid[opt_idx]:.2f}")
ax.set_xlabel("information purchased λ", fontsize=11)
ax.set_ylabel("value (entropy units)", fontsize=11)
ax.legend(fontsize=10)

ax2 = axes[1]
ax2.plot(λ_grid, marginal_vals, label="Marginal value I'(λ)",
         color="steelblue", linewidth=2)
ax2.plot(λ_grid, marginal_cost, label=r"Marginal cost $2c\lambda$",
         color="tomato", linestyle="--", linewidth=2)
ax2.axvline(λ_grid[opt_idx], color="green", linestyle=":", linewidth=2,
            label=f"λ* ≈ {λ_grid[opt_idx]:.2f}")
ax2.set_xlabel("information purchased λ", fontsize=11)
ax2.set_ylabel("marginal value / cost", fontsize=11)
ax2.legend(fontsize=10)

plt.tight_layout()
plt.show()
```

The optimal purchase λ^* occurs where marginal value equals marginal cost.

The marginal value first rises and then falls.

Near $\lambda = 0$ the value of information is convex: it grows only quadratically in the accuracy of the experiment.

Radner and Stiglitz [1984] show that this nonconcavity of the value of information is general, so that demand for information can jump from zero to a positive amount as its price falls.

Because accuracy has diminishing returns, the marginal value eventually falls below the rising marginal cost, which yields an interior optimum here.

Raising c shifts the marginal cost curve up and reduces λ^*, while a more asymmetric prior shifts the marginal value curve and changes the optimum.

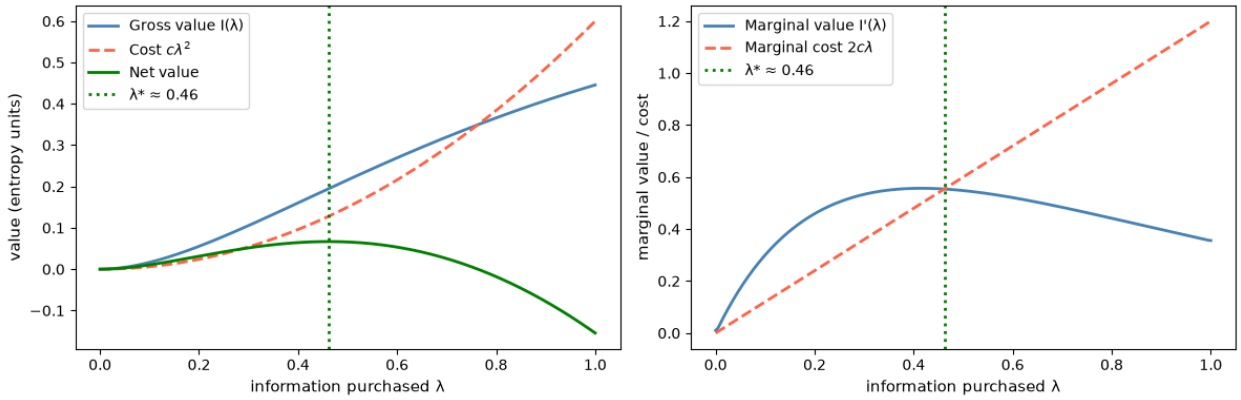


Fig. 30.6: Demand for information with a quadratic cost

30.9 Application 2: sequential experimental design

DeGroot [1962] applies the uncertainty-function framework to *sequential experimental design*.

Each period a statistician observes one draw and updates the posterior.

The question is which sequence of experiments minimizes cumulative expected uncertainty.

If one experiment is more informative than another at every stage, then the Blackwell order favors using the better experiment at every date.

A Problem that Stumped Milton Friedman and *A Bayesian Formulation of Friedman and Wald's Problem* study a complementary question: when to stop sampling from a fixed experiment.

We now simulate sequential belief updating for experiments of different quality.

```
def sequential_update(μ_matrix, prior, T=20, seed=0):
    """Simulate T sequential belief updates under experiment μ."""
    rng = np.random.default_rng(seed)
    N, M = μ_matrix.shape
    beliefs = np.zeros((T + 1, N))
    beliefs[0] = prior.copy()

    true_state = rng.choice(N, p=prior)

    for t in range(T):
        p = beliefs[t]
        signal = rng.choice(M, p=μ_matrix[true_state])
        unnorm = μ_matrix[:, signal] * p
        beliefs[t + 1] = unnorm / unnorm.sum()

    return beliefs, true_state


def plot_sequential_beliefs(θs_compare, prior2, T=25):
    fig, axes = plt.subplots(1, len(θs_compare), figsize=(14, 4), sharey=True)

    for ax, θ in zip(axes, θs_compare):
        μ_t = make_experiment(θ, N=2)
        for seed in range(15):
            beliefs, ts = sequential_update(μ_t, prior2, T=T, seed=seed)
```

(continues on next page)

(continued from previous page)

```
            c = "steelblue" if ts == 0 else "darkorange"
            ax.plot(beliefs[:, 0], alpha=0.35, color=c, linewidth=2)
        ax.axhline(prior2[0], linestyle="--", color="gray", linewidth=2,
                   label="prior")
        ax.axhline(1.0, linestyle=":", color="steelblue", linewidth=2)
        ax.axhline(0.0, linestyle=":", color="darkorange", linewidth=2)
        ax.set_xlabel(r"period $t$", fontsize=11)
        if θ == θs_compare[0]:
            ax.set_ylabel(r"posterior $p(s_1 \mid x^t)$", fontsize=11)
        ax.set_ylim(-0.05, 1.05)
        ax.text(0.03, 0.94, f"θ = {θ}", transform=ax.transAxes, va="top")
        ax.legend(fontsize=9)

    plt.tight_layout()
    plt.show()

plot_sequential_beliefs([0.2, 0.5, 0.9], prior2, T=30)
```

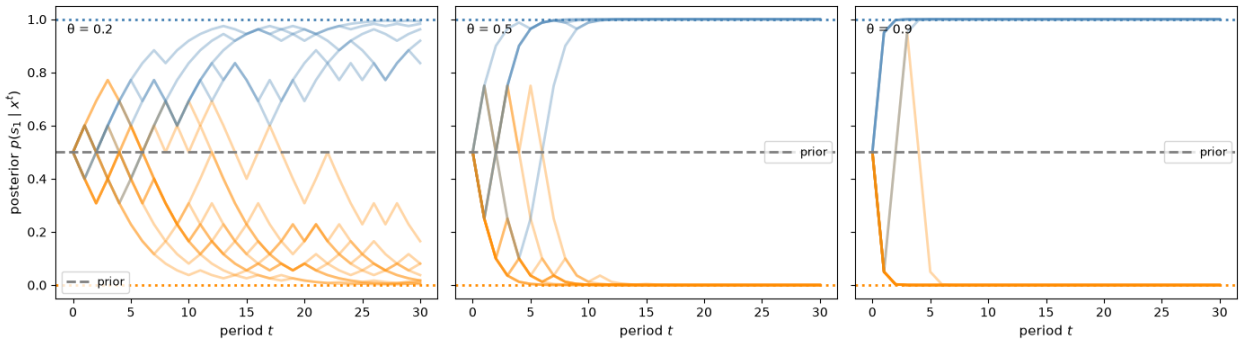


Fig. 30.7: Sequential posterior paths for different experiment qualities

More informative experiments make beliefs converge faster to the truth.

Likelihood Ratio Processes shows that the log likelihood ratio drifts at a rate equal to a Kullback–Leibler divergence between the two rows of the experiment, and that divergence rises with θ.

Under the correct prior, the posterior process is a martingale.

```
def check_martingale_mean(μ_matrix, prior, T=15, n_paths=2000, seed=0):
    """
    Simulate many belief paths and check E[p_t] = p_0.
    """
    rng = np.random.default_rng(seed)
    N, M = μ_matrix.shape
    all_paths = np.zeros((n_paths, T + 1, N))

    for k in range(n_paths):
        true_state = rng.choice(N, p=prior)
        p = prior.copy()
        all_paths[k, 0] = p
        for t in range(T):
            signal = rng.choice(M, p=μ_matrix[true_state])
            unnorm = μ_matrix[:, signal] * p
```

(continues on next page)

(continued from previous page)

```
            p = unnorm / unnorm.sum()
            all_paths[k, t + 1] = p

    mean_path = all_paths[:, :, 0].mean(axis=0)

    fig, ax = plt.subplots(figsize=(8, 4))
    ax.plot(mean_path, color="steelblue", linewidth=2,
            label=r"$\bar p_t(s_1)$ (mean over paths)")
    ax.axhline(prior[0], linestyle="--", color="tomato", linewidth=2,
               label=fr"Prior $p_0 = {prior[0]:.2f}$")
    ax.set_xlabel(r"period $t$", fontsize=12)
    ax.set_ylabel(r"$E[p_t(s_1)]$", fontsize=12)
    ax.legend(fontsize=11)
    ax.set_ylim(0, 1)
    plt.tight_layout()
    plt.show()

    print(f"Prior = {prior[0]:.4f}")
    print(f"Average mean belief across dates: {mean_path.mean():.4f}")

check_martingale_mean(μ_info, prior, T=20, n_paths=5000)
```

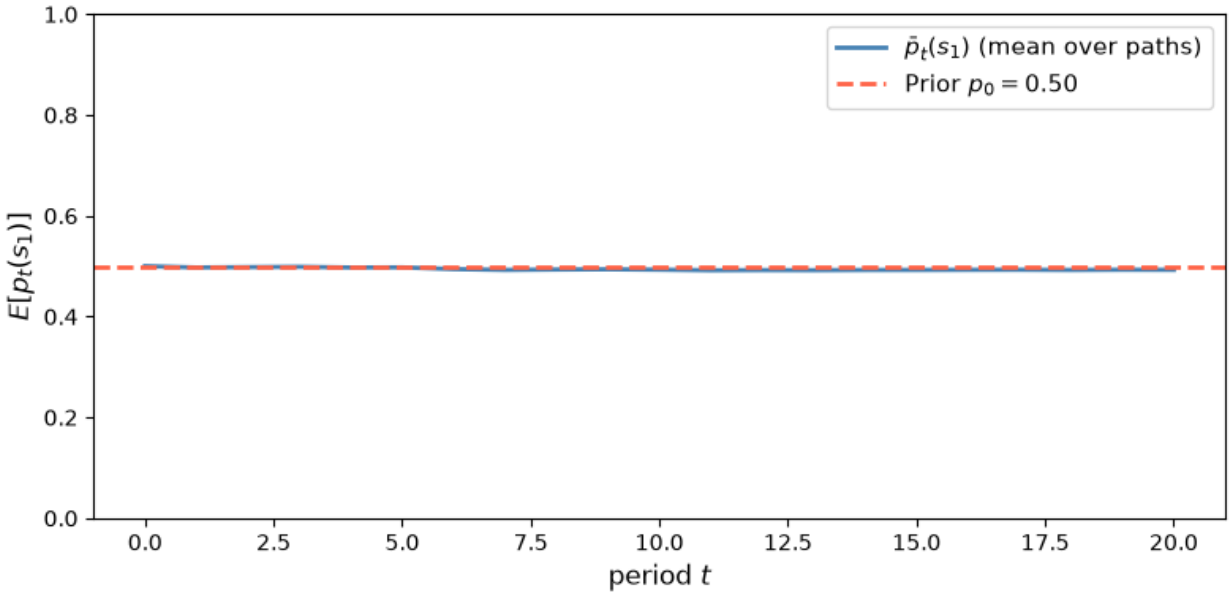


Fig. 30.8: Unconditional implication of the posterior martingale property

```
Prior = 0.5000
Average mean belief across dates: 0.4942
```

The simulated cross-sectional mean stays close to the prior at every date.

This is the unconditional implication of the posterior martingale property.

Merging of Opinions: The Blackwell–Dubins Theorem studies the almost-sure convergence of such belief martingales and the Blackwell–Dubins theorem on the merging of opinions.

30.10 The Data Processing Inequality and Coarse-Graining

Blackwell's condition that $\nu = \mu Q$ for some Markov kernel Q is the same mathematical operation that underlies the **data processing inequality** (DPI) and the **coarse-graining theorem** in information theory, information geometry, and machine learning.

30.10.1 The DPI for f-divergences

An **f-divergence** between two probability distributions P and Q over a finite space Ω is

$$D_f(P\|Q) = \sum_{\omega\in\Omega} q_\omega \, f\left(\frac{p_\omega}{q_\omega}\right),$$

where $f : (0,\infty) \to \mathbb{R}$ is a convex function with $f(1) = 0$.

Special cases include:

Divergence	Generator $f(t)$
KL-divergence	$t \log t$
Squared Hellinger H^2	$(\sqrt{t}-1)^2/2$
Total variation TV	$\lvert t-1\rvert/2$
Chi-squared χ^2	$(t-1)^2$

The class of f-divergences was introduced independently by Ali and Silvey [1966], Csiszár [1963], and Morimoto [1963]; see also Liese [2012].

Statistical Divergence Measures studies several of these divergences for continuous distributions.

Theorem 30.10.1 (Data Processing Inequality)

For any f-divergence D_f and any Markov kernel (stochastic transformation) κ, with $P\kappa$ denoting the image of P under κ, we have

$$D_f(P\|Q) \geq D_f(P\kappa\|Q\kappa).$$

If κ is induced by a sufficient statistic for the pair $\{P, Q\}$, then equality holds.

A converse of this form requires additional hypotheses; a clean binary-model characterization is given below.

The proof follows from Jensen's inequality applied to the convex function f, using the fact that κ is a stochastic matrix [Csiszár, 1963].

30.10.2 Connection to Blackwell's sufficiency condition

In Blackwell's framework, μ and ν are experiments over the same state space $S = \{s_1, \dots, s_N\}$.

For two states, each experiment has two rows: $\mu_1 = \mu(s_1, \cdot)$ and $\mu_2 = \mu(s_2, \cdot)$.

If $\nu = \mu Q$ (i.e., ν is a garbling of μ), then the pair $(\nu_1, \nu_2) = (\mu_1 Q, \mu_2 Q)$ is obtained by applying the Markov kernel Q to the pair (μ_1, μ_2).

The coarse-graining theorem then implies immediately:

$$D_f(\mu_1 \| \mu_2) \geq D_f(\nu_1 \| \nu_2) \quad \text{for every f-divergence } D_f,$$

whenever $\mu \geq \nu$ in the Blackwell order.

For two states the converse also holds, with the usual conventions for zero probabilities: if $D_f(\mu_1 \| \mu_2) \geq D_f(\nu_1 \| \nu_2)$ for every convex f, then $\mu \geq \nu$.

With more than two states, divergences between pairs of rows do not characterize Blackwell's order.

So a more informative experiment always produces *more separated* conditional signal distributions, in the sense of every f-divergence simultaneously.

The DPI is thus a statement about the *distinguishability* of states: garbling an experiment makes the states harder to tell apart under every statistical measure of separability.

For binary experiments, the equality condition links the DPI directly back to Blackwell: $D_f(\mu_1 Q \| \mu_2 Q) = D_f(\mu_1 \| \mu_2)$ for some strictly convex f if and only if Q is a sufficient statistic for (μ_1, μ_2).

Once sufficiency holds, equality follows for every convex f [Liese, 2012].

30.10.3 Information geometry: Chentsov's theorem

The DPI has an infinitesimal, differential-geometric companion.

Chentsov's theorem [Čencov, 1981] states that the **Fisher information metric** is, up to a constant factor, the *unique* Riemannian metric on the probability simplex that is invariant under Markov embeddings, the transformations induced by sufficient statistics.

The Fisher information matrix $I_F(\theta)$ also contracts under every Markov kernel:

$$I_F(\theta;\, \mu) \succeq I_F(\theta;\, \mu\kappa) \quad \text{for every differentiable family } \{\mu_\theta\} \text{ and every Markov kernel } \kappa.$$

Equality holds if and only if κ is a sufficient statistic for θ.

The uniqueness clause says that invariance under sufficient statistics singles out the Fisher metric among all Riemannian metrics.

In *Which Moments to Match?*, the Fisher information of the structural model sets the efficiency bound that score matching attains when the auxiliary model nests the structural model.

See Amari and Nagaoka [2000] for a thorough treatment of information geometry and its connections to sufficiency.

30.10.4 The information bottleneck in machine learning

The **information bottleneck** method of Tishby *et al.* [1999] provides a prominent application of the DPI in machine learning.

Given a joint distribution $p(X, Y)$ over an input X and a target Y, the goal is to find a compressed representation T, formed by a stochastic mapping $p(T \mid X)$, that retains as much information about Y as possible while using as few bits as possible to describe X.

The method minimizes the Lagrangian

$$\mathcal{L}[p(T \mid X)] = I(X;\, T) - \beta\, I(T;\, Y),$$

where $I(\cdot\,;\,\cdot)$ denotes mutual information and $\beta \geq 0$ governs the compression–relevance trade-off.

Because $Y - X - T$ forms a Markov chain (T is derived from X alone), the DPI implies

$$I(T;\, Y) \leq I(X;\, Y),$$

with equality if and only if T is a **sufficient statistic** of X for Y, that is, if $Y - T - X$ is also a Markov chain.

The Blackwell ordering explains why no deterministic or random post-processing of X can increase the mutual information with Y: any Markov kernel applied to X is a garbling in Blackwell's sense, and the DPI is the mutual-information form of the coarse-graining theorem.

In machine learning language the information bottleneck searches among all garblings of X for the one that best preserves relevant information about Y subject to a compression budget.

In a deep neural network with input X and target Y and layers $X \to T_1 \to T_2 \to \cdots \to T_L \to \hat{Y}$, each layer's representation is a garbling of the previous one.

The DPI then implies the chain of inequalities

$$I(X;\, Y) \geq I(T_1;\, Y) \geq I(T_2;\, Y) \geq \cdots \geq I(T_L;\, Y),$$

so successive layers can only lose, never gain, information about Y.

This observation was placed at the center of the study of what deep networks learn by [Shwartz-Ziv and Tishby, 2017].

Fig. 30.5 already illustrates this: as experiment quality θ increases, every measure of informativeness rises monotonically.

The DPI says the same thing in reverse: garbling (decreasing θ) can only contract these measures.

30.10.5 Summary of the DPI–Blackwell correspondence

The table below collects the precise correspondence between Blackwell's framework and the data-processing and coarse-graining literature.

Blackwell / DeGroot	Data processing / coarse-graining
Garbling $\nu = \mu Q$	Applying Markov kernel κ to a pair $(P, Q) = (\mu_1, \mu_2)$
$\mu \geq \nu$ in Blackwell order	$D_f(\mu_1 \mid \mu_2) \geq D_f(\nu_1 \mid \nu_2)$ for every f-divergence
Sufficiency (Q discards nothing)	Equality in DPI; in binary models, one strictly convex f already characterizes sufficiency
DeGroot value $I(\mu; U_H)$	Mutual information $I(\tilde{x}_\mu;\ \tilde{s})$ (Shannon DPI)
Posterior spreads under μ vs ν	D_f between rows larger under μ
Blackwell theorem for two states	D_f dominance for every $f \Leftrightarrow$ a single Markov kernel witnesses dominance
Chentsov's uniqueness theorem	Fisher information is the unique metric invariant under Markov embeddings (sufficient statistics); it also contracts under Markov kernels
Information bottleneck $I(T; Y) \leq I(X; Y)$	DPI for mutual information applied to Markov chain $Y - X - T$

30.11 Relation to Bayesian likelihood-ratio learning

The lecture *Likelihood Ratio Processes and Bayesian Learning* is a dynamic two-state special case of the framework developed here.

Let $S = \{s_1, s_2\}$ with $s_1 \leftrightarrow f$ and $s_2 \leftrightarrow g$, where f and g are the two candidate data-generating densities.

Then a single observation is a Blackwell experiment with rows $f(\cdot)$ and $g(\cdot)$, and the history $w^t = (w_1, \ldots, w_t)$ defines a richer experiment μ_t.

Here we date the first draw $t = 1$ and write the prior as π_0, as in *Likelihood Ratio Processes*; *Likelihood Ratio Processes and Bayesian Learning* dates the first draw $t = 0$ and writes the prior as π_{-1}.

Because one can always discard the last $t - s$ observations, μ_t Blackwell-dominates μ_s for every $t > s$.

The likelihood-ratio process

$$L(w^t) = \prod_{i=1}^{t} \frac{f(w_i)}{g(w_i)}$$

is a sufficient statistic for μ_t, and the posterior

$$\pi_t = \Pr(s_1 \mid w^t) = \frac{\pi_0 L(w^t)}{\pi_0 L(w^t) + 1 - \pi_0}$$

is Kihlstrom's standard experiment in this two-state setting.

Its martingale property, $E[\pi_t] = \pi_0$, is exactly the mean-preservation result proved above for posterior distributions.

Likewise, $\mu_t \geq \mu_s$ implies that the distribution of π_t is a mean-preserving spread of the distribution of π_s, so additional data pushes beliefs farther toward 0 and 1 while lowering expected uncertainty under every concave uncertainty function.

30.11.1 Summary table

The table below records the dictionary between the two lectures without repeating the earlier arguments.

Concept in *Likelihood Ratio Processes and Bayesian Learning*	Concept in this lecture
States $\{f, g\}$	State space $S = \{s_1, s_2\}$
Densities $f(\cdot)$, $g(\cdot)$	Rows of experiment matrix μ
Single draw w_t	Blackwell experiment with continuous signal space
History w^t of t IID draws	Richer experiment μ_t Blackwell-dominating μ_s, $s < t$
Likelihood ratio $L(w^t)$	Sufficient statistic for μ_t
Prior π_0	Prior $p \in P$ on the 1-simplex $[0, 1]$
Posterior π_t	Posterior on $P = [0, 1]$ (Kihlstrom's standard experiment)
Distribution of π_t across histories	$\hat{\mu}^c$ (Kihlstrom's posterior distribution)
Martingale property $E[\pi_t] = \pi_0$	Mean preservation of $\hat{\mu}^c$
$\pi_t \to 0$ or 1 almost surely	Posteriors spread to vertices (MPS in the limit)
Expected fall in the entropy of π_t	DeGroot value $I(\mu_t; U_H)$, the mutual information
More draws $\Rightarrow$ better for all decision makers	Blackwell ordering $\mu_t \geq \mu_s$
Garbling (discard last $t - s$ draws)	Stochastic transformation Q with $\mu_s = \mu_t Q$

30.12 Summary

Blackwell's theorem identifies a *partial order* on statistical experiments with three equivalent characterizations:

Criterion	Condition
Economic	Every decision maker weakly prefers μ to ν: $B(\mu, A, u) \supseteq B(\nu, A, u)$
Sufficiency	ν is a garbling of μ: $\nu = \mu Q$ for some Markov Q
Uncertainty	μ reduces expected uncertainty more for every prior p and every concave U

Kihlstrom's Bayesian exposition places the *posterior distribution* at the center.

A more informative experiment generates a more dispersed posterior distribution with the same mean prior.

The right probabilistic language is convex order on the simplex of posterior beliefs.

In the two-state case this reduces to the familiar SOSD / integrated-CDF test on $[0, 1]$.

DeGroot's contribution is to extend the comparison from particular utility functions to the full class of concave uncertainty functions.

Blackwell's order implies the data processing inequality for every f-divergence, and the likelihood-ratio learning of *Likelihood Ratio Processes and Bayesian Learning* is a dynamic special case of the framework.

30.13 Exercises

Exercise 30.13.1

The first section found that no stochastic transformation maps $\nu = \mu Q$ back into μ, but it offered only a numerical residual as evidence.

This exercise supplies a proof.

Suppose that $\mu = \nu Q'$ for some Markov matrix Q', so that $\tilde{x}_\mu$ is generated from $\tilde{x}_\nu$ by a randomization that does not depend on $\tilde{s}$.

1. Show that each posterior under μ is a weighted average of posteriors under ν:
$$p^\mu(x) = \sum_y \Pr(\tilde{x}_\nu = y \mid \tilde{x}_\mu = x)\, p^\nu(y).$$
2. Under the uniform prior, compute the posterior probabilities of s_1 that μ and ν from the first section generate.
3. Conclude that μ is not a garbling of ν.

Solution

For part 1, the state and the two signals form a Markov chain $\tilde{s} \to \tilde{x}_\nu \to \tilde{x}_\mu$, because Q' does not depend on the state.

Hence $\Pr(\tilde{s} = s_i \mid \tilde{x}_\nu = y, \tilde{x}_\mu = x) = \Pr(\tilde{s} = s_i \mid \tilde{x}_\nu = y)$, and

$$p^\mu_i(x) = \sum_y \Pr(\tilde{x}_\nu = y \mid \tilde{x}_\mu = x)\, \Pr(\tilde{s} = s_i \mid \tilde{x}_\nu = y) = \sum_y \Pr(\tilde{x}_\nu = y \mid \tilde{x}_\mu = x)\, p^\nu_i(y).$$

Every posterior of a garbled experiment therefore lies in the convex hull of the posteriors of the original experiment.

```
c2 = np.array([0.5, 0.5])
post_μ1, _ = compute_posteriors(μ, c2)
post_ν1, _ = compute_posteriors(ν, c2)

lo, hi = post_ν1[:, 0].min(), post_ν1[:, 0].max()
outside = (post_μ1[:, 0] < lo - 1e-12) | (post_μ1[:, 0] > hi + 1e-12)

print("posteriors on s1 under μ:", post_μ1[:, 0].round(4))
print("posteriors on s1 under ν:", post_ν1[:, 0].round(4))
print(f"interval spanned by ν's posteriors: [{lo:.4f}, {hi:.4f}]")
print("μ's posteriors outside that interval:", post_μ1[outside, 0].round(4))
```

```
posteriors on s1 under μ: [0.8571 0.5    0.1429]
posteriors on s1 under ν: [0.75 0.25]
interval spanned by ν's posteriors: [0.2500, 0.7500]
μ's posteriors outside that interval: [0.8571 0.1429]
```

Under ν the posteriors on s_1 are $3/4$ and $1/4$.

Under μ they are $6/7$, $1/2$ and $1/7$.

The posteriors $6/7$ and $1/7$ lie outside $[1/4, 3/4]$, so no Markov matrix Q' produces μ from ν.

The argument explains the large residual in the first section: a garbling can only pull posteriors toward the prior, and μ places posteriors farther from the prior than ν can reach.

Exercise 30.13.2

Blackwell's order is a partial order, and *The Design and Price of Information* builds on that fact.

Consider two experiments with two states and two signals,

$$\mu_a = \begin{bmatrix} 1 & 0 \\ 0.5 & 0.5 \end{bmatrix}, \qquad \mu_b = \begin{bmatrix} 0.5 & 0.5 \\ 0 & 1 \end{bmatrix},$$

and the uniform prior.

1. Compute the distribution of the posterior probability of s_1 under each experiment, and the mutual information between signal and state.
2. Use the call-option test $\mathbb{E}[\max(p-t,0)]$ to show that neither posterior distribution is a mean-preserving spread of the other.
3. For the concave uncertainty functions $U_t(p) = -\max(p_1 - t, 0)$ with $t = 0.2$ and $t = 0.6$, compute DeGroot's value of each experiment and interpret the result.

Solution

Under μ_a, signal x_2 occurs only in state s_2, so it reveals s_2.

Under μ_b, signal x_1 occurs only in state s_1, so it reveals s_1.

```
μ_a = np.array([[1.0, 0.0],
                [0.5, 0.5]])
μ_b = np.array([[0.5, 0.5],
                [0.0, 1.0]])

for name, m in [("μ_a", μ_a), ("μ_b", μ_b)]:
    posts, probs = compute_posteriors(m, c2)
    print(f"{name}: posteriors on s1 = {posts[:, 0].round(4)}, "
          f"probabilities = {probs.round(4)}, "
          f"mutual information = {degroot_value(m, c2, entropy):.4f}")
```

```
μ_a: posteriors on s1 = [0.6667 0.    ], probabilities = [0.75 0.25], mutual␣
↪information = 0.2158
μ_b: posteriors on s1 = [1.     0.3333], probabilities = [0.25 0.75], mutual␣
↪information = 0.2158
```

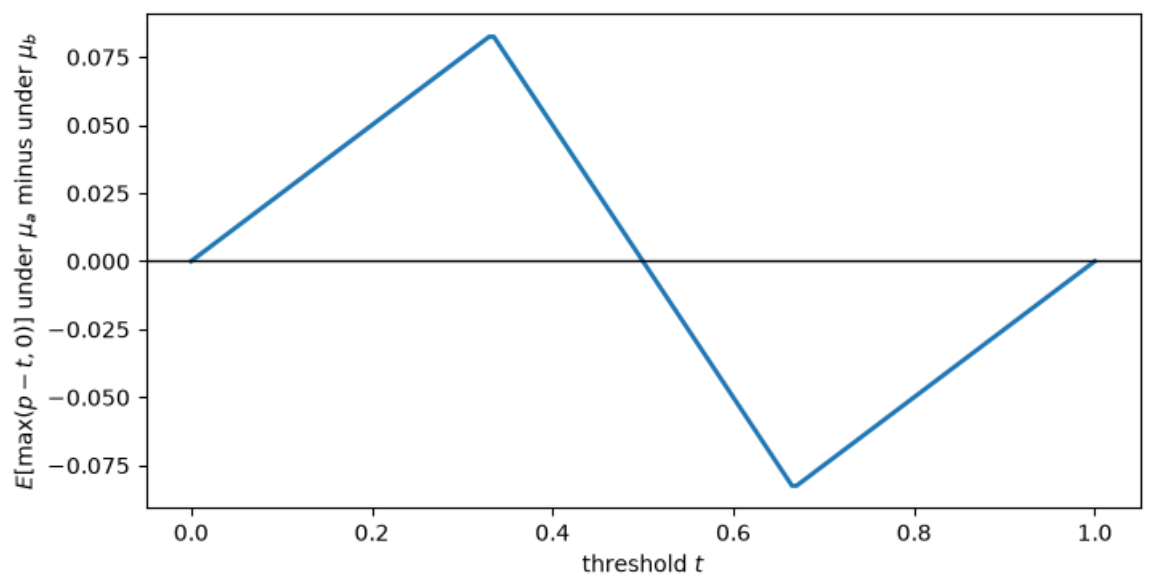
The two posterior distributions are mirror images.

Relabeling states and signals maps one experiment into the other, so the two have the same mutual information.

```
def call_values(m, prior, ts):
    "E[max(p_1 - t, 0)] under experiment m for each threshold t."
    posts, probs = compute_posteriors(m, prior)
    return np.array([(np.maximum(posts[:, 0] - t, 0) * probs).sum() for t in ts])

ts = np.linspace(0, 1, 201)
diff_ab = call_values(μ_a, c2, ts) - call_values(μ_b, c2, ts)

fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(ts, diff_ab, linewidth=2)
ax.axhline(0, color="black", linewidth=1)
ax.set_xlabel("threshold $t$")
ax.set_ylabel(r"$E[\max(p - t, 0)]$ under $\mu_a$ minus under $\mu_b$")
plt.show()
```



The difference changes sign.

Neither posterior distribution is a mean-preserving spread of the other, so by Kihlstrom's theorem neither experiment is Blackwell sufficient for the other.

```
def U_call(t):
    "The concave uncertainty function -max(p_1 - t, 0)."
    return lambda p: -max(p[0] - t, 0.0)

for t in (0.2, 0.6):
    print(f"t = {t}:  I(μ_a) = {degroot_value(μ_a, c2, U_call(t)):.4f},  "
          f"I(μ_b) = {degroot_value(μ_b, c2, U_call(t)):.4f}")
```

```
t = 0.2:  I(μ_a) = 0.0500,  I(μ_b) = 0.0000
t = 0.6:  I(μ_a) = 0.0500,  I(μ_b) = 0.1000
```

With $t = 0.2$, μ_a has value 0.05 and μ_b has value zero.

With $t = 0.6$, the ranking reverses.

Up to sign and an affine term, U_t is the value function of a decision maker who chooses between two actions and switches actions at the belief t.

A decision maker who switches at $t = 0.2$ values only signals that push the posterior below 0.2, and μ_b never does so.

A decision maker who switches at $t = 0.6$ values μ_b's signal that reveals s_1 more than μ_a's signal that raises the posterior to $2/3$.

Mutual information ranks the two experiments equally, yet decision makers disagree about them.

A single uncertainty function cannot settle a Blackwell comparison.

Exercise 30.13.3

This exercise connects Blackwell's order to the learning dynamics of *Likelihood Ratio Processes and Bayesian Learning*.

Let a single draw come from the experiment `make_experiment(θ)` with $\theta = 0.3$, so that the signal matches the state with probability $(1+\theta)/2$.

Observe t independent draws.

1. Explain why the number k of draws equal to x_1 is a sufficient statistic, and write the $2 \times (t+1)$ experiment matrix μ_t for k.
2. Construct the Markov matrix Q_t that discards the last of $t+1$ draws, and verify numerically that $\mu_t = \mu_{t+1} Q_t$.
3. Under the uniform prior, use the call-option test to verify that the posterior distribution after $t+1$ draws is a mean-preserving spread of the posterior distribution after t draws, for $t = 1, \ldots, 29$.
4. Plot the mutual information between k and the state against t, and compare it with its upper bound $\log 2$.

Solution

Write $a = (1+\theta)/2$ and $b = (1-\theta)/2$.

The likelihood of t draws containing k copies of x_1 is $a^k b^{t-k}$ in state s_1 and $b^k a^{t-k}$ in state s_2.

The likelihood ratio depends on the data only through k, so k is sufficient.

The rows of μ_t are the binomial distributions $\text{Bin}(t, a)$ and $\text{Bin}(t, b)$.

The draws are exchangeable, as in *Exchangeability and Bayesian Updating*.

Given k' copies of x_1 among $t+1$ draws, the last draw equals x_1 with probability $k'/(t+1)$ in either state.

Discarding it moves k' to $k'-1$ with that probability and leaves k' unchanged otherwise.

The kernel does not depend on the state, so it is a garbling.

```
from scipy.stats import binom

def draws_experiment(θ, t):
    "Experiment matrix for the number of x_1 signals in t independent draws."
    a, b = (1 + θ) / 2, (1 - θ) / 2
    k = np.arange(t + 1)
    return np.vstack([binom.pmf(k, t, a), binom.pmf(k, t, b)])

def discard_last(t):
    "Markov matrix from k' in {0, ..., t+1} to k in {0, ..., t}."
    Q_t = np.zeros((t + 2, t + 1))
    for kp in range(t + 2):
        if kp >= 1:
            Q_t[kp, kp - 1] = kp / (t + 1)
        if kp <= t:
            Q_t[kp, kp] = 1 - kp / (t + 1)
    return Q_t

θ_ex = 0.3
for t in (1, 4, 9):
    err = np.abs(draws_experiment(θ_ex, t + 1) @ discard_last(t)
                 - draws_experiment(θ_ex, t)).max()
    print(f"t = {t}:  max |μ_(t+1) Q_t - μ_t| = {err:.1e}")

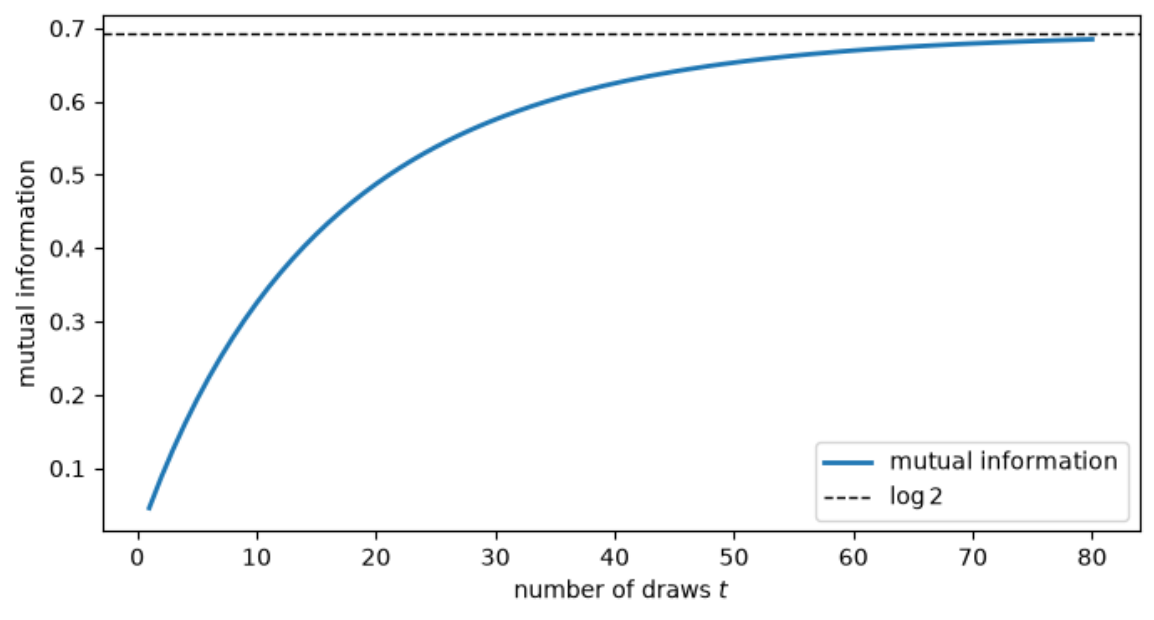
is_mps = all(
    (call_values(draws_experiment(θ_ex, t + 1), c2, ts)
     - call_values(draws_experiment(θ_ex, t), c2, ts)).min() >= -1e-12
    for t in range(1, 30))
print(f"\nposteriors after t + 1 draws are an MPS of those after t draws: {is_mps}
↪")
```

```
t = 1:  max |μ_(t+1) Q_t - μ_t| = 1.1e-16
t = 4:  max |μ_(t+1) Q_t - μ_t| = 9.7e-17
t = 9:  max |μ_(t+1) Q_t - μ_t| = 1.1e-16

posteriors after t + 1 draws are an MPS of those after t draws: True
```

```
t_grid = np.arange(1, 81)
mi = [degroot_value(draws_experiment(θ_ex, t), c2, entropy) for t in t_grid]

fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(t_grid, mi, linewidth=2, label="mutual information")
ax.axhline(np.log(2), color="black", linestyle="--", linewidth=1, label=r"$\log 2$
↪")
ax.set_xlabel("number of draws $t$")
ax.set_ylabel("mutual information")
ax.legend()
plt.show()
```



Each additional draw spreads the posterior distribution while its mean stays at the prior, the martingale property studied in *Likelihood Ratio Processes and Bayesian Learning*.

Mutual information rises toward log 2, the entropy of the uniform prior, as the posterior concentrates on the true state.

CHAPTER

THIRTYONE

INFORMATION AND MARKET EQUILIBRIUM

Contents

- *Information and Market Equilibrium*
  - *Overview*
  - *Setup*
  - *Price revelation*
  - *Invertibility and the elasticity of substitution*
  - *Bayesian price expectations in a dynamic economy*
  - *Simulating Bayesian learning from prices*
  - *Exercises*

31.1 Overview

This lecture studies two questions about the **informational role of prices** posed and answered by Kihlstrom and Mirman [1975].

1. *When do prices transmit inside information?*
   - An informed insider observes a private signal correlated with an unknown state of the world and adjusts demand accordingly.
   - Equilibrium prices shift.
   - Under what conditions can an outside observer *infer* the insider's posterior distribution from the equilibrium price?
2. *Do Bayesian price expectations converge?*
   - In a stationary stochastic exchange economy, an uninformed observer uses the history of market prices and Bayes' Law to form beliefs about the economy's structure and hence about its induced price distribution.
   - Do those expectations eventually agree with those of a fully informed observer?

Kihlstrom and Mirman's answers rely on two classical ideas from statistics:

- **Blackwell sufficiency**: a random variable $\tilde{y}$ is said to be *sufficient* for a random variable $\tilde{y}'$ with respect to an unknown state if knowing $\tilde{y}$ gives all the information about the state that $\tilde{y}'$ contains.

- **Bayesian consistency**: as the sample grows, posterior beliefs eliminate models that imply the wrong **price distribution**, so even when structure is not identified from prices the posterior mass on the true **reduced form** still converges to one.

Important findings of Kihlstrom and Mirman [1975] are:

- Equilibrium prices transmit inside information *if and only if* the map from the insider's posterior distribution to the equilibrium price is one-to-one on the set of posteriors that can actually arise from the signal.
  - For the two-state case ($S = 2$), invertibility holds when the informed agent's utility is homothetic and the elasticity of substitution is everywhere either below one or above one.
- In the dynamic economy, as information accumulates, Bayesian price expectations converge to **rational expectations**, even when the deep structure is not identified from prices alone.

Note

Kihlstrom and Mirman [1975] use the terms "structure" and "reduced form" as careful econometricians do.

A structure is a complete specification of the economy's primitives, here a parameter vector λ.

Each structure implies a probability distribution of observable variables, here the distribution of equilibrium prices, and that distribution is the reduced form.

Several structures can imply the same reduced form.

We call the set of structures that imply the same reduced form a reduced-form class; price data identify the class but not the structure within it.

The lecture is organized as follows.

1. Set up the static two-commodity model and define equilibrium.
2. State the price-revelation theorem and the invertibility conditions.
3. Illustrate invertibility and its failure with numerical examples using CES and Cobb-Douglas preferences.
4. Introduce the dynamic stochastic economy and derive the Bayesian convergence result.
5. Simulate Bayesian learning from price observations.

This lecture builds on ideas in *Blackwell's Theorem on Comparing Experiments* and *Likelihood Ratio Processes and Bayesian Learning*.

We start by importing some Python packages.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import brentq
from scipy.stats import norm
```

31.2 Setup

31.2.1 Preferences, endowments, and the unknown state

The economy has two goods.

Good 2 is the numeraire (price normalized to 1).

Good 1 trades at price $p > 0$.

An unknown parameter $\bar{a}$ affects the value of good 1.

Agent i's expected utility from a bundle (x_1^i, x_2^i) is

$$U^i(x_1^i, x_2^i) = \sum_{s=1}^{S} u^i(a_s x_1^i,\ x_2^i)\, P^i(\bar{a} = a_s),$$

where P^i is agent i's subjective probability distribution over the finite state space $A = \{a_1, \dots, a_S\}$.

Each agent starts with an endowment w^i of good 2 and a share θ^i of the representative firm.

In the paper's formal model, a single firm transforms good 2 into good 1 according to $y_1 = f(y_2)$ with $f' < 0$ and chooses production to maximize

$$\pi(p) = \max_{y_2 \le 0} \{p f(y_2) + y_2\}.$$

The firm's profit π is then distributed to households according to the shares θ^i.

Agent i's budget constraint is

$$p x_1^i + x_2^i = w^i + \theta^i \pi.$$

Agents maximize expected utility subject to their budget constraints.

A **competitive equilibrium** is a price $\hat{p}$ that clears both markets simultaneously.

Under the maintained convexity assumptions equilibrium exists, and following Kihlstrom and Mirman [1975] we assume the equilibrium price is unique.

We write $\hat{p} = p(\mu, P^2, \dots, P^n)$ for the equilibrium price when agent 1 holds beliefs μ and the other agents hold beliefs $P^2, \dots, P^n$.

The beliefs $P^2, \dots, P^n$ stay fixed throughout the static analysis, so we abbreviate the price function to $p(\mu)$.

For most of what follows, the production side matters only through the induced equilibrium price map, so when we turn to numerical illustrations we will suppress production and use a pure-exchange / portfolio interpretation to keep the calculations transparent.

31.2.2 The informed agent's problem

Suppose **agent 1** (the insider) observes a private signal $\tilde{y}$ correlated with $\bar{a}$ before trading, where $\tilde{y}$ takes values in a finite set Y.

Before the signal arrives, agent 1 has prior beliefs $\mu_0 = P^1$.

Upon observing $\tilde{y} = y$, agent 1 updates to the **posterior** $\mu_y = (\mu_{y1}, \dots, \mu_{yS})$ via Bayes' rule:

$$\mu_{ys} = P(\bar{a} = a_s \mid \tilde{y} = y).$$

Because agent 1's demand depends on μ_y, the new equilibrium price satisfies

$$\hat{p} = p(\mu_y).$$

Outside observers who see $\hat{p}$ but not $\tilde{y}$ can try to *back out* the insider's posterior from the price.

Define the set of realized posteriors

$$M = \{\mu_y : y \in Y,\ P(\tilde{y} = y) > 0\}.$$

The key question is whether the map $\mu \mapsto p(\mu)$ is one-to-one on M.

To answer that question, we now translate "information in prices" into Blackwell's language of sufficiency.

31.3 Price revelation

31.3.1 Blackwell sufficiency

The price variable $p(\mu_{\tilde{y}})$ *accurately transmits* the insider's private information if observing the equilibrium price is just as informative about $\bar{a}$ as observing the signal $\tilde{y}$ directly.

In Blackwell's language (Blackwell [1951] and Blackwell [1953]), this means $p(\mu_{\tilde{y}})$ is **sufficient** for $\tilde{y}$.

Definition 31.3.1 (Sufficiency)

A random variable $\tilde{y}$ is *sufficient* for $\tilde{y}'$ with respect to $\bar{a}$ if there exists a conditional distribution $P(y' \mid y)$, **independent of** $\bar{a}$, such that

$$\phi'_a(y') = \sum_{y \in Y} P(y' \mid y)\, \phi_a(y) \quad \text{for all } a \text{ and all } y',$$

where $\phi_a(y) = P(\tilde{y} = y \mid \bar{a} = a)$.

Thus, once $\tilde{y}$ is known, $\tilde{y}'$ provides no additional information about $\bar{a}$.

This is the garbling relation of *Definition 30.2.1* in *Blackwell's Theorem on Comparing Experiments*: the likelihoods $\phi_a(y)$ form the more informative experiment, the likelihoods $\phi'_a(y')$ form its garbling, and $P(y' \mid y)$ is the Markov matrix Q.

Kihlstrom and Mirman [1975] show that

Lemma 31.3.1 (Posterior Sufficiency)

The posterior distribution $\mu_{\tilde{y}}$ is a sufficient statistic for $\tilde{y}$.

Proof. (Sketch) The posterior $\mu_{\tilde{y}}$ satisfies

$$P(\bar{a} = a_s \mid \mu_{\tilde{y}} = \mu_y,\ \tilde{y} = y) = \mu_{ys} = P(\bar{a} = a_s \mid \mu_{\tilde{y}} = \mu_y).$$

This identity says that once the posterior is known, conditioning on the original signal $\tilde{y}$ does not change beliefs about $\bar{a}$.

Equivalently, the conditional law of $\tilde{y}$ given $\mu_{\tilde{y}}$ is independent of $\bar{a}$, so $\mu_{\tilde{y}}$ is sufficient for $\tilde{y}$ in Blackwell's sense.

The lemma is Kihlstrom's standard-experiment point, described under "Posteriors and standard experiments" in *Blackwell's Theorem on Comparing Experiments*: replacing each signal by the posterior it induces loses no information about $\bar{a}$.

Now let's think about the mapping from belief to price.

Theorem 31.3.1 (Price Revelation)

In the model outlined above, the price random variable $p(\mu_{\tilde{y}})$ is sufficient for the random variable $\tilde{y}$ if and only if the function $p(\cdot)$ is invertible on the set of prices

$$\mathcal{P} = \Big\{ p(\mu_y) : y \in Y,\ P(\tilde{y} = y) = \sum_{a \in A} \phi_a(y)\, \mu_0(a) > 0 \Big\}.$$

The logic is

$$\tilde{y} \quad \longrightarrow \quad \mu_{\tilde{y}} \quad \longrightarrow \quad p(\mu_{\tilde{y}}).$$

The first arrow loses no information about $\bar{a}$ by *Lemma 31.3.1*, and the theorem asks when the second arrow also loses no information.

The proof has two parts.

If $p(\cdot)$ is one-to-one on M, then observing the price is equivalent to observing the posterior itself because

$$P(\mu_{\tilde{y}} = \mu \mid p(\mu_{\tilde{y}}) = p) = \begin{cases} 1 & \text{if } \mu = p^{-1}(p), \\ 0 & \text{otherwise.} \end{cases}$$

This conditional distribution is independent of the state, so price is sufficient for the posterior; together with *Lemma 31.3.1*, price is therefore sufficient for the signal.

Conversely, suppose that two different posteriors $\mu_y \neq \mu_{y'}$ in M generate the same price p.

Conditional on that price, an observer's posterior about $\bar{a}$ is a mixture, with positive weights, of the posteriors μ_z for which $p(\mu_z) = p$.

A mixture with positive weight on two different posteriors cannot equal both of them, so learning $\tilde{y}$ after seeing p would change the observer's beliefs about $\bar{a}$ for at least one of y and y'.

The price is therefore not sufficient for $\tilde{y}$.

Before turning to invertibility itself, it helps to keep in mind the two economic interpretations emphasized in the paper.

31.3.2 Two interpretations

Insider trading in a stock market

Good 1 is a risky asset with random return $\bar{a}$; good 2 is "money".

An insider's demand reveals private information about the return.

If the invertibility condition holds, outside observers can read the insider's posterior distribution – the useful information the insider's signal carries about $\bar{a}$ – from the equilibrium stock price.

Price as a quality signal

Good 1 has uncertain quality $\bar{a}$.

Experienced consumers (who have sampled the good) observe a signal correlated with quality and buy accordingly.

Uninformed consumers can infer quality from the market price, provided invertibility holds.

31.3.3 Uninformed traders who learn from prices

Kihlstrom and Mirman's uninformed traders hold fixed beliefs $P^2, \ldots, P^n$ and do not learn from the price.

Theorem 31.3.1 describes what an outside observer who knows the map $p(\cdot)$ could infer.

If the uninformed traders themselves knew $p(\cdot)$ and conditioned on the price, their beliefs and demands would depend on the price, and equilibrium would become a rational expectations equilibrium.

Grossman [1976] constructs a rational expectations equilibrium in which the price fully reveals the informed traders' information.

Radner [1979] shows that with finitely many states fully revealing rational expectations equilibria exist generically.

Grossman and Stiglitz [1980] add a cost of acquiring information.

If the price revealed everything, no trader would pay for information; if no trader paid for information, the price would have nothing to reveal.

They resolve the paradox with a noisy supply of the risky asset, which makes the price reveal information only partially.

Rational Learning and Rational Expectations studies Bayesian traders who learn from prices inside a rational expectations equilibrium.

31.4 Invertibility and the elasticity of substitution

When does the belief-to-price map fail to be invertible?

Theorem 31.4.1 shows that for a two-state economy ($S = 2$), the answer depends on the **elasticity of substitution** σ of agent 1's utility function.

Before stating the theorem, it helps to see the two intermediate steps in the paper's argument.

Lemma 31.4.1 (Same Price Implies Same Allocation)

Assume that u^i has continuous first partial derivatives and that u^i is strictly quasi-concave.

Let $p \in \mathcal{P}$.

If there exist two measures μ^* and μ' in M such that $p(\mu^*) = p(\mu') = p$, then

$$x^i(\mu^*) = x^i(\mu'), \quad i = 1, \ldots, n,$$

where $x^i(\mu)$ is trader i's equilibrium bundle when agent 1's beliefs are μ and the other agents' beliefs are $P^2, \ldots, P^n$.

Fix the beliefs of all agents except agent 1.

The lemma says that if two posterior beliefs μ^* and μ' for agent 1 both support the same equilibrium price p, then they support the same equilibrium allocation for every trader.

The intuition is that when the price is unchanged, the demands of the uninformed traders are unchanged too, so market clearing forces the informed agent's bundle to be unchanged as well.

Strict quasi-concavity makes each trader's demand a single-valued function of the price, and the uninformed traders' beliefs $P^2, \ldots, P^n$ do not move with the price.

This lemma lets us define the informed agent's equilibrium bundle as a function of price alone:

$$x(p) = (x_1(p), x_2(p)).$$

Throughout, u^i_j denotes the partial derivative of u^i with respect to its j-th argument.

Whenever the informed agent consumes positive amounts of both goods, optimality of $x(p)$ under posterior μ gives the interior first-order condition

$$p = \frac{\sum_{s=1}^{S} a_s u_1^1(a_s x_1(p), x_2(p))\, \mu(a_s)}{\sum_{s=1}^{S} u_2^1(a_s x_1(p), x_2(p))\, \mu(a_s)}.$$

For a fixed price p, the bundle $x(p)$ is fixed too, so invertibility boils down to whether this equation admits a unique posterior μ.

Lemma 31.4.2 (Unique Posterior at a Given Price)

Assume that the first partial derivatives of u^1 exist and that u^1 is strictly quasi-concave.

Also assume that agent 1 always consumes positive quantities of both goods.

Then $p(\cdot)$ is invertible on $\mathcal{P}$ if for each $p \in \mathcal{P}$ there exists a unique probability measure $\mu \in M$ such that

$$\frac{\sum_{s=1}^{S} a_s\, u_1^1(a_s x_1(p), x_2(p))\, \mu(a_s)}{\sum_{s=1}^{S} u_2^1(a_s x_1(p), x_2(p))\, \mu(a_s)} = p.$$

If two different posteriors gave the same price, then by *Lemma 31.4.1* they would share the same bundle $x(p)$, contradicting uniqueness of the posterior that solves the first-order condition at that price.

31.4.1 The two-state first-order condition

With $S = 2$ and $\mu = (q,\ 1-q)$, define

$$\alpha_s(p) = a_s\, u_1^1(a_s x_1(p),\ x_2(p)), \qquad \beta_s(p) = u_2^1(a_s x_1(p),\ x_2(p)), \qquad s = 1, 2.$$

Then the first-order condition becomes

$$p = \frac{\alpha_1(p)\, q + \alpha_2(p)\, (1-q)}{\beta_1(p)\, q + \beta_2(p)\, (1-q)}.$$

At a fixed price p, the quantities $\alpha_s(p)$ and $\beta_s(p)$ are constants, so uniqueness of the posterior is the same as uniqueness of the scalar q solving this equation.

Theorem 31.4.1 (Invertibility Conditions)

Assume that the first partial derivatives of u^1 exist and that u^1 is quasi-concave and homothetic.

Also suppose that the informed agent always consumes positive quantities of both goods in all equilibrium allocations.

If $S = 2$ and the elasticity of substitution of u^1 is either always less than one or always greater than one, then $p(\cdot)$ is invertible on $\mathcal{P}$.

If u^1 is Cobb-Douglas (elasticity of substitution constant and equal to one), then $p(\cdot)$ is constant on M.

When $\sigma = 1$ the income and substitution effects exactly cancel, so agent 1's demand for good 1 does not respond to changes in beliefs about $\bar{a}$.

Because the demand is unchanged, the market-clearing price is unchanged too, and the price reveals nothing about the insider's signal.

31.4.2 CES utility

For concreteness we work with a simplified example with the **constant-elasticity-of-substitution** (CES) utility function

$$u(c_1, c_2) = \Big(\frac{c_1^\rho + c_2^\rho}{2}\Big)^{1/\rho}, \qquad \rho \in (-\infty, 0) \cup (0, 1),$$

whose elasticity of substitution is $\sigma = 1/(1-\rho)$.

- $\rho \to 0$: Cobb-Douglas ($\sigma = 1$), $u \to \sqrt{c_1 c_2}$.
- $\rho < 0$: $\sigma < 1$ (complements).
- $0 < \rho < 1$: $\sigma > 1$ (substitutes).

The factor $1/2$ inside the parentheses makes u converge to $\sqrt{c_1 c_2}$ as $\rho \to 0$; without it, u diverges as $\rho \downarrow 0$ and vanishes as $\rho \uparrow 0$.

The factor rescales u_1 and u_2 by a common constant, so it does not affect marginal rates of substitution or prices.

Pertinent partial derivatives are

$$u_1(c_1, c_2) = \frac{1}{2}\Big(\frac{c_1^\rho + c_2^\rho}{2}\Big)^{1/\rho - 1} c_1^{\rho-1}, \qquad u_2(c_1, c_2) = \frac{1}{2}\Big(\frac{c_1^\rho + c_2^\rho}{2}\Big)^{1/\rho - 1} c_2^{\rho-1}.$$

This CES example is only an illustration, because the theorem itself covers any homothetic utility with elasticity everywhere above one or everywhere below one.

With that example in hand, we can compute the equilibrium price directly as a function of the posterior.

31.4.3 Equilibrium price as a function of the posterior

We suppress production.

Uninformed traders own one unit of good 1, have no use for it, and sell it at any positive price, so at equilibrium agent 1 absorbs the unit and $x_1 = 1$.

Agent 1 owns W_1 units of good 2 and no shares of the firm ($\theta^1 = 0$), so his wealth $W_1 = w^1$ does not depend on p.

Agent 1's budget constraint then reduces to $x_2 = W_1 - p$, and the equilibrium price is the unique $p \in (0, W_1)$ satisfying the first-order condition

$$p\big[q\, u_2(a_1,\, W_1 - p) + (1-q)\, u_2(a_2,\, W_1 - p)\big] = q\, a_1\, u_1(a_1,\, W_1 - p) + (1-q)\, a_2\, u_1(a_2,\, W_1 - p).$$

For Cobb-Douglas utility ($\sigma = 1$), the first-order condition becomes $p = W_1 - p$, giving $p^* = W_1/2$ regardless of the posterior q, confirming that no information is transmitted through the price in the Cobb-Douglas case.

We compute first-order conditions numerically below.

```
def ces_derivatives(c1, c2, ρ):
    """
    Return CES marginal utilities.

    Use the Cobb-Douglas limit near rho = 0.
    """
    if abs(ρ) < 1e-4:
        u1 = 0.5 * np.sqrt(c2 / c1)
        u2 = 0.5 * np.sqrt(c1 / c2)
    else:
        common = 0.5 * ((c1**ρ + c2**ρ) / 2)**(1 / ρ - 1)
        u1 = common * c1**(ρ - 1)
        u2 = common * c2**(ρ - 1)
    return u1, u2


def eq_price(q, a1, a2, W1, ρ):
    """Return the equilibrium price for posterior q."""
    def residual(p):
        x2 = W1 - p
        u1_s1, u2_s1 = ces_derivatives(a1, x2, ρ)
        u1_s2, u2_s2 = ces_derivatives(a2, x2, ρ)
        lhs = p * (q * u2_s1 + (1 - q) * u2_s2)
        rhs = q * a1 * u1_s1 + (1 - q) * a2 * u1_s2
        return lhs - rhs

    try:
        return brentq(residual, 1e-6, W1 - 1e-6, xtol=1e-10)
    except ValueError:
        return np.nan
```

```
a1, a2 = 2.0, 0.5     # state values (a1 > a2)
W1 = 4.0

q_grid = np.linspace(0.05, 0.95, 200)

ρ_values = [-0.5, 0.0, 0.5]
ρ_labels = [
    r"$\rho = -0.5$ ($\sigma = 0.67$, complements)",
    r"$\rho = 0$ ($\sigma = 1$, Cobb-Douglas)",
    r"$\rho = 0.5$ ($\sigma = 2$, substitutes)",
]

fig, ax = plt.subplots(figsize=(8, 5))

for ρ, label in zip(ρ_values, ρ_labels):
    prices = [eq_price(q, a1, a2, W1, ρ) for q in q_grid]
    ax.plot(q_grid, prices, label=label, lw=2)

ax.set_xlabel(r"posterior probability $q = \Pr(\bar{a} = a_1)$", fontsize=12)
ax.set_ylabel("equilibrium price $p^*(q)$", fontsize=12)
ax.legend(fontsize=10)
plt.tight_layout()
plt.show()
```

The plot confirms *Theorem 31.4.1*.

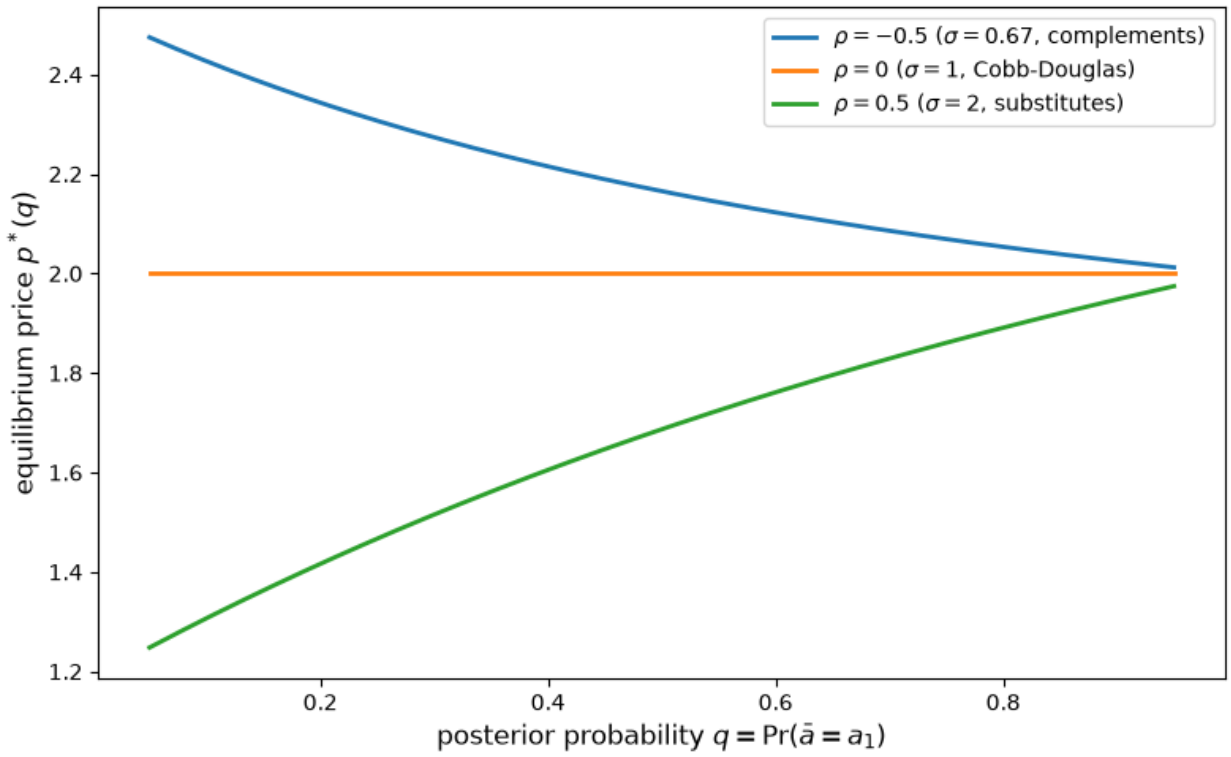


Fig. 31.1: equilibrium price vs posterior

For CES with $\sigma \neq 1$, the equilibrium price is strictly monotone in q.

The price falls with q when $\sigma < 1$ and rises with q when $\sigma > 1$.

When the goods are complements, good news about $\bar{a}$ lowers the price: a higher a_s lowers the marginal rate of substitution $a_s u_1/u_2$ (see *Why monotonicity depends on \sigma*), so the insider values the unit of good 1 less.

An outside observer who knows the equilibrium map $p^*(\cdot)$ can therefore invert the price uniquely to recover q, so the inside information is fully transmitted.

For Cobb-Douglas ($\sigma = 1$), the price is flat in q, so information is never transmitted through the market.

```
p_cd = [eq_price(q, a1, a2, W1, ρ=0.0) for q in q_grid]

print(f"Cobb-Douglas (rho=0): min p* = {min(p_cd):.6f}, "
      f"max p* = {max(p_cd):.6f}, "
      f"range = {max(p_cd)-min(p_cd):.2e}")
print(f"Analytical CD price  = W1/2 = {W1/2:.6f}")
```

```
Cobb-Douglas (rho=0): min p* = 2.000000, max p* = 2.000000, range = 8.88e-16
Analytical CD price  = W1/2 = 2.000000
```

Every entry equals $W_1/2 = 2.0$ up to rounding error, confirming numerically that the Cobb-Douglas equilibrium price is independent of q and of the state values a_1, a_2.

The numerical plot shows monotonicity, and the next subsection connects that pattern back to the proof of *Theorem 31.4.1*.

31.4.4 Why monotonicity depends on σ

Fix a price p and treat $\alpha_s(p)$ and $\beta_s(p)$ as constants.

The right-hand side of the two-state first-order condition

$$\frac{\alpha_1(p)\,q + \alpha_2(p)\,(1-q)}{\beta_1(p)\,q + \beta_2(p)\,(1-q)}$$

is then a function of q alone, with derivative

$$\frac{\partial}{\partial q}\,\frac{\alpha_1 q + \alpha_2(1-q)}{\beta_1 q + \beta_2(1-q)} = \frac{\alpha_1\beta_2 - \alpha_2\beta_1}{\left[\beta_1 q + \beta_2(1-q)\right]^2}.$$

So the sign is determined by $\alpha_1\beta_2 - \alpha_2\beta_1$, and if that sign is constant then for each fixed price there is at most one posterior weight q consistent with the first-order condition, which is exactly what *Theorem 31.4.1* requires.

Using

$$\frac{\alpha_s}{\beta_s} = \frac{a_s\,u_1(a_s x_1, x_2)}{u_2(a_s x_1, x_2)} = a_s^{(\sigma-1)/\sigma}\left(\frac{x_2}{x_1}\right)^{1/\sigma},$$

one can show

$$\frac{\partial}{\partial a}\,\frac{\alpha}{\beta} = \frac{(\sigma-1)}{\sigma}\,a^{-1/\sigma}\left(\frac{x_2}{x_1}\right)^{1/\sigma}.$$

For the CES specification, this derivative is positive when $\sigma > 1$, negative when $\sigma < 1$, and *zero when* $\sigma = 1$.

In other words, for CES utility the ratio α_s/β_s moves monotonically with the state value a_s unless $\sigma = 1$, which makes the fixed-price first-order-condition expression monotone in q and in turn delivers invertibility.

The vanishing derivative in the Cobb-Douglas case means the marginal rate of substitution is independent of a_s, so the informed agent's demand, and hence the equilibrium price, does not respond to changes in beliefs.

Let us visualize the ratio α_s/β_s as a function of a_s for different values of σ:

```
a_vals = np.linspace(0.3, 3.0, 300)
x1_fix, x2_fix = 1.0, 1.0

fig, ax = plt.subplots(figsize=(7, 4))
for ρ in [-0.5, -1e-6, 0.5]:
    σ = 1 / (1 - ρ) if abs(ρ) > 1e-8 else 1.0
    ratios = []
    for a in a_vals:
        u1, u2 = ces_derivatives(a * x1_fix, x2_fix, ρ)
        ratios.append(a * u1 / u2)
    ax.plot(a_vals, ratios, label=rf"$\sigma = {σ:.2f}$", lw=2)

ax.set_xlabel(r"state value $a_s$", fontsize=12)
ax.set_ylabel(r"$\alpha_s / \beta_s = a_s u_1 / u_2$", fontsize=12)
ax.axhline(y=1.0, color="black", lw=0.8, ls="--")
ax.legend(fontsize=10)
plt.tight_layout()
plt.show()
```

When $\sigma = 1$ the ratio is constant across all a_s values, so information about the state has no effect on the marginal rate of substitution.

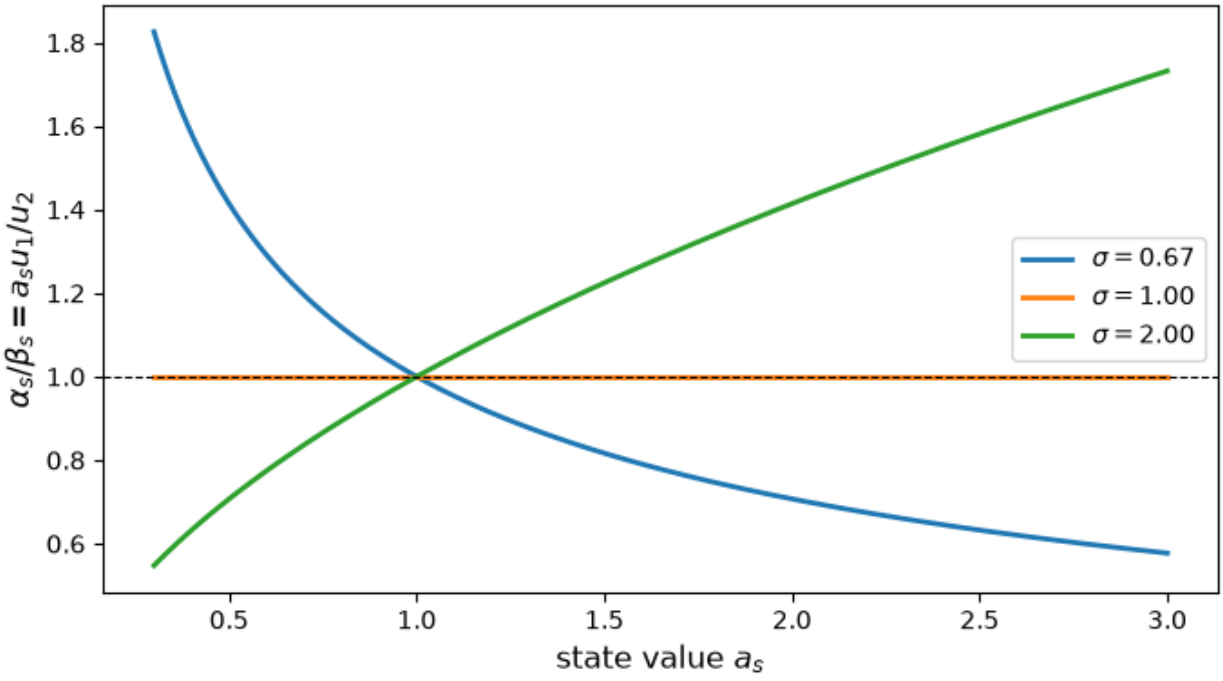


Fig. 31.2: marginal rate of substitution

For $\sigma < 1$ the ratio is decreasing in a_s, and for $\sigma > 1$ it is increasing, making the equilibrium price strictly monotone in the posterior q in both cases.

The static analysis asks whether a current price reveals current private information, whereas the next section asks what a whole history of prices reveals over time.

31.5 Bayesian price expectations in a dynamic economy

We now turn to a question addressed in Section 3 of Kihlstrom and Mirman [1975].

31.5.1 A stochastic exchange economy

Time is discrete: $t = 1, 2, \ldots$

In each period t:

1. Consumer i receives a random endowment ω_i^t.
2. Markets open; competitive prices $p^t = p(\omega^t)$ clear all markets.
3. Consumers trade and consume.

The structural parameter vector $\lambda = (\lambda_1, \ldots, \lambda_K)$ is *fixed but unknown*; it takes values in a finite or countable set $\Lambda \subset \mathbb{R}^K$.

Conditional on λ, the endowment vectors $\{\tilde{\omega}^t\}$ are **i.i.d.** with density $f(\omega^t \mid \lambda)$.

The equilibrium price at time t is a deterministic function of ω^t, so conditional on λ the prices $\{p^t\}$ are also i.i.d.

An observer who does not know λ and puts a prior over Λ regards the price sequence as exchangeable, a mixture of i.i.d. sequences, as in *Exchangeability and Bayesian Updating*.

Let P_λ denote the probability distribution of $\{\omega^t, p^t\}_{t\geq 1}$ under structure λ.

For any measurable set of prices B, let

$$W(B) = \{\omega^t : p(\omega^t) \in B\}.$$

Then

$$P_\lambda(p^t \in B) = P_\lambda(\omega^t \in W(B)) = \int_{W(B)} f(\omega^t \mid \lambda)\, d\omega^t.$$

The induced price density is denoted by $g(p^t \mid \lambda)$.

For a given structure λ, this density is the observable implication of the model, and when several structures imply the same density we group them into a single reduced-form class.

The next issue is therefore what an observer can and cannot infer about the structure from price data alone.

31.5.2 The identification problem

Because price observations identify only the induced price density $g(\cdot \mid \lambda)$, and because the structural-to-reduced-form map $\lambda \mapsto g(\cdot \mid \lambda)$ may be many-to-one, price data may identify only a reduced-form class rather than the exact structure.

In particular, it may be impossible to recover λ from $g(p \mid \lambda)$ even with infinite price data.

To handle this, partition Λ into equivalence classes η such that $\lambda \in \eta$ and $\lambda' \in \eta$ whenever $g(p \mid \lambda) = g(p \mid \lambda')$ for all p.

We use η for reduced-form classes to keep μ for the insider's posterior in the static model.

The equivalence class η containing the true λ is the **reduced form** relevant for price data.

An observer who knows the infinite price history learns η but not necessarily λ.

Once that distinction is clear, Bayesian updating can be written down directly.

31.5.3 Bayesian updating

An uninformed observer begins with a prior $h(\lambda)$ over $\lambda \in \Lambda$.

If the observer could see endowments directly, the posterior would be

$$h(\lambda \mid \omega^1, \dots, \omega^t) = \frac{h(\lambda)\, \prod_{\tau=1}^t f(\omega^\tau \mid \lambda)}{\sum_{\lambda' \in \Lambda} h(\lambda') \prod_{\tau=1}^t f(\omega^\tau \mid \lambda')},$$

and the paper appeals to a Bayesian consistency result to conclude that this posterior concentrates on the true structure $\bar{\lambda}$.

That conclusion requires $h(\bar{\lambda}) > 0$ and identification of the structure from endowments: $f(\cdot \mid \lambda) \neq f(\cdot \mid \lambda')$ whenever $\lambda \neq \lambda'$.

After observing the price sequence $(p^1, \dots, p^t)$, the observer's Bayesian posterior is

$$h(\lambda \mid p^1, \dots, p^t) = \frac{h(\lambda)\, \prod_{\tau=1}^t g(p^\tau \mid \lambda)}{\sum_{\lambda' \in \Lambda} h(\lambda') \prod_{\tau=1}^t g(p^\tau \mid \lambda')}.$$

Price data cannot distinguish structures inside the same reduced-form class.

Indeed, if λ and λ' belong to the same class η, then $g(\cdot \mid \lambda) = g(\cdot \mid \lambda')$, so

$$\frac{h(\lambda \mid p^1, \dots, p^t)}{h(\lambda' \mid p^1, \dots, p^t)} = \frac{h(\lambda)}{h(\lambda')}$$

for every sample history, so the relative odds within an observationally equivalent class never change.

At time t, the observer's price expectations for the next period are

$$g(p^{t+1} \mid p^1, \dots, p^t) = \sum_{\lambda \in \Lambda} g(p^{t+1} \mid \lambda)\, h(\lambda \mid p^1, \dots, p^t).$$

31.5.4 The convergence theorem

Theorem 31.5.1 (Bayesian Convergence)

Let $\bar{\lambda}$ be the true structural parameter and $\bar{\eta}$ the reduced form that contains $\bar{\lambda}$.

Assume the prior assigns positive probability to the reduced-form class $\bar{\eta}$.

Define the posterior mass on a reduced-form class by

$$H_t(\eta) = \sum_{\lambda \in \eta} h(\lambda \mid p^1, \dots, p^t).$$

Because all structures inside a class imply the same $g(\cdot \mid \lambda)$, the predictive density can equivalently be written as

$$g(p^{t+1} \mid p^1, \dots, p^t) = \sum_{\eta} g(p^{t+1} \mid \eta)\, H_t(\eta).$$

Then

$$\lim_{t \to \infty} H_t(\eta) = \begin{cases} 1 & \text{if } \eta = \bar{\eta}, \\ 0 & \text{otherwise,} \end{cases}$$

with $P_{\bar{\lambda}}$-probability one.

Consequently,

$$\lim_{t \to \infty} \left\| g(\cdot \mid p^1, \dots, p^t) - g(\cdot \mid \bar{\eta}) \right\|_{\mathrm{TV}} = 0,$$

where $g(\cdot \mid \bar{\eta})$ is the rational-expectations price distribution for a fully informed observer.

Note

The theorem only requires the prior to assign positive probability to the reduced-form class $\bar{\eta}$ that contains the true structure $\bar{\lambda}$.

This is the absolute-continuity condition of *Merging of Opinions: The Blackwell–Dubins Theorem*, explained below.

It is implied by, but weaker than, assigning positive probability to the true structural parameter $\bar{\lambda}$ itself.

A prior could place zero mass on $\bar{\lambda}$ while still placing positive mass on other structures inside $\bar{\eta}$.

The important distinction is that price observers need not learn $\bar{\lambda}$ itself.

They only learn which reduced-form class is correct.

That is enough for forecasting because every $\lambda \in \bar{\eta}$ generates the same price density $g(\cdot \mid \bar{\eta})$.

Rational price expectations emerge from learning the reduced form, not from identifying every structural detail of the economy.

Here "rational expectations" means that the observer's predictive distribution for next period's price matches the objective price distribution generated by the true reduced form.

31.5.5 Connection to merging of opinions

Theorem 31.5.1 is an instance of the merging theorem of Blackwell and Dubins [1962] studied in *Merging of Opinions: The Blackwell–Dubins Theorem*.

Let Q^* denote the distribution of the price sequence under the true structure, $P_{\bar{\lambda}}$, and let P denote its distribution under the observer's subjective mixture $\sum_\lambda h(\lambda) P_\lambda$.

Every structure in $\bar{\eta}$ induces the same distribution Q^* of prices, so $P(A) \geq H_0(\bar{\eta})\, Q^*(A)$ for every event A, where $H_0(\bar{\eta}) = \sum_{\lambda \in \bar{\eta}} h(\lambda)$.

Hence $Q^*(A) \leq P(A)/H_0(\bar{\eta})$, and $Q^* \ll P$ whenever $H_0(\bar{\eta}) > 0$.

Theorem 34.4.1 then implies that, Q^*-almost surely, the observer's conditional distribution of the entire future price sequence merges in total variation with the true one.

For a class $\eta \neq \bar{\eta}$, prices are i.i.d. with a density $g(\cdot \mid \eta)$ that differs from $g(\cdot \mid \bar{\eta})$.

Theorem 34.7.1 then makes the two distributions of the price sequence mutually singular, and the likelihood ratio $\prod_{\tau=1}^{t} g(p^\tau \mid \eta)/g(p^\tau \mid \bar{\eta})$ converges to zero Q^*-almost surely.

With finitely many classes this gives $H_t(\bar{\eta}) \to 1$.

Exercise 31.7.5 shows that the merging distance for the entire future equals $1 - H_t(\bar{\eta})$, so Blackwell and Dubins's theorem delivers $H_t(\bar{\eta}) \to 1$ with countably many classes as well.

Let's now turn to a simple simulation.

31.6 Simulating Bayesian learning from prices

We illustrate the theorem with a two-model example.

Two possible reduced forms η_1 and η_2 generate prices $p^t \sim N(\bar{p}_i, \sigma_p^2)$ for $i = 1, 2$ respectively.

The observer knows the two possible price distributions (the reduced forms) but not which one governs the data.

This is a **Bayesian model selection** problem we have seen in *Likelihood Ratio Processes and Bayesian Learning*.

With a prior h_0 on η_1 and the observed price p^t, the posterior weight on η_1 after period t is

$$h_t = \frac{h_{t-1}\, g(p^t \mid \eta_1)}{h_{t-1}\, g(p^t \mid \eta_1) + (1 - h_{t-1})\, g(p^t \mid \eta_2)}.$$

We consider a numerical example with two normal distributions with different means

```
def simulate_bayesian_learning(
    p_bar_true, p_bar_alt, σ_p, T, h0, n_paths, seed=42
):
    """Simulate posterior learning between two Gaussian reduced forms."""
```

(continues on next page)

(continued from previous page)

```
    rng = np.random.default_rng(seed)
    h_paths = np.zeros((n_paths, T + 1))
    h_paths[:, 0] = h0

    for path in range(n_paths):
        h = h0
        prices = rng.normal(p_bar_true, σ_p, size=T)
        for t, p in enumerate(prices):
            g_true = norm.pdf(p, loc=p_bar_true, scale=σ_p)
            g_alt = norm.pdf(p, loc=p_bar_alt, scale=σ_p)
            denom = h * g_true + (1 - h) * g_alt
            h = h * g_true / denom
            h_paths[path, t + 1] = h

    return h_paths


def plot_bayesian_learning(h_paths, p_bar_true, p_bar_alt, ax):
    """Plot posterior beliefs over time."""
    T = h_paths.shape[1] - 1
    t_grid = np.arange(T + 1)

    for path in h_paths:
        ax.plot(t_grid, path, alpha=0.25, lw=0.8, color="steelblue")

    median_path = np.median(h_paths, axis=0)
    ax.plot(t_grid, median_path, color="navy", lw=2, label="median posterior")

    ax.axhline(
        y=1.0,
        color="black",
        ls="--",
        lw=1.2,
        label="true model weight = 1",
    )
    ax.set_xlabel("period $t$", fontsize=12)
    ax.set_ylabel(r"$h_t$ = posterior weight on true model", fontsize=12)
    ax.legend(fontsize=10)
```

We consider two cases, one that is easy to learn and another one that is harder to learn, using $T = 300$ periods, $n = 40$ simulated paths, a diffuse prior $h_0 = 0.5$, and common standard deviation $\sigma_p = 0.4$.

- *Easy case*: true model $N(2.0,\, 0.4^2)$, alternative $N(1.2,\, 0.4^2)$.
- *Hard case*: true model $N(2.0,\, 0.4^2)$, alternative $N(1.8,\, 0.4^2)$.

Whether easy or hard to learn depends on "how close" the true distribution is compared to the alternative hypothesis, as measured by the Kullback–Leibler divergence of *Statistical Divergence Measures* (Exercise 31.7.2).

```
T = 300
h0 = 0.5      # diffuse prior
n_paths = 40
σ_p = 0.4

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

# Distinct reduced forms
```

(continues on next page)

(continued from previous page)

```
p_bar_true, p_bar_alt = 2.0, 1.2
h_paths = simulate_bayesian_learning(p_bar_true, p_bar_alt, σ_p, T, h0, n_paths)
plot_bayesian_learning(h_paths, p_bar_true, p_bar_alt, axes[0])

# Similar reduced forms
p_bar_true, p_bar_alt = 2.0, 1.8
h_paths_hard = simulate_bayesian_learning(
    p_bar_true, p_bar_alt, σ_p, T, h0, n_paths
)
plot_bayesian_learning(h_paths_hard, p_bar_true, p_bar_alt, axes[1])

plt.tight_layout()
plt.show()
```

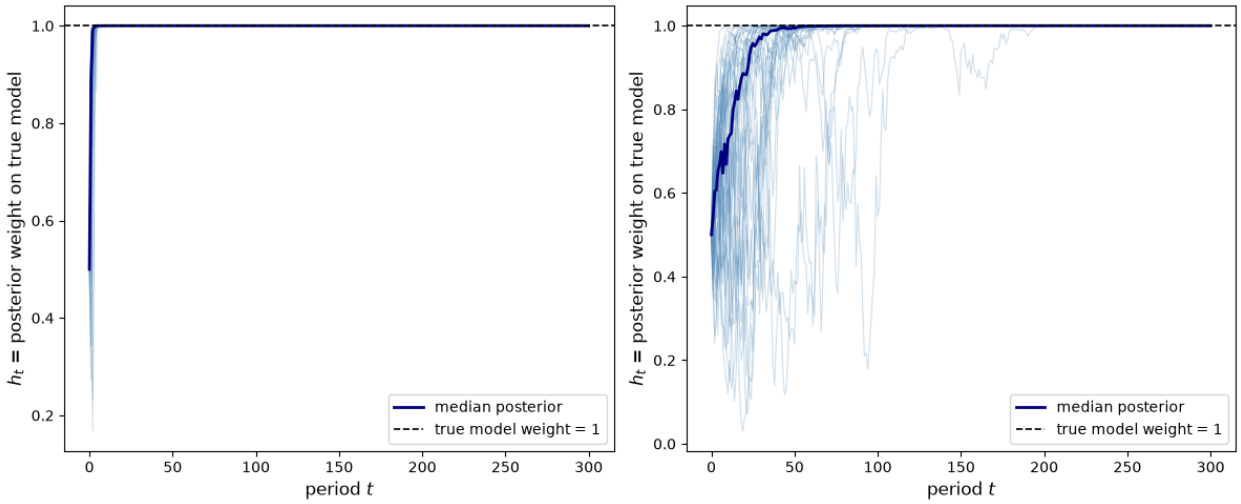


Fig. 31.3: bayesian learning across paths

In both panels the posterior weight on the true model converges to 1 with probability one, though convergence is slower when the two price distributions are similar (right panel).

31.6.1 Price expectations vs. rational expectations

We now verify that the observer's price expectations converge to the rational-expectations distribution $g(p \mid \bar{\eta})$.

We use the parameterization of the "hard-to-learn" example above ($\bar{p}_{\text{true}} = 2.0$, $\bar{p}_{\text{alt}} = 1.8$, $\sigma_p = 0.4$), extending to $T = 1{,}000$ periods with a single simulated path and prior $h_0 = 0.5$

```
def price_expectation(h_t, p_bar_true, p_bar_alt, σ_p, p_grid):
    """Return the predictive price density at posterior weight h_t."""
    return (
        h_t * norm.pdf(p_grid, loc=p_bar_true, scale=σ_p)
        + (1 - h_t) * norm.pdf(p_grid, loc=p_bar_alt, scale=σ_p)
    )
```

(continues on next page)

(continued from previous page)

```
p_bar_true, p_bar_alt = 2.0, 1.8
σ_p = 0.4
n_paths = 1
T_long = 1000

h_paths_long = simulate_bayesian_learning(
    p_bar_true, p_bar_alt, σ_p, T_long, h0=0.5, n_paths=n_paths, seed=7
)

p_grid = np.linspace(0.0, 3.5, 300)
re_density = norm.pdf(p_grid, loc=p_bar_true, scale=σ_p)

fig, ax = plt.subplots(figsize=(8, 5))
snapshots = [0, 25, 100, 300, 1000]
palette   = plt.cm.Blues(np.linspace(0.3, 1.0, len(snapshots)))

for t_snap, col in zip(snapshots, palette):
    h_t = h_paths_long[0, t_snap]
    dens = price_expectation(h_t, p_bar_true, p_bar_alt, σ_p, p_grid)
    ax.plot(
        p_grid,
        dens,
        color=col,
        lw=2,
        label=rf"$t = {t_snap}$, $h_t = {h_t:.3f}$",
    )

ax.plot(p_grid, re_density, "k--", lw=2,
        label=r"rational expectations $g(p \mid \bar{\eta})$")
ax.set_xlabel("price $p$", fontsize=12)
ax.set_ylabel("density", fontsize=12)
ax.legend(fontsize=9)
plt.tight_layout()
plt.show()
```

The sequence of predictive densities (shades of blue) converges to the rational-expectations density (dashed black line) as experience accumulates.

Convergence need not be monotone.

On this path $h_{25} = 0.183$, below the prior $h_0 = 0.5$, before h_t rises toward one, because the log posterior odds follow a random walk whose positive drift dominates only in long samples.

This illustrates *Theorem 31.5.1*.

We can now sharpen the point by looking at a case in which the reduced form is learned but the underlying structure is not.

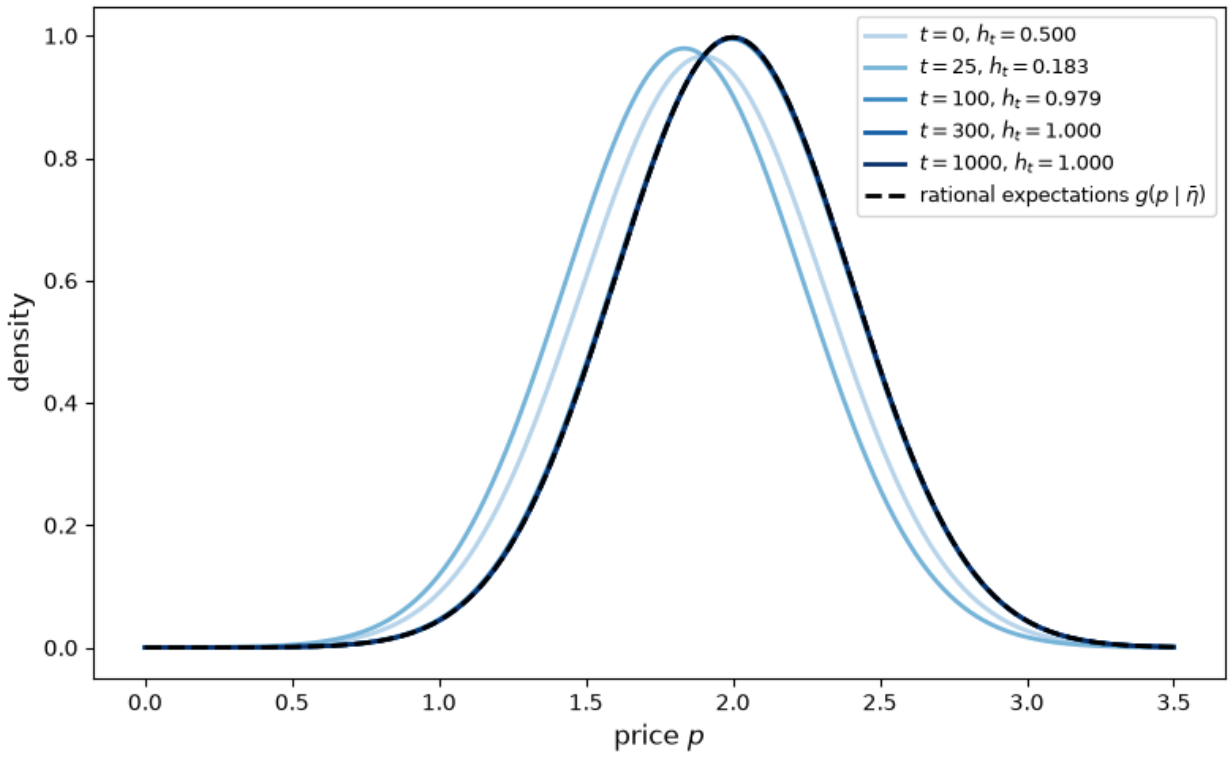


Fig. 31.4: price distribution convergence

31.6.2 Learning the reduced form without identifying the structure

The convergence result is particularly striking because the observer converges to *rational expectations* even when the underlying **structure** λ is *not identified* by prices.

To illustrate this, consider a case with *three* possible structures $\lambda^{(1)}, \lambda^{(2)}, \lambda^{(3)}$ but only *two* reduced forms $\eta_1 = \{\lambda^{(1)}, \lambda^{(2)}\}$ and $\eta_2 = \{\lambda^{(3)}\}$ (because $\lambda^{(1)}$ and $\lambda^{(2)}$ generate the same price distribution).

We continue with the hard-to-learn parameterization, so the three structures have price means $\bar{p}_1 = \bar{p}_2 = 2.0$ and $\bar{p}_3 = 1.8$, with common standard deviation $\sigma_p = 0.4$, a uniform prior $h_0 = (1/3, 1/3, 1/3)$, and $T = 400$ periods over 30 paths.

The true structure is $\lambda^{(1)}$.

```
def simulate_learning_3struct(
    T, h0_vec, p_bar_vec, σ_p, true_idx, n_paths, seed=0
):
    """Simulate learning with three structures and two reduced forms."""
    rng = np.random.default_rng(seed)
    h_paths = np.zeros((n_paths, T + 1, 3))
    h_paths[:, 0, :] = h0_vec

    for path in range(n_paths):
        h = np.array(h0_vec, dtype=float)
        prices = rng.normal(p_bar_vec[true_idx], σ_p, size=T)
        for t, p in enumerate(prices):
            likelihoods = norm.pdf(p, loc=p_bar_vec, scale=σ_p)
            h = h * likelihoods
```

(continues on next page)

(continued from previous page)

```
            h /= h.sum()
            h_paths[path, t + 1, :] = h

    return h_paths


# Structures 0 and 1 share the same reduced form
p_bar_vec = np.array([2.0, 2.0, 1.8])
h0_vec = np.array([1 / 3, 1 / 3, 1 / 3])
σ_p = 0.4
T = 400
true_idx = 0      # Structure 0 is observationally equivalent to 1

h_paths_3 = simulate_learning_3struct(
    T, h0_vec, p_bar_vec, σ_p, true_idx, n_paths=30
)
t_grid = np.arange(T + 1)

fig, axes = plt.subplots(1, 3, figsize=(13, 4), sharey=True)
struct_labels = [
    r"$\lambda^{(1)}$",
    r"$\lambda^{(2)}$",
    r"$\lambda^{(3)}$",
]

for k, (ax, label) in enumerate(zip(axes, struct_labels)):
    for path in h_paths_3:
        ax.plot(t_grid, path[:, k], alpha=0.25, lw=0.8, color="steelblue")
    ax.plot(t_grid, np.median(h_paths_3[:, :, k], axis=0),
            color="navy", lw=2, label=f"median weight on {label}")
    ax.set_xlabel("period $t$", fontsize=11)
    ax.legend(fontsize=9)

axes[0].set_ylabel("posterior weight", fontsize=11)
plt.tight_layout()
plt.show()
```

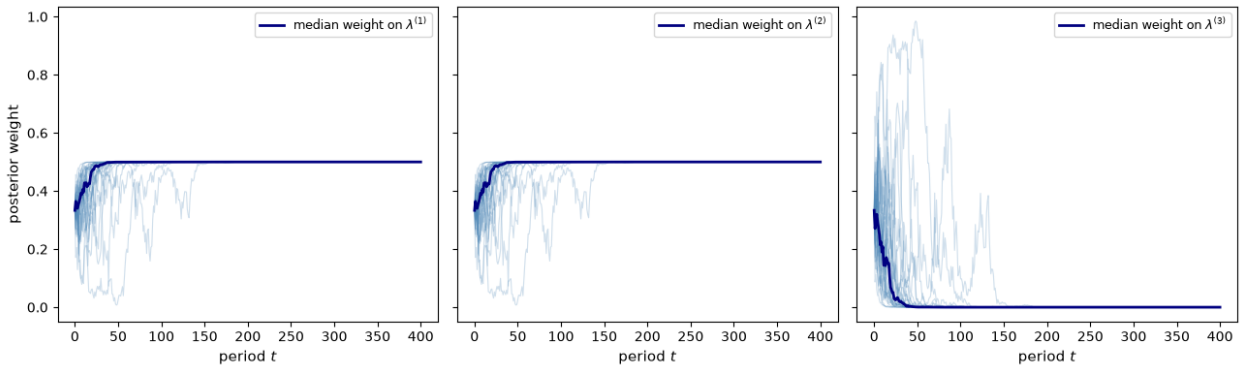


Fig. 31.5: learning with non-identification

The observer correctly rules out $\lambda^{(3)}$ (the wrong reduced form) with probability one, but cannot distinguish $\lambda^{(1)}$ from $\lambda^{(2)}$ because they generate an identical price distribution.

Nevertheless, the observer's **price expectations** converge to rational expectations because both structures imply the same reduced form $\bar{\eta}$.

The next figure plots the posterior mass on the true class, $H_t(\bar{\eta}) = h_t(\lambda^{(1)}) + h_t(\lambda^{(2)})$, and the share of that mass assigned to $\lambda^{(1)}$.

```
H_true = h_paths_3[:, :, 0] + h_paths_3[:, :, 1]
share_1 = h_paths_3[:, :, 0] / H_true

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
for path_H, path_s in zip(H_true, share_1):
    axes[0].plot(t_grid, path_H, alpha=0.25, lw=0.8, color="steelblue")
    axes[1].plot(t_grid, path_s, alpha=0.25, lw=0.8, color="steelblue")
axes[0].plot(t_grid, np.median(H_true, axis=0), color="navy", lw=2,
             label="median")
axes[0].set_ylabel(r"$H_t(\bar\eta)$", fontsize=11)
axes[1].set_ylabel(r"$h_t(\lambda^{(1)}) / H_t(\bar\eta)$", fontsize=11)
axes[1].set_ylim(0, 1)
for ax in axes:
    ax.set_xlabel("period $t$", fontsize=11)
axes[0].legend(fontsize=9)
plt.tight_layout()
plt.show()

print(f"max over paths of 1 - H_T:          {(1 - H_true[:, -1]).max():.1e}")
print(f"max deviation of share from 1/2:  {np.abs(share_1 - 0.5).max():.1e}")
```

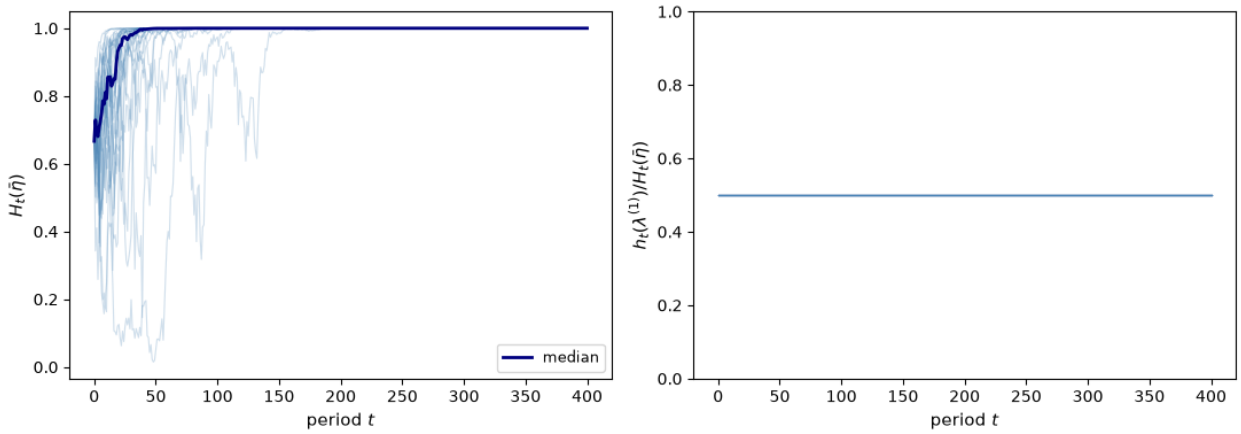


Fig. 31.6: posterior mass on the true reduced-form class

```
max over paths of 1 - H_T:          1.3e-13
max deviation of share from 1/2:  0.0e+00
```

$H_t(\bar{\eta})$ converges to one on every path.

The share $h_t(\lambda^{(1)})/H_t(\bar{\eta})$ stays at its prior value $1/2$ for every t, because $\lambda^{(1)}$ and $\lambda^{(2)}$ assign the same likelihood to every price.

31.7 Exercises

Exercise 31.7.1

CARA portfolio utility and the stock-market interpretation.

Consider a two-state economy ($a_1 = 2$, $a_2 = 0.5$) where the informed agent has **CARA** (constant absolute risk aversion) preferences over portfolio wealth:

$$u(W) = -e^{-\gamma W}, \quad W = x_2 + \bar{a}\, x_1.$$

The agent chooses x_1 to maximize

$$q\, u(W_1) + (1-q)\, u(W_2), \quad W_s = w - p\, x_1 + a_s\, x_1,$$

subject to the budget constraint $p\, x_1 + x_2 = w$.

Total supply of good 1 is $X_1 = 1$.

1. Derive the first-order condition for the informed agent's optimal x_1.
2. Use the market-clearing condition $x_1 = 1$ (the informed agent absorbs the entire supply) to obtain an implicit equation for the equilibrium price $p^*(q)$, and solve it numerically for $q \in (0,1)$ and several values of γ.
3. Show *analytically* that $p^*(q)$ admits the closed form

$$p^*(q) = \frac{a_2 + R(q,\gamma)\, a_1}{1 + R(q,\gamma)}, \qquad R(q,\gamma) = \frac{q}{1-q}\, e^{-\gamma(a_1 - a_2)},$$

and verify that $p^*(q)$ is strictly increasing in q.

Solution

For the first-order condition, define $W_s = w + (a_s - p)\, x_1$ for $s = 1, 2$.

Then the FOC is

$$q\,(a_1 - p)\,\gamma\, e^{-\gamma W_1} = (1-q)\,(p - a_2)\,\gamma\, e^{-\gamma W_2},$$

or equivalently, after dividing both sides by $\gamma e^{-\gamma w}$,

$$q\,(a_1 - p)\, e^{-\gamma(a_1 - p)x_1} = (1-q)\,(p - a_2)\, e^{\gamma(p - a_2)x_1}.$$

Setting $x_1 = 1$ (the informed agent absorbs all supply), this becomes a scalar root-finding problem in p:

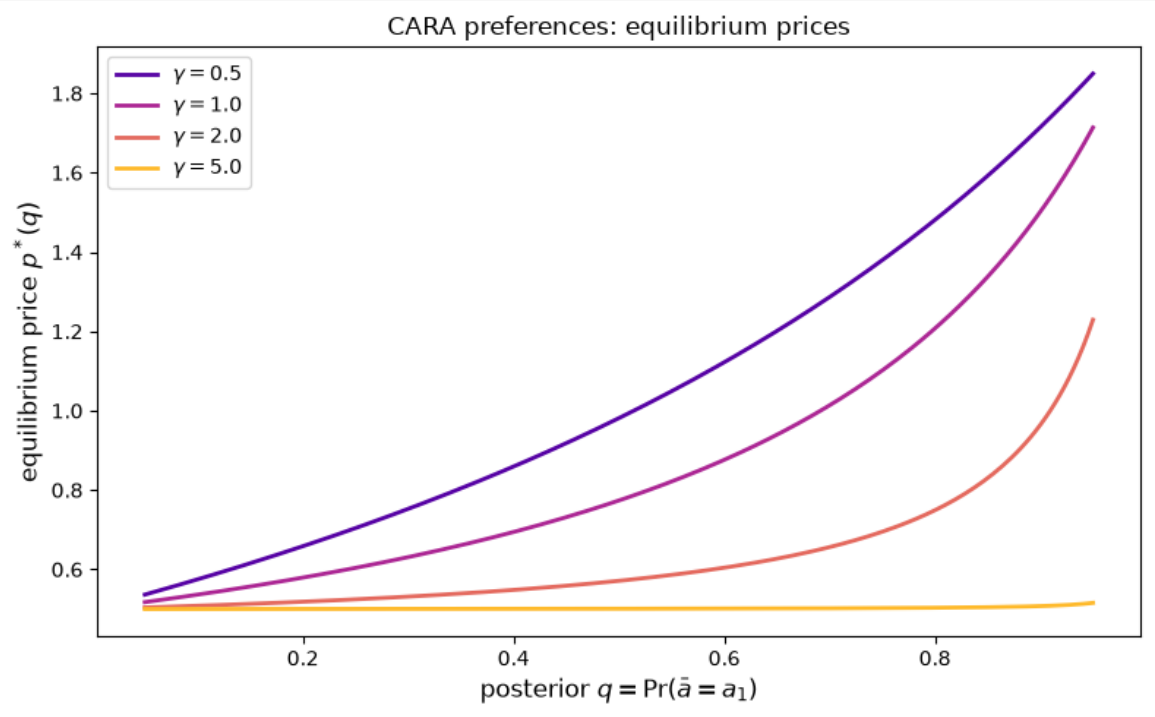
$$F(p;\, q, \gamma) \equiv q\,(a_1 - p)\, e^{-\gamma(a_1 - p)} - (1-q)\,(p - a_2)\, e^{\gamma(p - a_2)} = 0.$$

```
def F_cara(p, q, a1, a2, γ, x1=1.0):
    """Residual for the CARA equilibrium condition."""
    return (q * (a1 - p) * np.exp(-γ * (a1 - p) * x1)
            - (1 - q) * (p - a2) * np.exp(γ * (p - a2) * x1))

a1, a2 = 2.0, 0.5
q_grid = np.linspace(0.05, 0.95, 200)
γ_values = [0.5, 1.0, 2.0, 5.0]
colors_sol = plt.cm.plasma(np.linspace(0.15, 0.85, len(γ_values)))
```

```
fig, ax = plt.subplots(figsize=(8, 5))
for γ, color in zip(γ_values, colors_sol):
    p_eq = [brentq(F_cara, a2, a1,
                   args=(q, a1, a2, γ))
            for q in q_grid]
    ax.plot(q_grid, p_eq, lw=2, color=color,
            label=rf"$\gamma = {γ}$")

ax.set_xlabel(r"posterior $q = \Pr(\bar a = a_1)$", fontsize=12)
ax.set_ylabel("equilibrium price $p^*(q)$", fontsize=12)
ax.set_title("CARA preferences: equilibrium prices", fontsize=12)
ax.legend(fontsize=10)
plt.tight_layout()
plt.show()
```



The price is strictly increasing in q for every $\gamma > 0$.

For the closed form, start from the FOC at $x_1 = 1$, divide both sides by $(1-q)(p-a_2)e^{-\gamma(a_1-p)}$, and combine the exponentials:

$$\frac{q\,(a_1 - p)}{(1-q)\,(p-a_2)} = e^{\gamma(a_1 - a_2)}.$$

Rearranging gives

$$\frac{p-a_2}{a_1-p} = \frac{q}{1-q}\,e^{-\gamma(a_1-a_2)} \equiv R(q,\gamma),$$

and solving the resulting linear equation in p yields

$$p^*(q) = \frac{a_2 + R(q,\gamma)\,a_1}{1 + R(q,\gamma)}.$$

Since $R(q,\gamma)$ is strictly increasing in q and $dp^*/dR = (a_1 - a_2)/(1+R)^2 > 0$, the equilibrium price $p^*(q)$ is strictly increasing in q.

The closed form is an exponentially tilted expectation.

With $x_1 = 1$ we have $e^{-\gamma W_s} = e^{-\gamma(w-p)}e^{-\gamma a_s}$, so the first-order condition is $\sum_s \mu_s(a_s - p)e^{-\gamma a_s} = 0$ for a posterior $\mu = (\mu_1, \ldots, \mu_S)$, and

$$p^* = \frac{\sum_s \mu_s\, a_s\, e^{-\gamma a_s}}{\sum_s \mu_s\, e^{-\gamma a_s}}.$$

The price is the expected value of $\bar{a}$ under the probabilities $\tilde{\mu}_s \propto \mu_s e^{-\gamma a_s}$, which tilt the insider's posterior toward low-payoff states.

With $\mu = (q, 1-q)$, dividing the numerator and denominator by $(1-q)e^{-\gamma a_2}$ gives the formula for $p^*(q)$ above.

```
def p_tilted(μ, a, γ):
    """Equilibrium price as an exponentially tilted expectation."""
    weights = np.asarray(μ) * np.exp(-γ * np.asarray(a))
    return weights @ np.asarray(a) / weights.sum()


gap = max(abs(brentq(F_cara, a2, a1, args=(q, a1, a2, γ))
              - p_tilted([q, 1 - q], [a1, a2], γ))
          for γ in γ_values for q in q_grid)
print(f"largest gap between root-finding and closed form: {gap:.1e}")
```

```
largest gap between root-finding and closed form: 5.0e-13
```

This exercise uses the stock-market interpretation emphasized by Kihlstrom and Mirman [1975].

Portfolio wealth is $W = x_2 + \bar{a}\,x_1$, so ax_1 and x_2 are perfect substitutes in each state.

Hence the elasticity of substitution between the two arguments of $u(ax_1, x_2)$ is infinite, corresponding to the $\sigma > 1$ side of *Theorem 31.4.1*.

The difference is that this example is not the full equilibrium model the theorem analyzes, but rather a partial equilibrium model with a single informed agent and a fixed supply of the risky asset.

Exercise 31.7.2

In the Bayesian learning simulation, the speed of convergence to rational expectations is determined by the **Kullback-Leibler divergence** (*Statistical Divergence Measures*) between the two reduced forms.

The KL divergence $D_{KL}(\eta_1 \| \eta_2)$ from $g(\cdot \mid \eta_1)$ to $g(\cdot \mid \eta_2)$, for two normal distributions with means $\bar{p}_1$ and $\bar{p}_2$ and common variance σ_p^2, is

$$D_{KL}(\eta_1 \| \eta_2) = \frac{(\bar{p}_1 - \bar{p}_2)^2}{2\sigma_p^2},$$

which is symmetric in the two means under equal variances.

Let η_1, with $\bar{p}_1 = 2.0$, be the true reduced form, and set $\sigma_p = 0.4$ and $h_0 = 1/2$.

1. Compute D_{KL} for the "easy" case ($\bar{p}_2 = 1.2$), the "hard" case ($\bar{p}_2 = 1.8$), the intermediate cases $\bar{p}_2 \in \{1.4, 1.6\}$, and a harder case $\bar{p}_2 = 1.9$.
2. For each case simulate 1,000 paths and compute the first period $T_{0.99}$ at which $h_t \geq 0.99$.

Plot histograms of $T_{0.99}$ for the easy and hard cases.

3. Show that under the true model the log posterior odds $\ell_t = \log[h_t/(1-h_t)]$ follow a random walk with drift D_{KL} and per-period variance $2D_{KL}$.

   Use Wald's identity to argue that

$$\mathbb{E}[T_{0.99}] \approx \frac{\log[(0.99/0.01)/(h_0/(1-h_0))]}{D_{KL}}$$

   when D_{KL} is small, so that $T_{0.99} \approx C/D_{KL}$ with $C \approx \log 99 \approx 4.6$ for $h_0 = 1/2$.

4. Plot the mean and median of $T_{0.99}$ against D_{KL} on log-log scales, together with the line $\log 99/D_{KL}$, and explain the deviations.

Solution

Here is one solution.

Write $\Delta = \bar{p}_1 - \bar{p}_2$.

The log likelihood ratio of one price is

$$\xi = \log \frac{g(p \mid \eta_1)}{g(p \mid \eta_2)} = \frac{\Delta\,(2p - \bar{p}_1 - \bar{p}_2)}{2\sigma_p^2}.$$

Under the true model $p = \bar{p}_1 + \sigma_p \varepsilon$ with $\varepsilon \sim N(0,1)$, so $\xi = \Delta^2/(2\sigma_p^2) + (\Delta/\sigma_p)\varepsilon$, which has mean D_{KL} and variance $\Delta^2/\sigma_p^2 = 2D_{KL}$.

Bayes' law adds one such increment to ℓ_t each period, so ℓ_t is a random walk with drift D_{KL} and variance $2D_{KL}$ per period.

$T_{0.99}$ is the first time ℓ_t reaches $b = \log 99$.

Wald's identity gives $\mathbb{E}[\ell_T] = \ell_0 + D_{KL}\,\mathbb{E}[T]$, so

$$\mathbb{E}[T_{0.99}]\,D_{KL} = b - \ell_0 + \mathbb{E}[\ell_T - b].$$

The last term is the overshoot of the boundary.

It is of the order of one step's standard deviation $\sqrt{2D_{KL}}$, so it vanishes relative to $b - \ell_0$ as $D_{KL} \to 0$.

```
def log_odds_paths(p_true, p_alt, σ, h0, n_paths, T, seed=42):
    """Paths of log posterior odds on the true model, t = 1, ..., T."""
    rng = np.random.default_rng(seed)
    p = rng.normal(p_true, σ, size=(n_paths, T))
    ξ = norm.logpdf(p, p_true, σ) - norm.logpdf(p, p_alt, σ)
    return np.log(h0 / (1 - h0)) + np.cumsum(ξ, axis=1)


def first_passage(ell, b):
    """First period with ell_t >= b, and the value of ell_t then."""
    hit = ell >= b
    assert hit.any(axis=1).all()            # every path crosses within T
    T_hit = hit.argmax(axis=1)
    return T_hit + 1, ell[np.arange(len(ell)), T_hit]


σ_p, h0, b = 0.4, 0.5, np.log(0.99 / 0.01)
p_alts = [1.2, 1.4, 1.6, 1.8, 1.9]
results = {}

print(f"{'p_alt':>6}{'D_KL':>8}{'mean T':>9}{'median T':>10}"
      f"{'mean T*D':>10}{'log 99 + overshoot':>20}{'median T*D':>12}"
      f"{'sqrt(2D)':>10}")
for p_alt in p_alts:
    D = (2.0 - p_alt)**2 / (2 * σ_p**2)
    ell = log_odds_paths(2.0, p_alt, σ_p, h0, n_paths=1_000, T=3_000)
    T99, ell_T = first_passage(ell, b)
    results[p_alt] = (D, T99)
    print(f"{p_alt:6.1f}{D:8.4f}{T99.mean():9.2f}{np.median(T99):10.1f}"
          f"{T99.mean() * D:10.3f}{b + (ell_T - b).mean():20.3f}"
          f"{np.median(T99) * D:12.3f}{np.sqrt(2 * D):10.3f}")

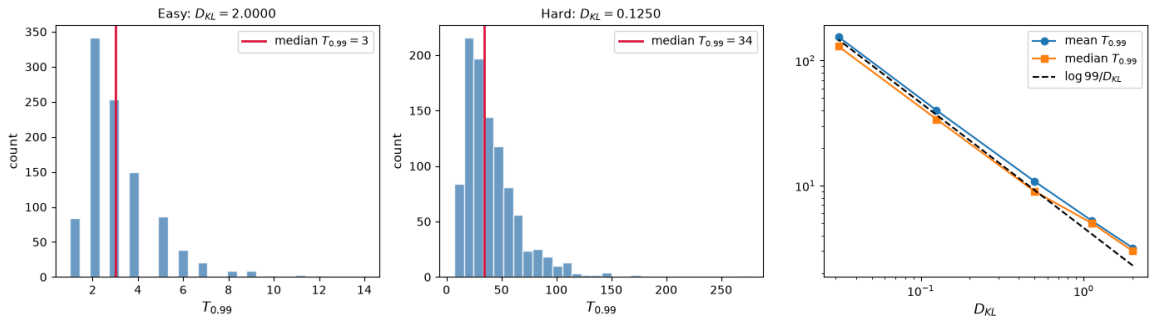
# Brownian limit: T is inverse Gaussian with mean b/D and shape b^2/(2D)
from scipy.stats import invgauss
print(f"\nlog 99 = {b:.3f};  median/mean in the Brownian limit = "
      f"{(b / 2) * invgauss(mu=2 / b).median():.3f}")
```

```
 p_alt    D_KL   mean T  median T  mean T*D  log 99 + overshoot  median T*D  ␣
↪sqrt(2D)
   1.2  2.0000     3.18       3.0     6.356               6.335       6.000  ␣
↪2.000
   1.4  1.1250     5.23       5.0     5.879               5.773       5.625  ␣
↪1.500
   1.6  0.5000    10.80       9.0     5.402               5.307       4.500  ␣
↪1.000
   1.8  0.1250    40.03      34.0     5.004               4.921       4.250  ␣
↪0.500
   1.9  0.0313   154.86     130.0     4.839               4.746       4.063  ␣
↪0.250

log 99 = 4.595;  median/mean in the Brownian limit = 0.825
```

```
fig, axes = plt.subplots(1, 3, figsize=(14, 4))
for ax, (name, p_alt) in zip(axes[:2], [("Easy", 1.2), ("Hard", 1.8)]):
    D, T99 = results[p_alt]
    ax.hist(T99, bins=30, color="steelblue", edgecolor="white", alpha=0.8)
    ax.axvline(np.median(T99), color="crimson", lw=2,
               label=fr"median $T_{{0.99}} = {np.median(T99):.0f}$")
    ax.set_title(f"{name}: $D_{{KL}} = {D:.4f}$", fontsize=11)
    ax.set_xlabel(r"$T_{0.99}$", fontsize=12)
    ax.set_ylabel("count", fontsize=11)
    ax.legend(fontsize=10)

D_vals = np.array([results[p][0] for p in p_alts])
axes[2].loglog(D_vals, [results[p][1].mean() for p in p_alts], "o-",
               label=r"mean $T_{0.99}$")
axes[2].loglog(D_vals, [np.median(results[p][1]) for p in p_alts], "s-",
               label=r"median $T_{0.99}$")
axes[2].loglog(D_vals, b / D_vals, "k--", label=r"$\log 99 / D_{KL}$")
axes[2].set_xlabel(r"$D_{KL}$", fontsize=12)
axes[2].legend(fontsize=10)
plt.tight_layout()
plt.show()
```



Both the mean and the median of $T_{0.99}$ fall roughly in proportion to $1/D_{KL}$, so learning is faster when the two reduced forms are easier to distinguish.

The product $\mathbb{E}[T_{0.99}]D_{KL}$ falls toward $\log 99 = 4.6$ as D_{KL} shrinks.

By Wald's identity the gap equals the expected overshoot, and the columns "mean T*D" and "log 99 + overshoot" agree up to sampling error.

In the easy case a single price moves ℓ_t by about two units, so the boundary is overshot by a wide margin and $\mathbb{E}[T_{0.99}]D_{KL}$ exceeds $\log 99$ substantially.

The median lies below the mean because the first-passage time is skewed to the right.

In the small-D_{KL} limit ℓ_t behaves like a Brownian motion with drift, $T_{0.99}$ is inverse Gaussian with mean $\log 99/D_{KL}$ and shape $(\log 99)^2/(2D_{KL})$, and its median is $0.825 \log 99/D_{KL}$.

Exercise 31.7.3

Theorem 31.5.1 requires the prior to assign positive probability to the true reduced-form class $\bar{\eta}$, equivalently to some structure that generates the true price distribution $g(\cdot \mid \bar{\eta})$.

In this exercise the true reduced form itself is excluded from the prior support, so we investigate what happens when no model in the prior generates the true price distribution.

Simulate $T = 1,000$ periods of prices from $N(2.0, 0.4^2)$ but use a prior that places equal weight on two *wrong* models: $N(1.5, 0.4^2)$ and $N(2.3, 0.4^2)$.

Plot the posterior weight on each model over time.

Discuss your findings.

Solution

Here is one solution:

```
def simulate_misspecified(
    T, p_bar_true, p_bar_wrong, σ_p, h0, n_paths, seed=0
):
    """Simulate learning under a misspecified two-model prior."""
    rng = np.random.default_rng(seed)
    h_paths = np.zeros((n_paths, T + 1, 2))
    h_paths[:, 0, :] = h0

    for path in range(n_paths):
        h = np.array(h0, dtype=float)
        prices = rng.normal(p_bar_true, σ_p, size=T)
        for t, price in enumerate(prices):
            likes = norm.pdf(price, loc=p_bar_wrong, scale=σ_p)
            h = h * likes
            h /= h.sum()
            h_paths[path, t + 1, :] = h

    return h_paths


def predictive_density(weights, means, σ_p, p_grid):
    """Return the predictive density under the current posterior weights."""
    density = np.zeros_like(p_grid)
    for weight, mean in zip(weights, means):
        density += weight * norm.pdf(p_grid, loc=mean, scale=σ_p)
    return density


T = 1000
p_true = 2.0
p_wrong = np.array([1.5, 2.3])
σ_p = 0.4
h0 = np.array([0.5, 0.5])
n_paths = 30

h_misspec = simulate_misspecified(T, p_true, p_wrong, σ_p, h0, n_paths)

kl_vals = (p_true - p_wrong)**2 / (2 * σ_p**2)
for mean, kl in zip(p_wrong, kl_vals):
    print(f"KL(true || N({mean:.1f}, σ^2)) = {kl:.4f}")

t_grid = np.arange(T + 1)
fig, axes = plt.subplots(1, 2, figsize=(12, 4))

labels = [r"$N(1.5, \sigma^2)$", r"$N(2.3, \sigma^2)$"]
for ax, k, label in zip(axes, [0, 1], labels):
```

```
    for path in h_misspec:
        ax.plot(t_grid, path[:, k], alpha=0.2, lw=0.8, color="steelblue")
    ax.plot(t_grid, np.median(h_misspec[:, :, k], axis=0),
            color="navy", lw=2, label="median")
    ax.set_title(f"Posterior weight on {label}", fontsize=11)
    ax.set_xlabel("period $t$", fontsize=11)
    ax.set_ylabel("posterior weight", fontsize=11)
    ax.legend(fontsize=9)

plt.tight_layout()
plt.show()

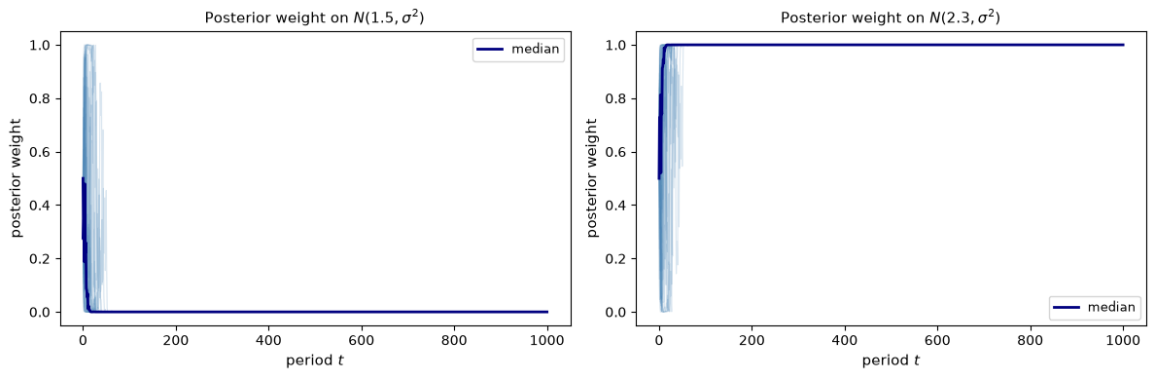
# Predictive density and mean along the median posterior path
median_path = np.median(h_misspec, axis=0)
p_grid = np.linspace(0.0, 3.5, 300)
closer_idx = np.argmin(kl_vals)

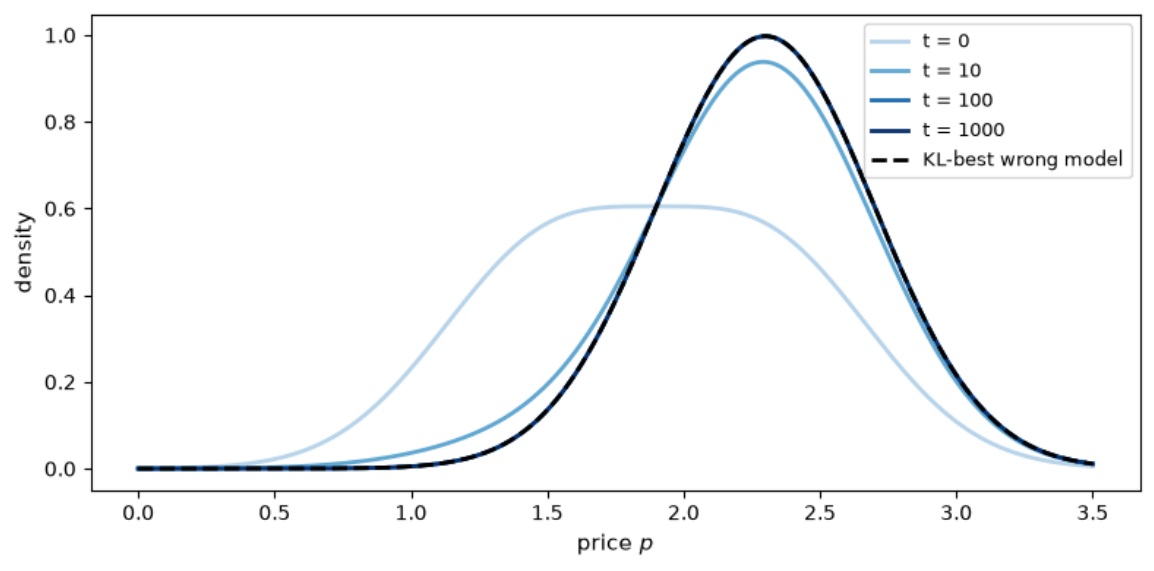
fig, ax = plt.subplots(figsize=(8, 4))
colors = plt.cm.Blues(np.linspace(0.3, 1.0, 4))
for t_snap, color in zip([0, 10, 100, T], colors):
    dens = predictive_density(median_path[t_snap], p_wrong, σ_p, p_grid)
    ax.plot(p_grid, dens, color=color, lw=2, label=f"t = {t_snap}")

ax.plot(
    p_grid,
    norm.pdf(p_grid, loc=p_wrong[closer_idx], scale=σ_p),
    "k--",
    lw=2,
    label="KL-best wrong model",
)
ax.set_xlabel("price $p$", fontsize=11)
ax.set_ylabel("density", fontsize=11)
ax.legend(fontsize=9)
plt.tight_layout()
plt.show()

pred_mean = np.median(
    h_misspec[:, :, 0] * p_wrong[0] + h_misspec[:, :, 1] * p_wrong[1], axis=0
)
print(f"True mean: {p_true}")
print(f"Predictive mean at T={T}: {pred_mean[-1]:.4f}")
print(f"Closer misspecified mean: {p_wrong[np.argmin(kl_vals)]:.1f}")
```

```
KL(true || N(1.5, σ^2)) = 0.7812
KL(true || N(2.3, σ^2)) = 0.2812
True mean: 2.0
Predictive mean at T=1000: 2.3000
Closer misspecified mean: 2.3
```





Here

$$D_{KL}(N(2.0, 0.4^2) \,\|\, N(2.3, 0.4^2)) < D_{KL}(N(2.0, 0.4^2) \,\|\, N(1.5, 0.4^2)),$$

so the model with mean 2.3 is the KL-best approximation among the two wrong models, and in the simulation posterior weight concentrates on that model.

Posterior odds are cumulative *likelihood ratios*.

Write $g_\star$ for the true price density $N(2.0, 0.4^2)$, write g_1 and g_2 for the wrong models $N(1.5, 0.4^2)$ and $N(2.3, 0.4^2)$, and write $h_t^{(1)}, h_t^{(2)}$ for the posterior weights on them.

The log posterior odds on g_2 against g_1 are

$$\log \frac{h_t^{(2)}}{h_t^{(1)}} = \log \frac{h_0^{(2)}}{h_0^{(1)}} + \sum_{\tau=1}^{t} \log \frac{g_2(p^\tau)}{g_1(p^\tau)}.$$

Under $g_\star$ each term in the sum has mean $D_{KL}(g_\star \| g_1) - D_{KL}(g_\star \| g_2) = 0.78125 - 0.28125 = 0.5 > 0$, by the formula in *A helpful formula* of *Likelihood Ratio Processes*.

So the odds drift toward g_2, the model that is closer to $g_\star$ in Kullback–Leibler divergence, and $h_t^{(2)} \to 1$.

The predictive density converges to g_2 rather than to $g_\star$: price expectations settle down, but not on rational expectations.

The prior puts no mass on the true reduced form, so the absolute-continuity condition behind *Theorem 31.5.1* fails.

Incorrect Models studies the same mechanism when the true distribution is a mixture that lies outside the agent's set of models.

Exercise 31.7.4

Theorem 31.4.1 covers two states.

This exercise shows what changes with three.

Use the CES economy of the lecture with $\sigma = 2$ ($\rho = 0.5$), $W_1 = 4$, and three states $(a_1, a_2, a_3) = (2, 1, 0.5)$.

1. Show that at a fixed price p the first-order condition is linear in the posterior μ, so the set of posteriors that support p is the intersection of a line with the simplex.
2. Write a function that computes the equilibrium price for any posterior μ, find two distinct posteriors that support the same price, and verify the result numerically.
3. Plot the level sets of $p(\mu)$ over the simplex.
4. Explain why $p(\cdot)$ can nonetheless be one-to-one on the finite set M of posteriors generated by a signal with finitely many values, and why the two-state theorem is not contradicted.

   Check the claim with randomly drawn signal likelihoods.

Solution

Here is one solution.

At a fixed price the bundle $x(p) = (1, W_1 - p)$ is fixed, and the first-order condition reads

$$\sum_{s=1}^{3} \mu_s \, c_s(p) = 0, \qquad c_s(p) = a_s u_1(a_s, W_1 - p) - p \, u_2(a_s, W_1 - p).$$

The numbers $c_s(p)$ do not depend on μ.

Together with $\sum_s \mu_s = 1$, this is two linear equations in three unknowns.

Their solutions form a line, and the line meets the simplex in a segment of posteriors that all support the price p.

With two states the simplex is itself a segment, and the single equation $q\,c_1(p) + (1-q)\,c_2(p) = 0$ has at most one solution unless $c_1(p) = c_2(p) = 0$, which a constant sign of $\alpha_1\beta_2 - \alpha_2\beta_1$ rules out.

```
def eq_price_S(μ, a, W1, ρ):
    """Equilibrium price when agent 1 holds posterior μ over states a."""
    μ, a = np.asarray(μ, dtype=float), np.asarray(a, dtype=float)

    def residual(p):
        u1, u2 = ces_derivatives(a, W1 - p, ρ)
        return μ @ (a * u1 - p * u2)

    return brentq(residual, 1e-6, W1 - 1e-6, xtol=1e-12)


a3, W1_3, ρ_3 = np.array([2.0, 1.0, 0.5]), 4.0, 0.5
μ_star = np.array([0.5, 0.2, 0.3])
p_star = eq_price_S(μ_star, a3, W1_3, ρ_3)

u1, u2 = ces_derivatives(a3, W1_3 - p_star, ρ_3)
c = a3 * u1 - p_star * u2
v = np.linalg.svd(c, np.ones(3))            # v @ c = 0 and v.sum() = 0

# range of t with μ_star + t v in the simplex
t_lo = max(-μ_star[i] / v[i] for i in range(3) if v[i] > 0)
t_hi = min(-μ_star[i] / v[i] for i in range(3) if v[i] < 0)
```

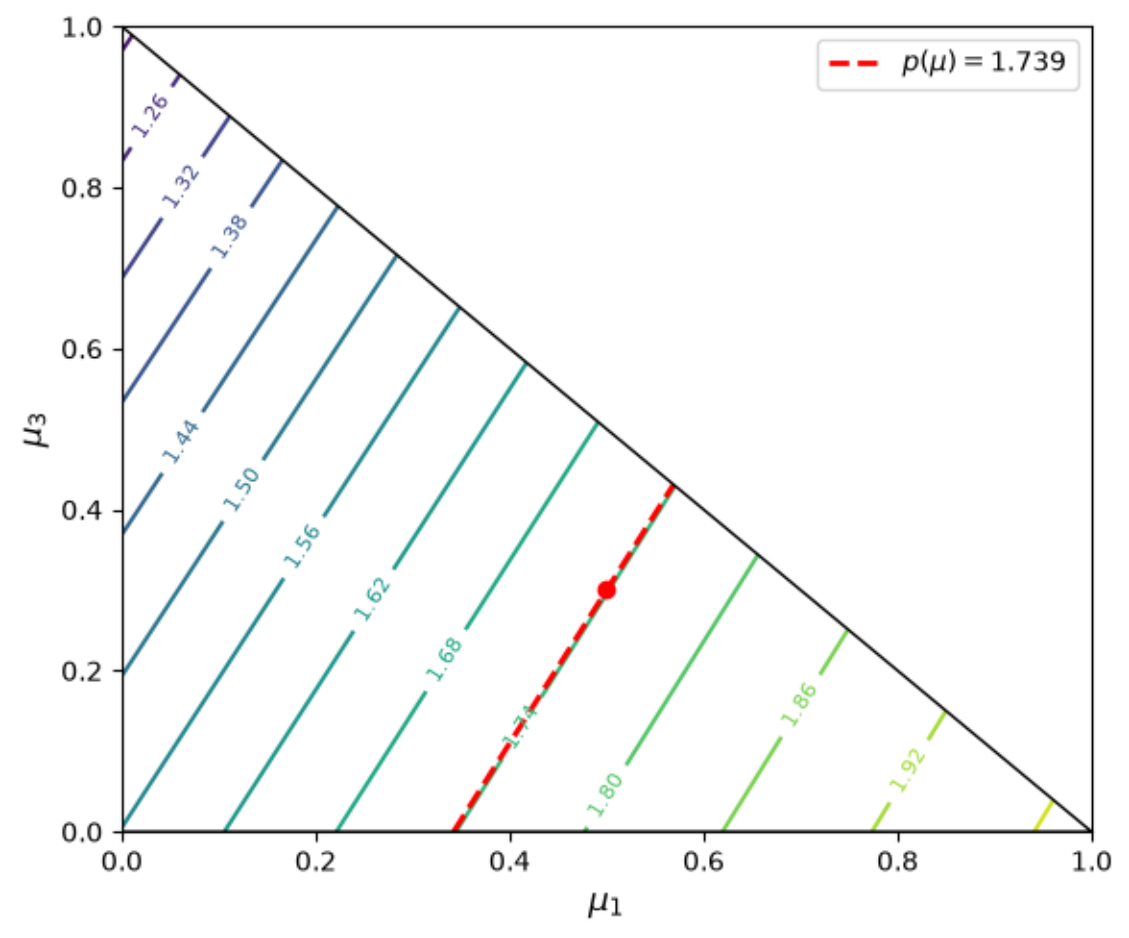
```
p(μ_star) = 1.7388112483

μ = [0.3737 0.5663 0.06  ]   p(μ) = 1.7388112483
μ = [0.4369 0.3831 0.18  ]   p(μ) = 1.7388112483
μ = [0.5 0.2 0.3]   p(μ) = 1.7388112483
μ = [0.5276 0.12   0.3524]   p(μ) = 1.7388112483
μ = [0.5552 0.04   0.4048]   p(μ) = 1.7388112483
```

Five distinct posteriors support the same price to ten decimal places.

```
n = 81
grid = np.linspace(0, 1, n)
P = np.full((n, n), np.nan)
for i, m1 in enumerate(grid):
    for j, m3 in enumerate(grid):
        if m1 + m3 <= 1:
            P[j, i] = eq_price_S([m1, 1 - m1 - m3, m3], a3, W1_3, ρ_3)

fig, ax = plt.subplots(figsize=(6, 5))
cs = ax.contour(grid, grid, P, levels=15, cmap="viridis")
ax.clabel(cs, fontsize=8, fmt="%.2f")
ends = np.array([μ_star + t * v for t in [t_lo, t_hi]])
ax.plot(ends[:, 0], ends[:, 2], "r--", lw=2, label=fr"$p(\mu) = {p_star:.3f}$")
ax.plot(μ_star[0], μ_star[2], "ro")
ax.plot([0, 1, 0, 0], [0, 0, 1, 0], color="black", lw=1)
ax.set_xlabel(r"$\mu_1$", fontsize=12)
ax.set_ylabel(r"$\mu_3$", fontsize=12)
ax.legend(fontsize=10)
plt.tight_layout()
plt.show()
```



The level sets of $p(\mu)$ are straight lines across the simplex, so $p(\cdot)$ is not one-to-one on the whole simplex.

Theorem 31.3.1 requires less: $p(\cdot)$ must be one-to-one on the finite set M of posteriors that the signal can generate.

Two posteriors in M share a price only if one lies on the level line through the other.

That coincidence imposes one equation on the signal likelihoods, and a small perturbation of the likelihoods destroys it.

```
rng = np.random.default_rng(1)
μ0 = np.ones(3) / 3
gaps = []
for _ in range(500):
    ψ = rng.dirichlet(np.ones(4), size=3)        # P(signal | state), 4 signals
    joint = μ0[:, None] * ψ
    posteriors = (joint / joint.sum(axis=0)).T   # one posterior per signal
    prices = np.sort([eq_price_S(m, a3, W1_3, ρ_3) for m in posteriors])
    gaps.append(np.diff(prices).min())

print(f"smallest price gap within a draw, over 500 draws: {min(gaps):.2e}")
```

```
smallest price gap within a draw, over 500 draws: 8.10e-05
```

In every draw the four posteriors support four distinct prices, so the price is sufficient for the signal.

With finitely many signal values, invertibility on M holds generically even though it fails on the simplex.

Exercise 31.7.5

This exercise measures merging of opinions in the two-model example.

Take the hard case: the true reduced form $\bar{\eta} = \eta_1$ has $\bar{p}_1 = 2.0$, the alternative η_2 has $\bar{p}_2 = 1.8$, $\sigma_p = 0.4$, and $h_0 = 1/2$.

Write $H_t = H_t(\bar{\eta})$.

1. Show that $\|g(\cdot \mid p^1, \ldots, p^t) - g(\cdot \mid \bar{\eta})\|_{\text{TV}} = (1 - H_t)\,\|g(\cdot \mid \eta_2) - g(\cdot \mid \bar{\eta})\|_{\text{TV}}$, and evaluate the last factor for two normal densities with a common variance.
2. Show that for the joint distribution of the next k prices the distance is $(1 - H_t)\bigl[2\Phi\bigl(\sqrt{k}\,|\bar{p}_1 - \bar{p}_2| / (2\sigma_p)\bigr) - 1\bigr]$, where Φ is the standard normal distribution function.

   Conclude that the distance d_t between the observer's and the true conditional distributions of the entire future price sequence is $1 - H_t$.
3. Compute and plot both distances along 40 simulated paths of length 300.
4. Relate the results to *Theorem 34.4.1* and *Theorem 34.7.1* in *Merging of Opinions: The Blackwell–Dubins Theorem*.

Solution

Here is one solution.

The predictive density is $H_t\, g(\cdot \mid \bar{\eta}) + (1 - H_t)\, g(\cdot \mid \eta_2)$, so

$$g(\cdot \mid p^1, \ldots, p^t) - g(\cdot \mid \bar{\eta}) = (1 - H_t)\bigl[g(\cdot \mid \eta_2) - g(\cdot \mid \bar{\eta})\bigr],$$

and the total variation distance scales by $1 - H_t$.

Two normal densities with a common variance cross at the midpoint of their means, so

$$\|N(\bar{p}_1, \sigma_p^2) - N(\bar{p}_2, \sigma_p^2)\|_{\text{TV}} = 2\Phi\Bigl(\frac{|\bar{p}_1 - \bar{p}_2|}{2\sigma_p}\Bigr) - 1.$$

Conditional on the reduced form, the next k prices are i.i.d., so the same argument applies to their joint density.

Their sum is a sufficient statistic for the mean, so the distance between the two k-fold products equals the distance between $N(k\bar{p}_1, k\sigma_p^2)$ and $N(k\bar{p}_2, k\sigma_p^2)$, which is $2\Phi(\sqrt{k}\,|\bar{p}_1 - \bar{p}_2|/(2\sigma_p)) - 1$.

As $k \to \infty$ this distance rises to one.

The two i.i.d. price laws are mutually singular, as *Theorem 34.7.1* asserts, so $d_t = 1 - H_t$.

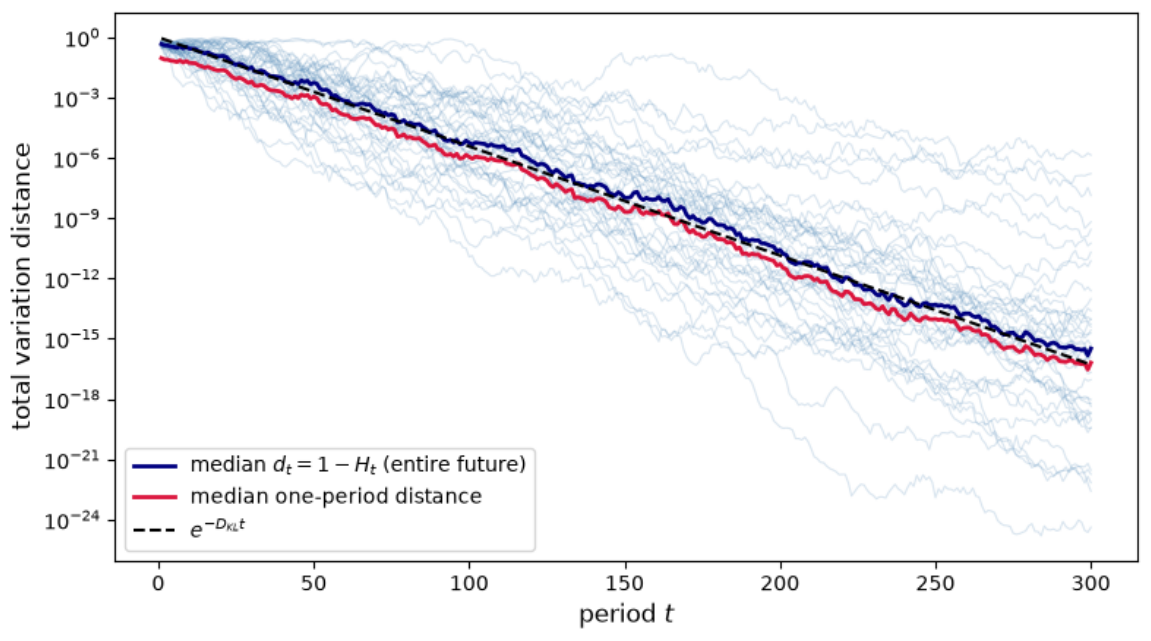
```
from scipy.special import expit

Δ = 2.0 - 1.8
tv_one = 2 * norm.cdf(Δ / (2 * σ_p)) - 1
print(f"one-period distance between the two models: {tv_one:.4f}")
for k in [1, 10, 100, 1_000]:
    print(f"  k = {k:5d}:  2Φ(√k Δ / 2σ) - 1 = "
          f"{2 * norm.cdf(np.sqrt(k) * Δ / (2 * σ_p)) - 1:.4f}")

ell = log_odds_paths(2.0, 1.8, σ_p, h0=0.5, n_paths=40, T=300)
one_minus_H = expit(-ell)                       # 1 - H_t, computed stably
t_vals = np.arange(1, 301)

fig, ax = plt.subplots(figsize=(8, 4.5))
for path in one_minus_H:
    ax.semilogy(t_vals, path, color="steelblue", alpha=0.2, lw=0.8)
ax.semilogy(t_vals, np.median(one_minus_H, axis=0), color="navy", lw=2,
            label=r"median $d_t = 1 - H_t$ (entire future)")
ax.semilogy(t_vals, tv_one * np.median(one_minus_H, axis=0), color="crimson",
            lw=2, label="median one-period distance")
D_hard = Δ**2 / (2 * σ_p**2)
ax.semilogy(t_vals, np.exp(-D_hard * t_vals), "k--", lw=1.5,
            label=r"$e^{-D_{KL} t}$")
ax.set_xlabel("period $t$", fontsize=12)
ax.set_ylabel("total variation distance", fontsize=12)
ax.legend(fontsize=10)
plt.tight_layout()
plt.show()
```

```
one-period distance between the two models: 0.1974
  k =     1:  2Φ(√k Δ / 2σ) - 1 = 0.1974
  k =    10:  2Φ(√k Δ / 2σ) - 1 = 0.5708
  k =   100:  2Φ(√k Δ / 2σ) - 1 = 0.9876
  k =  1000:  2Φ(√k Δ / 2σ) - 1 = 1.0000
```



The observer's subjective distribution of the price sequence is $P = h_0 Q^* + (1 - h_0) Q_2$, where Q^* and Q_2 make prices i.i.d. with densities $g(\cdot \mid \bar{\eta})$ and $g(\cdot \mid \eta_2)$.

Because $h_0 > 0$, $Q^* \ll P$, and *Theorem 34.4.1* asserts that $d_t \to 0$ almost surely under Q^*.

Because Q^* and Q_2 are mutually singular, $d_t = 1 - H_t$, so merging of opinions about the entire future and $H_t \to 1$ are the same event.

With countably many classes the argument is the same.

Each class $\eta \neq \bar{\eta}$ assigns probability zero to some event of Q^*-probability one, and the intersection of these countably many events still has Q^*-probability one, so again $d_t = 1 - H_t(\bar{\eta})$.

The one-period distance is smaller than d_t by the factor 0.197, because one price distinguishes $N(2.0, 0.4^2)$ from $N(1.8, 0.4^2)$ poorly.

Both distances fall at the exponential rate $D_{KL} = 0.125$ per period, the drift of the log posterior odds.

CHAPTER

THIRTYTWO

INCORRECT MODELS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

```
!pip install numpyro
```

32.1 Overview

This is a sequel to *Likelihood Ratio Processes and Bayesian Learning*.

We discuss two ways to create a compound lottery and their consequences.

A compound lottery can be said to create a *mixture distribution*.

Our two ways of constructing a compound lottery will differ in their *timing*.

- in one, mixing between two possible probability distributions will occur once and for all at the beginning of time
- in the other, mixing between the same two possible probability distributions will occur each period

The statistical setting is close but not identical to the problem studied in that lecture.

In that lecture, there were two IID processes that could possibly govern successive draws of a non-negative random variable W.

Nature decided once and for all whether to make a sequence of IID draws from either f or from g.

That lecture studied an agent who knew both f and g but did not know which distribution nature chose at time -1.

The agent represented that ignorance by assuming that nature had chosen f or g by flipping an unfair coin that put probability π_{-1} on probability distribution f.

That assumption allowed the agent to construct a subjective joint probability distribution over the random sequence $\{W_t\}_{t=0}^{\infty}$.

We studied how the agent would then use the laws of conditional probability and an observed history $w^t = \{w_s\}_{s=0}^t$ to form

$$\pi_t = \mathbb{P}\{q = f \mid w^t\}, \quad t = 0, 1, 2, ...$$

where $q \in \{f, g\}$ denotes the distribution that nature chose.

However, in the setting of this lecture, that rule imputes to the agent an incorrect model.

The reason is that now the wage sequence is actually described by a different statistical model.

Thus, we change the specification in *Likelihood Ratio Processes and Bayesian Learning* in the following way.

Now, **each period** $t \geq 0$, nature flips a possibly unfair coin that comes up f with probability α and g with probability $1 - \alpha$.

Thus, nature perpetually draws from the **mixture distribution** with CDF

$$H(w) = \alpha F(w) + (1 - \alpha)G(w), \quad \alpha \in (0, 1)$$

and density

$$h(w|\alpha) = \alpha f(w) + (1 - \alpha)g(w).$$

We'll study two agents who try to learn about the wage process, but who use different statistical models.

Both types of agent know f and g but neither knows α.

Our first type of agent erroneously thinks that at time -1, nature once and for all chose f or g and thereafter permanently draws from that distribution.

Our second type of agent knows, correctly, that nature mixes f and g with mixing probability $\alpha \in (0, 1)$ each period, though the agent doesn't know the mixing parameter.

Our first type of agent applies the learning algorithm described in *Likelihood Ratio Processes and Bayesian Learning*.

In the context of the statistical model that prevailed in that lecture, that was a good learning algorithm and it enabled the Bayesian learner eventually to learn the distribution that nature had drawn at time -1.

This is because the agent's statistical model was *correct* in the sense of being aligned with the data generating process.

But in the present context, our type 1 decision maker's model is incorrect because the model h that actually generates the data is neither f nor g and so is beyond the support of the models that the agent thinks are possible.

Nevertheless, we'll see that our first type of agent muddles through and eventually learns something interesting and useful, even though it is not *true*.

Instead, it turns out that our type 1 agent who is armed with a wrong statistical model ends up learning whichever probability distribution, f or g, is in a special sense *closest* to the h that actually generates the data.

We'll tell the sense in which it is closest.

Our second type of agent understands that nature mixes between f and g each period with a fixed mixing probability α.

But the agent doesn't know α.

The agent sets out to learn α using Bayes' law applied to his model.

His model is correct in the sense that it includes the actual data generating process h as a possible distribution.

Both agents' statistical models are exchangeable mixtures of IID sequences of the kind described by de Finetti's theorem in *Exchangeability and Bayesian Updating*.

The type 1 agent's mixing measure puts probability π_{-1} on IID draws from f and probability $1 - \pi_{-1}$ on IID draws from g.

The type 2 agent's mixing measure is a prior density $\pi_{-1}(\alpha)$ over IID draws from $h(\cdot|\alpha)$, $\alpha \in (0,1)$.

Nature's process, IID draws from $h(\cdot|\alpha)$, lies inside the support of the type 2 agent's mixing measure and outside the support of the type 1 agent's mixing measure.

That is the sense in which the type 1 agent's model is misspecified.

The two timings of mixing leave different marks on the data.

When nature mixes anew each period, successive draws are independent and $\text{Cov}(W_0, W_1) = 0$.

When nature mixes once and for all with probability π on f, $\text{Cov}(W_0, W_1) = \pi(1-\pi)(\mu_f - \mu_g)^2$, where μ_f and μ_g are the means of f and g; see Exercise Exercise **??** in *Exchangeability and Bayesian Updating*.

Merging of Opinions: The Blackwell–Dubins Theorem studies agents who, like our type 2 agent, mix IID sequences with a prior over a parameter.

In this lecture, we'll learn about

- how nature can *mix* between two distributions f and g to create a new distribution h.
- The Kullback-Leibler statistical divergence that governs statistical learning under an incorrect statistical model
- A useful Python function `numpy.searchsorted` that, in conjunction with a uniform random number generator, can be used to sample from an arbitrary distribution

As usual, we'll start by importing some Python tools.

```
import matplotlib.pyplot as plt
import seaborn as sns
import numpy as np
from scipy.integrate import quad
from scipy.optimize import brentq, minimize_scalar
from scipy.special import expit, logit

import numpyro
import jax
import numpyro.distributions as dist
from numpyro.infer import MCMC, NUTS
import jax.numpy as jnp
import jax.scipy.stats as jsp
from jax import random
from jax.scipy.special import gamma

# Enable JAX to use 64-bit float operations
jax.config.update("jax_enable_x64", True)

colors = sns.color_palette()
```

Let's use Python to generate two Beta distributions.

```
# Parameters in the two beta distributions.
F_a, F_b = 1, 1
G_a, G_b = 3, 1.2


@jax.jit
def p(x, a, b):
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x**(a-1) * (1 - x)**(b-1)

# The two density functions.
```

(continues on next page)

(continued from previous page)

```
f = jax.jit(lambda x: p(x, F_a, F_b))
g = jax.jit(lambda x: p(x, G_a, G_b))
```

32.2 Sampling from compound lottery H

We implement two methods to draw samples from our mixture model $\alpha F + (1-\alpha)G$.

We'll generate samples using each of them and verify that they match well.

Here is pseudo code for a direct "method 1" for drawing from our compound lottery:

- Step one:
  - use a uniform random draw to flip an unfair coin that selects distribution F with prob α and G with prob $1-\alpha$
- Step two:
  - draw from either F or G, as determined by the coin flip.
- Step three:
  - put the first two steps in a big loop and do them for each realization of w

Our second method uses a uniform distribution and the following fact that we also described and used in *Some Probability Distributions*:

- If a random variable X has CDF F, then a random variable $F^{-1}(U)$ also has CDF F, where U is a uniform random variable on $[0, 1]$.

In other words, if $X \sim F(x)$ we can generate a random sample from F by drawing a random sample from a uniform distribution on $[0, 1]$ and computing $F^{-1}(U)$.

We'll use this fact in conjunction with the `numpy.searchsorted` command to sample from H directly.

See the `numpy.searchsorted` documentation for details on the `searchsorted` function.

In the Python code below, we'll use both of our methods and confirm that each of them does a good job of sampling from our target mixture distribution.

```
def draw_lottery(key, α, N):
    """Draw from the compound lottery directly."""
    # split the key: one for the coin flip, one for each distribution
    key_uniform, key_beta1, key_beta2 = jax.random.split(key, 3)

    u = jax.random.uniform(key_uniform, shape=(N,))
    f_draws = jax.random.beta(key_beta1, F_a, F_b, shape=(N,))
    g_draws = jax.random.beta(key_beta2, G_a, G_b, shape=(N,))

    # select F with probability α, else G
    return jnp.where(u <= α, f_draws, g_draws)


# JIT-compile with `N` as a static argument
draw_lottery = jax.jit(draw_lottery, static_argnums=(2,))
```

```
def draw_lottery_MC(key, α, N):
    "Draw from the compound lottery using the Monte Carlo trick."

    xs = jnp.linspace(1e-8, 1 - (1e-8), 10000)
    CDF = α * jsp.beta.cdf(xs, F_a, F_b) + (1 - α) * jsp.beta.cdf(xs, G_a, G_b)

    Us = jax.random.uniform(key, shape=(N,))
    draws = xs[jnp.searchsorted(CDF[:-1], Us)]
    return draws
```

```
# verify
N = 100000
α = 0.6

key1, key2 = jax.random.split(jax.random.key(42))
sample1 = draw_lottery(key1, α, N)
sample2 = draw_lottery_MC(key2, α, N)

# plot draws and density function
plt.hist(sample1, 50, density=True, alpha=0.5, label="direct draws")
plt.hist(sample2, 50, density=True, alpha=0.5, label="MC draws")

xs = np.linspace(0, 1, 1000)
plt.plot(xs, α * f(xs) + (1 - α) * g(xs), color="red", label="density", lw=2)

plt.legend()
plt.show()
```

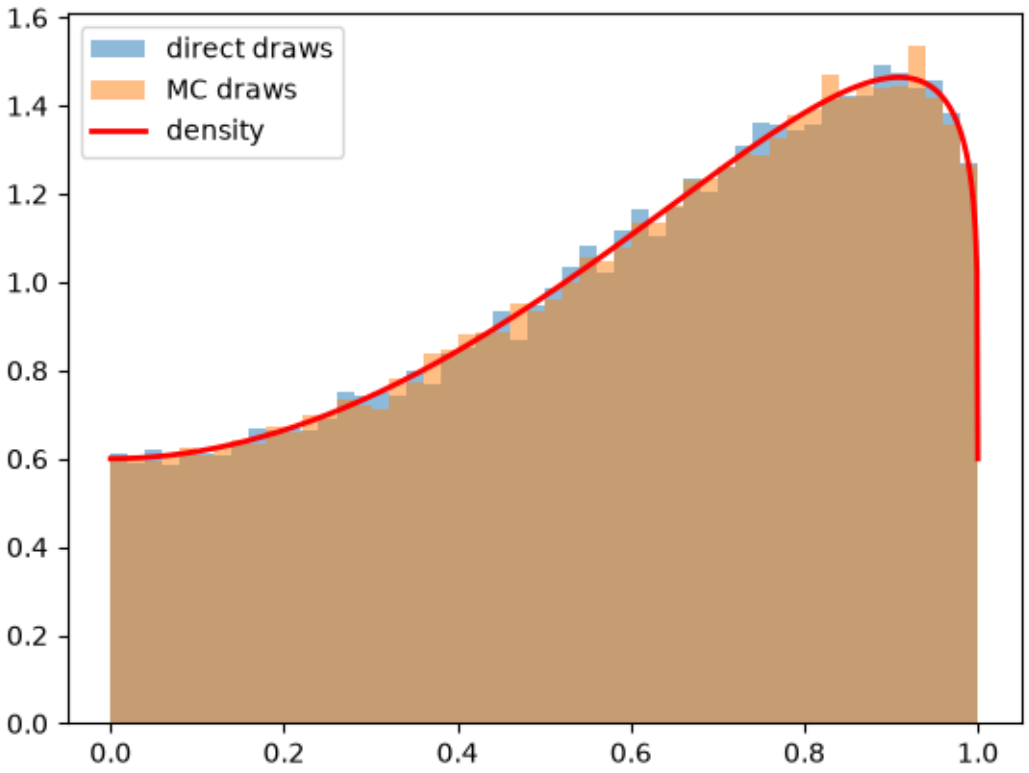


Fig. 32.1: Direct and Monte Carlo draws

32.3 Type 1 agent

We'll now study what our type 1 agent learns.

Remember that our type 1 agent uses the wrong statistical model, thinking that nature mixed between f and g once and for all at time -1.

The type 1 agent thus uses the learning algorithm studied in *Likelihood Ratio Processes and Bayesian Learning*.

We'll briefly review that learning algorithm now.

Let π_t be a Bayesian posterior defined as

$$\pi_t = \mathbb{P}\{q = f \mid w^t\}$$

The likelihood ratio process plays a principal role in the formula that governs the evolution of the posterior probability π_t, an instance of **Bayes' law**.

Bayes' law implies that $\{\pi_t\}$ obeys the recursion

$$\pi_t = \frac{\pi_{t-1}\ell(w_t)}{\pi_{t-1}\ell(w_t) + 1 - \pi_{t-1}} \tag{32.1}$$

with π_{-1} being a Bayesian prior probability that $q = f$, i.e., a personal or subjective belief about q based on our having seen no data.

Here $\ell(w) = f(w)/g(w)$ is the likelihood ratio.

Iterating on (32.1), as in *Likelihood Ratio Processes and Bayesian Learning*, expresses the time t posterior π_t as a function of the prior π_{-1} and the likelihood ratio process $L(w^t) = \prod_{s=0}^{t} \ell(w_s)$:

$$\pi_t = \frac{\pi_{-1} L\left(w^t\right)}{\pi_{-1} L\left(w^t\right) + 1 - \pi_{-1}}. \tag{32.2}$$

Formula (32.2) can be regarded as a one step revision of prior probability π_{-1} after seeing the batch of data $\{w_s\}_{s=0}^{t}$.

In terms of log odds, formula (32.2) is

$$\log \frac{\pi_t}{1-\pi_t} = \log \frac{\pi_{-1}}{1-\pi_{-1}} + \log L(w^t). \tag{32.3}$$

Our code computes π_t from (32.3).

Iterating on (32.1) in floating point arithmetic accumulates rounding errors once π_t gets close to 0 or 1.

32.4 What a type 1 agent learns when mixture H generates data

We now study what happens when the mixture distribution $h(\cdot|\alpha)$ truly generated the data each period.

For all but one value of α, the sequence π_t continues to converge, despite the agent's misspecified model, and the limit is either 0 or 1.

This is true even though in truth nature always mixes between f and g.

After verifying that claim about possible limit points of π_t sequences, we'll drill down and study what fundamental force determines the limiting value of π_t.

Let's set a value of α and then watch how π_t evolves.

```
def simulate_mixed(key, α, T=50, N=500):
    """
    Generate N sets of T observations of the likelihood ratio,
    return as N x T matrix, when the true density is h(·|α)
    """

    w_s = draw_lottery(key, α, N * T).reshape(N, T)
    l_arr = f(w_s) / g(w_s)
    return l_arr


def plot_π_seq(key, α, π1=0.2, π2=0.8, T=200):
    """
    Compute and plot π_t for two priors and the log likelihood ratio process
    when the mixture h(·|α) governs the data.
    """

    l_arr_mixed = simulate_mixed(key, α, T=T, N=1)
    log_L = np.cumsum(np.log(l_arr_mixed[0]))     # log L(w^t), t = 0, ..., T-1

    # plot
    fig, ax1 = plt.subplots()
    for π_init in (π1, π2):
        π_path = expit(logit(π_init) + log_L)
        ax1.plot(
            range(-1, T),
            np.concatenate(([π_init], π_path)),
            label=rf"$\pi_{{-1}}$={π_init}",
            lw=2,
        )

    ax1.plot(
        np.nan, np.nan, "--", color="b",
        label="Log likelihood ratio process", lw=2,
    )
    ax1.set_ylabel(r"$\pi_t$")
    ax1.set_xlabel("$t$")
    ax1.legend()

    ax2 = ax1.twinx()
    ax2.plot(range(T), log_L, "--", color="b", lw=2)
    ax2.set_ylabel(r"$\log(L(w^{t}))$")

    plt.show()
```

```
plot_π_seq(jax.random.key(42), α=0.6)
```

The above graph shows a sample path of the log likelihood ratio process as the blue dashed line, together with sample paths of π_t that start from two distinct initial conditions.

Let's see what happens when we change α.

```
plot_π_seq(jax.random.key(42), α=0.2)
```

Evidently, α is having a big effect on the destination of π_t as $t \to +\infty$.

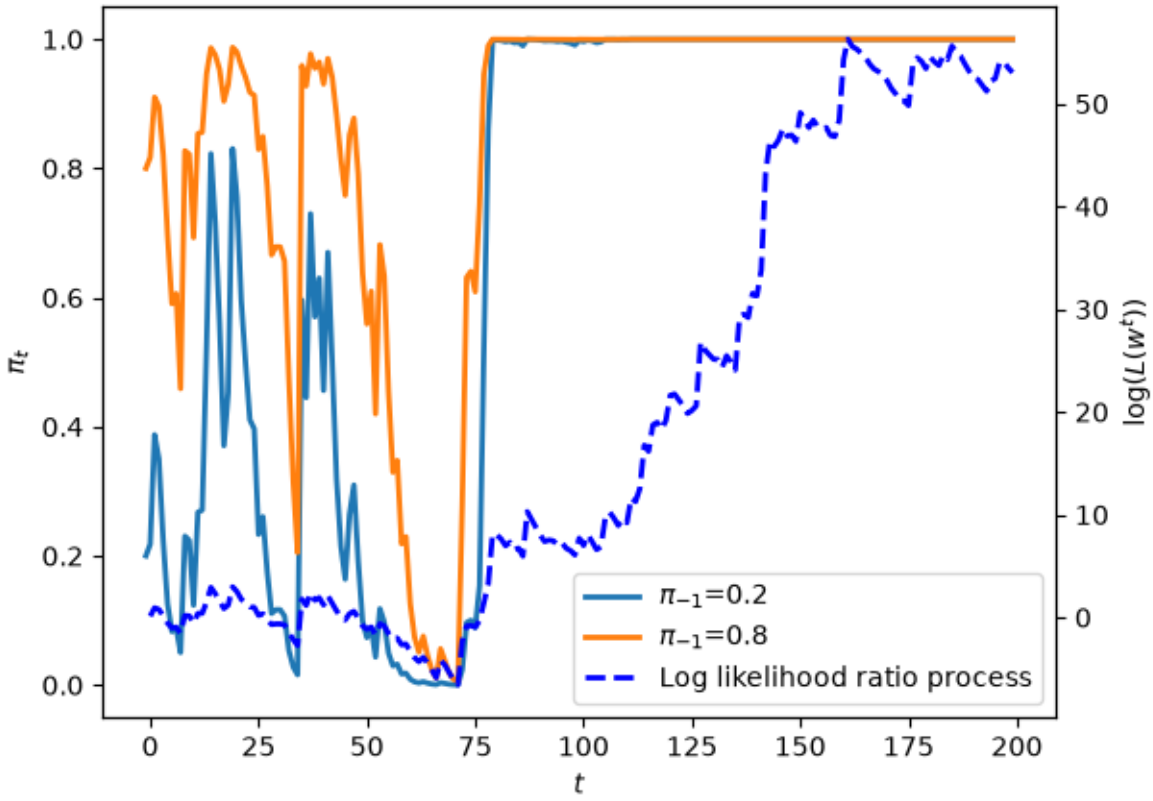


Fig. 32.2: Belief paths, $\alpha = 0.6$

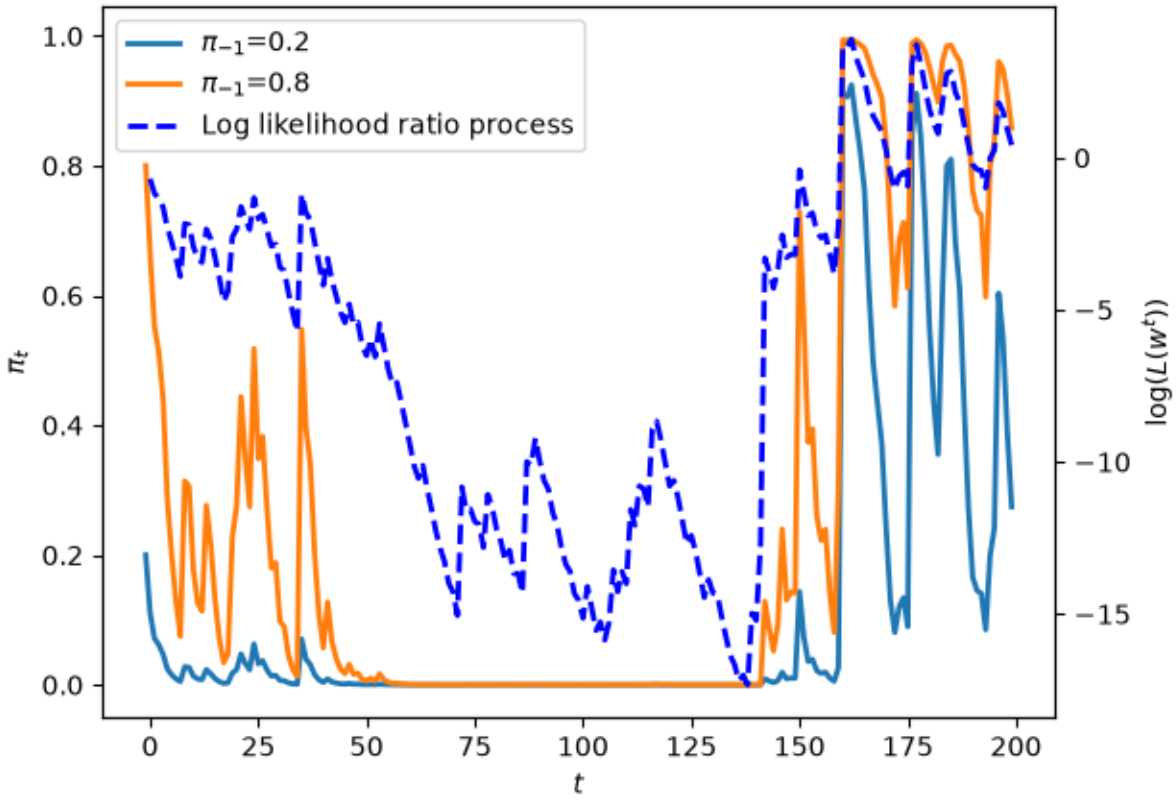


Fig. 32.3: Belief paths, $\alpha = 0.2$

32.5 Kullback-Leibler divergence governs limit of π_t

To understand what determines whether the limit point of π_t is 0 or 1 and how the answer depends on the true value of the mixing probability $\alpha \in (0, 1)$ that generates

$$h(w|\alpha) = \alpha f(w) + (1 - \alpha) g(w)$$

we shall compute the following two Kullback-Leibler divergences

$$KL_g(\alpha) = KL(h, g) = \int \log \left(\frac{h(w|\alpha)}{g(w)} \right) h(w|\alpha) dw$$

and

$$KL_f(\alpha) = KL(h, f) = \int \log \left(\frac{h(w|\alpha)}{f(w)} \right) h(w|\alpha) dw$$

We shall plot both of these functions against α as we use α to vary $h(w|\alpha)$.

Because $\log \ell(w) = \log \frac{h(w|\alpha)}{g(w)} - \log \frac{h(w|\alpha)}{f(w)}$,

$$E_h[\log \ell(w)] = KL_g(\alpha) - KL_f(\alpha),$$

where E_h denotes the expectation under $h(\cdot|\alpha)$; see *A helpful formula* and *A third distribution h* in *Likelihood Ratio Processes*.

By the law of large numbers, $\frac{1}{t+1} \log L(w^t)$ converges to $KL_g(\alpha) - KL_f(\alpha)$ with probability one.

Formula (32.3) then implies that the limit of π_t is determined by

$$\arg \min_{q \in \{f, g\}} KL(h, q).$$

As $t \to +\infty$, π_t goes to one if $KL_f(\alpha) < KL_g(\alpha)$ and goes to zero if $KL_f(\alpha) > KL_g(\alpha)$.

```
@jax.jit
def KL_g(α):
    "Compute the KL divergence KL(h, g)."
    err = 1e-8  # to avoid 0 at end points
    ws = jnp.linspace(err, 1 - err, 10000)
    gs, fs = g(ws), f(ws)
    hs = α * fs + (1 - α) * gs
    return jnp.sum(jnp.log(hs / gs) * hs) / 10000


KL_g_v = jax.vmap(KL_g)


@jax.jit
def KL_f(α):
    "Compute the KL divergence KL(h, f)."
    err = 1e-8  # to avoid 0 at end points
    ws = jnp.linspace(err, 1 - err, 10000)
    gs, fs = g(ws), f(ws)
    hs = α * fs + (1 - α) * gs
    return jnp.sum(jnp.log(hs / fs) * hs) / 10000
```

(continues on next page)

(continued from previous page)

```
KL_f_v = jax.vmap(KL_f)


@jax.jit
def π_lim(key, α, T=5000, π_init=0.4):
    """Compute π_{T-1} from the log odds formula eq:logodds_mix."""
    # Get lottery draws
    l_arr = simulate_mixed(key, α, T, N=1)[0]
    log_odds = jnp.log(π_init / (1 - π_init)) + jnp.sum(jnp.log(l_arr))
    return jax.nn.sigmoid(log_odds)


π_lim_v = jax.vmap(π_lim)
```

Let us first plot the KL divergences $KL_g(\alpha), KL_f(\alpha)$ for each α.

```
α_arr = np.linspace(0, 1, 100)
KL_g_arr = KL_g_v(α_arr)
KL_f_arr = KL_f_v(α_arr)

fig, ax = plt.subplots()

ax.plot(α_arr, KL_g_arr, label="KL(h, g)", lw=2)
ax.plot(α_arr, KL_f_arr, label="KL(h, f)", lw=2)
ax.set_ylabel("KL divergence")
ax.set_xlabel(r"$\alpha$")

ax.legend(loc="upper right")
plt.show()
```

Let's compute the α^* at which the KL divergence between h and g equals that between h and f.

Because $E_h[\log \ell(w)] = \alpha E_f[\log \ell(w)] + (1-\alpha)E_g[\log \ell(w)] = \alpha KL(f,g) - (1-\alpha)KL(g,f)$, the difference $KL_g(\alpha) - KL_f(\alpha)$ is linear in α and

$$\alpha^* = \frac{KL(g,f)}{KL(f,g) + KL(g,f)}.$$

```
def KL(q1, q2):
    "KL(q1, q2) computed with scipy's quad."
    return quad(lambda w: q1(w) * np.log(q1(w) / q2(w)), 0, 1)[0]

KL_fg, KL_gf = KL(f, g), KL(g, f)
α_star = KL_gf / (KL_fg + KL_gf)

h_star = lambda w: α_star * f(w) + (1 - α_star) * g(w)
print(f"α* = {α_star:.4f}")
print(f"KL(h, f) = {KL(h_star, f):.4f}, KL(h, g) = {KL(h_star, g):.4f} at α*")
```

```
α* = 0.3116
KL(h, f) = 0.1435, KL(h, g) = 0.1435 at α*
```

At α^*, f and g both minimize $KL(h,q)$ over $q \in \{f,g\}$.

There $\log L(w^t)$ is a random walk whose increments have mean zero.

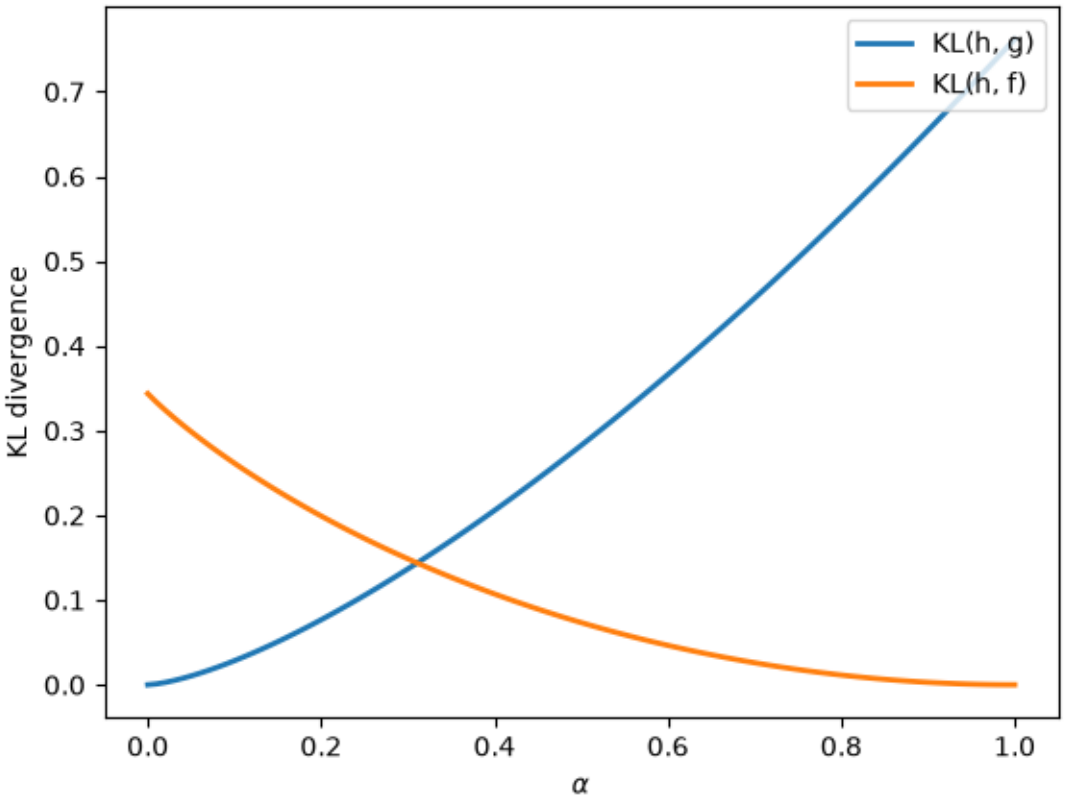


Fig. 32.4: KL divergences against α

Such a random walk is recurrent: with probability one it takes arbitrarily large positive and arbitrarily large negative values infinitely often.

So when $\alpha = \alpha^*$, π_t does not converge; it keeps returning to values near 0 and values near 1.

Exercise Exercise 32.8.2 studies this case.

We can compute and plot π_{T-1} with $T = 5000$ for each α on our grid to verify that the convergence is governed by the KL divergence.

The blue circles show the values of π_{T-1} that simulations discover for different values of α recorded on the x axis.

The dotted vertical line marks α^*.

```
keys = jax.random.split(jax.random.key(42), len(α_arr))
π_lim_arr = π_lim_v(keys, α_arr)

# plot
fig, ax = plt.subplots()

ax.plot(α_arr, KL_g_arr, label="KL(h, g)", lw=2)
ax.plot(α_arr, KL_f_arr, label="KL(h, f)", lw=2)
ax.axvline(α_star, color="k", linestyle=":", lw=1)
ax.set_ylabel("KL divergence")
ax.set_xlabel(r"$\alpha$")

# plot KL
ax2 = ax.twinx()
# plot limit point
ax2.scatter(
    α_arr,
    π_lim_arr,
    facecolors="none",
    edgecolors="tab:blue",
    label=r"$\pi_{T-1}$",
)
ax2.set_ylabel(r"$\pi_{T-1}$")

ax.legend(loc=[0.85, 0.8])
ax2.legend(loc=[0.85, 0.73])
plt.show()

wrong_side = (((α_arr < α_star) & (π_lim_arr > 0.5))
              | ((α_arr > α_star) & (π_lim_arr < 0.5)))
print("values of α at which π_{T-1} lies on the side opposite to its limit:",
      np.round(α_arr[wrong_side], 3))
```

```
values of α at which π_{T-1} lies on the side opposite to its limit: [0.293 0.313]
```

Evidently, our type 1 learner who applies Bayes' law to his misspecified set of statistical models eventually learns an approximating model that is as close as possible to the true model, as measured by its Kullback-Leibler divergence:

- When $\alpha < \alpha^*$, $KL_g < KL_f$, meaning the divergence of g from h is smaller than that of f and so the limit point of π_t is 0.
- When $\alpha > \alpha^*$, $KL_f < KL_g$, meaning the divergence of f from h is smaller than that of g and so the limit point of π_t is 1.

For α close to α^*, the drift $KL_g(\alpha) - KL_f(\alpha)$ of $\log L(w^t)$ is small relative to the standard deviation of $\log \ell(w)$, and 5000 draws are too few to settle where π_t is going.

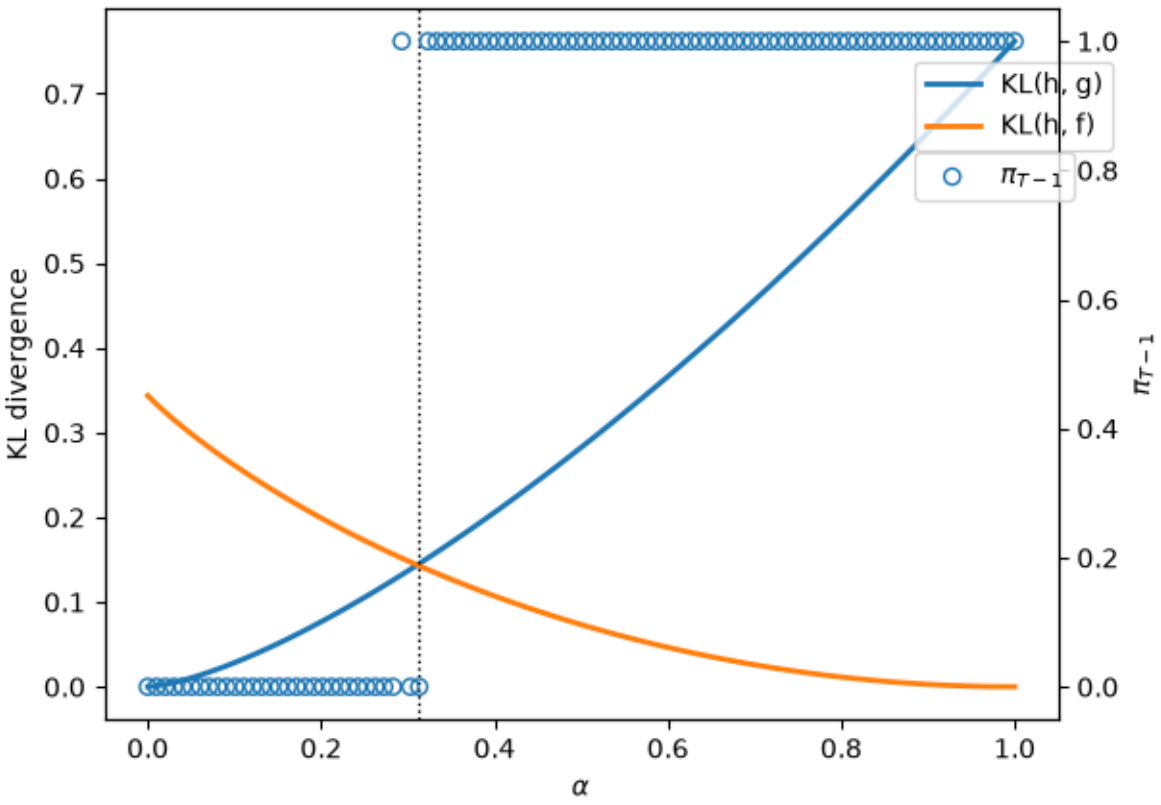


Fig. 32.5: Limit points and KL divergences

That is why the simulated values of π_{T-1} for the values of α printed above, all near α^*, lie on the side opposite to the eventual limit.

32.6 Type 2 agent

We now describe how our type 2 agent formulates his learning problem and what he eventually learns.

Our type 2 agent understands the correct statistical model but does not know α.

We apply Bayes' law to deduce an algorithm for learning α under the assumption that the agent knows that

$$h(w) = h(w|\alpha)$$

but does not know α.

We'll assume that the agent starts out with a prior probability density $\pi_{-1}(\alpha)$ on $\alpha \in (0, 1)$ where the prior has one of the forms that we deployed in *Non-Conjugate Priors*.

We'll fire up `numpyro` and apply it to the present situation.

Bayes' law now takes the form

$$\pi_{t+1}(\alpha) = \frac{h(w_{t+1}|\alpha)\pi_t(\alpha)}{\int h(w_{t+1}|\hat{\alpha})\pi_t(\hat{\alpha})d\hat{\alpha}}$$

where $\pi_t(\alpha)$ is the posterior density of α given w^t.

We'll use numpyro to approximate this equation.

We'll create graphs of the posterior $\pi_t(\alpha)$ as $t \to +\infty$ corresponding to ones presented in *Non-Conjugate Priors*.

We anticipate that a posterior distribution will collapse around the true α as $t \to +\infty$.

Let us try a uniform prior first.

We use the `Mixture` class in numpyro to construct the likelihood function.

```
a = 0.8
sizes = [5, 20, 50, 200, 1000, 25000]

# simulate data with true a
data = draw_lottery(jax.random.key(42), a, max(sizes))


def model(w):
    a = numpyro.sample("a", dist.Uniform(low=0.0, high=1.0))

    y_samp = numpyro.sample(
        "w",
        dist.Mixture(
            dist.Categorical(jnp.array([a, 1 - a])),
            [dist.Beta(F_a, F_b), dist.Beta(G_a, G_b)],
        ),
        obs=w,
    )


def MCMC_run(ws):
    "Compute posterior using MCMC with observed ws"
```

(continues on next page)

(continued from previous page)

```
    kernel = NUTS(model)
    mcmc = MCMC(kernel, num_samples=5000, num_warmup=1000, progress_bar=False)

    mcmc.run(rng_key=random.key(42), w=jnp.array(ws))
    sample = mcmc.get_samples()
    return sample["α"]
```

The following code generates the graph below that displays Bayesian posteriors for α at various history lengths.

```
fig, ax = plt.subplots()

for i in range(len(sizes)):
    sample = MCMC_run(data[: sizes[i]])
    sns.histplot(
        data=sample,
        kde=True,
        stat="density",
        alpha=0.2,
        ax=ax,
        color=colors[i],
        binwidth=0.02,
        linewidth=0.05,
        label=f"t={sizes[i]}",
    )
ax.legend()
ax.set_xlabel(r"$\alpha$")
plt.show()
```

Evidently, the Bayesian posterior narrows in on the true value $\alpha = 0.8$ of the mixing parameter as the length of a history of observations grows.

The posterior standard deviation shrinks in proportion to $1/\sqrt{t}$; Exercise Exercise 32.8.3 computes the constant of proportionality.

32.7 Concluding remarks

Our type 1 agent deploys an incorrect statistical model.

He believes that either f or g generated the w process, but just doesn't know which one.

That is wrong because nature is actually mixing each period with mixing probability α.

Our type 1 agent eventually believes that either f or g generated the w sequence, the outcome being determined by the model, either f or g, whose KL divergence relative to h is smaller.

At the knife-edge value α^*, the two divergences are equal and his beliefs never settle down.

Our type 2 agent has a different statistical model, one that is correctly specified.

He knows the parametric form of the statistical model but not the mixing parameter α.

He knows that he does not know it.

But by using Bayes' law in conjunction with his statistical model and a history of data, he eventually acquires a more and more accurate inference about α.

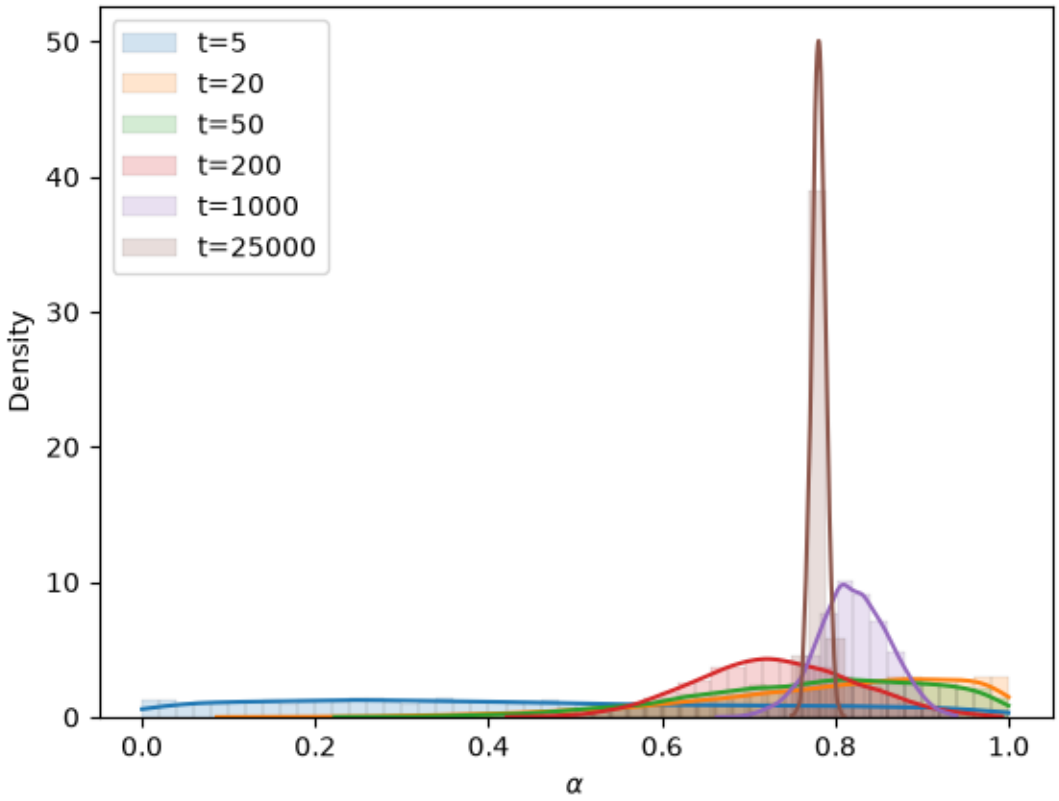


Fig. 32.6: Posterior for α as t grows

This little laboratory exhibits some important general principles that govern outcomes of Bayesian learning of misspecified models.

Thus, the following situation prevails quite generally in empirical work.

A scientist approaches the data with a manifold S of statistical models $s(X|\theta)$, where s is a probability distribution over a random vector X, $\theta \in \Theta$ is a vector of parameters, and Θ indexes the manifold of models.

The scientist with observations that he interprets as realizations x of the random vector X wants to solve an **inverse problem** of somehow *inverting* $s(x|\theta)$ to infer θ from x.

But the scientist's model is misspecified, being only an approximation to an unknown model h that nature uses to generate X.

If the scientist uses Bayes' law or a related likelihood-based method to infer θ from IID observations, then for large sample sizes his inferences concentrate on values of θ that minimize

$$KL(h \,\|\, s_\theta) = \int h(x) \log \frac{h(x)}{s(x|\theta)} dx$$

over $\theta \in \Theta$, where $s_\theta = s(\cdot|\theta)$.

Berk [1966] established this for Bayesian posterior distributions, and White [1982] established it for maximum likelihood estimators.

When the minimizer θ^* is unique, the posterior concentrates on θ^* and the maximum likelihood estimator converges to θ^*.

When several values of θ attain the minimum, the posterior concentrates on the set of minimizers but need not converge.

Our type 1 agent at $\alpha = \alpha^*$ is an instance.

The section on *KL divergence and maximum-likelihood estimation* in *Statistical Divergence Measures* shows that maximum likelihood minimizes the cross-entropy of the model relative to the empirical distribution of the data, the sample analogue of the KL divergence from the data-generating distribution to the model.

Exercise Exercise 32.8.4 studies a type 2 agent whose model is misspecified.

32.8 Exercises

Exercise 32.8.1

In *Likelihood Ratio Processes and Bayesian Learning*, we studied the consequences of applying a likelihood ratio and Bayes' law to a misspecified statistical model, namely, our type 1 agent's model, when the true data generating process is a mixture.

In this lecture, we used `numpyro` to compute posteriors of our type 2 agent, whose model is correctly specified.

This exercise asks you to replicate the type 2 agent's posterior with a grid approximation and compare it with the NUTS posterior.

Here is the algorithm.

First we specify a prior distribution for α given by $\alpha \sim \text{Beta}(a_0, b_0)$ with expectation $\mathbb{E}[\alpha] = \frac{a_0}{a_0+b_0}$.

The likelihood for a single observation w_t is $h(w_t|\alpha) = \alpha f(w_t) + (1-\alpha) g(w_t)$.

For a history $w^t = (w_0, \ldots, w_t)$, the likelihood is $\prod_{s=0}^{t} h(w_s|\alpha)$.

The posterior density is $\pi_t(\alpha) \propto \prod_{s=0}^{t} h(w_s|\alpha) \pi_{-1}(\alpha)$.

Recursively, $\pi_t(\alpha) \propto h(w_t|\alpha)\pi_{t-1}(\alpha)$.

Without a conjugate prior, we can approximate the posterior by discretizing α into a grid.

1. Implement this algorithm in Python.
2. With a uniform prior ($a_0 = b_0 = 1$), compute the posterior of α from the first $5, 20, 50, 200$, and 1000 observations in `data` and compare it with the NUTS posterior.
3. Let the true mixing probability be $\alpha = 0.5$. For a history of 200 draws, compute the paths of posterior means of α for three priors with means 0.25, 0.5, and 0.75. Then compute paths of posterior means with a uniform prior for 20 histories of $10{,}000$ draws.

Solution

Here is one solution.

We put a grid on α and update the log posterior on the grid one observation at a time.

```
α_grid = jnp.linspace(1e-3, 1 - 1e-3, 2000)
Δα = float(α_grid[1] - α_grid[0])


@jax.jit
def learn_α_bayesian(observations, a0, b0):
    """
    Sequential Bayesian learning of the mixing probability α on α_grid.
    Returns the path of posterior means, starting from the prior mean,
    and the final log posterior probabilities on the grid.
    """
    w = jnp.asarray(observations)

    # Log prior
    log_prior = (a0 - 1) * jnp.log(α_grid) + (b0 - 1) * jnp.log1p(-α_grid)
    log_prior = log_prior - jax.nn.logsumexp(log_prior)

    def scan_fn(log_post, wt):
        # h(w_t | α) = α f(w_t) + (1 - α) g(w_t)
        like = α_grid * f(wt) + (1 - α_grid) * g(wt)
        log_post = log_post + jnp.log(like)

        # Normalize using log-sum-exp trick
        log_post = log_post - jax.nn.logsumexp(log_post)

        # Compute posterior mean
        μ = α_grid @ jnp.exp(log_post)

        return log_post, μ

    # Scan over observations
    log_post, μ_path = jax.lax.scan(scan_fn, log_prior, w)

    # Prepend the prior mean
    μ_prior = a0 / (a0 + b0)
    return jnp.concatenate([jnp.array([μ_prior]), μ_path]), log_post
```

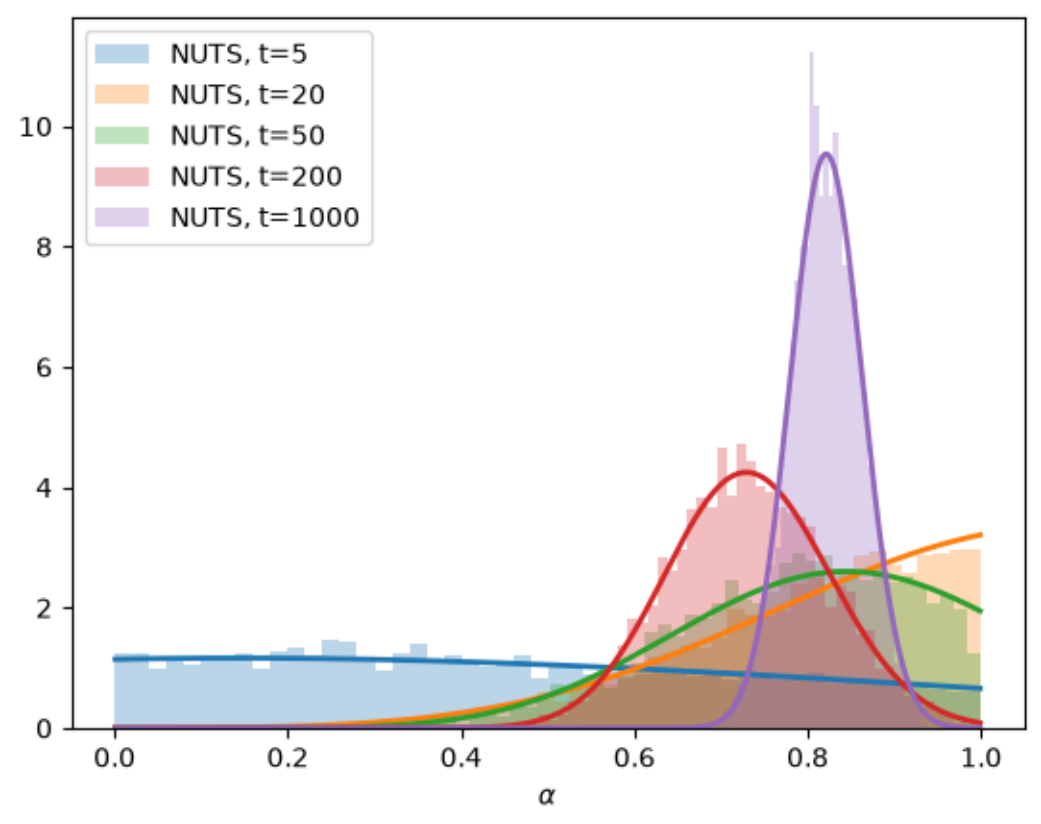
Now we compare the grid posterior densities with the NUTS draws for the data used in the lecture.

```
fig, ax = plt.subplots()

for i, n in enumerate(sizes[:-1]):
    _, log_post = learn_α_bayesian(data[:n], 1.0, 1.0)
    density = np.exp(log_post) / Δα
    sample = MCMC_run(data[:n])
    ax.hist(np.asarray(sample), bins=50, density=True, alpha=0.3,
            color=colors[i], label=f"NUTS, t={n}")
    ax.plot(α_grid, density, color=colors[i], lw=2)
    print(f"t = {n:4d}: grid mean = {float(α_grid @ jnp.exp(log_post)):.4f}, "
          f"NUTS mean = {float(sample.mean()):.4f}")

ax.set_xlabel(r"$\alpha$")
ax.legend()
plt.show()
```

```
t =    5: grid mean = 0.4568, NUTS mean = 0.4487
t =   20: grid mean = 0.7885, NUTS mean = 0.7897
t =   50: grid mean = 0.7753, NUTS mean = 0.7737
t =  200: grid mean = 0.7318, NUTS mean = 0.7302
t = 1000: grid mean = 0.8216, NUTS mean = 0.8206
```



The lines are the grid posteriors and the histograms are the NUTS draws.

They agree up to Monte Carlo error in the NUTS draws.

Next we study the paths of posterior means for three priors when $\alpha = 0.5$.

```
α_true = 0.5
T_mix = 200

# Three different priors with means 0.25, 0.5, 0.75
prior_params = [(1, 3), (1, 1), (3, 1)]
prior_means = [a0 / (a0 + b0) for a0, b0 in prior_params]

w_mix = draw_lottery(jax.random.key(42), α_true, T_mix)

α_posterior_means = [learn_α_bayesian(w_mix, float(a0), float(b0))[0]
                     for a0, b0 in prior_params]
```

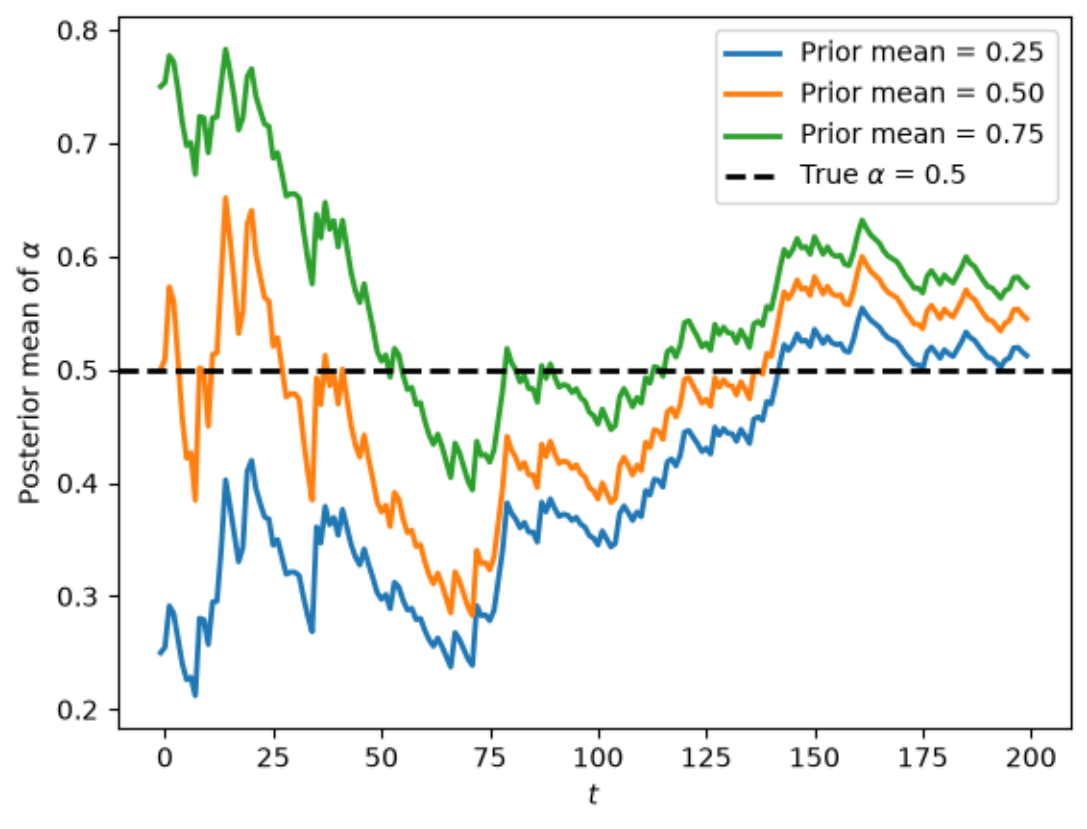
Let's visualize how the posterior mean of α evolves over time, starting from three different prior beliefs.

```
fig, ax = plt.subplots()

for i, (α_means, mean0) in enumerate(zip(α_posterior_means, prior_means)):
    ax.plot(
        range(-1, T_mix),
        α_means,
        label=rf"Prior mean = ${mean0:.2f}$",
        color=colors[i],
        linewidth=2,
    )

ax.axhline(
    y=α_true,
    color="black",
    linestyle="--",
    label=rf"True $\alpha$ = {α_true}",
    linewidth=2,
)
ax.set_xlabel("$t$")
ax.set_ylabel(r"Posterior mean of $\alpha$")
ax.legend()
plt.show()

for α_means, mean0 in zip(α_posterior_means, prior_means):
    print(f"prior mean {mean0:.2f}: posterior mean at t = {T_mix - 1} is {float(α_
↪means[-1]):.3f}")
```




```
prior mean 0.25: posterior mean at t = 199 is 0.512
prior mean 0.50: posterior mean at t = 199 is 0.545
prior mean 0.75: posterior mean at t = 199 is 0.573
```

All three posterior means move toward the true value $\alpha = 0.5$, and their differences shrink as t grows.

Next, let's look at multiple simulations with a longer time horizon, all starting from a uniform prior.

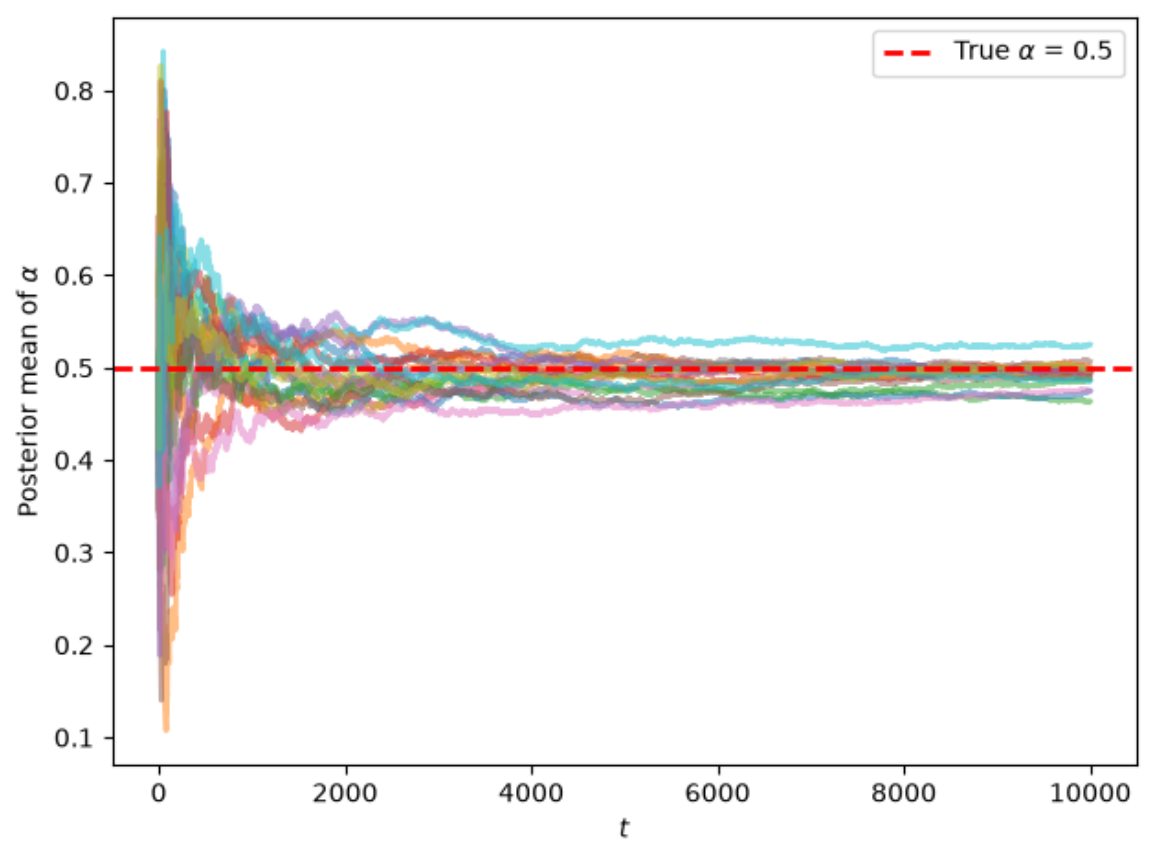
```
n_paths = 20
T_long = 10_000

keys = jax.random.split(jax.random.key(42), n_paths)

fig, ax = plt.subplots()

for j in range(n_paths):
    w_path = draw_lottery(keys[j], α_true, T_long)
    α_means, _ = learn_α_bayesian(w_path, 1.0, 1.0)  # Uniform prior
    ax.plot(range(-1, T_long), α_means, alpha=0.5, lw=2)

ax.axhline(
    y=α_true,
    color="red",
    linestyle="--",
    label=rf"True $\alpha$ = {α_true}",
    linewidth=2,
)
ax.set_ylabel(r"Posterior mean of $\alpha$")
ax.set_xlabel("$t$")
ax.legend()
plt.tight_layout()
plt.show()
```



We can see that the posterior mean of α converges to the true value $\alpha = 0.5$.

Exercise 32.8.2

This exercise studies the type 1 agent at the knife-edge value α^*.

1. Compute α^* by finding the root of $KL_f(\alpha) - KL_g(\alpha)$ and compare it with the formula $\alpha^* = KL(g,f)/(KL(f,g)+KL(g,f))$.
2. For $\alpha \in \{\alpha^* - 0.05, \alpha^*, \alpha^* + 0.05\}$, simulate 20 paths of $\log L(w^t)$ for $t = 0, \ldots, 19{,}999$ and plot them.
3. For each α, compute π_t from (32.3) with $\pi_{-1} = 0.5$, and count how often each path of π_t crosses $1/2$ during the second half of the sample.

Solution

We compute KL_f and KL_g with `quad` and find their crossing with `brentq`.

```
def KL_h(a, q):
    "KL(h, q) for h = h(·|a), computed with quad."
    h = lambda w: a * f(w) + (1 - a) * g(w)
    return KL(h, q)

a_root = brentq(lambda a: KL_h(a, f) - KL_h(a, g), 0.05, 0.95)
print(f"root of KL_f - KL_g: {a_root:.5f}")
```

```
print(f"formula:             {α_star:.5f}")
```

```
root of KL_f - KL_g: 0.31156
formula:             0.31156
```

Next we simulate paths of $\log L(w^t)$.

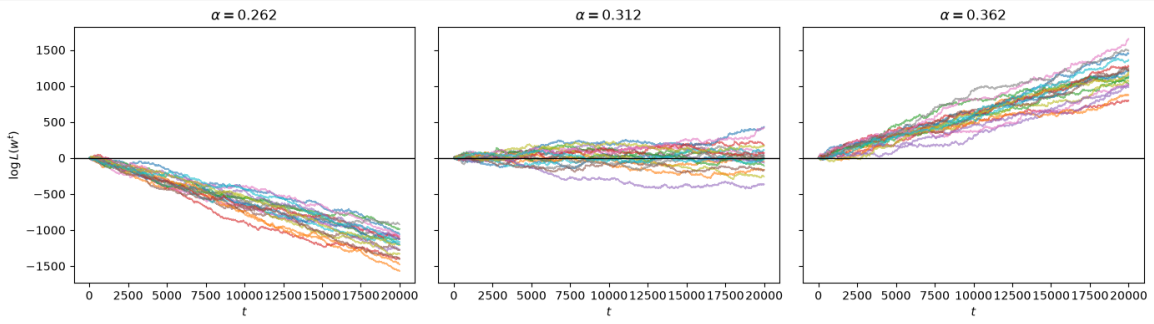
```
T_k, N_k = 20_000, 20
α_values = [α_star - 0.05, α_star, α_star + 0.05]

fig, axes = plt.subplots(1, 3, figsize=(14, 4), sharey=True)
for ax, α_k, key in zip(axes, α_values, jax.random.split(jax.random.key(1), 3)):
    log_L = np.cumsum(np.log(np.asarray(simulate_mixed(key, α_k, T=T_k, N=N_k))),␣
↪axis=1)
    ax.plot(log_L.T, lw=1, alpha=0.6)
    ax.axhline(0, color='k', lw=1)
    ax.set_title(rf"$\alpha = {α_k:.3f}$")
    ax.set_xlabel("$t$")

    π_paths = expit(log_L)            # π_{-1} = 0.5, so the prior log odds are␣
↪zero
    late = π_paths[:, T_k // 2:] > 0.5
    crossings = np.sum(late[:, 1:] != late[:, :-1], axis=1)
    print(f"α = {α_k:.3f}: crossings of 1/2 in second half, by path: {crossings}")

axes[0].set_ylabel(r"$\log L(w^t)$")
plt.tight_layout()
plt.show()
```

```
α = 0.262: crossings of 1/2 in second half, by path: [0 0 0 0 0 0 0 0 0 0 0 0 0 0␣
↪0 0 0 0 0 0 0]
α = 0.312: crossings of 1/2 in second half, by path: [58 14 11  0 63 44  0 12␣
↪23 39  0  0 41  0  0  0  0 93  0 42]
α = 0.362: crossings of 1/2 in second half, by path: [0 0 0 0 0 0 0 0 0 0 0 0 0 0␣
↪0 0 0 0 0 0 0]
```



At $\alpha^* \pm 0.05$, every path of $\log L(w^t)$ drifts away from zero and π_t stops crossing $1/2$.

At α^*, the paths wander without drift, and many paths of π_t still cross $1/2$ late in the sample.

A path that does not cross $1/2$ during the second half is not evidence of convergence.

A driftless random walk spends long stretches on one side of zero, and with probability one it returns.

Exercise 32.8.3

The Bernstein–von Mises theorem says that, for a correctly specified model, the posterior distribution of α after t observations is approximately normal with mean near the true α and standard deviation $1/\sqrt{tI(\alpha)}$, where

$$I(\alpha) = \int \frac{(f(w) - g(w))^2}{h(w|\alpha)} dw$$

is the Fisher information of one observation.

1. Derive $I(\alpha)$ from the log likelihood $\log h(w|\alpha)$.
2. For the data in `data` ($\alpha = 0.8$), compute the grid posterior standard deviation after 200, 1000, and 5000 observations with a uniform prior, and compare it with $1/\sqrt{tI(0.8)}$.

Solution

The score of one observation is $\frac{\partial}{\partial \alpha} \log h(w|\alpha) = \frac{f(w) - g(w)}{h(w|\alpha)}$.

Its variance under $h(\cdot|\alpha)$ is the Fisher information

$$I(\alpha) = \int \frac{(f(w) - g(w))^2}{h(w|\alpha)^2} h(w|\alpha) dw = \int \frac{(f(w) - g(w))^2}{h(w|\alpha)} dw,$$

because the score has mean $\int (f - g) dw = 0$.

```
def fisher_info(α):
    return quad(lambda w: (f(w) - g(w))**2 / (α * f(w) + (1 - α) * g(w)), 0, 1)[0]

I_08 = fisher_info(0.8)
print(f"I(0.8) = {I_08:.4f}")

for n in [200, 1000, 5000]:
    _, log_post = learn_α_bayesian(data[:n], 1.0, 1.0)
    post = jnp.exp(log_post)
    mean = α_grid @ post
    sd = jnp.sqrt(((α_grid - mean)**2) @ post)
    print(f"t = {n:4d}: posterior mean = {float(mean):.3f}, posterior sd =
↪{float(sd):.4f}, "
          f"1/sqrt(t I) = {1 / np.sqrt(n * I_08):.4f}")
```

```
I(0.8) = 0.5835
t =  200: posterior mean = 0.732, posterior sd = 0.0925, 1/sqrt(t I) = 0.0926
t = 1000: posterior mean = 0.822, posterior sd = 0.0418, 1/sqrt(t I) = 0.0414
t = 5000: posterior mean = 0.787, posterior sd = 0.0186, 1/sqrt(t I) = 0.0185
```

The posterior standard deviations are close to $1/\sqrt{tI(0.8)}$.

Exercise 32.8.4

Suppose that nature draws IID from $h(w|0.5) = 0.5f(w) + 0.5g(w)$ with $g = \text{Beta}(3, 1.2)$, as in the lecture.

A type 2 agent knows f but wrongly believes that the second component is $\tilde{g} = \text{Beta}(3, 2)$.

His model is $s(w|\alpha) = \alpha f(w) + (1 - \alpha)\tilde{g}(w)$, and he puts a uniform prior on α.

1. Compute $\tilde{\alpha} = \arg\min_\alpha KL(h, s_\alpha)$.

2. Simulate 5000 draws from $h(\cdot|0.5)$ and compute the agent's grid posterior after 200, 1000, and 5000 observations. Where does it concentrate?

Solution

```
g_tilde = jax.jit(lambda x: p(x, 3, 2))
α_true = 0.5
h_true = lambda w: α_true * f(w) + (1 - α_true) * g(w)

def KL_misspecified(α):
    s = lambda w: α * f(w) + (1 - α) * g_tilde(w)
    return KL(h_true, s)

res = minimize_scalar(KL_misspecified, bounds=(1e-3, 1 - 1e-3), method='bounded')
α_tilde = res.x
print(f"pseudo-true value α~ = {α_tilde:.3f}, min KL = {res.fun:.4f}")
```

```
pseudo-true value α~ = 0.616, min KL = 0.0460
```

The agent's grid posterior uses $\tilde{g}$ in place of g.

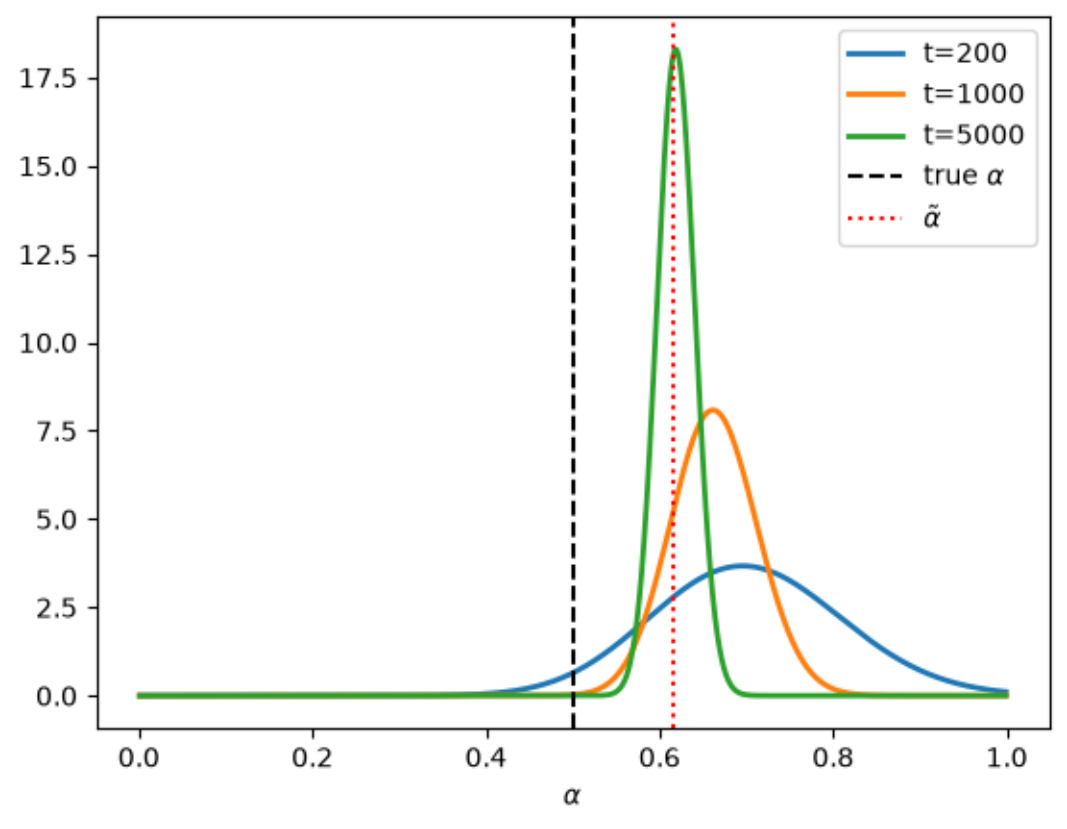
```
@jax.jit
def log_post_misspecified(observations):
    "Log posterior on α_grid, uniform prior, for the model α f + (1 - α) g~."
    w = jnp.asarray(observations)
    log_like = jnp.log(α_grid[:, None] * f(w) + (1 - α_grid[:, None]) * g_tilde(w))
    log_post = jnp.sum(log_like, axis=1)
    return log_post - jax.nn.logsumexp(log_post)

w_obs = draw_lottery(jax.random.key(7), α_true, 5000)

fig, ax = plt.subplots()
for i, n in enumerate([200, 1000, 5000]):
    post = jnp.exp(log_post_misspecified(w_obs[:n]))
    mean = float(α_grid @ post)
    sd = float(jnp.sqrt(((α_grid - mean)**2) @ post))
    ax.plot(α_grid, post / Δα, color=colors[i], lw=2, label=f"t={n}")
    print(f"t = {n:4d}: posterior mean = {mean:.3f}, posterior sd = {sd:.4f}")

ax.axvline(α_true, color="k", linestyle="--", label=r"true $\alpha$")
ax.axvline(α_tilde, color="r", linestyle=":", label=r"$\tilde \alpha$")
ax.set_xlabel(r"$\alpha$")
ax.legend()
plt.show()
```

```
t =  200: posterior mean = 0.701, posterior sd = 0.1066
t = 1000: posterior mean = 0.663, posterior sd = 0.0493
t = 5000: posterior mean = 0.619, posterior sd = 0.0218
```



The posterior concentrates at the pseudo-true value $\tilde{\alpha}$, not at the true $\alpha = 0.5$.

As Berk [1966] showed, a Bayesian with a misspecified model learns the member of his family of models that is closest to the truth in Kullback–Leibler divergence.

CHAPTER

THIRTYTHREE

BAYESIAN VERSUS FREQUENTIST DECISION RULES

Contents

- *Bayesian versus Frequentist Decision Rules*
  - *Overview*
  - *Setup*
  - *Frequentist Decision Rule*
  - *Bayesian Decision Rule*
  - *Was the Navy Captain's Hunch Correct?*
  - *More Details*
  - *Distribution of Bayesian Decision Rule's Time to Decide*
  - *Probability of Making Correct Decision*
  - *Distribution of Likelihood Ratios at Neyman-Pearson's t*
  - *Exercises*


```
import matplotlib.pyplot as plt
import numpy as np
from numba import jit, prange, float64, int64
from numba.experimental import jitclass
from math import gamma
from scipy.integrate import quad
```

33.1 Overview

This lecture follows up on ideas presented in the following lectures:

- *A Problem that Stumped Milton Friedman*
- *A Bayesian Formulation of Friedman and Wald's Problem*
- *Exchangeability and Bayesian Updating*
- *Likelihood Ratio Processes*
- *Likelihood Ratio Processes and Bayesian Learning*

A Problem that Stumped Milton Friedman described a problem that a Navy Captain presented to Milton Friedman during World War II.

The Navy had told the Captain to use a decision rule for comparing two designs of an ordnance device by firing test rounds.

In particular, the Navy had ordered the Captain to use an instance of a **frequentist decision rule**.

The Captain doubted that that rule was a good one.

Milton Friedman recognized the Captain's conjecture as posing a challenging statistical problem that he and other members of the US Government's Statistical Research Group at Columbia University proceeded to try to solve.

A member of the group, the great mathematician and economist Abraham Wald, soon solved the problem.

A good way to formulate the problem is to use some ideas from Bayesian statistics that we describe in *Exchangeability and Bayesian Updating* and in the section on *likelihood ratio processes and Bayes' law* of *Likelihood Ratio Processes and Bayesian Learning*, which describes the link between Bayesian updating and likelihood ratio processes.

The present lecture uses Python to generate simulations that evaluate expected losses under the Neyman-Pearson **frequentist** procedure that the Navy captain questioned and the **Bayesian** decision rule described in *A Bayesian Formulation of Friedman and Wald's Problem*.

We choose the frequentist rule's sample size and cutoff to minimize expected loss under the prior probability and losses that the Bayesian uses.

The frequentist rule then becomes a Bayes rule whose sample size is fixed in advance.

The Bayesian rule chooses its sample size sequentially.

A rule with a fixed sample size is a special sequential rule, so the best sequential rule has an expected loss no larger than the best fixed-sample rule.

Wald and Wolfowitz [1948] proved that, among all tests with given probabilities of type I and type II errors, Wald's sequential probability ratio test minimizes the expected number of draws under both hypotheses.

Arrow *et al.* [1949] showed how to compute Bayes solutions of sequential decision problems by backward induction.

Our simulations illustrate these theorems and measure how much the Navy Captain's hunch was worth.

33.2 Setup

To formalize the problem that had confronted the Navy Captain, we consider a setting with the following parts.

- Each period a decision maker draws a non-negative random variable Z. He knows that two probability distributions are possible, f_0 and f_1, and that which ever distribution it is remains fixed over time. The decision maker believes that before the beginning of time, nature once and for all had selected either f_0 or f_1 and that the probability that it selected f_0 is probability π^*.
- The decision maker observes a sample $\{z_i\}_{i=1}^t$ from the distribution chosen by nature.

The decision maker wants to decide which distribution actually governs Z.

He is worried about two types of errors and the losses that they will impose on him.

- a loss $\bar{L}_1$ from a **type I error** that occurs if he decides that $f = f_1$ when actually $f = f_0$
- a loss $\bar{L}_0$ from a **type II error** that occurs if he decides that $f = f_0$ when actually $f = f_1$

The decision maker pays a cost c for drawing another z.

Note

This lecture lets π denote the probability of f_0 and defines the likelihood ratio as $L = f_0/f_1$.

Under this convention, high values of π and of L favor f_0, and the Bayesian decision maker accepts f_0 when $\pi \geq A$.

A Bayesian Formulation of Friedman and Wald's Problem lets π denote the probability of f_1; its decision maker accepts f_1 when $\pi \geq A$.

A Problem that Stumped Milton Friedman defines the likelihood ratio as f_1/f_0.

We rewrite the code that we borrow from *A Bayesian Formulation of Friedman and Wald's Problem* for our convention: π and $1-\pi$ trade places in the Bayes update κ, in the expected losses from stopping, and in the Bellman operator $\mathbb{Q}$.

We mainly borrow parameters from the quantecon lecture *A Bayesian Formulation of Friedman and Wald's Problem* except that we increase both $\bar{L}_0$ and $\bar{L}_1$ from 25 to 100 to encourage the Bayesian decision rule to take more draws before deciding.

We set the cost c of taking one more draw at 1.25.

We set the probability distributions f_0 and f_1 to be beta distributions with $a_0 = b_0 = 1$, $a_1 = 3$, and $b_1 = 1.2$, respectively.

Below is some Python code that sets up these objects.

```
@jit
def p(x, a, b):
    "Beta distribution."

    r = gamma(a + b) / (gamma(a) * gamma(b))

    return r * x**(a-1) * (1 - x)**(b-1)
```

We start with defining a `jitclass` that stores parameters and functions we need to solve problems for both the Bayesian and frequentist Navy Captains.

```
wf_data = [
    ('c', float64),           # cost of drawing another z
    ('a0', float64),          # parameters of beta distribution
    ('b0', float64),
    ('a1', float64),
    ('b1', float64),
    ('L0', float64),          # cost of selecting f0 when f1 is true
    ('L1', float64),          # cost of selecting f1 when f0 is true
    ('π_grid', float64[:]),   # grid of beliefs π
    ('π_grid_size', int64),
    ('mc_size', int64),       # size of Monte Carlo simulation
    ('z0', float64[:]),       # sequence of random values
    ('z1', float64[:])        # sequence of random values
]
```

```
@jitclass(wf_data)
class WaldFriedman:

    def __init__(self,
                 c=1.25,
                 a0=1,
                 b0=1,
```

(continues on next page)

(continued from previous page)

```
                a1=3,
                b1=1.2,
                L0=100,
                L1=100,
                π_grid_size=200,
                mc_size=1000):

        self.c, self.π_grid_size = c, π_grid_size
        self.a0, self.b0, self.a1, self.b1 = a0, b0, a1, b1
        self.L0, self.L1 = L0, L1
        self.π_grid = np.linspace(0, 1, π_grid_size)
        self.mc_size = mc_size

        self.z0 = np.random.beta(a0, b0, mc_size)
        self.z1 = np.random.beta(a1, b1, mc_size)

    def f0(self, x):

        return p(x, self.a0, self.b0)

    def f1(self, x):

        return p(x, self.a1, self.b1)

    def κ(self, z, π):
        """
        Updates π using Bayes' rule and the current observation z
        """

        a0, b0, a1, b1 = self.a0, self.b0, self.a1, self.b1

        π_f0, π_f1 = π * p(z, a0, b0), (1 - π) * p(z, a1, b1)
        π_new = π_f0 / (π_f0 + π_f1)

        return π_new
```

The constructor of `WaldFriedman` draws `mc_size` Monte Carlo values of z from each density with numba's random number generator.

Calling `np.random.seed` from Python does not seed numba's generator.

We seed it inside a jitted function.

```
@jit
def seed_numba(seed):
    "Seed the random number generator used by numba-compiled code."
    np.random.seed(seed)

seed_numba(1234)
wf = WaldFriedman()

grid = np.linspace(0, 1, 50)

plt.figure()

plt.title("Two Distributions")
plt.plot(grid, wf.f0(grid), lw=2, label="$f_0$")
```

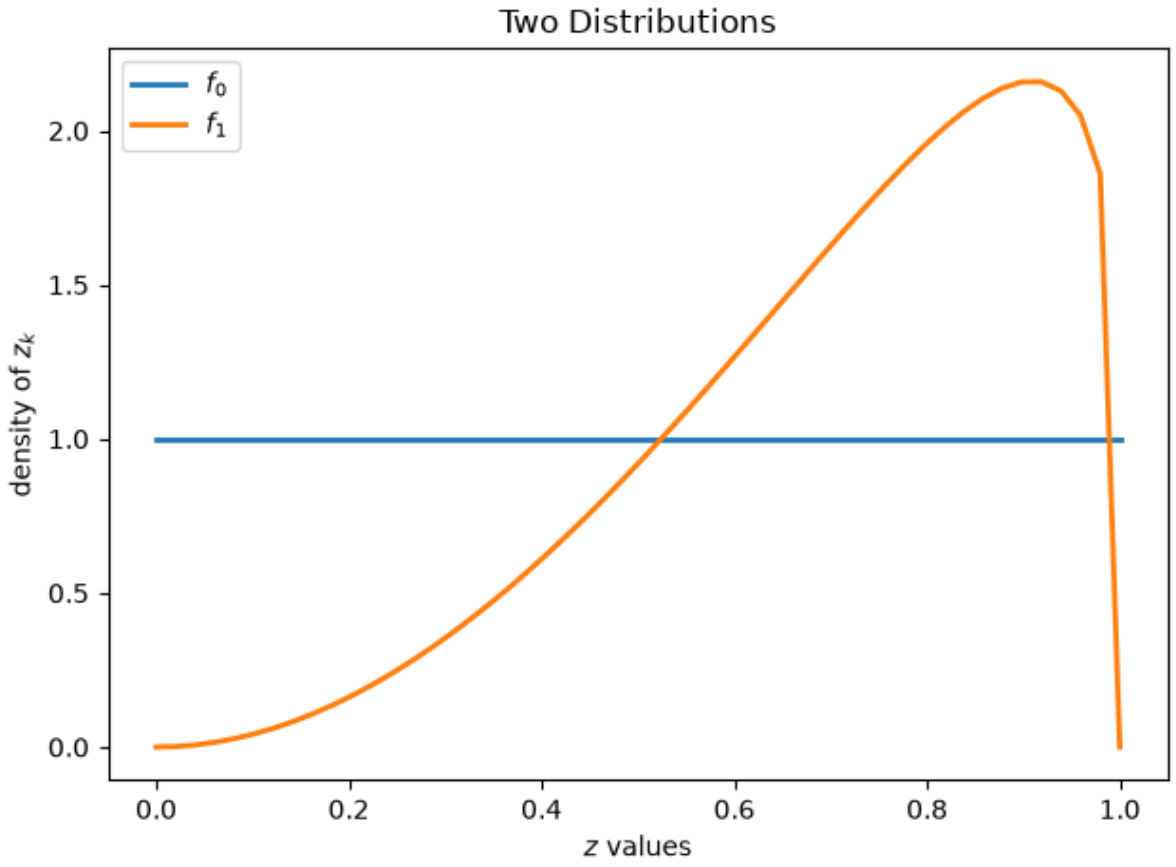
(continues on next page)

(continued from previous page)

```
plt.plot(grid, wf.f1(grid), lw=2, label="$f_1$")

plt.legend()
plt.xlabel("$z$ values")
plt.ylabel("density of $z_k$")

plt.tight_layout()
plt.show()
```



Above, we plot the two possible probability densities f_0 and f_1

33.3 Frequentist Decision Rule

The Navy told the Captain to use a Neyman-Pearson likelihood ratio decision rule.

That decision rule is characterized by

- a sample size t, and
- a cutoff value d of a likelihood ratio

Let $L\left(z^t\right) = \prod_{i=1}^t \frac{f_0(z_i)}{f_1(z_i)}$ be the likelihood ratio associated with observing the sequence $\{z_i\}_{i=1}^t$.

The decision rule associated with a sample size t is:

- decide that f_0 is the distribution if the likelihood ratio is greater than d

- decide that f_1 is the distribution if the likelihood ratio is less than d

For our purposes here, we want to compute an expected loss from using this rule, where we borrow loss parameters $\bar{L}_0$ and $\bar{L}_1$ from *A Bayesian Formulation of Friedman and Wald's Problem*.

Let null and alternative hypotheses be

- null: H_0: $f = f_0$,
- alternative H_1: $f = f_1$.

Given sample size t and cutoff d, under the model described above, the mathematical expectation of total loss is

$$\bar{V}_{fre}(t,d) = ct + \pi^* PFA \times \bar{L}_1 + (1-\pi^*)(1-PD) \times \bar{L}_0 \tag{33.1}$$

$$\begin{aligned} \text{where} \quad PFA &= \Pr\{L(z^t) < d \mid q = f_0\} \\ PD &= \Pr\{L(z^t) < d \mid q = f_1\} \end{aligned}$$

Here

- PFA denotes the probability of a **false alarm**, i.e., rejecting H_0 when it is true
- PD denotes the probability of **detection**, i.e., rejecting H_0 when H_1 is true (the power of the test); $1 - PD$ is the probability of a type II error

For a given sample size t, the pairs (PFA, PD) lie on a **receiver operating characteristic curve**.

- by choosing d, we select a particular pair (PFA, PD) along the curve for a given t

To see some receiver operating characteristic curves, please see this lecture *Likelihood Ratio Processes*.

To solve for $\bar{V}_{fre}(t,d)$ numerically, we first simulate sequences of z when either f_0 or f_1 generates data.

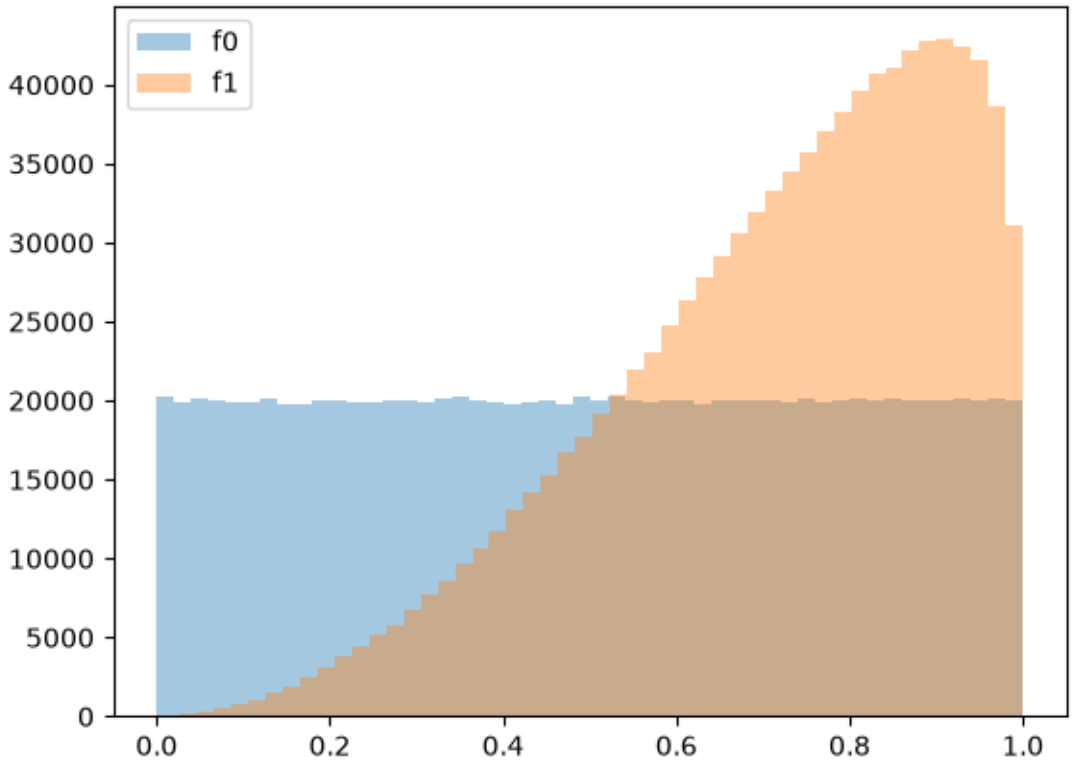
Let's plot empirical distributions, i.e., histograms, associated with f_0 and f_1.

```
N = 10000
T = 100
```

```
rng = np.random.default_rng(1234)

z0_arr = rng.beta(wf.a0, wf.b0, (N, T))
z1_arr = rng.beta(wf.a1, wf.b1, (N, T))
```

```
plt.hist(z0_arr.flatten(), bins=50, alpha=0.4, label='f0')
plt.hist(z1_arr.flatten(), bins=50, alpha=0.4, label='f1')
plt.legend()
plt.show()
```



We can compute sequences of likelihood ratios using simulated samples.

```
l = lambda z: wf.f0(z) / wf.f1(z)
```

```
l0_arr = l(z0_arr)
l1_arr = l(z1_arr)

L0_arr = np.cumprod(l0_arr, 1)
L1_arr = np.cumprod(l1_arr, 1)
```

Column `t-1` of `L0_arr` and `L1_arr` holds $L(z^t)$, the likelihood ratio after t draws.

With an empirical distribution of likelihood ratios in hand, we can draw **receiver operating characteristic curves** by enumerating (PFA, PD) pairs given each sample size t.

```
PFA = np.arange(0, 100, 1)

for t in range(1, 15, 4):
    percentile = np.percentile(L0_arr[:, t-1], PFA)
    PD = [np.sum(L1_arr[:, t-1] < p) / N for p in percentile]

    plt.plot(PFA / 100, PD, label=f"t={t}")

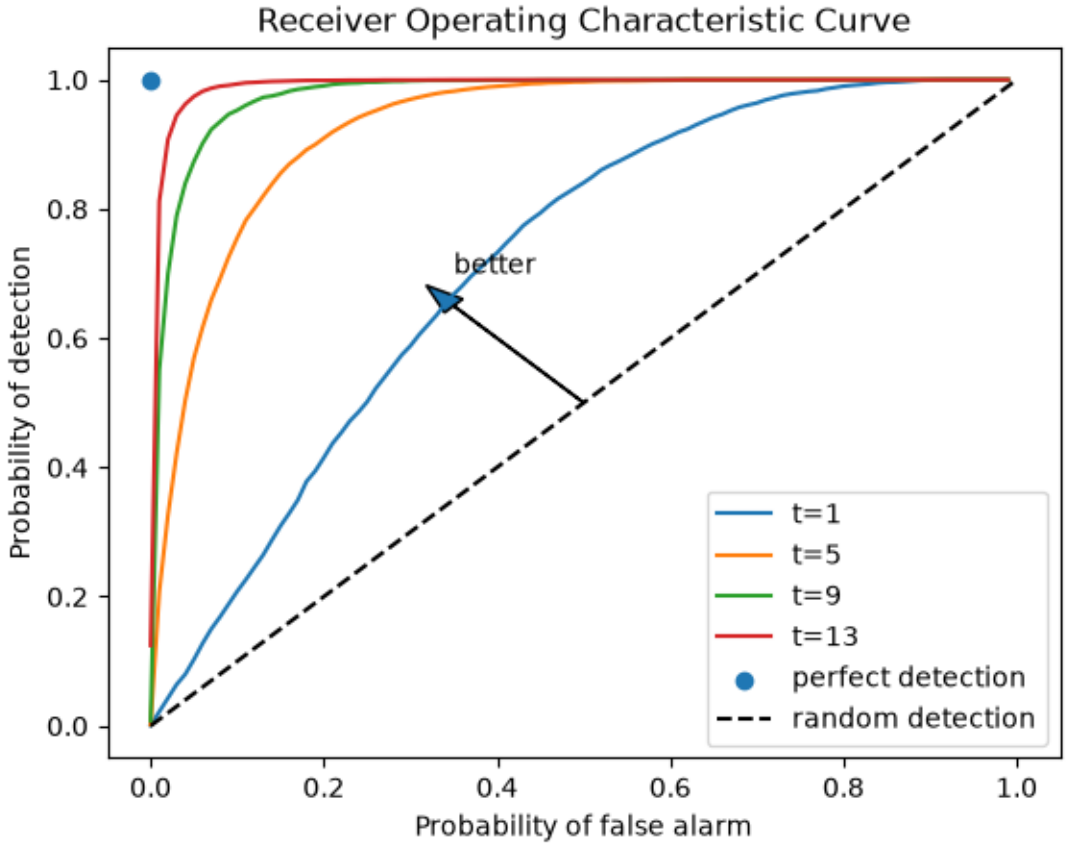
plt.scatter(0, 1, label="perfect detection")
plt.plot([0, 1], [0, 1], color='k', ls='--', label="random detection")

plt.arrow(0.5, 0.5, -0.15, 0.15, head_width=0.03)
plt.text(0.35, 0.7, "better")
plt.xlabel("Probability of false alarm")
plt.ylabel("Probability of detection")
```

(continues on next page)

(continued from previous page)

```
plt.legend()
plt.title("Receiver Operating Characteristic Curve")
plt.show()
```



We can minimize the expected total loss presented in equation (33.1) by choosing (t, d).

Doing that delivers an expected loss

$$\bar{V}_{fre} = \min_{t,d} \bar{V}_{fre}(t, d).$$

We let t range over $1, 2, \ldots, T$.[1]

We first consider the case in which $\pi^* = \Pr\{\text{nature selects } f_0\} = 0.5$.

We can solve the minimization problem in two steps.

First, we fix t and find the optimal cutoff d and consequently the minimal $\bar{V}_{fre}(t)$.

For every t, the cutoff

$$d^* = \frac{(1-\pi^*)\bar{L}_0}{\pi^*\bar{L}_1} \tag{33.2}$$

minimizes $\bar{V}_{fre}(t, d)$.

The rule "decide f_1 if $L(z^t) < d^*$" chooses the decision with the smaller posterior expected loss after t draws.

[1] A decision maker can also decide without drawing any z, at expected loss $\min\{\pi^*\bar{L}_1, (1-\pi^*)\bar{L}_0\}$. On our grid of values of π^*, that option beats every $t \geq 1$ at $\pi^* = 0.1, 0.86$, and 0.9. Exercise 33.10.3 lets the fixed-sample rule choose $t = 0$.

The frequentist rule with $d = d^*$ is therefore a Bayes rule whose sample size is fixed in advance.

Exercise 33.10.1 asks you to derive (33.2) and to verify it numerically.

Second, we minimize $\bar{V}_{fre}(t) = \bar{V}_{fre}(t, d^*)$ over t.

Here is Python code that does that and then plots a useful graph.

```
@jit
def V_fre_d_t(d, t, L0_arr, L1_arr, π_star, wf):

    N = L0_arr.shape[0]

    PFA = np.sum(L0_arr[:, t-1] < d) / N
    PD = np.sum(L1_arr[:, t-1] < d) / N

    V = π_star * PFA * wf.L1 + (1 - π_star) * (1 - PD) * wf.L0

    return V
```

```
def d_star(π_star, wf):
    "Cutoff for L = f0/f1 that minimizes expected loss for any t."

    return (1 - π_star) * wf.L0 / (π_star * wf.L1)


def V_fre_t(t, L0_arr, L1_arr, π_star, wf):

    d = d_star(π_star, wf)
    V = V_fre_d_t(d, t, L0_arr, L1_arr, π_star, wf)

    PFA = np.mean(L0_arr[:, t-1] < d)
    PD = np.mean(L1_arr[:, t-1] < d)

    return V, PFA, PD
```

```
def compute_V_fre(L0_arr, L1_arr, π_star, wf):

    T = L0_arr.shape[1]

    V_fre_arr = np.empty(T)
    PFA_arr = np.empty(T)
    PD_arr = np.empty(T)

    for t in range(1, T+1):
        V, PFA, PD = V_fre_t(t, L0_arr, L1_arr, π_star, wf)
        V_fre_arr[t-1] = wf.c * t + V
        PFA_arr[t-1] = PFA
        PD_arr[t-1] = PD

    return V_fre_arr, PFA_arr, PD_arr
```

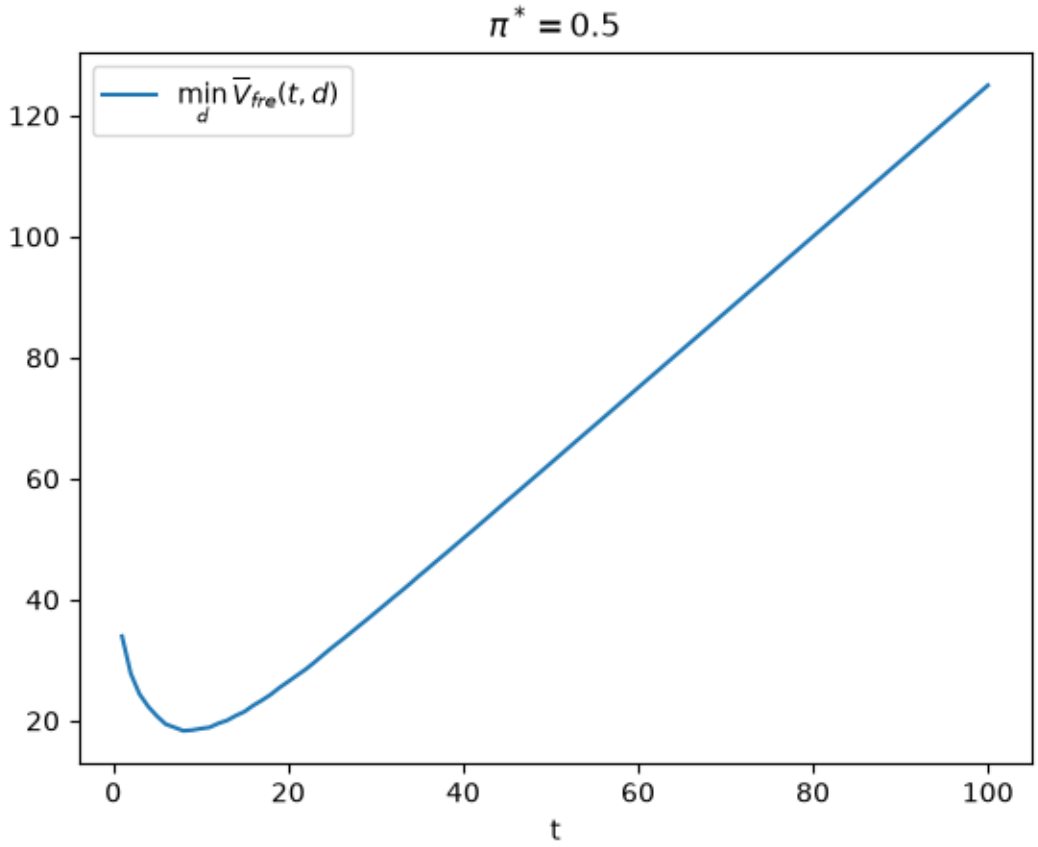
```
π_star = 0.5
V_fre_arr, PFA_arr, PD_arr = compute_V_fre(L0_arr, L1_arr, π_star, wf)

plt.plot(range(1, T+1), V_fre_arr, label=r'$\min_{d} \overline{V}_{fre}(t,d)$')
plt.xlabel('t')
plt.title(r'$\pi^*=0.5$')
```

(continues on next page)

(continued from previous page)

```
plt.legend()
plt.show()
```




```
t_optimal = np.argmin(V_fre_arr) + 1
```

The above graph illustrates how minimizing over t tells the frequentist to draw t_{optimal} observations and then decide.

Let's now change the value of π^* and watch how the decision rule changes.

```
n_π = 20
π_star_arr = np.linspace(0.1, 0.9, n_π)

V_fre_bar_arr = np.empty(n_π)
t_optimal_arr = np.empty(n_π)
PFA_optimal_arr = np.empty(n_π)
PD_optimal_arr = np.empty(n_π)

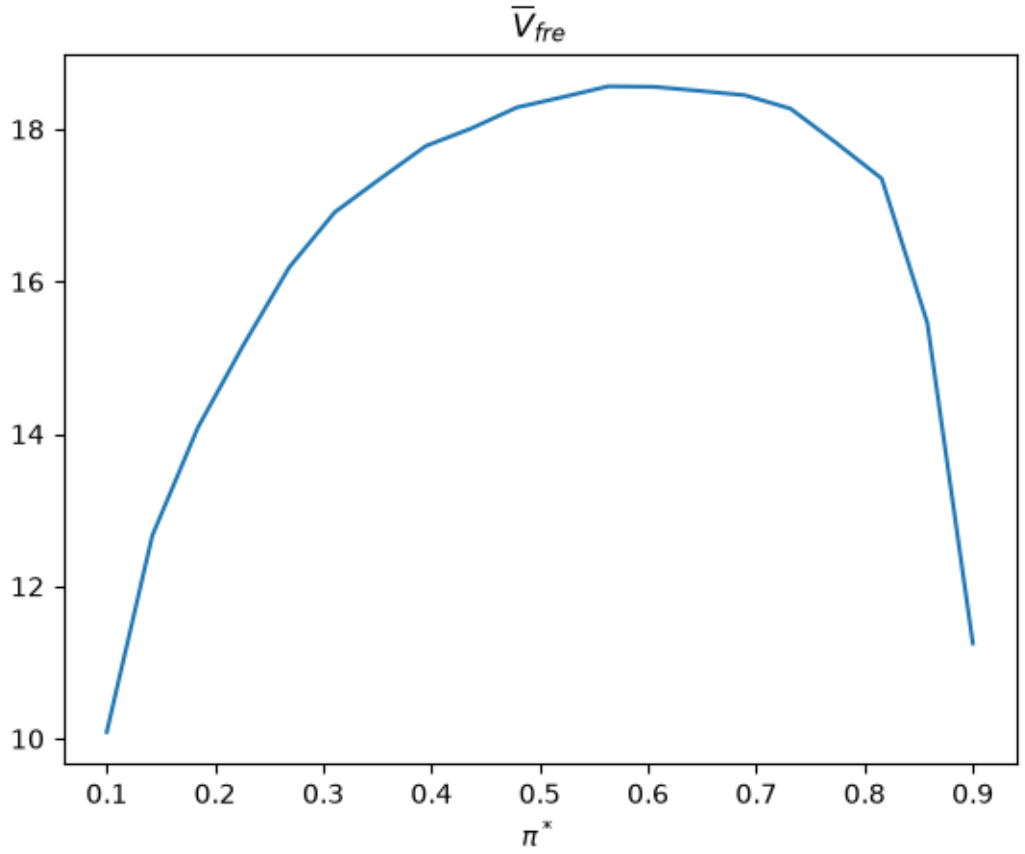
for i, π_star in enumerate(π_star_arr):
    V_fre_arr, PFA_arr, PD_arr = compute_V_fre(L0_arr, L1_arr, π_star, wf)
    t_idx = np.argmin(V_fre_arr)

    V_fre_bar_arr[i] = V_fre_arr[t_idx]
    t_optimal_arr[i] = t_idx + 1
    PFA_optimal_arr[i] = PFA_arr[t_idx]
    PD_optimal_arr[i] = PD_arr[t_idx]
```

```
plt.plot(π_star_arr, V_fre_bar_arr)
plt.xlabel(r'$\pi^*$')
plt.title(r'$\overline{V}_{fre}$')

plt.show()

V_no_data = np.minimum(π_star_arr * wf.L1, (1 - π_star_arr) * wf.L0)
print("π* at which deciding without data beats every t ≥ 1:",
      np.round(π_star_arr[V_no_data < V_fre_bar_arr], 2))
```




```
π* at which deciding without data beats every t ≥ 1: [0.1  0.86 0.9 ]
```

The following shows how optimal sample size t and targeted (PFA, PD) change as π^* varies.

```
fig, axs = plt.subplots(1, 2, figsize=(14, 5))

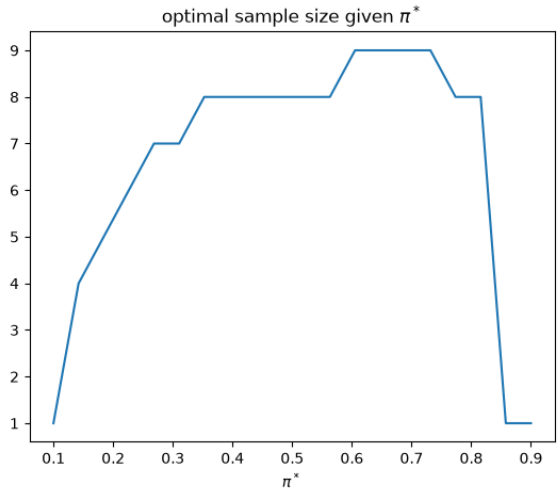
axs[0].plot(π_star_arr, t_optimal_arr)
axs[0].set_xlabel(r'$\pi^*$')
axs[0].set_title(r'optimal sample size given $\pi^*$')

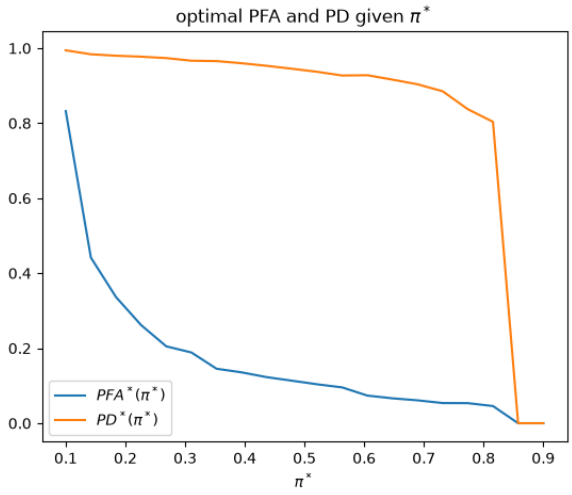
axs[1].plot(π_star_arr, PFA_optimal_arr, label=r'$PFA^*(\pi^*)$')
axs[1].plot(π_star_arr, PD_optimal_arr, label=r'$PD^*(\pi^*)$')
axs[1].set_xlabel(r'$\pi^*$')
axs[1].legend()
axs[1].set_title(r'optimal PFA and PD given $\pi^*$')
```

(continues on next page)

(continued from previous page)

```
plt.show()
```





33.4 Bayesian Decision Rule

In *A Problem that Stumped Milton Friedman*, we learned how Abraham Wald confirmed the Navy Captain's hunch that there is a better decision rule.

In *A Bayesian Formulation of Friedman and Wald's Problem* we presented a Bayesian procedure that makes decisions by comparing a Bayesian posterior probability π with cutoff probabilities called A and B.

To proceed, we borrow some Python code from the quantecon lecture *A Bayesian Formulation of Friedman and Wald's Problem* that computes optimal values of A and B.

We rewrite that code for our convention that π is the probability of f_0.

```
@jit(parallel=True)
def Q(h, wf):

    c, π_grid = wf.c, wf.π_grid
    L0, L1 = wf.L0, wf.L1
    z0, z1 = wf.z0, wf.z1
    mc_size = wf.mc_size

    κ = wf.κ

    h_new = np.empty_like(π_grid)
    h_func = lambda p: np.interp(p, π_grid, h)

    for i in prange(len(π_grid)):
        π = π_grid[i]

        # Find the expected value of J by integrating over z
        integral_f0, integral_f1 = 0, 0
        for m in range(mc_size):
            π_0 = κ(z0[m], π)  # Draw z from f0 and update π
            integral_f0 += min((1 - π_0) * L0, π_0 * L1, h_func(π_0))
```

(continues on next page)

(continued from previous page)

```
            π_1 = κ(z1[m], π)  # Draw z from f1 and update π
            integral_f1 += min((1 - π_1) * L0, π_1 * L1, h_func(π_1))

        integral = (π * integral_f0 + (1 - π) * integral_f1) / mc_size

        h_new[i] = c + integral

    return h_new
```

```
@jit
def solve_model(wf, tol=1e-4, max_iter=1000):
    """
    Compute the continuation value function

    * wf is an instance of WaldFriedman
    """

    # Set up loop
    h = np.zeros(len(wf.π_grid))
    i = 0
    error = tol + 1

    while i < max_iter and error > tol:
        h_new = Q(h, wf)
        error = np.max(np.abs(h - h_new))
        i += 1
        h = h_new

    if error > tol:
        print("Failed to converge!")

    return h_new
```

```
h_star = solve_model(wf)
```

```
def find_cutoff_rule(wf, h):
    """
    This function takes a continuation value function and returns the
    cutoffs (B, A): the decision maker accepts f1 if π ≤ B, accepts f0
    if π ≥ A, and draws another z if B < π < A
    """

    π_grid = wf.π_grid

    # Expected losses from accepting each model
    payoff_f0 = (1 - π_grid) * wf.L0
    payoff_f1 = π_grid * wf.L1

    # B is the largest π at which accepting f1 is optimal;
    # A is the smallest π at which accepting f0 is optimal
    B = π_grid[payoff_f1 <= np.minimum(payoff_f0, h)].max()
    A = π_grid[payoff_f0 <= np.minimum(payoff_f1, h)].min()

    return (B, A)

B, A = find_cutoff_rule(wf, h_star)
```

(continues on next page)

(continued from previous page)

```
print(f"B = {B:.3f}, A = {A:.3f}")

cost_L0 = (1 - wf.π_grid) * wf.L0
cost_L1 = wf.π_grid * wf.L1

fig, ax = plt.subplots(figsize=(10, 6))

ax.plot(wf.π_grid, h_star, label='continuation value')
ax.plot(wf.π_grid, cost_L1, label='choose f1')
ax.plot(wf.π_grid, cost_L0, label='choose f0')
ax.plot(wf.π_grid,
        np.amin(np.column_stack([h_star, cost_L0, cost_L1]),axis=1),
        lw=15, alpha=0.1, color='b', label='minimum cost')

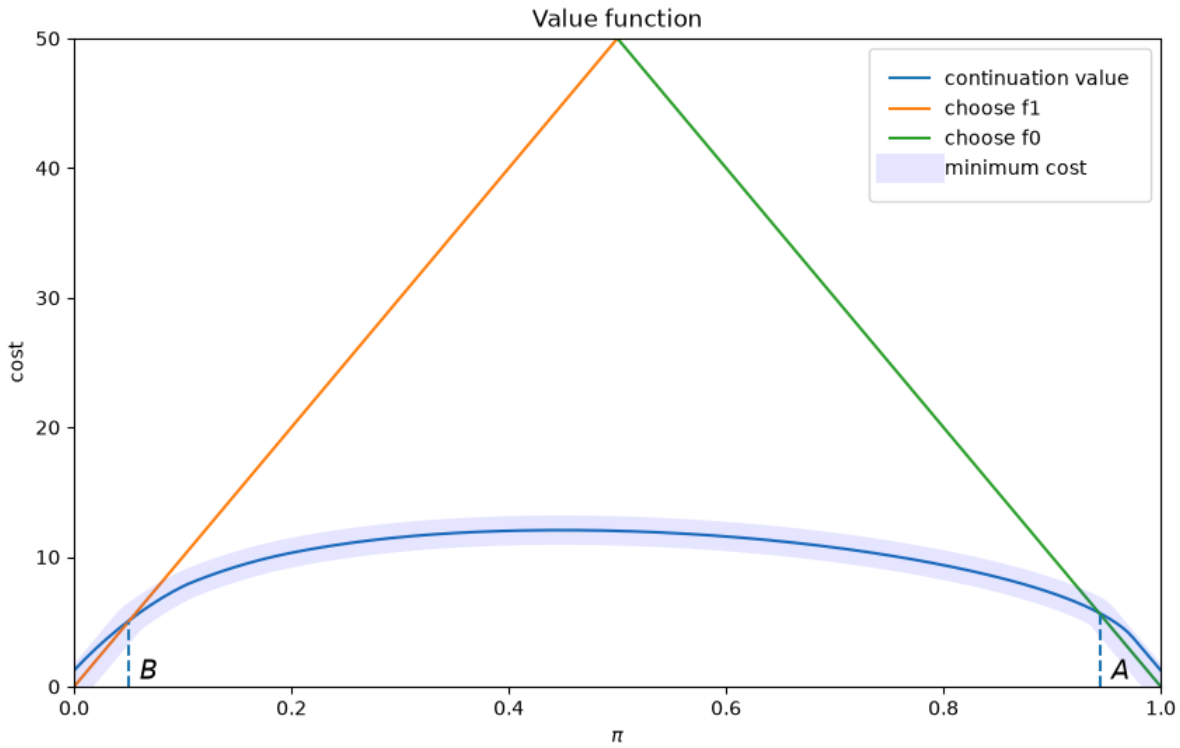
ax.annotate(r"$B$", xy=(B + 0.01, 0.5), fontsize=14)
ax.annotate(r"$A$", xy=(A + 0.01, 0.5), fontsize=14)

plt.vlines(B, 0, B * wf.L1, linestyle="--")
plt.vlines(A, 0, (1 - A) * wf.L0, linestyle="--")

ax.set(xlim=(0, 1), ylim=(0, 0.5 * max(wf.L0, wf.L1)), ylabel="cost",
       xlabel=r"$\pi$", title="Value function")

plt.legend(borderpad=1.1)
plt.show()
```

```
B = 0.050, A = 0.945
```



The above figure portrays the value function plotted against the decision maker's Bayesian posterior.

It also shows the cutoff probabilities A and B.

The Bayesian decision rule is:

- accept H_0 if $\pi \geq A$
- accept H_1 if $\pi \leq B$
- delay deciding and draw another z if $B < \pi < A$

We can calculate two "objective" loss functions under this situation conditioning on knowing for sure that nature has selected f_0, in the first case, or f_1, in the second case.

1. under f_0,

$$V^0(\pi) = \begin{cases} 0 & \text{if } A \leq \pi, \\ c + E_{f_0} V^0(\pi') & \text{if } B < \pi < A, \\ \bar{L}_1 & \text{if } \pi \leq B. \end{cases}$$

2. under f_1

$$V^1(\pi) = \begin{cases} \bar{L}_0 & \text{if } A \leq \pi, \\ c + E_{f_1} V^1(\pi') & \text{if } B < \pi < A, \\ 0 & \text{if } \pi \leq B. \end{cases}$$

where $\pi' = \frac{\pi f_0(z')}{\pi f_0(z') + (1-\pi) f_1(z')}$ and E_{f_i} integrates over z' drawn from f_i.

Given a prior probability π_0, the expected loss for the Bayesian is

$$\bar{V}_{Bayes}(\pi_0) = \pi^* V^0(\pi_0) + (1 - \pi^*) V^1(\pi_0).$$

Below we write some Python code that computes $V^0(\pi)$ and $V^1(\pi)$ numerically.

The code iterates on the two functional equations at the grid points in the continuation region (B, A).

It replaces each expectation with an average over the `mc_size` Monte Carlo draws stored in `wf`.

It uses the terminal losses exactly when π' leaves (B, A) and interpolates only between grid points inside (B, A).

```
@jit(parallel=True)
def V_q(wf, A, B, flag, tol=1e-5, max_iter=10_000):
    """
    Expected loss of the Bayesian cutoff rule (B, A) when f0 (flag=0)
    or f1 (flag=1) generates the data
    """

    π_grid = wf.π_grid

    if flag == 0:
        z_arr = wf.z0
        loss_f0, loss_f1 = 0.0, wf.L1   # losses from accepting f0, f1
    else:
        z_arr = wf.z1
        loss_f0, loss_f1 = wf.L0, 0.0

    m = len(z_arr)

    # Terminal losses outside the continuation region
    V = np.where(π_grid >= A, loss_f0, np.where(π_grid <= B, loss_f1, 0.0))
```

(continues on next page)

(continued from previous page)

```
    cont = (π_grid > B) & (π_grid < A)
    π_cont = π_grid[cont]
    n_cont = len(π_cont)
    if n_cont == 0:
        return V

    V_cont = np.zeros(n_cont)
    for it in range(max_iter):
        V_old = V_cont.copy()

        for i in prange(n_cont):
            total = 0.0
            for j in range(m):
                π_next = wf.κ(z_arr[j], π_cont[i])
                if π_next >= A:
                    total += loss_f0
                elif π_next <= B:
                    total += loss_f1
                else:
                    total += np.interp(π_next, π_cont, V_old)
            V_cont[i] = wf.c + total / m

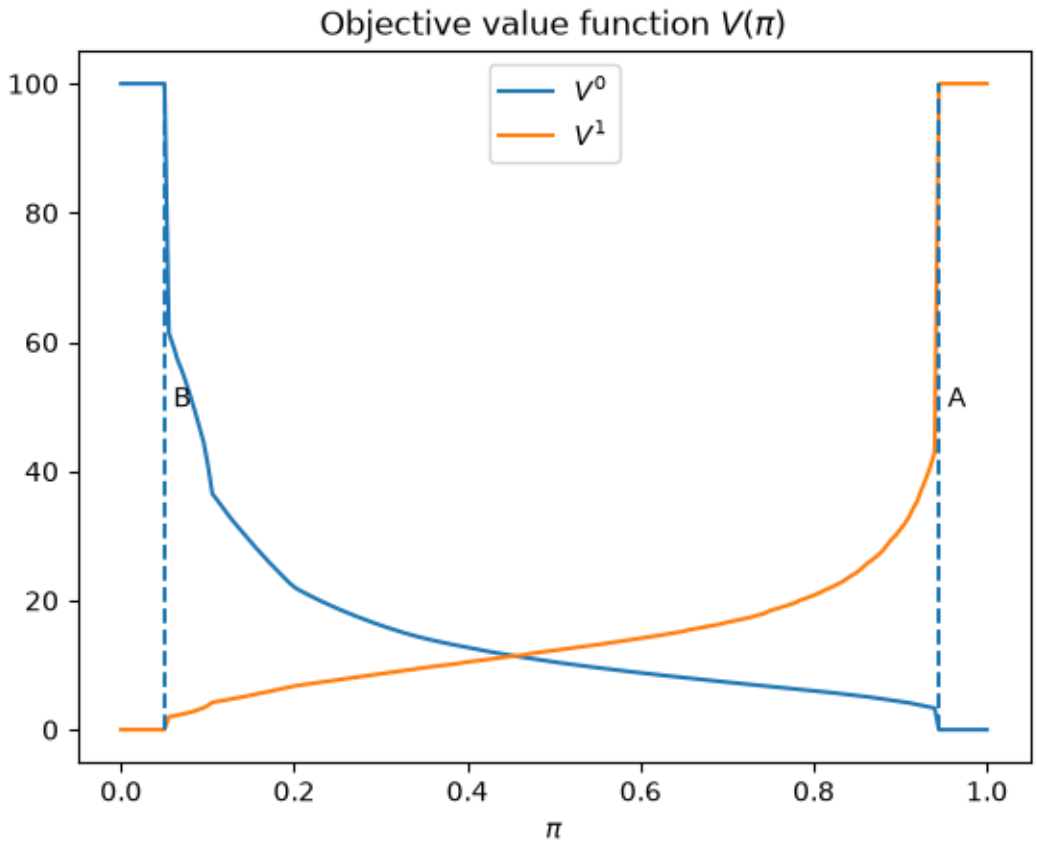
        if np.max(np.abs(V_cont - V_old)) < tol:
            break

    V[cont] = V_cont

    return V
```

```
V0 = V_q(wf, A, B, 0)
V1 = V_q(wf, A, B, 1)

plt.plot(wf.π_grid, V0, label='$V^0$')
plt.plot(wf.π_grid, V1, label='$V^1$')
plt.vlines(B, 0, wf.L0, linestyle='--')
plt.text(B+0.01, wf.L0/2, 'B')
plt.vlines(A, 0, wf.L0, linestyle='--')
plt.text(A+0.01, wf.L0/2, 'A')
plt.xlabel(r'$\pi$')
plt.title(r'Objective value function $V(\pi)$')
plt.legend()
plt.show()
```



Given an assumed value for $\pi^* = \Pr\{\text{nature selects } f_0\}$, we can then compute $\bar{V}_{Bayes}(\pi_0)$.

We can then determine an initial Bayesian prior π_0^* that minimizes this objective concept of expected loss.

A theorem tells us that $\pi_0^* = \pi^*$.

The Bayes rule that starts from the prior π^* minimizes π^*-weighted expected loss over all sequential decision rules.

A Bayesian who starts from another prior π_0 uses one of those rules.

Therefore $\bar{V}_{Bayes}(\pi_0) \geq \bar{V}_{Bayes}(\pi^*)$ for every π_0.

The figure below plots four cases corresponding to $\pi^* = 0.25, 0.3, 0.5, 0.7$.

The numerical minimizers lie near π^* but not exactly at it.

The function $\bar{V}_{Bayes}(\pi_0)$ is flat near its minimum, so small grid and Monte Carlo errors move the computed π_0^*.

```
def compute_V_baye_bar(π_star, V0, V1, wf):

    V_baye = π_star * V0 + (1 - π_star) * V1
    π_idx = np.argmin(V_baye)
    π_optimal = wf.π_grid[π_idx]
    V_baye_bar = V_baye[π_idx]

    return V_baye, π_optimal, V_baye_bar
```

```
π_star_arr = [0.25, 0.3, 0.5, 0.7]

fig, axs = plt.subplots(2, 2, figsize=(15, 10))
```

(continues on next page)

(continued from previous page)

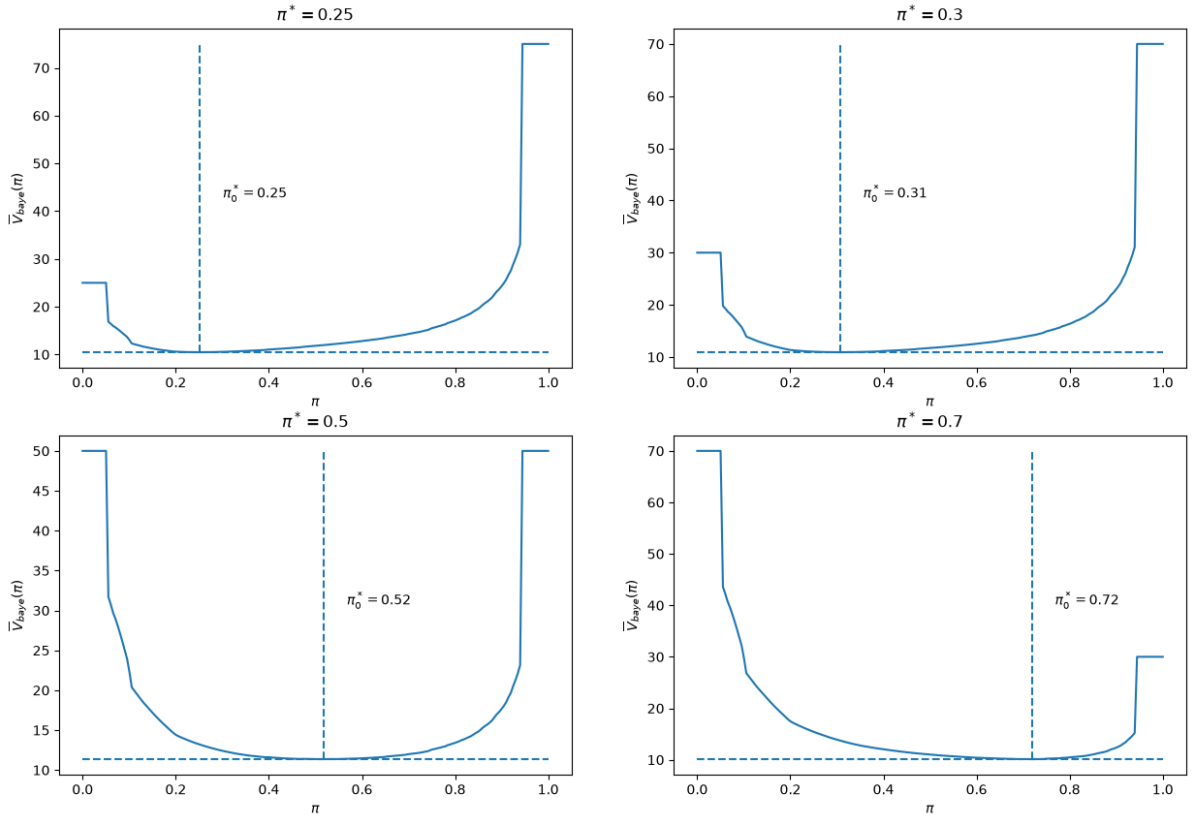
```
for i, π_star in enumerate(π_star_arr):
    row_i = i // 2
    col_i = i % 2

    V_baye, π_optimal, V_baye_bar = compute_V_baye_bar(π_star, V0, V1, wf)

    axs[row_i, col_i].plot(wf.π_grid, V_baye)
    axs[row_i, col_i].hlines(V_baye_bar, 0, 1, linestyle='--')
    axs[row_i, col_i].vlines(π_optimal, V_baye_bar, V_baye.max(), linestyle='--')
    axs[row_i, col_i].text(π_optimal+0.05, (V_baye_bar + V_baye.max()) / 2,
                           r'${\pi_0^*}=$'+f'{π_optimal:0.2f}')
    axs[row_i, col_i].set_xlabel(r'$\pi$')
    axs[row_i, col_i].set_ylabel(r'$\overline{V}_{baye}(\pi)$')
    axs[row_i, col_i].set_title(r'$\pi^*=$' + f'{π_star}')

fig.suptitle(r'$\overline{V}_{baye}(\pi)=\pi^*V^0(\pi) + (1-\pi^*)V^1(\pi)$',␣
 ↪fontsize=16)
plt.show()
```

$$\overline{V}_{baye}(\pi) = \pi^* V^0(\pi) + (1 - \pi^*) V^1(\pi)$$



The following Python code computes π_0^* and $\bar{V}_{Bayes} = \bar{V}_{Bayes}(\pi_0^*)$ for a grid of values of π^*.

```
π_star_arr = np.linspace(0.1, 0.9, n_π)
V_baye_bar_arr = np.empty_like(π_star_arr)
π_optimal_arr = np.empty_like(π_star_arr)
```

(continues on next page)

(continued from previous page)

```
for i, π_star in enumerate(π_star_arr):

    V_baye, π_optimal, V_baye_bar = compute_V_baye_bar(π_star, V0, V1, wf)

    V_baye_bar_arr[i] = V_baye_bar
    π_optimal_arr[i] = π_optimal

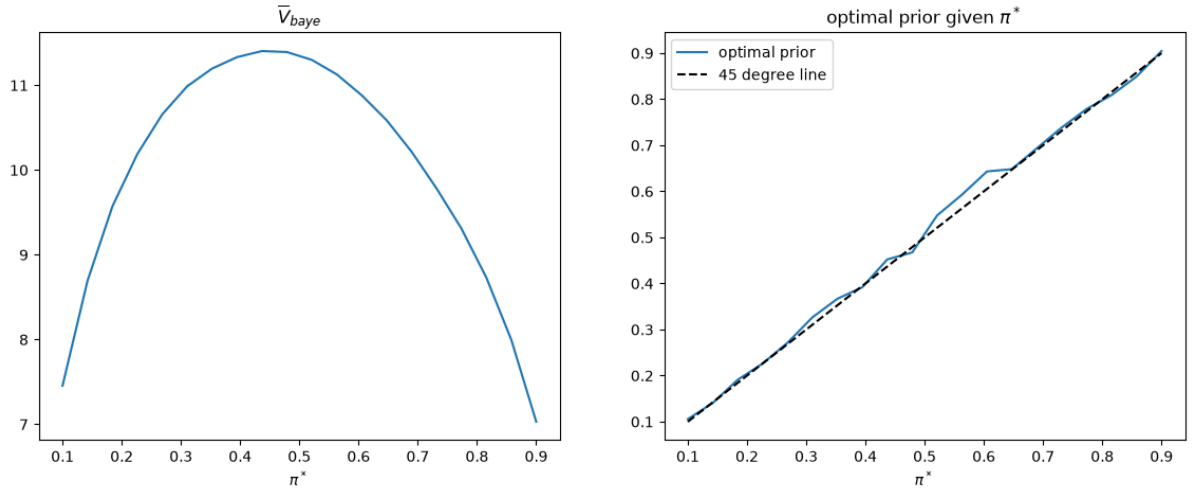
fig, axs = plt.subplots(1, 2, figsize=(14, 5))

axs[0].plot(π_star_arr, V_baye_bar_arr)
axs[0].set_xlabel(r'$\pi^*$')
axs[0].set_title(r'$\overline{V}_{baye}$')

axs[1].plot(π_star_arr, π_optimal_arr, label='optimal prior')
axs[1].plot([π_star_arr.min(), π_star_arr.max()],
            [π_star_arr.min(), π_star_arr.max()],
            c='k', linestyle='--', label='45 degree line')
axs[1].set_xlabel(r'$\pi^*$')
axs[1].set_title(r'optimal prior given $\pi^*$')
axs[1].legend()

plt.show()

print(f"largest |π_0^* - π^*| = {np.max(np.abs(π_optimal_arr - π_star_arr)):.3f}")
```




```
largest |π_0^* - π^*| = 0.038
```

The computed π_0^* stays within 0.04 of the 45 degree line.

33.5 Was the Navy Captain's Hunch Correct?

We now compare average losses obtained by our frequentist Neyman-Pearson and Bayesian decision rules.

The comparison is between a rule with a fixed sample size and a sequential rule.

Both rules use the prior π^* and the losses $\bar{L}_0$, $\bar{L}_1$, and c.

A fixed-sample rule (t, d) is a sequential rule that ignores the data until date t and then stops.

The Bayesian rule minimizes expected loss over all sequential rules.

Therefore $\bar{V}_{Bayes} \leq \bar{V}_{fre}$ for every π^*.

The simulations below illustrate this inequality and measure the gap.

As a starting point, let's compare average loss functions when $\pi^* = 0.5$.

```
π_star = 0.5
```

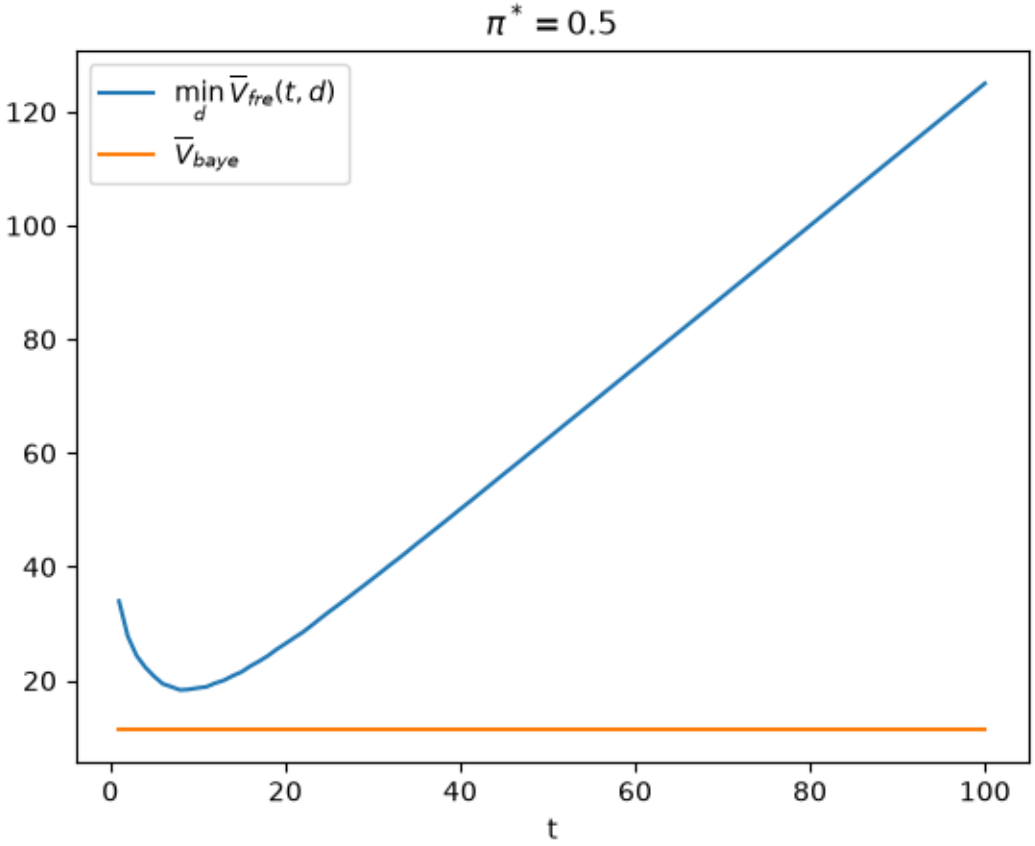
```
# frequentist
V_fre_arr, PFA_arr, PD_arr = compute_V_fre(L0_arr, L1_arr, π_star, wf)

# bayesian
V_baye = π_star * V0 + (1 - π_star) * V1
V_baye_bar = V_baye.min()

print(f"V_fre = {V_fre_arr.min():.2f} at t = {np.argmin(V_fre_arr) + 1}")
print(f"V_Bayes = {V_baye_bar:.2f}")
```

```
V_fre = 18.33 at t = 8
V_Bayes = 11.36
```

```
plt.plot(range(1, T+1), V_fre_arr, label=r'$\min_{d} \overline{V}_{fre}(t,d)$')
plt.plot([1, T], [V_baye_bar, V_baye_bar], label=r'$\overline{V}_{baye}$')
plt.xlabel('t')
plt.title(r'$\pi^*=0.5$')
plt.legend()
plt.show()
```



Evidently, there is no sample size t at which the Neyman-Pearson decision rule attains a lower loss function than does the Bayesian rule.

At $\pi^* = 0.5$, the best fixed-sample rule draws 8 observations and has expected loss 18.33; the Bayesian rule has expected loss 11.36.

The Bayesian value is computed with the $1{,}000$ Monte Carlo draws stored in `wf.z0` and `wf.z1`, so its last digits are not reliable; the simulation in a later section puts the Bayesian loss near 12.3, still far below 18.33.

The following graph compares the two rules for all values of π^*.

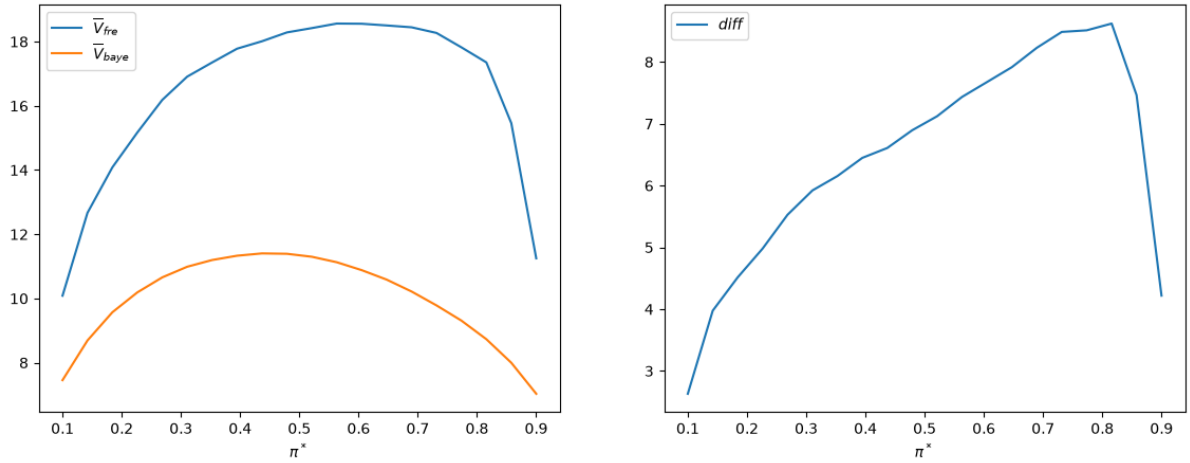
```
fig, axs = plt.subplots(1, 2, figsize=(14, 5))

axs[0].plot(π_star_arr, V_fre_bar_arr, label=r'$\overline{V}_{fre}$')
axs[0].plot(π_star_arr, V_baye_bar_arr, label=r'$\overline{V}_{baye}$')
axs[0].legend()
axs[0].set_xlabel(r'$\pi^*$')

axs[1].plot(π_star_arr, V_fre_bar_arr - V_baye_bar_arr, label='$diff$')
axs[1].legend()
axs[1].set_xlabel(r'$\pi^*$')

plt.show()

diff = V_fre_bar_arr - V_baye_bar_arr
print(f"min diff = {diff.min():.2f} at π* = {π_star_arr[np.argmin(diff)]:.2f}")
print(f"max diff = {diff.max():.2f} at π* = {π_star_arr[np.argmax(diff)]:.2f}")
```




```
min diff = 2.63 at π* = 0.10
max diff = 8.62 at π* = 0.82
```

The right panel of the above graph plots the difference $\bar{V}_{fre} - \bar{V}_{Bayes}$.

It is always positive, as the theorem requires.

The gap is smallest, 2.63, at $\pi^* = 0.1$ and largest, 8.62, at $\pi^* \approx 0.82$.

33.6 More Details

We can provide more insights by focusing on the case in which $\pi^* = 0.5 = \pi_0$.

```
π_star = 0.5
```

Recall that when $\pi^* = 0.5$, the frequentist Neyman-Pearson decision rule sets a sample size `t_optimal` **ex ante**.

For our parameter settings, we can compute its value:

```
t_optimal
```

```
np.int64(8)
```

For convenience, let's define `t_idx` as the Python array index corresponding to `t_optimal` sample size.

```
t_idx = t_optimal - 1
```

33.7 Distribution of Bayesian Decision Rule's Time to Decide

We use simulations to compute the frequency distribution of the time to decide for the Bayesian decision rule and compare that time to the frequentist rule's fixed t.

The following Python code creates a graph that shows the frequency distribution of Bayesian times to decide of Bayesian decision maker, conditional on distribution $q = f_0$ or $q = f_1$ generating the data.

The blue and orange dashed lines show averages for the Bayesian decision rule under f_0 and f_1, while the black dashed line shows the frequentist optimal sample size t.

On average the Bayesian rule decides **earlier** than the frequentist rule when $q = f_0$ and **later** when $q = f_1$.

The function `check_results` records time $T+1$ for a path that has not reached a decision by T.

```
@jit(parallel=True)
def check_results(L_arr, A, B, flag, π0):

    N, T = L_arr.shape

    time_arr = np.full(N, T + 1)                    # T + 1 means no decision by T
    correctness = np.zeros(N, dtype=np.bool_)

    π_arr = π0 * L_arr / (π0 * L_arr + 1 - π0)

    for i in prange(N):
        for t in range(T):
            if (π_arr[i, t] <= B) or (π_arr[i, t] >= A):
                time_arr[i] = t + 1
                correctness[i] = (flag == 0 and π_arr[i, t] >= A) or (flag == 1 and π_
↪arr[i, t] <= B)
                break

    return time_arr, correctness
```

```
time_arr0, correctness0 = check_results(L0_arr, A, B, 0, π_star)
time_arr1, correctness1 = check_results(L1_arr, A, B, 1, π_star)

# unconditional distribution
time_arr_u = np.concatenate((time_arr0, time_arr1))
correctness_u = np.concatenate((correctness0, correctness1))

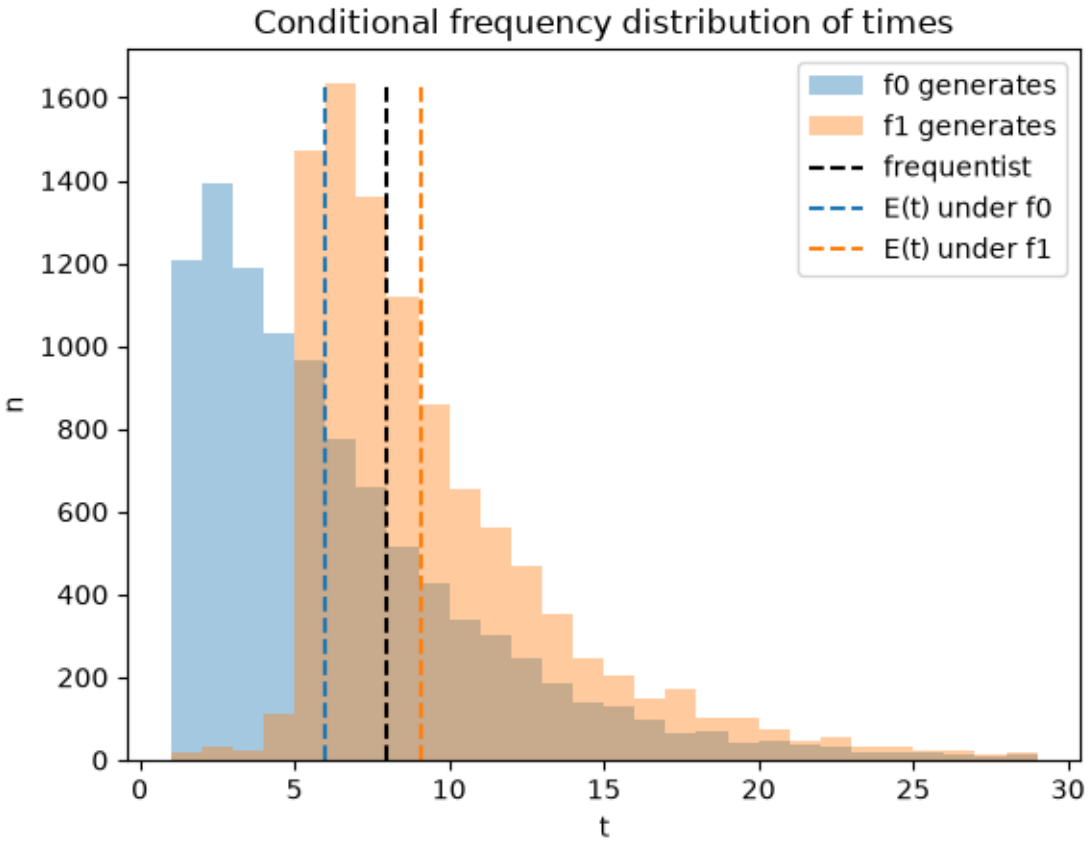
print("paths with no decision by T:", np.sum(time_arr_u > T))
print(f"mean time to decide: {time_arr0.mean():.2f} under f0, "
      f"{time_arr1.mean():.2f} under f1, {time_arr_u.mean():.2f} pooled")
```

```
paths with no decision by T: 0
mean time to decide: 5.97 under f0, 9.11 under f1, 7.54 pooled
```

```
n1 = plt.hist(time_arr0, bins=range(1, 30), alpha=0.4, label='f0 generates')[0]
n2 = plt.hist(time_arr1, bins=range(1, 30), alpha=0.4, label='f1 generates')[0]
plt.vlines(t_optimal, 0, max(n1.max(), n2.max()), linestyle='--',
           color='k', label='frequentist')
plt.vlines(np.mean(time_arr0), 0, max(n1.max(), n2.max()),
           linestyle='--', color='C0', label='E(t) under f0')
plt.vlines(np.mean(time_arr1), 0, max(n1.max(), n2.max()),
           linestyle='--', color='C1', label='E(t) under f1')
plt.legend();

plt.xlabel('t')
plt.ylabel('n')
plt.title('Conditional frequency distribution of times')

plt.show()
```



Later we'll figure out how these distributions ultimately affect objective expected values under the Neyman-Pearson and Bayesian decision rules.

To begin, let's look at simulations of the Bayesian's beliefs over time.

We can compute updated beliefs at any time t using the one-to-one mapping from L_t to π_t given π_0 described in the section on *likelihood ratio processes and Bayes' law* of *Likelihood Ratio Processes and Bayesian Learning*.

```
π0_arr = π_star * L0_arr / (π_star * L0_arr + 1 - π_star)
π1_arr = π_star * L1_arr / (π_star * L1_arr + 1 - π_star)
```

```
fig, axs = plt.subplots(1, 2, figsize=(14, 4))

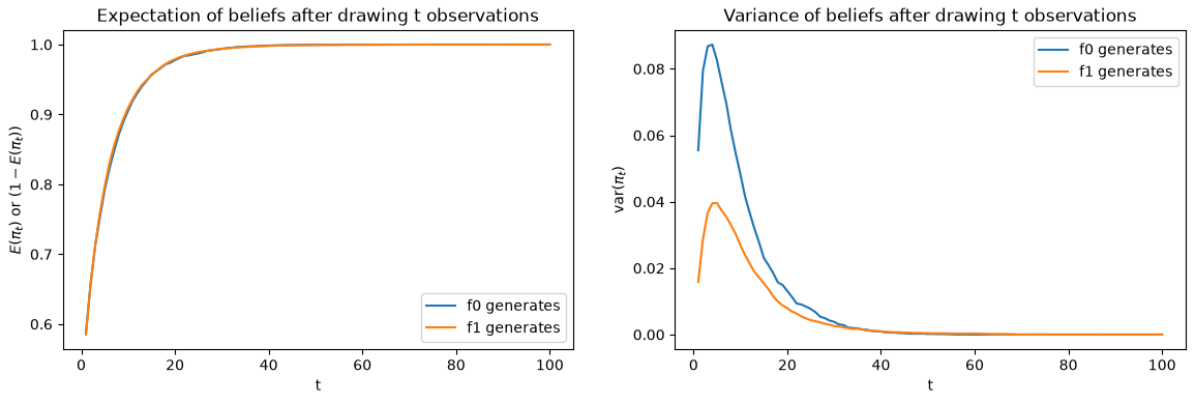
axs[0].plot(np.arange(1, π0_arr.shape[1]+1), np.mean(π0_arr, 0), label='f0 generates')
axs[0].plot(np.arange(1, π1_arr.shape[1]+1), 1 - np.mean(π1_arr, 0), label='f1␣
↪generates')
axs[0].set_xlabel('t')
axs[0].set_ylabel(r'$E(\pi_t)$ or ($1 - E(\pi_t)$)')
axs[0].set_title('Expectation of beliefs after drawing t observations')
axs[0].legend()

axs[1].plot(np.arange(1, π0_arr.shape[1]+1), np.var(π0_arr, 0), label='f0 generates')
axs[1].plot(np.arange(1, π1_arr.shape[1]+1), np.var(π1_arr, 0), label='f1 generates')
axs[1].set_xlabel('t')
axs[1].set_ylabel(r'var($\pi_t$)')
axs[1].set_title('Variance of beliefs after drawing t observations')
axs[1].legend()
```

(continues on next page)

(continued from previous page)

```
plt.show()
```



The above figures compare averages and variances of updated Bayesian posteriors after t draws.

The left graph compares $E_{f_0}(\pi_t)$ to $1 - E_{f_1}(\pi_t)$: they lie on top of each other.

The martingale property of π_t implies this outcome.

Under the Bayesian's subjective distribution, which mixes f_0 and f_1 with weights π_0 and $1 - \pi_0$, π_t is a martingale with mean π_0.

Therefore $\pi_0 E_{f_0}(\pi_t) + (1 - \pi_0) E_{f_1}(\pi_t) = \pi_0$ for every t.

With $\pi_0 = 1/2$, this identity becomes $E_{f_0}(\pi_t) = 1 - E_{f_1}(\pi_t)$.

The two simulated curves differ only by Monte Carlo error.

The right graph shows that the variance of π_t is lower under f_1 when t is small.

An asymmetry in Kullback-Leibler divergences explains why the Bayesian decision maker waits longer to decide when f_1 generates the data.

Statistical Divergence Measures describes Kullback-Leibler divergence.

Under f_0, $\log L_t$ rises on average by $KL(f_0 \| f_1) = E_{f_0}[\log(f_0(z)/f_1(z))]$ per draw.

Under f_1, $\log L_t$ falls on average by $KL(f_1 \| f_0) = E_{f_1}[\log(f_1(z)/f_0(z))]$ per draw.

With $\pi_0 = 1/2$, the Bayesian stops when $\log L_t$ leaves the interval $(\log L_B, \log L_A)$, where $L_A = A/(1 - A)$ and $L_B = B/(1 - B)$.

Wald's approximation to the expected number of draws ignores the amount by which $\log L_t$ overshoots a threshold:

$$E_{f_1}[N] \approx \frac{\Pr_{f_1}\{\text{accept } f_1\} \log L_B + \Pr_{f_1}\{\text{accept } f_0\} \log L_A}{-KL(f_1 \| f_0)}$$

and similarly under f_0.

```
def KL(f, g):
    "KL(f || g), the integral of f log(f/g) over [0, 1]."
    integrand = lambda z: f(z) * np.log(f(z) / g(z))
    return quad(integrand, 0, 1)[0]

KL_01 = KL(wf.f0, wf.f1)
KL_10 = KL(wf.f1, wf.f0)

LA = (1 - π_star) * A / (π_star * (1 - A))
```

(continues on next page)

(continued from previous page)

```
LB = (1 - π_star) * B / (π_star * (1 - B))

p0 = correctness0.mean()    # probability of accepting f0 under f0
p1 = correctness1.mean()    # probability of accepting f1 under f1

EN0_wald = (p0 * np.log(LA) + (1 - p0) * np.log(LB)) / KL_01
EN1_wald = (p1 * np.log(LB) + (1 - p1) * np.log(LA)) / (-KL_10)

print(f"KL(f0||f1) = {KL_01:.3f},  KL(f1||f0) = {KL_10:.3f}")
print(f"log L_A = {np.log(LA):.2f},  log L_B = {np.log(LB):.2f}")
print(f"E_f0[N]: Wald approximation {EN0_wald:.1f}, simulation {time_arr0.mean():.1f}
 ↪")
print(f"E_f1[N]: Wald approximation {EN1_wald:.1f}, simulation {time_arr1.mean():.1f}
 ↪")

# Average overshoot of log L_A under f0 and of log L_B under f1
logL0_stop = np.log(L0_arr[np.arange(N), time_arr0 - 1])
logL1_stop = np.log(L1_arr[np.arange(N), time_arr1 - 1])
print(f"mean overshoot: {np.mean(logL0_stop[correctness0] - np.log(LA)):.2f} "
      f"under f0, {np.mean(np.log(LB) - logL1_stop[correctness1]):.2f} under f1")
```

```
KL(f0||f1) = 0.759,  KL(f1||f0) = 0.344
log L_A = 2.84,  log L_B = -2.94
E_f0[N]: Wald approximation 3.4, simulation 6.0
E_f1[N]: Wald approximation 8.3, simulation 9.1
mean overshoot: 1.98 under f0, 0.31 under f1
```

The divergence $KL(f_0 \| f_1) \approx 0.759$ is more than twice $KL(f_1 \| f_0) \approx 0.344$.

The thresholds $\log L_A \approx 2.84$ and $\log L_B \approx -2.94$ lie about equally far from zero.

Under f_1, $\log L_t$ drifts toward $\log L_B$ at less than half the speed at which it drifts toward $\log L_A$ under f_0.

Wald's approximation gives $E_{f_1}[N] \approx 8.3$, close to the simulated 9.1.

It gives $E_{f_0}[N] \approx 3.4$, far below the simulated 6.0.

Under f_0, draws of z near zero make $\log(f_0(z)/f_1(z))$ very large, so $\log L_t$ overshoots $\log L_A$ by about 2 on average; under f_1, $\log L_t$ overshoots $\log L_B$ by about 0.3.

The code below plots outcomes of constructing an unconditional distribution by simply pooling the simulated data across the two possible distributions f_0 and f_1.

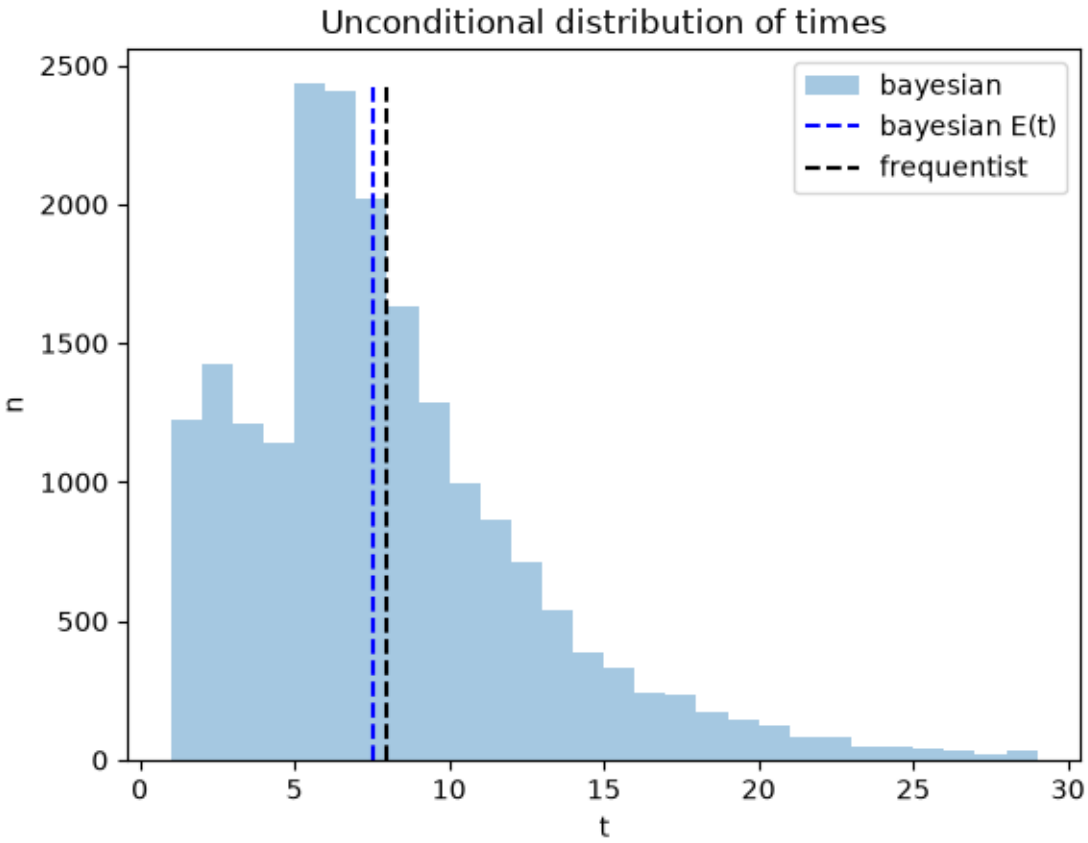
The pooled distribution describes a sense in which on average the Bayesian decides earlier, an outcome that seems at least partly to confirm the Navy Captain's hunch.

The pooled mean time to decide is 7.54 draws, against $t_{\text{optimal}} = 8$.

```
n = plt.hist(time_arr_u, bins=range(1, 30), alpha=0.4, label='bayesian')[0]
plt.vlines(np.mean(time_arr_u), 0, n.max(), linestyle='--',
           color='b', label='bayesian E(t)')
plt.vlines(t_optimal, 0, n.max(), linestyle='--', color='k', label='frequentist')
plt.legend()

plt.xlabel('t')
plt.ylabel('n')
plt.title('Unconditional distribution of times')

plt.show()
```



33.8 Probability of Making Correct Decision

Now we use simulations to compute the fractions of samples in which the Bayesian and the frequentist Neyman-Pearson decision rules decide correctly.

For the frequentist Neyman-Pearson rule, the probability of making the correct decision under f_1 is the optimal probability of detection given t that we defined earlier, and similarly it equals 1 minus the optimal probability of a false alarm under f_0.

Below we plot these two probabilities for the frequentist rule, along with the conditional probabilities that the Bayesian rule decides before t *and* that the decision is correct.

```
# optimal PFA and PD of frequentist with optimal sample size
V, PFA, PD = V_fre_t(t_optimal, L0_arr, L1_arr, π_star, wf)
```

```
plt.plot([1, 20], [PD, PD], linestyle='--', label='PD: fre. chooses f1 correctly')
plt.plot([1, 20], [1-PFA, 1-PFA], linestyle='--', label='1-PFA: fre. chooses f0␣
↪correctly')
plt.vlines(t_optimal, 0, 1, linestyle='--', label='frequentist optimal sample size')

T_arr = np.arange(1, 21)
plt.plot(T_arr, [np.mean(correctness0 & (time_arr0 <= t)) for t in T_arr],
         label='q=f0 and baye. choose f0')
plt.plot(T_arr, [np.mean(correctness1 & (time_arr1 <= t)) for t in T_arr],
         label='q=f1 and baye. choose f1')
```

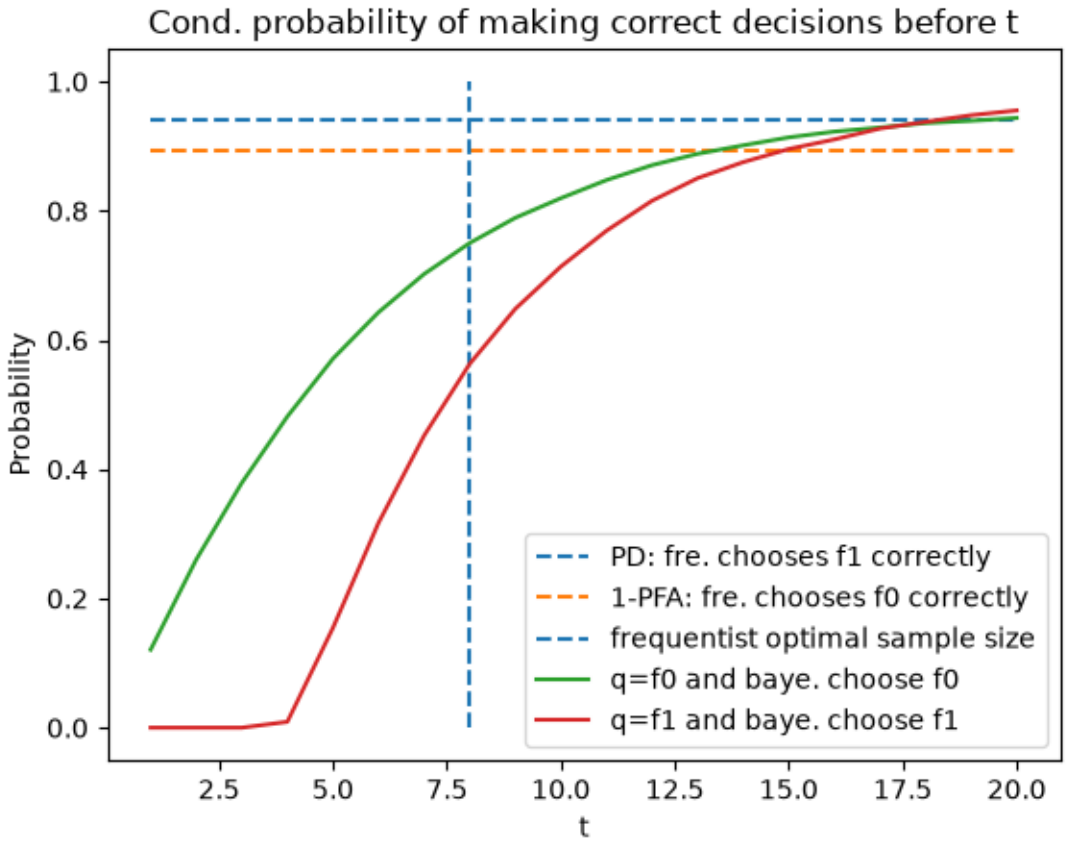
(continues on next page)

(continued from previous page)

```
plt.legend(loc=4)

plt.xlabel('t')
plt.ylabel('Probability')
plt.title('Cond. probability of making correct decisions before t')

plt.show()
```



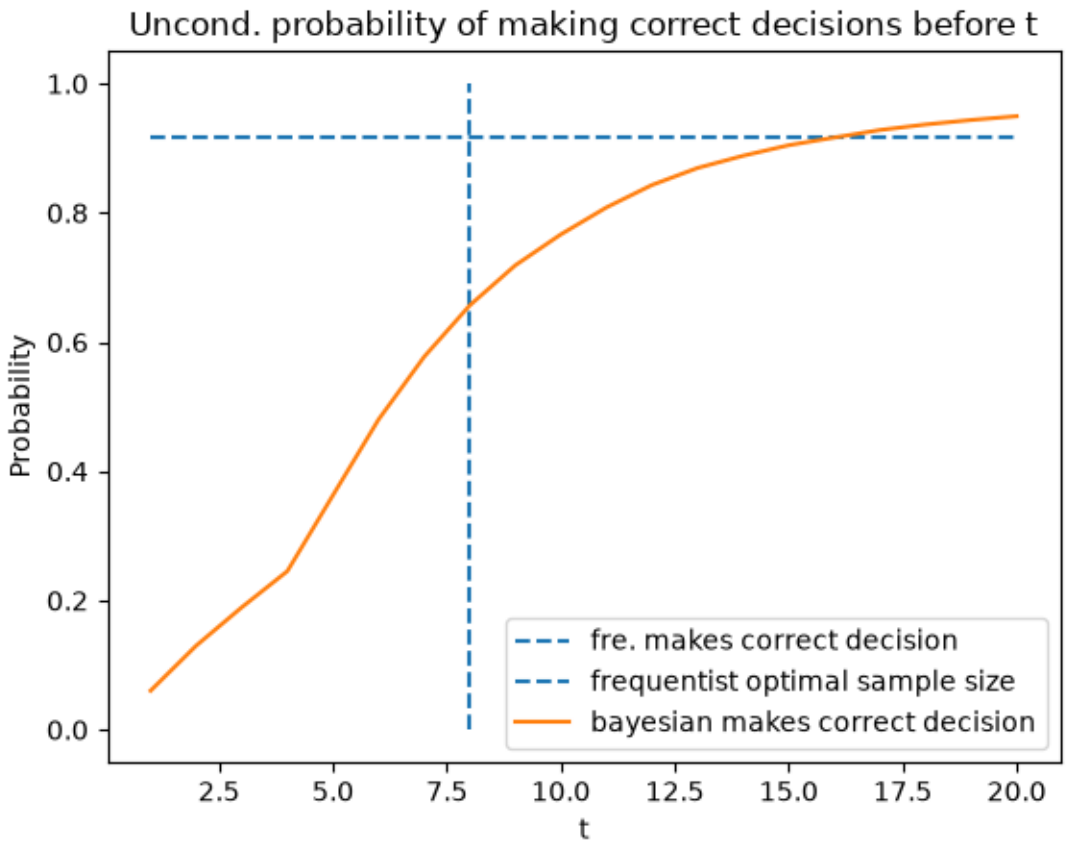
By averaging using π^*, we also plot the unconditional distribution.

```
plt.plot([1, 20], [(PD + 1 - PFA) / 2, (PD + 1 - PFA) / 2],
         linestyle='--', label='fre. makes correct decision')
plt.vlines(t_optimal, 0, 1, linestyle='--', label='frequentist optimal sample size')

plt.plot(T_arr, [np.mean(correctness_u & (time_arr_u <= t)) for t in T_arr],
         label="bayesian makes correct decision")
plt.legend()

plt.xlabel('t')
plt.ylabel('Probability')
plt.title('Uncond. probability of making correct decisions before t')

plt.show()
```



The simulated times to decide and probabilities of correct decisions give a second estimate of the Bayesian's expected loss.

This estimate uses the same simulated paths that we used for the frequentist rule.

```
V0_sim = wf.c * time_arr0.mean() + wf.L1 * (1 - correctness0.mean())
V1_sim = wf.c * time_arr1.mean() + wf.L0 * (1 - correctness1.mean())
V_baye_sim = π_star * V0_sim + (1 - π_star) * V1_sim

V_fre_sim = wf.c * t_optimal + V

print(f"Bayesian rule:     {V_baye_sim:.2f}")
print(f"fixed-sample rule: {V_fre_sim:.2f}")
```

```
Bayesian rule:     12.27
fixed-sample rule: 18.33
```

The simulated expected loss of the Bayesian rule, 12.27, exceeds the value 11.36 that `V_q` delivered.

The gap comes from Monte Carlo error in the 1,000 draws `wf.z0` and `wf.z1` that `V_q` uses to approximate expectations.

Both values lie far below the fixed-sample loss of 18.33.

33.9 Distribution of Likelihood Ratios at Neyman-Pearson's t

Next we use simulations to construct distributions of likelihood ratios after t draws.

To serve as useful reference points, we also show the likelihood ratios L_A and L_B that correspond to the Bayesian cutoffs A and B.

In order to exhibit the distribution more clearly, we report logarithms of likelihood ratios.

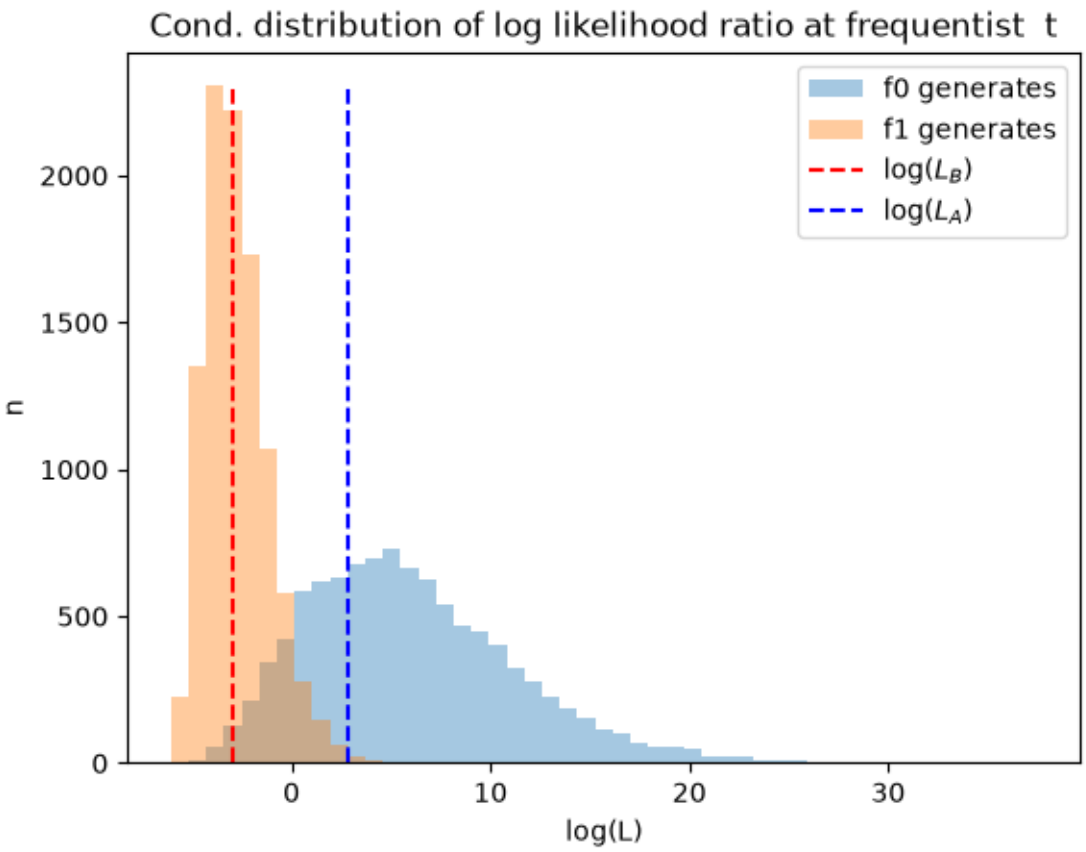
The graphs below reports two distributions, one conditional on f_0 generating the data, the other conditional on f_1 generating the data.

```
L_min = min(L0_arr[:, t_idx].min(), L1_arr[:, t_idx].min())
L_max = max(L0_arr[:, t_idx].max(), L1_arr[:, t_idx].max())
bin_range = np.linspace(np.log(L_min), np.log(L_max), 50)
n0 = plt.hist(np.log(L0_arr[:, t_idx]), bins=bin_range, alpha=0.4, label='f0 generates
 ↪')[0]
n1 = plt.hist(np.log(L1_arr[:, t_idx]), bins=bin_range, alpha=0.4, label='f1 generates
 ↪')[0]

plt.vlines(np.log(LB), 0, max(n0.max(), n1.max()), linestyle='--', color='r', label=
 ↪'log($L_B$)')
plt.vlines(np.log(LA), 0, max(n0.max(), n1.max()), linestyle='--', color='b', label=
 ↪'log($L_A$)')
plt.legend()

plt.xlabel('log(L)')
plt.ylabel('n')
plt.title('Cond. distribution of log likelihood ratio at frequentist  t')

plt.show()
```



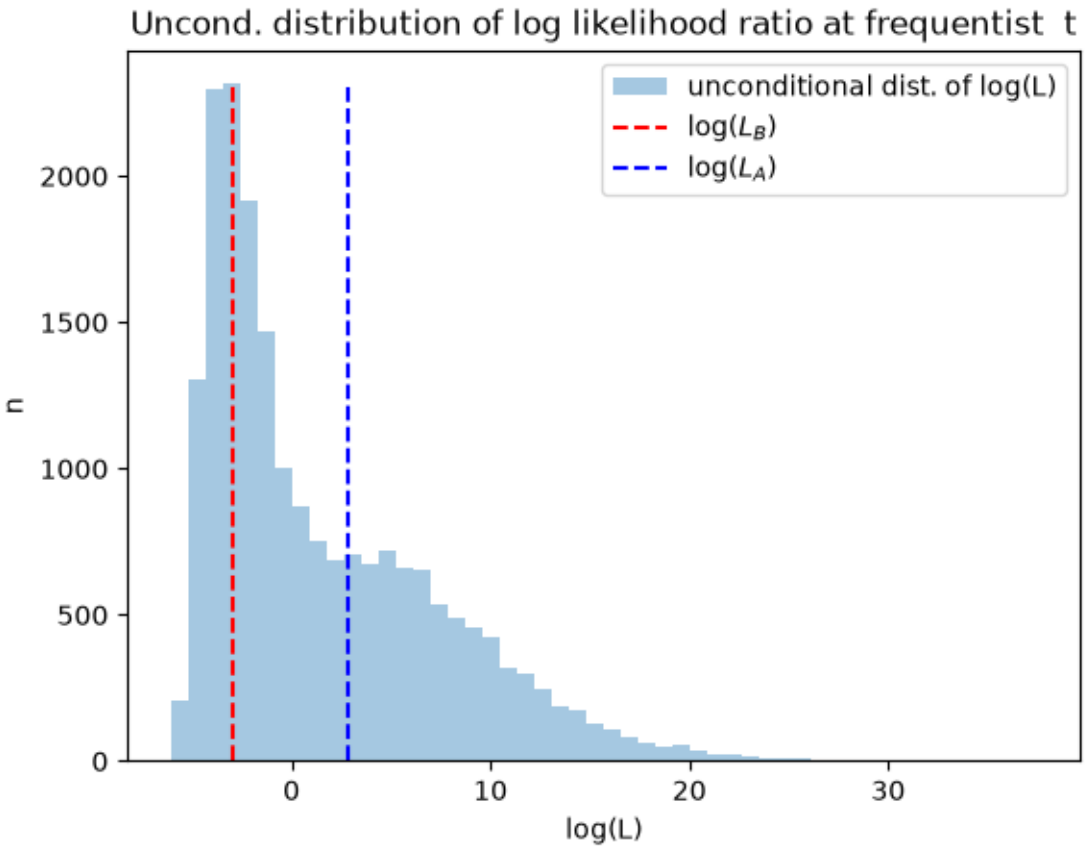
The next graph plots the unconditional distribution of log L at the frequentist t, constructed by pooling the two conditional distributions.

With $\pi^* = 0.5$, pooling equal numbers of paths from f_0 and f_1 weights them by π^* and $1 - \pi^*$.

```
n = plt.hist(np.log(np.concatenate([L0_arr[:, t_idx], L1_arr[:, t_idx]])),
             bins=50, alpha=0.4, label='unconditional dist. of log(L)')[0]
plt.vlines(np.log(LB), 0, n.max(), linestyle='--', color='r', label='log($L_B$)')
plt.vlines(np.log(LA), 0, n.max(), linestyle='--', color='b', label='log($L_A$)')
plt.legend()

plt.xlabel('log(L)')
plt.ylabel('n')
plt.title('Uncond. distribution of log likelihood ratio at frequentist  t')

plt.show()
```



33.10 Exercises

Exercise 33.10.1

Fix a sample size t and regard the expected loss (33.1) as a function of the cutoff d.

1. Show that $d^* = (1-\pi^*)\,\bar{L}_0 / \left(\pi^* \bar{L}_1\right)$ minimizes (33.1) for every t. Hint: write PFA and $1-PD$ as integrals over samples z^t and minimize the integrand at each z^t.
2. Let π_t be the posterior probability of f_0 after t draws from a prior π^*. Show that the rule "decide f_1 if $L(z^t) < d^*$" chooses the decision with the smaller posterior expected loss. Explain why this makes the frequentist rule of this lecture a Bayes rule with a sample size fixed in advance.
3. For the 20 values of π^* in `π_star_arr` and $t = 1, \ldots, 20$, compare the simulated loss `V_fre_d_t` at d^* with its minimum over a fine grid of values of d. Do the same for the loss attained by a Nelder-Mead search over d that starts at $d = 1$.

Solution

Write $f_i(z^t) = f_i(z_1) \cdots f_i(z_t)$.

The rule rejects H_0 on the set $R_d = \{z^t : L(z^t) < d\}$.

Then

$$PFA = \int_{R_d} f_0(z^t)\, dz^t, \qquad 1 - PD = \int_{R_d^c} f_1(z^t)\, dz^t$$

and the loss terms of (33.1) that depend on d equal

$$\int \left[\mathbf{1}\{z^t \in R_d\}\, \pi^* \bar{L}_1 f_0(z^t) + \mathbf{1}\{z^t \notin R_d\}\, (1-\pi^*) \bar{L}_0 f_1(z^t) \right] dz^t.$$

The integrand is smallest at each z^t if z^t belongs to R_d whenever $\pi^* \bar{L}_1 f_0(z^t) < (1-\pi^*) \bar{L}_0 f_1(z^t)$ and to its complement otherwise.

Dividing by $\pi^* \bar{L}_1 f_1(z^t)$ turns this condition into $L(z^t) < d^*$.

The set R_{d^*} minimizes the integrand at every z^t, so no other rejection region does better for any t.

Bayes' law gives the posterior

$$\pi_t = \frac{\pi^* L(z^t)}{\pi^* L(z^t) + 1 - \pi^*}.$$

Deciding f_1 has posterior expected loss $\pi_t \bar{L}_1$; deciding f_0 has posterior expected loss $(1-\pi_t)\bar{L}_0$.

The inequality $\pi_t \bar{L}_1 < (1-\pi_t)\bar{L}_0$ is equivalent to $\pi^* L(z^t) \bar{L}_1 < (1-\pi^*)\bar{L}_0$, which is $L(z^t) < d^*$.

After t draws, the frequentist rule with $d = d^*$ therefore takes the action that a Bayesian with prior π^* would take.

Minimizing $\bar{V}_{fre}(t, d^*)$ over t then chooses the best sample size before any data arrive.

The Bayesian sequential rule differs only in letting the data decide when to stop.

Here is the numerical check.

```
from scipy.optimize import minimize

d_grid = np.exp(np.linspace(-10, 10, 4001))

def V_grid_min(t, π_s):
    "Minimum over d_grid of the simulated loss V_fre_d_t."
    n = L0_arr.shape[0]
    PFA = np.searchsorted(np.sort(L0_arr[:, t-1]), d_grid) / n    # Pr{L < d | f0}
    PD = np.searchsorted(np.sort(L1_arr[:, t-1]), d_grid) / n     # Pr{L < d | f1}
    return np.min(π_s * PFA * wf.L1 + (1 - π_s) * (1 - PD) * wf.L0)

gap_closed, gap_nm = [], []
for π_s in π_star_arr:
    for t in range(1, 21):
        V_min = V_grid_min(t, π_s)
        V_closed = V_fre_d_t(d_star(π_s, wf), t, L0_arr, L1_arr, π_s, wf)
        res = minimize(V_fre_d_t, 1, args=(t, L0_arr, L1_arr, π_s, wf),
                       method='Nelder-Mead')
        gap_closed.append(V_closed - V_min)
        gap_nm.append(res.fun - V_min)

gap_closed, gap_nm = np.array(gap_closed), np.array(gap_nm)
print(f"closed form: largest excess over grid minimum {gap_closed.max():.3f}")
print(f"Nelder-Mead: largest excess over grid minimum {gap_nm.max():.3f}, "
      f"cases with excess > 0.5: {np.sum(gap_nm > 0.5)} of {gap_nm.size}")
```

```
closed form: largest excess over grid minimum 0.245
Nelder-Mead: largest excess over grid minimum 0.685, cases with excess > 0.5: 2␣
↪of 400
```

The closed form comes within 0.25 of the grid minimum in all 400 cases.

The grid minimum lies below the loss at d^* because it fits d to the simulated sample; d^* minimizes the population loss.

Nelder-Mead treats the step function `V_fre_d_t` as if it were smooth and can stop on a flat step.

Here it stops more than 0.5 above the grid minimum in 2 of the 400 cases and 0.69 above it in the worst case.

Exercise 33.10.2

A frequentist can test the Navy Captain's hunch without a prior.

Set $\pi^* = 0.5$ and take the error probabilities of the best fixed-sample rule: $\alpha = PFA$ and $\beta = 1 - PD$ at t_{optimal}.

Wald's sequential probability ratio test with targets (α, β) uses the thresholds $a(\alpha, \beta) = (1-\beta)/\alpha$ and $b(\alpha, \beta) = \beta/(1-\alpha)$ of *A Problem that Stumped Milton Friedman*.

That lecture applies them to the likelihood ratio $f_1/f_0 = 1/L$.

In terms of our $L_t = L(z^t)$, the test is

- accept f_0 if $L_t \geq A_L \equiv 1/b(\alpha, \beta) = (1-\alpha)/\beta$
- accept f_1 if $L_t \leq B_L \equiv 1/a(\alpha, \beta) = \alpha/(1-\beta)$
- draw another z if $B_L < L_t < A_L$

Simulate the test on the paths `L0_arr` and `L1_arr`.

Report its realized error probabilities, $E_{f_0}[N]$, and $E_{f_1}[N]$, and compare them with t_{optimal}.

Solution

With $\pi_0 = 1/2$, $\pi_t = L_t/(1+L_t)$.

Wald's test is therefore the rule in `check_results` with $\pi_0 = 1/2$ and probability cutoffs $A_L/(1+A_L)$ and $B_L/(1+B_L)$.

```
V_t, PFA_t, PD_t = V_fre_t(t_optimal, L0_arr, L1_arr, 0.5, wf)
α, β = PFA_t, 1 - PD_t

A_L = (1 - α) / β
B_L = α / (1 - β)

n0, accept0 = check_results(L0_arr, A_L / (1 + A_L), B_L / (1 + B_L), 0, 0.5)
n1, accept1 = check_results(L1_arr, A_L / (1 + A_L), B_L / (1 + B_L), 1, 0.5)

print(f"targets:   α = {α:.3f}, β = {β:.3f}")
print(f"realized:  α = {1 - accept0.mean():.3f}, β = {1 - accept1.mean():.3f}")
print(f"E_f0[N] = {n0.mean():.2f},  E_f1[N] = {n1.mean():.2f},  "
      f"t_optimal = {t_optimal}")
print("paths with no decision by T:", np.sum(n0 > T) + np.sum(n1 > T))
```

```
targets:   α = 0.108, β = 0.059
realized:  α = 0.087, β = 0.019
E_f0[N] = 5.42,  E_f1[N] = 6.84,  t_optimal = 8
paths with no decision by T: 0
```

Wald's test meets both targets.

Its realized error probabilities, 0.087 and 0.019, lie below $\alpha = 0.108$ and $\beta = 0.059$.

It draws 5.42 observations on average under f_0 and 6.84 under f_1, against $t_{\text{optimal}} = 8$ for the fixed-sample rule.

The sequential test has smaller error probabilities and fewer expected draws under both hypotheses, so it dominates the Navy's rule without any appeal to a prior.

Wald and Wolfowitz [1948] proved that no test whose error probabilities are no larger than those of a sequential probability ratio test has a smaller expected number of draws under either hypothesis.

Exercise 33.10.3

This exercise measures the value of sequential sampling.

Evaluate both rules on the simulated paths `L0_arr` and `L1_arr`, so that both use the same random numbers.

1. Set $\pi^* = 0.5$. For a grid of sampling costs c between 0.25 and 14, compute $\bar{V}_{fre}$ and $\bar{V}_{Bayes}$ and plot $\bar{V}_{fre} - \bar{V}_{Bayes}$ against c. For each c, solve the Bayesian's problem again to get new cutoffs A and B. Let the fixed-sample rule choose $t = 0$, which means deciding without data at expected loss $\min\{\pi^* \bar{L}_1, (1-\pi^*)\bar{L}_0\}$.
2. Set $c = 1.25$ and $\pi^* = 0.5$, and let the Bayesian start from a prior $\pi_0 \neq \pi^*$. Plot $\bar{V}_{fre} - \bar{V}_{Bayes}(\pi_0)$ against π_0. Do the same when the fixed-sample rule chooses (t, d) with π_0 in place of π^*.

Solution

The function below computes the expected loss of the Bayesian rule with cutoffs (B, A) and prior π_0 when nature selects f_0 with probability π^*.

A Bayesian whose prior lies outside (B, A) decides without data.

```
def V_baye_sim(A, B, π0, π_star, wf):

    if π0 >= A:
        return (1 - π_star) * wf.L0
    if π0 <= B:
        return π_star * wf.L1

    t0, c0 = check_results(L0_arr, A, B, 0, π0)
    t1, c1 = check_results(L1_arr, A, B, 1, π0)
    assert max(t0.max(), t1.max()) <= T      # every path decides by T

    V0 = wf.c * t0.mean() + wf.L1 * (1 - c0.mean())
    V1 = wf.c * t1.mean() + wf.L0 * (1 - c1.mean())

    return π_star * V0 + (1 - π_star) * V1
```

The misclassification losses of the fixed-sample rule do not depend on c, so we compute them once.

```
π_star = 0.5
t_grid = np.arange(1, T+1)
loss_t = np.array([V_fre_t(t, L0_arr, L1_arr, π_star, wf)[0] for t in t_grid])
loss_0 = min(π_star * wf.L1, (1 - π_star) * wf.L0)    # decide with t = 0

c_grid = np.array([0.25, 0.5, 1, 1.5, 2, 2.5, 3, 4, 6, 8, 10, 14])
V_fre_c = np.empty_like(c_grid)
```

```
V_baye_c = np.empty_like(c_grid)

for i, c in enumerate(c_grid):
    seed_numba(1234)
    wf_c = WaldFriedman(c=c)
    B_c, A_c = find_cutoff_rule(wf_c, solve_model(wf_c))

    V_fre_c[i] = min(loss_0, (c * t_grid + loss_t).min())
    V_baye_c[i] = V_baye_sim(A_c, B_c, π_star, π_star, wf_c)

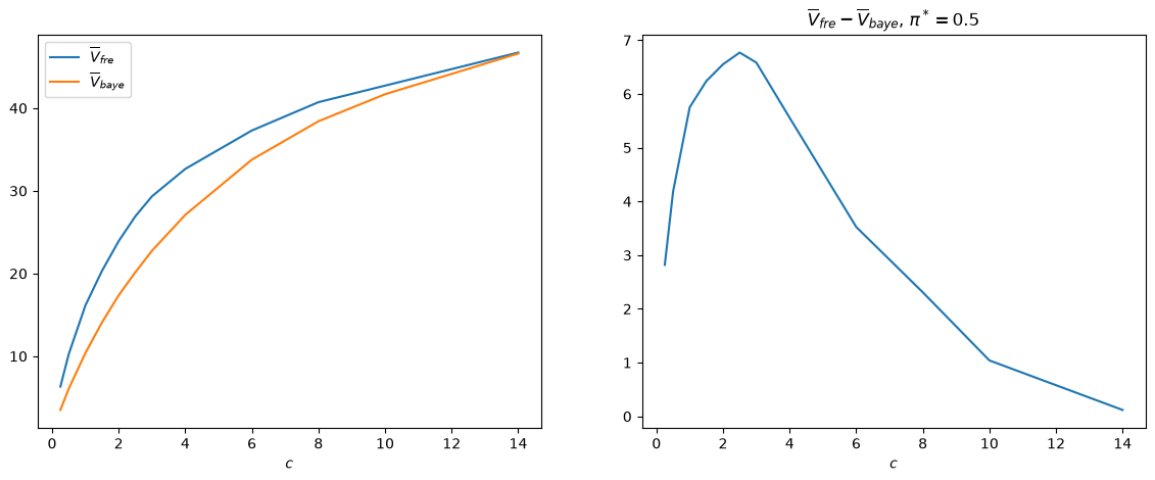
fig, axs = plt.subplots(1, 2, figsize=(14, 5))

axs[0].plot(c_grid, V_fre_c, label=r'$\overline{V}_{fre}$')
axs[0].plot(c_grid, V_baye_c, label=r'$\overline{V}_{baye}$')
axs[0].set_xlabel('$c$')
axs[0].legend()

axs[1].plot(c_grid, V_fre_c - V_baye_c)
axs[1].set_xlabel('$c$')
axs[1].set_title(r'$\overline{V}_{fre} - \overline{V}_{baye}$, $\pi^*=0.5$')

plt.show()

print(np.round(V_fre_c - V_baye_c, 2))
```




```
[2.82 4.18 5.75 6.24 6.55 6.77 6.58 5.56 3.52 2.31 1.04 0.12]
```

The gap rises from 2.82 at $c = 0.25$ to 6.77 at $c = 2.5$ and falls to 0.12 at $c = 14$.

When draws are cheap, both rules take many draws and rarely err.

When draws are expensive, both rules take about one draw.

At intermediate costs, the sequential rule stops early after informative draws, while the fixed-sample rule pays for draws that it does not need.

For the second part, `compute_V_fre` with π_0 in place of π^* gives the fixed-sample rule that a decision maker with prior π_0 would choose.

We evaluate that rule's error probabilities with the weights π^* and $1 - \pi^*$.

```
V_fre_arr, _, _ = compute_V_fre(L0_arr, L1_arr, π_star, wf)
V_fre_bar = V_fre_arr.min()

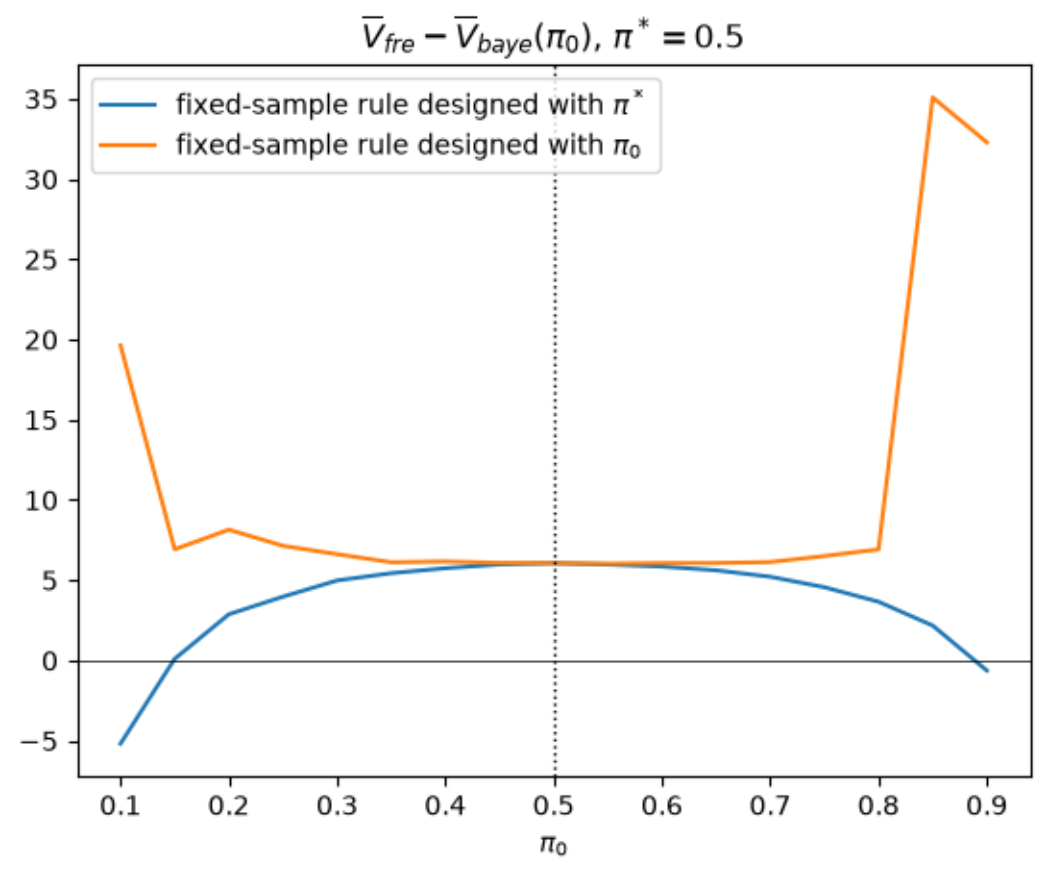
π0_grid = np.linspace(0.1, 0.9, 17)
V_baye_π0 = np.empty_like(π0_grid)
V_fre_π0 = np.empty_like(π0_grid)

for i, π0 in enumerate(π0_grid):
    V_baye_π0[i] = V_baye_sim(A, B, π0, π_star, wf)

    V_arr, PFA_arr, PD_arr = compute_V_fre(L0_arr, L1_arr, π0, wf)
    k = np.argmin(V_arr)
    V_fre_π0[i] = (wf.c * (k + 1) + π_star * PFA_arr[k] * wf.L1
                   + (1 - π_star) * (1 - PD_arr[k]) * wf.L0)

plt.plot(π0_grid, V_fre_bar - V_baye_π0,
         label=r'fixed-sample rule designed with $\pi^*$')
plt.plot(π0_grid, V_fre_π0 - V_baye_π0,
         label=r'fixed-sample rule designed with $\pi_0$')
plt.axhline(0, color='k', lw=0.5)
plt.axvline(π_star, color='k', ls=':', lw=1)
plt.xlabel(r'$\pi_0$')
plt.title(r'$\overline{V}_{fre} - \overline{V}_{baye}(\pi_0)$, $\pi^*=0.5$')
plt.legend()
plt.show()

print(np.round(V_fre_bar - V_baye_π0, 2))
print(np.round(V_fre_π0 - V_baye_π0, 2))
```




```
[-5.18  0.1   2.88  3.98  4.98  5.44  5.75  5.99  6.06  5.98  5.86  5.62
  5.22  4.56  3.66  2.17 -0.63]
[19.64  6.93  8.15  7.15  6.62  6.12  6.17  6.08  6.06  6.03  6.07  6.07
  6.13  6.51  6.92 35.09 32.29]
```

A Bayesian with a wrong prior still beats the correctly designed fixed-sample rule for $0.15 \leq \pi_0 \leq 0.85$.

At $\pi_0 = 0.1$ and $\pi_0 = 0.9$, the Bayesian starts close to a cutoff and often accepts the wrong distribution.

When the fixed-sample rule uses the same wrong prior, the sequential rule wins at every π_0 on the grid.

CHAPTER

THIRTYFOUR

MERGING OF OPINIONS: THE BLACKWELL–DUBINS THEOREM

Contents

- *Merging of Opinions: The Blackwell–Dubins Theorem*
  - *Overview*
  - *Probability measures on sequence spaces*
  - *The likelihood-ratio martingale*
  - *The Blackwell–Dubins theorem*
  - *The Beta–Bernoulli model*
  - *Failure of merging: mutual singularity*
  - *Kakutani's theorem: when does merging hold?*
  - *Extension to continuous time*
  - *Applications*
  - *The rate of merging*
  - *Summary and extensions*
  - *Exercises*

34.1 Overview

This lecture studies the merging-of-opinions theorem of Blackwell and Dubins [1962].

The theorem asks a simple question:

> If two agents hold different prior beliefs about a stochastic process but observe the same stream of data indefinitely, will their probability assessments eventually converge?

The answer is yes under an absolute-continuity condition.

If $Q \ll P$ (that is, P dominates Q), then the conditional distributions under P and Q over the entire future path merge in total variation, Q-almost surely.

If in addition $P \ll Q$ (so that $P \sim Q$), then the same conclusion holds under both agents' probabilities.

This result connects to several other ideas:

- Bayesian consistency: posterior predictions approach the truth when the prior lies in the right absolute-continuity class (*Likelihood Ratio Processes and Bayesian Learning*).
- Agreement results: Aumann [1976] studies agents with a common prior and private information; Blackwell–Dubins studies agents with different priors and common information.
- Kakutani's dichotomy: for product measures, equivalence versus singularity can be read from a Hellinger criterion.

We develop the theory in discrete time and then sketch the continuous-time analogue.

Throughout, we use the Beta–Bernoulli model as a running example.

Two agents observe the same stream of coin flips but start from different priors over the coin's bias.

Let us start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.stats import beta as beta_dist, binom, betabinom, norm
from scipy.special import betaln, beta as beta_fn, polygamma
```

34.2 Probability measures on sequence spaces

34.2.1 The sequence space and its filtration

Let $(S, \mathcal{S})$ be a standard Borel space (i.e., a measurable space isomorphic to a Borel subset of a complete separable metric space), called the signal space.

The standard Borel assumption guarantees the existence of regular conditional distributions, which the theorem requires.

Set $\Omega = S^{\mathbb{N}}$, the set of all infinite sequences $\omega = (x_1, x_2, \ldots)$ with $x_n \in S$, equipped with the product σ-algebra $\mathcal{F} = \mathcal{S}^{\otimes \mathbb{N}}$.

For each $n \geq 1$, define the *finite-horizon* σ-algebra

$$\mathcal{F}_n = \sigma(x_1, \ldots, x_n),$$

and set $\mathcal{F}_0 = \{\emptyset, \Omega\}$, so $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \cdots \subseteq \mathcal{F}$.

Define $\mathcal{F}_\infty = \sigma\big(\bigcup_{n \geq 1} \mathcal{F}_n\big)$, the σ-algebra of events that can eventually be learned.

Cylinder sets $\{x_1 \in A_1, \ldots, x_n \in A_n\}$ generate the product σ-algebra and belong to $\bigcup_n \mathcal{F}_n$, so $\mathcal{F}_\infty = \mathcal{F}$.

The σ-algebra $\mathcal{F}_\infty$ differs from the **tail σ-algebra** $\mathcal{T} = \bigcap_{n \geq 1} \sigma(x_n, x_{n+1}, \ldots)$, whose events, such as $\{x_n \in A \text{ for infinitely many } n\}$, do not depend on any finite stretch of data.

The collection $\{\mathcal{F}_n\}_{n \geq 1}$ is the **natural filtration** generated by the observation process; $\mathcal{F}_n$ encodes everything that can be learned from the first n data points.

Let P and Q denote two probability measures on $(\Omega, \mathcal{F})$.

Write $P_n = P|_{\mathcal{F}_n}$ and $Q_n = Q|_{\mathcal{F}_n}$ for their restrictions to the history up to time n.

34.2.2 Absolute continuity

Definition 34.2.1 (Absolute Continuity)

P is **absolutely continuous** with respect to Q, written $P \ll Q$, if $Q(A) = 0$ implies $P(A) = 0$ for every $A \in \mathcal{F}$.

They are **mutually absolutely continuous**, or **equivalent**, written $P \sim Q$, if both $P \ll Q$ and $Q \ll P$.

P is **locally absolutely continuous** with respect to Q if $P_n \ll Q_n$ for every $n \geq 1$.

Global absolute continuity $P \ll Q$ implies local absolute continuity, but not conversely.

Mutual absolute continuity means the two agents agree on which events are *possible*.

They can disagree about probabilities, but neither agent rules out an event the other deems possible.

34.2.3 Total variation distance

Definition 34.2.2 (Total Variation Distance)

For two probability measures μ and ν on $(E, \mathcal{E})$,

$$\|\mu - \nu\|_{\text{TV}} = \sup_{A \in \mathcal{E}} |\mu(A) - \nu(A)| = \frac{1}{2} \int_E \left| \frac{d\mu}{d\lambda} - \frac{d\nu}{d\lambda} \right| d\lambda,$$

where λ is any **dominating measure**, meaning $\mu \ll \lambda$ and $\nu \ll \lambda$ (for example, $\lambda = \mu + \nu$).

The distance $\|\mu - \nu\|_{\text{TV}}$ lies in $[0, 1]$; it equals 0 if and only if $\mu = \nu$ and 1 if and only if $\mu \perp \nu$ (mutual singularity).

When $\mu \ll \nu$ with $f = d\mu/d\nu$,

$$\|\mu - \nu\|_{\text{TV}} = \mathbb{E}_\nu[(f-1)^+] = 1 - \mathbb{E}_\nu[\min(f, 1)].$$

Exercise 34.2.1

Show the identity above.

Hint: Start from $\|\mu - \nu\|_{\text{TV}} = \frac{1}{2}\,\mathbb{E}_\nu[|f-1|]$ (which follows from taking ν as the dominating measure) and use the fact that $\mathbb{E}_\nu[f] = 1$.

Solution

Since $\mu \ll \nu$, we can use ν as the dominating measure, so $d\mu/d\nu = f$ and $d\nu/d\nu = 1$, giving

$$\|\mu - \nu\|_{\text{TV}} = \tfrac{1}{2}\,\mathbb{E}_\nu[|f-1|].$$

Write $|f-1| = (f-1)^+ + (1-f)^+$.

Since μ is a probability measure, $\mathbb{E}_\nu[f] = 1$, so the two parts contribute equally: $\mathbb{E}_\nu[(f-1)^+] = \mathbb{E}_\nu[(1-f)^+]$.

Therefore $\frac{1}{2}\,\mathbb{E}_\nu[|f-1|] = \mathbb{E}_\nu[(f-1)^+]$.

Next, note that $(f-1)^+ = f - \min(f,1)$, so $\mathbb{E}_\nu[(f-1)^+] = \mathbb{E}_\nu[f] - \mathbb{E}_\nu[\min(f,1)] = 1 - \mathbb{E}_\nu[\min(f,1)]$.

Total variation is one of the strongest standard notions of distance between probability measures.

If two measures are close in total variation, then their probabilities of every event are close.

34.2.4 The merging question

The Blackwell–Dubins theorem studies the conditional distribution of the *future* given the *past*.

At time n, after observing $(x_1, \dots, x_n)$, each agent forms a conditional distribution over all future events:

$$P(\cdot \mid \mathcal{F}_n)(\omega), \qquad Q(\cdot \mid \mathcal{F}_n)(\omega).$$

These are probability measures on the whole future path, not just the next observation.

The merging question asks whether

$$d_n \;:=\; \big\| P(\cdot \mid \mathcal{F}_n) - Q(\cdot \mid \mathcal{F}_n) \big\|_{\mathrm{TV}} \;\longrightarrow\; 0$$

almost surely as $n \to \infty$.

At $n = 0$ the conditioning is trivial and $d_0 = \|P - Q\|_{\mathrm{TV}}$.

34.3 The likelihood-ratio martingale

Our main tool is the Radon–Nikodym derivative process.

34.3.1 The likelihood ratio

Since $Q \ll P$ implies $Q_n \ll P_n$ for every n, the Radon–Nikodym theorem guarantees the existence of the likelihood ratio

$$Z_n = \frac{dQ_n}{dP_n}, \qquad Z_n \geq 0 \;\; P\text{-a.s.}, \qquad \mathbb{E}_P[Z_n] = 1.$$

The key structural property is that global absolute continuity $Q \ll P$ implies the existence of an overall Radon–Nikodym derivative $Z = dQ/dP$ on all of $(\Omega, \mathcal{F})$, and

$$Z_n = \mathbb{E}_P[Z \mid \mathcal{F}_n] \qquad P\text{-a.s.}$$

That is, $\{Z_n, \mathcal{F}_n\}_{n \geq 0}$ is a non-negative, uniformly integrable P-martingale with $Z_0 = 1$.

Throughout the lecture, Z and Z_n denote densities of Q with respect to P.

Posterior probabilities are martingales of the same kind; *Blackwell's Theorem on Comparing Experiments* uses that property to compare experiments.

Lemma 34.3.1 (Martingale Convergence)

The likelihood-ratio process $\{Z_n\}$ satisfies:

1. $Z_n \to Z_\infty$ P-almost surely as $n \to \infty$.
2. $Z_\infty = \mathbb{E}_P[Z \mid \mathcal{F}_\infty] = Z$ P-a.s.
3. $Z_n \to Z_\infty$ in $L^1(P)$: $\mathbb{E}_P[|Z_n - Z_\infty|] \to 0$.

Proof sketch. Non-negativity and the martingale property give boundedness in $L^1(P)$.

Then almost-sure convergence follows from Doob's martingale convergence theorem [Doob, 1953].

Uniform integrability (which follows from $Z \in L^1(P)$ via the conditional Jensen inequality) upgrades this to $L^1(P)$ convergence.

Lévy's upward theorem identifies the limit as $\mathbb{E}_P[Z \mid \mathcal{F}_\infty]$, which equals Z because $\mathcal{F}_\infty = \mathcal{F}$. □

From here on we write Z_∞ for $Z = dQ/dP$.

34.3.2 Connecting conditional measures to the likelihood ratio

The following identity connects the likelihood ratio to the conditional distributions.

On the set $\{Z_n > 0\}$, the Radon–Nikodym derivative of $Q(\cdot \mid \mathcal{F}_n)$ with respect to $P(\cdot \mid \mathcal{F}_n)$ is

$$\frac{d\,Q(\cdot \mid \mathcal{F}_n)}{d\,P(\cdot \mid \mathcal{F}_n)} = \frac{Z_\infty}{Z_n} \qquad P\text{-a.s. on } \{Z_n > 0\}.$$

This is the abstract Bayes formula $Q(A \mid \mathcal{F}_n) = \mathbb{E}_P[Z_\infty \mathbf{1}_A \mid \mathcal{F}_n] / \mathbb{E}_P[Z_\infty \mid \mathcal{F}_n]$.

The set $\{Z_n > 0\}$ has Q-probability one, because $Q(Z_n = 0) = \mathbb{E}_P[Z_n \mathbf{1}\{Z_n = 0\}] = 0$.

Applying the total-variation formula with $f = Z_\infty / Z_n$ then gives

$$d_n = \mathbb{E}_{P(\cdot \mid \mathcal{F}_n)}\left[\left(\frac{Z_\infty}{Z_n} - 1\right)^+\right] = 1 - \mathbb{E}_{P(\cdot \mid \mathcal{F}_n)}\left[\min\left(\frac{Z_\infty}{Z_n},\, 1\right)\right].$$

Because Z_n is $\mathcal{F}_n$-measurable, the same quantity is a ratio of conditional expectations:

$$d_n = \frac{\mathbb{E}_P[(Z_\infty - Z_n)^+ \mid \mathcal{F}_n]}{Z_n} = \frac{\mathbb{E}_P[|Z_\infty - Z_n| \mid \mathcal{F}_n]}{2 Z_n} \qquad Q\text{-a.s.} \tag{34.1}$$

The second equality holds because $\mathbb{E}_P[Z_\infty - Z_n \mid \mathcal{F}_n] = 0$, so the positive and negative parts of $Z_\infty - Z_n$ have equal conditional means.

Multiplying through by Z_n and taking the P-expectation (then using $\mathbb{E}_P[Z_n\, g] = \mathbb{E}_Q[g]$ for non-negative $\mathcal{F}_n$-measurable g):

$$2\,\mathbb{E}_Q[d_n] \;=\; \mathbb{E}_P[|Z_\infty - Z_n|].$$

So the $L^1(P)$ convergence of the martingale gives $\mathbb{E}_Q[d_n] \to 0$, and its rate governs how fast $\mathbb{E}_Q[d_n]$ goes to zero.

34.4 The Blackwell–Dubins theorem

Theorem 34.4.1 (Blackwell–Dubins (1962))

Let P and Q be probability measures on $(\Omega, \mathcal{F})$ with $Q \ll P$.

Define

$$d_n = \left\| P(\cdot \mid \mathcal{F}_n) - Q(\cdot \mid \mathcal{F}_n) \right\|_{\text{TV}}.$$

Then $d_n \to 0$ Q-almost surely.

The identity $\mathbb{E}_Q[d_n] = \frac{1}{2}\,\mathbb{E}_P[|Z_\infty - Z_n|] \to 0$ gives convergence in Q-probability.

The proof of almost-sure convergence has three steps.

Step 1. A formula for d_n.

Equation (34.1) writes d_n as the ratio of $\mathbb{E}_P[|Z_\infty - Z_n| \,|\, \mathcal{F}_n]$ to $2Z_n$, Q-a.s.

Step 2. The denominator converges to a positive limit.

By *Lemma 34.3.1*, $Z_n \to Z_\infty$ P-a.s., hence Q-a.s. because $Q \ll P$.

The limit is positive Q-a.s., because $Q(Z_\infty = 0) = \mathbb{E}_P[Z_\infty \mathbf{1}\{Z_\infty = 0\}] = 0$.

Step 3. The numerator converges to zero.

Fix $M > 0$ and split $Z_\infty = U + V$, where $U = \min(Z_\infty, M)$ and $V = (Z_\infty - M)^+$.

Let $U_n = \mathbb{E}_P[U \mid \mathcal{F}_n]$ and $V_n = \mathbb{E}_P[V \mid \mathcal{F}_n]$, so $Z_n = U_n + V_n$.

The triangle inequality gives $|Z_\infty - Z_n| \le |U - U_n| + V + V_n$, and conditioning on $\mathcal{F}_n$ gives

$$\mathbb{E}_P[|Z_\infty - Z_n| \,|\, \mathcal{F}_n] \le \mathbb{E}_P[|U - U_n| \,|\, \mathcal{F}_n] + 2V_n.$$

Lévy's upward theorem gives $V_n \to V$ P-a.s.

For the bounded part, fix m and let $W_m = \sup_{k \ge m} |U - U_k| \le 2M$.

For $n \ge m$, $\mathbb{E}_P[|U - U_n| \,|\, \mathcal{F}_n] \le \mathbb{E}_P[W_m \mid \mathcal{F}_n]$, and Lévy's upward theorem sends the right side to W_m as $n \to \infty$.

Lévy's upward theorem also gives $U_k \to U$ P-a.s., so $W_m \to 0$ as $m \to \infty$.

Combining these limits,

$$\limsup_{n \to \infty} \mathbb{E}_P[|Z_\infty - Z_n| \,|\, \mathcal{F}_n] \le 2\,(Z_\infty - M)^+ \qquad P\text{-a.s.}$$

for every M.

Letting $M \to \infty$ through the integers sends the right side to zero, because Z_∞ is finite.

So the numerator converges to zero P-a.s., hence Q-a.s.

Steps 1–3 give $d_n \to 0$ Q-a.s. □

The argument for the bounded part U is Hunt's lemma.

The proof uses no monotonicity of $\{d_n\}$; a *later section* shows that $\{d_n\}$ need not be a supermartingale.

Remark 34.4.1 (One-Sided vs. Mutual Absolute Continuity)

The theorem requires only $Q \ll P$, not $P \ll Q$.

One-sided absolute continuity $Q \ll P$ gives merging Q-almost surely.

Since $Q \ll P$ means every P-null set is also Q-null, Q-a.s. convergence does *not* automatically imply P-a.s. convergence.

To conclude that $d_n \to 0$ under *both* agents' measures, one needs mutual absolute continuity $P \sim Q$.

With $P \ll Q$ added, the proof can be run with the roles of P and Q swapped, yielding $d_n \to 0$ P-a.s. as well.

Remark 34.4.2 (Sharpness)

Absolute continuity matters.

When P and Q are singular, merging can fail completely.

Suppose $P \perp Q$ and $P_n \sim Q_n$ for every n.

Take $S \in \mathcal{F}$ with $P(S) = 1$ and $Q(S) = 0$.

Then $P(S \mid \mathcal{F}_n) = 1$ P-a.s. and $Q(S \mid \mathcal{F}_n) = 0$ Q-a.s.

The events $\{P(S \mid \mathcal{F}_n) = 1\}$ and $\{Q(S \mid \mathcal{F}_n) = 0\}$ belong to $\mathcal{F}_n$, so local equivalence makes each hold both P-a.s. and Q-a.s.

Hence $d_n \geq |P(S \mid \mathcal{F}_n) - Q(S \mid \mathcal{F}_n)| = 1$ for every n, under both measures.

The point-mass example below has $d_n = 1$ for every n.

For product measures, Kakutani's theorem later gives a sharp equivalence-versus-singularity dichotomy.

34.5 The Beta–Bernoulli model

Before turning to Python, we introduce the main example used throughout the simulations.

34.5.1 Model

Suppose the data stream $(x_1, x_2, ...)$ consists of IID Bernoulli draws with unknown probability $p^* \in (0, 1)$.

Agent i has a Beta prior:

$$p \sim \text{Beta}(\alpha_i, \beta_i), \qquad i = 1, 2.$$

After observing n draws with k successes, Bayes' rule yields the posterior

$$p \mid x^n \ \sim \ \text{Beta}(\alpha_i + k, \ \beta_i + n - k),$$

and the one-step-ahead predictive probability is

$$\hat{p}_i^n = \mathbb{E}[p \mid x^n] = \frac{\alpha_i + k}{\alpha_i + \beta_i + n}.$$

By the strong law of large numbers, $k/n \to p^*$ almost surely, so both $\hat{p}_1^n$ and $\hat{p}_2^n$ converge to p^* regardless of the agents' initial priors (α_i, β_i).

34.5.2 The marginal likelihood and likelihood ratio

For each fixed value of $p \in (0,1)$, let P_p denote the IID Bernoulli(p) probability law on infinite sequences.

Agent i does not know p.

Instead, agent i places the prior density π_i on p, which induces a probability measure P_i on data sequences via

$$P_i(A) = \int_0^1 P_p(A)\,\pi_i(p)\,dp \qquad \text{for every event } A.$$

So P_i is the agent's marginal probability measure over histories after averaging over uncertainty about p.

In particular, if x^n is an exact observed history with k successes, then $P_i(x^n)$ means the probability that agent i assigns to that history under this mixture measure.

To compute it, start from the Beta density

$$\pi_i(p) = \frac{p^{\alpha_i - 1}(1-p)^{\beta_i - 1}}{B(\alpha_i, \beta_i)}, \qquad 0 < p < 1.$$

Given p, the probability of that ordered history is $p^k(1-p)^{n-k}$.

Therefore

$$\begin{aligned} P_i(x^n) &= \int_0^1 p^k (1-p)^{n-k} \pi_i(p)\,dp \\ &= \frac{1}{B(\alpha_i, \beta_i)} \int_0^1 p^{\alpha_i + k - 1}(1-p)^{\beta_i + n - k - 1}\,dp \\ &= \frac{B(\alpha_i + k,\ \beta_i + n - k)}{B(\alpha_i,\ \beta_i)}. \end{aligned}$$

where $B(a,b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$ is the beta function.

This expression is the probability of the ordered history x^n.

It depends on the data only through the count k, so histories with the same number of successes receive the same probability.

The likelihood ratio at time n is therefore

$$Z_n = \frac{P_{1,n}(x^n)}{P_{2,n}(x^n)} = \frac{B(\alpha_2,\ \beta_2)}{B(\alpha_1,\ \beta_1)} \cdot \frac{B(\alpha_1 + k,\ \beta_1 + n - k)}{B(\alpha_2 + k,\ \beta_2 + n - k)}.$$

In the notation of the theorem, $Q = P_1$ and $P = P_2$.

The process $\{Z_n\}$ is a martingale under P_2 (agent 2's probability) and converges almost surely to a finite positive limit Z_∞, reflecting the fact that $P_1 \sim P_2$ for any Beta priors with positive parameters.

The limit has a closed form.

Let $\bar{p} = \lim_n k_n/n$, which exists P_i-a.s. because, conditional on p, the strong law of large numbers applies.

Then

$$Z_\infty = \frac{dP_1}{dP_2} = \frac{\pi_1(\bar{p})}{\pi_2(\bar{p})}.$$

Exercise 34.12.3 asks for a derivation.

For the priors Beta$(1,8)$ and Beta$(8,1)$ used below, $\pi_1(p) = 8(1-p)^7$ and $\pi_2(p) = 8p^7$, so

$$\log Z_\infty = 7 \log \frac{1-\bar{p}}{\bar{p}}.$$

34.5.3 The exact Blackwell–Dubins distance

For the Beta–Bernoulli model, there is a clean formula for d_n.

By construction, agent i's conditional distribution of the future sequence $(x_{n+1}, x_{n+2}, \ldots)$ given x^n is a mixture of IID Bernoulli(p) processes, with p drawn from agent i's posterior Beta distribution.

Exchangeability and Bayesian Updating describes de Finetti's theorem, which delivers such a mixture for every exchangeable process; here the mixture holds by construction.

The Bernoulli$(p)^\infty$ measures for different p are mutually singular, and the limiting frequency $\bar{p}$ recovers p.

The map $\omega \mapsto \bar{p}$ sends each agent's conditional distribution of the future to that agent's posterior, so the TV distance between the conditional distributions is at least the TV distance between the posteriors.

Each conditional distribution is the image of a posterior under the Markov kernel $p \mapsto \text{Bernoulli}(p)^\infty$, so the TV distance between the conditional distributions is at most the TV distance between the posteriors.

Therefore

$$d_n = \left\| \text{Beta}(\alpha_1 + k_n,\, \beta_1 + n - k_n) - \text{Beta}(\alpha_2 + k_n,\, \beta_2 + n - k_n) \right\|_{\text{TV}}.$$

Both posteriors concentrate around k_n/n with standard deviations of order $n^{-1/2}$.

Concentration alone does not deliver $d_n \to 0$.

Two distributions with standard deviation s_n have a TV distance that goes to zero only if their means differ by $o(s_n)$.

When $\alpha_1 + \beta_1 = \alpha_2 + \beta_2 = 9$, as below, the posterior means differ by $\delta_n = (\alpha_2 - \alpha_1)/(n+9) = 7/(n+9)$, which is $O(1/n)$.

Both posterior standard deviations are approximately $s_n = \sqrt{\hat{p}_n(1-\hat{p}_n)/n}$, where $\hat{p}_n = k_n/n$.

Approximating both posteriors by normal distributions with common standard deviation s_n and means δ_n apart gives

$$d_n \approx 2\,\Phi\left(\frac{\delta_n}{2s_n}\right) - 1 \approx \frac{7}{\sqrt{2\pi n\,\hat{p}_n(1-\hat{p}_n)}},$$

where Φ is the standard normal CDF.

At $\hat{p}_n = 0.65$ the right side is $5.85/\sqrt{n}$.

The section on the rate of merging checks this approximation.

When $\alpha_1 + \beta_1 = \alpha_2 + \beta_2$, the ratio of the two posterior densities is proportional to $((1-p)/p)^{\beta_1 - \beta_2}$, which is monotone in p.

The two densities cross once, at a point x^*, and $d_n = |F_1(x^*) - F_2(x^*)|$, where F_i is agent i's posterior CDF.

The following code implements the Beta–Bernoulli updating, predictive probabilities, TV distance, and likelihood-ratio computations described above.

```
def beta_bernoulli_update(data, a0, b0):
    """
    Sequential Beta-Bernoulli Bayesian updating.
    """
    n = len(data)
    cum_k = np.concatenate([[0], np.cumsum(data)])    # cumulative successes
    ns    = np.arange(n + 1)                           # 0, 1, ..., n
    a_post = a0 + cum_k
    b_post = b0 + (ns - cum_k)
    return a_post, b_post
```

(continues on next page)

(continued from previous page)

```
def predictive_prob(a_post, b_post):
    """One-step-ahead predictive probability P(X=1 | data)."""
    return a_post / (a_post + b_post)


def tv_distance_beta(a1, b1, a2, b2):
    """
    Exact TV distance between Beta(a1, b1) and Beta(a2, b2)
    when a1 + b1 = a2 + b2 and a1 != a2.

    The densities cross once, at x_star, where
    log(x_star / (1 - x_star))
        = (betaln(a2, b2) - betaln(a1, b1)) / (b1 - b2).
    Works elementwise on arrays.
    """
    logit_x = (betaln(a2, b2) - betaln(a1, b1)) / (b1 - b2)
    x_star = 1 / (1 + np.exp(-logit_x))
    return np.abs(beta_dist.cdf(x_star, a1, b1)
                  - beta_dist.cdf(x_star, a2, b2))


def log_likelihood_ratio(data, a1, b1, a2, b2):
    """
    Log likelihood ratio log Z_n = log P1_n(data) - log P2_n(data)
    for every prefix of `data`.

    Returns an array of length len(data) + 1, starting at 0 (before data).
    """
    a1p, b1p = beta_bernoulli_update(data, a1, b1)
    a2p, b2p = beta_bernoulli_update(data, a2, b2)
    log_P1 = betaln(a1p, b1p) - betaln(a1, b1)
    log_P2 = betaln(a2p, b2p) - betaln(a2, b2)
    return log_P1 - log_P2


def run_simulation(p_true, a1, b1, a2, b2, n_steps, seed=0):
    """
    Simulate one realisation of the merging experiment.

    Returns a dict with arrays of length n_steps + 1 (index 0 = prior).
    """
    rng  = np.random.default_rng(seed)
    data = rng.binomial(1, p_true, n_steps)

    a1p, b1p = beta_bernoulli_update(data, a1, b1)
    a2p, b2p = beta_bernoulli_update(data, a2, b2)

    pred1    = predictive_prob(a1p, b1p)
    pred2    = predictive_prob(a2p, b2p)
    tv_1step = np.abs(pred1 - pred2)

    # TV between posterior Betas; in this model this equals d_n
    tv_beta = tv_distance_beta(a1p, b1p, a2p, b2p)

    log_Z = log_likelihood_ratio(data, a1, b1, a2, b2)
```

(continues on next page)

(continued from previous page)

```
    return dict(data=data, pred1=pred1, pred2=pred2,
                tv_1step=tv_1step, tv_beta=tv_beta, log_Z=log_Z)
```

34.5.4 Simulation

We choose two agents with very different beliefs about the bias of a coin.

- Agent 1 (skeptic): prior Beta$(1, 8)$, so $\hat{p}_1^0 = 1/9 \approx 0.11$.
- Agent 2 (optimist): prior Beta$(8, 1)$, so $\hat{p}_2^0 = 8/9 \approx 0.89$.

Both priors are supported on all of $(0, 1)$, so $P_1 \sim P_2$.

Blackwell–Dubins guarantees merging on P_1-almost every path and on P_2-almost every path.

The figure below follows one path generated by a coin with fixed probability of heads $p^* = 0.65$.

The law $P_{0.65}$ of such a path is singular with respect to P_1 and P_2: the event $\{k_n/n \to 0.65\}$ has $P_{0.65}$-probability one and P_i-probability zero.

Blackwell–Dubins therefore says nothing about this path.

Merging at a fixed p^* follows from the Bernstein–von Mises theorem.

Under P_{p^*}, each agent's posterior is close in total variation to the normal distribution with mean k_n/n and variance $p^*(1-p^*)/n$, so the two posteriors are close to each other.

The normal approximation above makes this concrete.

The next subsection draws p from agent 1's prior on each path, which puts us in the setting of the theorem.

```
p_true = 0.65
a1, b1 = 1.0, 8.0    # skeptic
a2, b2 = 8.0, 1.0    # optimist
n_steps = 600

sim = run_simulation(p_true, a1, b1, a2, b2, n_steps, seed=7)
steps = np.arange(n_steps + 1)

fig, axes = plt.subplots(2, 2, figsize=(11, 7))
ax = axes[0, 0]
ax.plot(steps, sim['pred1'], color='steelblue', lw=2,
        label=r'Agent 1 $\hat p_1^n$ (prior: skeptic)')
ax.plot(steps, sim['pred2'], color='firebrick', lw=2,
        label=r'Agent 2 $\hat p_2^n$ (prior: optimist)')
ax.axhline(p_true, color='black', lw=1.0, ls='--',
           label=f'Truth $p^*={p_true}$')
ax.set_xlabel('observations $n$')
ax.set_ylabel('predictive probability')
ax.set_title('(a) posterior predictive means')
ax.legend(fontsize=8)
ax.set_ylim(0, 1)

ax = axes[0, 1]
ax.semilogy(steps, sim['tv_beta'] + 1e-10, color='mediumpurple', lw=2)
ax.set_xlabel('observations $n$')
ax.set_ylabel(
    r'$d_n = \|P(\cdot|\mathscr{F}_n)'
    r' - Q(\cdot|\mathscr{F}_n)\|_{\mathrm{TV}}$'
```

(continues on next page)

(continued from previous page)

```
)
ax.set_title(r'(b) total-variation distance $d_n$')
ax.set_ylim(0.1, 1.1)

ax = axes[1, 0]
log_Z_limit = 7 * np.log((1 - p_true) / p_true)
ax.plot(steps, sim['log_Z'], color='darkorange', lw=2,
        label=r'$\log Z_n$')
ax.axhline(log_Z_limit, color='black', lw=1.0, ls='--',
           label=r'$7 \log((1 - p^*)/p^*)$')
ax.axhline(0, color='black', lw=0.8, ls=':')
ax.set_xlabel('observations $n$')
ax.set_ylabel(r'$\log Z_n$')
ax.set_title(r'(c) log likelihood ratio')
ax.legend(fontsize=8)

ax = axes[1, 1]
xs = np.linspace(0.01, 0.99, 500)
epochs = [0, 20, 100, n_steps]
colors = plt.cm.viridis(np.linspace(0.2, 0.85, len(epochs)))

for epoch, col in zip(epochs, colors):
    k_e = int(np.sum(sim['data'][:epoch]))
    pdf1 = beta_dist.pdf(xs, a1 + k_e, b1 + epoch - k_e)
    pdf2 = beta_dist.pdf(xs, a2 + k_e, b2 + epoch - k_e)
    ax.plot(xs, pdf1, color=col, lw=2, ls='-')
    ax.plot(xs, pdf2, color=col, lw=2, ls='--')

ax.axvline(p_true, color='black', lw=1.0, ls=':', label=f'$p^*={p_true}$')
ax.set_xlabel('$p$')
ax.set_ylabel('posterior density')
ax.set_title('(d) posterior densities')

from matplotlib.lines import Line2D
handles = [
    Line2D([0], [0], color='black', lw=2, label='agent 1'),
    Line2D([0], [0], color='black', lw=2, ls='--', label='agent 2'),
]
for epoch, col in zip(epochs, colors):
    handles.append(Line2D([0], [0], color=col, lw=2, label=f'$n={epoch}$'))
handles.append(
    Line2D([0], [0], color='black', lw=1.0, ls=':', label=f'$p^*={p_true}$')
)
ax.legend(handles=handles, fontsize=8)
ax.set_ylim(bottom=0)

plt.tight_layout()
plt.show()

p_hat = sim['data'].sum() / n_steps
print(f"k_n/n at n = {n_steps}: {p_hat:.3f}")
print(f"d_n at n = {n_steps}: {sim['tv_beta'][-1]:.3f}")
print(f"log Z_n at n = {n_steps}: {sim['log_Z'][-1]:.2f}")
print(f"7 log((1 - k_n/n) / (k_n/n)): "
      f"{7 * np.log((1 - p_hat) / p_hat):.2f}")
print(f"7 log((1 - p*) / p*): {log_Z_limit:.2f}")
```

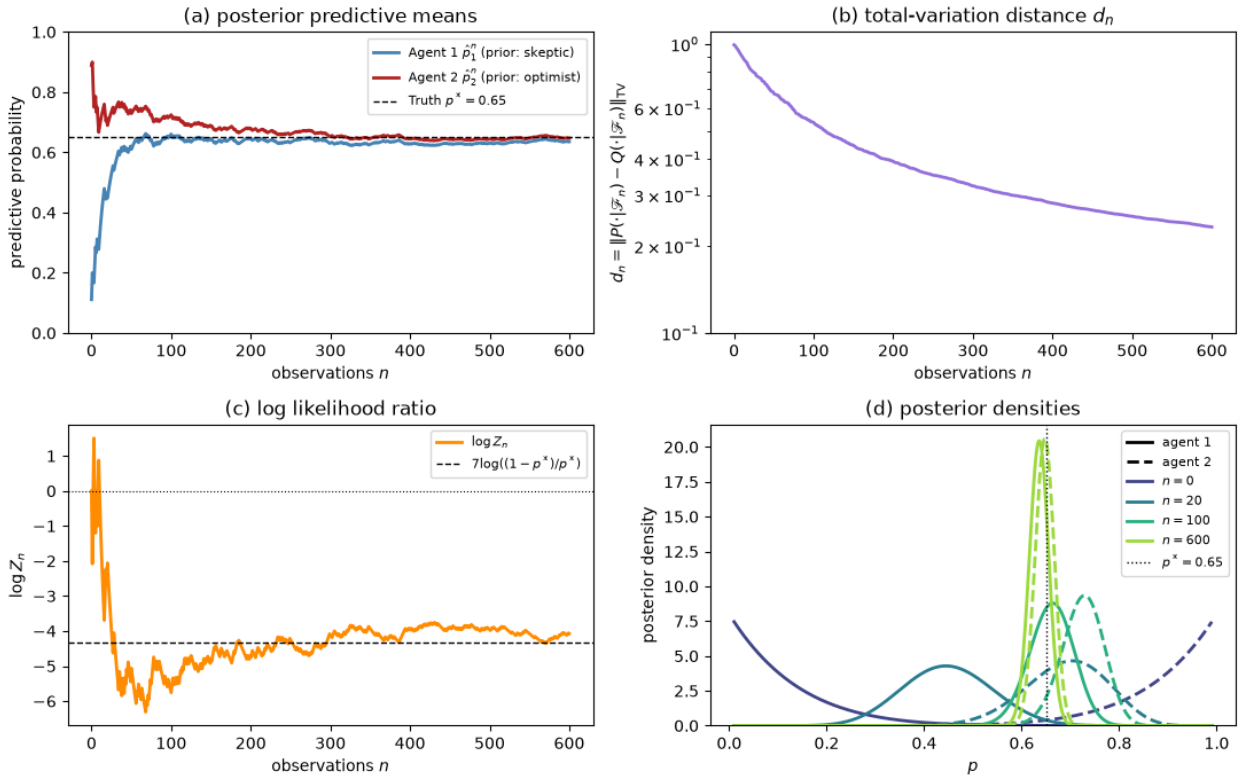


Fig. 34.1: Merging in the Beta–Bernoulli example. The four panels show posterior predictive means, the total-variation distance d_n, the log likelihood ratio $\log Z_n$, and posterior densities at selected horizons.

```
k_n/n at n = 600: 0.643
d_n at n = 600: 0.233
log Z_n at n = 600: -4.07
7 log((1 - k_n/n) / (k_n/n)): -4.13
7 log((1 - p*) / p*): -4.33
```

The four panels show:

- Panel (a): Starting from $\hat{p}_1^0 \approx 0.11$ and $\hat{p}_2^0 \approx 0.89$, both agents' predictive probabilities converge to $p^* = 0.65$.
- Panel (b): The total-variation distance d_n falls from 0.99 to 0.23 at $n = 600$.
- Panel ©: The log likelihood ratio $\log Z_n$ settles near $7 \log((1-p^*)/p^*) = -4.33$, the closed-form limit at $\bar{p} = p^*$. At $n = 600$ the path has $k_n/n = 0.643$, $\log Z_n = -4.07$, and $7 \log((1 - k_n/n)/(k_n/n)) = -4.13$.
- Panel (d): The posterior Beta densities for the two agents start far apart (one near 0, one near 1) and progressively concentrate to the same distribution centred on the truth.

34.5.5 Almost-sure convergence across many paths

To illustrate the almost-sure character of the theorem, we run many independent replications.

Each replication draws p from agent 1's prior $\pi_1 = \text{Beta}(1, 8)$ and then draws $x_1, x_2, \ldots$ IID Bernoulli(p).

Each path is then a draw from P_1.

With $Q = P_1$ and $P = P_2$, the theorem says that $d_n \to 0$ on Q-almost every path.

```
N_paths = 80
n_steps = 2000

rng_p = np.random.default_rng(1234)
p_draws = rng_p.beta(a1, b1, N_paths)     # p drawn from agent 1's prior

tv_all   = np.empty((N_paths, n_steps + 1))
logZ_all = np.empty((N_paths, n_steps + 1))
steps    = np.arange(n_steps + 1)

for i in range(N_paths):
    s = run_simulation(p_draws[i], a1, b1, a2, b2, n_steps, seed=i)
    tv_all[i]   = s['tv_beta']
    logZ_all[i] = s['log_Z']

log_Z_limits = 7 * np.log((1 - p_draws) / p_draws)
norm_p = plt.Normalize(0, p_draws.max())
cmap = plt.cm.viridis

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
ax_tv, ax_log = axes

for i in range(N_paths):
    col = cmap(norm_p(p_draws[i]))
    ax_tv.semilogy(steps, tv_all[i], color=col, lw=0.8, alpha=0.5)
    ax_log.plot(steps, logZ_all[i], color=col, lw=0.8, alpha=0.5)
ax_tv.semilogy(steps, tv_all.mean(axis=0),
               color='black', lw=2, label='mean across paths')
ax_tv.set_xlabel('observations $n$')
ax_tv.set_ylabel(r'$d_n$ (log scale)')
```

(continues on next page)

(continued from previous page)

```
ax_tv.legend()

ax_log.scatter(np.full(N_paths, n_steps), log_Z_limits,
               c=p_draws, cmap=cmap, norm=norm_p, s=12, zorder=3,
               label=r'limit $7 \log((1-p)/p)$')
ax_log.axhline(0, color='gray', lw=0.8, ls=':')
ax_log.set_xlabel('observations $n$')
ax_log.set_ylabel(r'$\log Z_n$')
ax_log.legend()

fig.colorbar(plt.cm.ScalarMappable(norm=norm_p, cmap=cmap),
             ax=axes, label='$p$ drawn from $\\pi_1$')
plt.show()

for n in (500, n_steps):
    print(f"n = {n}: fraction of paths with d_n < 0.30 = "
          f"{np.mean(tv_all[:, n] < 0.30):.2f}, "
          f"mean d_n = {tv_all[:, n].mean():.3f}")
slow = np.argsort(tv_all[:, -1])[-5:]
print(f"p on the five slowest paths: {np.round(p_draws[slow], 4)}")
```

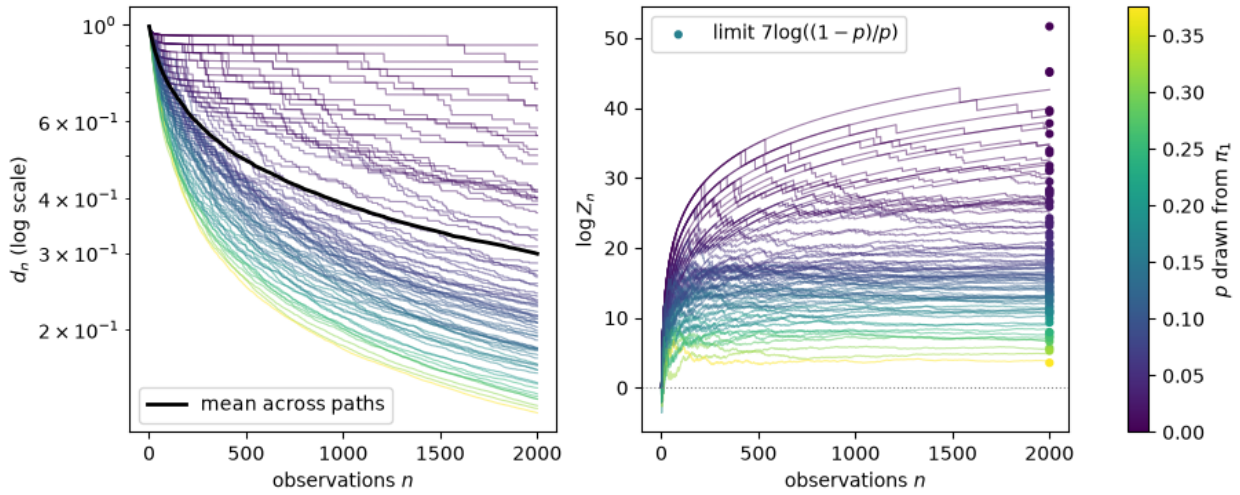


Fig. 34.2: Almost-sure merging across many sample paths drawn from P_1. The left panel plots the total-variation distance and the right panel plots the log likelihood ratio $\log Z_n$. Line colors record the value of p drawn on each path; dots mark the limits $7\log((1-p)/p)$.

```
n = 500: fraction of paths with d_n < 0.30 = 0.14, mean d_n = 0.489
n = 2000: fraction of paths with d_n < 0.30 = 0.68, mean d_n = 0.299
p on the five slowest paths: [0.0036 0.0016 0.0035 0.0006 0.0016]
```

On every path d_{2000} is below d_0, but on many paths d_n falls slowly.

At $n = 500$, 14 percent of paths have $d_n < 0.30$ and the mean of d_n is 0.489; at $n = 2000$ the figures are 68 percent and 0.299.

The slowest paths have p below 0.004.

Agent 2's prior density $8p^7$ is tiny near $p = 0$, so agent 2 needs many observations to move posterior mass there.

These finite-horizon numbers are consistent with the theorem, because almost-sure convergence is an asymptotic state-

ment.

34.5.6 Is d_n a supermartingale?

The proof of the theorem uses no monotonicity of $\{d_n\}$.

In general, $\{d_n\}$ is neither a P-supermartingale nor a Q-supermartingale.

A five-point example shows this.

Let $\Omega = \{1, 2, 3, 4, 5\}$, let P be uniform, and let $Z_\infty = dQ/dP = (0, 1, 1, 1, 2)$.

A first observation reveals whether $\omega \in \{1, 2\}$ or $\omega \in \{3, 4, 5\}$, and a second observation reveals ω.

```
P_pts = np.full(5, 0.2)
Q_pts = np.array([0, 1, 1, 1, 2]) * P_pts
blocks = [[0, 1], [2, 3, 4]]      # cells revealed by the first observation

d_0 = 0.5 * np.abs(P_pts - Q_pts).sum()
E_Q_d1 = E_P_d1 = 0.0
for b in blocks:
    P_b, Q_b = P_pts[b].sum(), Q_pts[b].sum()
    d_1 = 0.5 * np.abs(P_pts[b] / P_b - Q_pts[b] / Q_b).sum()
    E_Q_d1 += Q_b * d_1
    E_P_d1 += P_b * d_1

print(f"d_0 = {d_0:.3f}, E_Q[d_1] = {E_Q_d1:.3f}, E_P[d_1] = {E_P_d1:.3f}")
```

```
d_0 = 0.200, E_Q[d_1] = 0.233, E_P[d_1] = 0.300
```

Here $d_0 = 0.200$, $\mathbb{E}_Q[d_1] = 0.233$, and $\mathbb{E}_P[d_1] = 0.300$.

The first observation raises the expected distance under both measures.

The second observation reveals ω, so $d_2 = 0$.

The identity $2\,\mathbb{E}_Q[d_n] = \mathbb{E}_P[|Z_\infty - Z_n|]$ locates the source of the failure: the L^1 distance between Z_∞ and its conditional expectation given $\mathcal{F}_n$ need not decrease with n.

When $\mathbb{E}_P[Z_\infty^2] < \infty$, the L^2 distance does decrease, because $\mathbb{E}_P[(Z_\infty - Z_n)^2] = \mathbb{E}_P[Z_\infty^2] - \mathbb{E}_P[Z_n^2]$ and $\mathbb{E}_P[Z_n^2]$ is nondecreasing in n.

In the Beta–Bernoulli model with the priors above, the inequality $\mathbb{E}[d_{n+1} \mid \mathcal{F}_n] \le d_n$ holds at every state we check.

The next cell computes $\mathbb{E}_Q[d_{n+1} \mid \mathcal{F}_n] - d_n$ and $\mathbb{E}_P[d_{n+1} \mid \mathcal{F}_n] - d_n$ exactly at every state (n, k_n) with $n < 400$, using each agent's one-step predictive probability.

```
def conditional_increments(n, a1, b1, a2, b2):
    """
    E_Q[d_{n+1} | n, k] - d_n and E_P[d_{n+1} | n, k] - d_n
    for k = 0, ..., n, with Q = P_1 and P = P_2.
    """
    k = np.arange(n + 1)
    d_now  = tv_distance_beta(a1 + k, b1 + n - k, a2 + k, b2 + n - k)
    d_head = tv_distance_beta(a1 + k + 1, b1 + n - k,
                              a2 + k + 1, b2 + n - k)
    d_tail = tv_distance_beta(a1 + k, b1 + n - k + 1,
                              a2 + k, b2 + n - k + 1)
    q1 = (a1 + k) / (a1 + b1 + n)         # agent 1's predictive probability
    q2 = (a2 + k) / (a2 + b2 + n)         # agent 2's predictive probability
```

(continues on next page)

(continued from previous page)

```
    inc_Q = q1 * d_head + (1 - q1) * d_tail - d_now
    inc_P = q2 * d_head + (1 - q2) * d_tail - d_now
    return inc_Q, inc_P

n_check = 400
max_inc_Q = np.empty(n_check)
max_inc_P = np.empty(n_check)
for n in range(n_check):
    inc_Q, inc_P = conditional_increments(n, a1, b1, a2, b2)
    max_inc_Q[n], max_inc_P[n] = inc_Q.max(), inc_P.max()

fig, axes = plt.subplots(1, 2, figsize=(11, 4))

ax = axes[0]
for n, col in zip([20, 100, 399], ['steelblue', 'darkorange', 'firebrick']):
    inc_Q, _ = conditional_increments(n, a1, b1, a2, b2)
    ax.plot(np.arange(n + 1) / n, inc_Q, color=col, lw=2, label=f'$n={n}$')
ax.axhline(0, color='black', lw=0.8, ls='--')
ax.set_xlabel('$k_n / n$')
ax.set_ylabel(r'$\mathbb{E}_Q[d_{n+1} | \mathscr{F}_n] - d_n$')
ax.legend(fontsize=8)

ax = axes[1]
ax.plot(max_inc_Q, color='purple', lw=2, label='under $Q = P_1$')
ax.plot(max_inc_P, color='gray', lw=2, ls='--', label='under $P = P_2$')
ax.axhline(0, color='black', lw=0.8, ls=':')
ax.set_xlabel('observations $n$')
ax.set_ylabel(r'largest increment over $k_n$')
ax.legend(fontsize=8)

plt.tight_layout()
plt.show()

print(f"Largest increment under Q over all states: {max_inc_Q.max():.2e}")
print(f"Largest increment under P over all states: {max_inc_P.max():.2e}")
```

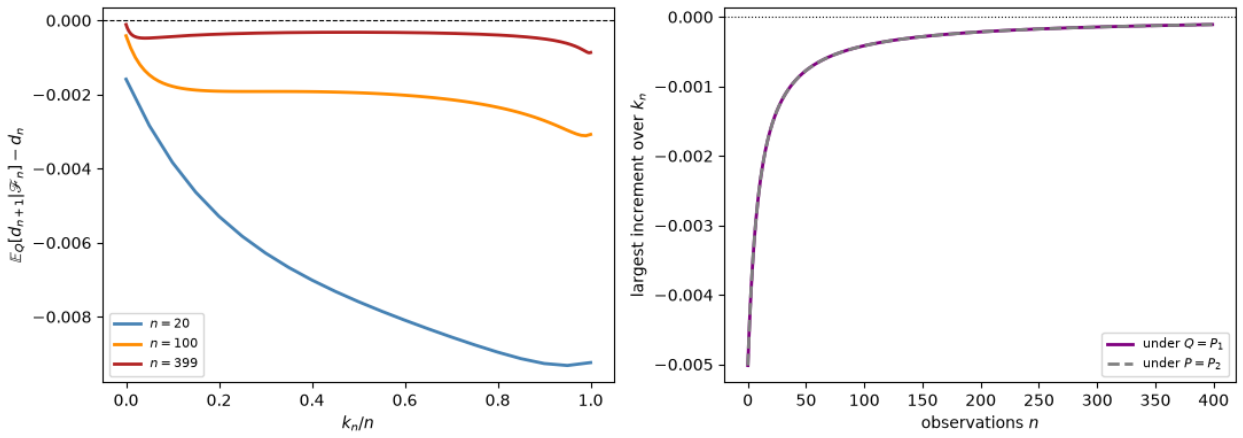


Fig. 34.3: Exact one-step conditional increments of d_n in the Beta–Bernoulli model. The left panel plots $\mathbb{E}_Q[d_{n+1}|\mathcal{F}_n] - d_n$ against k_n/n for selected n. The right panel plots the largest increment over all k_n at each n.

```
Largest increment under Q over all states: -1.10e-04
Largest increment under P over all states: -1.10e-04
```

Every increment is negative, so in this model $\{d_n\}$ behaves as a supermartingale under both P_1 and P_2 over the states we check.

The two curves in the right panel coincide because $\pi_2(p) = \pi_1(1-p)$: exchanging heads and tails exchanges the two agents.

The supermartingale property is specific to the model; the five-point example shows that it fails in general.

34.6 Failure of merging: mutual singularity

What happens when the hypothesis $Q \ll P$ fails?

The singular case is the cleanest counterexample.

34.6.1 Point-mass priors

Suppose both agents hold degenerate (point-mass) priors:

- Agent P: certain that $p = p_P = 0.30$.
- Agent Q: certain that $p = p_Q = 0.75$.

Since P charges only sequences whose empirical frequency converges to 0.30, and Q charges only sequences whose empirical frequency converges to 0.75, the two measures are mutually singular: $P \perp Q$.

The conditional distributions do not update, because both agents are already certain of their model.

For the theorem's object, namely the conditional law of the entire future path,

$$\|P(\cdot \mid \mathcal{F}_n) - Q(\cdot \mid \mathcal{F}_n)\|_{\text{TV}} = \|P - Q\|_{\text{TV}} = 1 \quad \text{for all } n.$$

This equality holds because the infinite-product Bernoulli measures with distinct success probabilities are singular.

If we look only one step ahead, the predictive distance is $|p_P - p_Q| = 0.45$.

That is smaller than one, but it is not the quantity that appears in Blackwell–Dubins.

```
p_P = 0.30
p_Q = 0.75
n_steps = 500

tv_singular_full = np.ones(n_steps + 1)
tv_singular_1step = np.full(n_steps + 1, np.abs(p_P - p_Q))

sim_abs_cont = run_simulation(
    p_Q, 1.0, 8.0, 8.0, 1.0, n_steps, seed=1
)

fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(np.arange(n_steps + 1), tv_singular_full,
        color='firebrick', lw=2,
        label=r'singular: full-path $d_n = 1$')
ax.plot(np.arange(n_steps + 1), tv_singular_1step,
        color='gray', lw=2, ls=':',
```

(continues on next page)

(continued from previous page)

```
        label=r'one-step gap $= |p_P - p_Q|$')
ax.plot(np.arange(n_steps + 1),
        sim_abs_cont['tv_beta'],
        color='steelblue', lw=2,
        label=(r'$\mathrm{Beta}(1,8)$ vs'
               r' $\mathrm{Beta}(8,1)$'))
ax.set_xlabel('observations $n$')
ax.set_ylabel(r'$d_n$')
ax.legend(fontsize=8)
ax.set_ylim(0, 1.05)

plt.tight_layout()
plt.show()
```

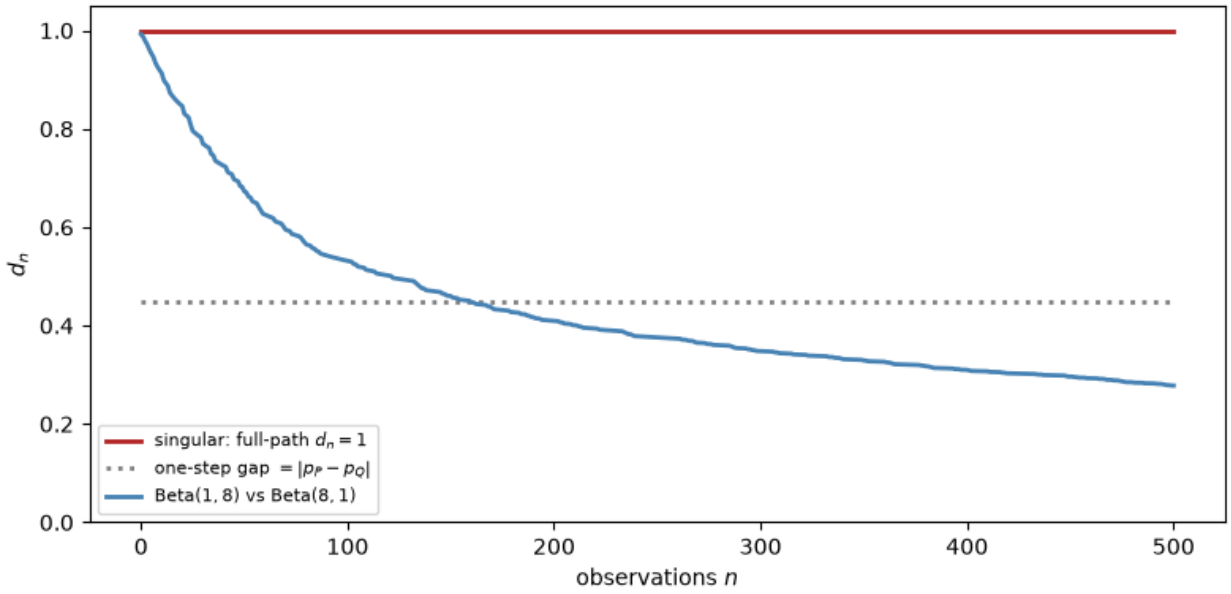


Fig. 34.4: Failure of merging under singular priors. The full future-path distance stays at one, while the one-step predictive gap stays at $|p_P - p_Q|$.

The contrast is sharp.

With mutually absolutely continuous priors, d_n decays to zero on a path generated by a coin with $p = 0.75$.

With singular point-mass priors, the full future-path distance stays at one forever.

More data does not reconcile the agents, because each rules out paths the other assigns positive probability.

34.7 Kakutani's theorem: when does merging hold?

A natural question is: for which product measures does the Blackwell–Dubins hypothesis $Q \ll P$ hold?

For infinite product measures, the answer is given by a classical result of Kakutani [1948].

34.7.1 Hellinger affinities

In this section $P^{(m)}$ and $Q^{(m)}$ denote probability measures on the signal space $(S, \mathcal{S})$, the factors of product measures.

The notation P_n and Q_n keeps its earlier meaning: restrictions of P and Q to $\mathcal{F}_n$.

Definition 34.7.1 (Hellinger Affinity)

For probability measures $P^{(m)}$ and $Q^{(m)}$ on $(S, \mathcal{S})$ with common dominating measure λ, the **Hellinger affinity** is

$$\rho_m = \int_S \sqrt{\frac{dP^{(m)}}{d\lambda} \cdot \frac{dQ^{(m)}}{d\lambda}}\, d\lambda \;\in\; [0, 1].$$

$\rho_m = 1$ if and only if $P^{(m)} = Q^{(m)}$; $\rho_m = 0$ if and only if $P^{(m)} \perp Q^{(m)}$.

The Hellinger affinity is the Bhattacharyya coefficient, the Chernoff coefficient at $\phi = 1/2$ in *Statistical Divergence Measures*.

For two specific one-dimensional families:

- Gaussian: $P^{(m)} = \mathcal{N}(0, 1)$ vs $Q^{(m)} = \mathcal{N}(\mu_m, 1)$:

$$\rho_m^{\text{Gauss}} = \exp\left(-\frac{\mu_m^2}{8}\right).$$

- Bernoulli: $P^{(m)} = \text{Bernoulli}(p)$ vs $Q^{(m)} = \text{Bernoulli}(q)$:

$$\rho_m^{\text{Bern}} = \sqrt{pq} + \sqrt{(1-p)(1-q)}.$$

34.7.2 Kakutani's dichotomy

Theorem 34.7.1 (Kakutani (1948))

Let $P = \bigotimes_{m=1}^{\infty} P^{(m)}$ and $Q = \bigotimes_{m=1}^{\infty} Q^{(m)}$ be infinite product measures whose factors are pairwise equivalent: $P^{(m)} \sim Q^{(m)}$ for every m.

Then either $P \sim Q$ or $P \perp Q$; there is no intermediate case.

Specifically,

$$P \sim Q \quad \Longleftrightarrow \quad \prod_{m=1}^{\infty} \rho_m > 0 \quad \Longleftrightarrow \quad \sum_{m=1}^{\infty} (1 - \rho_m) < \infty.$$

If $\prod_{m=1}^{\infty} \rho_m = 0$, then $P \perp Q$.

Proof idea. A standard proof studies the likelihood-ratio martingale $Z_N = \prod_{m=1}^{N} (dQ^{(m)}/dP^{(m)})$ together with the identity $\mathbb{E}_P[\sqrt{Z_N}] = \prod_{m=1}^{N} \rho_m$.

If the product is positive, $\{\sqrt{Z_N}\}$ is Cauchy in $L^2(P)$, so Z_N converges in $L^1(P)$ and $Q \ll P$; the same argument with P and Q exchanged gives $P \ll Q$.

If the product is zero, Fatou's lemma gives $\mathbb{E}_P[\sqrt{Z_\infty}] = 0$, so $Z_\infty = 0$ P-a.s. and $P \perp Q$.

□

34.7.3 Implication for merging

For IID-type sequences, Kakutani's theorem gives the following picture:

Scenario	$\sum_m (1-\rho_m)$	Conclusion	Merging?
$P^{(m)} = Q^{(m)}$ for all m	0	$P = Q$	Trivially yes
$P^{(m)} \neq Q^{(m)}$ with $\sum_m (1-\rho_m) < \infty$	Finite	$P \sim Q$	Yes; Blackwell–Dubins applies
$P^{(m)} = f \neq g = Q^{(m)}$ for all m	∞	$P \perp Q$	No

The IID case with different fixed marginals is the standard no-merging example.

If two agents assign permanently different distributions to each observation, they end up in disjoint probability worlds.

Likelihood Ratio Processes shows the likelihood ratio process converging to zero when nature permanently draws from one of two fixed densities.

For product measures, the conditional distribution of the future given $\mathcal{F}_n$ does not depend on the past.

So $d_n = \left\| \bigotimes_{m>n} P^{(m)} - \bigotimes_{m>n} Q^{(m)} \right\|_{\text{TV}}$ is deterministic.

Each ρ_m is positive because the factors are equivalent, so dropping finitely many factors does not change whether $\prod_m \rho_m$ is positive.

When $P \perp Q$, the products over $m > n$ are singular, and $d_n = 1$ for every n.

When $P \sim Q$, $d_n \to 0$ by Blackwell–Dubins.

Exercise 34.12.2 computes d_n exactly for Gaussian products.

34.7.4 A Gaussian product-measure example

We illustrate Kakutani's dichotomy with Gaussian product measures.

Take $P = \mathcal{N}(0,1)^{\otimes \mathbb{N}}$ as the reference measure and $Q = \bigotimes_m \mathcal{N}(\mu_m, 1)$ as the alternative.

Then $\log Z_N = \sum_{m=1}^{N} (\mu_m x_m - \mu_m^2/2)$, and the code below draws x_m from $\mathcal{N}(0,1)$, that is, under P.

Three choices of μ_m:

1. $\mu_m = \mu > 0$ constant ($\sum(1-\rho_m) = \infty$) $\Rightarrow P \perp Q$.
2. $\mu_m = c/\sqrt{m}$ ($\sum(1-\rho_m) \approx \sum c^2/(8m) = \infty$) $\Rightarrow P \perp Q$.
3. $\mu_m = c/m$ ($\sum(1-\rho_m) \approx \sum c^2/(8m^2) < \infty$) $\Rightarrow P \sim Q$.

```
N_max  = 2000
ms     = np.arange(1, N_max + 1)
c      = 2.0
N_plot = 400
rng    = np.random.default_rng(0)
```

(continues on next page)

(continued from previous page)

```
cases = [
    (r'$\mu_m = c$ (constant)', np.full(N_max, c)),
    (r'$\mu_m = c/\sqrt{m}$', c / np.sqrt(ms)),
    (r'$\mu_m = c/m$', c / ms),
]
```

With constant drift, $\log Z_N$ drifts to $-\infty$ under P, so $Z_N \to 0$ P-a.s. and $P \perp Q$.

```
label, μ_seq = cases[0]
x = rng.standard_normal(N_plot)
log_Z_inc = μ_seq[:N_plot] * x - μ_seq[:N_plot]**2 / 2
log_Z = np.concatenate([[0], np.cumsum(log_Z_inc)])

fig, ax = plt.subplots(figsize=(8, 3))
ax.plot(np.arange(N_plot + 1), log_Z,
        color='darkorange', lw=2, label=label)
ax.axhline(0, color='black', lw=0.8, ls=':')
ax.set_xlabel('horizon $N$')
ax.set_ylabel(r'$\log Z_N$ under $P$')
ax.legend(fontsize=8)
plt.tight_layout()
plt.show()
```

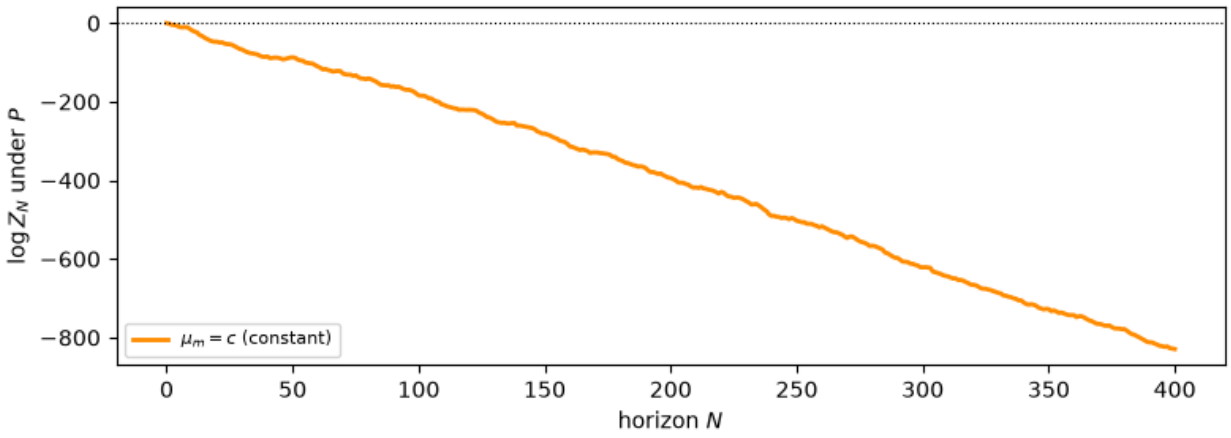


Fig. 34.5: Constant drift $\mu_m = c$: the likelihood ratio collapses ($P \perp Q$).

The $\mu_m = c/\sqrt{m}$ case shows the same qualitative picture: despite the drift vanishing, it does so too slowly.

```
label, μ_seq = cases[1]
x = rng.standard_normal(N_plot)
log_Z_inc = μ_seq[:N_plot] * x - μ_seq[:N_plot]**2 / 2
log_Z = np.concatenate([[0], np.cumsum(log_Z_inc)])

fig, ax = plt.subplots(figsize=(8, 3))
ax.plot(np.arange(N_plot + 1), log_Z,
        color='purple', lw=2, label=label)
ax.axhline(0, color='black', lw=0.8, ls=':')
ax.set_xlabel('horizon $N$')
```

(continues on next page)

(continued from previous page)

```
ax.set_ylabel(r'$\log Z_N$ under $P$')
ax.legend(fontsize=8)
plt.tight_layout()
plt.show()
```

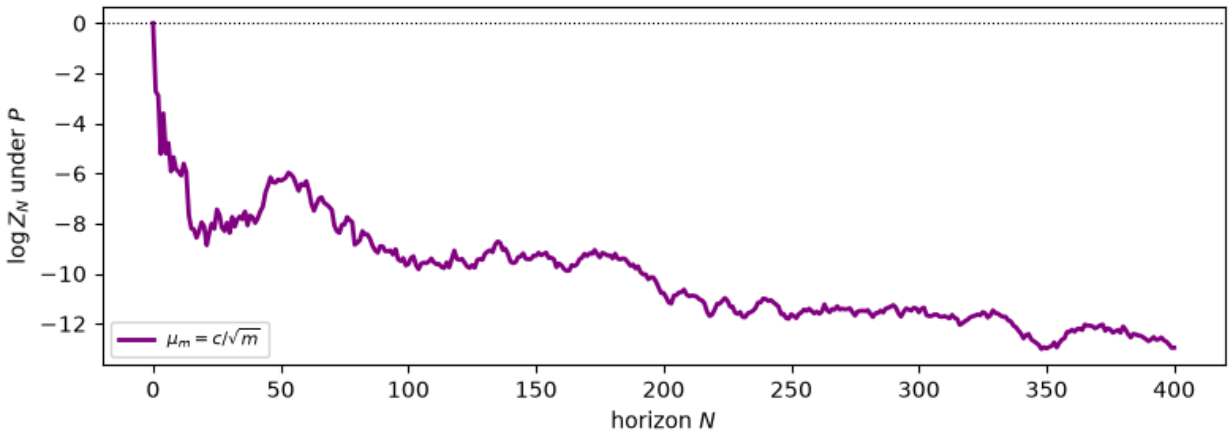


Fig. 34.6: Drift $\mu_m = c/\sqrt{m}$: still singular ($P \perp Q$).

Only with $\mu_m = c/m$ does $\sum(1 - \rho_m) < \infty$ hold, so the likelihood ratio remains nondegenerate and $P \sim Q$.

Blackwell–Dubins applies only in this case.

```
label, μ_seq = cases[2]
x = rng.standard_normal(N_plot)
log_Z_inc = μ_seq[:N_plot] * x - μ_seq[:N_plot]**2 / 2
log_Z = np.concatenate([[0], np.cumsum(log_Z_inc)])

fig, ax = plt.subplots(figsize=(8, 3))
ax.plot(np.arange(N_plot + 1), log_Z,
        color='steelblue', lw=2, label=label)
ax.axhline(0, color='black', lw=0.8, ls=':')
ax.set_xlabel('horizon $N$')
ax.set_ylabel(r'$\log Z_N$ under $P$')
ax.legend(fontsize=8)
plt.tight_layout()
plt.show()
```

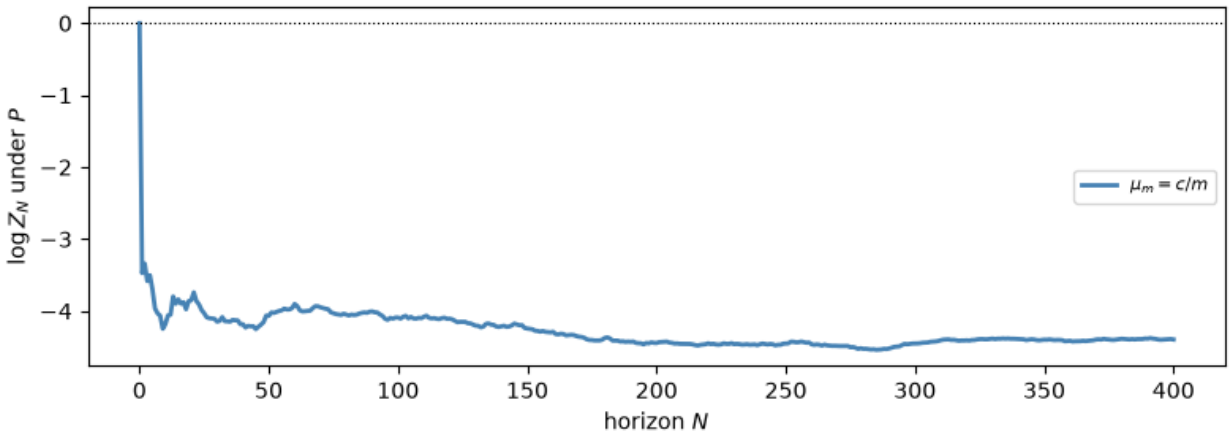


Fig. 34.7: Drift $\mu_m = c/m$: the likelihood ratio stabilises ($P \sim Q$).

34.8 Extension to continuous time

The same logic extends to continuous time.

34.8.1 Girsanov's theorem and the likelihood-ratio process

On the canonical Wiener space with P the Wiener measure (standard Brownian motion W), suppose agent Q believes the process has an additional drift $\theta = \{\theta_s\}_{s \geq 0}$:

$$W_t = \widetilde{W}_t + \int_0^t \theta_s \, ds,$$

where $\widetilde{W}$ is a Q-Brownian motion.

The Girsanov–Cameron–Martin theorem [Girsanov, 1960] gives the likelihood-ratio process $Z_t = dQ_t/dP_t$ as the stochastic exponential

$$Z_t = \exp\left(\int_0^t \theta_s \, dW_s - \frac{1}{2} \int_0^t \theta_s^2 \, ds \right).$$

Z_t is always a non-negative P-local martingale; it is a true martingale if and only if $\mathbb{E}_P[Z_t] = 1$ for all t.

Novikov's condition [Novikov, 1972], $\mathbb{E}_P\left[\exp\left(\frac{1}{2}\int_0^T \theta_s^2 \, ds\right)\right] < \infty$ for all T, is sufficient.

34.8.2 The dichotomy at infinity

A key subtlety on $[0, +\infty)$ is that local absolute continuity does *not* imply global absolute continuity on $\mathcal{F}_\infty$.

Remark 34.8.1 (Infinite-Horizon Subtlety)

Suppose Z_t is a true P-martingale for every finite horizon and let $Z_t \to Z_\infty$ P-a.s.

If $\{Z_t\}$ is uniformly integrable on $[0, \infty)$, then $Q \ll P$ on $\mathcal{F}_\infty$ with $dQ/dP = Z_\infty$.

That is the hypothesis of Blackwell–Dubins, so merging holds Q-a.s.

If in addition $Z_\infty > 0$ P-a.s., then $P \ll Q$ as well, and merging holds P-a.s.

Merging P-a.s. requires only $P \ll Q$, because d_t is symmetric in P and Q and the theorem applies with their roles exchanged.

If uniform integrability fails, then global absolute continuity on $\mathcal{F}_\infty$ can fail.

In many standard examples, including a non-zero constant drift, the measures are in fact singular on $\mathcal{F}_\infty$.

For a deterministic drift θ, the **energy condition**

$$\int_0^\infty \theta_s^2 \, ds < \infty$$

separates the two cases.

If it holds, $\mathbb{E}_P[Z_t^2] = \exp\Big(\int_0^t \theta_s^2 \, ds\Big)$ is bounded in t, $\{Z_t\}$ is uniformly integrable, $Z_\infty > 0$, and $P \sim Q$ on $\mathcal{F}_\infty$.

If it fails, $\log Z_t \to -\infty$ P-a.s. and $P \perp Q$ on $\mathcal{F}_\infty$.

Informally, $\int_0^\infty \theta_s^2 \, ds$ measures the total amount of information separating the two measures over the infinite horizon.

Under the energy condition, Blackwell–Dubins applies and merging holds under both measures.

When θ is a non-zero constant, the condition fails, the measures are singular on $\mathcal{F}_\infty$, and merging does not occur.

Survival and Long-Run Dynamics under Recursive Preferences studies agents whose beliefs differ from the truth by a constant drift of this kind.

When θ depends on the path of W, $\int_0^\infty \theta_s^2 \, ds$ is a random variable.

In that case $P \sim Q$ on $\mathcal{F}_\infty$ if and only if $\int_0^\infty \theta_s^2 \, ds < \infty$ both P-a.s. and Q-a.s.

Under P, $\log Z_t = \int_0^t \theta_s \, dW_s - \frac{1}{2}\int_0^t \theta_s^2 \, ds$ tends to $-\infty$ on the event $\{\int_0^\infty \theta_s^2 \, ds = \infty\}$.

Under Q, $\log Z_t = \int_0^t \theta_s \, d\widetilde{W}_s + \frac{1}{2}\int_0^t \theta_s^2 \, ds$ tends to $+\infty$ on the same event.

On the complementary event, $\log Z_t$ converges to a finite limit under both measures.

Since $Q \ll P$ on $\mathcal{F}_\infty$ if and only if $Z_\infty < \infty$ Q-a.s., and $P \ll Q$ if and only if $Z_\infty > 0$ P-a.s., the claim follows.

Whenever $Q \ll P$ on $\mathcal{F}_\infty$ is established, the discrete-time proof carries over.

The conditional distance satisfies $d_t = \mathbb{E}_P[|Z_\infty - Z_t| \,|\, \mathcal{F}_t]/(2Z_t)$, $Z_t \to Z_\infty > 0$ Q-a.s., and the truncation argument in Step 3 sends the numerator to zero.

The identity $\mathbb{E}_Q[d_t] = \frac{1}{2}\mathbb{E}_P[|Z_t - Z_\infty|]$ also carries over.

34.9 Applications

34.9.1 Bayesian learning

The most direct application is Bayesian inference.

Suppose data $(x_1, x_2, \ldots)$ are drawn from the true measure Q^*.

An agent holds a prior π over a family $\{Q_\theta : \theta \in \Theta\}$, inducing a marginal $P = \int Q_\theta \, \pi(d\theta)$.

If $Q^* \ll P$ (i.e., the agent's marginal model dominates the truth), then Blackwell–Dubins gives

$$\left\| P(\cdot \mid x_1, \ldots, x_n) - Q^*(\cdot \mid x_1, \ldots, x_n) \right\|_{\mathrm{TV}} \to 0 \quad Q^*\text{-a.s.}$$

This is a strong form of Bayesian consistency: the agent's predictions merge with the truth under the true measure.

A prior assigning positive mass to a neighbourhood of the true parameter typically guarantees *local* absolute continuity $Q_n^* \ll P_n$ for every finite horizon n, but not the global condition $Q^* \ll P$ on $\mathcal{F}_\infty$ that Blackwell–Dubins requires.

For example, in the Beta–Bernoulli model with a non-atomic prior π, the mixture $P = \int \mathrm{Bernoulli}(p)^\infty \, \pi(dp)$ satisfies $Q_n^* \ll P_n$ for every n, yet $Q^* \not\ll P$ globally because the set $\{\lim k_n/n = p^*\}$ has Q^*-measure one but P-measure zero (different Bernoulli product measures are mutually singular).

In an identifiable IID family, the measures Q_θ for distinct θ are mutually singular, so $Q_{\theta^*} \ll P$ if and only if $\pi(\{\theta^*\}) > 0$.

Global absolute continuity therefore holds when π has an atom at θ^*, for example when Θ is finite.

Likelihood Ratio Processes and Bayesian Learning studies an agent whose prior puts positive probability on each of two models, one of which is true; there $Q^* \ll P$.

With a non-atomic prior, the agent's forecasts of the entire future never merge with the truth.

In the Beta–Bernoulli model, the argument that gave the closed form for d_n gives $\|P_i(\cdot \mid \mathcal{F}_n) - P_{p^*}(\cdot \mid \mathcal{F}_n)\|_{\mathrm{TV}} = 1$ for every n, because agent i's posterior has no atom at p^*.

Each agent's one-step-ahead forecast still converges to p^*, and the two agents' forecasts merge with each other.

Doob's consistency theorem [Doob, 1949] states that the posterior concentrates at the true θ for π-almost every θ.

Diaconis and Freedman [1986] study consistency at a fixed parameter value, where Doob's theorem is silent.

They show that in infinite-dimensional models the exceptional π-null set can matter: in a location problem with a Dirichlet-process prior on the error distribution, the posterior can be inconsistent at parameter values in the support of the prior.

In *Incorrect Models*, nature draws IID from a mixture h of densities f and g, while the agent believes that nature drew once and for all from f or from g.

There $P \perp Q^*$, the agent's forecasts stay wrong, and the agent's posterior concentrates on the model closer to h in Kullback–Leibler divergence.

34.9.2 Rational expectations and heterogeneous priors

In macroeconomics, rational-expectations models typically impose a common prior.

Blackwell–Dubins gives a dynamic justification for weaker initial agreement.

If two agents start with equivalent priors and observe the same history, their conditional forecasts eventually agree on every event.

Aumann [1976]'s agreement theorem complements this result.

Aumann's agents share a common prior and see different private information; if their posterior probabilities of an event are common knowledge, the posteriors are equal.

Blackwell and Dubins's agents hold different priors and see the same information.

When merging fails, disagreement persists and has economic consequences.

Heterogeneous Beliefs and Financial Markets studies how wealth moves between agents with different beliefs, including an agent whose prior rules out the true model.

Survival and Long-Run Dynamics under Recursive Preferences studies which agents survive when beliefs differ permanently.

34.9.3 Ergodic Markov chains

For a Markov chain with transition kernel Π and two initial distributions μ and ν, the n-step distributions are $\mu\Pi^n$ and $\nu\Pi^n$.

If Π is ergodic with unique stationary distribution π, both converge to π, so

$$\|\mu\Pi^n - \nu\Pi^n\|_{\mathrm{TV}} \le \|\mu\Pi^n - \pi\|_{\mathrm{TV}} + \|\nu\Pi^n - \pi\|_{\mathrm{TV}} \to 0.$$

This convergence does *not* require absolute continuity: μ and ν can be point masses at different states.

It concerns a different object from Blackwell–Dubins.

The measures $\mu\Pi^n$ and $\nu\Pi^n$ are unconditional distributions of the state at date n, and their convergence says that an ergodic chain forgets its initial condition.

Blackwell–Dubins concerns conditional distributions of the entire future given a common observed history.

Two agents who agree on Π and observe the current state make identical conditional forecasts of the future, whether or not Π is ergodic.

34.10 The rate of merging

Blackwell–Dubins is qualitative.

It tells us that $d_n \to 0$, but not how fast.

The identity

$$\mathbb{E}_Q[d_n] = \tfrac{1}{2}\,\mathbb{E}_P[|Z_n - Z_\infty|]$$

shows that the rate of merging in Q-mean is governed by the $L^1(P)$ convergence rate of the likelihood-ratio martingale.

In regular parametric examples, one often sees $n^{-1/2}$-type behavior.

In the Beta–Bernoulli model, the normal approximation derived earlier gives $d_n \approx 7/\sqrt{2\pi n\,\hat{p}_n(1-\hat{p}_n)}$.

Under a fixed coin P_{p^*}, $\hat{p}_n \to p^*$, so $\mathbb{E}_{P_{p^*}}[d_n] \approx C(p^*)/\sqrt{n}$ with $C(p^*) = 7/\sqrt{2\pi p^*(1-p^*)}$ and $C(0.65) = 5.85$.

Under $Q = P_1$, the coin p is drawn from π_1, and averaging over p gives $\mathbb{E}_Q[d_n] \approx C_Q/\sqrt{n}$ with

$$C_Q = \frac{7}{\sqrt{2\pi}}\,\mathbb{E}_{\pi_1}\left[(p(1-p))^{-1/2}\right] = \frac{7}{\sqrt{2\pi}}\,\frac{B(1/2,\,15/2)}{B(1,\,8)} = 14.7.$$

The next cell computes both expectations exactly by summing d_n over the number of successes k_n, which is binomial under P_{p^*} and beta-binomial under $Q = P_1$.

```
def expected_d(n):
    """Exact E_{P_{p*}}[d_n] and E_Q[d_n], with Q = P_1."""
    k = np.arange(n + 1)
    d = tv_distance_beta(a1 + k, b1 + n - k, a2 + k, b2 + n - k)
    return binom.pmf(k, n, p_true) @ d, betabinom.pmf(k, n, a1, b1) @ d

n_grid = np.unique(np.logspace(0, 4, 60).astype(int))
mean_fixed, mean_Q = np.array([expected_d(n) for n in n_grid]).T

C_fixed = 7 / np.sqrt(2 * np.pi * p_true * (1 - p_true))
C_Q = 7 / np.sqrt(2 * np.pi) * beta_fn(0.5, 7.5) / beta_fn(a1, b1)
```

(continues on next page)

(continued from previous page)

```
fig, ax = plt.subplots(figsize=(8, 4))
ax.loglog(n_grid, mean_fixed, color='steelblue', lw=2,
          label=r'$\mathbb{E}_{P_{0.65}}[d_n]$')
ax.loglog(n_grid, C_fixed / np.sqrt(n_grid), color='steelblue',
          lw=1.5, ls='--', label=rf'${C_fixed:.2f}/\sqrt{{n}}$')
ax.loglog(n_grid, mean_Q, color='firebrick', lw=2,
          label=r'$\mathbb{E}_Q[d_n]$, $Q = P_1$')
ax.loglog(n_grid, C_Q / np.sqrt(n_grid), color='firebrick',
          lw=1.5, ls='--', label=rf'${C_Q:.1f}/\sqrt{{n}}$')
ax.set_ylim(1e-2, 1.5)
ax.set_xlabel('sample size $n$')
ax.set_ylabel('mean of $d_n$')
ax.legend()
plt.tight_layout()
plt.show()

for n in (100, 500, 1000, 10_000):
    e_fixed, e_Q = expected_d(n)
    δ, s = 7 / (n + 9), np.sqrt(p_true * (1 - p_true) / n)
    print(f"n = {n:>6}: E_P0.65[d_n] = {e_fixed:.3f}, "
          f"normal approx = {2 * norm.cdf(δ / (2 * s)) - 1:.3f}, "
          f"C/sqrt(n) = {C_fixed / np.sqrt(n):.3f} | "
          f"E_Q[d_n] / (C_Q/sqrt(n)) = {e_Q / (C_Q / np.sqrt(n)):.2f}")

late = n_grid >= 1000
for name, m in (("P_0.65", mean_fixed), ("Q", mean_Q)):
    slope = np.polyfit(np.log(n_grid[late]), np.log(m[late]), 1)[0]
    print(f"log-log slope under {name}, 1000 <= n <= 10^4: {slope:.2f}")
```

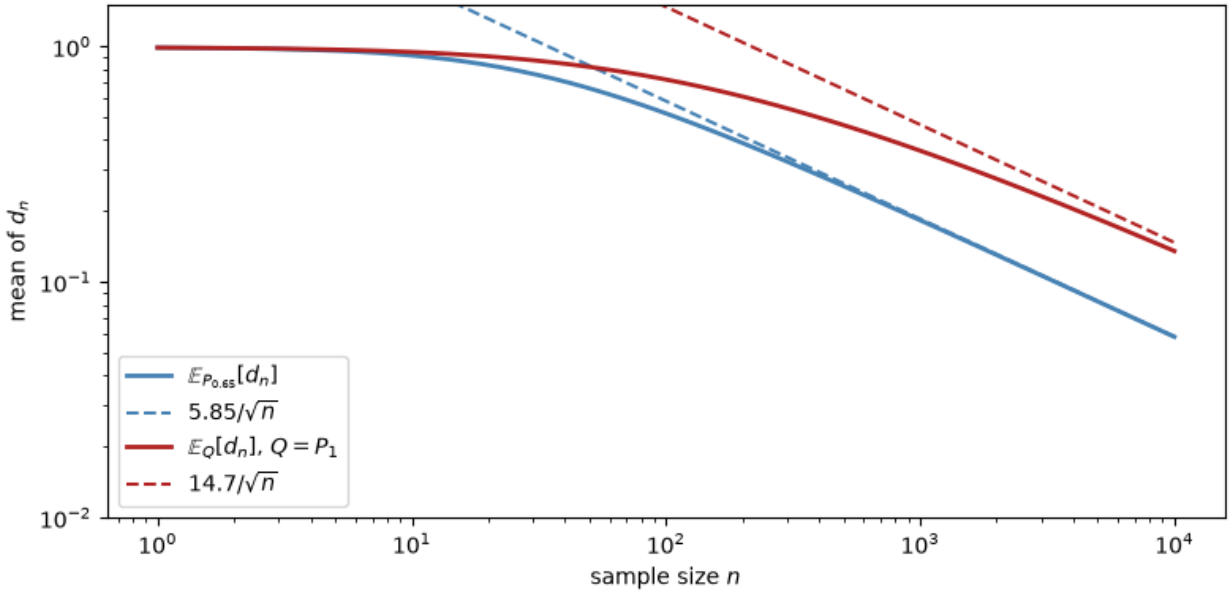


Fig. 34.8: Exact means of d_n in the Beta–Bernoulli model under a fixed coin $P_{0.65}$ and under $Q = P_1$. Dashed lines are the normal approximations $C/\sqrt{n}$.

```
n =    100: E_P0.65[d_n] = 0.519, normal approx = 0.499, C/sqrt(n) = 0.585 | E_Q[d_
↪n] / (C_Q/sqrt(n)) = 0.49
n =    500: E_P0.65[d_n] = 0.255, normal approx = 0.253, C/sqrt(n) = 0.262 | E_Q[d_
↪n] / (C_Q/sqrt(n)) = 0.70
n =   1000: E_P0.65[d_n] = 0.183, normal approx = 0.182, C/sqrt(n) = 0.185 | E_Q[d_
↪n] / (C_Q/sqrt(n)) = 0.78
n =  10000: E_P0.65[d_n] = 0.058, normal approx = 0.058, C/sqrt(n) = 0.059 | E_Q[d_
↪n] / (C_Q/sqrt(n)) = 0.92
log-log slope under P_0.65, 1000 <= n <= 10^4: -0.50
log-log slope under Q, 1000 <= n <= 10^4: -0.43
```

Under $P_{0.65}$, the exact mean approaches $5.85/\sqrt{n}$ from below.

At $n = 500$, the exact mean is 0.255, the normal approximation $2\Phi(\delta_n/(2s_n)) - 1$ at $\hat{p}_n = 0.65$ is 0.253, and $5.85/\sqrt{n}$ is 0.262.

The fitted log-log slope over $1000 \leq n \leq 10^4$ is -0.50.

Under Q, the approach to $14.7/\sqrt{n}$ is slow.

The ratio of $\mathbb{E}_Q[d_n]$ to $14.7/\sqrt{n}$ is 0.70 at $n = 500$ and 0.92 at $n = 10^4$, and the fitted slope over $1000 \leq n \leq 10^4$ is -0.43.

The prior π_1 puts much of its mass near $p = 0$, where the normal approximation needs np to be large.

34.11 Summary and extensions

The logical flow underlying the Blackwell–Dubins theorem is:

$$Q \ll P \implies Z = \frac{dQ}{dP} \in L^1(P) \implies Z_n = \mathbb{E}_P[Z \,|\, \mathcal{F}_n] \to Z \text{ and } \mathbb{E}_P[|Z - Z_n| \,|\, \mathcal{F}_n] \to 0 \;\; P\text{-a.s.} \implies d_n = \frac{\mathbb{E}_P[|Z - Z_n| \,|\, \mathcal{F}_n]}{2Z_n} \xrightarrow{Q\text{-a}}$$

Takeaways:

1. One-sided absolute continuity $Q \ll P$ gives merging Q-almost surely. For merging under *both* measures, one needs mutual absolute continuity $P \sim Q$.
2. The likelihood-ratio martingale $Z_n = \mathbb{E}_P[Z|\mathcal{F}_n]$ drives the result: in $d_n = \mathbb{E}_P[|Z - Z_n| \,|\, \mathcal{F}_n]/(2Z_n)$, the numerator converges to zero and Z_n converges to $Z > 0$, Q-a.s.
3. More data can only increase the ability to distinguish two hypotheses: $\|P_n - Q_n\|_{\text{TV}}$ on $\mathcal{F}_n$ is nondecreasing in n. The conditional distance d_n need not decrease, even in expectation. It converges to zero when $Q \ll P$.
4. For infinite product measures, Kakutani's theorem gives a sharp equivalence-versus-singularity dichotomy: either $P \sim Q$ (when $\sum_m (1 - \rho_m) < \infty$) or $P \perp Q$ (when the sum diverges), with no intermediate case.
5. When $P \sim Q$, Blackwell–Dubins applies and merging occurs under both measures; when $P \perp Q$ and $P_n \sim Q_n$ for every n, $d_n = 1$ for every n and disagreement persists forever.

34.11.1 Applications in economics

Some influential applications and extensions are:

- [Kalai and Lehrer, 1993]: repeated-game learning drives play toward Nash behavior when priors are absolutely continuous with respect to the truth.
- [Kalai and Lehrer, 1993]: subjective and objective equilibria coincide asymptotically under the same condition.
- [Kalai and Lehrer, 1994]: weak and strong notions of merging are introduced for environments where full total-variation convergence is too strong.
- [Kalai *et al.*, 1999]: merging is linked to calibrated forecasting.
- [Jackson *et al.*, 1999]: de Finetti-style representations are connected to Bayesian learning and posterior convergence.
- [Jackson and Kalai, 1999]: social learning erodes reputational effects that rely on persistent disagreement across cohorts.
- [Sandroni, 1998]: near-absolute-continuity conditions are shown to suffice for Nash-type convergence in repeated games.
- [Miller and Sanchirico, 1999]: gives an alternative proof and an economic interpretation of persistent disagreement in terms of mutually favorable bets.
- [Lehrer and Smorodinsky, 1996]: studies broader compatibility notions beyond Blackwell–Dubins absolute continuity.
- [Lehrer and Smorodinsky, 1996]: surveys merging and learning in repeated strategic environments.
- [Nyarko, 1994]: relates Bayesian learning under absolute continuity to convergence toward correlated equilibrium.
- [Pomatto *et al.*, 2014]: extends the theorem to finitely additive probabilities and connects merging to test manipulability.
- [Acemoglu *et al.*, 2016]: shows how disagreement can persist when agents are uncertain about the signal structure itself.

34.12 Exercises

Exercise 34.12.1

Let agent 1's prior be $\pi_1 = \text{Unif}(0, \frac{1}{2})$ and agent 2's prior be $\pi_2 = \text{Unif}(0, 1)$, and let $P_i = \int_0^1 P_p \, \pi_i(p) \, dp$ as in the Beta–Bernoulli section.

1. Show that $P_1 \ll P_2$ but that $P_2 \ll P_1$ fails.
2. Show that $d_n = \Pr_2(p > \frac{1}{2} \mid x^n)$, agent 2's posterior probability that $p > \frac{1}{2}$.
3. Draw paths from P_2: draw p from π_2, then draw $x_1, x_2, \ldots$ IID Bernoulli(p). Compute and plot d_n along each path, and verify that $d_n \to 0$ on paths with $p < \frac{1}{2}$ and $d_n \to 1$ on paths with $p > \frac{1}{2}$.
4. Reconcile part 3 with the Blackwell–Dubins theorem.

Hint: For part 2, agent 1's posterior is agent 2's posterior conditioned on $p < \frac{1}{2}$.

Solution

Part 1.

The density of π_1 is $2 \cdot \mathbf{1}\{p < \frac{1}{2}\} \leq 2\pi_2(p)$, so $P_1(A) \leq 2P_2(A)$ for every event A, and $P_1 \ll P_2$.

Let $\bar{p} = \limsup_n k_n/n$.

The event $\{\bar{p} > \frac{1}{2}\}$ has P_2-probability $\frac{1}{2}$ and P_1-probability zero, so $P_2 \ll P_1$ fails.

Exercise 34.12.3 gives $dP_1/dP_2 = \pi_1(\bar{p})/\pi_2(\bar{p}) = 2 \cdot \mathbf{1}\{\bar{p} < \frac{1}{2}\}$.

Part 2.

Agent 2's posterior after x^n is $\mu_n = \text{Beta}(k_n + 1, n - k_n + 1)$.

Agent 1's posterior has density proportional to $p^{k_n}(1-p)^{n-k_n}\mathbf{1}\{p < \frac{1}{2}\}$, so it equals $\mu_n(\cdot \mid A)$ with $A = (0, \frac{1}{2})$.

For a probability measure μ, an event A with $\mu(A) > 0$, and any event B,

$$\mu(B \mid A) - \mu(B) = \mu(B \cap A)\left(\frac{1}{\mu(A)} - 1\right) - \mu(B \cap A^c),$$

which lies between $-\mu(A^c)$ and $\mu(A^c)$, with the upper bound attained at $B = A$.

Hence $\|\mu(\cdot \mid A) - \mu\|_{\text{TV}} = \mu(A^c)$.

The argument in the Beta–Bernoulli section equates d_n with the TV distance between the two agents' posteriors, so

$$d_n = \mu_n((\tfrac{1}{2}, 1)) = 1 - F_{\text{Beta}(k_n+1,\, n-k_n+1)}(\tfrac{1}{2}).$$

Because $\bar{p} = p$ P_2-a.s., $d_n = P_2(\bar{p} > \frac{1}{2} \mid \mathcal{F}_n)$.

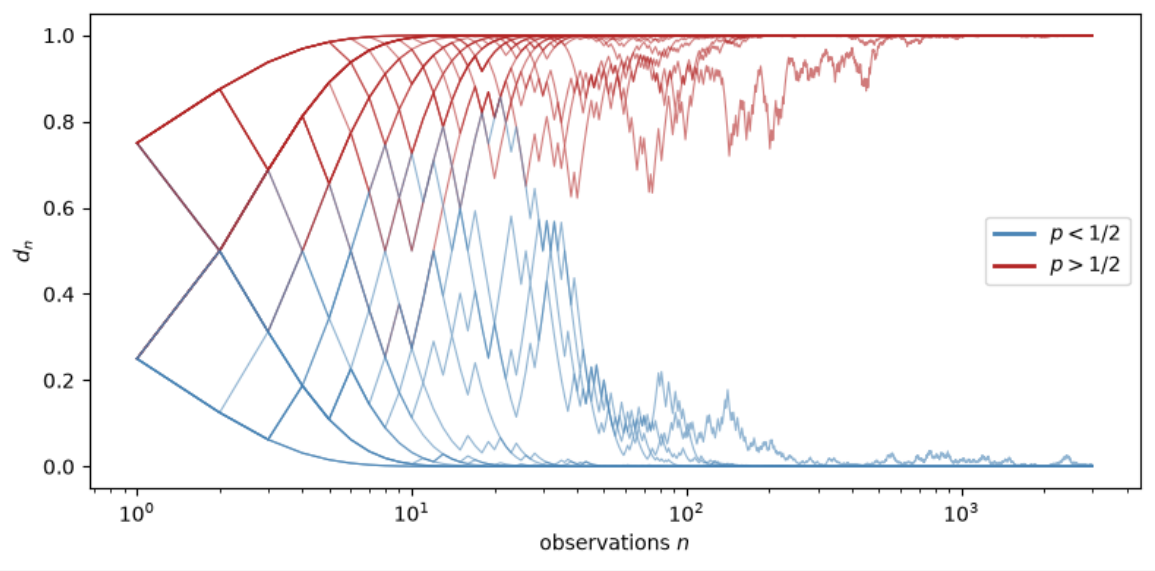
So $\{d_n\}$ is a P_2-martingale, and Lévy's upward theorem gives $d_n \to \mathbf{1}\{\bar{p} > \frac{1}{2}\}$ P_2-a.s.

Part 3.

```
n_steps_ex = 3000
N_paths_ex = 40
rng_ex = np.random.default_rng(0)
p_paths = rng_ex.uniform(0, 1, N_paths_ex)      # p drawn from π_2
steps_ex = np.arange(n_steps_ex + 1)

fig, ax = plt.subplots(figsize=(8, 4))
d_final = np.empty(N_paths_ex)
for i, p in enumerate(p_paths):
    x = rng_ex.binomial(1, p, n_steps_ex)
    k = np.concatenate([[0], np.cumsum(x)])
    d = beta_dist.sf(0.5, k + 1, steps_ex - k + 1)    # Pr_2(p > 1/2 | x^n)
    d_final[i] = d[-1]
    ax.semilogx(steps_ex[1:], d[1:], lw=0.8, alpha=0.6,
                color='firebrick' if p > 0.5 else 'steelblue')
ax.set_xlabel('observations $n$')
ax.set_ylabel('$d_n$')
ax.legend(handles=[
    Line2D([0], [0], color='steelblue', lw=2, label=r'$p < 1/2$'),
    Line2D([0], [0], color='firebrick', lw=2, label=r'$p > 1/2$'),
])
plt.tight_layout()
plt.show()

err = np.abs(d_final - (p_paths > 0.5))
print(f"largest |d_n - 1{{p > 1/2}}| at n = {n_steps_ex}: {err.max():.4f}, "
      f"on the path with p = {p_paths[err.argmax()]:.3f}")
```




```
largest |d_n - 1{p > 1/2}| at n = 3000: 0.0034, on the path with p = 0.486
```

Every path ends near $\mathbf{1}\{p > \frac{1}{2}\}$; the paths that converge slowly have p near $\frac{1}{2}$.

Part 4.

With $Q = P_1$ and $P = P_2$, the hypothesis $Q \ll P$ holds, and the theorem gives $d_n \to 0$ P_1-a.s.

Because $dP_1/dP_2 = 2 \cdot \mathbf{1}\{\bar{p} < \frac{1}{2}\}$, P_1 is P_2 conditioned on $\{\bar{p} < \frac{1}{2}\}$.

So the theorem's conclusion is $d_n \to 0$ P_2-a.s. on $\{\bar{p} < \frac{1}{2}\}$.

The theorem is silent on $\{\bar{p} > \frac{1}{2}\}$, an event with P_2-probability $\frac{1}{2}$ and P_1-probability zero.

There $d_n \to 1$: agent 1's prior rules out $p > \frac{1}{2}$, and no data move agent 1's posterior mass above $\frac{1}{2}$.

Exercise 34.12.2

Let $P = \mathcal{N}(0,1)^{\otimes\mathbb{N}}$ and $Q = \bigotimes_m \mathcal{N}(\mu_m, 1)$, as in the Gaussian product-measure example, and let $\sigma_n^2 = \sum_{m>n} \mu_m^2$.

1. Show that $d_n = 2\Phi(\sigma_n/2) - 1$, with $d_n = 1$ when $\sigma_n = \infty$.
2. For $\mu_m = c/m$, show that $d_n \approx c/\sqrt{2\pi n}$ for large n. Compare the exact and approximate values at $n = 100$ with $c = 2$.
3. For $\mu_m = c/\sqrt{m}$, show that $d_n = 1$ for every n, while the distance between the two agents' one-step-ahead forecasts, $\|\mathcal{N}(0,1) - \mathcal{N}(\mu_{n+1}, 1)\|_{\mathrm{TV}}$, converges to zero.

Kalai and Lehrer [1994] call convergence of forecasts over every fixed finite horizon **weak merging** and convergence of d_n **strong merging**.

Part 3 gives weak merging without strong merging.

Solution

Part 1.

For product measures, the future is independent of the past, so d_n is the TV distance between $\bigotimes_{m>n} \mathcal{N}(0,1)$ and $\bigotimes_{m>n} \mathcal{N}(\mu_m, 1)$.

Suppose $\sigma_n < \infty$.

Then $\sum_{m>n}(1-\rho_m) \leq \sum_{m>n} \mu_m^2/8 < \infty$, so the two tail products are equivalent, and their likelihood ratio is

$$L = \exp\left(S - \tfrac{1}{2}\sigma_n^2\right), \qquad S = \sum_{m>n} \mu_m x_m.$$

Under P, $S \sim \mathcal{N}(0, \sigma_n^2)$; under Q, $S \sim \mathcal{N}(\sigma_n^2, \sigma_n^2)$.

Because L is a function of S, the TV distance between the tail products equals the TV distance between the two laws of S.

Those two normal densities cross at $\sigma_n^2/2$, so

$$d_n = \Phi\left(\frac{\sigma_n}{2}\right) - \Phi\left(-\frac{\sigma_n}{2}\right) = 2\Phi\left(\frac{\sigma_n}{2}\right) - 1.$$

Suppose $\sigma_n = \infty$.

If $\mu_m \not\to 0$, the terms $1-\rho_m = 1 - e^{-\mu_m^2/8}$ do not vanish; if $\mu_m \to 0$, then $1-\rho_m \geq \mu_m^2/16$ for large m.

Either way $\sum_{m>n}(1-\rho_m) = \infty$, Kakutani's theorem makes the tail products singular, and $d_n = 1$.

Part 2.

For $\mu_m = c/m$, $\sigma_n^2 = c^2 \sum_{m>n} m^{-2}$, and $1/(n+1) \leq \sum_{m>n} m^{-2} \leq 1/n$.

So $\sigma_n \approx c/\sqrt{n}$, and $2\Phi(x) - 1 \approx x\sqrt{2/\pi}$ for small x gives $d_n \approx c/\sqrt{2\pi n}$.

The tail sum $\sum_{m>n} m^{-2}$ is the trigamma function $\psi'(n+1)$, which `polygamma(1, n + 1)` computes.

Part 3.

For $\mu_m = c/\sqrt{m}$, $\sum_{m>n} c^2/m = \infty$ for every n, so $d_n = 1$.

The one-step forecast distance is $2\Phi(c/(2\sqrt{n+1})) - 1 \to 0$.

For any fixed k, the distance between the two agents' forecasts of $(x_{n+1}, \ldots, x_{n+k})$ is $2\Phi(\frac{1}{2}(\sum_{m=n+1}^{n+k} c^2/m)^{1/2}) - 1 \leq 2\Phi(c\sqrt{k}/(2\sqrt{n+1})) - 1 \to 0$.

The agents' forecasts of any fixed number of future observations merge, but their forecasts of the entire future never do.

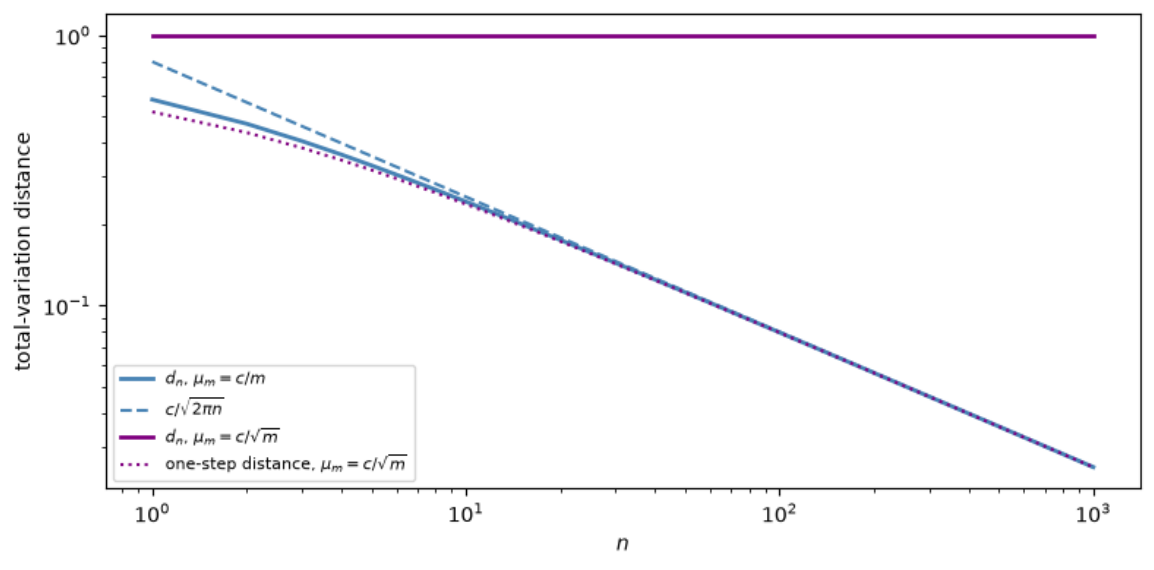
```
c = 2.0
n_vals = np.arange(1, 1001)

σ_inv_m = c * np.sqrt(polygamma(1, n_vals + 1))      # μ_m = c/m
d_inv_m = 2 * norm.cdf(σ_inv_m / 2) - 1
d_inv_m_approx = c / np.sqrt(2 * np.pi * n_vals)
one_step_sqrt = 2 * norm.cdf(c / (2 * np.sqrt(n_vals + 1))) - 1    # μ_m = c/√m

fig, ax = plt.subplots(figsize=(8, 4))
ax.loglog(n_vals, d_inv_m, color='steelblue', lw=2,
          label=r'$d_n$, $\mu_m = c/m$')
ax.loglog(n_vals, d_inv_m_approx, color='steelblue', lw=1.5, ls='--',
          label=r'$c/\sqrt{2\pi n}$')
ax.loglog(n_vals, np.ones_like(n_vals, dtype=float), color='purple',
          lw=2, label=r'$d_n$, $\mu_m = c/\sqrt{m}$')
ax.loglog(n_vals, one_step_sqrt, color='purple', lw=1.5, ls=':',
```

```
          label=r'one-step distance, $\mu_m = c/\sqrt{m}$')
ax.set_xlabel('$n$')
ax.set_ylabel('total-variation distance')
ax.legend(fontsize=8)
plt.tight_layout()
plt.show()

n = 100
print(f"n = {n}: exact d_n = {d_inv_m[n - 1]:.4f}, "
      f"approximation = {d_inv_m_approx[n - 1]:.4f}")
```




```
n = 100: exact d_n = 0.0795, approximation = 0.0798
```

At $n = 100$ with $c = 2$, the exact distance is 0.0795 and the approximation is 0.0798.

The one-step distance for $\mu_m = c/\sqrt{m}$ nearly coincides with d_n for $\mu_m = c/m$, because $\mu_{n+1}^2 = c^2/(n+1)$ and $c^2 \sum_{m>n} m^{-2}$ are both close to c^2/n.

Exercise 34.12.3

Let π_1 and π_2 be prior densities on $(0,1)$ with $\pi_2 > 0$ wherever $\pi_1 > 0$, let $P_i = \int_0^1 P_p \, \pi_i(p)\, dp$, and let $\bar{p} = \limsup_n k_n/n$.

1. Show that $dP_1/dP_2 = \pi_1(\bar{p})/\pi_2(\bar{p})$ on $\mathcal{F}$, so that $Z_\infty = \pi_1(\bar{p})/\pi_2(\bar{p})$.
2. For the priors $\mathrm{Beta}(1,8)$ and $\mathrm{Beta}(8,1)$, simulate several paths from P_2. Compare $\log Z_n$ at large n with $7\log((1-\hat{p}_n)/\hat{p}_n)$, where $\hat{p}_n = k_n/n$, and with $7\log((1-p)/p)$, where p is the coin drawn on that path.

Solution

Part 1.

Under P_p, the strong law of large numbers gives $\bar{p} = p$ almost surely.

Set $\pi_1/\pi_2 = 0$ where $\pi_2 = 0$.

For every $A \in \mathcal{F}$,

$$P_1(A) = \int_0^1 P_p(A)\,\pi_1(p)\,dp = \int_0^1 \mathbb{E}_{P_p}[\mathbf{1}_A]\,\frac{\pi_1(p)}{\pi_2(p)}\,\pi_2(p)\,dp$$
$$= \int_0^1 \mathbb{E}_{P_p}\left[\mathbf{1}_A\,\frac{\pi_1(\bar{p})}{\pi_2(\bar{p})}\right]\pi_2(p)\,dp = \mathbb{E}_{P_2}\left[\mathbf{1}_A\,\frac{\pi_1(\bar{p})}{\pi_2(\bar{p})}\right].$$

The second equality uses $\pi_2 > 0$ wherever $\pi_1 > 0$.

The third replaces p with $\bar{p}$, which equals p P_p-a.s.

So $\pi_1(\bar{p})/\pi_2(\bar{p})$ is a density of P_1 with respect to P_2 on $\mathcal{F}$, and *Lemma 34.3.1* gives $Z_\infty = \pi_1(\bar{p})/\pi_2(\bar{p})$.

A Laplace approximation explains the convergence of Z_n.

The ratio $Z_n = \int_0^1 p^{k_n}(1-p)^{n-k_n}\pi_1(p)\,dp \,/\, \int_0^1 p^{k_n}(1-p)^{n-k_n}\pi_2(p)\,dp$ weights both priors by a likelihood that concentrates near $\hat{p}_n$, so $Z_n \approx \pi_1(\hat{p}_n)/\pi_2(\hat{p}_n)$.

Part 2.

```
n_ex = 20_000
rng_z = np.random.default_rng(5)

print(f"{'p':>6} {'p_hat':>7} {'log Z_n':>9} "
      f"{'gap to p_hat formula':>21} {'gap to p formula':>17}")
for j in range(6):
    p = rng_z.beta(a2, b2)                      # p drawn from π_2
    data = rng_z.binomial(1, p, n_ex)
    log_Z_n = log_likelihood_ratio(data, a1, b1, a2, b2)[-1]
    p_hat = data.mean()
    gap_hat = log_Z_n - 7 * np.log((1 - p_hat) / p_hat)
    gap_true = log_Z_n - 7 * np.log((1 - p) / p)
    print(f"{p:6.3f} {p_hat:7.4f} {log_Z_n:9.3f} "
          f"{gap_hat:21.3f} {gap_true:17.3f}")
```

```
     p   p_hat   log Z_n  gap to p_hat formula  gap to p formula
 0.813  0.8105   -10.169                 0.006             0.113
 0.918  0.9191   -16.996                 0.016            -0.057
 0.877  0.8740   -13.545                 0.009             0.211
 0.917  0.9141   -16.534                 0.015             0.254
 0.856  0.8536   -12.334                 0.008             0.118
 0.986  0.9862   -29.811                 0.099             0.199
```

At $n = 20{,}000$, $\log Z_n$ differs from $7\log((1-\hat{p}_n)/\hat{p}_n)$ by at most 0.10 across these paths.

That gap is the $O(1/n)$ error of the Laplace approximation, and it is largest on the path with p closest to one, where the derivatives of $\log(\pi_1/\pi_2)$ are large.

The gap between $\log Z_n$ and $7\log((1-p)/p)$ reaches 0.25, because $\hat{p}_n$ differs from p by $O(n^{-1/2})$; that gap vanishes as $\hat{p}_n \to p$.

On the path in the first figure of the lecture, $\log Z_{600} = -4.07$ and $7\log((1-\hat{p}_{600})/\hat{p}_{600}) = -4.13$.

CHAPTER

THIRTYFIVE

SURVIVAL AND LONG-RUN DYNAMICS UNDER RECURSIVE PREFERENCES

Contents

- *Survival and Long-Run Dynamics under Recursive Preferences*
  - *Overview*
  - *Environment*
  - *Planner's problem*
  - *Survival conditions*
  - *Wealth dynamics decomposition*
  - *Survival regions*
  - *Three survival channels*
  - *Varying the IES*
  - *Asymptotic results*
  - *The separable case*
  - *Asset pricing implications*
  - *Optimistic and pessimistic distortions*
  - *Long-run consumption distribution*
  - *Summary*

35.1 Overview

This lecture studies the theory of long-run survival in Borovička [2020].

The classical **market selection hypothesis** says that agents with less accurate beliefs are driven out of the market in the long run.

Sandroni [2000] and Blume and Easley [2006] established this result for economies with separable expected-utility preferences and aggregate endowment bounded above and away from zero.

Borovička shows that the conclusion can fail under Epstein-Zin recursive preferences.

With recursive preferences, agents with distorted beliefs can survive and can even dominate.

The key mechanism is that recursive preferences separate risk aversion from the intertemporal elasticity of substitution.

Three channels matter for survival:

1. The *risk premium channel* rewards the more optimistic agent for holding more of the risky asset.
2. The *speculative volatility channel* penalizes aggressive positions through log-return volatility.
3. The *saving channel* changes consumption and saving decisions when the IES differs from one.

Under separable preferences, risk aversion equals the inverse of the IES, and that restriction ties the three channels together.

Their sum then depends only on the relative entropies of the two agents' beliefs, and the agent with more accurate beliefs dominates.

Under recursive preferences, the comparison of risk aversion γ with the inverse IES ρ determines which long-run outcomes can occur.

Both agents can survive only if $\gamma > \rho$, which means that agents prefer early resolution of uncertainty.

Each agent can dominate with positive probability only if $\gamma < \rho$.

An IES above one strengthens the saving channel of an optimistic agent, but coexistence does not require it.

Note

The paper builds on the continuous-time recursive utility formulation of Duffie and Epstein [1992], using the planner's problem approach of Dumas *et al.* [2000].

Important foundations for the market selection hypothesis were laid by De Long *et al.* [1991] and Blume and Easley [1992].

We start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from math import erf, sqrt
from scipy.optimize import brentq
```

35.2 Environment

The economy contains two infinitely lived agents, indexed by $n \in \{1, 2\}$.

The agents have identical recursive preferences but different beliefs about aggregate endowment growth.

We write Borovička's belief distortions u^n as ω^n.

35.2.1 Aggregate endowment

Under the true probability measure P, aggregate endowment satisfies

$$d \log Y_t = \mu_Y dt + \sigma_Y dW_t, \quad Y_0 > 0 \tag{35.1}$$

where W is a standard Brownian motion, μ_Y is the drift, and $\sigma_Y > 0$ is the volatility.

35.2.2 Heterogeneous beliefs

Agent n believes that the drift is $\mu_Y + \omega^n \sigma_Y$ instead of μ_Y.

The parameter ω^n measures optimism when $\omega^n > 0$ and pessimism when $\omega^n < 0$.

Agent n's subjective probability measure Q^n is defined by the Radon–Nikodym derivative

$$M_t^n = \left. \frac{dQ^n}{dP} \right|_t = \exp\left(-\frac{1}{2} |\omega^n|^2 t + \omega^n W_t \right) \tag{35.2}$$

M^n is a likelihood ratio process, a continuous-time counterpart of the processes studied in *Likelihood Ratio Processes*.

Under Q^n, the process $W_t^n = W_t - \omega^n t$ is a Brownian motion, and agent n perceives

$$d \log Y_t = (\mu_Y + \omega^n \sigma_Y) dt + \sigma_Y dW_t^n.$$

An agent with $\omega^n > 0$ is optimistic about endowment growth, while an agent with $\omega^n < 0$ is pessimistic.

Beliefs are dogmatic: agent n never revises ω^n as she observes Y.

For each finite t, Q^n and P are equivalent on $\mathcal{F}_t$.

Under P, $\log M_t^n$ has constant drift $-\frac{1}{2}(\omega^n)^2$, so $M_t^n \to 0$ P-almost surely when $\omega^n \neq 0$, and Q^n and P are mutually singular on $\mathcal{F}_\infty$.

The Blackwell–Dubins theorem in *Merging of Opinions: The Blackwell–Dubins Theorem* requires absolute continuity on $\mathcal{F}_\infty$, so it does not apply here.

An agent with $\omega^n \neq 0$ never comes to agree with P, and agents with $\omega^1 \neq \omega^2$ never come to agree with each other.

Market selection replaces merging: the open question is whether the market eliminates the agent whose beliefs are wrong.

35.2.3 Recursive preferences

Both agents have Epstein-Zin recursive preferences.

We use $\gamma > 0$ for relative risk aversion, $\rho > 0$ for the inverse of the IES, and $\beta > 0$ for the time-preference rate.

The Duffie-Epstein-Zin felicity function is

$$F(C, \nu) = \beta \frac{C^{1-\gamma}}{1-\gamma} \cdot \left(\frac{(1-\gamma) - (1-\rho)\nu/\beta}{\rho - \gamma} \right)^{(\gamma-\rho)/(1-\rho)} \tag{35.3}$$

where ν is the endogenous discount rate.

Note

In discrete time, Epstein-Zin preferences aggregate current consumption with a certainty equivalent of future utility via a CES aggregator (see Doubts or Variability?).

In continuous time there is no "next-period V_{t+1}," so Duffie and Epstein [1992] recast the recursion as a felicity function $F(C, \nu)$ that depends on the agent's own continuation-value rate ν.

The two formulations encode the same separation of risk aversion γ from the inverse IES ρ.

When $\gamma = \rho$, preferences reduce to the standard separable CRRA case.

35.3 Planner's problem

Following Dumas *et al.* [2000], we study equilibrium allocations through a social planner's problem.

The planner chooses consumption shares z^1 and $z^2 = 1 - z^1$ and discount-rate processes ν^n for the two agents.

35.3.1 Modified discount factors

It is convenient to absorb belief distortions into the modified discount factors $\tilde{\lambda}^n = \lambda^n M^n$, where M^n is the Radon-Nikodym derivative (35.2).

These processes satisfy

$$d \log \tilde{\lambda}_t^n = -\left(\nu_t^n + \frac{1}{2}(\omega^n)^2\right) dt + \omega^n dW_t. \tag{35.4}$$

Exercise 35.3.1

Derive (35.4).

Hint: Use $\log \tilde{\lambda}^n = \log \lambda^n + \log M^n$. The Pareto weight λ^n evolves as $d \log \lambda_t^n = -\nu_t^n \, dt$, and $\log M_t^n$ is given by (35.2).

Solution

From the definition $\tilde{\lambda}^n = \lambda^n M^n$, we have

$$\log \tilde{\lambda}_t^n = \log \lambda_t^n + \log M_t^n.$$

The Pareto weight satisfies $d \log \lambda_t^n = -\nu_t^n \, dt$.

From (35.2), $\log M_t^n = -\frac{1}{2}|\omega^n|^2 t + \omega^n W_t$, so

$$d \log M_t^n = -\tfrac{1}{2}(\omega^n)^2 \, dt + \omega^n \, dW_t.$$

Adding the two:

$$d \log \tilde{\lambda}_t^n = -\nu_t^n \, dt - \tfrac{1}{2}(\omega^n)^2 \, dt + \omega^n \, dW_t = -\left(\nu_t^n + \tfrac{1}{2}(\omega^n)^2\right) dt + \omega^n \, dW_t.$$

35.3.2 State variable: Pareto share

The key state variable is the Pareto share of agent 1:

$$\upsilon = \frac{\tilde{\lambda}^1}{\tilde{\lambda}^1 + \tilde{\lambda}^2} \in (0,1) \tag{35.5}$$

It captures the relative weight of agent 1 in the planner's allocation.

Define the log-odds ratio $\vartheta = \log(\upsilon/(1-\upsilon))$.

Its dynamics are

$$d\vartheta_t = \underbrace{\left[\nu_t^2 + \frac{1}{2}(\omega^2)^2 - \nu_t^1 - \frac{1}{2}(\omega^1)^2\right]}_{m_\vartheta(\upsilon_t)} dt + (\omega^1 - \omega^2) dW_t \tag{35.6}$$

The drift $m_\vartheta(\upsilon)$ determines the long-run behavior of the Pareto share.

Exercise 35.3.2

Derive (35.6) from (35.4) and the definition $\vartheta = \log(\upsilon/(1-\upsilon))$.

Hint: First show that $\vartheta = \log\tilde{\lambda}^1 - \log\tilde{\lambda}^2$, then subtract the two SDEs.

Solution

Since $\upsilon = \tilde{\lambda}^1/(\tilde{\lambda}^1 + \tilde{\lambda}^2)$, we have $1 - \upsilon = \tilde{\lambda}^2/(\tilde{\lambda}^1 + \tilde{\lambda}^2)$, so

$$\vartheta = \log\frac{\upsilon}{1-\upsilon} = \log\tilde{\lambda}^1 - \log\tilde{\lambda}^2.$$

From (35.4), the two log-discount-factor SDEs are

$$d\log\tilde{\lambda}_t^1 = -\left(\nu_t^1 + \tfrac{1}{2}(\omega^1)^2\right)dt + \omega^1 dW_t,$$

$$d\log\tilde{\lambda}_t^2 = -\left(\nu_t^2 + \tfrac{1}{2}(\omega^2)^2\right)dt + \omega^2 dW_t.$$

Subtracting the second from the first:

$$d\vartheta_t = \left[\nu_t^2 + \tfrac{1}{2}(\omega^2)^2 - \nu_t^1 - \tfrac{1}{2}(\omega^1)^2\right]dt + (\omega^1 - \omega^2)dW_t.$$

35.3.3 HJB equation

Homotheticity reduces the planner's problem to a nonlinear ODE in the single state variable υ.

Because each agent's utility is homogeneous of degree $1-\gamma$ in consumption, the planner's value function is homogeneous of degree one in the weights $(\tilde{\lambda}^1, \tilde{\lambda}^2)$ and of degree $1-\gamma$ in Y.

It factors as $J(\tilde{\lambda}, Y) = (\tilde{\lambda}^1 + \tilde{\lambda}^2)Y^{1-\gamma}\tilde{J}(\upsilon)$, eliminating Y as a state variable.

From discrete to continuous time

In discrete time, a planner maximizes a weighted sum of agents' utilities by choosing allocations at each date.

The Bellman equation is

$$\tilde{J}(\upsilon) = \max_{z^1, z^2} \left\{ \upsilon\, u(z^1) + (1-\upsilon)\, u(z^2) + \delta\, \mathbb{E}\left[\tilde{J}(\upsilon')\right] \right\}.$$

Here $\delta \in (0,1)$ is a one-period discount factor; elsewhere in this lecture β is a continuous-time rate of time preference.

In continuous time, the period length shrinks to dt.

The "flow payoff" over $[t, t+dt)$ becomes $\left[\upsilon F(z^1,\nu^1) + (1-\upsilon)F(z^2,\nu^2)\right] dt$, where F is the Duffie-Epstein-Zin felicity (35.3).

The expected change in the value function over dt is captured by the **infinitesimal generator** $\mathcal{L}$.

For a diffusion $d\upsilon = m\, dt + s\, dW$, Itô's lemma gives

$$\mathcal{L}\tilde{J}(\upsilon) = m(\upsilon)\, \tilde{J}'(\upsilon) + \tfrac{1}{2} s(\upsilon)^2\, \tilde{J}''(\upsilon),$$

where m and s are the drift and diffusion of the Pareto share.

This is the continuous-time analogue of $\delta\, \mathbb{E}[\tilde{J}(\upsilon')] - \tilde{J}(\upsilon)$: it measures how the value function drifts and fluctuates as υ evolves.

Setting flow payoff plus expected capital gain equal to zero gives the schematic HJB equation:

$$0 = \sup_{(z^1, z^2, \nu^1, \nu^2)} \left\{ \upsilon F(z^1,\nu^1) + (1-\upsilon)F(z^2,\nu^2) + \mathcal{L}\tilde{J}(\upsilon) \right\} \tag{35.7}$$

subject to $z^1 + z^2 \leq 1$.

Exact reduced ODE

Proposition 2.3 of Borovička [2020] gives the exact HJB equation after substituting the homogeneity reduction $J(\tilde{\lambda}, Y) = (\tilde{\lambda}^1 + \tilde{\lambda}^2) Y^{1-\gamma} \tilde{J}(\upsilon)$ and the dynamics of υ and Y:

$$\begin{aligned} 0 = \sup_{(z^1, z^2, \nu^1, \nu^2)} & \; \upsilon\, F(z^1,\nu^1) + (1-\upsilon)\, F(z^2,\nu^2) \\ & + \left[-\upsilon\nu^1 - (1-\upsilon)\nu^2 + (\upsilon\omega^1 + (1-\upsilon)\omega^2)(1-\gamma)\sigma_Y + (1-\gamma)\mu_Y + \tfrac{1}{2}(1-\gamma)^2\sigma_Y^2 \right] \tilde{J}(\upsilon) \\ & + \upsilon(1-\upsilon)\left[\nu^2 - \nu^1 + (\omega^1 - \omega^2)(1-\gamma)\sigma_Y \right] \tilde{J}'(\upsilon) \\ & + \tfrac{1}{2}\upsilon^2 (1-\upsilon)^2 (\omega^1 - \omega^2)^2\, \tilde{J}''(\upsilon) \end{aligned} \tag{35.8}$$

subject to $z^1 + z^2 \leq 1$.

The first line is the flow payoff from the two agents' felicity functions.

The second line multiplies $\tilde{J}(\upsilon)$ by a term that combines the agents' discount rates, belief-weighted endowment drift, and a variance correction – these arise from absorbing the $Y^{1-\gamma}$ factor via Itô's lemma.

The third line multiplies $\tilde{J}'(\upsilon)$ by the drift of the Pareto share, which depends on the difference in discount rates and the belief-weighted response to endowment risk.

The fourth line multiplies $\tilde{J}''(\upsilon)$ by the squared diffusion of the Pareto share.

The boundary conditions are $\tilde{J}(0) = \tilde{V}^2$ and $\tilde{J}(1) = \tilde{V}^1$, where $\tilde{V}^n$ is the continuation value in the homogeneous economy populated by agent n alone.

This is the continuous-time counterpart of the discrete-time planner's problem in Blume and Easley [2006] (see also *Heterogeneous Beliefs and Financial Markets*).

35.4 Survival conditions

The central result characterizes survival by the boundary behavior of $m_\vartheta(\upsilon)$.

Proposition 35.4.1

Define the following repelling conditions (i) and (ii) and their attracting counterparts (i') and (ii'):

$$\text{(i)} \lim_{\upsilon \searrow 0} m_\vartheta(\upsilon) > 0, \qquad \text{(i')} \lim_{\upsilon \searrow 0} m_\vartheta(\upsilon) < 0$$

$$\text{(ii)} \lim_{\upsilon \nearrow 1} m_\vartheta(\upsilon) < 0, \qquad \text{(ii')} \lim_{\upsilon \nearrow 1} m_\vartheta(\upsilon) > 0$$

Then:

(a) If (i) and (ii) hold, both agents survive under P.

(b) If (i) and (ii') hold, agent 1 dominates in the long run under P.

© If (i') and (ii) hold, agent 2 dominates in the long run under P.

(d) If (i') and (ii') hold, each agent dominates with strictly positive probability.

The proof uses the Feller classification of boundary behavior for diffusions, as in Karlin and Taylor [1981].

Condition (i) says that when agent 1 is close to extinction, there is a force pushing her share back up.

Condition (ii) says that when agent 1 is close to absorbing the whole economy, there is a force pushing her share back down.

When both forces are present, the Pareto share is recurrent and both agents survive.

35.5 Wealth dynamics decomposition

We now rewrite the survival conditions from *Proposition 35.4.1* in terms of equilibrium wealth dynamics.

Agent 1 survives near extinction if and only if her wealth grows faster than agent 2's when she is negligibly small.

When $\upsilon \searrow 0$, prices are set entirely by agent 2, as if the economy were homogeneous.

Agent 1 is a price-taker in agent 2's economy.

Let $m_A^n(\upsilon)$ denote the expected log growth rate of agent n's wealth A^n, so that $d \log A_t^n = m_A^n(\upsilon_t)dt + \sigma_A^n(\upsilon_t)dW_t$.

Proposition 3.3 of Borovička [2020] links the boundary drifts of the log-odds ratio to these growth rates:

$$\lim_{\upsilon \to \bar{\upsilon}} \gamma^{-1} m_\vartheta(\upsilon) = \lim_{\upsilon \to \bar{\upsilon}} [m_A^1(\upsilon) - m_A^2(\upsilon)], \qquad \bar{\upsilon} \in \{0, 1\}. \tag{35.9}$$

Let R^n denote the cumulative return index of agent n's portfolio, with $d \log R_t^n = m_R^n(\upsilon_t)dt + \sigma_R^n(\upsilon_t)dW_t$.

Thus m_R^n is the expected *log* return on agent n's portfolio, and $m_R^n + \frac{1}{2}(\sigma_R^n)^2$ is its expected arithmetic return.

Let $(y^n)^{-1}$ denote agent n's consumption-wealth ratio.

The difference in wealth growth rates decomposes into two channels:

$$\lim_{\upsilon \searrow 0} [m_A^1(\upsilon) - m_A^2(\upsilon)] = \underbrace{\lim_{\upsilon \searrow 0} [m_R^1(\upsilon) - m_R^2(\upsilon)]}_{\text{portfolio returns}} + \underbrace{\lim_{\upsilon \searrow 0} [(y^2(\upsilon))^{-1} - (y^1(\upsilon))^{-1}]}_{\text{consumption-wealth ratios}} \tag{35.10}$$

The first term measures how much faster agent 1's portfolio grows.

The second measures how much less agent 1 consumes out of wealth – a lower consumption-wealth ratio means more saving and faster wealth accumulation.

When this total difference is positive, condition (i) holds and agent 1 survives near extinction.

When it is negative, the boundary $\upsilon = 0$ attracts.

Agent 1 then becomes extinct almost surely if condition (ii) also holds (case ©), and with strictly positive probability if condition (ii') holds (case (d)).

Exercise 35.5.1

Derive (35.10).

Let A^n denote agent n's wealth and C^n her consumption.

Let R^n be the cumulative return index of her portfolio, with $d \log R^n = m_R^n \, dt + \sigma_R^n \, dW$, so that m_R^n is her expected log portfolio return.

The budget constraint is $dA^n = A^n \, dR^n / R^n - C^n dt$.

Define the consumption-wealth ratio $(y^n)^{-1} = C^n / A^n$.

Show that $d \log A^n = (m_R^n - (y^n)^{-1}) \, dt + \sigma_R^n \, dW$, so the difference in expected log wealth growth is $m_A^1 - m_A^2 = (m_R^1 - m_R^2) + [(y^2)^{-1} - (y^1)^{-1}]$.

Solution

By Itô's lemma applied to $R^n = \exp(\log R^n)$, the instantaneous return is

$$\frac{dR^n}{R^n} = \left(m_R^n + \tfrac{1}{2}(\sigma_R^n)^2\right) dt + \sigma_R^n \, dW.$$

Dividing the budget constraint by A^n:

$$\frac{dA^n}{A^n} = \left(m_R^n + \tfrac{1}{2}(\sigma_R^n)^2 - (y^n)^{-1}\right) dt + \sigma_R^n \, dW.$$

By Itô's lemma, $d \log A^n = \frac{dA^n}{A^n} - \frac{1}{2}\left(\frac{dA^n}{A^n}\right)^2$, and $(dA^n / A^n)^2 = (\sigma_R^n)^2 dt$.

Then

$$d \log A^n = (m_R^n - (y^n)^{-1}) \, dt + \sigma_R^n \, dW,$$

so $m_A^n = m_R^n - (y^n)^{-1}$.

Taking the difference for agents 1 and 2 gives

$$m_A^1 - m_A^2 = (m_R^1 - m_R^2) + \left[(y^2)^{-1} - (y^1)^{-1}\right],$$

and passing to the limit $\upsilon \searrow 0$ gives (35.10).

35.5.1 Portfolio returns

At the boundary $\upsilon \searrow 0$, the difference in expected log portfolio returns is

$$\lim_{\upsilon \searrow 0}[m_R^1 - m_R^2] = \underbrace{\frac{\omega^1 - \omega^2}{\gamma\sigma_Y}}_{\text{difference in risky shares}} \cdot \underbrace{(\gamma\sigma_Y^2 - \omega^2\sigma_Y)}_{\text{risk premium}} - \underbrace{\frac{\omega^1 - \omega^2}{\gamma}\left(\sigma_Y + \frac{\omega^1 - \omega^2}{2\gamma}\right)}_{\text{volatility term}} \tag{35.11}$$

An optimistic agent ($\omega^1 > \omega^2$) overweights the risky asset by $(\omega^1 - \omega^2)/(\gamma\sigma_Y)$ relative to agent 2 and earns the equity risk premium on that extra exposure.

The subtracted *volatility penalty* reflects the cost of holding a more extreme portfolio: higher variance of log returns drags down expected log wealth growth.

This term depends on risk aversion γ but not on the IES, because portfolio choice is determined by risk aversion alone.

Exercise 35.5.2

Derive (35.11).

At the boundary $\upsilon \searrow 0$, agent n's optimal risky-asset share is $\pi^n = 1 + (\omega^n - \omega^2)/(\gamma\sigma_Y)$ (see (35.15)).

Let r denote the risk-free rate and $\bar{\mu}_R = r + \gamma\sigma_Y^2 - \omega^2\sigma_Y$ the expected arithmetic return on the risky asset under P.

(Using (35.13) and (35.14), $\bar{\mu}_R$ equals the expected capital gain $\mu_Y + \frac{1}{2}\sigma_Y^2$ plus the dividend yield $(y(0))^{-1}$.)

The continuously rebalanced portfolio has expected log return $m_R^n = r + \pi^n(\bar{\mu}_R - r) - \frac{1}{2}(\pi^n)^2\sigma_Y^2$.

Compute $m_R^1 - m_R^2$ and simplify.

Solution

Using $m_R^n = r + \pi^n(\bar{\mu}_R - r) - \frac{1}{2}(\pi^n)^2\sigma_Y^2$, the difference is

$$m_R^1 - m_R^2 = (\pi^1 - \pi^2)(\bar{\mu}_R - r) - \tfrac{1}{2}[(\pi^1)^2 - (\pi^2)^2]\sigma_Y^2.$$

The difference in risky shares is $\pi^1 - \pi^2 = (\omega^1 - \omega^2)/(\gamma\sigma_Y)$.

The arithmetic equity premium is $\bar{\mu}_R - r = \gamma\sigma_Y^2 - \omega^2\sigma_Y$, so:

$$(\pi^1 - \pi^2)(\bar{\mu}_R - r) = \frac{\omega^1 - \omega^2}{\gamma\sigma_Y} \cdot (\gamma\sigma_Y^2 - \omega^2\sigma_Y).$$

For the volatility term, write $(\pi^1)^2 - (\pi^2)^2 = (\pi^1 - \pi^2)(\pi^1 + \pi^2)$ and use $\pi^2 = 1$, so $\pi^1 + \pi^2 = 2 + (\omega^1 - \omega^2)/(\gamma\sigma_Y)$.

After simplification:

$$\tfrac{1}{2}[(\pi^1)^2 - (\pi^2)^2]\sigma_Y^2 = \frac{\omega^1 - \omega^2}{\gamma}\left(\sigma_Y + \frac{\omega^1 - \omega^2}{2\gamma}\right).$$

Combining the two pieces gives (35.11).

35.5.2 Consumption-wealth ratios

The difference in consumption-wealth ratios at the boundary is

$$\lim_{\upsilon \searrow 0}[(y^2)^{-1} - (y^1)^{-1}] = \frac{1-\rho}{\rho}\left[(\omega^1 - \omega^2)\sigma_Y + \frac{(\omega^1 - \omega^2)^2}{2\gamma}\right] \tag{35.12}$$

The term in brackets is the gap between the *subjective* certainty-equivalent returns that the two agents expect to earn on their optimal portfolios.

Under her own beliefs Q^n, agent n perceives the Sharpe ratio $\lambda^n = \gamma\sigma_Y + \omega^n - \omega^2$ on the risky asset.

Her optimal portfolio earns the certainty-equivalent return $r + (\lambda^n)^2/(2\gamma)$, and the term in brackets equals $[(\lambda^1)^2 - (\lambda^2)^2]/(2\gamma)$.

The bracket also equals half the gap between the agents' subjective expected arithmetic portfolio returns.

The factor $(1-\rho)/\rho$ translates this perceived return advantage into a saving response.

- When IES > 1 ($\rho < 1$), the factor is positive: a higher perceived return makes the agent save more, because the substitution effect dominates the income effect.
- When IES < 1 ($\rho > 1$), the factor is negative: the income effect dominates and the agent saves less, working against survival.
- When IES $= 1$ ($\rho = 1$), the two effects cancel and the saving channel vanishes entirely.

This is the channel through which recursive preferences alter survival outcomes by separating γ from ρ.

Exercise 35.5.3

Derive (35.12).

At the boundary $\upsilon \searrow 0$, both agents take as given the risk-free rate r and a risky asset with volatility σ_Y and expected excess return $\gamma\sigma_Y^2 - \omega^2\sigma_Y$ under P.

1. Show that under Q^n agent n perceives the Sharpe ratio $\lambda^n = \gamma\sigma_Y + \omega^n - \omega^2$. A portfolio with risky share π has subjective certainty-equivalent return $r + \pi\sigma_Y\lambda^n - \frac{1}{2}\gamma\pi^2\sigma_Y^2$. Show that the optimal share reproduces (35.15) and that the maximized certainty-equivalent return is $\mathrm{CE}^n = r + (\lambda^n)^2/(2\gamma)$.
2. Let $c^n = (y^n)^{-1}$ and let $\mu_V^n = \mathrm{CE}^n - c^n$ be the subjective certainty-equivalent growth rate of agent n's wealth, which is also that of her consumption. Use the Hamilton-Jacobi-Bellman equation of an Epstein-Zin investor facing constant investment opportunities to show that $c^n = \beta - (1-\rho)\mu_V^n$, and hence $\mu_V^n = (\mathrm{CE}^n - \beta)/\rho$.
3. Show that $\mu_V^2 = \mu_Y + \omega^2\sigma_Y + \frac{1}{2}(1-\gamma)\sigma_Y^2$, so that c^2 reproduces (35.14).
4. Deduce (35.12).

Hint: For step 2, use the normalized Duffie-Epstein aggregator

$$f(C,V) = \frac{\beta}{1-\rho}(1-\gamma)V\left[\left(\frac{C}{((1-\gamma)V)^{1/(1-\gamma)}}\right)^{1-\rho} - 1\right]$$

and guess $V(A) = (\kappa A)^{1-\gamma}/(1-\gamma)$ for a constant $\kappa > 0$.

Solution

Step 1. Under Q^n, $dW = dW^n + \omega^n dt$.

The risky asset's instantaneous return, $(r + \gamma\sigma_Y^2 - \omega^2\sigma_Y)dt + \sigma_Y dW$, becomes $(r + \gamma\sigma_Y^2 - \omega^2\sigma_Y + \omega^n\sigma_Y)dt + \sigma_Y dW^n$.

Its subjective Sharpe ratio is therefore $\lambda^n = \gamma\sigma_Y + \omega^n - \omega^2$.

The certainty-equivalent return $r + \pi\sigma_Y\lambda^n - \frac{1}{2}\gamma\pi^2\sigma_Y^2$ is maximized at $\pi^n = \lambda^n/(\gamma\sigma_Y) = 1 + (\omega^n - \omega^2)/(\gamma\sigma_Y)$, which is (35.15).

The maximized value is $\text{CE}^n = r + (\lambda^n)^2/(2\gamma)$.

Step 2. Under Q^n, wealth satisfies $dA = A\left[(r + \pi\sigma_Y\lambda^n - c)dt + \pi\sigma_Y dW^n\right]$.

With $V(A) = (\kappa A)^{1-\gamma}/(1-\gamma)$, we have $((1-\gamma)V)^{1/(1-\gamma)} = \kappa A$ and the HJB equation is

$$0 = \max_{c,\pi}\left\{f(cA, V(A)) + V'(A)A(r + \pi\sigma_Y\lambda^n - c) + \tfrac{1}{2}V''(A)A^2\pi^2\sigma_Y^2\right\}.$$

Dividing by $(\kappa A)^{1-\gamma}$ gives

$$0 = \max_{c,\pi}\left\{\frac{\beta}{1-\rho}\left[\left(\frac{c}{\kappa}\right)^{1-\rho} - 1\right] + r + \pi\sigma_Y\lambda^n - c - \tfrac{1}{2}\gamma\pi^2\sigma_Y^2\right\}.$$

The first-order condition for c is $\beta(c/\kappa)^{-\rho}/\kappa = 1$, which implies $(c/\kappa)^{1-\rho} = c/\beta$.

Substituting this and the optimal π into the HJB equation gives

$$0 = \frac{c^n - \beta}{1-\rho} + \text{CE}^n - c^n,$$

which rearranges to $c^n = \beta - (1-\rho)(\text{CE}^n - c^n) = \beta - (1-\rho)\mu_V^n$.

Solving the same equation for c^n gives $\rho c^n = \beta - (1-\rho)\text{CE}^n$, so $\mu_V^n = \text{CE}^n - c^n = (\text{CE}^n - \beta)/\rho$.

The quantity μ_V^n is the subjective certainty-equivalent growth rate of wealth.

Under Q^n, $d\log A$ has drift $r + \pi\sigma_Y\lambda^n - c - \frac{1}{2}\pi^2\sigma_Y^2$ and volatility $\pi\sigma_Y$.

Adding the risk adjustment $\frac{1}{2}(1-\gamma)\pi^2\sigma_Y^2$ and evaluating at π^n gives $\text{CE}^n - c^n$.

Step 3. For agent 2, $\lambda^2 = \gamma\sigma_Y$, so $\text{CE}^2 = r + \frac{1}{2}\gamma\sigma_Y^2$.

Substituting (35.13) gives $\text{CE}^2 = \beta + \rho\left(\mu_Y + \omega^2\sigma_Y + \frac{1}{2}(1-\gamma)\sigma_Y^2\right)$, so $\mu_V^2 = \mu_Y + \omega^2\sigma_Y + \frac{1}{2}(1-\gamma)\sigma_Y^2$.

Then $c^2 = \beta - (1-\rho)\mu_V^2$ is (35.14).

Step 4. Taking differences,

$$(y^2)^{-1} - (y^1)^{-1} = c^2 - c^1 = (1-\rho)(\mu_V^1 - \mu_V^2) = \frac{1-\rho}{\rho}\left(\text{CE}^1 - \text{CE}^2\right).$$

Since $\lambda^1 - \lambda^2 = \omega^1 - \omega^2$ and $\lambda^1 + \lambda^2 = 2\gamma\sigma_Y + \omega^1 - \omega^2$,

$$\text{CE}^1 - \text{CE}^2 = \frac{(\lambda^1)^2 - (\lambda^2)^2}{2\gamma} = (\omega^1 - \omega^2)\sigma_Y + \frac{(\omega^1 - \omega^2)^2}{2\gamma},$$

which gives (35.12).

The factor $1/\rho$ reflects feedback from consumption to wealth growth: c^n depends on μ_V^n, and $\mu_V^n = \text{CE}^n - c^n$ depends on c^n.

Here is a numerical check at $\gamma = 5$, $\rho = 0.67$, $\omega^1 = 0.25$, $\omega^2 = 0$, $\sigma_Y = 0.02$, $\mu_Y = 0.02$, and $\beta = 0.05$.

```
γ, ρ, β, μ_y, σ_y = 5.0, 0.67, 0.05, 0.02, 0.02
ω_1, ω_2 = 0.25, 0.0

r = β + ρ * (μ_y + ω_2 * σ_y + 0.5 * (1 - γ) * σ_y**2) - 0.5 * γ * σ_y**2

def consumption_rate(ω_n):
    λ_n = γ * σ_y + ω_n - ω_2
    CE_n = r + λ_n**2 / (2 * γ)
    return (β - (1 - ρ) * CE_n) / ρ

c_1, c_2 = consumption_rate(ω_1), consumption_rate(ω_2)
y0_inv = β - (1 - ρ) * (μ_y + ω_2 * σ_y + 0.5 * (1 - γ) * σ_y**2)
Δω = ω_1 - ω_2
formula = (1 - ρ) / ρ * (Δω * σ_y + Δω**2 / (2 * γ))

print(f"c^2 from the HJB solution:      {c_2:.6f}")
print(f"(y(0))^(-1) from eq:wc_ratio:   {y0_inv:.6f}")
print(f"c^2 - c^1 from the HJB:         {c_2 - c_1:.6f}")
print(f"eq:consumption_rates:           {formula:.6f}")
```

```
c^2 from the HJB solution:      0.043664
(y(0))^(-1) from eq:wc_ratio:   0.043664
c^2 - c^1 from the HJB:         0.005541
eq:consumption_rates:           0.005541
```

35.5.3 Two comparative statics

Survival depends on γ, ρ, and the signal-to-noise ratios ω^1/σ_Y and ω^2/σ_Y, not on ω^1, ω^2, and σ_Y separately.

The survival conditions do not depend on β or μ_Y, which affect the level of consumption and prices but not relative wealth dynamics at the boundary.

```
def portfolio_return_diff(ω_1, ω_2, γ, σ_y):
    """
    Difference in expected log portfolio returns at the boundary.
    """
    Δω = ω_1 - ω_2
    risky_share_diff = Δω / (γ * σ_y)
    risk_premium = γ * σ_y**2 - ω_2 * σ_y
    volatility_term = (Δω / γ) * (σ_y + 0.5 * Δω / γ)
    return risky_share_diff * risk_premium - volatility_term


def saving_channel(ω_1, ω_2, γ, ρ, σ_y):
    """
    Difference in consumption-wealth ratios at the boundary.
    """
    Δω = ω_1 - ω_2
    subjective_return_diff = Δω * σ_y + Δω**2 / (2 * γ)
    return (1 - ρ) / ρ * subjective_return_diff


def boundary_drift(ω_1, ω_2, γ, ρ, σ_y):
    """
    Boundary drift m_9(0+) when agent 1 becomes negligible.
```

(continues on next page)

(continued from previous page)

```
    By eq:drift_wealth, this is y times the wealth-growth
    differential in eq:wealth_decomp.
    Positive drift means agent 1 survives (repelling boundary).
    """
    return y * (
        portfolio_return_diff(ω_1, ω_2, y, σ_y)
        + saving_channel(ω_1, ω_2, y, ρ, σ_y)
    )
```

The case $v \nearrow 1$ is symmetric.

Swapping the agents' labels turns ϑ into $-\vartheta$, so $\lim_{v \nearrow 1} m_\vartheta(v)$ equals `-boundary_drift(ω_2, ω_1, y, ρ, σ_y)`.

35.5.4 Risk aversion versus the IES

The two boundary drifts satisfy a simple identity.

Proposition 35.5.1

$$\lim_{v \searrow 0} m_\vartheta(v) - \lim_{v \nearrow 1} m_\vartheta(v) = (\omega^1 - \omega^2)^2 \left(\frac{1}{\rho} - \frac{1}{\gamma} \right).$$

To derive it, let $\Delta = \omega^1 - \omega^2$.

Combining (35.9), (35.10), (35.11), and (35.12) gives

$$\lim_{v \searrow 0} m_\vartheta(v) = \Delta(\gamma \sigma_Y - \omega^2) - \Delta \sigma_Y - \frac{\Delta^2}{2\gamma} + \frac{1-\rho}{\rho} \left(\gamma \Delta \sigma_Y + \frac{\Delta^2}{2} \right).$$

At $v \nearrow 1$ the roles of the agents are reversed, which replaces Δ by $-\Delta$, ω^2 by ω^1, and m_ϑ by $-m_\vartheta$:

$$\lim_{v \nearrow 1} m_\vartheta(v) = -\left[-\Delta(\gamma \sigma_Y - \omega^1) + \Delta \sigma_Y - \frac{\Delta^2}{2\gamma} + \frac{1-\rho}{\rho} \left(-\gamma \Delta \sigma_Y + \frac{\Delta^2}{2} \right) \right].$$

In the difference, the terms linear in σ_Y cancel.

The risk premium terms add to $\Delta(\omega^1 - \omega^2) = \Delta^2$, the volatility terms add to $-\Delta^2/\gamma$, and the saving terms add to $\Delta^2(1-\rho)/\rho$.

The sum is $\Delta^2(1/\rho - 1/\gamma)$.

The identity sorts the (γ, ρ) plane:

- Outcome (a), coexistence, requires $\lim_{v \searrow 0} m_\vartheta > 0 > \lim_{v \nearrow 1} m_\vartheta$, so it requires $\gamma > \rho$.
- Outcome (d), in which each agent dominates with positive probability, requires the reverse inequalities, so it requires $\gamma < \rho$.
- On the diagonal $\gamma = \rho$ the two boundary drifts are equal, so when $|\omega^1| \neq |\omega^2|$ exactly one agent dominates.
- Outcomes (b) and © can occur on either side of the diagonal.

Epstein-Zin agents prefer early resolution of uncertainty when $\gamma > \rho$ [Epstein and Zin, 1989, Kreps and Porteus, 1978].

Both agents can survive only in economies populated by agents with that preference.

Borovička [2020] notes that the coexistence region in the top-left panel of his Figure 2 lies where $\gamma > \rho$; *Proposition 35.5.1* shows that this holds for all belief distortions and all σ_Y.

Neither the IES alone nor risk aversion alone settles the question.

An IES above one is not sufficient for coexistence, because $\gamma < \rho < 1$ rules it out.

An IES above one is not necessary either, as the next section shows.

Exercise 35.5.4

1. Carry out the algebra behind *Proposition 35.5.1*, starting from (35.11) and (35.12).
2. Show that when $\gamma = \rho$ both boundary drifts equal $\frac{1}{2}\left[(\omega^2)^2 - (\omega^1)^2\right]$.
3. Use `boundary_drift` to verify the identity at 10,000 randomly drawn parameter vectors $(\omega^1, \omega^2, \gamma, \rho, \sigma_Y)$.

Solution

Part 1. Multiplying (35.11) by γ gives

$$\gamma \lim_{\upsilon \searrow 0}[m_R^1 - m_R^2] = \frac{\Delta}{\sigma_Y}(\gamma\sigma_Y^2 - \omega^2\sigma_Y) - \Delta\left(\sigma_Y + \frac{\Delta}{2\gamma}\right) = \Delta(\gamma\sigma_Y - \omega^2) - \Delta\sigma_Y - \frac{\Delta^2}{2\gamma}.$$

Multiplying (35.12) by γ gives $\frac{1-\rho}{\rho}\left(\gamma\Delta\sigma_Y + \frac{\Delta^2}{2}\right)$.

Their sum is $\lim_{\upsilon \searrow 0} m_\vartheta$ by (35.9) and (35.10).

At $\upsilon \nearrow 1$, agent 2 is the negligible agent and agent 1 sets prices.

Applying the same formulas with the labels swapped gives $\gamma \lim_{\upsilon \nearrow 1}[m_A^2 - m_A^1]$, which equals $-\lim_{\upsilon \nearrow 1} m_\vartheta$.

Swapping labels replaces Δ by $-\Delta$ and ω^2 by ω^1, so

$$-\lim_{\upsilon \nearrow 1} m_\vartheta = -\Delta(\gamma\sigma_Y - \omega^1) + \Delta\sigma_Y - \frac{\Delta^2}{2\gamma} + \frac{1-\rho}{\rho}\left(-\gamma\Delta\sigma_Y + \frac{\Delta^2}{2}\right).$$

Adding the two expressions term by term:

$$\Delta(\omega^1 - \omega^2) - \frac{\Delta^2}{\gamma} + \frac{1-\rho}{\rho}\Delta^2 = \Delta^2\left(1 - \frac{1}{\gamma} + \frac{1}{\rho} - 1\right) = \Delta^2\left(\frac{1}{\rho} - \frac{1}{\gamma}\right).$$

Part 2. Setting $\rho = \gamma$, the saving term becomes $(1-\gamma)\Delta\sigma_Y + (1-\gamma)\Delta^2/(2\gamma)$.

Adding it to the portfolio term, the terms in σ_Y and in $\Delta^2/(2\gamma)$ cancel and leave

$$\lim_{\upsilon \searrow 0} m_\vartheta = -\Delta\omega^2 - \frac{\Delta^2}{2} = -\frac{(\omega^1 - \omega^2)(\omega^1 + \omega^2)}{2} = \frac{1}{2}\left[(\omega^2)^2 - (\omega^1)^2\right].$$

By Part 1, the drift at $\upsilon \nearrow 1$ is the same.

When $\gamma = \rho$ the saving term is not zero unless $\rho = 1$; the three channels net to a constant that does not depend on γ or σ_Y.

Part 3.

```
rng = np.random.default_rng(0)
n_draws = 10_000
max_err = 0.0
for _ in range(n_draws):
    ω_1, ω_2 = rng.normal(0.0, 0.3, size=2)
    γ = rng.uniform(0.05, 30.0)
    ρ = rng.uniform(0.05, 5.0)
    σ_y = rng.uniform(0.005, 0.2)
    m_0 = boundary_drift(ω_1, ω_2, γ, ρ, σ_y)
    m_1 = -boundary_drift(ω_2, ω_1, γ, ρ, σ_y)
    gap = (ω_1 - ω_2)**2 * (1 / ρ - 1 / γ)
    max_err = max(max_err, abs(m_0 - m_1 - gap))

print(f"largest error over {n_draws} draws: {max_err:.2e}")

for γ in [0.5, 2.0, 7.0]:
    m_sep = boundary_drift(0.25, -0.1, γ, γ, 0.02)
    print(f"γ = ρ = {γ}: m_ϑ(0+) = {m_sep:.6f},  "
          f"[(ω^2)^2 - (ω^1)^2] / 2 = {0.5 * (0.1**2 - 0.25**2):.6f}")
```

```
largest error over 10000 draws: 1.07e-14
γ = ρ = 0.5: m_ϑ(0+) = -0.026250,  [(ω^2)^2 - (ω^1)^2] / 2 = -0.026250
γ = ρ = 2.0: m_ϑ(0+) = -0.026250,  [(ω^2)^2 - (ω^1)^2] / 2 = -0.026250
γ = ρ = 7.0: m_ϑ(0+) = -0.026250,  [(ω^2)^2 - (ω^1)^2] / 2 = -0.026250
```

35.6 Survival regions

A central contribution of Borovička [2020] is the characterization of survival regions in the (γ, ρ) plane.

Under separable preferences, $\gamma = \rho$, the agent with more accurate beliefs always dominates.

Under recursive preferences, all four outcomes in *Proposition 35.4.1* can occur.

Figure 2 in the paper studies the case where agent 2 has correct beliefs, so $\omega^2 = 0$.

The next cell follows that figure.

For the limiting panel $|\omega^1|/\sigma_Y \to \infty$ we use the limits of $m_\vartheta(0+)/(\omega^1)^2$ and $m_\vartheta(1-)/(\omega^1)^2$, which Exercise 35.6.1 derives.

```
def limit_boundary_drifts(γ, ρ):
    """
    Limits of m_ϑ(0+) / (ω^1)^2 and m_ϑ(1-) / (ω^1)^2
    as |ω^1| / σ_y -> ∞, with ω^2 = 0.
    """
    m_0 = (γ - (1 + γ) * ρ) / (2 * γ * ρ)
    m_1 = ((1 - γ) * ρ - γ) / (2 * γ * ρ)
    return m_0, m_1
```

Each panel plots two curves in the (γ, ρ) plane for a different value of agent 1's belief distortion ω^1 (agent 2 has correct beliefs, $\omega^2 = 0$).

- The dashed curve (blue) is where the boundary drift at $\upsilon = 0$ equals zero – condition (i) in *Proposition 35.4.1*.
- The solid curve (red) is where the boundary drift at $\upsilon = 1$ equals zero – condition (ii).

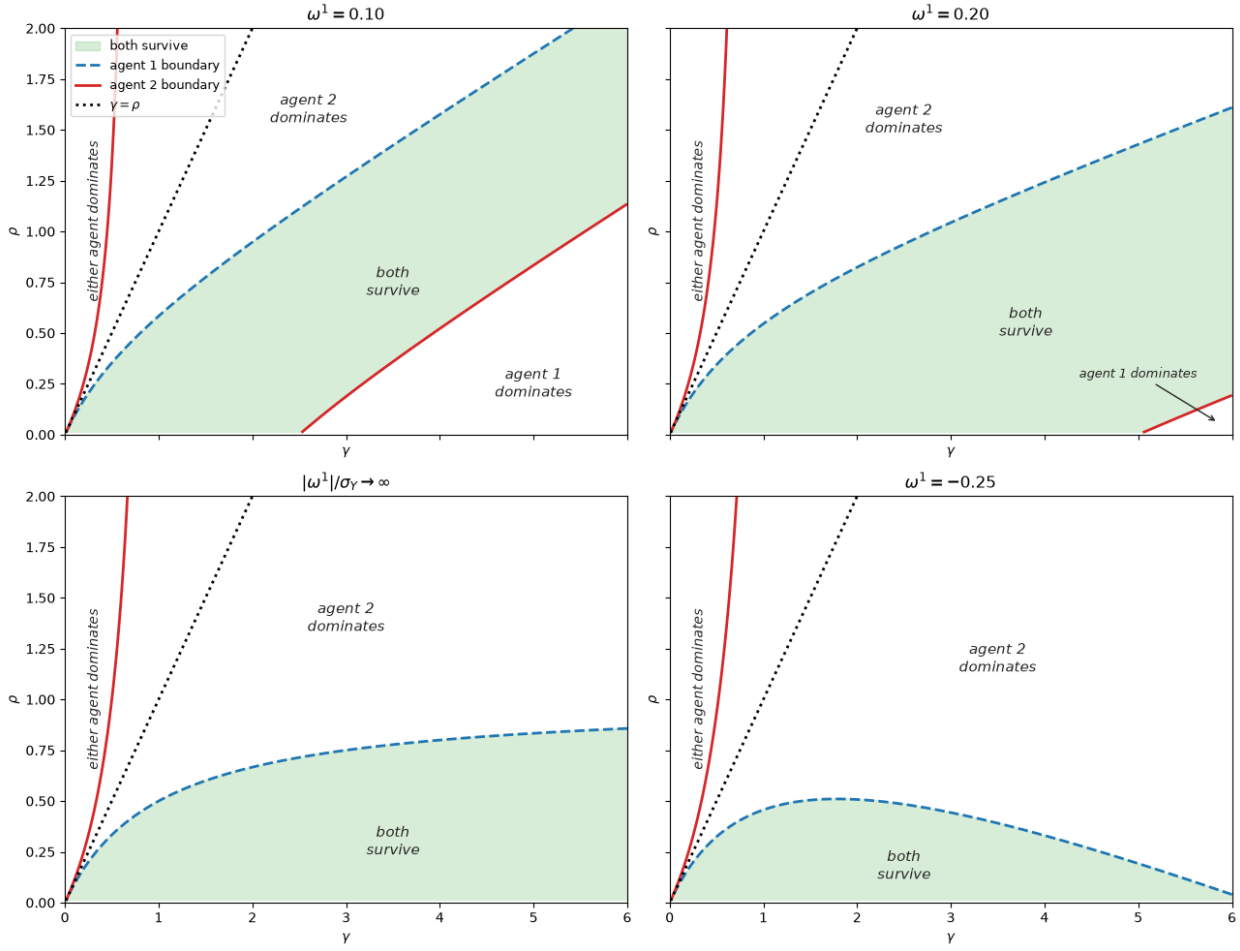


Fig. 35.1: Survival regions corresponding to Figure 2 in Borovička (2020)

- The shaded region is where both agents survive.
- The dotted diagonal $\gamma = \rho$ is the separable CRRA case, along which the agent with more accurate beliefs always dominates.

The next cell reports the share of the plotted window occupied by each outcome, the share of the coexistence region with $\rho > 1$ (IES < 1), and the largest ρ in the coexistence region.

```
names = ["ω¹ = 0.10", "ω¹ = 0.20", "|ω¹|/σ_y → ∞", "ω¹ = -0.25"]
print(f"{'panel':<15}{'both':>7}{'agent 1':>9}{'agent 2':>9}{'either':>8}"
      f"{'share ρ > 1':>13}{'max ρ':>8}")
n_wrong_side = 0
for idx, name in enumerate(names):
    both, ag1_dom, ag2_dom, either = region_masks[idx]
    print(f"{name:<15}{both.mean():7.3f}{ag1_dom.mean():9.3f}"
          f"{ag2_dom.mean():9.3f}{either.mean():8.3f}"
          f"{(both & (R > 1)).sum() / both.sum():13.3f}{R[both].max():8.3f}")
    n_wrong_side += (both & (R >= G)).sum() + (either & (R <= G)).sum()

print(f"\ngrid points with coexistence and ρ ≥ γ, "
      f"or with outcome (d) and ρ ≤ γ: {n_wrong_side}")
```

```
panel             both  agent 1  agent 2  either  share ρ > 1   max ρ
ω¹ = 0.10        0.439    0.171    0.323   0.066         0.421   2.000
ω¹ = 0.20        0.486    0.008    0.436   0.070         0.173   1.606
|ω¹|/σ_y → ∞     0.335    0.000    0.590   0.075         0.000   0.853
ω¹ = -0.25       0.167    0.000    0.755   0.078         0.000   0.509

grid points with coexistence and ρ ≥ γ, or with outcome (d) and ρ ≤ γ: 0
```

Every shaded region lies below the diagonal, and every region where either agent can dominate lies above it, as *Proposition 35.5.1* requires.

Moderate optimism ($\omega^1 = 0.10$) produces a wide coexistence region that extends across most of the γ range.

Coexistence does not require IES > 1: 42 percent of the shaded area in this panel has $\rho > 1$, including the point $\gamma = 5$, $\rho = 1.5$.

Stronger optimism ($\omega^1 = 0.20$) rotates both curves clockwise.

The agent 1 curve falls, so agent 2 dominates in a larger part of the plane above it.

The lower branch of the agent 2 curve meets the axis $\rho = 0$ at $\gamma = \omega^1/(2\sigma_Y)$, which is 2.5 when $\omega^1 = 0.10$ and 5 when $\omega^1 = 0.20$.

At that value of γ, the two agents' subjective certainty-equivalent returns at the boundary $v = 1$ are equal, and the saving channel changes sign there.

Agent 1 therefore dominates only in a sliver at the bottom right of the $\omega^1 = 0.20$ panel, and the coexistence region grows from 44 to 49 percent of the plotted window.

In the limit $|\omega^1|/\sigma_Y \to \infty$ (bottom-left), the boundary curves are $\rho = \gamma/(1+\gamma)$ and, for $\gamma < 1$, $\rho = \gamma/(1-\gamma)$ (see Exercise 35.6.1).

Both agents survive exactly when $\rho < \gamma/(1+\gamma)$, the region below the agent 1 curve.

That curve rises toward $\rho = 1$ as γ grows, so in this limit coexistence requires IES > 1.

The agent 2 curve separates outcome © from outcome (d), and agent 1 never dominates.

Pessimistic distortions ($\omega^1 = -0.25$, bottom-right) can also survive, but only in a much narrower part of the parameter space: 17 percent of the plotted window, all of it with $\rho < 0.51$.

The coexistence region ends at $\gamma = |\omega^1|/(2\sigma_Y) = 6.25$, just beyond the plotted range.

For larger γ her short position is too small, $\pi^1 > -1$, for her subjective certainty-equivalent return to exceed agent 2's.

With IES > 1 the saving channel then works against her, and with IES < 1 it is too weak to offset her portfolio losses.

Exercise 35.6.1

Let $\omega^2 = 0$.

1. Show that, holding γ and ρ fixed, $m_\vartheta(0+)/(\omega^1)^2$ and $m_\vartheta(1-)/(\omega^1)^2$ depend on ω^1 and σ_Y only through σ_Y/ω^1, and that as $\sigma_Y/\omega^1 \to 0$ they converge to the expressions in `limit_boundary_drifts`.
2. Deduce that the boundary curves in the limit are $\rho = \gamma/(1+\gamma)$ and, for $\gamma < 1$, $\rho = \gamma/(1-\gamma)$. Classify the four outcomes in the limit and show that agent 1 never dominates.
3. Verify numerically that the zero curve of `boundary_drift` at $\upsilon = 0$ converges to $\rho = \gamma/(1+\gamma)$ as ω^1/σ_Y grows.

Hint: $m_\vartheta(0+)$ is affine in $1/\rho$, so for each γ its zero can be computed in closed form.

Solution

Part 1. Let $\Delta = \omega^1$ and $s = \sigma_Y/\omega^1$.

From the expression for $\lim_{\upsilon \searrow 0} m_\vartheta$ in the derivation of *Proposition 35.5.1*, with $\omega^2 = 0$,

$$\frac{m_\vartheta(0+)}{\Delta^2} = (\gamma - 1)s - \frac{1}{2\gamma} + \frac{1-\rho}{\rho}\left(\gamma s + \frac{1}{2}\right) \longrightarrow \frac{\gamma - (1+\gamma)\rho}{2\gamma\rho}.$$

Similarly,

$$\frac{m_\vartheta(1-)}{\Delta^2} = (\gamma - 1)s - 1 + \frac{1}{2\gamma} + \frac{1-\rho}{\rho}\left(\gamma s - \frac{1}{2}\right) \longrightarrow \frac{(1-\gamma)\rho - \gamma}{2\gamma\rho}.$$

The difference of the two limits is $(\gamma - \rho)/(\gamma\rho) = 1/\rho - 1/\gamma$, as *Proposition 35.5.1* requires.

Part 2. The first limit is positive if and only if $\rho < \gamma/(1+\gamma)$.

The second limit is negative if and only if $(1-\gamma)\rho < \gamma$.

That inequality holds for every $\rho > 0$ when $\gamma \geq 1$, and holds if and only if $\rho < \gamma/(1-\gamma)$ when $\gamma < 1$.

Since $\gamma/(1+\gamma) < \gamma/(1-\gamma)$ for $\gamma \in (0,1)$, condition (i) implies condition (ii).

Hence:

- (a) both survive if and only if $\rho < \gamma/(1+\gamma)$;
- (b) agent 1 dominates nowhere, because (i) and (ii') cannot hold together;
- © agent 2 dominates if $\rho > \gamma/(1+\gamma)$ and either $\gamma \geq 1$ or $\rho < \gamma/(1-\gamma)$;
- (d) either agent dominates if $\gamma < 1$ and $\rho > \gamma/(1-\gamma)$.

Part 3. Write $m_\vartheta(0+) = \gamma\,[P + (1/\rho - 1)S]$, where P is the portfolio-return difference and S is the bracket in (35.12).

The drift vanishes at $1/\rho = 1 - P/S$, so $\rho = S/(S-P)$.

```
def lower_zero_curve(ω_1, ω_2, σ_y, γ):
    """
    The ρ at which m_ϑ(0+) = 0, using that m_ϑ(0+) is affine in 1/ρ.
    """
    Δω = ω_1 - ω_2
    P = portfolio_return_diff(ω_1, ω_2, γ, σ_y)
    S = Δω * σ_y + Δω**2 / (2 * γ)
    return S / (S - P)


γ_check = np.array([0.5, 1.0, 2.0, 5.0, 20.0])
σ_y = 0.02
print("γ:               ", γ_check)
for ratio in [5, 50, 500, 5000]:
    ρ_zero = lower_zero_curve(ratio * σ_y, 0.0, σ_y, γ_check)
    print(f"ω¹/σ_y = {ratio:>5}: ", np.round(ρ_zero, 4))
print("γ / (1 + γ):     ", np.round(γ_check / (1 + γ_check), 4))
```

```
γ:               [ 0.5  1.   2.   5.  20. ]
ω¹/σ_y =     5:  [0.3529 0.5833 0.9474 1.875  6.2069]
ω¹/σ_y =    50:  [0.3355 0.5098 0.7013 0.9677 1.6514]
ω¹/σ_y =   500:  [0.3336 0.501  0.6702 0.8472 1.0247]
ω¹/σ_y =  5000:  [0.3334 0.5001 0.667  0.8347 0.9596]
γ / (1 + γ):     [0.3333 0.5    0.6667 0.8333 0.9524]
```

35.7 Three survival channels

The decomposition above can be visualized directly.

```
def decompose_survival(ω_1, ω_2, γ_grid, ρ, σ_y):
    """
    Decompose the wealth-growth differential in proposition 3.4.

    The total equals lim (m_A^1 - m_A^2) = m_ϑ(0+) / γ as υ -> 0.
    """
    Δω = ω_1 - ω_2
    risk_premium_term = Δω * (γ_grid * σ_y - ω_2) / γ_grid
    volatility_term = -(Δω / γ_grid) * (σ_y + 0.5 * Δω / γ_grid)
    saving_term = (1 - ρ) / ρ * (Δω * σ_y + Δω**2 / (2 * γ_grid))
    total = risk_premium_term + volatility_term + saving_term
    return risk_premium_term, volatility_term, saving_term, total


ω_1 = 0.25
ω_2 = 0.0
ρ = 0.67
σ_y = 0.02
γ_grid = np.linspace(0.5, 15.0, 300)

risk_term, vol_term, save_term, total = decompose_survival(
    ω_1, ω_2, γ_grid, ρ, σ_y
)

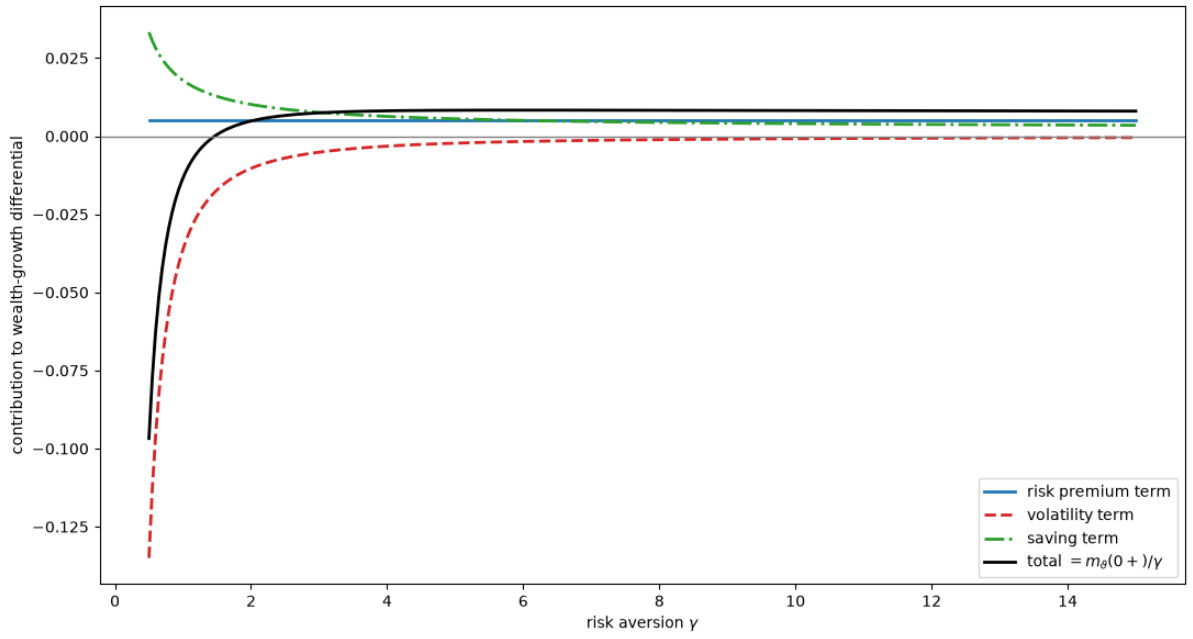
fig, ax = plt.subplots(figsize=(11, 6))
```

(continues on next page)

(continued from previous page)

```
ax.plot(γ_grid, risk_term, color="C0", lw=2, label="risk premium term")
ax.plot(γ_grid, vol_term, "--", color="C3", lw=2, label="volatility term")
ax.plot(γ_grid, save_term, "-.", color="C2", lw=2, label="saving term")
ax.plot(γ_grid, total, color="black", lw=2,
        label=r"total $= m_\vartheta(0+)/\gamma$")
ax.axhline(0, color="gray", lw=1)
ax.set_xlabel(r"risk aversion $\gamma$")
ax.set_ylabel("contribution to wealth-growth differential")
ax.legend()
plt.tight_layout()
plt.show()

γ_star = brentq(
    lambda γ: decompose_survival(ω_1, ω_2, γ, ρ, σ_y)[3], 0.5, 15.0
)
print(f"the total crosses zero at γ = {γ_star:.3f}")
```




```
the total crosses zero at γ = 1.465
```

This figure decomposes the wealth-growth differential $\lim_{\upsilon \searrow 0}[m_A^1 - m_A^2] = \lim_{\upsilon \searrow 0} m_\vartheta/\gamma$ into three terms for an optimistic agent ($\omega^1 = 0.25, \omega^2 = 0$) with IES $= 1/\rho \approx 1.49$ and $\sigma_Y = 0.02$.

- The risk premium term (blue) is positive throughout because the optimistic agent overweights the risky asset and earns the equity premium.
- The volatility term (red dashed) is negative and large at low γ, reflecting the cost of holding a volatile portfolio.
- The saving term (green dash-dot) is positive when IES > 1 because the optimistic agent perceives a high return on wealth and saves more aggressively.
- The total (black) crosses zero at $\gamma \approx 1.47$; below that value the volatility penalty dominates, condition (i) fails, and the boundary $\upsilon = 0$ attracts.

35.8 Varying the IES

The sign of the saving term is pinned down by the IES.

```
fig, axes = plt.subplots(1, 3, figsize=(16, 4.5), sharey=True)

ω_1 = 0.25
ω_2 = 0.0
σ_y = 0.02
y_grid = np.linspace(0.5, 25.0, 300)

ies_values = [0.5, 1.0, 1.5]

for idx, ies in enumerate(ies_values):
    ρ = 1.0 / ies
    risk_term, vol_term, save_term, total = decompose_survival(
        ω_1, ω_2, y_grid, ρ, σ_y
    )

    ax = axes[idx]
    ax.plot(y_grid, risk_term, color="C0", lw=2, label="risk premium")
    ax.plot(y_grid, vol_term, "--", color="C3", lw=2, label="volatility")
    ax.plot(y_grid, save_term, "-.", color="C2", lw=2, label="saving")
    ax.plot(y_grid, total, color="black", lw=2, label="total")
    ax.axhline(0, color="gray", lw=1)
    ax.set_title(f"IES = {ies:.1f}", fontsize=12)
    ax.set_xlabel(r"risk aversion $\gamma$")
    ax.set_ylabel("contribution")

axes[0].legend(fontsize=9)
plt.tight_layout()
plt.show()

for ies in ies_values:
    y_star = brentq(
        lambda y: decompose_survival(ω_1, ω_2, y, 1 / ies, σ_y)[3], 0.5, 25.0
    )
    print(f"IES = {ies:.1f}: the total crosses zero at y = {y_star:.2f}")
```

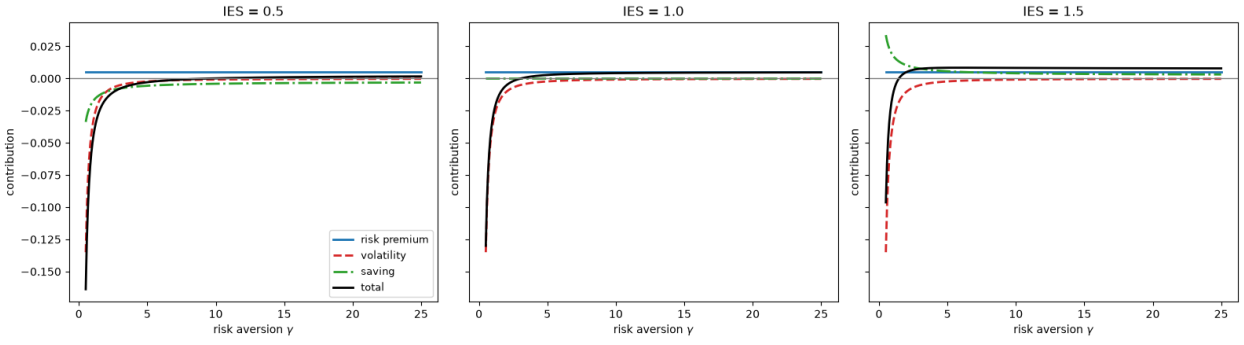


Fig. 35.2: Boundary decomposition for different IES values

```
IES = 0.5: the total crosses zero at y = 9.56
IES = 1.0: the total crosses zero at y = 3.05
IES = 1.5: the total crosses zero at y = 1.45
```

Each panel shows the same three-term decomposition as the previous figure, but now for three different values of the IES ($\omega^1 = 0.25$, $\omega^2 = 0$, $\sigma_Y = 0.02$).

- Left panel (IES $= 0.5$): the saving term is negative, so the optimistic agent actually saves less, working against survival. The risk premium term still makes the total positive for $\gamma > 9.56$, so condition (i) holds with IES below one.
- Center panel (IES $= 1.0$): the saving term vanishes entirely, so only the portfolio return and volatility channels remain.
  - This eliminates the saving channel but does not by itself reproduce the full separable CRRA benchmark, which requires $\gamma = \rho$ (i.e., IES $= 1/\gamma$), not merely $\rho = 1$.
- Right panel (IES $= 1.5$): the saving term is positive and shifts the total drift upward, expanding the range of γ values for which the optimistic agent survives.

The zero crossing moves from $\gamma = 9.56$ to 3.05 to 1.45 as the IES rises from 0.5 to 1 to 1.5.

35.9 Asymptotic results

Borovička derives several useful asymptotic results.

1. As $\gamma \searrow 0$, each agent dominates with strictly positive probability.
2. As $\gamma \nearrow \infty$, the relatively more optimistic agent dominates.
3. As $\rho \searrow 0$, the relatively more optimistic agent always survives.
   - The relatively more pessimistic agent can also survive when risk aversion is sufficiently low.
4. As $\rho \nearrow \infty$, a nondegenerate long-run equilibrium cannot exist.

The next figure illustrates the first result by plotting both boundary drifts as γ becomes small.

```
ω_1 = 0.25
ω_2 = 0.0
ρ = 0.67
σ_y = 0.02
γ_grid = np.linspace(0.05, 5.0, 300)

drift_at_0 = np.array([boundary_drift(ω_1, ω_2, γ, ρ, σ_y) for γ in γ_grid])
drift_at_1 = np.array([-boundary_drift(ω_2, ω_1, γ, ρ, σ_y) for γ in γ_grid])

fig, ax = plt.subplots(figsize=(10, 5))
ax.plot(γ_grid, drift_at_0, color="C0", lw=2, label=r"$\upsilon \to 0$")
ax.plot(γ_grid, drift_at_1, "--", color="C3", lw=2, label=r"$\upsilon \to 1$")
ax.axhline(0, color="gray", lw=1)
ax.set_xlabel(r"risk aversion $\gamma$")
ax.set_ylabel("boundary drift")
ax.legend()
plt.tight_layout()
plt.show()
```

This figure plots the two boundary drifts as a function of γ ($\omega^1 = 0.25$, $\omega^2 = 0$, IES ≈ 1.49).

- The solid blue curve is the drift m_ϑ at $\upsilon \to 0$ (agent 1 near extinction); coexistence requires this to be positive (condition (i)).

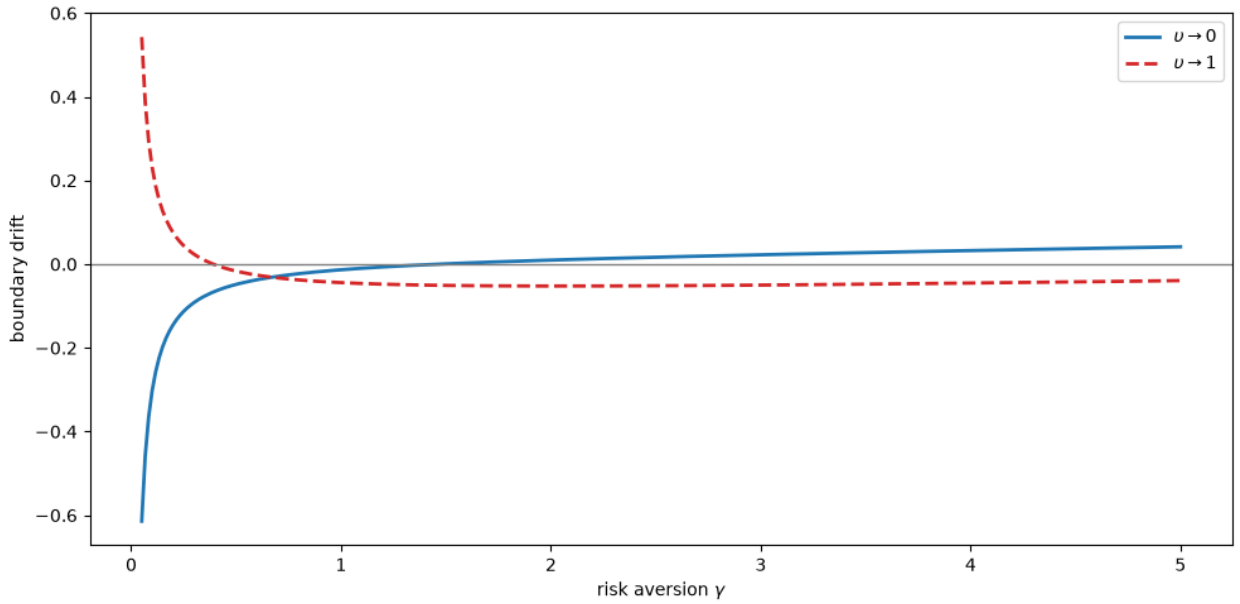


Fig. 35.3: Boundary drifts for small risk aversion

- The dashed red curve is the drift m_ϑ at $\upsilon \to 1$ (agent 2 near extinction); coexistence requires this to be negative (condition (ii)).

The figure illustrates asymptotic result 1.

For small γ, the blue curve is negative and the red curve is positive.

Both boundaries are attracting: near $\upsilon = 0$ the negative drift pulls υ toward 0, and near $\upsilon = 1$ the positive drift pushes υ toward 1.

This is outcome (d) in *Proposition 35.4.1*: neither boundary is repelling, so whichever agent happens to get ahead early will dominate, with each agent having strictly positive probability of dominance depending on the realized Brownian path.

The red curve crosses zero first, at $\gamma \approx 0.39$, while the blue curve is still negative.

For γ between about 0.39 and 1.47, the boundary $\upsilon = 1$ repels and the boundary $\upsilon = 0$ attracts, so agent 2 dominates – outcome ©.

The two curves cross each other at $\gamma = \rho = 0.67$, as *Proposition 35.5.1* requires; their common value there is $-\frac{1}{2}(\omega^1)^2 = -0.03125$.

The blue curve crosses zero at $\gamma \approx 1.47$.

Beyond that point both boundaries repel and both agents survive – outcome (a).

Outside the plotted range, the red curve crosses zero again at $\gamma \approx 10.7$, and for larger γ agent 1 dominates – outcome (b), in line with asymptotic result 2.

Exercise 35.9.1

Take the parameters of Fig. 35.3: $\omega^1 = 0.25$, $\omega^2 = 0$, $\rho = 0.67$, $\sigma_Y = 0.02$.

1. Show that $2\gamma \lim_{\upsilon \searrow 0} m_\vartheta$ and $2\gamma \lim_{\upsilon \nearrow 1} m_\vartheta$ are quadratic polynomials in γ, and use them to find every $\gamma > 0$ at which a boundary drift changes sign.
2. Confirm the roots with `boundary_drift`.

3. Label the outcome of *Proposition 35.4.1* on each interval between consecutive roots.

Solution

Let $\Delta = \omega^1, \omega^2 = 0$, and $k = (1-\rho)/\rho$.

Multiplying the two boundary drifts from the derivation of *Proposition 35.5.1* by 2γ gives

$$2\gamma \lim_{\upsilon \searrow 0} m_\vartheta = \frac{2\Delta\sigma_Y}{\rho}\gamma^2 + \left(k\Delta^2 - 2\Delta\sigma_Y\right)\gamma - \Delta^2,$$

$$2\gamma \lim_{\upsilon \nearrow 1} m_\vartheta = \frac{2\Delta\sigma_Y}{\rho}\gamma^2 - \left(2\Delta\sigma_Y + \frac{1+\rho}{\rho}\Delta^2\right)\gamma + \Delta^2.$$

When $\Delta > 0$, the first quadratic has one positive root because its constant term is negative.

The second has two positive roots when its discriminant is positive, because its constant term and its leading coefficient are both positive.

```
ω_1, ω_2, ρ, σ_y = 0.25, 0.0, 0.67, 0.02
Δ, k = ω_1 - ω_2, (1 - ρ) / ρ

roots_0 = np.roots([2 * Δ * σ_y / ρ, k * Δ**2 - 2 * Δ * σ_y, -Δ**2])
roots_1 = np.roots([2 * Δ * σ_y / ρ,
                    -(2 * Δ * σ_y + (1 + ρ) / ρ * Δ**2), Δ**2])
roots_0 = np.sort(roots_0[roots_0 > 0])
roots_1 = np.sort(roots_1[roots_1 > 0])

m_0 = lambda γ: boundary_drift(ω_1, ω_2, γ, ρ, σ_y)
m_1 = lambda γ: -boundary_drift(ω_2, ω_1, γ, ρ, σ_y)

print("roots of m_ϑ(0+):", roots_0, " check:", [m_0(g) for g in roots_0])
print("roots of m_ϑ(1-):", roots_1, " check:", [m_1(g) for g in roots_1])

outcome = {(True, True): "(a) both survive",
           (True, False): "(b) agent 1 dominates",
           (False, True): "(c) agent 2 dominates",
           (False, False): "(d) either agent dominates"}

cuts = np.concatenate(([0.0], np.sort(np.concatenate((roots_0, roots_1))),
                       [np.inf]))
for lo, hi in zip(cuts[:-1], cuts[1:]):
    γ_mid = 0.5 * (lo + hi) if np.isfinite(hi) else 2 * lo
    print(f"γ in ({lo:6.3f}, {hi:6.3f}):  "
          f"{outcome[(m_0(γ_mid) > 0, m_1(γ_mid) < 0)]}")
```

```
roots of m_ϑ(0+): [1.46529206]  check: [np.float64(5.0837530595323925e-18)]
roots of m_ϑ(1-): [ 0.39074321 10.71675679]  check: [np.float64(-2.
↪4401931409719703e-17), np.float64(-6.971478593011729e-18)]
γ in ( 0.000,  0.391):  (d) either agent dominates
γ in ( 0.391,  1.465):  (c) agent 2 dominates
γ in ( 1.465, 10.717):  (a) both survive
γ in (10.717,    inf):  (b) agent 1 dominates
```

The outcomes run from (d) to © to (a) to (b) as γ rises, consistent with asymptotic results 1 and 2.

35.10 The separable case

When $\gamma = \rho$, the model collapses to the separable CRRA benchmark.

Blume and Easley [2006] study this benchmark in discrete time; *Heterogeneous Beliefs and Financial Markets* presents their analysis.

The two agents' discount rates ν^1 and ν^2 then coincide, and the log-odds process (35.6) becomes

$$d\vartheta_t = \frac{1}{2}\left[(\omega^2)^2 - (\omega^1)^2\right] dt + (\omega^1 - \omega^2) dW_t.$$

Its drift is the common value of the two boundary drifts in *Proposition 35.5.1* when $\gamma = \rho$ (see Exercise 35.5.4).

The drift is constant and depends only on the relative entropy of the two belief distortions.

Under P, the relative entropy of P with respect to Q^n on $\mathcal{F}_t$ is $E^P[-\log M_t^n] = \frac{1}{2}(\omega^n)^2 t$, the Kullback–Leibler divergence described in *Kullback–Leibler divergence*.

The drift of ϑ is the difference between the growth rates of the two agents' relative entropies.

The agent with the smaller $|\omega^n|$ dominates under P.

If the two agents have equal magnitudes of belief distortions, neither becomes extinct almost surely, but no nondegenerate stationary wealth distribution exists.

```
def simulate_crra_pareto(ω_1, ω_2, T, dt, n_paths, seed=42):
    """
    Simulate Pareto-share dynamics in the separable benchmark.
    """
    rng = np.random.default_rng(seed)
    n_steps = int(T / dt)
    t_grid = np.linspace(0, T, n_steps + 1)

    drift = 0.5 * (ω_2**2 - ω_1**2)
    volatility = ω_1 - ω_2

    θ = np.zeros((n_paths, n_steps + 1))
    dW = rng.normal(0.0, np.sqrt(dt), size=(n_paths, n_steps))

    for t in range(n_steps):
        θ[:, t + 1] = θ[:, t] + drift * dt + volatility * dW[:, t]

    υ_paths = 1.0 / (1.0 + np.exp(-θ))
    return t_grid, υ_paths


ω_1 = 0.10
ω_2 = 0.0
T, dt = 2000, 0.05
t_grid, υ_paths = simulate_crra_pareto(ω_1, ω_2, T=T, dt=dt, n_paths=50)

fig, ax = plt.subplots(figsize=(11, 5))

for i in range(20):
    ax.plot(t_grid, υ_paths[i], color="C0", alpha=0.25, lw=1)

ax.axhline(0.5, color="gray", linestyle=":", lw=1)
ax.set_xlabel("time")
ax.set_ylabel(r"Pareto share $\upsilon_t$")
```

(continues on next page)

(continued from previous page)

```
ax.set_ylim(0, 1)
plt.tight_layout()
plt.show()

drift = 0.5 * (ω_2**2 - ω_1**2)
vol = abs(ω_1 - ω_2)
print(f"{'t':>6}{'share of paths with υ_t < 1/2':>32}{'P(υ_t < 1/2)':>16}")
for t in [200, 500, 1000, 2000]:
    k = int(round(t / dt))
    share = np.mean(υ_paths[:, k] < 0.5)
    prob = 0.5 * (1 + erf(-drift * sqrt(t) / vol / sqrt(2)))
    print(f"{t:>6}{share:>32.2f}{prob:>16.3f}")
```

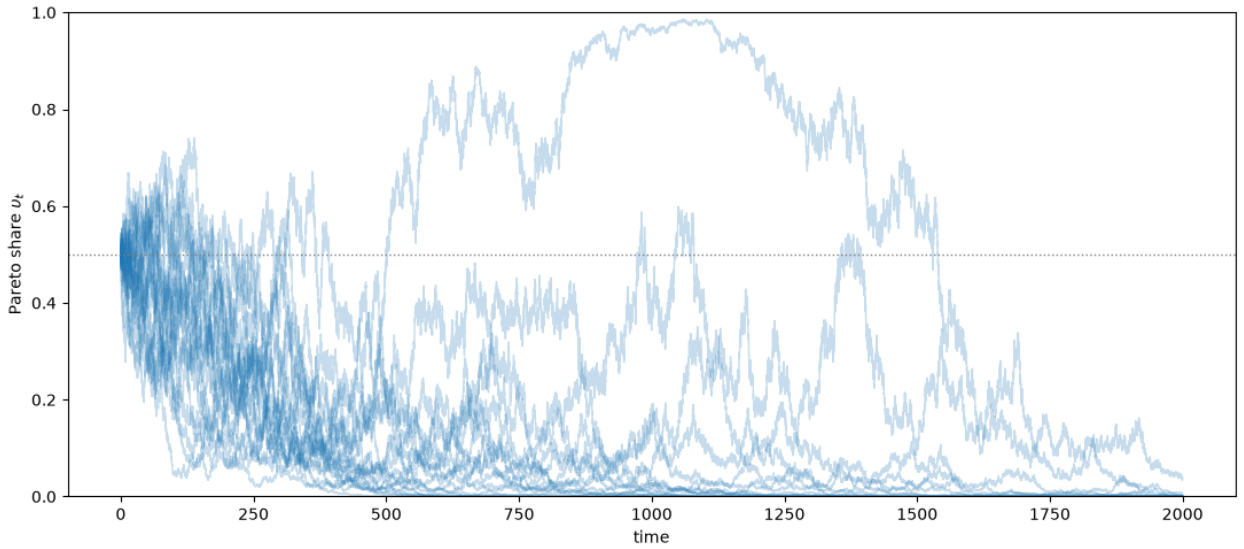


Fig. 35.4: Pareto-share paths in the separable benchmark

```
     t   share of paths with υ_t < 1/2    P(υ_t < 1/2)
   200                             0.78           0.760
   500                             0.84           0.868
  1000                             0.94           0.943
  2000                             1.00           0.987
```

This figure simulates 50 sample paths of the Pareto share υ_t under separable CRRA preferences ($\gamma = \rho$) with $\omega^1 = 0.10$ and $\omega^2 = 0$.

Agent 2 has correct beliefs, so ϑ has drift $-\frac{1}{2}(\omega^1)^2 = -0.005$ and $\upsilon_t \to 0$ almost surely.

Agent 1 is driven to extinction – the classical market-selection result of Blume and Easley [2006].

Selection is slow.

The volatility of ϑ is 0.1, so at $t = 200$ the mean of ϑ_t, -1, is smaller in magnitude than its standard deviation, 1.41.

The table compares the share of the 50 simulated paths with $\upsilon_t < 1/2$ with the probability $P(\vartheta_t < 0) = \Phi(0.05\sqrt{t})$, where Φ is the standard normal distribution function.

About three quarters of the paths have $\upsilon_t < 1/2$ at $t = 200$; nearly all do by $t = 2000$.

35.11 Asset pricing implications

As one agent becomes negligible, current prices converge to those of the homogeneous economy populated by the large agent.

When agent 2 is the large agent, Proposition 5.1 in Borovička [2020] implies

$$\lim_{\upsilon \searrow 0} r(\upsilon) = \beta + \rho \left(\mu_Y + \omega^2 \sigma_Y + \frac{1}{2}(1-\gamma)\sigma_Y^2 \right) - \frac{1}{2}\gamma\sigma_Y^2 \tag{35.13}$$

and

$$\lim_{\upsilon \searrow 0} y(\upsilon) = \left[\beta - (1-\rho) \left(\mu_Y + \omega^2 \sigma_Y + \frac{1}{2}(1-\gamma)\sigma_Y^2 \right) \right]^{-1}. \tag{35.14}$$

The aggregate wealth dynamics also converge to those of the homogeneous economy:

$$\lim_{\upsilon \searrow 0} m_A(\upsilon) = \mu_Y, \qquad \lim_{\upsilon \searrow 0} \sigma_A(\upsilon) = \sigma_Y.$$

Proposition 5.3 then gives the negligible agent's own consumption-saving and portfolio choices.

Her consumption-wealth ratio converges to

$$\lim_{\upsilon \searrow 0} (y^1(\upsilon))^{-1} = (y(0))^{-1} - \frac{1-\rho}{\rho} \left[(\omega^1 - \omega^2)\sigma_Y + \frac{(\omega^1 - \omega^2)^2}{2\gamma} \right].$$

The small agent's risky-asset share converges to

$$\lim_{\upsilon \searrow 0} \pi^1(\upsilon) = 1 + \frac{\omega^1 - \omega^2}{\gamma \sigma_Y}. \tag{35.15}$$

Hence optimism implies leverage, while sufficiently strong pessimism implies shorting.

```
ω_2 = 0.0
σ_y = 0.02
ω_grid = np.linspace(-0.5, 1.0, 300)

fig, ax = plt.subplots(figsize=(10, 5))

for γ in [2, 5, 10, 20]:
    π_1 = 1 + (ω_grid - ω_2) / (γ * σ_y)
    ax.plot(ω_grid, π_1, lw=2, label=rf"$\gamma = {γ}$")

ax.axhline(1.0, color="gray", linestyle=":", lw=1)
ax.axhline(0.0, color="gray", linestyle=":", lw=1)
ax.axvline(0.0, color="gray", linestyle=":", lw=1)
ax.set_xlabel(r"belief distortion $\omega^1$")
ax.set_ylabel(r"risky share $\pi^1$")
ax.legend()
plt.tight_layout()
plt.show()
```

This figure plots the limiting risky-asset share π^1 of the negligible agent as a function of her belief distortion ω^1 ($\omega^2 = 0$, $\sigma_Y = 0.02$), for four levels of risk aversion.

At $\omega^1 = 0$ the agent agrees with agent 2 and holds the market portfolio ($\pi^1 = 1$).

Optimism ($\omega^1 > 0$) leads to leverage ($\pi^1 > 1$), while sufficient pessimism ($\omega^1 < 0$) leads to shorting ($\pi^1 < 0$).

Higher risk aversion compresses these deviations toward one.

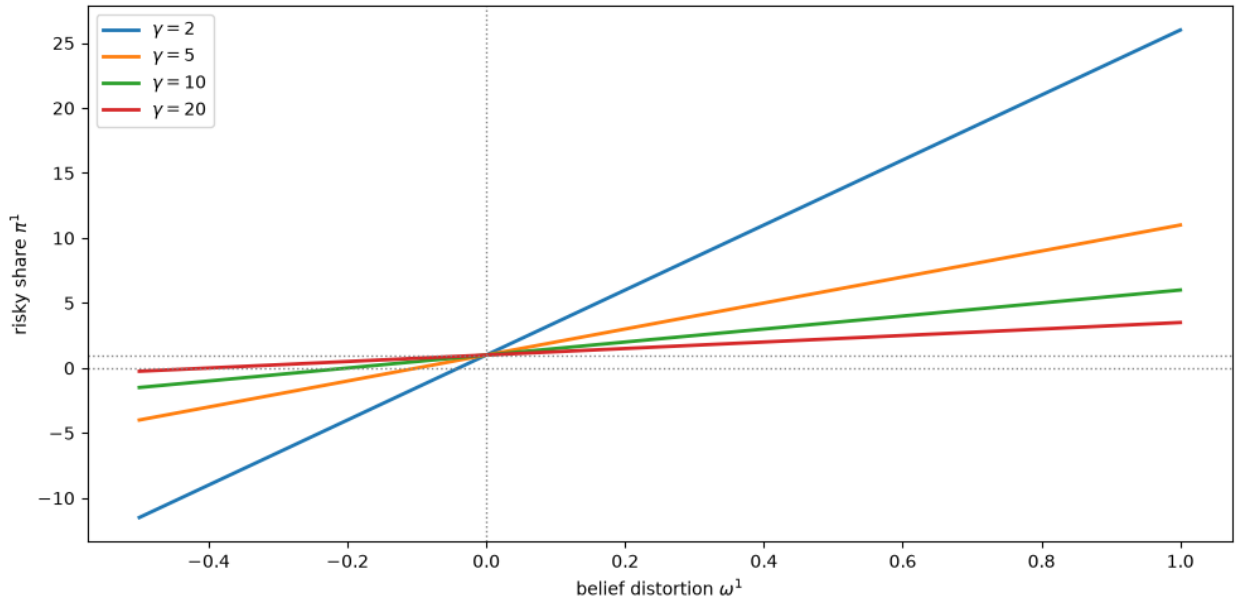


Fig. 35.5: Limiting risky-asset shares of the small agent

35.12 Optimistic and pessimistic distortions

Optimistic and pessimistic beliefs affect survival asymmetrically.

An optimistic agent benefits from the risk premium term and, when IES > 1, from the saving term as well.

A pessimistic agent holds less of the risky asset than agent 2, who holds the market portfolio, $\pi^2 = 1$.

With $\omega^2 = 0$, the portfolio that maximizes the expected log return under P has risky share γ.

When $\gamma \geq 1$, a smaller risky share moves the pessimist away from that portfolio, so she earns a lower expected log return than agent 2.

She can then survive only if the saving effect is strong enough to offset that loss.

With IES > 1 the saving term helps her only when she holds a large short position, $\pi^1 < -1$; with IES < 1 it helps her only when $-1 < \pi^1 < 1$.

When $\gamma < 1$, the market portfolio is more leveraged than the portfolio that maximizes the expected log return.

A mildly pessimistic agent with $2\gamma - 1 < \pi^1 < 1$ then earns a higher expected log return than agent 2 without help from the saving channel.

For example, with $\rho = 1$, $\gamma = 0.5$, $\sigma_Y = 0.02$, and $\omega^1 = -0.005$, both boundary drifts are positive and the pessimist dominates.

```
print(f"m_0(0+) = {boundary_drift(-0.005, 0.0, 0.5, 1.0, 0.02):.2e}, "
      f"m_0(1-) = {-boundary_drift(0.0, -0.005, 0.5, 1.0, 0.02):.2e}")
```

```
m_0(0+) = 2.50e-05, m_0(1-) = 5.00e-05
```

```
fig, axes = plt.subplots(1, 2, figsize=(14, 5), sharey=True)

σ_y = 0.02
```

(continues on next page)

(continued from previous page)

```
ω_2 = 0.0
ρ = 0.67
y_grid = np.linspace(0.5, 25.0, 300)

ax = axes[0]
for ω_1 in [0.1, 0.25, 0.5, 1.0]:
    _, _, _, total = decompose_survival(ω_1, ω_2, y_grid, ρ, σ_y)
    ax.plot(y_grid, total, lw=2, label=rf"$\omega^1 = {ω_1}$")
ax.axhline(0, color="gray", lw=1)
ax.set_title("optimistic", fontsize=12)
ax.set_xlabel(r"risk aversion $\gamma$")
ax.set_ylabel(r"$m_\vartheta(0+)/\gamma$")
ax.set_ylim(-0.03, 0.06)
ax.legend(fontsize=9)

ax = axes[1]
for ω_1 in [-0.1, -0.25, -0.5, -1.0]:
    _, _, _, total = decompose_survival(ω_1, ω_2, y_grid, ρ, σ_y)
    ax.plot(y_grid, total, lw=2, label=rf"$\omega^1 = {ω_1}$")
ax.axhline(0, color="gray", lw=1)
ax.set_title("pessimistic", fontsize=12)
ax.set_xlabel(r"risk aversion $\gamma$")
ax.legend(fontsize=9)

plt.tight_layout()
plt.show()

f = lambda y: decompose_survival(-1.0, ω_2, y, ρ, σ_y)[3]
total = f(y_grid)
idx = np.where(np.diff(np.sign(total)) != 0)[0]
roots = [brentq(f, y_grid[i], y_grid[i + 1]) for i in idx]
print("ω¹ = -1.0: the curve changes sign at γ =", np.round(roots, 2))
```

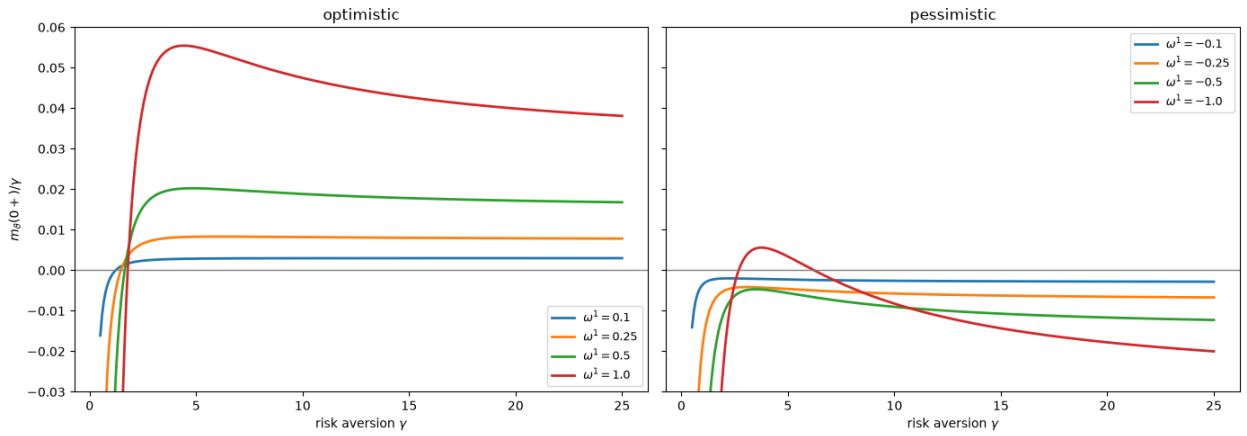


Fig. 35.6: Wealth-growth differential at the boundary for optimistic and pessimistic distortions

```
ω¹ = -1.0: the curve changes sign at γ = [2.69 6.23]
```

Both panels plot the wealth-growth differential $\lim_{\upsilon \searrow 0}[m_A^1 - m_A^2] = \lim_{\upsilon \searrow 0} m_\vartheta/\gamma$ as a function of γ (IES ≈ 1.49, $\omega^2 = 0$).

We clip the vertical axis; every curve falls steeply as $\gamma \to 0$.

Where the curve is positive, agent 1 survives near extinction.

- Left panel (optimistic agent): larger ω^1 means a bigger bet on the risky asset, so the volatility penalty dominates at low γ but the drift turns positive once γ is large enough.
- Right panel (pessimistic agent): the curves for $\omega^1 = -0.1, -0.25, -0.5$ are negative throughout. Only the most pessimistic agent, $\omega^1 = -1$, survives near extinction, for γ between about 2.69 and 6.23. There she shorts the risky asset heavily ($\pi^1 < -1$), expects a high return on her position, and saves enough to offset her portfolio losses.

35.13 Long-run consumption distribution

When both agents survive, the Pareto share keeps moving across the whole interval $(0, 1)$.

The next simulation is only a toy approximation.

It interpolates the drift between its two boundary values, so it illustrates the recurrence logic without solving the full equilibrium ODE.

```
def simulate_pareto_share_toy(ω_1, ω_2, γ, ρ, σ_y, T, dt, n_paths=20, seed=42):
    """
    Simulate a toy Pareto-share process by interpolating boundary drifts.
    """
    rng = np.random.default_rng(seed)
    n_steps = int(T / dt)
    t_grid = np.linspace(0, T, n_steps + 1)

    volatility = ω_1 - ω_2
    m_0 = boundary_drift(ω_1, ω_2, γ, ρ, σ_y)
    m_1 = -boundary_drift(ω_2, ω_1, γ, ρ, σ_y)

    θ = np.zeros((n_paths, n_steps + 1))
    dW = rng.normal(0.0, np.sqrt(dt), size=(n_paths, n_steps))

    for t in range(n_steps):
        υ = 1.0 / (1.0 + np.exp(-θ[:, t]))
        drift = m_0 * (1 - υ) + m_1 * υ
        θ[:, t + 1] = θ[:, t] + drift * dt + volatility * dW[:, t]

    υ_paths = 1.0 / (1.0 + np.exp(-θ))
    return t_grid, υ_paths


ω_1 = 0.25
ω_2 = 0.0
γ = 5.0
ρ = 0.67
σ_y = 0.02

t_grid, υ_paths = simulate_pareto_share_toy(
    ω_1, ω_2, γ, ρ, σ_y, T=500, dt=0.05, n_paths=50, seed=42
)

fig, axes = plt.subplots(1, 2, figsize=(14, 5))

ax = axes[0]
for i in range(20):
    ax.plot(t_grid, υ_paths[i], color="C0", alpha=0.25, lw=1)
```

(continues on next page)

(continued from previous page)

```
ax.axhline(0.5, color="gray", linestyle=":", lw=1)
ax.set_title("sample paths", fontsize=12)
ax.set_xlabel("time")
ax.set_ylabel(r"Pareto share $\upsilon_t$")
ax.set_ylim(0, 1)

ax = axes[1]
_, υ_long = simulate_pareto_share_toy(
    ω_1, ω_2, γ, ρ, σ_y, T=2000, dt=0.05, n_paths=5, seed=123
)
υ_stationary = υ_long[:, υ_long.shape[1] // 2:].ravel()
ax.hist(υ_stationary, bins=80, density=True, color="steelblue",
        edgecolor="white", alpha=0.7)
ax.set_title("approximate stationary density", fontsize=12)
ax.set_xlabel(r"Pareto share $\upsilon$")
ax.set_ylabel("density")
ax.set_xlim(0, 1)

plt.tight_layout()
plt.show()
```

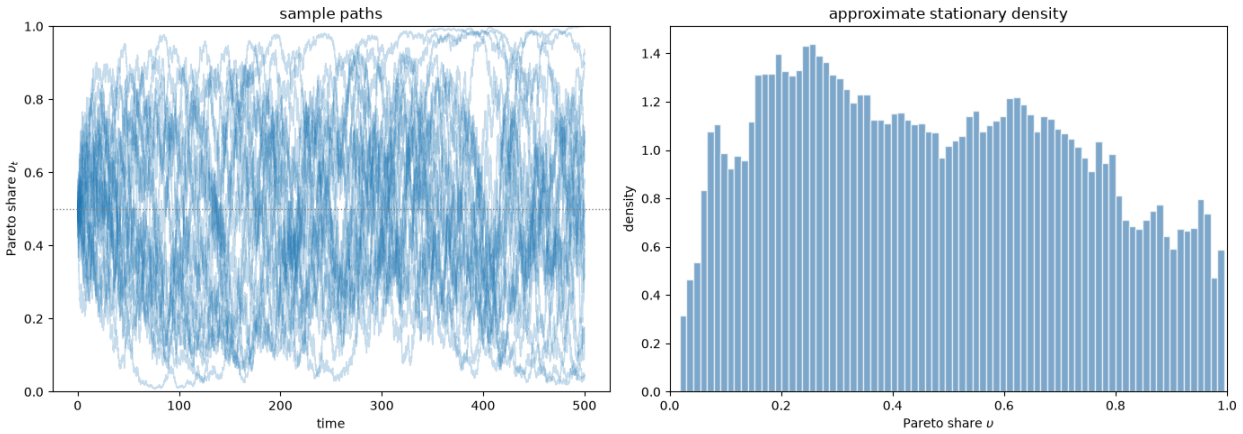


Fig. 35.7: A toy stationary Pareto-share simulation

The left panel shows 50 sample paths of the Pareto share υ_t under parameters inside the coexistence region ($\omega^1 = 0.25$, $\omega^2 = 0$, $\gamma = 5$, IES ≈ 1.49).

Unlike the separable case in Fig. 35.4, the paths do not drift to zero – they repeatedly visit a wide range of values, bouncing between the two repelling boundaries.

The right panel approximates the stationary density by pooling the second half of longer simulations.

The interior mode is consistent with neither agent being driven to extinction.

However, this toy interpolation only illustrates the recurrence logic; it does not reproduce the quantitative stationary consumption-share density in Figure 4 of Borovička [2020], which requires solving the full interior equilibrium ODE.

35.14 Summary

Recursive preferences weaken the classical market-selection result.

The portfolio return channel still rewards more optimistic beliefs.

The volatility channel still penalizes aggressive positions.

The saving channel rewards an agent who expects a high return on her portfolio when IES > 1 and penalizes her when IES < 1.

The two boundary drifts differ by $(\omega^1 - \omega^2)^2(1/\rho - 1/\gamma)$.

Both agents can survive only when $\gamma > \rho$, which means that agents prefer early resolution of uncertainty.

Each agent can dominate with positive probability only when $\gamma < \rho$.

Coexistence does not require IES > 1: with $\omega^1 = 0.10$, $\omega^2 = 0$, and $\sigma_Y = 0.02$, both agents survive at $\gamma = 5$, $\rho = 1.5$.

Recursive-preference economies with $\gamma > \rho$ can support stationary long-run wealth distributions with persistent heterogeneity in beliefs and portfolio positions.

CHAPTER

THIRTYSIX

THE DESIGN AND PRICE OF INFORMATION

Contents

- *The Design and Price of Information*
  - *Overview*
  - *The decision problem*
  - *Experiments and their value*
  - *Blackwell's order is only partial*
  - *The seller's problem*
  - *Two types*
  - *A continuum of types*
  - *Concluding remarks*
  - *Exercises*

36.1 Overview

Earlier lectures in this section asked *which* statistical experiment a decision maker should prefer.

Blackwell's Theorem on Comparing Experiments gave the classic answer: experiment μ is at least as informative as experiment ν when *every* Bayesian decision maker attains weakly higher expected utility with μ.

This lecture asks a different question.

Suppose somebody *owns* the data and wants to sell it.

What should she sell, and at what price?

We study Bergemann *et al.* [2018], who analyze a monopolist data seller facing a buyer who already has some private information of his own.

The buyer's private information is exactly what he would like to hide, because it determines his willingness to pay.

So the seller screens by offering a *menu* of statistical experiments, degrading the information sold to some buyers in order to charge more to others.

The central finding is that degrading information is not simply a matter of adding noise.

Blackwell's order is a *partial* order, so two experiments can be ranked differently by different decision makers.

The seller exploits precisely those incomparable pairs: information has a **vertical** dimension, its quality, and a **horizontal** dimension, its position.

That horizontal dimension has no counterpart in ordinary monopoly screening over quality or quantity, and it is what allows the seller to extract rents that would otherwise be impossible to reach.

Along the way we will

- compute the value of an arbitrary experiment to an arbitrary belief type,
- verify numerically that Blackwell's order fails to rank the experiments the seller wants to use,
- solve the two-type screening problem by brute force and check it against the paper's closed forms,
- solve the continuum-of-types problem as a **linear program**, which reproduces the paper's ironing and pooling results without any need to implement ironing by hand.

Let's start with imports.

```
import matplotlib.pyplot as plt
import numpy as np
from scipy import stats
from scipy.optimize import linprog

plt.rcParams['figure.figsize'] = (10, 5)
np.set_printoptions(precision=4, suppress=True)
```

36.2 The decision problem

A data buyer must choose an action a from a finite set A without knowing the state ω, which lives in a finite set Ω.

We work throughout with the **matching** case in which the buyer wants his action to match the state,

$$u(\omega_i, a_j) = \mathbb{1}[i = j] \cdot u_i, \tag{36.1}$$

so that matching state ω_i pays $u_i > 0$ and any mismatch pays zero.

From here on we take two states and two actions, $\Omega = \{\omega_1, \omega_2\}$ and $A = \{a_1, a_2\}$, which is the case Bergemann *et al.* [2018] solve completely.

The buyer's **type** is his interim belief

$$\theta = \Pr[\omega = \omega_1] \in [0, 1],$$

which is private information.

The seller knows only the distribution F from which θ is drawn.

Without extra information the buyer picks the better of two constant actions, so his reservation utility is

$$u(\theta) = \max\{\theta u_1,\ (1-\theta) u_2\}. \tag{36.2}$$

The type that is *least* sure what to do is the one where the two terms are equal,

$$\theta^* = \frac{u_2}{u_1 + u_2}. \tag{36.3}$$

Types above θ^* would choose a_1 on their own, types below would choose a_2.

Note

The buyer's belief θ can be generated from a common prior together with a privately observed signal, exactly as in *Likelihood Ratio Processes and Bayesian Learning* and *Exchangeability and Bayesian Updating*.

A buyer with a very precise private signal has θ near 0 or 1; a buyer who has learned nothing sits at the common prior, which equals θ^* only by coincidence.

A "high type" in this lecture is a buyer whose belief lies near θ^*, and he is willing to pay the most.

When the common prior equals θ^*, high types are the badly informed buyers.

36.3 Experiments and their value

A statistical experiment is a stochastic matrix mapping states into signals.

With two states and two actions it suffices to consider two signals, and we write

$$E = \begin{pmatrix} \pi_1 & 1-\pi_1 \\ 1-\pi_2 & \pi_2 \end{pmatrix}, \tag{36.4}$$

where row i gives the signal distribution in state ω_i.

Thus $\pi_1 = \Pr[s_1 \mid \omega_1]$ and $\pi_2 = \Pr[s_2 \mid \omega_2]$.

We adopt the normalization $\pi_1 + \pi_2 \geq 1$, which just says that signal s_1 is relatively more likely in state ω_1 than in state ω_2.

The **fully informative** experiment $\overline{E}$ has $\pi_1 = \pi_2 = 1$.

After seeing signal s_k the buyer picks the action with the highest expected payoff, so his gross value is obtained by summing the best he can do signal by signal.

Subtracting his reservation utility (36.2) gives the **net value of information**

$$V(E,\theta) = \max\{\theta\pi_1 u_1,\ (1-\theta)(1-\pi_2)u_2\} + \max\{\theta(1-\pi_1)u_1,\ (1-\theta)\pi_2 u_2\} - \max\{\theta u_1,\ (1-\theta)u_2\}. \tag{36.5}$$

```
def value(pi1, pi2, theta, u1=1.0, u2=1.0):
    """Net value of experiment (pi1, pi2) to a buyer with belief theta."""
    theta = np.asarray(theta, dtype=float)
    s1 = np.maximum(theta * pi1 * u1, (1 - theta) * (1 - pi2) * u2)
    s2 = np.maximum(theta * (1 - pi1) * u1, (1 - theta) * pi2 * u2)
    return s1 + s2 - np.maximum(theta * u1, (1 - theta) * u2)
```

If the buyer simply obeys the recommendation implicit in each signal, taking a_1 after s_1 and a_2 after s_2, the value collapses to

$$V(E,\theta) = \max\{\theta\pi_1 u_1 + (1-\theta)\pi_2 u_2 - \max\{\theta u_1, (1-\theta)u_2\},\ 0\}, \tag{36.6}$$

which is the expression the paper works with.

```
def value_obedient(pi1, pi2, theta, u1=1.0, u2=1.0):
    """Value when the buyer follows the recommendation, or ignores the signal."""
    theta = np.asarray(theta, dtype=float)
    return np.maximum(theta * pi1 * u1 + (1 - theta) * pi2 * u2
                      - np.maximum(theta * u1, (1 - theta) * u2), 0.0)
```

The two expressions agree exactly under the normalization $\pi_1 + \pi_2 \geq 1$ and can differ sharply without it, which is what the normalization is for.

```
grid = np.linspace(0, 1, 2001)
worst_ok = worst_bad = 0.0
for p1 in np.linspace(0, 1, 51):
    for p2 in np.linspace(0, 1, 51):
        gap = np.abs(value(p1, p2, grid) - value_obedient(p1, p2, grid)).max()
        if p1 + p2 >= 1:
            worst_ok = max(worst_ok, gap)
        else:
            worst_bad = max(worst_bad, gap)

print(f'largest gap where pi1 + pi2 >= 1:  {worst_ok:.2e}')
print(f'largest gap where pi1 + pi2 <  1:  {worst_bad:.4f}')
```

```
largest gap where pi1 + pi2 >= 1:  1.11e-16
largest gap where pi1 + pi2 <  1:  0.5000
```

Under the normalization, reversing the recommendation, taking a_2 after s_1 and a_1 after s_2, is dominated by ignoring the signal, so every type either obeys or ignores.

That is why the two formulas coincide for every θ, including types who misreport and buy an experiment designed for somebody else.

We use the general form (36.5) from here on because it does not rely on the normalization.

Here is the value of information as a function of the buyer's type.

```
theta = np.linspace(0, 1, 1001)

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
for ax, (p1, p2), ttl in zip(
        axes, [(1.0, 1.0), (0.5, 1.0)],
        [r'fully informative $(\pi_1,\pi_2)=(1,1)$',
         r'partially informative $(\pi_1,\pi_2)=(1/2,1)$']):
    ax.plot(theta, value(p1, p2, theta), lw=2)
    ax.axvline(0.5, color='0.6', ls='--', lw=1)
    ax.set(xlabel=r'$\theta$', ylabel=r'$V(E,\theta)$', title=ttl)
fig.suptitle('Value of information, $u_1 = u_2 = 1$')
fig.tight_layout()
plt.show()
```

Three features of these pictures drive everything that follows.

The value is **piecewise linear** in θ, because types are probabilities and expected utilities are linear in probabilities.

The value is **highest at** θ^* and falls to zero at $\theta \in \{0, 1\}$: the buyer who already knows the state will pay nothing, and the buyer who knows least will pay most.

The partially informative experiment in the right panel is worth **nothing at all** to types above $2/3$, even though it is worth a great deal to types just below $1/2$.

That last property is the seller's main tool.

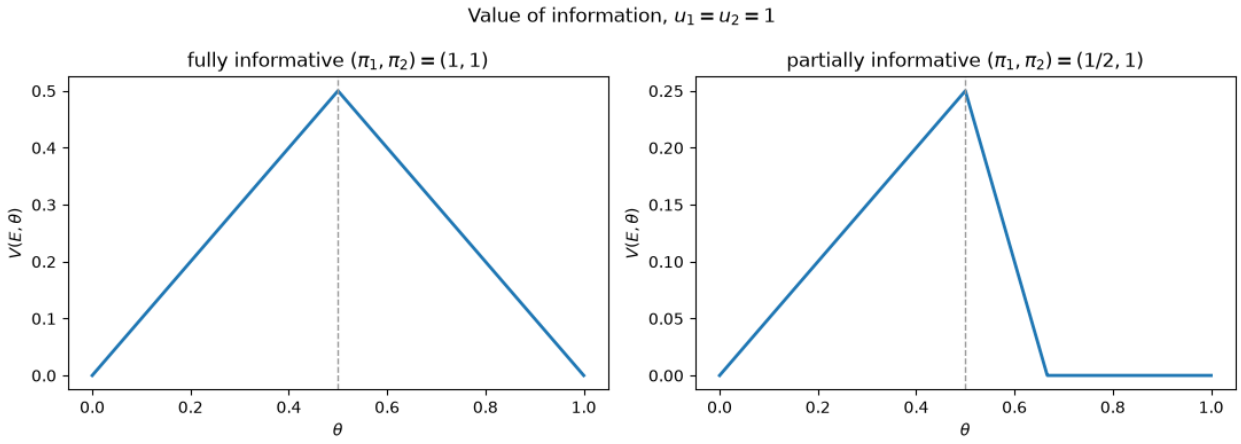


Fig. 36.1: Value of full and partial information

36.4 Blackwell's order is only partial

Blackwell's Theorem on Comparing Experiments establishes that experiment E is at least as informative as E' in Blackwell's sense exactly when E' is a **garbling** of E.

In the notation of *Definition 30.2.1*, with $\mu = E$ and $\nu = E'$, there is a Markov matrix Q with

$$E' = EQ. \tag{36.7}$$

When E is invertible this is easy to check: solve $Q = E^{-1}E'$ and ask whether every entry of Q is nonnegative.

The rows of Q sum to one automatically: $EQ\mathbf{1} = E'\mathbf{1} = \mathbf{1} = E\mathbf{1}$, and E is invertible, so $Q\mathbf{1} = \mathbf{1}$.

```
def experiment(pi1, pi2):
    return np.array([[pi1, 1 - pi1], [1 - pi2, pi2]])


def garbling(E, Ep, tol=1e-9):
    """Return Q with Ep = E @ Q if Ep is a garbling of E, else None."""
    if abs(np.linalg.det(E)) < tol:
        return None
    Q = np.linalg.solve(E, Ep)           # rows of Q sum to one automatically
    if (Q > -tol).all():
        return Q
    return None


pairs = [((1, 1), (0.8, 1)), ((1, 1), (1, 0.8)),
         ((0.9, 0.9), (0.8, 0.8)),
         ((0.8, 1), (1, 0.8)), ((1, 0.8), (0.8, 1))]
for a, b in pairs:
    ok = garbling(experiment(*a), experiment(*b)) is not None
    print(f'  is {b} a garbling of {a}?  {"yes" if ok else "no"}')
```

```
  is (0.8, 1) a garbling of (1, 1)?  yes
  is (1, 0.8) a garbling of (1, 1)?  yes
  is (0.8, 0.8) a garbling of (0.9, 0.9)?  yes
  is (1, 0.8) a garbling of (0.8, 1)?  no
  is (0.8, 1) a garbling of (1, 0.8)?  no
```

The fully informative experiment garbles into everything, and $(0.9, 0.9)$ garbles into the uniformly noisier $(0.8, 0.8)$.

Those are the **vertical** comparisons, and Blackwell's theorem says every type agrees about them.

But $(0.8, 1)$ and $(1, 0.8)$ garble into each other in neither direction.

Blackwell's order simply does not rank them, and that means different types are free to rank them differently.

```
va, vb = value(0.8, 1, theta), value(1, 0.8, theta)

fig, ax = plt.subplots()
ax.plot(theta, va, lw=2, label=r'$E_a = (0.8, 1)$')
ax.plot(theta, vb, lw=2, label=r'$E_b = (1, 0.8)$')
ax.fill_between(theta, va, vb, where=va > vb, alpha=0.15, color='C0')
ax.fill_between(theta, va, vb, where=vb > va, alpha=0.15, color='C1')
ax.axvline(0.5, color='0.4', ls='--', lw=1)
ax.set(xlabel=r'$\theta$', ylabel=r'$V(E,\theta)$',
       title='Types below $1/2$ prefer $E_a$, types above prefer $E_b$')
ax.legend()
fig.tight_layout()
plt.show()

for t in [0.2, 0.35, 0.65, 0.8]:
    pref = 'E_a' if value(0.8, 1, t) > value(1, 0.8, t) else 'E_b'
    print(f'  theta = {t}:  V(E_a) = {value(0.8, 1, t):.4f},'
          f'  V(E_b) = {value(1, 0.8, t):.4f}   prefers {pref}')
```

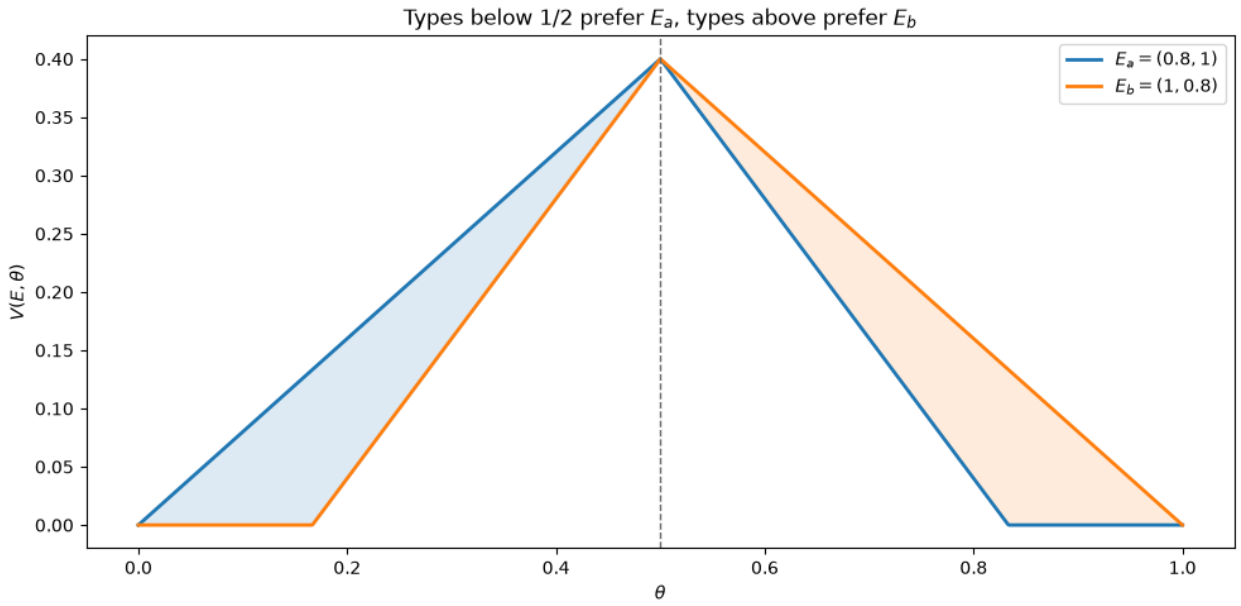


Fig. 36.2: Two experiments that Blackwell's order does not rank

```
theta = 0.2:  V(E_a) = 0.1600,  V(E_b) = 0.0400   prefers E_a
theta = 0.35:  V(E_a) = 0.2800,  V(E_b) = 0.2200   prefers E_a
theta = 0.65:  V(E_a) = 0.2200,  V(E_b) = 0.2800   prefers E_b
theta = 0.8:  V(E_a) = 0.0400,  V(E_b) = 0.1600   prefers E_b
```

Signal s_1 of E_a occurs only in state ω_1, so it rules out ω_2; signal s_2 of E_b occurs only in state ω_2, so it rules out ω_1.

A buyer who leans toward ω_1 takes a_1 on his own and switches to a_2 only on strong evidence for ω_2.

E_b supplies conclusive evidence for ω_2, so he values E_b.

A buyer who leans toward ω_2 values E_a for the same reason.

This is the **horizontal** dimension of information.

In ordinary nonlinear pricing over quality or quantity, all types agree on the ranking of products and the seller can only move up and down a single ladder.

Here the seller has a second dial, and turning it lets her hand one type something that is worthless to another.

36.5 The seller's problem

The seller commits to a menu $\{E(\theta), t(\theta)\}$ assigning an experiment and a price to each reported type.

Payments cannot be made contingent on the state, the signal, or the buyer's action, so the value of an experiment to a buyer is determined by his belief alone.

Writing $V(\theta) = V(E(\theta), \theta) - t(\theta)$ for the buyer's rent, the seller solves

$$\max_{\{E(\theta),\, t(\theta)\}} \int t(\theta)\, dF(\theta) \tag{36.8}$$

subject to incentive compatibility and individual rationality,

$$V(\theta) \geq V(E(\theta'), \theta) - t(\theta') \;\; \forall \theta, \theta', \qquad V(\theta) \geq 0 \;\; \forall \theta. \tag{36.9}$$

Bergemann *et al.* [2018] establish structural results that every menu we compute confirms.

Proposition 36.5.1

1. The fully informative experiment $\overline{E}$ is part of an optimal menu.
2. Every experiment in any optimal menu is **nondispersed**, meaning $\pi_{ij} = 0$ for some $i \neq j$.
3. In the matching case every experiment in any optimal menu is **concentrated**, meaning $\pi_{ii} = 1$ for some i.

Part 3 says that in our binary setting every experiment on the menu has $\pi_1 = 1$ or $\pi_2 = 1$.

With two states and two signals, a zero off-diagonal entry forces a one on the diagonal of the same row, so "nondispersed" and "concentrated" coincide.

An optimal experiment keeps one signal that perfectly reveals one state; the other signal is noisy.

36.6 Two types

Take two types θ^L and θ^H, with θ^H the *high value* type in the sense that he values the fully informative experiment more,

$$V(\overline{E}, \theta^H) \geq V(\overline{E}, \theta^L).$$

With $u_1 = u_2$ this says $|\theta^H - 1/2| \leq |\theta^L - 1/2|$, so the high type is the one who is *less* well informed to begin with.

Let $\gamma = \Pr[\theta = \theta^H]$.

The types are **congruent** if $\theta^* < \theta^H < \theta^L$, so both would take the same action without extra information, and **noncongruent** if $\theta^L < \theta^* < \theta^H$.

An optimal menu has the familiar shape: the high type buys $\overline{E}$, the low type's participation constraint binds, and the high type's incentive constraint binds.

Those three facts pin down both prices once the low type's experiment is chosen.

To search over *all* low-type experiments we also need prices for experiments that the high type values less than the low type does.

For those, the price that makes the high type's incentive constraint bind exceeds his value of $\overline{E}$, so his participation constraint binds instead.

The function below charges the high type the smaller of the two prices.

```
def two_type_revenue(pi1, pi2, tL, tH, gamma, u1=1.0, u2=1.0):
    """Revenue when the high type buys E_bar and the low type buys (pi1, pi2)."""
    VbarH, VbarL = value(1, 1, tH, u1, u2), value(1, 1, tL, u1, u2)
    VL_L, VL_H = value(pi1, pi2, tL, u1, u2), value(pi1, pi2, tH, u1, u2)
    t_low = VL_L                                    # low type's IR binds
    t_high = min(VbarH - VL_H + t_low,              # high type's IC
                 VbarH)                             # high type's IR
    if VbarL - t_high > 1e-12:                      # low type must not deviate
        return -np.inf
    return gamma * t_high + (1 - gamma) * t_low
```

36.6.1 Noncongruent types

Set $\theta^L = 1/5$ and $\theta^H = 7/10$ with $u_1 = u_2 = 1$, so that $\theta^* = 1/2$ lies between them.

Because the two types would take *different* actions on their own, the seller can build an experiment that is valuable to one and worthless to the other.

Choosing $\pi_2 = 1$ and

$$\pi_1' = \frac{u_1\theta^H - u_2(1-\theta^H)}{u_1\theta^H} \tag{36.10}$$

leaves the high type exactly indifferent between his two actions after signal s_2, so the experiment is worth nothing to him while the low type values it strictly.

That is feasible but not optimal.

The seller does better by making the high type's incentive constraint bind instead, which gives

$$\pi_1'' = \frac{u_1\theta^H - u_2(1-\theta^H)}{u_1(\theta^H - \theta^L)}. \tag{36.11}$$

```
tL, tH, u1, u2 = 0.2, 0.7, 1.0, 1.0
pi1_zero = (u1 * tH - u2 * (1 - tH)) / (u1 * tH)
pi1_opt = (u1 * tH - u2 * (1 - tH)) / (u1 * (tH - tL))

print(f'zero-value-to-high experiment   pi1 = {pi1_zero:.4f}  (= 4/7)')
print(f'binding-IC experiment            pi1 = {pi1_opt:.4f}  (= 4/5)')
print(f'\n  V(E_zero, theta_H) = {value(pi1_zero, 1, tH):.4f}')
print(f'  V(E_opt,  theta_L) = {value(pi1_opt, 1, tL):.4f}'
      f'    V(E_opt, theta_H) = {value(pi1_opt, 1, tH):.4f}')
```

```
zero-value-to-high experiment   pi1 = 0.5714   (= 4/7)
binding-IC experiment           pi1 = 0.8000   (= 4/5)

  V(E_zero, theta_H) = 0.0000
  V(E_opt,  theta_L) = 0.1600    V(E_opt, theta_H) = 0.1600
```

The second experiment gives *both* types the same net value of information, 0.16.

The seller charges the low type 0.16 for it and charges the high type his full value of $\overline{E}$.

The high type's incentive constraint binds, and so do both participation constraints.

The seller extracts the entire surplus from both types; the only cost of screening is the distortion of the low type's experiment.

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4.2))
for ax, p1, ttl in zip(axes, [pi1_zero, pi1_opt],
                       ['suboptimal menu: partial experiment worth zero to $\\theta^H$
↪',
                        'optimal menu: high type indifferent between the two items']):
    t_hi = value(1, 1, tH)                                    # price of E_bar
    t_lo = value(p1, 1, tL)                                   # price of partial item
    ax.plot(theta, value(1, 1, theta) - t_hi, lw=2, label='fully informative')
    ax.plot(theta, value(p1, 1, theta) - t_lo, lw=2, ls='--',
            label=f'partial, $\\pi_1={p1:.3f}$')
    ax.axhline(0, color='0.3', lw=1)
    for t, nm in [(tL, r'$\theta^L$'), (tH, r'$\theta^H$')]:
        ax.axvline(t, color='0.7', ls=':', lw=1)
        ax.annotate(nm, (t, ax.get_ylim()[0]), fontsize=9)
    ax.set(xlabel=r'$\theta$', ylabel=r'$V - t$', title=ttl, ylim=(-0.35, 0.25))
    ax.legend(fontsize=8, loc='upper left')
fig.tight_layout()
plt.show()
```

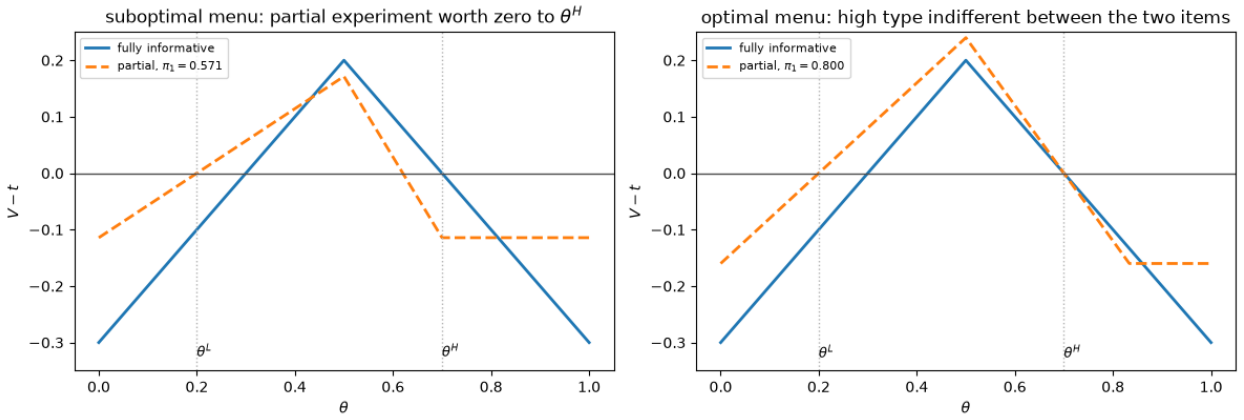


Fig. 36.3: Net value of the two menus as a function of the buyer's type

In the left panel the high type's net value of the fully informative experiment lies strictly above his net value of the partial one, so his incentive constraint is slack and the seller is leaving money on the table.

In the right panel the two curves meet exactly at θ^H.

Now we check the closed form (36.11) against a brute-force search over *all* experiments.

```
def brute_force(tL, tH, gamma, n=301, u1=1.0, u2=1.0):
    """Search over all (pi1, pi2) for the best low-type experiment."""
    g = np.linspace(0, 1, n)
    best, arg = -np.inf, None
    for p1 in g:
        for p2 in g:
            if p1 + p2 < 1:
                continue
            r = two_type_revenue(p1, p2, tL, tH, gamma, u1, u2)
            if r > best:
                best, arg = r, (p1, p2)
    return best, arg


print(f'{"gamma":>7s}{"brute force":>13s}{"argmax":>18s}'
      f'{"closed form":>14s}{"E_bar to both":>15s}')
for gamma in [0.10, 0.25, 0.30, 0.50, 0.90]:
    best, arg = brute_force(tL, tH, gamma)
    closed = two_type_revenue(pi1_opt, 1.0, tL, tH, gamma)
    both = two_type_revenue(1.0, 1.0, tL, tH, gamma)
    print(f'{gamma:7.2f}{best:13.5f}   ({arg[0]:.3f}, {arg[1]:.3f})'
          f'{closed:14.5f}{both:15.5f}')
print(f'\nthe paper: discriminate iff gamma > theta_L / theta_H = {tL / tH:.4f}')
```

```
  gamma  brute force            argmax   closed form  E_bar to both
   0.10      0.20000   (1.000, 1.000)       0.17400        0.20000
   0.25      0.20000   (1.000, 1.000)       0.19500        0.20000
   0.30      0.20200   (0.800, 1.000)       0.20200        0.20000
   0.50      0.23000   (0.800, 1.000)       0.23000        0.20000
   0.90      0.28600   (0.800, 1.000)       0.28600        0.20000

the paper: discriminate iff gamma > theta_L / theta_H = 0.2857
```

The brute-force optimum sits at $(\pi_1, \pi_2) = (0.8, 1)$ whenever discrimination pays, matching (36.11) exactly, and at $(1, 1)$ otherwise.

The switch happens right at $\gamma = \theta^L/\theta^H$.

When low types are common the seller prefers to sell everyone the fully informative experiment cheaply; when high types are common she prefers to protect the high price by degrading what the low type gets.

Note also that both experiments in the optimal menu have $\pi_2 = 1$, confirming part 3 of *Proposition 36.5.1*.

36.7 A continuum of types

Now let θ be distributed on $[0, 1]$ with density f and distribution F.

The key simplification is that the value of a *concentrated* experiment depends on (π_1, π_2) only through the scalar

$$q = \pi_1 u_1 - \pi_2 u_2 \in [-u_2,\ u_1], \tag{36.12}$$

which Bergemann *et al.* [2018] call the **differential informativeness** of the experiment.

By part 3 of *Proposition 36.5.1*, restricting the seller to concentrated experiments loses nothing.

For other experiments q does not determine the value: with $u_1 = u_2 = 1$, the experiments $(0.8, 0.8)$, $(0.9, 0.9)$ and $(1, 1)$ all have $q = 0$, but at $\theta = 1/2$ they are worth 0.3, 0.4 and 0.5.

In terms of q the value becomes

$$V(q,\theta) = \max\{\theta q + u_2 + \min\{u_1 - u_2 - q,\ 0\} - \max\{\theta u_1,\ (1-\theta)u_2\},\ 0\}. \tag{36.13}$$

The fully informative experiment is $q = u_1 - u_2$.

The two endpoints $q = -u_2$ and $q = u_1$ are the experiments in which one signal occurs with probability one in both states, so they convey nothing.

```
def value_q(q, theta, u1=1.0, u2=1.0):
    """Value of the experiment with differential informativeness q."""
    theta = np.asarray(theta, dtype=float)
    gross = theta * q + u2 + np.minimum(u1 - u2 - q, 0.0)
    return np.maximum(gross - np.maximum(theta * u1, (1 - theta) * u2), 0.0)


def q_to_experiment(q, u1=1.0, u2=1.0):
    """Recover (pi1, pi2) from q using pi1 = 1 or pi2 = 1."""
    return (1.0, (u1 - q) / u2) if q >= u1 - u2 else ((q + u2) / u1, 1.0)


for q in [-1.0, -0.5, 0.0, 0.5, 1.0]:
    p1, p2 = q_to_experiment(q)
    print(f'  q = {q:+.2f}  ->  (pi1, pi2) = ({p1:.3f}, {p2:.3f}),'
          f'   max value over types = {value_q(q, theta).max():.4f}')
```

```
  q = -1.00  ->  (pi1, pi2) = (0.000, 1.000),   max value over types = 0.0000
  q = -0.50  ->  (pi1, pi2) = (0.500, 1.000),   max value over types = 0.2500
  q = +0.00  ->  (pi1, pi2) = (1.000, 1.000),   max value over types = 0.5000
  q = +0.50  ->  (pi1, pi2) = (1.000, 0.500),   max value over types = 0.2500
  q = +1.00  ->  (pi1, pi2) = (1.000, 0.000),   max value over types = 0.0000
```

A menu is now a function $q(\theta)$, and incentive compatibility requires it to be non-decreasing.

Types who think ω_1 is more likely want experiments with higher q, which deliver sharper evidence about the state they consider *less* likely.

There is a second, less familiar restriction.

Because information is worthless to types $\theta \in \{0, 1\}$, applying the envelope theorem separately on $[0, \theta^*]$ and $[\theta^*, 1]$ and matching the two expressions for the rent of the pivotal type θ^* forces

$$\int_0^1 q(\theta)\, d\theta = u_1 - u_2. \tag{36.14}$$

Note that this integral is taken with respect to $d\theta$, not $dF(\theta)$.

With those two constraints, the seller's problem reduces to

$$\max_{q(\cdot)} \int_0^1 \Big[(\theta f(\theta) + F(\theta))q(\theta) + \min\big\{(u_1 - u_2 - q(\theta))f(\theta),\ 0\big\}\Big] d\theta \tag{36.15}$$

subject to q non-decreasing and (36.14).

36.7.1 Solving it as a linear program

The integrand of (36.15) is **concave and piecewise linear** in q, since $\min\{(d-q)f, 0\} = -f\max\{q-d, 0\}$ with $d = u_1 - u_2$ and $f \geq 0$.

Maximizing a concave piecewise-linear objective subject to linear constraints is a linear program.

Introducing $z(\theta) \geq \max\{q(\theta) - d,\ 0\}$ and discretizing θ on a grid gives

$$\max_{q,z} \sum_n w_n \Big[(\theta_n f_n + F_n) q_n - f_n z_n\Big]$$

subject to $z_n \geq q_n - d$, $z_n \geq 0$, $q_{n+1} \geq q_n$, $-u_2 \leq q_n \leq u_1$, and $\sum_n w_n q_n = d$.

```
def solve_menu(theta, f, u1=1.0, u2=1.0):
    """Solve the seller's problem on a grid of types by linear programming."""
    N = len(theta)
    dth = theta[1] - theta[0]
    f = f / (np.sum(f) * dth)                                  # density on the grid
    F = np.cumsum(f) * dth
    w = np.full(N, dth)
    d = u1 - u2

    c = np.concatenate([-(theta * f + F) * w, f * w])          # linprog minimizes
    A_ub = np.hstack([np.eye(N), -np.eye(N)])                  # q - z <= d
    b_ub = np.full(N, d)
    D = np.zeros((N - 1, 2 * N))                               # q_n - q_{n+1} <= 0
    rows = np.arange(N - 1)
    D[rows, rows], D[rows, rows + 1] = 1.0, -1.0
    A_ub = np.vstack([A_ub, D])
    b_ub = np.concatenate([b_ub, np.zeros(N - 1)])
    A_eq = np.concatenate([w, np.zeros(N)])[None, :]           # integral constraint
    bounds = [(-u2, u1)] * N + [(0, None)] * N

    res = linprog(c, A_ub=A_ub, b_ub=b_ub, A_eq=A_eq, b_eq=np.array([d]),
                  bounds=bounds, method='highs')
    return res.x[:N], res
```

The linear program handles the monotonicity constraint automatically.

This matters, because the alternative is to implement Myerson's **ironing** procedure by hand: form the virtual values

$$\phi^-(\theta) = \theta f(\theta) + F(\theta), \qquad \phi^+(\theta) = (\theta - 1) f(\theta) + F(\theta), \tag{36.16}$$

replace them by the derivatives of the convex hulls of their integrals, and then find the multiplier on (36.14) (see Myerson [1981] and Toikka [2011]).

Denote the ironed virtual values by $\bar{\phi}^-$ and $\bar{\phi}^+$.

Bergemann *et al.* [2018] show that a menu q^* is optimal if and only if it satisfies (36.14), it is constant on every interval where the relevant ironed virtual value is constant, and for some multiplier $\lambda^* > 0$ and every θ, $q^*(\theta)$ maximizes

$$\bar{\phi}^-(\theta)\min\{q, d\} + \bar{\phi}^+(\theta)\max\{0, q - d\} - \lambda^* q \tag{36.17}$$

over $q \in [-u_2, u_1]$.

Where $\bar{\phi}^+(\theta) \leq \bar{\phi}^-(\theta)$, the maximizer is $q^* = -u_2$ if $\bar{\phi}^-(\theta) < \lambda^*$, $q^* = d$ if $\bar{\phi}^+(\theta) < \lambda^* < \bar{\phi}^-(\theta)$, and $q^* = u_1$ if $\bar{\phi}^+(\theta) > \lambda^*$.

The linear program does all of that implicitly.

The multiplier λ^* on (36.14) is the dual value of the equality constraint, `-res.eqlin.marginals[0]`; the sign flips because `linprog` minimizes minus revenue.

We also want the prices, which follow from the requirement that a buyer at the boundary between two items be indifferent between them.

```
def menu_items(theta, q, u1=1.0, u2=1.0, tol=1e-4, min_width=0.01):
    """Distinct items in the menu, with the interval of types served and the price.

    Values of q taken on a negligible set of types are transition artifacts of the
    grid, not items on the menu, so we drop them.
    """
    qr = np.round(q / tol) * tol
    vals = [v for v in np.unique(qr)
            if theta[qr == v].max() - theta[qr == v].min() >= min_width]
    items = sorted([(v, theta[qr == v].min(), theta[qr == v].max()) for v in vals],
                   key=lambda x: x[1])
    out, prev_v, prev_t = [], None, 0.0
    for v, lo, hi in items:
        if value_q(v, theta, u1, u2).max() < 1e-9:      # uninformative item
            price = 0.0
        elif prev_v is None:
            price = 0.0
        else:
            price = float(value_q(v, lo, u1, u2)
                          - value_q(prev_v, lo, u1, u2) + prev_t)
        out.append((v, lo, hi, price))
        prev_v, prev_t = v, price
    return out
```

36.7.2 Uniformly distributed types

With $u_1 = u_2 = 1$ and θ uniform, the virtual values are $\phi^-(\theta) = 2\theta$ and $\phi^+(\theta) = 2\theta - 1$.

Both are strictly increasing, so no ironing is required and the optimal menu should contain a single informative item.

```
N = 2001
theta_g = np.linspace(0, 1, N)
q_unif, res = solve_menu(theta_g, np.ones(N))

print('LP status:', res.message)
print('distinct values of q:', np.unique(np.round(q_unif, 4)))
for v, lo, hi, p in menu_items(theta_g, q_unif):
    p1, p2 = q_to_experiment(v)
    print(f'   q = {v:+.4f}  (pi1, pi2) = ({p1:.3f}, {p2:.3f})'
          f'   types [{lo:.3f}, {hi:.3f}]   price {p:.4f}')
```

```
LP status: Optimization terminated successfully. (HiGHS Status 7: Optimal)
distinct values of q: [-1. -0.  1.]
   q = -1.0000  (pi1, pi2) = (0.000, 1.000)   types [0.000, 0.249]   price 0.0000
   q = -0.0000  (pi1, pi2) = (1.000, 1.000)   types [0.250, 0.750]   price 0.2500
   q = +1.0000  (pi1, pi2) = (1.000, 0.000)   types [0.751, 1.000]   price 0.0000
```

The seller offers the fully informative experiment to the middle range of types at a single price and nothing to anyone else.

The cutoffs and the price match the analytic solution of Bergemann *et al.* [2018] exactly: full information to $\theta \in [1/4, 3/4]$ at a price of $1/4$.

This is the "no-haggling" outcome of Riley and Zeckhauser [1983] adapted to information.

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4))
axes[0].plot(theta_g, 2 * theta_g, lw=2, label=r'$\phi^-(\theta) = 2\theta$')
axes[0].plot(theta_g, 2 * theta_g - 1, lw=2, label=r'$\phi^+(\theta) = 2\theta - 1$')
axes[0].axhline(0.5, color='0.4', ls='--', lw=1, label=r'$\lambda^* = 1/2$')
axes[0].set(xlabel=r'$\theta$', title='virtual values, both strictly increasing')
axes[0].legend(fontsize=9)

axes[1].step(theta_g, q_unif, lw=2, where='mid')
axes[1].set(xlabel=r'$\theta$', ylabel=r'$q^*(\theta)$', ylim=(-1.15, 1.15),
            title='optimal menu: one informative item')
axes[1].annotate('no information', (0.06, -0.85), fontsize=9)
axes[1].annotate('full information', (0.38, 0.12), fontsize=9)
axes[1].annotate('no information', (0.78, 0.85), fontsize=9)
fig.tight_layout()
plt.show()
```

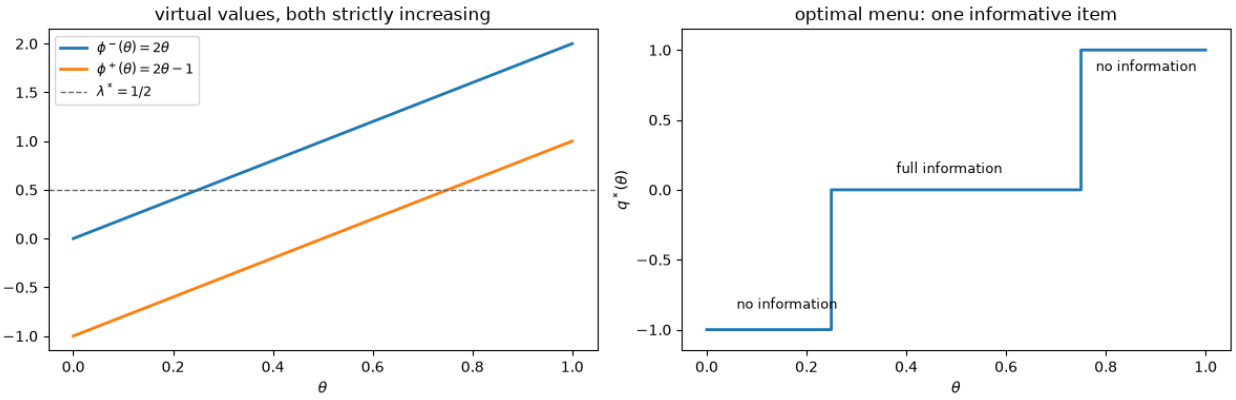


Fig. 36.4: Optimal menu with uniformly distributed types

36.7.3 Bimodal types and the case for versioning

Corollary 1 of Bergemann *et al.* [2018] says that a second experiment is offered only when the virtual values require ironing.

Differentiating (36.16) gives

$$\phi^{-\prime}(\theta) = 2f(\theta) + \theta f'(\theta), \qquad \phi^{+\prime}(\theta) = 2f(\theta) - (1-\theta)f'(\theta). \tag{36.18}$$

So ϕ^- is increasing if and only if $\theta^2 f(\theta)$ is increasing, and ϕ^+ is increasing if and only if $(1-\theta)^2 f(\theta)$ is decreasing.

Regularity fails where the density falls faster than θ^{-2} or rises faster than $(1-\theta)^{-2}$, that is, where the density is steep.

Bergemann *et al.* [2018] remark informally that ironing becomes necessary when most buyers are well informed ex ante, with most types close to 0 or 1.

That remark needs qualification.

For a Beta(a, b) density both virtual values are increasing if and only if $a \leq 1$ and $b \leq 1$ (Exercise 36.9.4), so U-shaped Beta populations of well-informed buyers are regular and are offered a single item.

Steep clusters of beliefs are irregular whether they sit near the ends of $[0, 1]$ or in its interior, and Exercise 36.9.4 exhibits a well-informed population of that kind that is offered two items.

Irregularity is necessary for a second item but not sufficient.

Beta$(30, 30)$ is irregular, but it is symmetric about $1/2$, and with $u_1 = u_2$ Corollary 1 of Bergemann *et al.* [2018] then delivers a single item.

We follow the paper and take an equal mixture of Beta$(8, 30)$ and Beta$(60, 30)$.

Each component concentrates beliefs in a narrow range, so the mixture has two steep clusters, near $\theta \approx 0.2$ and $\theta \approx 0.67$.

```
f_bimodal = (0.5 * stats.beta(8, 30).pdf(theta_g)
             + 0.5 * stats.beta(60, 30).pdf(theta_g))

def falling_intervals(theta, phi, min_width=0.01):
    """Intervals of types on which phi decreases."""
    down = np.concatenate([[0], (np.diff(phi) < 0).astype(int), [0]])
    edges = np.flatnonzero(np.diff(down))
    return [(theta[a], theta[b]) for a, b in zip(edges[::2], edges[1::2])
            if theta[b] - theta[a] >= min_width]


dth = theta_g[1] - theta_g[0]
F_bi = np.cumsum(f_bimodal) * dth
for name, phi in [('phi-', theta_g * f_bimodal + F_bi),
                  ('phi+', (theta_g - 1) * f_bimodal + F_bi)]:
    print(f'{name} decreases on',
          ', '.join(f'[{a:.3f}, {b:.3f}]'
                    for a, b in falling_intervals(theta_g, phi)))

q_bi, res_bi = solve_menu(theta_g, f_bimodal)

print('\nLP status:', res_bi.message)
print('distinct values of q:', np.unique(np.round(q_bi, 3)))
print(f'multiplier lambda* = {-res_bi.eqlin.marginals[0]:.4f}')
print()
for v, lo, hi, p in menu_items(theta_g, q_bi):
    p1, p2 = q_to_experiment(v)
    label = 'no information' if abs(p) < 1e-9 else (
        'full information' if abs(v) < 1e-6 else 'partial information')
    print(f'   q = {v:+.4f}  (pi1, pi2) = ({p1:.3f}, {p2:.3f})'
          f'   types [{lo:.3f}, {hi:.3f}]   price {p:.4f}   {label}')
```

```
phi- decreases on [0.237, 0.472], [0.678, 0.955]
phi+ decreases on [0.000, 0.184], [0.475, 0.655]

LP status: Optimization terminated successfully. (HiGHS Status 7: Optimal)
distinct values of q: [-1.    -0.214 -0.     1.   ]
multiplier lambda* = 0.6684

   q = -1.0000  (pi1, pi2) = (0.000, 1.000)   types [0.000, 0.175]   price 0.0000 ␣
↪ no information
   q = -0.2141  (pi1, pi2) = (0.786, 1.000)   types [0.175, 0.550]   price 0.1375 ␣
↪ partial information
   q = -0.0000  (pi1, pi2) = (1.000, 1.000)   types [0.551, 0.745]   price 0.2555 ␣
↪ full information
   q = +1.0000  (pi1, pi2) = (1.000, 0.000)   types [0.745, 1.000]   price 0.0000 ␣
↪ no information
```

ϕ^- decreases on the upper flank of each cluster, where the density falls steeply; ϕ^+ decreases on the lower flank of each cluster, where it rises steeply.

Both virtual values require ironing.

Now the menu contains **two** informative items, in line with *Proposition 36.5.1* and with the result of Bergemann *et al.* [2018] that some optimal menu contains at most two.

The partial item has $\pi_2 = 1$, so signal s_1 occurs only in state ω_1 and perfectly reveals it, while signal s_2 leaves the buyer uncertain.

Types $\theta \in [0.175, 0.550]$ buy it.

That interval contains $\theta^* = 1/2$, the type least sure what to do, who values information most.

The type indifferent between the two informative items is $\theta \approx 0.55$.

Selling $\overline{E}$ to θ^* would require a further distortion of the partial item and a lower price for it, and that would cost revenue from the many types near $\theta \approx 0.2$.

The partial item lets the seller charge a high price for $\overline{E}$ to the large mass of buyers near $\theta \approx 0.7$.

```
items = menu_items(theta_g, q_bi)

fig, axes = plt.subplots(1, 2, figsize=(12, 4))
axes[0].plot(theta_g, f_bimodal, lw=2, color='C2')
axes[0].fill_between(theta_g, f_bimodal, alpha=0.2, color='C2')
axes[0].set(xlabel=r'$\theta$', ylabel='density',
            title='beliefs packed into two steep clusters')

axes[1].step(theta_g, q_bi, lw=2, where='mid')
axes[1].set(xlabel=r'$\theta$', ylabel=r'$q^*(\theta)$', ylim=(-1.15, 1.15),
            title='optimal menu: two informative items')
for v, lo, hi, p in items:
    if p > 1e-9:
        axes[1].annotate(f'price {p:.3f}', ((lo + hi) / 2, v + 0.12),
                         ha='center', fontsize=9)
fig.tight_layout()
plt.show()
```

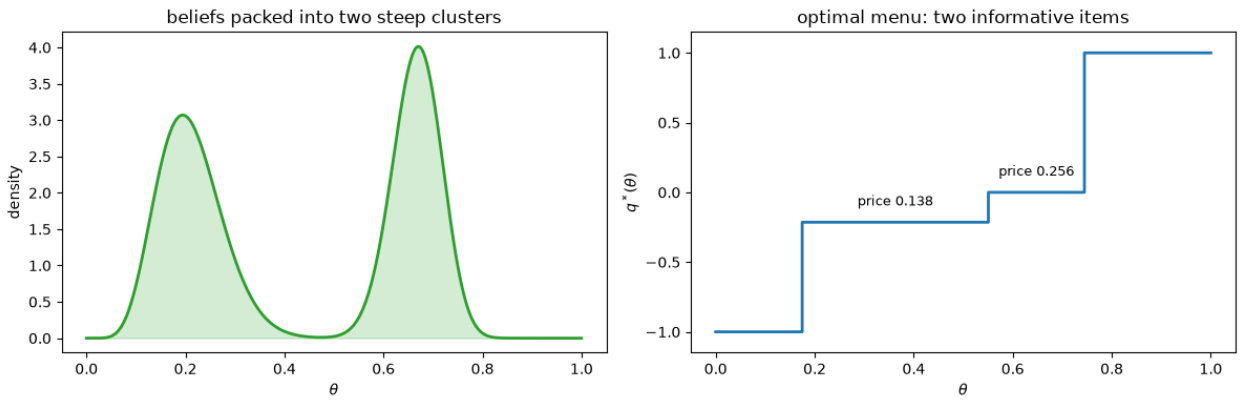


Fig. 36.5: Bimodal type density and the resulting two-item menu

We can see directly why the seller bothers.

Every type values $\overline{E}$ at least as much as any other experiment, so the best single-item menu sells $\overline{E}$ at a posted price p.

With $u_1 = u_2 = 1$, $V(\overline{E}, \theta) = \min\{\theta, 1-\theta\}$, so types $\theta \in [p, 1-p]$ buy and revenue is $p\,[F(1-p) - F(p)]$.

```
def revenue(theta, q, f, u1=1.0, u2=1.0):
    """Expected revenue from the menu q under density f."""
    dth = theta[1] - theta[0]
    price = np.zeros_like(theta)
    for v, lo, hi, p in menu_items(theta, q, u1, u2):
        price[(theta >= lo) & (theta <= hi)] = p
    return np.sum(price * f) * dth / (np.sum(f) * dth)


def F_mix(x):
    """Distribution function of the bimodal mixture."""
    return 0.5 * stats.beta(8, 30).cdf(x) + 0.5 * stats.beta(60, 30).cdf(x)


p_post = np.linspace(0, 0.5, 5001)
rev_post = p_post * (F_mix(1 - p_post) - F_mix(p_post))
k = rev_post.argmax()
rev_two = revenue(theta_g, q_bi, f_bimodal)

print(f'revenue, optimal two-item menu      {rev_two:.4f}')
print(f'revenue, best posted price for E_bar {rev_post[k]:.4f}'
      f'   at p = {p_post[k]:.3f}')
print(f'gain from versioning                {100 * (rev_two / rev_post[k] - 1):.1f}%')
```

```
revenue, optimal two-item menu      0.1675
revenue, best posted price for E_bar 0.1547   at p = 0.223
gain from versioning                8.3%
```

Adding the partial item raises revenue by about 8 percent over the best posted price for full information.

The partial experiment is not a noisier version of the same product; it is a *differently positioned* one, cheap enough for the well-informed types and useless enough to the ill-informed ones that it does not undercut the high price.

36.8 Concluding remarks

Blackwell's theorem tells us when *all* decision makers agree that one experiment beats another.

Read as a design principle, its real content is the size of the set where it is silent.

Bergemann *et al.* [2018] show that a monopolist selling data lives in that set, since screening by belief requires products that different types rank differently.

Two lessons carry beyond the model.

First, optimal degradation of information is structured rather than random: every experiment on the menu has a signal that perfectly reveals one state.

Bergemann *et al.* [2018] conclude that no product on an optimal menu should add unbiased noise to the seller's data, unless buyers compete with each other downstream or the data must be anonymized.

Second, versioning requires virtual values that must be ironed, and steep densities of beliefs produce them.

How well informed the buyers are does not settle the matter: a U-shaped Beta population of well-informed buyers is offered a single item, while steep clusters of beliefs, near the ends of $[0, 1]$ or in its interior, can be offered two.

Selling information to imperfectly informed buyers has a long history.

Admati and Pfleiderer [1986] study a seller facing a continuum of *ex ante identical* traders who then trade a common-value asset, and find that the seller wants to supply noisy and *idiosyncratic* information, so that each trader retains a local monopoly on what he knows.

Their traders compete downstream, so the first lesson does not apply to them.

The heterogeneity there is created by the seller; here it is the buyer's own prior information, and that is what turns the problem into one of screening.

Bergemann and Bonatti [2015] study a single data provider who sets a price per "cookie", a record on an individual consumer, and an advertiser with no private information who chooses which cookies to buy.

A useful contrast is Kamenica and Gentzkow [2011], where a sender also commits to an information structure but has no monetary transfers and cares directly about the receiver's action; here the seller cares only about revenue and cannot condition payments on the state, the signal, or the buyer's action.

Readers who want the statistical background can return to *Blackwell's Theorem on Comparing Experiments* for the equivalence between the economic, sufficiency, and uncertainty-reduction criteria, to *Likelihood Ratio Processes and Bayesian Learning* for how private signals generate the interim beliefs that are the buyer types here, and to *Information and Market Equilibrium* for what happens when information is transmitted by prices rather than sold directly.

36.9 Exercises

Exercise 36.9.1

This exercise studies the **congruent** case, in which both types would take the same action without extra information.

Set $u_1 = u_2 = 1$, $\theta^L = 0.9$ and $\theta^H = 0.7$, so that $\theta^* = 1/2 < \theta^H < \theta^L$.

Because both types would choose a_1 on their own, the seller has no reason to degrade what the low type learns about ω_1, so set $\pi_1 = 1$ and treat π_2 as the only choice variable.

1. Plot the seller's revenue against $\pi_2 \in [0, 1]$ for several values of γ, and confirm that it is *linear*.
2. Conclude that the optimum is always at an endpoint, so the low type receives either full information or none.
3. Bergemann *et al.* [2018] show that the low type receives the fully informative experiment if and only if $\gamma \leq (1 - \theta^L)/(1 - \theta^H)$.

   Locate the switch point numerically by bisection and compare.

Why is the answer extremal here, when the noncongruent case in the lecture produced an interior $\pi_1 = 4/5$?

Solution

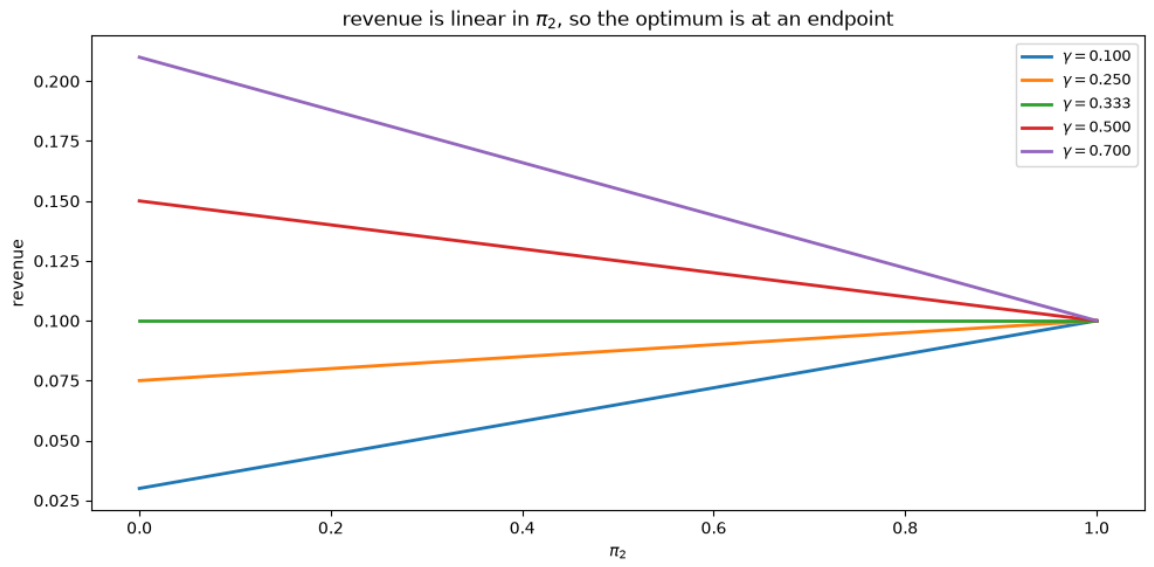
Here is one solution:

```
tL_c, tH_c = 0.9, 0.7
p2_grid = np.linspace(0, 1, 401)

fig, ax = plt.subplots()
for gamma in [0.1, 0.25, 1/3, 0.5, 0.7]:
    r = np.array([two_type_revenue(1.0, p2, tL_c, tH_c, gamma) for p2 in p2_grid])
    dev = np.abs(r - np.interp(p2_grid, [0, 1], [r[0], r[-1]])).max()
    ax.plot(p2_grid, r, lw=2, label=rf'$\gamma = {gamma:.3f}$')
    print(f'gamma = {gamma:.3f}:  revenue at pi2=0 is {r[0]:.5f}, '
```

```
          f'at pi2=1 is {r[-1]:.5f},  deviation from linear {dev:.1e}')
ax.set(xlabel=r'$\pi_2$', ylabel='revenue',
       title='revenue is linear in $\pi_2$, so the optimum is at an endpoint')
ax.legend(fontsize=9)
fig.tight_layout()
plt.show()
```

```
gamma = 0.100:  revenue at pi2=0 is 0.03000, at pi2=1 is 0.10000,  deviation␣
 ↪from linear 8.3e-17
gamma = 0.250:  revenue at pi2=0 is 0.07500, at pi2=1 is 0.10000,  deviation␣
 ↪from linear 8.3e-17
gamma = 0.333:  revenue at pi2=0 is 0.10000, at pi2=1 is 0.10000,  deviation␣
 ↪from linear 9.7e-17
gamma = 0.500:  revenue at pi2=0 is 0.15000, at pi2=1 is 0.10000,  deviation␣
 ↪from linear 9.7e-17
gamma = 0.700:  revenue at pi2=0 is 0.21000, at pi2=1 is 0.10000,  deviation␣
 ↪from linear 1.1e-16
```




```
lo, hi = 0.0, 1.0
for _ in range(60):
    mid = (lo + hi) / 2
    if two_type_revenue(1, 1, tL_c, tH_c, mid) >= two_type_revenue(1, 0, tL_c, tH_
 ↪c, mid):
        lo = mid
    else:
        hi = mid

print(f'numerical switch point            gamma = {lo:.6f}')
print(f'(1 - theta_L) / (1 - theta_H)           = {(1 - tL_c) / (1 - tH_c):.6f}')
```

```
numerical switch point            gamma = 0.333333
(1 - theta_L) / (1 - theta_H)           = 0.333333
```

Revenue is linear in π_2 to machine precision, so the optimum is always at $\pi_2 \in \{0, 1\}$, and the switch occurs exactly at $\gamma = 1/3$ as predicted.

The reason for the extremal answer is that with congruent beliefs both types would choose a_1 anyway, so the only question is how much the seller reveals about ω_2.

Both types then value the experiment through the same term $(1-\theta)\pi_2 u_2$, which is why the objective and constraints

are linear in the single variable π_2 and why the no-haggling logic of Riley and Zeckhauser [1983] applies.

With noncongruent beliefs the two types take different actions on their own, the kink in the value function lies between them, and the seller can position an experiment so that it is worth much to one type and little to the other.

That possibility is what makes an interior distortion optimal.

Exercise 36.9.2

Corollary 1 of Bergemann *et al.* [2018] states that the optimal menu contains a single item whenever both virtual values (36.16) are strictly increasing, and that for uniformly distributed types this holds **irrespective of the payoffs** (u_1, u_2).

1. Verify this by solving the seller's problem with uniform types for several asymmetric payoff pairs, for instance $(u_1, u_2) \in \{(1,1), (2,1), (1,3), (5,1)\}$.
2. For each case report θ^*, the interval of types served, and the price.
3. Confirm that the fully informative experiment is always the item offered, as *Proposition 36.5.1* requires.
4. Use (36.17) and (36.14) to show analytically that the multiplier is $\lambda^* = \theta^*$, that types $\theta \in [\theta^*/2,\ (1+\theta^*)/2]$ are served, and that the price is $u_1 u_2/(2(u_1+u_2))$.

   Print these values next to the ones from the linear program.

Solution

Here is one solution:

```
print(f'{"u1":>3s}{"u2":>3s}{"theta*":>8s}{"q":>6s}{"":3s}'
      f'{"served (LP)":<16s}{"served (theory)":<17s}'
      f'{"price":>7s}{"theory":>8s}{"lambda*":>9s}{"theory":>8s}')
for u1_, u2_ in [(1, 1), (2, 1), (1, 3), (5, 1)]:
    q_a, res_a = solve_menu(theta_g, np.ones(N), u1_, u2_)
    star = u2_ / (u1_ + u2_)
    served = [it for it in menu_items(theta_g, q_a, u1_, u2_) if it[3] > 1e-9]
    assert len(served) == 1                      # a single informative item
    v, lo, hi, p = served[0]
    assert abs(v - (u1_ - u2_)) < 1e-3           # the item is fully informative
    lam = -res_a.eqlin.marginals[0]
    print(f'{u1_:3d}{u2_:3d}{star:8.4f}{v:6.1f}{"":3s}'
          f'{f"[{lo:.3f}, {hi:.3f}]":<16s}'
          f'{f"[{star / 2:.3f}, {(1 + star) / 2:.3f}]":<17s}'
          f'{p:7.4f}{u1_ * u2_ / (2 * (u1_ + u2_)):8.4f}'
          f'{lam:9.4f}{star:8.4f}')
print('\nevery menu contains exactly one informative item, '
      'and it is the fully informative one')
```

```
u1 u2  theta*     q   served (LP)      served (theory)     price  theory ␣
↪lambda*  theory
 1  1  0.5000  -0.0   [0.250, 0.750]   [0.250, 0.750]    0.2500  0.2500   0.
↪5002  0.5000
 2  1  0.3333   1.0   [0.167, 0.666]   [0.167, 0.667]    0.3340  0.3333   0.
↪3343  0.3333
 1  3  0.7500  -2.0   [0.375, 0.875]   [0.375, 0.875]    0.3750  0.3750   0.
↪7501  0.7500
 5  1  0.1667   4.0   [0.084, 0.584]   [0.083, 0.583]    0.4200  0.4167   0.
↪1674  0.1667

every menu contains exactly one informative item, and it is the fully
↪informative one
```

The virtual values for a uniform density are $\phi^-(\theta) = 2\theta$ and $\phi^+(\theta) = 2\theta - 1$ whatever the payoffs, because u_1 and u_2 enter the seller's problem only through $d = u_1 - u_2$ and the bounds on q, not through f or F.

Both are strictly increasing, so no ironing is needed and a single item is optimal.

The payoffs do move θ^* and hence which types are served and at what price, but they never make versioning worthwhile under a uniform density.

For the analytic benchmark, no ironing is needed, so (36.17) gives $q^* = -u_2$ for $\theta < \lambda^*/2$, $q^* = u_1 - u_2$ for $\lambda^*/2 < \theta < (1+\lambda^*)/2$, and $q^* = u_1$ for $\theta > (1+\lambda^*)/2$.

Substituting into (36.14) gives

$$-u_2 \frac{\lambda^*}{2} + (u_1 - u_2)\frac{1}{2} + u_1 \frac{1-\lambda^*}{2} = u_1 - u_2,$$

so $\lambda^*(u_1 + u_2) = u_2$ and $\lambda^* = \theta^*$.

The lowest type served, $\theta^*/2$, is indifferent between $\overline{E}$ and nothing.

For $\theta \leq \theta^*$ we have $V(\overline{E}, \theta) = \theta u_1$, so the price is $u_1\theta^*/2 = u_1 u_2/(2(u_1+u_2))$.

The highest type served, $(1+\theta^*)/2$, values $\overline{E}$ at $(1-\theta)u_2 = u_1 u_2/(2(u_1+u_2))$, the same price.

The linear program matches all three objects up to the discretization of the type grid.

Exercise 36.9.3

This exercise connects the lecture back to *Blackwell's Theorem on Comparing Experiments*.

Blackwell's theorem says that if E' is a garbling of E then *every* decision maker weakly prefers E.

1. Draw many random pairs of binary experiments with $\pi_1 + \pi_2 \geq 1$.
2. For each pair, use `garbling` to decide whether one is a garbling of the other, and separately compute whether one dominates the other in value at every type.

   Because $V(E, \theta)$ is piecewise linear in θ, it suffices to compare the two values at $\theta \in \{0, 1\}$ and at the types where either value function has a kink.
3. Confirm that garbling implies unanimous preference, and report what fraction of random pairs Blackwell's order fails to rank.
4. Among the unranked pairs, verify that some types prefer one experiment and some the other.
5. Explain why, for two-state experiments, the pairs that Blackwell's order leaves unranked should be exactly the pairs on which types disagree.

   What goes wrong if dominance is checked only on a finite grid of types?

Solution

Here is one solution:

```
def kinks(pi1, pi2, u1=1.0, u2=1.0):
    """Endpoints and types at which V(E, theta) can change slope."""
    k = [0.0, 1.0, u2 / (u1 + u2)]
    for a, b in [(pi1 * u1, (1 - pi2) * u2), ((1 - pi1) * u1, pi2 * u2)]:
        if a + b > 0:
            k.append(b / (a + b))        # indifference after one of the signals
```

```
    return np.array(k)


rng = np.random.default_rng(0)
grid_t = np.linspace(0.001, 0.999, 999)

n_pairs, n_garble, n_unranked, violations = 4000, 0, 0, 0
n_disagree, n_disagree_grid, n_unanimous_unranked = 0, 0, 0
for _ in range(n_pairs):
    (a1, a2), (b1, b2) = rng.uniform(0, 1, 2), rng.uniform(0, 1, 2)
    if a1 + a2 < 1 or b1 + b2 < 1:
        continue
    Ea, Eb = experiment(a1, a2), experiment(b1, b2)
    t = np.concatenate([kinks(a1, a2), kinks(b1, b2)])
    va, vb = value(a1, a2, t), value(b1, b2, t)
    va_g, vb_g = value(a1, a2, grid_t), value(b1, b2, grid_t)

    a_garbles_b = garbling(Ea, Eb) is not None      # Eb is a garbling of Ea
    b_garbles_a = garbling(Eb, Ea) is not None
    a_dominates = np.all(va >= vb - 1e-12)
    b_dominates = np.all(vb >= va - 1e-12)

    if a_garbles_b:
        n_garble += 1
        if not a_dominates:
            violations += 1
    if b_garbles_a:
        n_garble += 1
        if not b_dominates:
            violations += 1
    if not (a_garbles_b or b_garbles_a):
        n_unranked += 1
        if not (a_dominates or b_dominates):
            n_disagree += 1
        if not (np.all(va_g >= vb_g - 1e-9) or np.all(vb_g >= va_g - 1e-9)):
            n_disagree_grid += 1
    if (a_dominates and not a_garbles_b) or (b_dominates and not b_garbles_a):
        n_unanimous_unranked += 1

print(f'garbling relations found              {n_garble}')
print(f'violations of Blackwell               {violations}')
print(f'unanimous rankings without garbling   {n_unanimous_unranked}')
print(f'pairs unranked by Blackwell           {n_unranked}')
print(f'  of which types disagree (kinks)     {n_disagree}')
print(f'  of which types disagree (grid)      {n_disagree_grid}')
```

```
garbling relations found              674
violations of Blackwell               0
unanimous rankings without garbling   0
pairs unranked by Blackwell           318
  of which types disagree (kinks)     318
  of which types disagree (grid)      317
```

Blackwell's theorem is never violated: whenever one experiment garbles into the other, every type prefers the garbling source.

The converse also holds in every draw: whenever all types rank two experiments the same way, one is a garbling of the other.

A large share of random pairs is left unranked, and for every one of them some types prefer one experiment and some the other.

That unranked region is exactly the room the data seller needs.

If Blackwell's order were complete, every buyer would agree on the ranking of all information products, the seller's problem would collapse to standard nonlinear pricing over a single quality index, and the horizontal screening described in this lecture would be impossible.

The match between garbling and unanimous preference is exact for two-state experiments.

Fix a reference prior $c \in (0,1)$ and let x denote the posterior probability of ω_1 that a signal induces under c.

The gross value of E to type θ with $u_1 = u_2 = 1$ equals the expectation, under c, of the convex function $k_\theta(x) = \max\{\theta x/c,\ (1-\theta)(1-x)/(1-c)\}$, which has a single kink.

As θ ranges over $(0,1)$, the kink ranges over $(0,1)$.

Every convex function on $[0,1]$ is an affine function plus a nonnegative mixture of such one-kink functions, and the affine part has the same expectation under every experiment.

By *Theorem 30.5.1*, E' is a garbling of E if and only if $V(E,\theta) \geq V(E',\theta)$ for every θ.

A check on a finite grid of types can miss a disagreement confined to a short interval between grid points.

The grid of 999 types does so once in the draws above.

Evaluating at kinks is exact, because the difference of two piecewise-linear functions attains its maximum and minimum at kinks or endpoints.

Exercise 36.9.4

This exercise characterizes the type densities for which the virtual values (36.16) need no ironing.

1. Use (36.18) to show that on $(0,1)$, ϕ^- is strictly increasing if and only if $\theta^2 f(\theta)$ is strictly increasing, and ϕ^+ is strictly increasing if and only if $(1-\theta)^2 f(\theta)$ is strictly decreasing.
2. Deduce that for a Beta(a,b) density both virtual values are strictly increasing if and only if $a \leq 1$ and $b \leq 1$.
3. Confirm the condition numerically on a grid of (a,b) pairs.
4. Use `solve_menu` to show that the U-shaped densities Beta$(0.5,0.5)$, Beta$(0.3,0.3)$ and Beta$(0.5,0.8)$ are offered a single item, that the irregular Beta$(30,30)$ is also offered a single item, and that the irregular Beta$(60,30)$ is offered two informative items.
5. Show that an equal mixture of Beta$(2,60)$ and Beta$(40,4)$, a population in which most buyers are well informed, is irregular and is offered two informative items.

Solution

Here is one solution.

For $\theta \in (0,1)$,

$$\frac{d}{d\theta}\left[\theta^2 f(\theta)\right] = \theta\left[2f(\theta) + \theta f'(\theta)\right] = \theta\,\phi^{-\prime}(\theta), \qquad \frac{d}{d\theta}\left[(1-\theta)^2 f(\theta)\right] = -(1-\theta)\,\phi^{+\prime}(\theta).$$

So $\phi^{-\prime}$ has the sign of the derivative of $\theta^2 f$, and $\phi^{+\prime}$ has the sign opposite to the derivative of $(1-\theta)^2 f$.

For $f(\theta) \propto \theta^{a-1}(1-\theta)^{b-1}$,

$$\frac{d}{d\theta}\log[\theta^2 f(\theta)] = \frac{a+1}{\theta} - \frac{b-1}{1-\theta}, \qquad \frac{d}{d\theta}\log[(1-\theta)^2 f(\theta)] = \frac{a-1}{\theta} - \frac{b+1}{1-\theta}.$$

The first expression is positive on all of $(0,1)$ if and only if $b \leq 1$, since for $b > 1$ it diverges to $-\infty$ as $\theta \to 1$.

The second is negative on all of $(0,1)$ if and only if $a \leq 1$, since for $a > 1$ it diverges to $+\infty$ as $\theta \to 0$.

```
th_in = theta_g[1:-1]                    # Beta densities can be infinite at 0 and 1


def mixture(components):
    """Density and distribution function of an equal mixture of Beta laws."""
    f_c = np.mean([stats.beta(a, b).pdf(th_in) for a, b in components], axis=0)
    F_c = np.mean([stats.beta(a, b).cdf(th_in) for a, b in components], axis=0)
    return f_c, F_c


def regular(f_c, F_c):
    """True if both virtual values increase on the grid."""
    return (np.all(np.diff(th_in * f_c + F_c) > 0)
            and np.all(np.diff((th_in - 1) * f_c + F_c) > 0))


ab_grid = [0.3, 0.5, 0.8, 1.0, 1.2, 2.0, 5.0]
mismatch = [(a, b) for a in ab_grid for b in ab_grid
            if regular(*mixture([(a, b)])) != (a <= 1 and b <= 1)]
print(f'(a, b) pairs checked: {len(ab_grid)**2},  mismatches: {mismatch}')

for comps in [[(0.5, 0.5)], [(0.3, 0.3)], [(0.5, 0.8)], [(30, 30)], [(60, 30)],
              [(2, 60), (40, 4)]]:
    f_c, F_c = mixture(comps)
    near_ends = np.mean([stats.beta(a, b).cdf(0.15) + stats.beta(a, b).sf(0.85)
                         for a, b in comps])
    q_c, _ = solve_menu(th_in, f_c)
    informative = [it for it in menu_items(th_in, q_c) if it[3] > 1e-9]
    print('\n' + ' + '.join(f'Beta{ab}' for ab in comps)
          + f':  regular = {regular(f_c, F_c)},'
          f'  mass within 0.15 of 0 or 1 = {near_ends:.2f},'
          f'  informative items = {len(informative)}')
    for v, lo, hi, p in informative:
        print(f'   q = {v:+.3f}   types [{lo:.3f}, {hi:.3f}]   price {p:.4f}')
```

```
(a, b) pairs checked: 49,  mismatches: []

Beta(0.5, 0.5):  regular = True,  mass within 0.15 of 0 or 1 = 0.51, ␣
↪informative items = 1
   q = -0.000   types [0.237, 0.764]   price 0.2365

Beta(0.3, 0.3):  regular = True,  mass within 0.15 of 0 or 1 = 0.64, ␣
↪informative items = 1
   q = -0.000   types [0.230, 0.770]   price 0.2295

Beta(0.5, 0.8):  regular = True,  mass within 0.15 of 0 or 1 = 0.46, ␣
↪informative items = 1
   q = -0.000   types [0.240, 0.761]   price 0.2395

Beta(30, 30):  regular = False,  mass within 0.15 of 0 or 1 = 0.00, ␣
↪informative items = 1
```

```
   q = -0.000   types [0.384, 0.617]   price 0.3835

Beta(60, 30):  regular = False,  mass within 0.15 of 0 or 1 = 0.00, ␣
↪informative items = 2
   q = -0.821   types [0.164, 0.289]   price 0.0294
   q = -0.000   types [0.289, 0.733]   price 0.2670

Beta(2, 60) + Beta(40, 4):  regular = False,  mass within 0.15 of 0 or 1 = 0.95,
↪  informative items = 2
   q = -0.058   types [0.027, 0.719]   price 0.0254
   q = -0.000   types [0.719, 0.933]   price 0.0670
```

The numerical check agrees with the condition $a \leq 1, b \leq 1$ at every pair.

The three U-shaped densities are regular and are offered only $\overline{E}$.

Beta$(30, 30)$ is irregular but symmetric about $\theta^* = 1/2$, and it is also offered only $\overline{E}$.

The steep, asymmetric Beta$(60, 30)$ is offered a partial item with $q < 0$ as well as $\overline{E}$.

The mixture of Beta$(2, 60)$ and Beta$(40, 4)$ puts more of its mass within 0.15 of 0 or 1 than any of the U-shaped densities.

Its two clusters are steep, it is irregular, and it is offered two informative items.

Both kinds of population consist mostly of well-informed buyers; the steepness of the density separates them.

Exercise 36.9.5

This exercise irons the virtual values of the bimodal example by hand and checks the menu that the linear program found.

Set $u_1 = u_2 = 1$, so that $d = 0$, and use the density `f_bimodal`.

1. Compute $\Phi^{\pm}(\theta) = \int_0^\theta \phi^{\pm}(x)\,dx$ on the grid, their greatest convex minorants, and the ironed virtual values $\bar{\phi}^{\pm}$ as the slopes of the minorants.
2. By (36.17), types with $\bar{\phi}^- < \lambda^*$ get $q = -1$, types with $\bar{\phi}^+ < \lambda^* < \bar{\phi}^-$ get $q = 0$, types with $\bar{\phi}^+ > \lambda^*$ get $q = 1$, and types on an interval where $\bar{\phi}^- = \lambda^*$ share a common $q_m \in [-1, 0]$.

   Find λ^* and q_m from the integral constraint (36.14).
3. Compare λ^*, q_m, the interval of types who buy the partial item, and the cutoff above which types buy nothing with the solution of the linear program.
4. Plot $\phi^{\pm}$, $\bar{\phi}^{\pm}$ and λ^*.

Solution

Here is one solution.

```
def ironed(theta, phi):
    """Slopes of the greatest convex minorant of the integral of phi."""
    dth = theta[1] - theta[0]
    Phi = np.concatenate([[0.0], np.cumsum((phi[1:] + phi[:-1]) / 2) * dth])
    hull = [0]                                          # lower convex hull
    for i in range(1, len(theta)):
        while len(hull) >= 2:
            j, k = hull[-2], hull[-1]
```

```
            if ((Phi[k] - Phi[j]) * (theta[i] - theta[j])
                    >= (Phi[i] - Phi[j]) * (theta[k] - theta[j])):
                hull.pop()
            else:
                break
        hull.append(i)
    phi_bar = np.empty_like(phi)
    for j, k in zip(hull[:-1], hull[1:]):
        phi_bar[j:k + 1] = (Phi[k] - Phi[j]) / (theta[k] - theta[j])
    return phi_bar


dth = theta_g[1] - theta_g[0]
f_n = f_bimodal / (np.sum(f_bimodal) * dth)
F_n = np.cumsum(f_n) * dth
phi_m, phi_p = theta_g * f_n + F_n, (theta_g - 1) * f_n + F_n
bar_m, bar_p = ironed(theta_g, phi_m), ironed(theta_g, phi_p)
print('phi_bar+ <= phi_bar- everywhere:', np.all(bar_p <= bar_m + 1e-12))


def integral_q(lam):
    """Integral of q when types with bar_m = lam are assigned q = 0."""
    return (-(bar_m < lam).sum() + (bar_p > lam).sum()) * dth


lo, hi = bar_m.min(), bar_m.max()                    # integral_q decreases in lam
for _ in range(200):
    mid = (lo + hi) / 2
    lo, hi = (mid, hi) if integral_q(mid) > 0 else (lo, mid)
lam_star = hi

pooled = np.abs(bar_m - lam_star) < 1e-9
th_a, th_b = theta_g[pooled].min(), theta_g[pooled].max()
th_c = theta_g[bar_p > lam_star].min()
q_m = (th_a - (1 - th_c)) / (th_b - th_a)

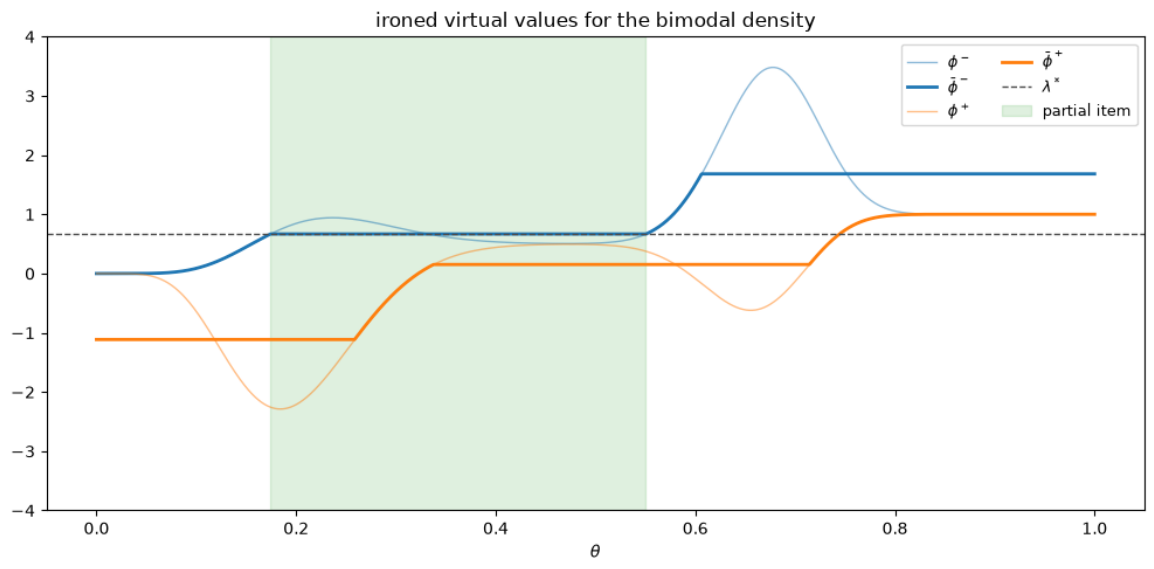
lp_items = [it for it in menu_items(theta_g, q_bi) if it[3] > 1e-9]
(q_lp, a_lp, b_lp, _), (_, _, c_lp, _) = lp_items
print(f'\n{"":24s}{"by hand":>10s}{"LP":>10s}')
print(f'{"multiplier lambda*":24s}{lam_star:10.4f}{-res_bi.eqlin.marginals[0]:10.
 ↪4f}')
print(f'{"partial item q_m":24s}{q_m:10.4f}{q_lp:10.4f}')
print(f'{"lowest buyer":24s}{th_a:10.4f}{a_lp:10.4f}')
print(f'{"partial / full cutoff":24s}{th_b:10.4f}{b_lp:10.4f}')
print(f'{"highest buyer":24s}{th_c:10.4f}{c_lp:10.4f}')
```

```
phi_bar+ <= phi_bar- everywhere: True

                          by hand        LP
multiplier lambda*         0.6684    0.6684
partial item q_m          -0.2141   -0.2141
lowest buyer               0.1745    0.1750
partial / full cutoff      0.5505    0.5505
highest buyer              0.7450    0.7445
```

```
fig, ax = plt.subplots()
ax.plot(theta_g, phi_m, lw=1, color='C0', alpha=0.5, label=r'$\phi^-$')
ax.plot(theta_g, bar_m, lw=2, color='C0', label=r'$\bar\phi^-$')
```

```
ax.plot(theta_g, phi_p, lw=1, color='C1', alpha=0.5, label=r'$\phi^+$')
ax.plot(theta_g, bar_p, lw=2, color='C1', label=r'$\bar\phi^+$')
ax.axhline(lam_star, color='0.3', ls='--', lw=1, label=r'$\lambda^*$')
ax.axvspan(th_a, th_b, color='C2', alpha=0.15, label='partial item')
ax.set(xlabel=r'$\theta$', ylim=(-4, 4),
       title='ironed virtual values for the bimodal density')
ax.legend(fontsize=9, ncol=2)
fig.tight_layout()
plt.show()
```



The multiplier is the level at which the integral of q jumps across zero.

Just below λ^* the pooled types get $q = 0$ and the integral is positive; just above λ^* they get $q = -1$ and the integral is negative.

On the pooled interval $[\theta_a, \theta_b]$ the integral constraint reads $-\theta_a + q_m(\theta_b - \theta_a) + (1 - \theta_c) = 0$, which gives $q_m = (\theta_a - (1 - \theta_c))/(\theta_b - \theta_a)$.

The pooled interval is the interval on which ironing flattens ϕ^- at the level λ^*.

Its upper end, $\theta \approx 0.55$, is the type indifferent between the partial item and $\overline{E}$.

The hand calculation and the linear program agree up to the grid spacing.

CHAPTER

THIRTYSEVEN

MARKET DIFFUSION WITH TWO-SIDED LEARNING

Contents

- *Market Diffusion with Two-Sided Learning*
  - *Overview*
  - *The market*
  - *Two-sided learning*
  - *Efficient experimentation*
  - *Equilibrium*
  - *Too much experimentation, then too little*
  - *Diffusion over time*
  - *Martingale properties*
  - *Concluding remarks*
  - *Exercises*

37.1 Overview

In *The Design and Price of Information* a monopolist *owned* information and sold it.

This lecture studies a market in which nobody sells information and everybody produces it.

We follow Bergemann and Välimäki [1997], who study a duopoly in which an established firm competes on price with a firm selling a new product of unknown quality.

Buyers learn what the new product is worth only by using it, and the aggregate record of their experience is public.

So every purchase of the new product is simultaneously a consumption decision and an experiment, and its informational value spills over to everyone.

Both sides of the market learn from the same public record, which is what "two-sided learning" means here: buyers and sellers hold identical beliefs at every date, and no asymmetric information ever arises.

Three results organize the lecture.

First, both firms *want* more information, but only the new firm's sales produce it.

That asymmetry softens price competition: the established firm prices less aggressively than it would in a one-shot game, and the entrant captures a larger market share early on.

Second, equilibrium experimentation is **excessive** when beliefs are pessimistic and **insufficient** when they are optimistic, with a single crossing in between.

Third, the noise-free diffusion path of a successful new product is **S-shaped**, matching a long empirical tradition, and its inflection occurs at a belief we can pin down exactly.

Note

The connection to the rest of this section runs through the *value of information*.

Blackwell's Theorem on Comparing Experiments shows that a more informative experiment spreads the posterior in the convex order.

A single decision maker's value is a maximum of functions that are linear in the belief, so it is convex, and every decision maker weakly gains from a more informative experiment.

A player's equilibrium value in a game need not be convex, because the other players' strategies also move with the belief.

Here beliefs are a martingale and experimentation controls how fast they spread, so each firm's gain from experimentation is governed by the convexity of its equilibrium value function.

Both equilibrium value functions are convex.

The belief itself is driven by a log-likelihood-ratio process of the kind studied in *Likelihood Ratio Processes*, now run in continuous time.

Let's start with imports.

```
import matplotlib.pyplot as plt
import numpy as np

plt.rcParams['figure.figsize'] = (10, 5)
np.set_printoptions(precision=4, suppress=True)
```

37.2 The market

Buyers are distributed uniformly on $[0, 1]$ and each demands one unit per instant.

The established product delivers value

$$s_n = s + nh \tag{37.1}$$

to buyer n, and the new product delivers

$$\mu_n = \mu + (1 - n)h. \tag{37.2}$$

The parameter $h > 0$ measures horizontal differentiation, so buyers near $n = 0$ are naturally drawn to the new product and buyers near $n = 1$ to the established one.

This is the standard Hotelling structure, with one twist: the vertical quality μ of the new product is **unknown** and can take one of two values,

$$\mu \in \{\mu_L, \mu_H\}, \qquad 0 < s - h < \mu_L < s < \mu_H < s + h. \tag{37.3}$$

The inner inequalities say the new product may be better or worse than the established one.

The outer inequalities say that under full information both firms would retain a positive share of the market, so the innovation is not drastic.

Marginal cost is zero for both firms.

If the new firm serves the buyers in $[0, n]$, the average flow value delivered by each product is

$$\bar{\mu}(n) = \mu + \frac{(2-n)h}{2}, \qquad \bar{s}(n) = s + \frac{(1+n)h}{2}, \tag{37.4}$$

so total surplus per unit of time is $n\bar{\mu}(n) + (1-n)\bar{s}(n)$.

Writing $\mu(\alpha)$ for the expected quality under belief $\alpha = \Pr[\mu = \mu_H]$, a little algebra puts the flow surplus in a convenient quadratic form,

$$F(n, \alpha) = s + \frac{h}{2} + n(\mu(\alpha) - s + h) - n^2 h. \tag{37.5}$$

```
class Market:
    """The duopoly of Bergemann and Valimaki (1997)."""

    def __init__(self, s=4.0, h=1.0, mu_L=3.1, mu_H=4.9, sigma=1.0):
        self.s, self.h = s, h
        self.mu_L, self.mu_H, self.sigma = mu_L, mu_H, sigma
        assert 0 < s - h < mu_L < s < mu_H < s + h, 'condition (4) fails'

    def mu(self, a):
        """Expected quality of the new product under belief a."""
        return (1 - a) * self.mu_L + a * self.mu_H

    def flow_surplus(self, n, a):
        return (self.s + self.h / 2 + n * (self.mu(a) - self.s + self.h)
                - n ** 2 * self.h)
```

37.3 Two-sided learning

A buyer's individual experience is a noisy draw on μ, and since each buyer has measure zero, only the *aggregate* record matters.

When a fraction n of buyers uses the new product, the cumulative market outcome X evolves as

$$dX = n\mu\, dt + \sigma\sqrt{n}\, dB, \tag{37.6}$$

so both the drift and the variance scale with the size of the experiment n.

Everyone observes X, so beliefs stay common.

Since μ takes only two values, the belief $\alpha_t = \Pr[\mu = \mu_H \mid \mathcal{F}_t]$ is a sufficient statistic.

Proposition 37.3.1 (Posterior belief)

The belief α_t is a martingale with zero drift and instantaneous variance

$$n\Sigma^2(\alpha) = n\left[\frac{\alpha(1-\alpha)(\mu_H - \mu_L)}{\sigma}\right]^2. \tag{37.7}$$

This is the standard filtering result for a two-point prior observed through a diffusion; see Liptser and Shiryaev [1977].

Likelihood Ratio Processes and Bayesian Learning and *A Bayesian Formulation of Friedman and Wald's Problem* study the same two-point learning problem in discrete time.

Two features of (37.7) drive everything.

The variance is **linear in** n, so information arrives in proportion to the size of the experiment, and only the new firm's sales generate it.

The variance is proportional to $\alpha^2(1-\alpha)^2$, so learning is fastest when beliefs are most diffuse and grinds to a halt as α approaches 0 or 1.

37.3.1 Learning as a likelihood ratio process

It is worth seeing where (37.7) comes from, because the mechanism is the one studied in *Likelihood Ratio Processes*, transplanted to continuous time.

Over a short interval of length Δ the increment ΔX is normal with mean $n\mu\Delta$ and variance $\sigma^2 n\Delta$ under either hypothesis, so the increment to the **log likelihood ratio** is

$$\Delta\ell = \log\frac{f_H(\Delta X)}{f_L(\Delta X)} = \frac{(\mu_H-\mu_L)\,\Delta X - \frac{1}{2}n\Delta(\mu_H^2-\mu_L^2)}{\sigma^2}. \tag{37.8}$$

Beliefs then follow from Bayes' rule in its log-odds form, exactly as in *Likelihood Ratio Processes and Bayesian Learning*,

$$\log\frac{\alpha_{t+\Delta}}{1-\alpha_{t+\Delta}} = \log\frac{\alpha_t}{1-\alpha_t} + \Delta\ell. \tag{37.9}$$

We implement (37.8) and (37.9) directly, which gives an *exact* Bayesian update at each step rather than a discretization of a stochastic differential equation.

```
def simulate_beliefs(mkt, alpha0, T, dt, mu_true, rng, policy):
    """Simulate beliefs by exact Bayesian updating of the log odds.

    `mu_true` holds the true quality for each path, so the paths run in parallel.
    Returns an array of shape (number of paths, number of steps + 1).
    """
    mu_true = np.atleast_1d(np.asarray(mu_true, dtype=float))
    M, steps = len(mu_true), int(T / dt)
    a = np.empty((M, steps + 1))
    a[:, 0] = alpha0
    ell = np.full(M, np.log(alpha0 / (1 - alpha0)))
    dmu, half = mkt.mu_H - mkt.mu_L, (mkt.mu_H ** 2 - mkt.mu_L ** 2) / 2
    for k in range(steps):
        n = policy(a[:, k])
        dX = n * mu_true * dt + mkt.sigma * np.sqrt(n * dt) * rng.standard_normal(M)
        ell += (dmu * dX - n * dt * half) / mkt.sigma ** 2
        a[:, k + 1] = 1 / (1 + np.exp(-ell))
    return a
```

Before using it, we check *Proposition 37.3.1* by Monte Carlo.

```
def Sigma2(mkt, a):
    return (a * (1 - a) * (mkt.mu_H - mkt.mu_L) / mkt.sigma) ** 2


mkt = Market()
```

(continues on next page)

(continued from previous page)

```
rng = np.random.default_rng(0)
dt, n_draw = 1e-4, 400_000

print(f'{"alpha":>7s}{"simulated var/dt":>19s}{"formula n*Sigma^2":>20s}'
      f'{"mean/dt (s.e.)":>22s}')
for a0 in [0.2, 0.5, 0.8]:
    n = 0.5                                  # hold the experiment size fixed
    ell0 = np.log(a0 / (1 - a0))
    steps = []
    for mu_true, w in [(mkt.mu_H, a0), (mkt.mu_L, 1 - a0)]:
        k = int(n_draw * w)
        dX = n * mu_true * dt + mkt.sigma * np.sqrt(n * dt) * rng.standard_normal(k)
        ell = ell0 + ((mkt.mu_H - mkt.mu_L) * dX
                      - n * dt * (mkt.mu_H ** 2 - mkt.mu_L ** 2) / 2) / mkt.sigma ** 2
        steps.append(1 / (1 + np.exp(-ell)) - a0)
    d = np.concatenate(steps)
    se = d.std() / np.sqrt(len(d)) / dt
    print(f'{a0:7.2f}{d.var() / dt:19.6f}{n * Sigma2(mkt, a0):20.6f}'
          f'{d.mean() / dt:14.4f} ({se:.3f})')
```

```
  alpha   simulated var/dt   formula n*Sigma^2        mean/dt (s.e.)
   0.20           0.041586            0.041472        0.0034 (0.032)
   0.50           0.101145            0.101250        0.0606 (0.050)
   0.80           0.041428            0.041472        0.0096 (0.032)
```

The simulated variance matches (37.7), and the mean increment is indistinguishable from zero, confirming that beliefs form a martingale.

37.4 Efficient experimentation

A planner choosing $n(\alpha)$ trades current surplus against the information that sales generate.

Bergemann and Välimäki [1997] avoid the nonlinear differential equations that discounting would produce by working with the **undiscounted** limit, using the strong long-run average criterion of Dutta [1991].

The optimal policies in this limit are the limits of the discounted policies as the discount rate goes to zero, so the intertemporal tradeoff survives.

The Bellman equation becomes

$$\max_{n} \left\{ F(n,\alpha) - v(\alpha) + \tfrac{1}{2} n \Sigma^2(\alpha) V''(\alpha) \right\} = 0, \tag{37.10}$$

where $v(\alpha)$ is the long-run average attainable under full information and the last term is the **value of information**: the size of the experiment n times the speed of learning Σ^2 times the shadow price V''.

Because the belief is a martingale, no first-derivative term appears.

Every policy we consider keeps n bounded away from zero, so information accumulates without bound.

The prior puts positive probability on each of the two possible values of μ, so the truth is absolutely continuous with respect to the prior predictive distribution, and the posterior converges to the truth almost surely; see the discussion of Bayesian learning with a finite parameter space in *Merging of Opinions: The Blackwell–Dubins Theorem*.

Since μ is eventually learned, v is just the linear interpolation of the two full-information values,

$$v(\alpha) = \frac{s + \mu(\alpha) + \frac{3}{2}h}{2} + (1-\alpha)\frac{(\mu_L - s)^2}{4h} + \alpha\frac{(\mu_H - s)^2}{4h}. \tag{37.11}$$

The clever step is that the maximized bracket in (37.10) equals zero, so we may divide through by n without changing the maximizer.

Doing so removes V'' from the first-order condition entirely and leaves

$$\max_n \left\{ \frac{s + \frac{h}{2} - v(\alpha)}{n} - hn \right\} + \text{terms free of } n,$$

whose first-order condition gives the efficient policy in closed form.

Proposition 37.4.1 (Efficient experimentation)

The efficient market share of the new product is

$$n^*(\alpha) = \sqrt{\frac{v(\alpha) - s - \frac{h}{2}}{h}}. \tag{37.12}$$

The myopic planner, who ignores the informational value of sales, instead sets $m^*(\alpha) = \arg\max_n F(n, \alpha)$.

```
def v_social(mkt, a):
    s, h = mkt.s, mkt.h
    return ((s + mkt.mu(a) + 1.5 * h) / 2
            + (1 - a) * (mkt.mu_L - s) ** 2 / (4 * h)
            + a * (mkt.mu_H - s) ** 2 / (4 * h))


def n_star(mkt, a):
    """Efficient share, equation (nstar)."""
    return np.sqrt((v_social(mkt, a) - mkt.s - mkt.h / 2) / mkt.h)


def m_star(mkt, a):
    """Myopically efficient share."""
    return (mkt.mu(a) - mkt.s + mkt.h) / (2 * mkt.h)
```

At $\alpha \in \{0, 1\}$ there is nothing left to learn, so the two must agree, and they do.

```
for a, mu_i in [(0.0, mkt.mu_L), (1.0, mkt.mu_H)]:
    direct = (mu_i - mkt.s + mkt.h) / (2 * mkt.h)
    print(f'alpha = {a}:  n* = {n_star(mkt, a):.6f}   '
          f'full-information share = {direct:.6f}')

A = np.linspace(1e-6, 1 - 1e-6, 4001)          # full grid, for plotting
A_int = np.linspace(0.05, 0.95, 1801)          # strictly interior grid
```

```
alpha = 0.0:  n* = 0.050000   full-information share = 0.050000
alpha = 1.0:  n* = 0.950000   full-information share = 0.950000
```

The planner always experiments **more** than the myopic benchmark.

Let $m_i = (\mu_i - s + h)/(2h)$ be the efficient full-information share when $\mu = \mu_i$, and let E_α average over $i \in \{L, H\}$ with weights $1 - \alpha$ and α.

Substituting m_i into (37.11) gives $v(\alpha) - s - \frac{h}{2} = hE_\alpha[m_i^2]$, so

$$n^*(\alpha) = \sqrt{E_\alpha[m_i^2]} \geq E_\alpha[m_i] = m^*(\alpha)$$

by Jensen's inequality, with strict inequality for $\alpha \in (0,1)$ because $m_L \neq m_H$.

The gap is the intertemporal value of information showing up as extra sales of the new product.

```
def E_alpha(f_L, f_H, a):
    """Average of full-information quantities under belief a."""
    return (1 - a) * f_L + a * f_H


m_L, m_H = [(mu_i - mkt.s + mkt.h) / (2 * mkt.h) for mu_i in (mkt.mu_L, mkt.mu_H)]
print(f'max |n*(alpha)^2 - E[m_i^2]|  '
      f'{np.abs(n_star(mkt, A) ** 2 - E_alpha(m_L ** 2, m_H ** 2, A)).max():.1e}')
print(f'max |m*(alpha) - E[m_i]|      '
      f'{np.abs(m_star(mkt, A) - E_alpha(m_L, m_H, A)).max():.1e}')
print(f'n*(0.5) - m*(0.5) = {float(n_star(mkt, 0.5) - m_star(mkt, 0.5)):.5f}')
```

```
max |n*(alpha)^2 - E[m_i^2]|  1.4e-15
max |m*(alpha) - E[m_i]|      4.4e-16
n*(0.5) - m*(0.5) = 0.17268
```

37.5 Equilibrium

Now let the two firms set prices p_1 and p_2 and let buyers choose.

The marginal buyer n is indifferent when $s + nh - p_1 = \mu(\alpha) + (1-n)h - p_2$, which pins the market share to prices.

Each firm solves a dynamic program in which its own value of information appears, and the same divide-by-n trick removes the second derivatives from the first-order conditions.

Proposition 37.5.1 (Equilibrium)

There is a unique Markov-perfect equilibrium, with

$$p_1(\alpha) = \tfrac{2}{3}(s - \mu(\alpha)) + \sqrt{2h v_2(\alpha)}, \qquad p_2(\alpha) = \tfrac{1}{3}(\mu(\alpha) - s) + h, \tag{37.13}$$

and market share of the new firm

$$n(\alpha) = \sqrt{\frac{v_2(\alpha)}{2h}}, \tag{37.14}$$

where $v_i(\alpha)$ is firm i's full-information long-run average revenue.

```
def v1(mkt, a):
    s, h = mkt.s, mkt.h
    return ((1 - a) * ((s - mkt.mu_L) / 3 + h) ** 2 / (2 * h)
            + a * ((s - mkt.mu_H) / 3 + h) ** 2 / (2 * h))


def v2(mkt, a):
    s, h = mkt.s, mkt.h
    return ((1 - a) * ((mkt.mu_L - s) / 3 + h) ** 2 / (2 * h)
            + a * ((mkt.mu_H - s) / 3 + h) ** 2 / (2 * h))
```

(continues on next page)

(continued from previous page)

```
def n_eq(mkt, a):
    return np.sqrt(v2(mkt, a) / (2 * mkt.h))


def p1(mkt, a):
    return 2 / 3 * (mkt.s - mkt.mu(a)) + np.sqrt(2 * mkt.h * v2(mkt, a))


def p2(mkt, a):
    return (mkt.mu(a) - mkt.s) / 3 + mkt.h


def p1_myopic(mkt, a):
    return (mkt.s - mkt.mu(a)) / 3 + mkt.h


def n_myopic(mkt, a):
    return ((mkt.mu(a) - mkt.s) / 3 + mkt.h) / (2 * mkt.h)
```

Comparing the dynamic equilibrium with the static one played period by period reveals the asymmetry at the heart of the paper.

In the static game at belief α, the share of the new firm is $(\mu(\alpha) - s + h + p_1 - p_2)/(2h)$, and the first-order conditions of the two revenue-maximizing firms form a linear system in (p_1, p_2).

We solve that system directly and compare its solution with the dynamic equilibrium.

```
def static_equilibrium(mkt, a):
    """Prices and share in the one-shot Hotelling game at belief a."""
    d = mkt.mu(a) - mkt.s
    # firm 1:  -2 p1 + p2 = d - h;   firm 2:  p1 - 2 p2 = -(d + h)
    coef = np.array([[-2.0, 1.0], [1.0, -2.0]])
    p1_s, p2_s = np.linalg.solve(coef, np.stack([d - mkt.h, -(d + mkt.h)]))
    n_s = (d + mkt.h + p1_s - p2_s) / (2 * mkt.h)
    return p1_s, p2_s, n_s


p1_s, p2_s, n_s = static_equilibrium(mkt, A_int)
print('checking the static formulas, on [0.05, 0.95]')
print(f'  max |p1_myopic - p1_static|  {np.abs(p1_myopic(mkt, A_int) - p1_s).max():.
↪1e}')
print(f'  max |n_myopic - n_static|    {np.abs(n_myopic(mkt, A_int) - n_s).max():.1e}
↪')
print('comparing the dynamic equilibrium with the static one, on [0.05, 0.95]')
print(f'  max |p2 - p2_static|         {np.abs(p2(mkt, A_int) - p2_s).max():.1e}')
print(f'  min (p1 - p1_static)         {(p1(mkt, A_int) - p1_s).min():.5f}')
print(f'  min (n_eq - n_static)        {(n_eq(mkt, A_int) - n_s).min():.5f}')
```

```
checking the static formulas, on [0.05, 0.95]
  max |p1_myopic - p1_static|  2.2e-16
  max |n_myopic - n_static|    1.1e-16
comparing the dynamic equilibrium with the static one, on [0.05, 0.95]
  max |p2 - p2_static|         2.2e-16
  min (p1 - p1_static)         0.00671
  min (n_eq - n_static)        0.00336
```

The new firm's price is *exactly* its myopic price, a knife-edge consequence of the linear preference structure and the

absence of discounting.

The established firm charges *more* than it would in a one-shot game, and so concedes market share.

That is the striking result: the incumbent softens competition, not out of weakness, but because the entrant's sales are the only source of information and the incumbent wants the information.

37.5.1 Who values information more?

The Bellman equations imply that each firm's value of information equals the gap between its expected full-information revenue and its current revenue.

```
voi_1 = v1(mkt, A) - (1 - n_eq(mkt, A)) * p1(mkt, A)
voi_2 = v2(mkt, A) - n_eq(mkt, A) * p2(mkt, A)

print(f'established firm, minimum value of information  {voi_1.min():.3e}')
print(f'new firm, minimum value of information          {voi_2.min():.3e}')
print(f'ratio voi_1 / voi_2:  min {np.min(voi_1 / voi_2):.6f}, '
      f'max {np.max(voi_1 / voi_2):.6f}')
```

```
established firm, minimum value of information  1.800e-07
new firm, minimum value of information          9.000e-08
ratio voi_1 / voi_2:  min 2.000000, max 2.000000
```

Both are positive, so both value functions are convex in the belief.

That is the *Blackwell's Theorem on Comparing Experiments* logic at work: beliefs are a martingale, more experimentation spreads them further, and a firm with a convex value function gains from the spread.

More surprisingly, the ratio is exactly 2 at every belief: the **established** firm values information twice as much as the entrant.

Let n_i be the full-information equilibrium share when $\mu = \mu_i$ and let $\bar{n}(\alpha) = E_\alpha[n_i]$, which equals the static equilibrium share `n_myopic`.

Full-information prices are $p_{1,i} = 2h(1-n_i)$ and $p_{2,i} = 2hn_i$, so $v_2 = 2hE_\alpha[n_i^2] = 2hn^2$ and $v_1 = 2hE_\alpha[(1-n_i)^2] = 2h(1-2\bar{n}+n^2)$.

The entrant charges its static price $p_2 = 2h\bar{n}$, while the incumbent charges $p_1 = 2h(1-\bar{n}) + 2h(n-\bar{n})$, above its static price by $2h(n-\bar{n})$.

Substituting,

$$v_2 - np_2 = 2h\,n(n-\bar{n}), \qquad v_1 - (1-n)p_1 = 4h\,n(n-\bar{n}).$$

Jensen's inequality, $n = \sqrt{E_\alpha[n_i^2]} > E_\alpha[n_i] = \bar{n}$ for $\alpha \in (0,1)$, makes both gaps positive.

```
n_bar = n_myopic(mkt, A)
print(f'max |voi_2 - 2h n (n - n_bar)|  '
      f'{np.abs(voi_2 - 2 * mkt.h * n_eq(mkt, A) * (n_eq(mkt, A) - n_bar)).max():.1e}
↪')
print(f'max |voi_1 - 4h n (n - n_bar)|  '
      f'{np.abs(voi_1 - 4 * mkt.h * n_eq(mkt, A) * (n_eq(mkt, A) - n_bar)).max():.1e}
↪')
```

```
max |voi_2 - 2h n (n - n_bar)|  2.0e-16
max |voi_1 - 4h n (n - n_bar)|  8.6e-16
```

37.6 Too much experimentation, then too little

We can now compare the equilibrium share with the efficient one.

```
gap = n_star(mkt, A) - n_eq(mkt, A)
cross = A[np.argmin(np.abs(gap))]

fig, axes = plt.subplots(1, 2, figsize=(12, 4.2))
axes[0].plot(A, n_star(mkt, A), lw=2, label=r'efficient $n^*(\alpha)$')
axes[0].plot(A, n_eq(mkt, A), lw=2, label=r'equilibrium $n(\alpha)$')
axes[0].plot(A, m_star(mkt, A), lw=1.5, ls='--', color='0.5',
             label=r'myopic planner $m^*(\alpha)$')
axes[0].set(xlabel=r'$\alpha$', ylabel='market share of the new firm',
            title='experimentation policies')
axes[0].legend(fontsize=9)

axes[1].plot(A, gap, lw=2, color='C3')
axes[1].axhline(0, color='0.3', lw=1)
axes[1].axvline(cross, color='0.6', ls=':', lw=1.5)
axes[1].fill_between(A, gap, 0, where=gap < 0, alpha=0.15, color='C3')
axes[1].fill_between(A, gap, 0, where=gap > 0, alpha=0.15, color='C0')
axes[1].annotate('equilibrium\nexperiments too much', (0.05, gap.min() / 2),
                 fontsize=9)
axes[1].annotate('too little', (0.75, gap.max() / 2), fontsize=9)
axes[1].set(xlabel=r'$\alpha$', ylabel=r'$n^*(\alpha) - n(\alpha)$',
            title=f'single crossing at ' + rf'$\alpha = {cross:.3f}$')
fig.tight_layout()
plt.show()

n_L, n_H = n_eq(mkt, 0.0), n_eq(mkt, 1.0)
print(f'n*^2 - n^2 at alpha = 0:  m_L^2 - n_L^2 = {m_L ** 2 - n_L ** 2:+.4f}')
print(f'n*^2 - n^2 at alpha = 1:  m_H^2 - n_H^2 = {m_H ** 2 - n_H ** 2:+.4f}')
```

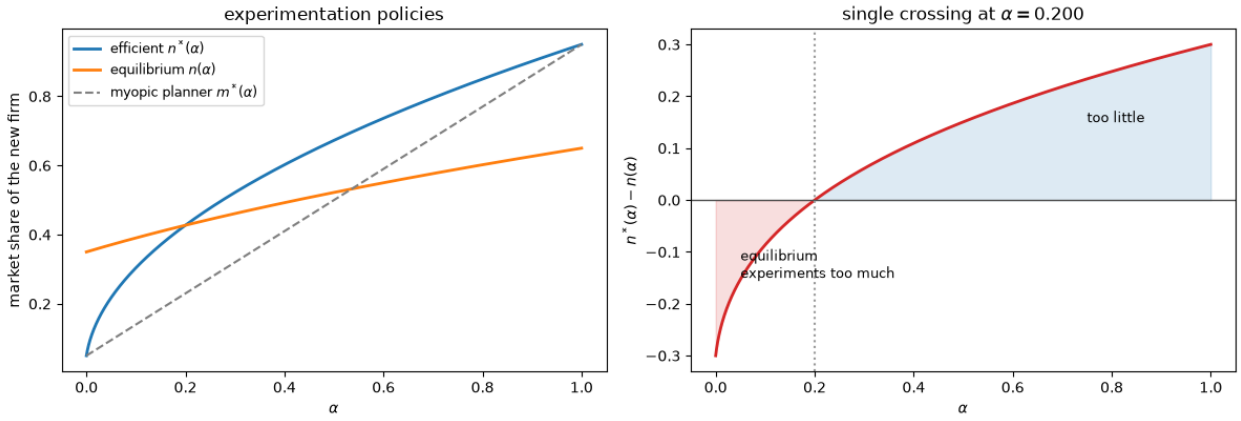


Fig. 37.1: Equilibrium versus efficient experimentation

```
n*^2 - n^2 at alpha = 0:  m_L^2 - n_L^2 = -0.1200
n*^2 - n^2 at alpha = 1:  m_H^2 - n_H^2 = +0.4800
```

The single crossing follows from the affine structure already used.

Both $n^*(\alpha)^2 = E_\alpha[m_i^2]$ and $n(\alpha)^2 = E_\alpha[n_i^2]$ are affine in α, so $n^*(\alpha)^2 - n(\alpha)^2$ is affine and changes sign at most once.

The gap $n^* - n$ has the same sign.

Condition (37.3) makes $m_L < n_L$ and $m_H > n_H$, so the gap is negative at $\alpha = 0$, positive at $\alpha = 1$, and crosses zero exactly once.

The intuition is about who has to cut price to gain a buyer.

At pessimistic beliefs the entrant is small, so attracting one more buyer costs it little in inframarginal revenue, while the incumbent is large and unwilling to defend its share by cutting price on everyone.

The entrant therefore expands aggressively and the market over-experiments.

At optimistic beliefs the positions are reversed, the incumbent fights harder, and experimentation falls short of the efficient level.

37.7 Diffusion over time

So far everything is a function of the state α.

To follow a product over calendar time we need the law of motion of the belief when the product really is good.

Conditional on $\mu = \mu_H$, the belief acquires an upward drift, since the data are generated by μ_H while the market still puts weight $1 - \alpha$ on μ_L,

$$d\alpha = \frac{n(\alpha)(\mu_H - \mu_L)^2 \alpha(1-\alpha)^2}{\sigma^2}\, dt + \frac{(\mu_H - \mu_L)\alpha(1-\alpha)\sqrt{n(\alpha)}}{\sigma}\, dB. \tag{37.15}$$

Equivalently, the log odds $\log(\alpha/(1-\alpha))$ drift up at rate $n(\alpha)(\mu_H - \mu_L)^2/(2\sigma^2)$.

The factor $(\mu_H - \mu_L)^2/(2\sigma^2)$ is the Kullback–Leibler divergence per unit of time between the signal distributions under μ_H and under μ_L (see *Kullback–Leibler divergence*), and $n(\alpha)$ scales it by the size of the experiment; Exercise 37.10.5 derives this drift and uses it to bound the expected time needed to learn.

Following Bergemann and Välimäki [1997], we drop the noise term.

The belief path $\hat{\alpha}(t)$ then solves the ordinary differential equation

$$\frac{d\hat{\alpha}}{dt} = \frac{n(\hat{\alpha})(\mu_H - \mu_L)^2 \hat{\alpha}(1-\hat{\alpha})^2}{\sigma^2}, \qquad \hat{\alpha}(0) = \alpha_0, \tag{37.16}$$

and $\hat{n}(t) = n(\hat{\alpha}(t))$ is the corresponding market share.

Bergemann and Välimäki [1997] call $\hat{\alpha}$ the mean posterior.

Because the drift in (37.15) is nonlinear in α, $\hat{\alpha}(t)$ is not the conditional mean $E[\alpha_t \mid \mu = \mu_H]$, and $\hat{n}(t)$ is not $E[n(\alpha_t) \mid \mu = \mu_H]$.

We call $\hat{\alpha}$ and $\hat{n}$ the noise-free paths.

```
def noise_free_belief_path(mkt, alpha0, T, dt, policy):
    """Euler solution of the noise-free belief equation when mu = mu_H."""
    steps = int(T / dt)
    a = np.empty(steps + 1)
    a[0] = alpha0
    dmu2 = (mkt.mu_H - mkt.mu_L) ** 2 / mkt.sigma ** 2
    for k in range(steps):
        drift = policy(a[k]) * dmu2 * a[k] * (1 - a[k]) ** 2
        a[k + 1] = min(max(a[k] + drift * dt, 1e-12), 1 - 1e-12)
    return a
```

Proposition 37.7.1 (S-shaped diffusion)

Let $\hat{\alpha}(t)$ solve (37.16) with $\alpha_0 \in (0,1)$ and let $\hat{n}(t) = n(\hat{\alpha}(t))$.

The noise-free market share $\hat{n}(t)$ is increasing over time.

Its rate of increase is itself increasing while $\hat{\alpha}(t) < 1/3$ and decreasing thereafter.

The proof uses the affine structure again.

Because $n(\alpha)^2 = n_L^2 + (n_H^2 - n_L^2)\alpha$, the product $n(\alpha)n'(\alpha)$ equals the constant $(n_H^2 - n_L^2)/2$, and

$$\frac{d\hat{n}}{dt} = n'(\hat{\alpha})\frac{d\hat{\alpha}}{dt} = \frac{(n_H^2 - n_L^2)(\mu_H - \mu_L)^2}{2\sigma^2}\,\hat{\alpha}(1-\hat{\alpha})^2.$$

Since $\hat{\alpha}$ increases with t and $\alpha(1-\alpha)^2$ increases in α exactly when $\alpha < 1/3$, the rate $d\hat{n}/dt$ rises until $\hat{\alpha} = 1/3$ and falls thereafter.

The composition of two forces produces the S.

Learning accelerates as beliefs move away from zero, which speeds up the growth of the entrant's share; but the equilibrium share $n(\alpha)$ is concave, so further belief improvements translate into ever smaller share gains.

The S-shape and the inflection at $\hat{\alpha} = 1/3$ are properties of the noise-free path.

Exercise 37.10.4 computes $E[\alpha_t \mid \mu = \mu_H]$ and $E[n_t \mid \mu = \mu_H]$ by Monte Carlo and compares them with $\hat{\alpha}$ and $\hat{n}$.

```
T, dt, alpha0 = 8.0, 1e-3, 0.03
policy = lambda a: n_eq(mkt, a)

a_hat = noise_free_belief_path(mkt, alpha0, T, dt, policy)
t_grid = np.linspace(0, T, len(a_hat))

rng = np.random.default_rng(12)
paths = simulate_beliefs(mkt, alpha0, T, dt, np.full(6, mkt.mu_H), rng, policy)

fig, axes = plt.subplots(1, 2, figsize=(12, 4.2))
for pth in paths:
    axes[0].plot(t_grid, pth, lw=0.7, alpha=0.55, color='C0')
axes[0].plot(t_grid, a_hat, lw=2.5, color='C3', label='noise-free path')
axes[0].axhline(1 / 3, color='0.5', ls=':', lw=1.5)
axes[0].set(xlabel='time', ylabel=r'$\alpha(t)$', title='beliefs')
axes[0].legend(fontsize=9)

for pth in paths:
    axes[1].plot(t_grid, n_eq(mkt, pth), lw=0.7, alpha=0.55, color='C0')
axes[1].plot(t_grid, n_eq(mkt, a_hat), lw=2.5, color='C3', label='noise-free path')
axes[1].set(xlabel='time', ylabel=r'$n(t)$',
            title='market share of the new firm')
axes[1].legend(fontsize=9)
fig.tight_layout()
plt.show()
```

The inflection point of the noise-free path is exactly where *Proposition 37.7.1* says it is.

```
n_hat = n_eq(mkt, a_hat)
growth = np.gradient(n_hat, t_grid)
```

(continues on next page)

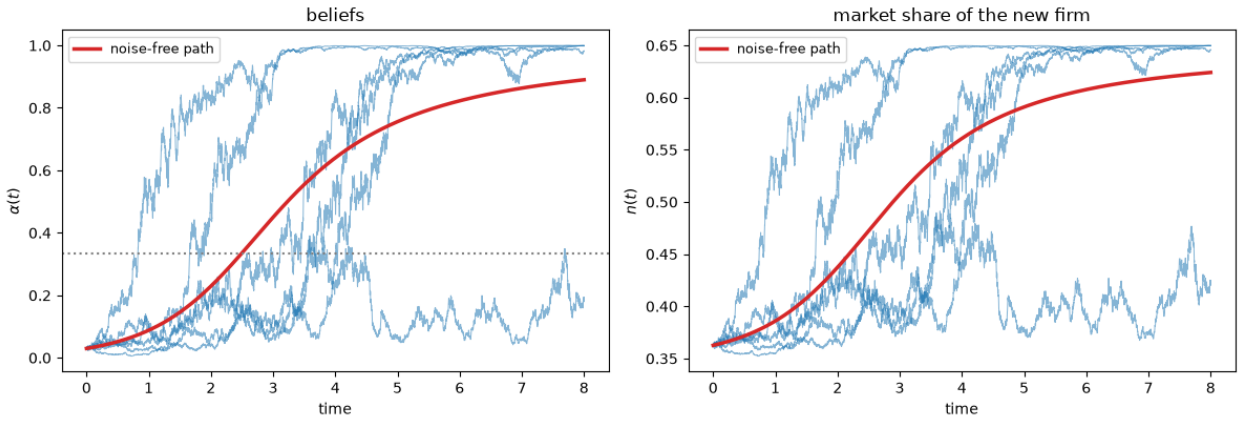


Fig. 37.2: The S-shaped diffusion path of a successful new product

(continued from previous page)

```
k = np.argmax(growth)
print(f'noise-free share grows fastest at t = {t_grid[k]:.3f}, '
      f'where alpha_hat = {a_hat[k]:.4f}   (theory: 1/3)')

drift = policy(A) * (mkt.mu_H - mkt.mu_L) ** 2 * A * (1 - A) ** 2
print(f'belief drift peaks at alpha = {A[np.argmax(drift)]:.4f}   '
      f'(theory: between 1/3 and 2/3)')
```

```
noise-free share grows fastest at t = 2.490, where alpha_hat = 0.3334   (theory: 1/
↪3)
belief drift peaks at alpha = 0.3830   (theory: between 1/3 and 2/3)
```

Prices move in step with shares.

```
fig, ax = plt.subplots()
for pth in paths:
    ax.plot(t_grid, p1(mkt, pth), lw=0.7, alpha=0.5, color='C0')
    ax.plot(t_grid, p2(mkt, pth), lw=0.7, alpha=0.5, color='C1')
ax.plot(t_grid, p1(mkt, a_hat), lw=2.5, color='C0',
        label=r'$p_1$, established firm (noise-free path)')
ax.plot(t_grid, p2(mkt, a_hat), lw=2.5, color='C1',
        label=r'$p_2$, new firm (noise-free path)')
ax.set(xlabel='time', ylabel='price',
       title='the incumbent retreats as the entrant is vindicated')
ax.legend()
fig.tight_layout()
plt.show()
```

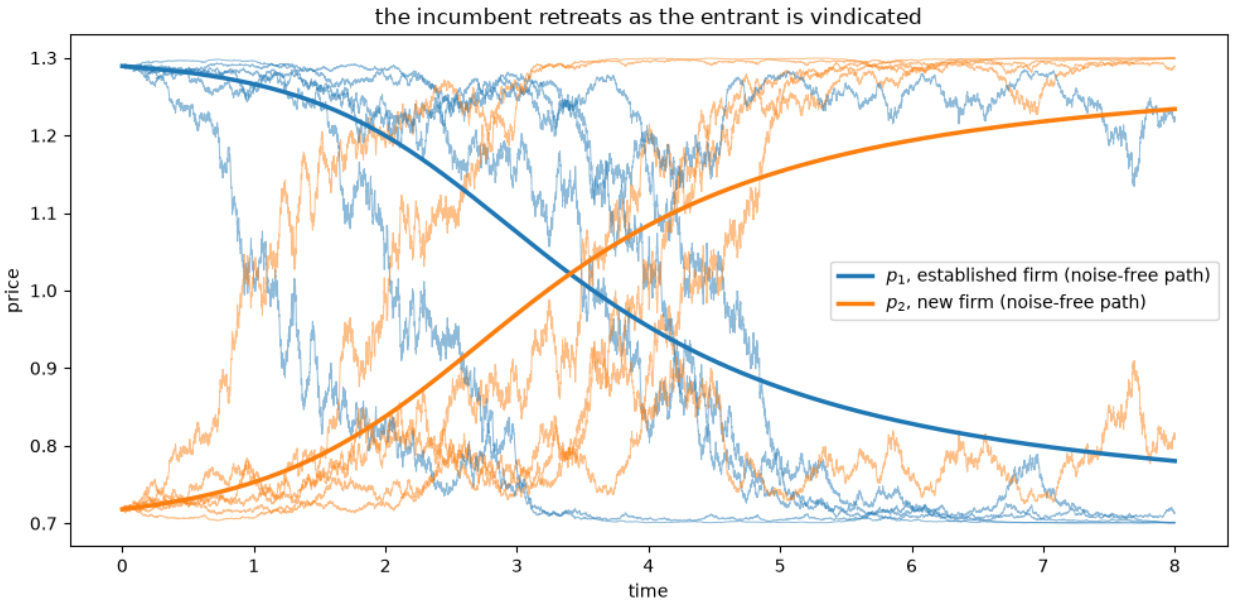


Fig. 37.3: Price paths of the two firms

37.8 Martingale properties

Bergemann and Välimäki [1997] characterize the equilibrium objects probabilistically: the entrant's price is a martingale, the incumbent's price and the entrant's share are supermartingales, and both revenues are submartingales.

Because α is a martingale, each of these follows from the shape of the corresponding function of α, and we can check them all by simulation.

```
rng = np.random.default_rng(3)
a0, M = 0.5, 20_000

# draw the true quality from the prior, one value per path
mu_true = np.where(rng.random(M) < a0, mkt.mu_H, mkt.mu_L)
ends = simulate_beliefs(mkt, a0, 4.0, 2e-3, mu_true, rng, policy)[:, -1]

rows = [
    ('belief', a0, ends.mean(), 'martingale'),
    ('share of new firm', n_eq(mkt, a0), n_eq(mkt, ends).mean(), 'supermartingale'),
    ('price of new firm', p2(mkt, a0), p2(mkt, ends).mean(), 'martingale'),
    ('price of incumbent', p1(mkt, a0), p1(mkt, ends).mean(), 'supermartingale'),
    ('revenue of incumbent', (1 - n_eq(mkt, a0)) * p1(mkt, a0),
     ((1 - n_eq(mkt, ends)) * p1(mkt, ends)).mean(), 'submartingale'),
    ('revenue of new firm', n_eq(mkt, a0) * p2(mkt, a0),
     (n_eq(mkt, ends) * p2(mkt, ends)).mean(), 'submartingale')]

print(f'{"":24s}{"t = 0":>10s}{"E[t = 4]":>11s}{"change":>10s}   prediction')
for name, x0, xT, pred in rows:
    print(f'{name:24s}{x0:10.4f}{xT:11.4f}{xT - x0:+10.4f}   {pred}')
```

```
                         t = 0   E[t = 4]    change   prediction
belief                  0.5000     0.5010   +0.0010   martingale
share of new firm       0.5220     0.5072   -0.0148   supermartingale
```

(continues on next page)

(continued from previous page)

```
price of new firm          1.0000     1.0006   +0.0006   martingale
price of incumbent         1.0440     1.0132   -0.0308   supermartingale
revenue of incumbent       0.4990     0.5310   +0.0320   submartingale
revenue of new firm        0.5220     0.5389   +0.0168   submartingale
```

Every sign comes out as predicted.

The entrant's expected share *falls* over time even though its share rises conditional on the product being good, because the early aggression reflects the value of information rather than confidence in the product.

Both firms expect to earn more later, which is the sense in which they sacrifice current profit to buy information.

37.9 Concluding remarks

Two lectures in this section now feature information whose value is entirely instrumental.

In *The Design and Price of Information* a seller designs and prices experiments, and the interesting economics comes from the fact that Blackwell's order is incomplete.

Here nobody prices information at all, and the interesting economics comes from the fact that only one firm's sales produce it.

Both rest on the same foundation from *Blackwell's Theorem on Comparing Experiments*: more informative experiments spread posterior beliefs in the convex order, and a decision maker whose value is convex in the belief gains from the spread.

A single decision maker's value is always convex in the belief, but an equilibrium value in a game need not be.

The distinctive lesson of Bergemann and Välimäki [1997] is that this convexity is shared by *competitors*.

Because both firms would rather face a market that has sorted out the quality of the new product, uncertainty about vertical quality relaxes price competition much as deterministic differentiation does in Shaked and Sutton [1982].

That is why the incumbent lets the entrant in cheaply at first and why a successful product diffuses along an S-shaped path.

The continuous-time technique used here, and in particular the device of taking the undiscounted limit to keep the Bellman equations tractable, comes from Bolton and Harris [1999], who were the first to study strategic experimentation in continuous time.

A companion paper, Bergemann and Välimäki [2000], studies the same duopoly with a continuum of *identical* consumers.

Homogeneity there rules out market sharing, so the horizontal differentiation that generates the diffusion path in this lecture is absent and the analysis concentrates instead on how informational externalities affect market efficiency.

37.10 Exercises

Exercise 37.10.1

Condition (37.3) requires $|\mu_i - s| < h$ for both quality levels.

1. Show algebraically that this confines the full-information equilibrium share of the new firm to the interval $(1/3, 2/3)$, and hence that $n(\alpha) \in (1/3, 2/3)$ for every belief.
2. Verify this numerically for several admissible (μ_L, μ_H) pairs.

3. Bergemann and Välimäki [1997] draw their diffusion figures with $s = 4$, $h = 1$, $\mu_L = 2$ and $\mu_H = 6$.

   Check whether these satisfy (37.3), compute the equilibrium shares at $\alpha \in \{0, 1\}$, and compute the myopically efficient share m^* at each quality level.

   What goes wrong, and which of the lecture's results still hold?

Solution

Here is one solution:

Under full information with quality μ_i the equilibrium share is $n_i = (\frac{1}{3}(\mu_i - s) + h)/(2h)$.

Condition (37.3) gives $-h < \mu_i - s < h$, so $\frac{1}{3}(\mu_i - s) \in (-h/3, h/3)$ and therefore $n_i \in (\frac{2h/3}{2h}, \frac{4h/3}{2h}) = (1/3, 2/3)$.

Since $n(\alpha)^2$ is a convex combination of n_0^2 and n_1^2, the equilibrium share lies between n_0 and n_1 for every α.

```
for mu_L, mu_H in [(3.1, 4.9), (3.4, 4.6), (3.9, 4.1)]:
    m_ = Market(mu_L=mu_L, mu_H=mu_H)
    lo, hi = n_eq(m_, 0.0), n_eq(m_, 1.0)
    print(f'(mu_L, mu_H) = ({mu_L}, {mu_H}):  n_eq ranges over '
          f'[{lo:.4f}, {hi:.4f}]   inside (1/3, 2/3): {1/3 < lo and hi < 2/3}')
```

```
(mu_L, mu_H) = (3.1, 4.9):  n_eq ranges over [0.3500, 0.6500]   inside (1/3, 2/
↪3): True
(mu_L, mu_H) = (3.4, 4.6):  n_eq ranges over [0.4000, 0.6000]   inside (1/3, 2/
↪3): True
(mu_L, mu_H) = (3.9, 4.1):  n_eq ranges over [0.4833, 0.5167]   inside (1/3, 2/
↪3): True
```

```
class LooseMarket(Market):
    def __init__(self, **kw):                # skip the assertion
        self.s, self.h = kw['s'], kw['h']
        self.mu_L, self.mu_H, self.sigma = kw['mu_L'], kw['mu_H'], kw.get('sigma',␣
↪1.0)


paper = LooseMarket(s=4, h=1, mu_L=2, mu_H=6)
print(f'condition (4) needs  s - h < mu_L:  {paper.s - paper.h} < {paper.mu_L}?  '
      f'{paper.s - paper.h < paper.mu_L}')
print(f'condition (4) needs  mu_H < s + h:  {paper.mu_H} < {paper.s + paper.h}?  '
      f'{paper.mu_H < paper.s + paper.h}')
print(f'\nequilibrium shares:  n(0) = {n_eq(paper, 0.0):.4f}, '
      f'n(1) = {n_eq(paper, 1.0):.4f}')
for mu_i, nm in [(paper.mu_L, 'mu_L'), (paper.mu_H, 'mu_H')]:
    print(f'myopically efficient share at {nm}: '
          f'{(mu_i - paper.s + paper.h) / (2 * paper.h):+.4f}')
```

```
condition (4) needs  s - h < mu_L:  3 < 2?  False
condition (4) needs  mu_H < s + h:  6 < 5?  False

equilibrium shares:  n(0) = 0.1667, n(1) = 0.8333
myopically efficient share at mu_L: -0.5000
myopically efficient share at mu_H: +1.5000
```

The paper's figure parameters violate (37.3) at both ends.

The consequence is that the *efficient* allocation is at a corner: it would assign every buyer to the established product

when $\mu = \mu_L$ and every buyer to the new product when $\mu = \mu_H$, so the interior formula (37.12) no longer applies and the efficiency comparison of *Proposition 37.4.1* breaks down.

Everything about the *equilibrium* survives, because equilibrium shares remain strictly interior at $1/6$ and $5/6$.

That is why those parameters are fine for drawing diffusion paths, which is all the paper uses them for, and why they buy a much more dramatic S-curve than any admissible parameter set could.

Exercise 37.10.2

The lecture found a single belief at which equilibrium experimentation switches from excessive to insufficient.

1. Write a function that locates this crossing point by bisection.
2. Compute it as the quality spread $\mu_H - \mu_L$ widens, holding the midpoint $\frac{1}{2}(\mu_L + \mu_H) = s$ fixed, and again as the horizontal differentiation parameter h varies.
3. Both experiments produce the same numbers whenever the ratio $(\mu_H - \mu_L)/h$ agrees.

   Guess the closed form for the crossing point and check it numerically.
4. Does your formula survive when the quality midpoint is moved away from s?

Solution

Here is one solution:

```
def crossing(mkt, tol=1e-13):
    """Belief at which n*(alpha) = n(alpha), by bisection."""
    lo, hi = 1e-12, 1 - 1e-12
    f = lambda a: n_star(mkt, a) - n_eq(mkt, a)
    if f(lo) > 0 or f(hi) < 0:
        return np.nan
    while hi - lo > tol:
        mid = (lo + hi) / 2
        lo, hi = (mid, hi) if f(mid) < 0 else (lo, mid)
    return (lo + hi) / 2


print('widening the quality spread, midpoint fixed at s = 4, h = 1')
for spread in [0.4, 0.8, 1.2, 1.6, 1.9]:
    m_ = Market(s=4, h=1, mu_L=4 - spread / 2, mu_H=4 + spread / 2)
    print(f'   (mu_H - mu_L)/h = {spread / 1:.3f}:  crossing = {crossing(m_):.6f}')

print('\nvarying horizontal differentiation, mu = (3.4, 4.6)')
for h_ in [0.65, 0.8, 1.0, 1.5, 2.5]:
    m_ = Market(s=4, h=h_, mu_L=3.4, mu_H=4.6)
    print(f'   (mu_H - mu_L)/h = {1.2 / h_:.3f}:  crossing = {crossing(m_):.6f}')
```

```
widening the quality spread, midpoint fixed at s = 4, h = 1
   (mu_H - mu_L)/h = 0.400:  crossing = 0.433333
   (mu_H - mu_L)/h = 0.800:  crossing = 0.366667
   (mu_H - mu_L)/h = 1.200:  crossing = 0.300000
   (mu_H - mu_L)/h = 1.600:  crossing = 0.233333
   (mu_H - mu_L)/h = 1.900:  crossing = 0.183333

varying horizontal differentiation, mu = (3.4, 4.6)
   (mu_H - mu_L)/h = 1.846:  crossing = 0.192308
   (mu_H - mu_L)/h = 1.500:  crossing = 0.250000
   (mu_H - mu_L)/h = 1.200:  crossing = 0.300000
   (mu_H - mu_L)/h = 0.800:  crossing = 0.366667
   (mu_H - mu_L)/h = 0.480:  crossing = 0.420000
```

Sorted by the ratio $(\mu_H - \mu_L)/h$ the two tables line up, which suggests that the crossing point depends on the parameters only through that ratio.

The numbers fall on a straight line with slope $-1/6$ through $1/2$.

```
print(f'{"(mu_H-mu_L)/h":>15s}{"bisection":>12s}{"1/2 - ratio/6":>16s}{"error":>
 ↪12s}')
for mu_L_, mu_H_, h_ in [(3.4, 4.6, 1.0), (3.1, 4.9, 1.0), (3.8, 4.2, 1.0),
                         (3.4, 4.6, 1.5), (3.4, 4.6, 0.8), (3.05, 4.95, 1.0)]:
    m_ = Market(s=4, h=h_, mu_L=mu_L_, mu_H=mu_H_)
    r = (mu_H_ - mu_L_) / h_
    c, pred = crossing(m_), 0.5 - r / 6
    print(f'{r:15.4f}{c:12.6f}{pred:16.6f}{c - pred:12.1e}')
```

```
  (mu_H-mu_L)/h   bisection   1/2 - ratio/6       error
         1.2000    0.300000        0.300000    -1.5e-14
         1.8000    0.200000        0.200000    -8.2e-15
         0.4000    0.433333        0.433333    -4.9e-15
         0.8000    0.366667        0.366667     1.9e-14
         1.5000    0.250000        0.250000     1.7e-14
         1.9000    0.183333        0.183333    -1.6e-14
```

So when the two quality levels straddle s symmetrically, the switch occurs at

$$\alpha^{\text{cross}} = \frac{1}{2} - \frac{\mu_H - \mu_L}{6h},$$

which condition (37.3) keeps strictly inside $(1/6, 1/2)$, since that condition forces $\mu_H - \mu_L < 2h$.

The region of excessive experimentation therefore *shrinks* as the quality spread widens relative to h.

A wider spread means more is at stake in learning, and the efficient policy responds by experimenting a great deal; the equilibrium, driven by each firm's private revenue motive, does not keep up except at the most pessimistic beliefs.

Raising h works in the opposite direction, since strongly attached buyers blunt the price instrument and let the entrant expand more freely than a planner would choose.

The symmetric formula does not survive a move of the midpoint, but a general closed form does.

Let $d_i = (\mu_i - s)/h$.

The efficient and equilibrium full-information shares are $m_i = (1 + d_i)/2$ and $n_i = (1 + d_i/3)/2$, so $m_i^2 - n_i^2 = \varphi(d_i)/9$ with $\varphi(d) = d(3 + 2d)$.

Because $n^*(\alpha)^2 - n(\alpha)^2 = E_\alpha[m_i^2 - n_i^2]$ is affine in α, it vanishes at

$$\alpha^{\text{cross}} = \frac{\varphi(d_L)}{\varphi(d_L) - \varphi(d_H)}.$$

Condition (37.3) gives $d_L \in (-1, 0)$ and $d_H \in (0, 1)$, so $\varphi(d_L) < 0 < \varphi(d_H)$ and $\alpha^{\text{cross}} \in (0, 1)$.

When $d_H = -d_L = (\mu_H - \mu_L)/(2h)$, the formula reduces to $\frac{1}{2} - (\mu_H - \mu_L)/(6h)$.

```
def crossing_formula(mkt):
    φ = lambda d: d * (3 + 2 * d)
    d_L, d_H = (mkt.mu_L - mkt.s) / mkt.h, (mkt.mu_H - mkt.s) / mkt.h
    return φ(d_L) / (φ(d_L) - φ(d_H))


print('moving the quality midpoint away from s, with mu = (3.4, 4.6), h = 1')
for s_ in [3.9, 4.0, 4.1]:
    m_ = Market(s=s_, h=1, mu_L=3.4, mu_H=4.6)
    print(f'   s = {s_}:  bisection = {crossing(m_):.6f}'
          f'   general formula = {crossing_formula(m_):.6f}'
          f'   symmetric formula = {0.5 - 1.2 / 6:.6f}')
```

```
moving the quality midpoint away from s, with mu = (3.4, 4.6), h = 1
   s = 3.9:  bisection = 0.245098   general formula = 0.245098   symmetric␣
↪formula = 0.300000
   s = 4.0:  bisection = 0.300000   general formula = 0.300000   symmetric␣
↪formula = 0.300000
   s = 4.1:  bisection = 0.358974   general formula = 0.358974   symmetric␣
↪formula = 0.300000
```

The general formula matches the bisection at every s; the symmetric formula matches only at $s = 4$.

Exercise 37.10.3

This exercise makes the link with *Blackwell's Theorem on Comparing Experiments* precise.

In that lecture, a more informative experiment produces a mean-preserving spread of the posterior, and a decision maker whose value is convex in the belief gains from it.

Here the belief is a martingale and experimentation controls the speed at which it spreads, so the same logic applies to each firm, provided its equilibrium value function is convex.

1. Plot each firm's value of information, $v_i(\alpha)$ minus its current equilibrium revenue, against α.
2. Confirm that both are positive everywhere in the interior and vanish at $\alpha \in \{0, 1\}$, and explain why they must vanish there.
3. The value of information also equals $\frac{1}{2} n(\alpha) \Sigma^2(\alpha) V_i''(\alpha)$.

   Use this to recover $V_i''(\alpha)$ and confirm that both value functions are convex.

Solution

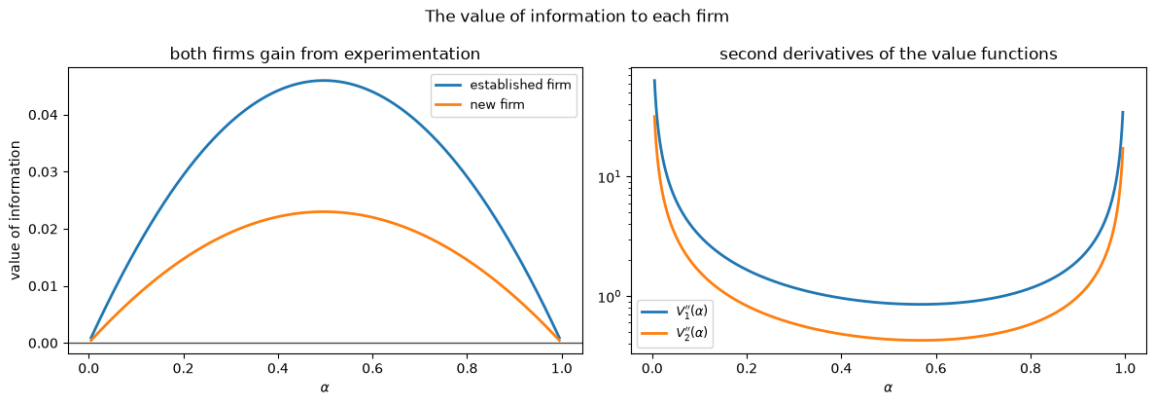
Here is one solution:

```
Ai = np.linspace(0.005, 0.995, 2001)
voi_1 = v1(mkt, Ai) - (1 - n_eq(mkt, Ai)) * p1(mkt, Ai)
voi_2 = v2(mkt, Ai) - n_eq(mkt, Ai) * p2(mkt, Ai)

fig, axes = plt.subplots(1, 2, figsize=(12, 4.2))
axes[0].plot(Ai, voi_1, lw=2, label='established firm')
axes[0].plot(Ai, voi_2, lw=2, label='new firm')
axes[0].axhline(0, color='0.3', lw=1)
axes[0].set(xlabel=r'$\alpha$', ylabel='value of information',
            title='both firms gain from experimentation')
axes[0].legend(fontsize=9)

V1pp = 2 * voi_1 / (n_eq(mkt, Ai) * Sigma2(mkt, Ai))
V2pp = 2 * voi_2 / (n_eq(mkt, Ai) * Sigma2(mkt, Ai))
axes[1].plot(Ai, V1pp, lw=2, label=r"$V_1''(\alpha)$")
axes[1].plot(Ai, V2pp, lw=2, label=r"$V_2''(\alpha)$")
axes[1].set(xlabel=r'$\alpha$', yscale='log',
            title='second derivatives of the value functions')
axes[1].legend(fontsize=9)
fig.suptitle('The value of information to each firm')
fig.tight_layout()
plt.show()
```

```
print(f'minimum value of information, established firm  {voi_1.min():.3e}')
print(f'minimum value of information, new firm          {voi_2.min():.3e}')
print(f'minimum of V1\'\'  {V1pp.min():.4f}    minimum of V2\'\'  {V2pp.min():.4f}
↪')
```




```
minimum value of information, established firm  8.957e-04
minimum value of information, new firm          4.479e-04
minimum of V1''  0.8531    minimum of V2''  0.4265
```

Both curves are strictly positive on the interior and both second derivatives are strictly positive, so both value functions are convex.

The value of information vanishes at $\alpha \in \{0, 1\}$ for two reinforcing reasons.

There is nothing left to learn, so the equilibrium coincides with the full-information equilibrium and the revenue gap closes.

And the speed of learning $\Sigma^2(\alpha) \propto \alpha^2(1-\alpha)^2$ vanishes as well, so even a convex value function earns nothing from an experiment that reveals nothing.

The second derivatives do not vanish at the endpoints; they diverge.

Near $\alpha = 0$ the value of information of the new firm is approximately $h(n_H - n_L)^2\alpha$, while $\Sigma^2(\alpha) \approx (\mu_H - \mu_L)^2\alpha^2/\sigma^2$.

Hence $V_2''(\alpha) \approx 2h(n_H - n_L)^2\sigma^2/(n_L(\mu_H - \mu_L)^2\alpha)$, so $\alpha V_2''(\alpha)$ converges to a positive constant, and the same holds for V_1'' and at $\alpha = 1$ with $1 - \alpha$ in place of α.

The shadow value V_i'' of spreading beliefs grows like $1/\alpha$, but the speed Σ^2 at which experimentation spreads them falls like α^2.

Their product, the value of information, therefore vanishes linearly.

```
n_L, n_H = n_eq(mkt, 0.0), n_eq(mkt, 1.0)
limit = 2 * mkt.h * (n_H - n_L) ** 2 * mkt.sigma ** 2 / (
    n_L * (mkt.mu_H - mkt.mu_L) ** 2)
for a in [1e-2, 1e-3, 1e-4]:
    voi_2a = v2(mkt, a) - n_eq(mkt, a) * p2(mkt, a)
    V2pp_a = 2 * voi_2a / (n_eq(mkt, a) * Sigma2(mkt, a))
    print(f'alpha = {a:.0e}:  alpha * V2\'\' = {a * V2pp_a:.4f}')
print(f'limit:            {limit:.4f}')
```

```
alpha = 1e-02:  alpha * V2'' = 0.1587
alpha = 1e-03:  alpha * V2'' = 0.1587
alpha = 1e-04:  alpha * V2'' = 0.1587
limit:            0.1587
```

Exercise 37.10.4

Proposition 37.7.1 describes the noise-free path $\hat{\alpha}(t)$, not the conditional mean of the belief.

1. Use `simulate_beliefs` to estimate $E[\alpha_t \mid \mu = \mu_H]$ and $E[n(\alpha_t) \mid \mu = \mu_H]$ for $t \in [0, 8]$, starting from $\alpha_0 = 0.03$, and plot them against $\hat{\alpha}(t)$ and $\hat{n}(t)$.
2. Find the time at which $E[n(\alpha_t) \mid \mu = \mu_H]$ grows fastest and the value of $E[\alpha_t \mid \mu = \mu_H]$ at that time.
3. Explain why the noise-free path runs ahead of the conditional mean.

Hint: Differentiating a Monte Carlo average with respect to time amplifies sampling noise.

Instead use $\frac{d}{dt}E[n(\alpha_t)] = E[\mathcal{G}n(\alpha_t)]$, where $\mathcal{G}n = n' \times \text{drift} + \frac{1}{2}n'' \times \text{variance}$ is the generator of the belief process (37.15) applied to n.

Solution

Here is one solution.

Write $n(\alpha)^2 = n_L^2 + b\alpha$ with $b = n_H^2 - n_L^2$, and let $k = (\mu_H - \mu_L)^2/\sigma^2$.

Then $n' = b/(2n)$ and $n'' = -b^2/(4n^3)$, and the coefficients of (37.15) give

$$\mathcal{G}n(\alpha) = \frac{bk}{2}\,\alpha(1-\alpha)^2\left[1 - \frac{b\alpha}{4n(\alpha)^2}\right].$$

The first term is the rate $d\hat{n}/dt$ along the noise-free path; the second is an Itô correction that comes from the concavity of n.

We simulate 10,000 paths in five batches to save memory.

```
alpha0, T4, dt4 = 0.03, 8.0, 5e-3
n_L, n_H = n_eq(mkt, 0.0), n_eq(mkt, 1.0)
b = n_H ** 2 - n_L ** 2
k2 = (mkt.mu_H - mkt.mu_L) ** 2 / mkt.sigma ** 2


def generator_n(a):
    """Generator of the belief process under mu_H applied to n(alpha)."""
    n = n_eq(mkt, a)
    return 0.5 * b * k2 * a * (1 - a) ** 2 * (1 - b * a / (4 * n ** 2))


rng = np.random.default_rng(2024)
M_batch, n_batches = 2_000, 5
sum_a = sum_n = sum_g = 0.0
for _ in range(n_batches):
    pth = simulate_beliefs(mkt, alpha0, T4, dt4, np.full(M_batch, mkt.mu_H),
                           rng, policy)
    sum_a = sum_a + pth.sum(axis=0)
    sum_n = sum_n + n_eq(mkt, pth).sum(axis=0)
    sum_g = sum_g + generator_n(pth).sum(axis=0)
M_total = M_batch * n_batches
E_a, E_n, E_g = sum_a / M_total, sum_n / M_total, sum_g / M_total

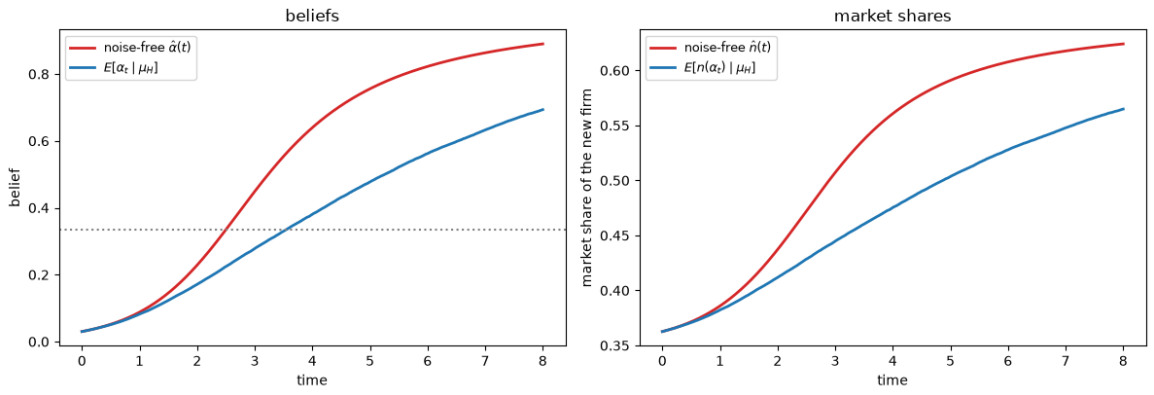
t4 = np.linspace(0, T4, len(E_a))
a_hat4 = noise_free_belief_path(mkt, alpha0, T4, dt4, policy)
```

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4.2))
axes[0].plot(t4, a_hat4, lw=2, color='C3', label=r'noise-free $\hat\alpha(t)$')
axes[0].plot(t4, E_a, lw=2, color='C0', label=r'$E[\alpha_t \mid \mu_H]$')
axes[0].axhline(1 / 3, color='0.5', ls=':', lw=1.5)
axes[0].set(xlabel='time', ylabel='belief', title='beliefs')
axes[0].legend(fontsize=9)
axes[1].plot(t4, n_eq(mkt, a_hat4), lw=2, color='C3',
             label=r'noise-free $\hat n(t)$')
axes[1].plot(t4, E_n, lw=2, color='C0', label=r'$E[n(\alpha_t) \mid \mu_H]$')
axes[1].set(xlabel='time', ylabel='market share of the new firm',
            title='market shares')
axes[1].legend(fontsize=9)
fig.tight_layout()
plt.show()

print(f'{"t":>4}{"alpha_hat":>11}{"E[alpha]":>10}{"n_hat":>8}{"E[n]":>8}')
for t in [2, 4, 6, 8]:
    j = int(round(t / dt4))
    print(f'{t:>4}{a_hat4[j]:11.3f}{E_a[j]:10.3f}'
          f'{n_eq(mkt, a_hat4[j]):8.3f}{E_n[j]:8.3f}')

j = np.argmax(E_g)
near = t4[E_g > 0.99 * E_g.max()]
print(f'\nE[n_t] grows fastest at t = {t4[j]:.2f}, where E[alpha_t] = {E_a[j]:.3f}
 ↪')
print(f'its growth rate is within 1% of the maximum for t in '
      f'[{near.min():.2f}, {near.max():.2f}]')

A_c = np.linspace(1e-3, 1 - 1e-3, 9_999)
drift_c = n_eq(mkt, A_c) * A_c * (1 - A_c) ** 2
curv = np.gradient(np.gradient(drift_c, A_c), A_c)
print(f'the belief drift is concave for alpha below {A_c[np.argmax(curv > 0)]:.2f}
 ↪')
```




```
   t  alpha_hat  E[alpha]   n_hat    E[n]
   2      0.229     0.171   0.437   0.412
   4      0.640     0.381   0.561   0.475
   6      0.823     0.562   0.608   0.528
   8      0.890     0.693   0.624   0.565

E[n_t] grows fastest at t = 2.41, where E[alpha_t] = 0.214
its growth rate is within 1% of the maximum for t in [2.02, 2.78]
the belief drift is concave for alpha below 0.72
```

The noise-free path runs well ahead of the conditional mean: at $t = 4$ it puts the belief at 0.64 and the share at 0.561, while the conditional means are about 0.38 and 0.48.

The conditional mean share is also S-shaped.

It grows fastest at about the same date as the noise-free share, near $t = 2.4$, but the conditional mean belief there is about 0.21, not 1/3.

The growth rate of $E[n_t \mid \mu_H]$ is within one percent of its maximum between $t \approx 2.0$ and $t \approx 2.8$, so the date of fastest growth is imprecisely located.

Two forces hold the mean back.

The drift in (37.15) is concave in α for α below about 0.72, so while most paths lie in that region the average drift across paths falls short of the drift at the average belief.

Paths that receive bad news early sink toward $\alpha = 0$, where both the drift and the volatility vanish and learning nearly stops; those paths hold the average down.

Exercise 37.10.5

Let $\ell_t = \log(\alpha_t/(1-\alpha_t))$ and $\kappa = (\mu_H - \mu_L)^2/(2\sigma^2)$.

1. Use (37.6) and (37.8) to show that, conditional on $\mu = \mu_H$,

$$d\ell = \kappa\, n(\alpha)\, dt + \frac{\mu_H - \mu_L}{\sigma}\sqrt{n(\alpha)}\, dB,$$

   and that conditional on $\mu = \mu_L$ the drift is $-\kappa\, n(\alpha)$. Interpret κ as a Kullback–Leibler divergence rate.

2. When $\mu = \mu_H$, let τ be the first time the belief reaches 0.9, starting from $\alpha_0 = 0.03$. Use part 1, the optional stopping theorem, and $n(\alpha) \in (1/3, 2/3)$ to bound $E[\tau]$, and check the bounds by simulation.

3. Repeat conditional on $\mu = \mu_L$. Explain why the expected time for the belief to rise from 0.03 to 0.9 is now infinite, and bound instead the expected time for it to fall from 0.97 to 0.1. Explain why the two simulated expected times differ although the bounds are the same.

Solution

Here is one solution.

Part 1. In continuous time (37.8) becomes $d\ell = [(\mu_H - \mu_L)\, dX - \frac{1}{2} n(\mu_H^2 - \mu_L^2)\, dt]/\sigma^2$.

Conditional on $\mu = \mu_H$, $dX = n\mu_H\, dt + \sigma\sqrt{n}\, dB$, so the drift of ℓ is

$$\frac{n}{\sigma^2}\left[(\mu_H - \mu_L)\mu_H - \tfrac{1}{2}(\mu_H - \mu_L)(\mu_H + \mu_L)\right] = \frac{n(\mu_H - \mu_L)^2}{2\sigma^2} = \kappa n$$

and its diffusion coefficient is $(\mu_H - \mu_L)\sqrt{n}/\sigma$.

Conditional on $\mu = \mu_L$, replacing μ_H by μ_L in the first term of the bracket gives drift $-\kappa n$.

Over an interval of length Δ with $n = 1$, the increment ΔX is $N(\mu_H\Delta, \sigma^2\Delta)$ under μ_H and $N(\mu_L\Delta, \sigma^2\Delta)$ under μ_L.

The Kullback–Leibler divergence between these two distributions, in either direction, is $(\mu_H - \mu_L)^2\Delta/(2\sigma^2) = \kappa\Delta$ (see *Kullback–Leibler divergence*).

So the log odds drift toward the truth at n times the KL divergence rate, the continuous-time version of the link between expected log likelihood ratios and KL divergence in *Likelihood Ratio Processes*.

Part 2. The process $\ell_t - \ell_0 - \kappa \int_0^t n(\alpha_s)\, ds$ is a martingale.

Because $n \geq n_L > 0$, the drift of ℓ is bounded below by a positive constant and $E[\tau] < \infty$, so optional stopping applies.

Because ℓ has continuous paths, $\ell_\tau = \ell^* = \log 9$.

Hence

$$\kappa\, E\left[\int_0^\tau n(\alpha_s)\, ds\right] = \ell^* - \ell_0,$$

and $\frac{1}{3}E[\tau] < E\left[\int_0^\tau n(\alpha_s)\, ds\right] < \frac{2}{3}E[\tau]$ gives

$$\frac{3}{2}\,\frac{\ell^* - \ell_0}{\kappa} < E[\tau] < 3\,\frac{\ell^* - \ell_0}{\kappa}.$$

With the lecture's parameters, n actually stays in $[n_L, n_H] = [0.35, 0.65]$, which tightens the bounds to $(\ell^* - \ell_0)/(\kappa n_H) \leq E[\tau] \leq (\ell^* - \ell_0)/(\kappa n_L)$.

Part 3. Conditional on $\mu = \mu_L$, Itô's lemma and part 1 show that the odds $e^{\ell_t} = \alpha_t/(1-\alpha_t)$ form a martingale: the drift of ℓ, $-\kappa n$, is minus one half of its instantaneous variance, $2\kappa n$.

The odds are the prior odds times a likelihood ratio process of the kind studied in *Likelihood Ratio Processes*.

Stopped when the belief first reaches 0.9, the odds are bounded by 9, and $\ell_t \to -\infty$ on paths that never get there.

Optional stopping gives a probability $(0.03/0.97)/9 \approx 0.0034$ that the belief ever reaches 0.9 from 0.03, so with probability close to one it never does, and the expected hitting time is infinite.

The time for the belief to fall from 0.97 to 0.1 obeys the argument of part 2 applied to $-\ell$, with the same distance $|\ell^* - \ell_0| = \log 9 + \log(0.97/0.03)$, so the bounds are identical.

```
def first_passage_times(mkt, alpha0, alpha_target, mu_true, M, dt, T, rng):
    """First times at which the belief reaches alpha_target (nan if not by T)."""
    dmu, half = mkt.mu_H - mkt.mu_L, (mkt.mu_H ** 2 - mkt.mu_L ** 2) / 2
    ell = np.full(M, np.log(alpha0 / (1 - alpha0)))
    ell_target = np.log(alpha_target / (1 - alpha_target))
    up = alpha_target > alpha0
    tau = np.full(M, np.nan)
    alive = np.ones(M, dtype=bool)
    for j in range(int(T / dt)):
        n = policy(1 / (1 + np.exp(-ell[alive])))
        dX = (n * mu_true * dt
              + mkt.sigma * np.sqrt(n * dt) * rng.standard_normal(alive.sum()))
        ell[alive] += (dmu * dX - n * dt * half) / mkt.sigma ** 2
        hit = alive & ((ell >= ell_target) if up else (ell <= ell_target))
        tau[hit] = (j + 1) * dt
        alive &= ~hit
        if not alive.any():
            break
    return tau


κ = (mkt.mu_H - mkt.mu_L) ** 2 / (2 * mkt.sigma ** 2)
rng = np.random.default_rng(7)
cases = [('good product, 0.03 -> 0.9', 0.03, 0.9, mkt.mu_H),
         ('bad product,  0.97 -> 0.1', 0.97, 0.1, mkt.mu_L)]
```

```
print(f'{"":27}{"bounds (1/3, 2/3)":>20}{"bounds (n_L, n_H)":>20}{"simulated":>16}
 ↪')
for name, a_start, a_end, mu_true in cases:
    dist = abs(np.log(a_end / (1 - a_end)) - np.log(a_start / (1 - a_start)))
    tau = first_passage_times(mkt, a_start, a_end, mu_true, 4_000, 5e-3, 100.0,␣
 ↪rng)
    print(f'{name:27}{1.5 * dist / κ:9.2f},{3 * dist / κ:9.2f} '
          f'{dist / (κ * n_H):9.2f},{dist / (κ * n_L):9.2f} '
          f'{np.nanmean(tau):9.2f} ({np.nanstd(tau) / np.sqrt(len(tau)):.2f})')
```

```
                            bounds (1/3, 2/3)   bounds (n_L, n_H)       ␣
 ↪simulated
good product, 0.03 -> 0.9       5.25,    10.51      5.39,    10.01      8.27 (0.
 ↪08)
bad product,  0.97 -> 0.1       5.25,    10.51      5.39,    10.01      6.15 (0.
 ↪05)
```

Both simulated means lie inside the bounds.

Optional stopping pins down $E\left[\int_0^\tau n(\alpha_s)\,ds\right]$, the expected cumulative size of the experiment, not $E[\tau]$ itself.

A good product that starts from a pessimistic belief is sold to a small share of buyers, near n_L, while the market learns about it.

A bad product that starts from an optimistic belief is sold to a large share, near n_H.

The market therefore learns faster that a popular product is bad than that an unpopular product is good.

CHAPTER

THIRTYEIGHT

WHICH MOMENTS TO MATCH?

Contents

- *Which Moments to Match?*
  - *Overview*
  - *The estimation problem*
  - *The auxiliary model and its score*
  - *The estimator*
  - *Why the score is the right thing to match*
  - *Efficiency*
  - *Comparison with matching the auxiliary parameters*
  - *Choosing the score generator*
  - *A first example: an AR(1) inside an AR(2)*
  - *A second example: stochastic volatility*
  - *Testing the structural model*
  - *Applications*
  - *Recovering the latent state*
  - *Related lectures*
  - *Exercises*

38.1 Overview

Some structural economic models have the following two properties:

- it is possible to compute an equilibrium and simulate it, given a known parameter vector
- it is not possible to write down a likelihood function analytically

We can solve "the direct problem," which maps a parameter vector into a probability distribution of the data.

We cannot use maximum likelihood to solve "the inverse problem," which maps an observed data set into an estimate of the parameter vector.

Researchers with such models turn to the generalized method of moments of Hansen [1982] and compute the moments by simulation [Duffie and Singleton, 1993, McFadden, 1989, Pakes and Pollard, 1989].

They must decide which moments to match.

A. Ronald Gallant and George Tauchen studied this question in Gallant and Tauchen [1996].

They put two models on the table:

- the economist's **structural** model, indexed by a parameter vector ρ, whose likelihood function cannot be written down
- a distinct, descriptive **auxiliary** model, indexed by a parameter vector θ, whose likelihood function can be written down and maximized

The two parameter vectors usually differ in dimension and always differ in meaning.

- ρ specifies the preferences, technologies, constraint sets and other features of the agents inside the structural model. The economist wants to learn ρ.
- θ summarizes the observed data. Its components answer no economic question. The economist estimates θ only as a step toward estimating ρ.

	structural model	auxiliary model
density	$p(y_t \mid x_t, \rho)$	$f(y_t \mid x_t, \theta)$
parameter	$\rho \in \mathbb{R}^{k_\rho}$	$\theta \in \mathbb{R}^{k_\theta}$
interpretation	economic	descriptive
likelihood	unavailable	available
role	object of inference	source of moment conditions

Here x_t collects lagged values of y_t, possibly the whole history.

In the stochastic volatility example below, $\rho = (a, r_w, r_y)$ governs a latent log-volatility process, while $\theta = (\mu, \omega, \alpha, \beta)$ holds the parameters of a GARCH(1,1) model.

No component of θ corresponds to a component of ρ.

Gallant and Tauchen's method uses the **scores** associated with the likelihood function of the auxiliary model as the moment vector for GMM.

They fit θ to the observed data once.

They then choose ρ so that, in data simulated from the structural model, the auxiliary score evaluated at that fitted θ averages as close to zero as possible.

Note

Section *Why the score is the right thing to match* shows that the score moment conditions are the negative gradient of a Kullback-Leibler divergence. *Statistical Divergence Measures* and *Likelihood Ratio Processes* study properties of the KL divergence.

Gallant and Tauchen prove that if the auxiliary model nests the structural model, their estimator is as efficient as maximum likelihood.

They argue that if the auxiliary model closely approximates the distribution of the data, the estimator is nearly as efficient.

Practitioners call their method the **efficient method of moments**, or EMM.

Gallant and Tauchen [2010] surveys the method and its applications.

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from numba import jit
from scipy.optimize import minimize
from scipy.stats import chi2
```

38.2 The estimation problem

A structural model is a sequence of probability densities for an observed process $\{y_t\}$,

$$p(y_t \mid x_t, \rho), \qquad \rho \in \mathcal{R} \subset \mathbb{R}^{k_\rho},$$

where x_t collects whatever the model conditions on, usually lagged values of y_t.

We call ρ the **structural parameter vector**.

Nature generated the observed data $\{\tilde{y}_t, \tilde{x}_t\}_{t=1}^n$ from ρ^0.

Tildes mark observed data and statistics computed from them.

Hats mark simulated data and the final estimator $\hat{\rho}_n$.

We assume two things about the structural model.

Simulation is feasible: given ρ, we can generate a long artificial series $\{\hat{y}_\tau, \hat{x}_\tau\}_{\tau=1}^N$.

The likelihood is unavailable: $\prod_t p(\tilde{y}_t \mid \tilde{x}_t, \rho)$ has no expression that we can evaluate and maximize.

Gallant and Tauchen take a stationary Markov structural model as the leading case.

A single long simulation then suffices to compute any expectation under the model.

Their paper also covers independent cross-sectional data and processes whose laws of motion depend on t.

38.3 The auxiliary model and its score

An auxiliary model is a second sequence of densities for the same observed process,

$$f(y_t \mid x_t, \theta), \qquad \theta \in \Theta \subset \mathbb{R}^{k_\theta}.$$

We call θ the **auxiliary parameter vector**.

Gallant and Tauchen call the auxiliary model the **score generator**.

Two properties make an auxiliary model useful.

Its density has an analytic expression, so we can write down its likelihood and differentiate it.

It describes the data well.

The first property is required.

Consistency requires only that the scores identify ρ, as Exercise 38.15.1 and Exercise 38.15.3 show.

Section *Choosing the score generator* explains what a close fit buys.

A GARCH model of asset returns, a vector autoregression, and a Hermite series expansion of a conditional density all qualify.

The **score** of the auxiliary model is the derivative of its log density with respect to its own parameters,

$$s(y, x, \theta) = \frac{\partial}{\partial \theta} \ln f(y \mid x, \theta).$$

The score is a $k_\theta \times 1$ vector that we can evaluate at any (y, x, θ).

38.4 The estimator

The estimator takes three steps.

Step 1. Projection. Fit the auxiliary model to the observed data by quasi-maximum likelihood,

$$\tilde{\theta}_n = \arg\max_{\theta \in \Theta} \frac{1}{n} \sum_{t=1}^{n} \ln f(\tilde{y}_t \mid \tilde{x}_t, \theta).$$

This step uses the data and the auxiliary model alone, so it happens once, before any structural parameter is estimated.

Step 2. Moments by simulation. Define the moment function $m(\rho, \theta)$ as the expectation *under the structural model* of the auxiliary score,

$$m(\rho, \theta) = \iint s(y, x, \theta)\, p(y \mid x, \rho)\, dy\, p(x \mid \rho)\, dx, \tag{38.1}$$

where $p(x \mid \rho)$ is the stationary density of x_t under the structural model.

A long simulation approximates the double integral at $\theta = \tilde{\theta}_n$,

$$m(\rho, \tilde{\theta}_n) \approx \frac{1}{N} \sum_{\tau=1}^{N} s(\hat{y}_\tau, \hat{x}_\tau, \tilde{\theta}_n) \tag{38.2}$$

The right side of (38.2) averages the auxiliary score over a long simulation from the structural model at ρ.

Each evaluation requires a simulation and no optimization.

Step 3. GMM. Estimate ρ by making the moment vector as close to zero as the GMM metric allows,

$$\hat{\rho}_n = \arg\min_{\rho \in \mathcal{R}} \; m(\rho, \tilde{\theta}_n)^\top \, (\tilde{\mathcal{J}}_n)^{-1} \, m(\rho, \tilde{\theta}_n), \tag{38.3}$$

where

$$\tilde{\mathcal{J}}_n = \frac{1}{n} \sum_{t=1}^{n} s(\tilde{y}_t, \tilde{x}_t, \tilde{\theta}_n)\, s(\tilde{y}_t, \tilde{x}_t, \tilde{\theta}_n)^\top. \tag{38.4}$$

The weighting matrix (38.4) is the outer product of the auxiliary scores at the data.

It estimates $\mathcal{J}^0$, the asymptotic variance of $n^{-1/2} \sum_{t=1}^{n} s(\tilde{y}_t, \tilde{x}_t, \theta^0)$, where θ^0 is the probability limit of $\tilde{\theta}_n$.

It depends on the auxiliary model and the data alone.

The outer product is a valid estimate when the scores are serially uncorrelated.

An auxiliary model that describes the conditional distribution of the data correctly delivers serially uncorrelated scores.

Gaussian quasi-maximum likelihood scores are also serially uncorrelated when the auxiliary model specifies the conditional mean and variance correctly [Bollerslev and Wooldridge, 1992].

Otherwise Gallant and Tauchen recommend a heteroskedasticity and autocorrelation consistent estimate with Parzen weights [Gallant, 1987], and they cite Andrews [1991] for alternatives.

Exercise 38.15.4 shows how far the outer product can understate $\mathcal{J}^0$ when the scores are serially correlated.

A researcher computes either estimate once, before the search over ρ begins.

GMM usually requires a preliminary estimate of the structural parameters to form an efficient weighting matrix, and with either of these estimates that requirement disappears.

The order condition for identification is $k_\theta \geq k_\rho$.

The order condition is necessary but not sufficient, and Exercise 38.15.3 constructs a failure.

38.5 Why the score is the right thing to match

Statistical Divergence Measures shows that maximum likelihood minimizes the cross-entropy of the parametric model relative to the empirical distribution of the data, the sample analogue of the KL divergence from the data-generating distribution to the model.

Step 1 projects the data onto the auxiliary family by minimizing KL divergence.

For each ρ, define the **pseudo-true value**

$$\theta^0(\rho) = \arg\min_{\theta \in \Theta} \; D_{KL}(p(\cdot \mid \rho) \,\|\, f(\cdot \mid \theta)),$$

the point in the auxiliary family closest to the structural model at ρ [White, 1982][1].

The function $\theta^0(\cdot)$ maps structural parameters into auxiliary parameters.

As n grows, $\tilde{\theta}_n$ converges to $\theta^0 = \theta^0(\rho^0)$.

Now differentiate the divergence with respect to θ:

$$\frac{\partial}{\partial \theta} D_{KL}(p(\cdot \mid \rho) \,\|\, f(\cdot \mid \theta)) = \frac{\partial}{\partial \theta} \left[\int p \ln p - \int p \ln f_\theta \right] = -\, m(\rho, \theta) \tag{38.5}$$

The moment function (38.1) is the negative gradient of a Kullback-Leibler divergence.

Setting $m(\rho, \theta) = 0$ states the first-order condition for a KL projection.

The estimator (38.3) searches for the structural parameter whose KL projection onto the auxiliary family lands where the data's projection landed, or as close as the GMM metric allows when $k_\theta > k_\rho$.

This reading explains why the structural model need not nest inside the auxiliary model.

At $\rho = \rho^0$ the structural model generated the data, the two projections coincide, and $m(\rho^0, \theta^0) = 0$.

Exercise 38.15.1 asks you to verify this chain of reasoning.

[1] Here $p(\cdot \mid \rho)$ and $f(\cdot \mid \theta)$ are conditional densities of y_t given x_t. The divergence averages $\int p(y \mid x, \rho) \ln[p(y \mid x, \rho)/f(y \mid x, \theta)]\, dy$ over the stationary density $p(x \mid \rho)$.

38.6 Efficiency

Consistency and asymptotic normality hold under weak regularity conditions,

$$\sqrt{n}\,(\hat{\rho}_n - \rho^0) \to N\Big(0, \big[(M^0)^\top (\mathcal{J}^0)^{-1} M^0\big]^{-1}\Big), \qquad M^0 = \frac{\partial}{\partial \rho^\top} m(\rho^0, \theta^0),$$

which is Theorem 1 of Gallant and Tauchen [1996].

Here $\mathcal{J}^0$ is the variance of the auxiliary score that $\tilde{\mathcal{J}}_n$ in (38.4) estimates.

The formula assumes that the weighting matrix converges to $(\mathcal{J}^0)^{-1}$, as $(\tilde{\mathcal{J}}_n)^{-1}$ does when $\tilde{\mathcal{J}}_n$ estimates $\mathcal{J}^0$ consistently.

With any other limiting weighting matrix W, the asymptotic variance takes the sandwich form

$$\big[(M^0)^\top W M^0\big]^{-1} (M^0)^\top W \, \mathcal{J}^0 \, W M^0 \big[(M^0)^\top W M^0\big]^{-1}.$$

Exercise 38.15.4 uses it.

Gallant and Tauchen assume that N is large enough to make simulation error negligible, so the formula treats m as computed exactly.

A simulation of length N adds independent noise to m and multiplies the asymptotic variance of $\sqrt{n}\,m$, and hence of $\hat{\rho}_n$, by $1 + n/N$.

Gallant and Tauchen say that the structural model is **smoothly embedded** within the auxiliary model when a twice continuously differentiable map g carries structural parameters into auxiliary parameters,

$$p(y \mid x, \rho) = f(y \mid x, g(\rho)) \qquad \text{for every } \rho \text{ near } \rho^0.$$

Smooth embedding says that the auxiliary model nests the structural model.

In general $p(\cdot \mid \rho)$ lies outside the auxiliary family, and no g satisfies this equation.

The function $\theta^0(\rho)$ always exists.

Under smooth embedding $\theta^0(\rho) = g(\rho)$.

Write $G^0 = \partial g(\rho^0)/\partial \rho^\top$.

Their Theorem 2 states that smooth embedding implies

$$\sqrt{n}\,(\hat{\rho}_n - \rho^0) \to N\Big(0, \big[(G^0)^\top \mathcal{J}^0 G^0\big]^{-1}\Big),$$

which is the asymptotic distribution of the maximum likelihood estimator of ρ.

The proof turns on the identity $M^0 = \mathcal{J}^0 G^0$, which converts the GMM variance into the ML variance.

The estimator attains maximum likelihood efficiency without the structural likelihood.

Smooth embedding is a strong assumption, and consistency and asymptotic normality do not require it.

Gallant and Tauchen argue that an auxiliary model that approximates the distribution of the data closely without nesting the structural model delivers an estimator that is nearly fully efficient.

Gallant and Tauchen read this as instruction about how to choose the auxiliary model: make it fit.

38.7 Comparison with matching the auxiliary parameters

Indirect inference names the class of estimators that compare observed and simulated data through an auxiliary model.

Gourieroux *et al.* [1993] and Smith [1993] match the auxiliary parameters.

Gourieroux *et al.* [1993] call $\theta^0(\rho)$ the **binding function**.

Their estimator is

$$\hat{\rho}_{II} = \arg\min_{\rho} \left[\tilde{\theta}_n - \theta^0(\rho)\right]^\top \Omega \left[\tilde{\theta}_n - \theta^0(\rho)\right]$$

for a weighting matrix Ω.

Under smooth embedding $\theta^0 = g$, and the choice $\Omega = \mathcal{J}^0$ yields the minimum chi-square estimator, which is as efficient as maximum likelihood.

The binding function is unavailable whenever the structural likelihood is unavailable, so a practitioner replaces $\theta^0(\rho)$ by the average of auxiliary estimates computed on H simulated data sets.

That replacement is expensive.

Each trial value of ρ requires H separate maximizations of the auxiliary likelihood, one per simulated data set.

The score-based estimator requires exactly one maximization of the auxiliary likelihood, in Step 1, and never repeats it.

With suitably chosen weighting matrices, the parameter-based and score-based estimators are asymptotically equivalent [Gourieroux *et al.*, 1993].

The advantage of matching scores is computational.

Finite-sample behavior can differ.

Duffee and Stanton [2008] study highly persistent Gaussian autoregressions and find that score matching yields confidence intervals that are wider than those of parameter matching yet less likely to cover the true value.

They trace the problem to the large role that the population variance plays in the score moment conditions.

38.8 Choosing the score generator

A model known to describe the data accurately should be the auxiliary model.

Lacking one, Gallant and Tauchen recommend two general purpose families.

The **seminonparametric** or SNP density of Gallant and Nychka [1987] and Gallant and Tauchen [1989] multiplies a Gaussian density by a squared Hermite polynomial whose coefficients depend on lagged values.

Tuning parameters nest an IID Gaussian, a Gaussian VAR, Gaussian ARCH and GARCH, and fully nonlinear nonparametric specifications, so a specification search selects the auxiliary model from the data [Gallant and Tauchen, 2010].

The **neural net** score approximates any nonlinear autoregression, including deterministic chaos.

Both families are series expansions, so adding terms drives the approximation error to zero.

Adding terms also drives the conditional mean of the score toward zero, so the scores approach martingale differences.

The outer product (38.4) then estimates $\mathcal{J}^0$ accurately, and the weighting matrix needs no correction for serial correlation.

A close fit therefore buys two things: near efficiency and a valid outer-product weighting matrix.

Exercise 38.15.4 shows what happens with a score generator that lacks the second.

The number of terms governs a trade-off between bias and variance.

More terms reduce the approximation error of the auxiliary model and raise the sampling variance of $\tilde{\theta}_n$.

Gallant and Tauchen warn against taking too many terms and then claiming that standard asymptotics apply.

38.9 A first example: an AR(1) inside an AR(2)

An example in which we *can* compute the maximum likelihood estimator lets us check Theorem 2.

Take as the structural model a Gaussian first-order autoregression,

$$y_t = \rho_1 y_{t-1} + \sigma \epsilon_t, \qquad \epsilon_t \sim N(0,1),$$

with $\rho = (\rho_1, \sigma)^\top$.

Take as the auxiliary model a Gaussian second-order autoregression,

$$y_t = \theta_1 y_{t-1} + \theta_2 y_{t-2} + \varsigma e_t, \qquad e_t \sim N(0,1),$$

with $\theta = (\theta_1, \theta_2, \varsigma)^\top$.

The AR(2) nests the AR(1) through $g(\rho_1, \sigma) = (\rho_1, 0, \sigma)$, so the structural model is smoothly embedded and $k_\theta = 3$ exceeds $k_\rho = 2$.

The components θ_1 and ρ_1 share an interpretation only because the AR(2) nests the AR(1).

Theorem 2 predicts that our estimator and the maximum likelihood estimator have the same asymptotic distribution, up to the simulation factor $1 + n/N$.

Writing $u_t = y_t - \theta_1 y_{t-1} - \theta_2 y_{t-2}$, the auxiliary score is

$$s(y_t, x_t, \theta) = \begin{bmatrix} u_t y_{t-1}/\varsigma^2 \\ u_t y_{t-2}/\varsigma^2 \\ -1/\varsigma + u_t^2/\varsigma^3 \end{bmatrix}.$$

```
@jit(nopython=True)
def simulate_ar1(ρ1, σ, ε):
    "Simulate the structural AR(1) from a fixed array of shocks ε."
    n = ε.size
    y = np.empty(n)
    y[0] = σ * ε[0] / np.sqrt(1 - ρ1**2)
    for t in range(1, n):
        y[t] = ρ1 * y[t-1] + σ * ε[t]
    return y

def ar2_scores(θ, y):
    "Per-observation scores of the Gaussian AR(2) auxiliary model."
    θ1, θ2, ς = θ
    yt, y1, y2 = y[2:], y[1:-1], y[:-2]
    u = yt - θ1 * y1 - θ2 * y2
    return np.column_stack([u * y1 / ς**2, u * y2 / ς**2, -1/ς + u**2 / ς**3])

def ar2_qmle(y):
    "Quasi-maximum likelihood estimates of the AR(2), available in closed form."
    X = np.column_stack([y[1:-1], y[:-2]])
    b = np.linalg.lstsq(X, y[2:], rcond=None)[0]
    ς = np.sqrt(np.mean((y[2:] - X @ b)**2))
    return np.array([b[0], b[1], ς])
```

Let's generate data from the structural model and compute the conditional maximum likelihood estimator, which for this model is least squares.

```
rng = np.random.default_rng(0)
T, ρ1_0, σ_0 = 1000, 0.8, 1.0
y_data = simulate_ar1(ρ1_0, σ_0, rng.standard_normal(T))

X1 = y_data[:-1]
ρ1_mle = float(X1 @ y_data[1:] / (X1 @ X1))
σ_mle = float(np.sqrt(np.mean((y_data[1:] - ρ1_mle * X1)**2)))
print(f"maximum likelihood:  ρ1 = {ρ1_mle:.4f}   σ = {σ_mle:.4f}")
```

```
maximum likelihood:  ρ1 = 0.8022   σ = 0.9784
```

Step 1 fits the auxiliary model and forms the weighting matrix.

```
θ_tilde = ar2_qmle(y_data)
S = ar2_scores(θ_tilde, y_data)
I_tilde = S.T @ S / S.shape[0]
W = np.linalg.inv(I_tilde)

print(f"auxiliary QMLE: θ1 = {θ_tilde[0]:.4f}, θ2 = {θ_tilde[1]:.4f}, ς = {θ_
↪tilde[2]:.4f}")
```

```
auxiliary QMLE: θ1 = 0.8083, θ2 = -0.0075, ς = 0.9789
```

The estimate of θ_2 sits near zero, as the nesting map g anticipates.

Steps 2 and 3 simulate and minimize.

We draw the shocks ϵ once and reuse them at every trial value of ρ.

We call them **common random numbers**, and a figure below shows their effect on the criterion.

```
N = 20_000
ε_sim = rng.standard_normal(N)      # common random numbers

def moments_ar1(ρ):
    "Expectation under the structural model of the auxiliary score."
    y_sim = simulate_ar1(ρ[0], ρ[1], ε_sim)
    return ar2_scores(θ_tilde, y_sim).mean(axis=0)

def criterion_ar1(ρ):
    if not (abs(ρ[0]) < 0.999 and ρ[1] > 1e-6):
        return 1e10
    m = moments_ar1(ρ)
    return float(m @ W @ m)

res = minimize(criterion_ar1, np.array([0.5, 1.5]), method='Nelder-Mead',
               options=dict(xatol=1e-6, fatol=1e-12, maxiter=2000))

print(f"true values:         ρ1 = {ρ1_0:.4f}   σ = {σ_0:.4f}")
print(f"maximum likelihood:  ρ1 = {ρ1_mle:.4f}   σ = {σ_mle:.4f}")
print(f"score matching:      ρ1 = {res.x[0]:.4f}   σ = {res.x[1]:.4f}")
```

```
true values:         ρ1 = 0.8000   σ = 1.0000
maximum likelihood:  ρ1 = 0.8022   σ = 0.9784
score matching:      ρ1 = 0.8009   σ = 0.9830
```

The two sets of estimates nearly coincide, differing in the third decimal place.

The next figure evaluates the criterion on a grid of ρ_1 values with σ held at its estimate.

One curve uses the common random numbers `ε_sim`, and the other draws fresh shocks at each grid point.

```
rng_fresh = np.random.default_rng(1)
ρ1_grid = np.linspace(0.77, 0.83, 61)
σ_hat = res.x[1]

def criterion_fresh(ρ):
    "The criterion with a fresh draw of shocks at each evaluation."
    y_sim = simulate_ar1(ρ[0], ρ[1], rng_fresh.standard_normal(N))
    m = ar2_scores(θ_tilde, y_sim).mean(axis=0)
    return float(m @ W @ m)

crn = [criterion_ar1(np.array([r1, σ_hat])) for r1 in ρ1_grid]
fresh = [criterion_fresh(np.array([r1, σ_hat])) for r1 in ρ1_grid]

fig, ax = plt.subplots()
ax.plot(ρ1_grid, crn, lw=2, label='common random numbers')
ax.plot(ρ1_grid, fresh, lw=2, label='fresh shocks at each point')
ax.set_xlabel(r'$\rho_1$')
ax.set_ylabel('criterion')
ax.legend()
plt.show()
```

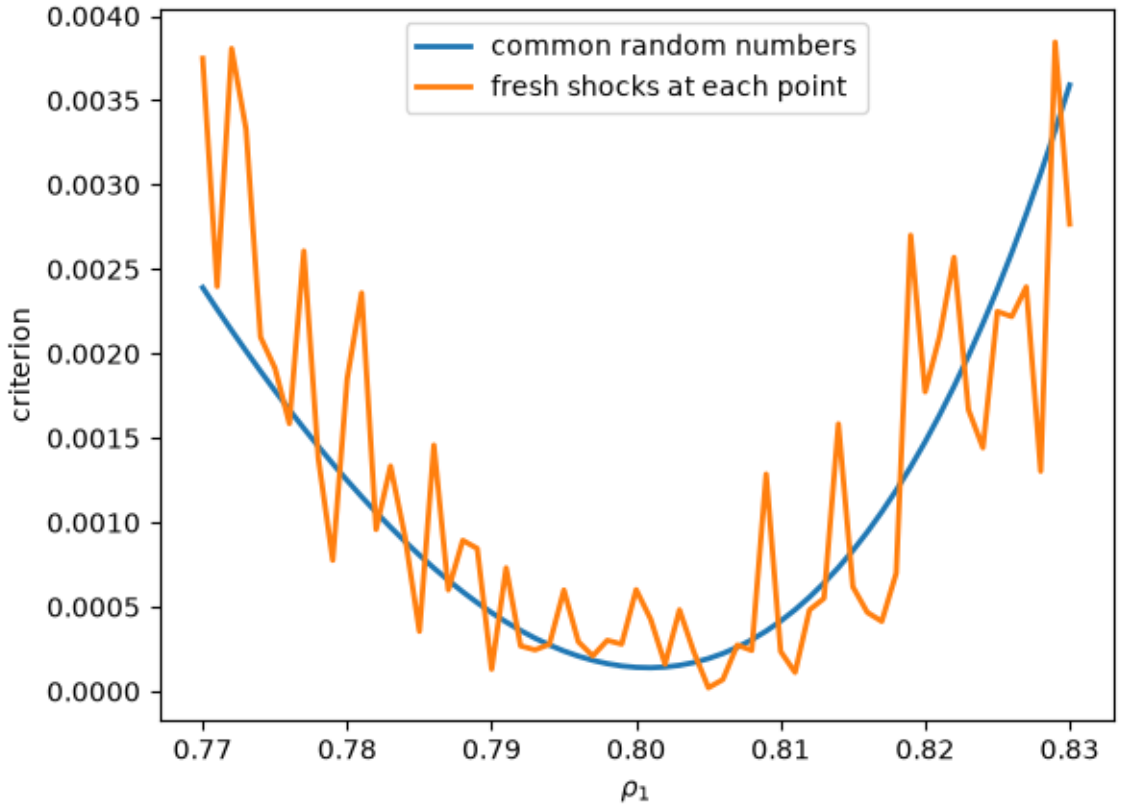


Fig. 38.1: Criterion with common and fresh shocks

Fresh shocks make the criterion jagged, with many local minima.

Common random numbers make it a smooth function of ρ_1 with a single minimum.

The minimized criterion tests the $k_\theta - k_\rho = 1$ overidentifying restriction that the structural model places on the auxiliary model.

Under correct specification, n times the minimized criterion, divided by the simulation factor $1 + n/N$, is asymptotically distributed as $\chi^2(k_\theta - k_\rho)$[2].

The code writes the sample size n as `T`.

```
J = T * criterion_ar1(res.x) / (1 + T / N)
df = 3 - 2
print(f"J = {J:.3f},  χ²({df}) 5% critical value = {chi2.ppf(0.95, df):.3f}")
```

```
J = 0.133,  χ²(1) 5% critical value = 3.841
```

Theorem 1 delivers standard errors.

We compute M^0 by central differences of the simulated moment vector, which common random numbers make smooth in ρ, and we multiply the variance by the simulation factor $1 + n/N$.

For the AR(1), the maximum likelihood standard errors are $\sqrt{(1-\rho_1^2)/n}$ for ρ_1 and $\sigma/\sqrt{2n}$ for σ.

```
def jacobian(moments, ρ, h=1e-4):
    "Central-difference Jacobian of a moment function with respect to ρ."
    cols = []
    for j in range(ρ.size):
        e = np.zeros(ρ.size)
        e[j] = h * max(1.0, abs(ρ[j]))
        cols.append((moments(ρ + e) - moments(ρ - e)) / (2 * e[j]))
    return np.column_stack(cols)

def std_errors(moments, ρ, W, n, N):
    "Asymptotic standard errors from Theorem 1, including the simulation factor."
    M = jacobian(moments, ρ)
    V = np.linalg.inv(M.T @ W @ M) * (1 + n / N) / n
    return np.sqrt(np.diag(V))

se_emm = std_errors(moments_ar1, res.x, W, T, N)
se_mle = np.array([np.sqrt((1 - ρ1_mle**2) / T), σ_mle / np.sqrt(2 * T)])

print(f"{'':20}{'se ρ1':>10}{'se σ':>10}")
print(f"{'maximum likelihood':20}{se_mle[0]:10.4f}{se_mle[1]:10.4f}")
print(f"{'score matching':20}{se_emm[0]:10.4f}{se_emm[1]:10.4f}")

e_data = y_data[1:] - ρ1_mle * y_data[:-1]
kurt_e = ((e_data - e_data.mean())**4).mean() / e_data.var()**2
print(f"\nkurtosis of the residuals = {kurt_e:.2f}")
```

```
                      se ρ1      se σ
maximum likelihood   0.0189    0.0219
score matching       0.0188    0.0241

kurtosis of the residuals = 3.29
```

[2] The factor is 1.05 in the AR(1) example and 1.04 in the stochastic volatility example. It matters when N is not large relative to n.

The two standard errors for ρ_1 agree to three decimal places.

The score matching standard error for σ exceeds the maximum likelihood one by about 10 percent.

The simulation factor $1 + n/N = 1.05$ accounts for about 2.5 percent.

The weighting matrix $\tilde{\mathcal{J}}_n$ accounts for the rest: it estimates the variance of the σ score from the sample fourth moment of the residuals, whose kurtosis is 3.3 in this sample and 3 in the population.

Exercise 38.15.2 asks you to compare the sampling distributions of the two estimators.

38.10 A second example: stochastic volatility

Now we turn to a model whose likelihood function is intractable.

Section 4.2 of Gallant and Tauchen [1996] discusses a stochastic volatility model whose mean equation also contains an intercept and a lagged return.

We drop those two terms and study

$$
\begin{aligned}
y_t &= \exp(w_t)\, r_y\, z_t, \\
w_t &= a\, w_{t-1} + r_w\, \eta_t,
\end{aligned}
$$

with z_t and η_t independent standard normals and $\rho = (a, r_w, r_y)^\top$.

The observed return y_t has a volatility $\exp(w_t) r_y$ driven by a latent autoregression w_t.

Writing the likelihood of $\{y_t\}_{t=1}^n$ requires integrating the latent volatility path out of an n-dimensional integral,

$$
p(y_1, \ldots, y_n \mid \rho) = \int \cdots \int \prod_{t=1}^{n} \phi(y_t; 0, e^{2w_t} r_y^2)\, p(w_1, \ldots, w_n \mid a, r_w)\, dw_1 \cdots dw_n,
$$

which no analytic expression evaluates.

Simulating the model takes two lines of code.

The model generates fat tails and persistent volatility.

Gallant *et al.* [1997] fit the full model and extensions of it to 16,127 daily observations on the Standard and Poor's composite index, using SNP score generators.

Section 4.1 of Gallant and Tauchen [1996] notes that ARCH and GARCH models fit returns data well.

We choose as the auxiliary model a Gaussian GARCH(1,1),

$$
y_t = \mu + u_t, \qquad u_t = \sigma_t \nu_t, \qquad \sigma_t^2 = \omega + \alpha u_{t-1}^2 + \beta \sigma_{t-1}^2,
$$

with $\theta = (\mu, \omega, \alpha, \beta)^\top$ and $\nu_t \sim N(0, 1)$.

The GARCH conditional density depends on the entire past of y_t through the recursion for σ_t^2, so here x_t is the whole history.

Gallant and Tauchen call such a score generator non-Markov.

The structural parameters (a, r_w, r_y) describe a latent log-volatility process; the auxiliary parameters $(\mu, \omega, \alpha, \beta)$ describe a conditional variance recursion in observed returns.

The auxiliary model needs no economic interpretation.

Its credential is a good fit to the data.

We differentiate the GARCH log density numerically.

Numerical differentiation keeps the method general: an auxiliary model need only support evaluation of its log likelihood.

The code fits the GARCH model in the unconstrained coordinates $(\mu, \log\omega, \operatorname{logit}\alpha, \operatorname{logit}\beta)$.

The scores, the weighting matrix, and the moment vector all refer to those coordinates.

The reparameterization changes neither $\hat{\rho}_n$ nor the test statistic.

It multiplies the score by a fixed invertible matrix, and the criterion (38.3) is invariant to that transformation.

```
@jit(nopython=True)
def simulate_sv(a, r_w, r_y, z, η):
    "Simulate the structural stochastic volatility model."
    n = z.size
    y = np.empty(n)
    w = r_w * η[0] / np.sqrt(1.0 - a * a)
    y[0] = np.exp(w) * r_y * z[0]
    for t in range(1, n):
        w = a * w + r_w * η[t]
        y[t] = np.exp(w) * r_y * z[t]
    return y

@jit(nopython=True)
def garch_ll(μ, ω, α, β, y):
    "Per-observation log densities of the Gaussian GARCH(1,1) auxiliary model."
    n = y.size
    ll = np.empty(n)
    s2 = ω / max(1e-8, 1.0 - α - β)
    for t in range(n):
        u = y[t] - μ
        ll[t] = -0.5 * (np.log(2 * np.pi) + np.log(s2) + u * u / s2)
        s2 = ω + α * u * u + β * s2
    return ll

def unpack(θ):
    "Map unconstrained parameters to ω > 0 and α, β in (0, 1)."
    μ, lω, lα, lβ = θ
    return μ, np.exp(lω), 1 / (1 + np.exp(-lα)), 1 / (1 + np.exp(-lβ))

def garch_negll(θ, y):
    μ, ω, α, β = unpack(θ)
    if α + β > 0.999:
        return 1e10
    return -garch_ll(μ, ω, α, β, y).sum()

def garch_scores(θ, y):
    "Per-observation auxiliary scores by central differences."
    p, n = θ.size, y.size
    S = np.empty((n, p))
    for j in range(p):
        h = 1e-5 * max(1.0, abs(θ[j]))
        θp, θm = θ.copy(), θ.copy()
        θp[j] += h
        θm[j] -= h
        S[:, j] = (garch_ll(*unpack(θp), y) - garch_ll(*unpack(θm), y)) / (2 * h)
    return S
```

Section 4.2 reports no parameter values, so we choose our own.

Let's generate a long return series from the structural model.

```
rng = np.random.default_rng(5)
T = 2000
a_0, r_w0, r_y0 = 0.80, 0.20, 0.65
y_sv = simulate_sv(a_0, r_w0, r_y0, rng.standard_normal(T), rng.standard_normal(T))

kurt = ((y_sv - y_sv.mean())**4).mean() / y_sv.var()**2
print(f"standard deviation = {y_sv.std():.3f},  kurtosis = {kurt:.2f}")
```

```
standard deviation = 0.738,  kurtosis = 3.93
```

The population kurtosis of y_t is $3\exp(4r_w^2/(1-a^2)) = 4.7$ at these parameter values, above the Gaussian value of 3.

The sample kurtosis understates it, as sample kurtoses of heavy-tailed series often do.

```
fig, ax = plt.subplots()
ax.plot(y_sv, lw=0.6)
ax.set_xlabel('$t$')
ax.set_ylabel('$y_t$')
plt.show()
```

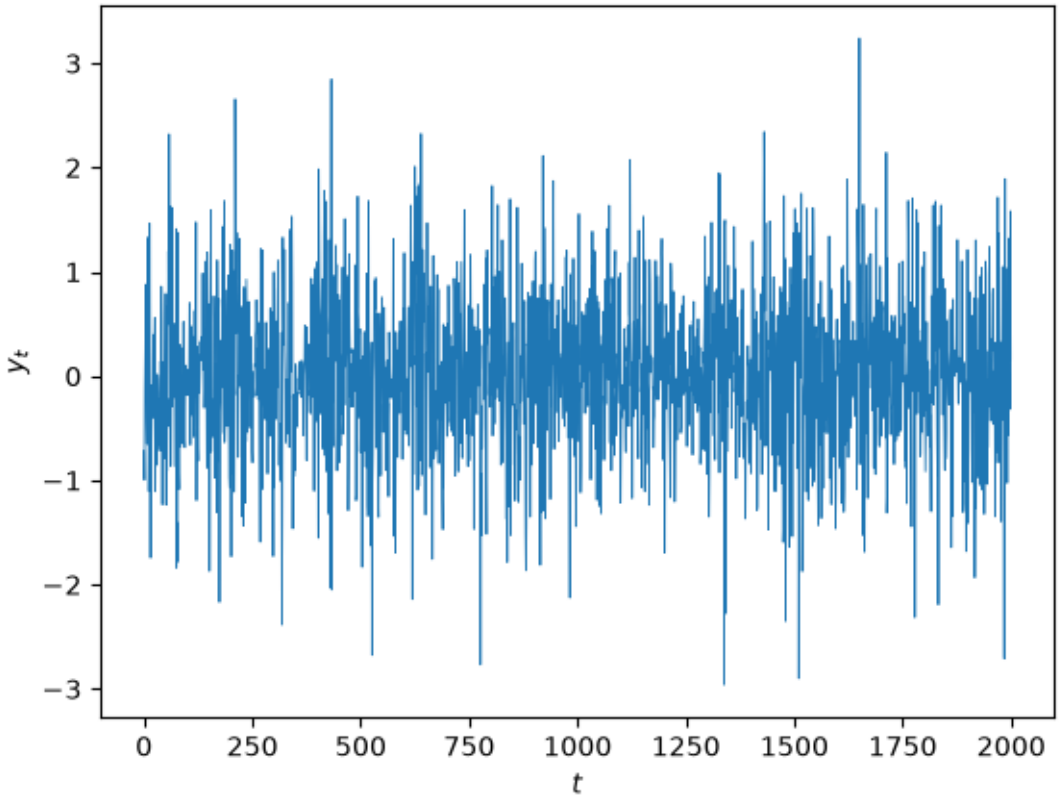


Fig. 38.2: Simulated returns from the SV model

Step 1 fits the GARCH auxiliary model to these data.

```
θ_start = np.array([0.0, np.log(0.05), -2.0, 1.5])
r = minimize(garch_negll, θ_start, args=(y_sv,), method='Nelder-Mead',
```

(continues on next page)

(continued from previous page)

```
            options=dict(maxiter=8000, xatol=1e-8, fatol=1e-8))
θ_sv = r.x
μ_h, ω_h, α_h, β_h = unpack(θ_sv)
print(f"auxiliary GARCH QMLE: μ = {μ_h:.4f}, ω = {ω_h:.4f}, "
      f"α = {α_h:.4f}, β = {β_h:.4f}   (α + β = {α_h + β_h:.4f})")

S_sv = garch_scores(θ_sv, y_sv)
W_sv = np.linalg.inv(S_sv.T @ S_sv / T)
```

```
auxiliary GARCH QMLE: μ = 0.0150, ω = 0.0922, α = 0.1209, β = 0.7127   (α + β = 0.
↪8337)
```

An $\alpha + \beta$ of 0.83 records the persistence of volatility that the auxiliary model reads off the data.

Steps 2 and 3 estimate the three structural parameters.

We search over unconstrained coordinates and reuse one pair of shock arrays throughout.

```
N = 50_000
z_sim, η_sim = rng.standard_normal(N), rng.standard_normal(N)

def moments_sv(ρ):
    y_s = simulate_sv(ρ[0], ρ[1], ρ[2], z_sim, η_sim)
    return garch_scores(θ_sv, y_s).mean(axis=0)

def to_ρ(u):
    "Map unconstrained coordinates to a in (0,1) and r_w, r_y > 0."
    return np.array([1 / (1 + np.exp(-u[0])), np.exp(u[1]), np.exp(u[2])])

def criterion_sv(u):
    if np.max(np.abs(u)) > 25:
        return 1e10
    m = moments_sv(to_ρ(u))
    return float(m @ W_sv @ m)

u_0 = np.array([np.log(0.6 / 0.4), np.log(0.3), np.log(0.5)])
res_sv = minimize(criterion_sv, u_0, method='Nelder-Mead',
                  options=dict(maxiter=400, xatol=1e-4, fatol=1e-10))
ρ_hat = to_ρ(res_sv.x)
```

Theorem 1 delivers standard errors.

```
se_sv = std_errors(moments_sv, ρ_hat, W_sv, T, N)

print(f"{'':17}{'a':>9}{'r_w':>9}{'r_y':>9}")
print(f"{'true values':17}{a_0:9.4f}{r_w0:9.4f}{r_y0:9.4f}")
print(f"{'score matching':17}" + "".join(f"{x:9.4f}" for x in ρ_hat))
print(f"{'standard errors':17}" + "".join(f"{x:9.4f}" for x in se_sv))
```

```
                         a       r_w       r_y
true values         0.8000    0.2000    0.6500
score matching      0.8447    0.1794    0.6700
standard errors     0.0404    0.0372    0.0244
```

The estimator locates all three parameters of a model whose likelihood function we never wrote down.

Each estimate lies within about one standard error of its true value.

The data pin down the scale r_y most precisely.

38.11 Testing the structural model

The auxiliary model has four parameters and the structural model has three, so one overidentifying restriction remains to be tested.

```
J_sv = T * criterion_sv(res_sv.x) / (1 + T / N)
df_sv = 4 - 3
print(f"J = {J_sv:.3f},  χ²({df_sv}) 5% critical value = {chi2.ppf(0.95, df_sv):.3f}")
```

```
J = 0.606,  χ²(1) 5% critical value = 3.841
```

The data give the structural model no trouble, as they should, because the structural model generated them.

A misspecified structural model fares differently.

Let's impose $r_w = 0$, which shuts down stochastic volatility and leaves a Gaussian white noise with one free parameter.

```
def moments_const_vol(v):
    "Moments of the constant volatility model r_w = 0, with v = log r_y; a has no␣
 ↪effect."
    return moments_sv(np.array([0.5, 1e-8, np.exp(v[0])]))

def criterion_const_vol(v):
    m = moments_const_vol(v)
    return float(m @ W_sv @ m)

r_cv = minimize(criterion_const_vol, np.array([0.0]), method='Nelder-Mead',
                options=dict(maxiter=200, xatol=1e-6))
J_cv = T * criterion_const_vol(r_cv.x) / (1 + T / N)
df_cv = 4 - 1
print(f"constant volatility model: J = {J_cv:.1f}, "
      f"χ²({df_cv}) 5% critical value = {chi2.ppf(0.95, df_cv):.3f}")
```

```
constant volatility model: J = 21.0, χ²(3) 5% critical value = 7.815
```

The test rejects the restricted model at the 1% level[3].

Gallant and Tauchen emphasize this second use of the machinery.

An auxiliary model that approximates the data well makes the test powerful against structural model error, and the individual score components say *which* features of the data a rejected model fails to match.

Gallant and Tauchen's **quasi-t-ratios** divide each component of $\sqrt{n}\, m(\hat{\rho}_n, \tilde{\theta}_n)$ by its standard error.

The asymptotic variance of $\sqrt{n}\, m(\hat{\rho}_n, \tilde{\theta}_n)$ is

$$\left(1 + \frac{n}{N}\right)\left[\mathcal{J}^0 - M^0\left((M^0)^\top (\mathcal{J}^0)^{-1} M^0\right)^{-1}(M^0)^\top\right].$$

Estimating ρ removes the variance in the directions that the columns of M^0 span.

The unconstrained coordinates transform each GARCH parameter separately through an increasing function.

[3] The $\chi^2(3)$ critical value is approximate here. Under the constant volatility null, the pseudo-true GARCH value has $\alpha = 0$, on the boundary of the parameter space, and β is then unidentified. In the code's coordinates, logit $\alpha \to -\infty$. Standard asymptotics do not apply. The weighting matrix raises no separate concern: under the null, the GARCH model with $\alpha = 0$ describes the data correctly, so its scores are martingale differences and the outer product is valid. Gallant and Tauchen observe that null hypotheses usually deliver such conditions.

The t-ratios are therefore the same as in the coordinates $(\mu, \omega, \alpha, \beta)$.

```
I_sv = S_sv.T @ S_sv / T
M_cv = jacobian(moments_const_vol, r_cv.x)
V_cv = (I_sv - M_cv @ np.linalg.inv(M_cv.T @ W_sv @ M_cv) @ M_cv.T) * (1 + T / N)
t_cv = np.sqrt(T) * moments_const_vol(r_cv.x) / np.sqrt(np.diag(V_cv))

for name, t in zip(['μ', 'ω', 'α', 'β'], t_cv):
    print(f"{name:>4}: {t:6.2f}")
```

```
   μ:  -0.75
   ω:  -3.04
   α:  -3.95
   β:  -2.79
```

The components for ω, α and β, which govern the conditional variance, carry the rejection.

The μ component does not.

Each negative sign says that, holding the other auxiliary parameters fixed, data from the constant volatility model would pull that parameter below the value that the observed data choose.

The constant volatility model matches the mean of returns and fails to reproduce the dependence of today's variance on yesterday's squared return.

Gallant *et al.* [1997] use the score vector of a rejected stochastic volatility model to diagnose its defects, and they then repair the model by thickening the tails of z_t and lengthening the memory of the volatility equation.

38.12 Applications

Gallant and Tauchen [2010] collects applications of the method to continuous-time models in finance.

Andersen *et al.* [2002] estimate stochastic volatility diffusions for the Standard and Poor's 500 index and reject them, and jump diffusions pass the same tests.

Other applications fit multifactor stochastic volatility models, term structure models of interest rates, and exchange rate models.

Bansal *et al.* [1995] fit a two-country general equilibrium monetary model to high-frequency currency market data, and Gallant and Tauchen [1996] propose fitting a one-sector stochastic growth model to consumption and stock returns.

Empirical auction models form another class of applications.

Game theory delivers the density of the winning bid under independent private values, and that density has a tractable form only for special valuation distributions.

Simulating from it is easy for any valuation distribution, so score matching brings all of the restrictions of the theory to bear.

38.13 Recovering the latent state

Having estimated ρ, a researcher often wants the latent state, such as the volatility path behind a return series.

Reprojection recovers it [Gallant and Tauchen, 2010].

Simulate a long path of states and observables at $\hat{\rho}_n$, calibrate inside the simulation the conditional distribution of the latent state given current and lagged observables, then evaluate those calibrated functions at the observed data.

Reprojection is a simulation-based nonlinear filter.

38.14 Related lectures

Statistical Divergence Measures develops the Kullback-Leibler divergence that appears in (38.5) and shows its link to maximum likelihood.

Likelihood Ratio Processes studies likelihood ratio processes and expresses their drift in terms of two KL divergences, a representation closely related to the projection step of this lecture.

Maximum Likelihood Estimation estimates parameters by maximum likelihood in settings where the likelihood is available.

Estimating Euler Equations by Generalized Method of Moments applies the generalized method of moments to Euler equations of an asset pricing model.

Mean of a Likelihood Ratio Process describes importance sampling, an alternative route to expectations that a model cannot deliver analytically.

38.15 Exercises

Exercise 38.15.1

Let the structural model at ρ have density $p(\cdot \mid \rho)$ and let the auxiliary family be $\{f(\cdot \mid \theta) : \theta \in \Theta\}$.

a. Show that

$$\frac{\partial}{\partial \theta} D_{KL}(p(\cdot \mid \rho) \,\|\, f(\cdot \mid \theta)) = -\, m(\rho, \theta),$$

where $m(\rho, \theta) = \mathbb{E}_{p(\cdot|\rho)}[\partial \ln f(\cdot \mid \theta)/\partial \theta]$.

b. Conclude that the pseudo-true value $\theta^0(\rho)$ satisfies $m(\rho, \theta^0(\rho)) = 0$.

c. Explain why $m(\rho^0, \theta^0) = 0$, and why the structural model does not need to nest inside the auxiliary family for this to hold.

Solution

For part a, write the divergence as

$$D_{KL}(p(\cdot \mid \rho) \,\|\, f(\cdot \mid \theta)) = \int p(y \mid \rho) \ln p(y \mid \rho)\, dy - \int p(y \mid \rho) \ln f(y \mid \theta)\, dy.$$

The first term does not involve θ.

Differentiating the second term under the integral sign gives

$$\frac{\partial}{\partial\theta} D_{KL} = -\int p(y \mid \rho) \frac{\partial}{\partial\theta} \ln f(y \mid \theta)\, dy = -m(\rho, \theta).$$

For part b, the pseudo-true value minimizes the divergence over θ.

An interior minimum sets the gradient to zero, so $m(\rho, \theta^0(\rho)) = 0$.

For part c, nature generated the data from ρ^0, so the quasi-maximum likelihood estimator $\tilde{\theta}_n$ converges to $\theta^0 = \theta^0(\rho^0)$, the KL projection of the true data generating process.

Part b evaluated at $\rho = \rho^0$ then gives $m(\rho^0, \theta^0) = 0$.

Nesting plays no role in this argument.

The pseudo-true value is defined as a projection whether or not the divergence at its minimum is zero, and the divergence is zero only in the nested case.

A misspecified auxiliary model still delivers valid moment conditions.

It costs efficiency.

It preserves consistency.

Exercise 38.15.2

Theorem 2 of Gallant and Tauchen [1996] says that the score matching estimator and the maximum likelihood estimator have the same asymptotic distribution when the structural model is smoothly embedded within the auxiliary model and simulation error is negligible.

Check this claim by Monte Carlo in the AR(1) and AR(2) example.

Generate 60 independent samples of length $T = 500$ from the structural AR(1) with $\rho_1 = 0.8$ and $\sigma = 1$.

For each sample compute both estimators, then compare their means, their standard deviations, and the correlation between them across replications.

Compare the Monte Carlo standard deviations with the asymptotic standard deviations that Theorem 2 delivers.

Solution

```
T_mc, N_mc, R = 500, 20_000, 60
rng_mc = np.random.default_rng(7)
ε_mc = rng_mc.standard_normal(N_mc)          # common random numbers

def emm_estimate(y, ε):
    "Score matching estimate of (ρ1, σ) from the sample y with the shock array ε."
    θ_r = ar2_qmle(y)
    S_r = ar2_scores(θ_r, y)
    W_r = np.linalg.inv(S_r.T @ S_r / S_r.shape[0])

    def crit_r(ρ):
        if not (abs(ρ[0]) < 0.999 and ρ[1] > 1e-6):
            return 1e10
        m = ar2_scores(θ_r, simulate_ar1(ρ[0], ρ[1], ε)).mean(axis=0)
        return float(m @ W_r @ m)
```

```
    return minimize(crit_r, np.array([0.5, 1.2]), method='Nelder-Mead',
                    options=dict(xatol=1e-6, fatol=1e-12)).x

samples = [simulate_ar1(ρ1_0, σ_0, rng_mc.standard_normal(T_mc)) for _ in range(R)]
mle_draws = np.empty((R, 2))
emm_draws = np.empty((R, 2))

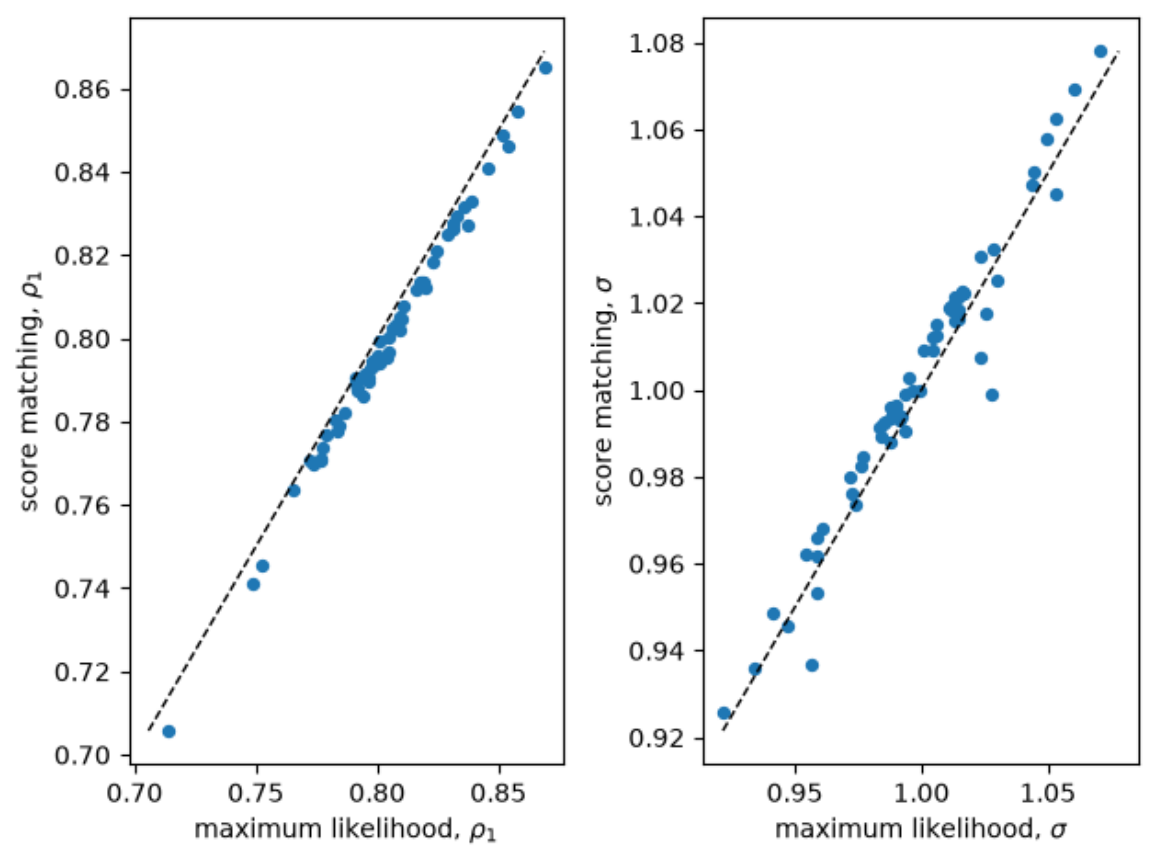
for r_, y_r in enumerate(samples):
    X_r = y_r[:-1]
    ρ_m = float(X_r @ y_r[1:] / (X_r @ X_r))
    σ_m = float(np.sqrt(np.mean((y_r[1:] - ρ_m * X_r)**2)))
    mle_draws[r_] = [ρ_m, σ_m]
    emm_draws[r_] = emm_estimate(y_r, ε_mc)

print(f"{'':20}{'mean ρ1':>10}{'sd ρ1':>10}{'mean σ':>10}{'sd σ':>10}")
print(f"{'maximum likelihood':20}{mle_draws[:,0].mean():10.4f}{mle_draws[:,0].
↪std():10.4f}"
      f"{mle_draws[:,1].mean():10.4f}{mle_draws[:,1].std():10.4f}")
print(f"{'score matching':20}{emm_draws[:,0].mean():10.4f}{emm_draws[:,0].std():10.
↪4f}"
      f"{emm_draws[:,1].mean():10.4f}{emm_draws[:,1].std():10.4f}")
sd_theory = [np.sqrt((1 - ρ1_0**2) / T_mc), σ_0 / np.sqrt(2 * T_mc)]
print(f"{'asymptotic theory':20}{ρ1_0:10.4f}{sd_theory[0]:10.4f}"
      f"{σ_0:10.4f}{sd_theory[1]:10.4f}")
print(f"\ncorrelation across replications: ρ1 "
      f"{np.corrcoef(emm_draws[:,0], mle_draws[:,0])[0,1]:.4f}, "
      f"σ {np.corrcoef(emm_draws[:,1], mle_draws[:,1])[0,1]:.4f}")
```

```
                        mean ρ1     sd ρ1    mean σ      sd σ
maximum likelihood       0.8038    0.0274    0.9981    0.0315
score matching           0.7991    0.0275    1.0015    0.0321
asymptotic theory        0.8000    0.0268    1.0000    0.0316

correlation across replications: ρ1 0.9976, σ 0.9748
```

```
fig, axes = plt.subplots(1, 2)
for j, name in enumerate([r'$\rho_1$', r'$\sigma$']):
    axes[j].scatter(mle_draws[:, j], emm_draws[:, j], s=18)
    lo = min(mle_draws[:, j].min(), emm_draws[:, j].min())
    hi = max(mle_draws[:, j].max(), emm_draws[:, j].max())
    axes[j].plot([lo, hi], [lo, hi], 'k--', lw=1)
    axes[j].set_xlabel(f'maximum likelihood, {name}')
    axes[j].set_ylabel(f'score matching, {name}')
plt.tight_layout()
plt.show()
```



The two estimators have nearly identical standard deviations.

Both standard deviations are close to the asymptotic values $\sqrt{(1-\rho_1^2)/T}$ and $\sigma/\sqrt{2T}$ that Theorem 2 delivers.

The simulation factor $1 + T/N$ raises the score matching values by about 1 percent.

The correlation between the two estimators across replications sits very close to one, and the scatter plots hug the 45 degree line.

The two estimators are nearly the same random variable in each sample.

The means of the two estimators of ρ_1 differ by about 0.005.

Every replication shares the shock array `ε_mc`, so its simulation error shifts all 60 score matching estimates together.

The next cell repeats the score matching estimates of ρ_1 with three other shock arrays.

```
gap = (emm_draws[:, 0] - mle_draws[:, 0]).mean()
print(f"shock array ε_mc:       mean gap in ρ1 = {gap:7.4f}")
for seed in (100, 101, 102):
    ε_alt = np.random.default_rng(seed).standard_normal(N_mc)
    ρ1_alt = np.array([emm_estimate(y_r, ε_alt)[0] for y_r in samples])
    gap = (ρ1_alt - mle_draws[:, 0]).mean()
    print(f"shock array seed {seed}:  mean gap in ρ1 = {gap:7.4f}")
```

```
shock array ε_mc:       mean gap in ρ1 = -0.0047
shock array seed 100:  mean gap in ρ1 =  0.0013
shock array seed 101:  mean gap in ρ1 =  0.0041
shock array seed 102:  mean gap in ρ1 =  0.0004
```

The gap changes sign across shock arrays, so simulation error accounts for it.

Exercise 38.15.3

The order condition $k_\theta \geq k_\rho$ does not guarantee identification.

Gallant and Tauchen warn that the parameters of the mixing process of a stochastic volatility model need third and fourth moment information, and that the score of a Gaussian vector autoregression will not supply it.

Investigate their warning.

Replace the GARCH auxiliary model of the stochastic volatility example by a Gaussian AR(2), which has $k_\theta = 3$ parameters against $k_\rho = 3$ structural parameters, so the order condition holds.

With $k_\theta = k_\rho$ the model leaves no overidentifying restriction to test.

Minimize the resulting criterion from several different starting values and report the estimates together with the criterion values.

Explain what you find, and say which function of (a, r_w, r_y) the Gaussian AR(2) score does pin down.

Solution

```
θ_ar = ar2_qmle(y_sv)
S_ar = ar2_scores(θ_ar, y_sv)
W_ar = np.linalg.inv(S_ar.T @ S_ar / S_ar.shape[0])

def criterion_ar_score(u):
    if np.max(np.abs(u)) > 25:
        return 1e10
    ρ = to_ρ(u)
    if ρ[1] / np.sqrt(1 - ρ[0]**2) > 1.0:   # keep sd(w) where second moments stay␣
 ↪reliable
        return 1e10
    m = ar2_scores(θ_ar, simulate_sv(ρ[0], ρ[1], ρ[2], z_sim, η_sim)).mean(axis=0)
    return float(m @ W_ar @ m)

def implied_var(ρ):
    "Unconditional variance of y implied by the structural parameters."
    a, r_w, r_y = ρ
    return r_y**2 * np.exp(2 * r_w**2 / (1 - a**2))

print(f"variance of the data = {y_sv.var():.3f}\n")
print(f"{'start a':>9}{'start r_w':>11}{'â':>9}{'r̂_w':>9}{'r̂_y':>9}"
      f"{'implied var':>13}{'criterion':>12}")
for a_s in [0.3, 0.6, 0.9]:
    for rw_s in [0.15, 0.30]:
        u_s = np.array([np.log(a_s / (1 - a_s)), np.log(rw_s), np.log(0.6)])
        rr = minimize(criterion_ar_score, u_s, method='Nelder-Mead',
                      options=dict(maxiter=600, xatol=1e-5, fatol=1e-14))
        ρ_s = to_ρ(rr.x)
        print(f"{a_s:>9.2f}{rw_s:>11.2f}{ρ_s[0]:>9.4f}{ρ_s[1]:>9.4f}"
              f"{ρ_s[2]:>9.4f}{implied_var(ρ_s):>13.3f}{rr.fun:>12.2e}")

bad = np.array([np.log(0.95/0.05), np.log(0.20), np.log(2.0)])
print(f"\ncriterion when the variance of y is badly wrong: {criterion_ar_
```

```
↪score(bad):.2e}")
```

```
variance of the data = 0.544

  start a  start r_w        â     f̂_w     f̂_y  implied var   criterion
     0.30       0.15   0.4818   0.8250   0.3157        0.587    2.08e-03
     0.30       0.30   0.4818   0.8250   0.3157        0.587    2.08e-03
     0.60       0.15   0.4818   0.8250   0.3157        0.587    2.08e-03
     0.60       0.30   0.4818   0.8250   0.3157        0.587    2.08e-03
     0.90       0.15   0.8681   0.4965   0.2897        0.620    1.57e-03
     0.90       0.30   0.8643   0.5030   0.2891        0.618    1.56e-03

criterion when the variance of y is badly wrong: 6.95e+01
```

The estimates wander all over the parameter space while the criterion stays near its floor.

A criterion of order 10^{-3} at each of these very different parameter vectors, against a value near 70 when the variance of y is badly wrong, reveals a flat valley in the criterion surface.

In this stochastic volatility model z_t is independent of the volatility path, so y_t is serially uncorrelated, though not independent, at every (a, r_w, r_y).

Write γ_j for the jth autocovariance of y_t under the structural model.

With $\gamma_1 = \gamma_2 = 0$, the three components of the Gaussian AR(2) score have expectations

$$-\frac{\theta_1 \gamma_0}{\varsigma^2}, \qquad -\frac{\theta_2 \gamma_0}{\varsigma^2}, \qquad -\frac{1}{\varsigma} + \frac{(1 + \theta_1^2 + \theta_2^2)\,\gamma_0}{\varsigma^3}.$$

All three depend on ρ only through γ_0, the variance

$$\mathbb{V}(y_t) = r_y^2 \exp\left(\frac{2 r_w^2}{1 - a^2}\right),$$

so one equation must determine three parameters.

In population, any (a, r_w, r_y) holding that expression fixed fits the Gaussian AR(2) score equally well.

The criterion does not fall to zero.

The sample autocorrelations of the data are not exactly zero, so $\tilde{\theta}_1$ and $\tilde{\theta}_2$ are not zero either, and no value of γ_0 sets all three expectations to zero.

In a finite simulation, Monte Carlo error in the simulated autocovariances perturbs the floor of the valley, so the minimized criterion differs a little across rows.

The table bears this out: the implied variances cluster near the variance of the data while the three parameters themselves differ wildly.

The code bounds the standard deviation of the log volatility process, $r_w / \sqrt{1 - a^2}$, which holds the search inside the region where simulated second moments are reliable.

A log volatility with a very large standard deviation makes $\exp(2w_t)$ so heavy tailed that the optimizer can drive the simulated moments to zero by exploiting Monte Carlo error alone.

Fat tails and volatility clustering are fourth moment phenomena, and a Gaussian autoregression records neither.

The GARCH auxiliary model used earlier does record both, through the dependence of its log density on u_t^2 / σ_t^2, and it identifies all three parameters.

Gallant and Tauchen recommend this diagnostic: flat spots in the criterion signal a score generator too small for the structural model.

Exercise 38.15.4

The stochastic volatility example sets $a = 0.8$.

Rerun it with $a = 0.95$, $r_w = 0.2$ and $r_y = 0.65$.

Log volatility then has standard deviation $r_w/\sqrt{1-a^2} = 0.64$, and the population kurtosis of y_t is about 15.5.

a. Compute the first three autocorrelations of each component of the GARCH score at the data.

b. Estimate $\mathcal{J}^0$ two more ways.

First, use Gallant and Tauchen's heteroskedasticity and autocorrelation consistent estimate. It adds to the outer product the sample autocovariances of the score at lags $\tau = 1, \ldots, \ell$, with $\ell = [n^{1/5}]$, weighted by the Parzen kernel $w(\tau/\ell)$, where

$$w(x) = \begin{cases} 1 - 6x^2 + 6x^3 & 0 \le x \le 1/2, \\ 2(1-x)^3 & 1/2 < x \le 1. \end{cases}$$

Second, at the estimate $\hat{\rho}_n$, simulate 300 independent data sets of length n from the structural model, compute $\sqrt{n}$ times the average GARCH score at $\tilde{\theta}_n$ in each, and take the covariance matrix across data sets.

Compare the diagonals of the three estimates.

c. Re-estimate ρ with each estimate of $\mathcal{J}^0$, and compare the Theorem 1 standard errors.

The Theorem 1 formula $\left[M^\top W M\right]^{-1}$ assumes that the weighting matrix W equals $(\mathcal{J}^0)^{-1}$.

For the outer-product and Parzen estimators, also compute the sandwich standard errors

$$\left(1 + \frac{n}{N}\right)\frac{1}{n}(M^\top W M)^{-1} M^\top W \mathcal{J} W M (M^\top W M)^{-1},$$

with the simulated estimate of $\mathcal{J}^0$ in place of $\mathcal{J}$.

d. Explain what you find.

Solution

```
rng4 = np.random.default_rng(5)
a_4, r_w4, r_y4 = 0.95, 0.20, 0.65
y_4 = simulate_sv(a_4, r_w4, r_y4, rng4.standard_normal(T), rng4.standard_
 ↪normal(T))

r_4 = minimize(garch_negll, θ_start, args=(y_4,), method='Nelder-Mead',
               options=dict(maxiter=8000, xatol=1e-8, fatol=1e-8))
θ_4 = r_4.x
S_4 = garch_scores(θ_4, y_4)
I_op = S_4.T @ S_4 / T

S_c = S_4 - S_4.mean(axis=0)
names = ['μ', 'ω', 'α', 'β']
print("autocorrelations of the GARCH score at lags 1, 2, 3")
for j, name in enumerate(names):
    ac = [np.corrcoef(S_c[:-k, j], S_c[k:, j])[0, 1] for k in (1, 2, 3)]
    print(f"{name:>4}: " + "".join(f"{x:8.3f}" for x in ac))
```

```
    autocorrelations of the GARCH score at lags 1, 2, 3
       μ:    0.015  -0.020   0.015
       ω:    0.078   0.065   0.057
       α:    0.283   0.370   0.218
       β:    0.213   0.270   0.160
```

```
def parzen_hac(S):
    "Parzen-weighted estimate of I^0 with lag truncation [n^(1/5)]."
    n = S.shape[0]
    ℓ = int(n ** 0.2)
    V = S.T @ S / n
    for τ in range(1, ℓ + 1):
        x = τ / ℓ
        w = 1 - 6 * x**2 + 6 * x**3 if x <= 0.5 else 2 * (1 - x)**3
        Γ = S[τ:].T @ S[:-τ] / n
        V += w * (Γ + Γ.T)
    return V

z_4, η_4 = rng4.standard_normal(N), rng4.standard_normal(N)

def moments_4(ρ):
    return garch_scores(θ_4, simulate_sv(ρ[0], ρ[1], ρ[2], z_4, η_4)).mean(axis=0)

def estimate_4(W):
    "Score matching estimate and Theorem 1 standard errors for weighting matrix W."
    def crit(u):
        if np.max(np.abs(u)) > 25:
            return 1e10
        m = moments_4(to_ρ(u))
        return float(m @ W @ m)
    r_ = minimize(crit, u_0, method='Nelder-Mead',
                  options=dict(maxiter=400, xatol=1e-4, fatol=1e-10))
    ρ = to_ρ(r_.x)
    return ρ, std_errors(moments_4, ρ, W, T, N)

I_pz = parzen_hac(S_4)
ρ_op, se_op = estimate_4(np.linalg.inv(I_op))
ρ_pz, se_pz = estimate_4(np.linalg.inv(I_pz))

R_4 = 300
draws = np.empty((R_4, 4))
for i in range(R_4):
    y_i = simulate_sv(ρ_op[0], ρ_op[1], ρ_op[2],
                      rng4.standard_normal(T), rng4.standard_normal(T))
    draws[i] = np.sqrt(T) * garch_scores(θ_4, y_i).mean(axis=0)
I_sim = np.cov(draws.T)
ρ_sim, se_sim = estimate_4(np.linalg.inv(I_sim))

def sandwich_se(ρ, W, I):
    "Standard errors of the estimator with weighting matrix W when the score␣
 ↪variance is I."
    M = jacobian(moments_4, ρ)
    A = np.linalg.inv(M.T @ W @ M)
    V = A @ M.T @ W @ I @ W @ M @ A * (1 + T / N) / T
    return np.sqrt(np.diag(V))

sw_op = sandwich_se(ρ_op, np.linalg.inv(I_op), I_sim)
sw_pz = sandwich_se(ρ_pz, np.linalg.inv(I_pz), I_sim)
```

```
print(f"{'diagonal of':24}" + "".join(f"{n:>9}" for n in names))
for label, I in [('outer product', I_op), ('Parzen', I_pz), ('simulated data sets',
 ↪ I_sim)]:
    print(f"{'  ' + label:24}" + "".join(f"{x:9.3f}" for x in np.diag(I)))

print(f"\n{'':24}{'a':>9}{'r_w':>9}{'r_y':>9}")
print(f"{'true values':24}{a_4:9.4f}{r_w4:9.4f}{r_y4:9.4f}")
for label, ρ, se, sw in [('outer product', ρ_op, se_op, sw_op),
                         ('Parzen', ρ_pz, se_pz, sw_pz),
                         ('simulated', ρ_sim, se_sim, None)]:
    print(f"{label + ': estimate':24}" + "".join(f"{x:9.4f}" for x in ρ))
    print(f"{'   std error':24}" + "".join(f"{x:9.4f}" for x in se))
    if sw is not None:
        print(f"{'   sandwich std error':24}" + "".join(f"{x:9.4f}" for x in sw))
```

```
diagonal of                      μ        ω        α        β
  outer product              3.059    0.045    1.182    1.701
  Parzen                     3.100    0.051    1.855    2.437
  simulated data sets        4.292    0.172    2.682    3.547

                                 a      r_w      r_y
true values                 0.9500   0.2000   0.6500
outer product: estimate     0.9453   0.1946   0.5791
   std error                0.0225   0.0264   0.0938
   sandwich std error       0.0335   0.0319   0.1356
Parzen: estimate            0.9452   0.1947   0.5785
   std error                0.0276   0.0282   0.1136
   sandwich std error       0.0336   0.0319   0.1351
simulated: estimate         0.9447   0.1956   0.5750
   std error                0.0342   0.0316   0.1331
```

The α and β components of the score have autocorrelations between 0.16 and 0.37 at lags 1 to 3.

A Gaussian GARCH model is a misspecified score generator for these data, and its scores are not martingale differences.

The outer product omits the autocovariances that the long-run variance $\mathcal{J}^0$ contains.

The simulated estimate of $\mathcal{J}^0$ is two to four times the outer product for the ω, α and β components.

The Parzen estimate closes less than half of that gap.

With $n = 2000$ its lag truncation is $\ell = 4$, and the Parzen weights at lags 1, 2 and 3 are 0.72, 0.25 and 0.03, so it discards much of the persistent autocovariance.

With the outer product as W, the Theorem 1 formula understates the standard errors of a and r_y by about a third.

The Parzen estimate raises those two standard errors by about a fifth.

The sandwich standard errors of the outer-product and Parzen estimators nearly equal the standard errors of the estimator that uses the simulated estimate of $\mathcal{J}^0$.

The three estimators are almost equally precise, and the estimates barely move.

The damage from a poor weighting matrix falls on the reported precision.

Here the estimator itself loses almost nothing.

The simulated estimate is itself imprecise.

The variance of the GARCH score depends on eighth moments of y_t, and those are enormous at these parameter values.

The simulated estimate also requires a preliminary estimate of ρ, the requirement that the outer product avoided.

Gallant and Tauchen's advice addresses the source of the problem: choose a score generator that fits.

Section *Choosing the score generator* explains why: adding terms to a series expansion pushes its scores toward martingale differences, and then the outer product is valid.

Part V

Linear Programming

CHAPTER

THIRTYNINE

OPTIMAL TRANSPORT

39.1 Overview

The **transportation** or **optimal transport** problem is interesting both because of its many applications and because of its important role in the history of economic theory.

In this lecture, we describe the problem, tell how linear programming is a key tool for solving it, and then provide some examples.

We will provide other applications in followup lectures.

The optimal transport problem was studied in early work about linear programming, as summarized for example by [Dorfman *et al.*, 1958]. A modern reference about applications in economics is [Galichon, 2016].

Below, we show how to solve the optimal transport problem using several implementations of linear programming, including, in order,

1. the linprog solver from SciPy,
2. the linprog_simplex solver from QuantEcon and
3. the simplex-based solvers included in the Python Optimal Transport package.

```
!pip install --upgrade quantecon
!pip install --upgrade POT
```

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import linprog
from quantecon.optimize.linprog_simplex import linprog_simplex
import ot
from scipy.stats import betabinom
import networkx as nx
```

39.2 The Optimal Transport Problem

Suppose that m factories produce goods that must be sent to n locations.

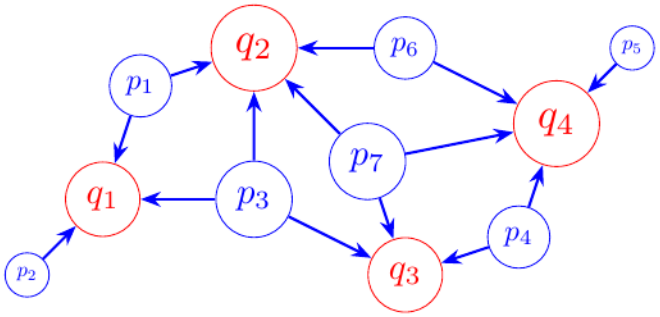
Let

- x_{ij} denote the quantity shipped from factory i to location j
- c_{ij} denote the cost of shipping one unit from factory i to location j
- p_i denote the capacity of factory i and q_j denote the amount required at location j.
- $i = 1, 2, \ldots, m$ and $j = 1, 2, \ldots, n$.

A planner wants to minimize total transportation costs subject to the following constraints:

- The amount shipped **from** each factory must equal its capacity.
- The amount shipped **to** each location must equal the quantity required there.

The figure below shows one visualization of this idea, when factories and target locations are distributed in the plane.



The size of the vertices in the figure are proportional to

- capacity, for the factories, and
- demand (amount required) for the target locations.

The arrows show one possible transport plan, which respects the constraints stated above.

The planner's problem can be expressed as the following constrained minimization problem:

$$\begin{aligned} \min_{x_{ij}} \ & \sum_{i=1}^{m} \sum_{j=1}^{n} c_{ij} x_{ij} \\ \text{subject to} \ & \sum_{j=1}^{n} x_{ij} = p_i, \quad i = 1, 2, \ldots, m \\ & \sum_{i=1}^{m} x_{ij} = q_j, \quad j = 1, 2, \ldots, n \\ & x_{ij} \geq 0 \end{aligned} \tag{39.1}$$

This is an **optimal transport problem** with

- mn decision variables, namely, the entries x_{ij} and
- $m + n$ constraints.

Summing the q_j's across all j's and the p_i's across all i's indicates that the total capacity of all the factories equals total requirements at all locations:

$$\sum_{j=1}^{n} q_j = \sum_{j=1}^{n} \sum_{i=1}^{m} x_{ij} = \sum_{i=1}^{m} \sum_{j=1}^{n} x_{ij} = \sum_{i=1}^{m} p_i \tag{39.2}$$

The presence of the restrictions in (39.2) will be the source of one redundancy in the complete set of restrictions that we describe below.

More about this later.

39.3 The Linear Programming Approach

In this section we discuss using using standard linear programming solvers to tackle the optimal transport problem.

39.3.1 Vectorizing a Matrix of Decision Variables

A *matrix* of decision variables x_{ij} appears in problem (39.1).

The SciPy function `linprog` expects to see a *vector* of decision variables.

This situation impels us to rewrite our problem in terms of a *vector* of decision variables.

Let

- X, C be $m \times n$ matrices with entries x_{ij}, c_{ij},
- p be m-dimensional vector with entries p_i,
- q be n-dimensional vector with entries q_j.

With $\mathbf{1}_n$ denoting the n-dimensional column vector $(1, 1, \ldots, 1)'$, our problem can now be expressed compactly as:

$$\begin{aligned} \min_X \ & \text{tr}(C'X) \\ \text{subject to } & X\,\mathbf{1}_n = p \\ & X'\,\mathbf{1}_m = q \\ & X \geq 0 \end{aligned}$$

We can convert the matrix X into a vector by stacking all of its columns into a column vector.

Doing this is called **vectorization**, an operation that we denote $\text{vec}(X)$.

Similarly, we convert the matrix C into an mn-dimensional vector $\text{vec}(C)$.

The objective function can be expressed as the inner product between $\text{vec}(C)$ and $\text{vec}(X)$:

$$\text{vec}(C)' \cdot \text{vec}(X).$$

To express the constraints in terms of $\text{vec}(X)$, we use a **Kronecker product** denoted by $\otimes$ and defined as follows.

Suppose A is an $m \times s$ matrix with entries (a_{ij}) and that B is an $n \times t$ matrix.

The **Kronecker product** of A and B is defined, in block matrix form, by

$$A \otimes B = \begin{bmatrix} a_{11}B & a_{12}B & \ldots & a_{1s}B \\ a_{21}B & a_{22}B & \ldots & a_{2s}B \\ & & \vdots & \\ a_{m1}B & a_{m2}B & \ldots & a_{ms}B \end{bmatrix}.$$

$A \otimes B$ is an $mn \times st$ matrix.

It has the property that for any $m \times n$ matrix X

$$\text{vec}(A'XB) = (B' \otimes A')\,\text{vec}(X). \tag{39.3}$$

We can now express our constraints in terms of $\text{vec}(X)$.

Let $A = \mathbf{I}'_m, B = \mathbf{1}_n$.

By equation (39.3)

$$X\,\mathbf{1}_n = \text{vec}(X\,\mathbf{1}_n) = \text{vec}(\mathbf{I}_m X\,\mathbf{1}_n) = (\mathbf{1}'_n \otimes \mathbf{I}_m)\,\text{vec}(X).$$

where $\mathbf{I}_m$ denotes the $m \times m$ identity matrix.

Constraint $X\,\mathbf{1}_n = p$ can now be written as:

$$(\mathbf{1}'_n \otimes \mathbf{I}_m)\,\text{vec}(X) = p.$$

Similarly, the constraint $X'\,\mathbf{1}_m = q$ can be rewriten as:

$$(\mathbf{I}_n \otimes \mathbf{1}'_m)\,\text{vec}(X) = q.$$

With $z := \text{vec}(X)$, our problem can now be expressed in terms of an mn-dimensional vector of decision variables:

$$\begin{aligned} \min_z \ & \text{vec}(C)'z \\ \text{subject to } \ & Az = b \\ & z \geq 0 \end{aligned} \tag{39.4}$$

where

$$A = \begin{bmatrix} \mathbf{1}'_n \otimes \mathbf{I}_m \\ \mathbf{I}_n \otimes \mathbf{1}'_m \end{bmatrix} \quad \text{and} \quad b = \begin{bmatrix} p \\ q \end{bmatrix}$$

39.3.2 An Application

We now provide an example that takes the form (39.4) that we'll solve by deploying the function `linprog`.

The table below provides numbers for the requirements vector q, the capacity vector p, and entries c_{ij} of the cost-of-shipping matrix C.

The numbers in the above table tell us to set $m = 3$, $n = 5$, and construct the following objects:

$$p = \begin{bmatrix} 50 \\ 100 \\ 150 \end{bmatrix}, \quad q = \begin{bmatrix} 25 \\ 115 \\ 60 \\ 30 \\ 70 \end{bmatrix} \quad \text{and} \quad C = \begin{bmatrix} 10 & 15 & 20 & 20 & 40 \\ 20 & 40 & 15 & 30 & 30 \\ 30 & 35 & 40 & 55 & 25 \end{bmatrix}.$$

Let's write Python code that sets up the problem and solves it.

```
# Define parameters
m = 3
n = 5

p = np.array([50.0, 100.0, 150.0])
```

(continues on next page)

(continued from previous page)

```
q = np.array([25.0, 115.0, 60.0, 30.0, 70.0])

C = np.array([[10.0, 15.0, 20.0, 20.0, 40.0],
              [20.0, 40.0, 15.0, 30.0, 30.0],
              [30.0, 35.0, 40.0, 55.0, 25.0]])

# Vectorize matrix C
C_vec = C.reshape((m*n, 1), order='F')

# Construct matrix A by Kronecker product
A1 = np.kron(np.ones((1, n)), np.identity(m))
A2 = np.kron(np.identity(n), np.ones((1, m)))
A = np.vstack([A1, A2])

# Construct vector b
b = np.hstack([p, q])

# Solve the primal problem
res = linprog(C_vec, A_eq=A, b_eq=b)

# Print results
print("message:", res.message)
print("nit:", res.nit)
print("fun:", res.fun)
print("z:", res.x)
print("X:", res.x.reshape((m,n), order='F'))
```

```
message: Optimization terminated successfully. (HiGHS Status 7: Optimal)
nit: 8
fun: 7225.0
z: [ 0. 10. 15. 50.  0. 65.  0. 60.  0.  0. 30.  0.  0.  0. 70.]
X: [[ 0. 50.  0.  0.  0.]
 [10.  0. 60. 30.  0.]
 [15. 65.  0.  0. 70.]]
```

Notice how, in the line `C_vec = C.reshape((m*n, 1), order='F')`, we are careful to vectorize using the flag `order='F'`.

This is consistent with converting C into a vector by stacking all of its columns into a column vector.

Here `'F'` stands for "Fortran", and we are using Fortran style column-major order.

(For an alternative approach, using Python's default row-major ordering, see this lecture by Alfred Galichon.)

Interpreting the solver behavior:

Looking at matrix A, we can see that it is rank deficient.

```
np.linalg.matrix_rank(A) < min(A.shape)
```

```
np.True_
```

This indicates that the linear program has been set up to include one or more redundant constraints.

Here, the source of the redundancy is the structure of restrictions (39.2).

Let's explore this further by printing out A and staring at it.

```
A
```

```
array([[1., 0., 0., 1., 0., 0., 1., 0., 0., 1., 0., 0., 1., 0., 0.],
       [0., 1., 0., 0., 1., 0., 0., 1., 0., 0., 1., 0., 0., 1., 0.],
       [0., 0., 1., 0., 0., 1., 0., 0., 1., 0., 0., 1., 0., 0., 1.],
       [1., 1., 1., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0.],
       [0., 0., 0., 1., 1., 1., 0., 0., 0., 0., 0., 0., 0., 0., 0.],
       [0., 0., 0., 0., 0., 0., 1., 1., 1., 0., 0., 0., 0., 0., 0.],
       [0., 0., 0., 0., 0., 0., 0., 0., 0., 1., 1., 1., 0., 0., 0.],
       [0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 1., 1., 1.]])
```

The singularity of A reflects that the first three constraints and the last five constraints both require that "total requirements equal total capacities" expressed in (39.2).

One equality constraint here is redundant.

Fortunately, SciPy's `linprog` function handles the redundant constraints automatically without explicitly warning about rank deficiency.

But we can drop one of the equality constraints, and use only 7 of them.

After doing this, we attain the same minimized cost.

However, we find a different transportation plan.

Though it is a different plan, it attains the same cost!

```
linprog(C_vec, A_eq=A[:-1], b_eq=b[:-1])
```

```
        message: Optimization terminated successfully. (HiGHS Status 7: Optimal)
        success: True
         status: 0
            fun: 7225.0
              x: [ 0.000e+00  1.000e+01 ...  0.000e+00  7.000e+01]
            nit: 8
          lower:  residual: [ 0.000e+00  1.000e+01 ...  0.000e+00
                              7.000e+01]
                 marginals: [ 0.000e+00  0.000e+00 ...  1.500e+01
                              0.000e+00]
          upper:  residual: [       inf        inf ...        inf
                                    inf]
                 marginals: [ 0.000e+00  0.000e+00 ...  0.000e+00
                              0.000e+00]
          eqlin:  residual: [ 0.000e+00  0.000e+00  0.000e+00  0.000e+00
                              0.000e+00  0.000e+00  0.000e+00]
                 marginals: [ 5.000e+00  1.500e+01  2.500e+01  5.000e+00
                              1.000e+01 -0.000e+00  1.500e+01]
        ineqlin:  residual: []
                 marginals: []
 mip_node_count: 0
 mip_dual_bound: 0.0
        mip_gap: 0.0
```

```
%time linprog(C_vec, A_eq=A[:-1], b_eq=b[:-1])
```

```
CPU times: user 2.59 ms, sys: 0 ns, total: 2.59 ms
Wall time: 2.24 ms
```

```
        message: Optimization terminated successfully. (HiGHS Status 7: Optimal)
        success: True
         status: 0
            fun: 7225.0
              x: [ 0.000e+00  1.000e+01 ...  0.000e+00  7.000e+01]
            nit: 8
          lower:  residual: [ 0.000e+00  1.000e+01 ...  0.000e+00
                              7.000e+01]
                 marginals: [ 0.000e+00  0.000e+00 ...  1.500e+01
                              0.000e+00]
          upper:  residual: [       inf        inf ...        inf
                                    inf]
                 marginals: [ 0.000e+00  0.000e+00 ...  0.000e+00
                              0.000e+00]
          eqlin:  residual: [ 0.000e+00  0.000e+00  0.000e+00  0.000e+00
                              0.000e+00  0.000e+00  0.000e+00]
                 marginals: [ 5.000e+00  1.500e+01  2.500e+01  5.000e+00
                              1.000e+01 -0.000e+00  1.500e+01]
        ineqlin:  residual: []
                 marginals: []
 mip_node_count: 0
 mip_dual_bound: 0.0
        mip_gap: 0.0
```

```
%time linprog(C_vec, A_eq=A, b_eq=b)
```

```
CPU times: user 2.57 ms, sys: 67 µs, total: 2.64 ms
Wall time: 2.28 ms
```

```
        message: Optimization terminated successfully. (HiGHS Status 7: Optimal)
        success: True
         status: 0
            fun: 7225.0
              x: [ 0.000e+00  1.000e+01 ...  0.000e+00  7.000e+01]
            nit: 8
          lower:  residual: [ 0.000e+00  1.000e+01 ...  0.000e+00
                              7.000e+01]
                 marginals: [ 0.000e+00  0.000e+00 ...  1.500e+01
                              0.000e+00]
          upper:  residual: [       inf        inf ...        inf
                                    inf]
                 marginals: [ 0.000e+00  0.000e+00 ...  0.000e+00
                              0.000e+00]
          eqlin:  residual: [ 0.000e+00  0.000e+00  0.000e+00  0.000e+00
                              0.000e+00  0.000e+00  0.000e+00  0.000e+00]
                 marginals: [ 1.000e+01  2.000e+01  3.000e+01 -0.000e+00
                              5.000e+00 -5.000e+00  1.000e+01 -5.000e+00]
        ineqlin:  residual: []
                 marginals: []
 mip_node_count: 0
 mip_dual_bound: 0.0
        mip_gap: 0.0
```

Evidently, it is slightly quicker to work with the system that removed a redundant constraint.

Let's drill down and do some more calculations to help us understand whether or not our finding **two** different optimal transport plans reflects our having dropped a redundant equality constraint.

Hint

It will turn out that dropping a redundant equality isn't really what mattered.

To verify our hint, we shall simply use **all** of the original equality constraints (including a redundant one), but we'll just shuffle the order of the constraints.

```
arr = np.arange(m+n)
```

```
sol_found = []
cost = []
rng = np.random.default_rng()

# simulate 1000 times
for i in range(1000):

    rng.shuffle(arr)
    res_shuffle = linprog(C_vec, A_eq=A[arr], b_eq=b[arr])

    # if find a new solution
    sol = tuple(res_shuffle.x)
    if sol not in sol_found:
        sol_found.append(sol)
        cost.append(res_shuffle.fun)
```

```
for i in range(len(sol_found)):
    print(f"transportation plan {i}: ", sol_found[i])
    print(f"     minimized cost {i}: ", cost[i])
```

```
transportation plan 0:  (np.float64(0.0), np.float64(10.0), np.float64(15.0), np.
↪float64(50.0), np.float64(0.0), np.float64(65.0), np.float64(0.0), np.float64(60.
↪0), np.float64(0.0), np.float64(0.0), np.float64(30.0), np.float64(0.0), np.
↪float64(0.0), np.float64(0.0), np.float64(70.0))
    minimized cost 0:  7225.0
```

Ah hah! As you can see, putting constraints in different orders in this case uncovers two optimal transportation plans that achieve the same minimized cost.

These are the same two plans computed earlier.

Next, we show that leaving out the first constraint "accidentally" leads to the initial plan that we computed.

```
linprog(C_vec, A_eq=A[1:], b_eq=b[1:])
```

```
        message: Optimization terminated successfully. (HiGHS Status 7: Optimal)
        success: True
         status: 0
            fun: 7225.0
              x: [ 0.000e+00  1.000e+01 ...  0.000e+00  7.000e+01]
            nit: 8
          lower:  residual: [ 0.000e+00  1.000e+01 ...  0.000e+00
                              7.000e+01]
                 marginals: [ 0.000e+00  0.000e+00 ...  1.500e+01
                              0.000e+00]
          upper:  residual: [       inf        inf ...        inf
```

(continues on next page)

(continued from previous page)

```
                                  inf]
                marginals: [ 0.000e+00  0.000e+00 ...  0.000e+00
                             0.000e+00]
          eqlin:  residual: [ 0.000e+00  0.000e+00  0.000e+00  0.000e+00
                             0.000e+00  0.000e+00  0.000e+00]
                marginals: [ 1.000e+01  2.000e+01  1.000e+01  1.500e+01
                             5.000e+00  2.000e+01  5.000e+00]
        ineqlin:  residual: []
                marginals: []
 mip_node_count: 0
 mip_dual_bound: 0.0
        mip_gap: 0.0
```

Let's compare this transport plan with

```
res.x
```

```
array([ 0., 10., 15., 50.,  0., 65.,  0., 60.,  0.,  0., 30.,  0.,  0.,
        0., 70.])
```

Here the matrix X contains entries x_{ij} that tell amounts shipped **from** factor $i = 1, 2, 3$ **to** location $j = 1, 2, \ldots, 5$.

The vector z evidently equals vec(X).

The minimized cost from the optimal transport plan is given by the fun variable.

39.3.3 Using a Just-in-Time Compiler

We can also solve optimal transportation problems using a powerful tool from QuantEcon, namely, `quantecon.optimize.linprog_simplex`.

While `scipy.optimize.linprog` uses the HiGHS solver by default, `quantecon.optimize.linprog_simplex` implements the simplex algorithm accelerated by using a just-in-time compiler shipped in the `numba` library.

As you will see very soon, by using `quantecon.optimize.linprog_simplex` the time required to solve an optimal transportation problem can be reduced significantly.

```
# construct matrices/vectors for linprog_simplex
c = C.flatten()

# Equality constraints
A_eq = np.zeros((m+n, m*n))
for i in range(m):
    for j in range(n):
        A_eq[i, i*n+j] = 1
        A_eq[m+j, i*n+j] = 1

b_eq = np.hstack([p, q])
```

Since `quantecon.optimize.linprog_simplex` does maximization instead of minimization, we need to put a negative sign before vector `c`.

```
res_qe = linprog_simplex(-c, A_eq=A_eq, b_eq=b_eq)
```

While the two LP solvers use different algorithms (HiGHS vs. simplex), both should find optimal solutions.

The solutions differs since there are multiple optimal solutions, but the objective values are the same

```
np.allclose(-res_qe.fun, res.fun)
```

```
True
```

```
res_qe.x.reshape((m, n), order='C')
```

```
array([[15., 35.,  0.,  0.,  0.],
       [10.,  0., 60., 30.,  0.],
       [ 0., 80.,  0.,  0., 70.]])
```

```
res.x.reshape((m, n), order='F')
```

```
array([[ 0., 50.,  0.,  0.,  0.],
       [10.,  0., 60., 30.,  0.],
       [15., 65.,  0.,  0., 70.]])
```

Let's do a speed comparison between `scipy.optimize.linprog` and `quantecon.optimize.linprog_simplex`.

```
# scipy.optimize.linprog
%time res = linprog(C_vec, A_eq=A[:-1, :], b_eq=b[:-1])
```

```
CPU times: user 2.35 ms, sys: 17 µs, total: 2.37 ms
Wall time: 2.18 ms
```

```
# quantecon.optimize.linprog_simplex
%time out = linprog_simplex(-c, A_eq=A_eq, b_eq=b_eq)
```

```
CPU times: user 60 µs, sys: 0 ns, total: 60 µs
Wall time: 62.9 µs
```

As you can see, the `quantecon.optimize.linprog_simplex` is much faster.

(Note however, that the SciPy version is probably more stable than the QuantEcon version, having been tested more extensively over a longer period of time.)

39.4 The Dual Problem

Let u, v denotes vectors of dual decision variables with entries $(u_i), (v_j)$.

The **dual** to **minimization** problem (39.1) is the **maximization** problem:

$$\max_{u_i, v_j} \sum_{i=1}^{m} p_i u_i + \sum_{j=1}^{n} q_j v_j$$
$$\text{subject to } u_i + v_j \leq c_{ij},\ i = 1, 2, \dots, m;\ j = 1, 2, \dots, n \tag{39.5}$$

The dual problem is also a linear programming problem.

It has $m + n$ dual variables and mn constraints.

Vectors u and v of **values** are attached to the first and the second sets of primal constraits, respectively.

Thus, u is attached to the constraints

- $(\mathbf{1}'_n \otimes \mathbf{I}_m)\,\text{vec}(X) = p$

and v is attached to constraints

- $(\mathbf{I}_n \otimes \mathbf{1}'_m)\,\text{vec}(X) = q$.

Components of the vectors u and v of per unit **values** are **shadow prices** of the quantities appearing on the right sides of those constraints.

We can write the dual problem as

$$\begin{aligned} &\max_{u_i, v_j} pu + qv \\ &\text{subject to } A' \begin{bmatrix} u \\ v \end{bmatrix} = \text{vec}(C) \end{aligned} \tag{39.6}$$

For the same numerical example described above, let's solve the dual problem.

```
# Solve the dual problem
res_dual = linprog(-b, A_ub=A.T, b_ub=C_vec,
                   bounds=[(None, None)]*(m+n))

#Print results
print("message:", res_dual.message)
print("nit:", res_dual.nit)
print("fun:", res_dual.fun)
print("u:", res_dual.x[:m])
print("v:", res_dual.x[-n:])
```

```
message: Optimization terminated successfully. (HiGHS Status 7: Optimal)
nit: 9
fun: -7225.0
u: [-20. -10.   0.]
v: [30. 35. 25. 40. 25.]
```

`quantecon.optimize.linprog_simplex` computes and returns the dual variables alongside the primal solution.

The dual variables (shadow prices) can be extracted directly from the primal solution:

```
# The dual variables are returned by linprog_simplex
print("Dual variables from linprog_simplex:")
print("u:", -res_qe.lambd[:m])
print("v:", -res_qe.lambd[m:])
```

```
Dual variables from linprog_simplex:
u: [-20. -10.  -0.]
v: [30. 35. 25. 40. 25.]
```

We can verify these match the dual solution from SciPy:

```
print("Dual variables from SciPy linprog:")
print("u:", res_dual.x[:m])
print("v:", res_dual.x[-n:])
```

```
Dual variables from SciPy linprog:
u: [-20. -10.   0.]
v: [30. 35. 25. 40. 25.]
```

39.4.1 Interpretation of dual problem

By **strong duality** (please see this lecture Linear Programming), we know that:

$$\sum_{i=1}^{m}\sum_{j=1}^{n} c_{ij}x_{ij} = \sum_{i=1}^{m} p_i u_i + \sum_{j=1}^{n} q_j v_j$$

One unit more capacity in factory i, i.e. p_i, results in u_i more transportation costs.

Thus, u_i describes the cost of shipping one unit **from** factory i.

Call this the ship-out cost of one unit shipped from factory i.

Similarly, v_j is the cost of shipping one unit **to** location j.

Call this the ship-in cost of one unit to location j.

Strong duality implies that total transprotation costs equals total ship-out costs **plus** total ship-in costs.

It is reasonable that, for one unit of a product, ship-out cost u_i **plus** ship-in cost v_j should equal transportation cost c_{ij}.

This equality is assured by **complementary slackness** conditions that state that whenever $x_{ij} > 0$, meaning that there are positive shipments from factory i to location j, it must be true that $u_i + v_j = c_{ij}$.

39.5 The Python Optimal Transport Package

There is an excellent Python package for optimal transport that simplifies some of the steps we took above.

In particular, the package takes care of the vectorization steps before passing the data out to a linear programming routine.

(That said, the discussion provided above on vectorization remains important, since we want to understand what happens under the hood.)

39.5.1 Replicating Previous Results

The following line of code solves the example application discussed above using linear programming.

```
X = ot.emd(p, q, C)
X
```

```
array([[15., 35.,  0.,  0.,  0.],
       [10.,  0., 60., 30.,  0.],
       [ 0., 80.,  0.,  0., 70.]])
```

Sure enough, we have the same solution and the same cost

```
total_cost = np.vdot(X, C)
total_cost
```

```
np.float64(7225.0)
```

Here we use np.vdot for the trace inner product of X and C

39.5.2 A Larger Application

Now let's try using the same package on a slightly larger application.

The application has the same interpretation as above but we will also give each node (i.e., vertex) a location in the plane.

This will allow us to plot the resulting transport plan as edges in a graph.

The following class defines a node by

- its location $(x, y) \in \mathbb{R}^2$,
- its group (factory or location, denoted by `p` or `q`) and
- its mass (e.g., p_i or q_j).

```
class Node:

    def __init__(self, x, y, mass, group, name):

        self.x, self.y = x, y
        self.mass, self.group = mass, group
        self.name = name
```

Next we write a function that repeatedly calls the class above to build instances.

It allocates to the nodes it creates their location, mass, and group.

Locations are assigned randomly.

```
def build_nodes_of_one_type(group='p', n=100, seed=123):

    nodes = []
    rng = np.random.default_rng(seed)

    for i in range(n):

        if group == 'p':
            m = 1/n
            x = rng.uniform(-2, 2)
            y = rng.uniform(-2, 2)
        else:
            m = betabinom.pmf(i, n-1, 2, 2)
            x = 0.6 * rng.uniform(-1.5, 1.5)
            y = 0.6 * rng.uniform(-1.5, 1.5)

        name = group + str(i)
        nodes.append(Node(x, y, m, group, name))

    return nodes
```

Now we build two lists of nodes, each one containing one type (factories or locations)

```
n_p = 32
n_q = 32
p_list = build_nodes_of_one_type(group='p', n=n_p)
q_list = build_nodes_of_one_type(group='q', n=n_q)

p_probs = [p.mass for p in p_list]
q_probs = [q.mass for q in q_list]
```

For the cost matrix C, we use the Euclidean distance between each factory and location.

```
c = np.empty((n_p, n_q))
for i in range(n_p):
    for j in range(n_q):
        x0, y0 = p_list[i].x, p_list[i].y
        x1, y1 = q_list[j].x, q_list[j].y
        c[i, j] = np.sqrt((x0-x1)**2 + (y0-y1)**2)
```

Now we are ready to apply the solver

```
%time pi = ot.emd(p_probs, q_probs, c)
```

```
CPU times: user 527 µs, sys: 25 µs, total: 552 µs
Wall time: 368 µs
```

Finally, let's plot the results using `networkx`.

In the plot below,

- node size is proportional to probability mass
- an edge (arrow) from i to j is drawn when a positive transfer is made from i to j under the optimal transport plan.

```
g = nx.DiGraph()
g.add_nodes_from([p.name for p in p_list])
g.add_nodes_from([q.name for q in q_list])

for i in range(n_p):
    for j in range(n_q):
        if pi[i, j] > 0:
            g.add_edge(p_list[i].name, q_list[j].name, weight=pi[i, j])

node_pos_dict={}
for p in p_list:
    node_pos_dict[p.name] = (p.x, p.y)

for q in q_list:
    node_pos_dict[q.name] = (q.x, q.y)

node_color_list = []
node_size_list = []
scale = 8_000
for p in p_list:
    node_color_list.append('blue')
    node_size_list.append(p.mass * scale)
for q in q_list:
    node_color_list.append('red')
    node_size_list.append(q.mass * scale)


fig, ax = plt.subplots(figsize=(7, 10))
plt.axis('off')

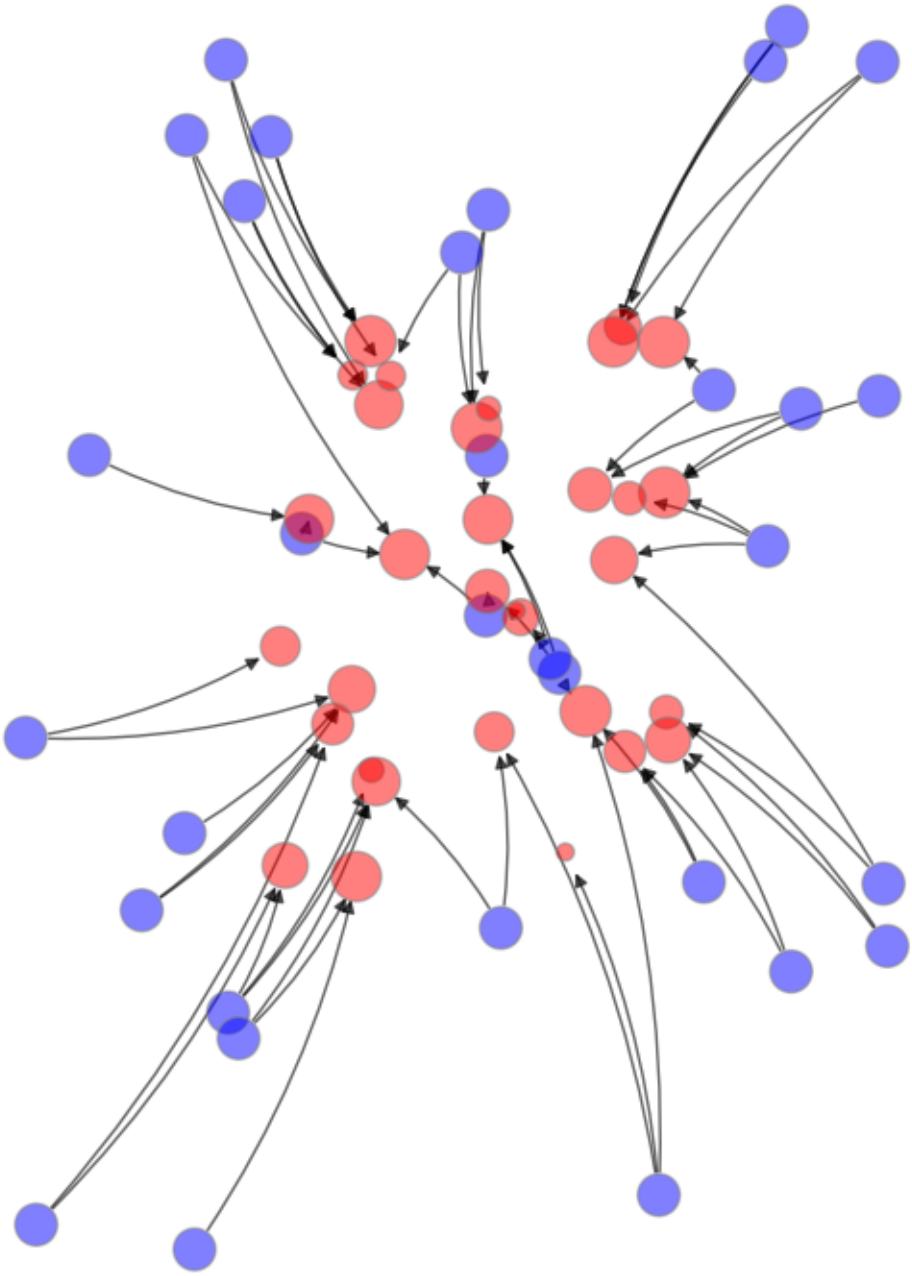
nx.draw_networkx_nodes(g,
                       node_pos_dict,
                       node_color=node_color_list,
                       node_size=node_size_list,
                       edgecolors='grey',
```

(continues on next page)

(continued from previous page)

```
                        linewidths=1,
                        alpha=0.5,
                        ax=ax)

nx.draw_networkx_edges(g,
                        node_pos_dict,
                        arrows=True,
                        connectionstyle='arc3,rad=0.1',
                        alpha=0.6)
plt.show()
```



CHAPTER

FORTY

VON NEUMANN GROWTH MODEL (AND A GENERALIZATION)

Contents

- *Von Neumann Growth Model (and a Generalization)*
  - *Notation*
  - *Model Ingredients and Assumptions*
  - *Dynamic Interpretation*
  - *Duality*
  - *Interpretation as Two-player Zero-sum Game*

This lecture uses the class `Neumann` to calculate key objects of a linear growth model of John von Neumann [von Neumann, 1937] that was generalized by Kemeny, Morgenstern and Thompson [Kemeny *et al.*, 1956].

Objects of interest are the maximal expansion rate (α), the interest factor (β), the optimal intensities (x), and prices (p).

In addition to watching how the towering mind of John von Neumann formulated an equilibrium model of price and quantity vectors in balanced growth, this lecture shows how fruitfully to employ the following important tools:

- a zero-sum two-player game
- linear programming
- the Perron-Frobenius theorem

We'll begin with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import fsolve, linprog
from textwrap import dedent

np.set_printoptions(precision=2)
```

The code below provides the `Neumann` class

```
class Neumann:

    """
    This class describes the Generalized von Neumann growth model as it was
    discussed in Kemeny et al. (1956, ECTA) and Gale (1960, Chapter 9.5):
```

(continues on next page)

(continued from previous page)

```
    Let:
    n ... number of goods
    m ... number of activities
    A ... input matrix is m-by-n
        a_{i,j} - amount of good j consumed by activity i
    B ... output matrix is m-by-n
        b_{i,j} - amount of good j produced by activity i

    x ... intensity vector (m-vector) with non-negative entries
        x'B - the vector of goods produced
        x'A - the vector of goods consumed
    p ... price vector (n-vector) with non-negative entries
        Bp - the revenue vector for every activity
        Ap - the cost of each activity

    Both A and B have non-negative entries. Moreover, we assume that
    (1) Assumption I (every good which is consumed is also produced):
        for all j, b_{.,j} > 0, i.e. at least one entry is strictly positive
    (2) Assumption II (no free lunch):
        for all i, a_{i,.} > 0, i.e. at least one entry is strictly positive

    Parameters
    ----------
    A : array_like or scalar(float)
        Part of the state transition equation.  It should be `n x n`
    B : array_like or scalar(float)
        Part of the state transition equation.  It should be `n x k`
    """

    def __init__(self, A, B):

        self.A, self.B = list(map(self.convert, (A, B)))
        self.m, self.n = self.A.shape

        # Check if (A, B) satisfy the basic assumptions
        assert self.A.shape == self.B.shape, 'The input and output matrices \
              must have the same dimensions!'
        assert (self.A >= 0).all() and (self.B >= 0).all(), 'The input and \
              output matrices must have only non-negative entries!'

        # (1) Check whether Assumption I is satisfied:
        if (np.sum(B, 0) <= 0).any():
            self.AI = False
        else:
            self.AI = True

        # (2) Check whether Assumption II is satisfied:
        if (np.sum(A, 1) <= 0).any():
            self.AII = False
        else:
            self.AII = True

    def __repr__(self):
        return self.__str__()

    def __str__(self):
```

(continues on next page)

(continued from previous page)

```
        me = """
        Generalized von Neumann expanding model:
          - number of goods          : {n}
          - number of activities     : {m}

        Assumptions:
          - AI:  every column of B has a positive entry    : {AI}
          - AII: every row of A has a positive entry       : {AII}

        """
        # Irreducible                                       : {irr}
        return dedent(me.format(n=self.n, m=self.m,
                                AI=self.AI, AII=self.AII))

    def convert(self, x):
        """
        Convert array_like objects (lists of lists, floats, etc.) into
        well-formed 2D NumPy arrays
        """
        return np.atleast_2d(np.asarray(x))


    def bounds(self):
        """
        Calculate the trivial upper and lower bounds for alpha (expansion rate)
        and beta (interest factor). See the proof of Theorem 9.8 in Gale (1960)
        """

        n, m = self.n, self.m
        A, B = self.A, self.B

        f = lambda α: ((B - α * A) @ np.ones((n, 1))).max()
        g = lambda β: (np.ones((1, m)) @ (B - β * A)).min()

        UB = fsolve(f, 1).item()  # Upper bound for α, β
        LB = fsolve(g, 2).item()  # Lower bound for α, β

        return LB, UB


    def zerosum(self, γ, dual=False):
        """
        Given gamma, calculate the value and optimal strategies of a
        two-player zero-sum game given by the matrix

                M(gamma) = B - gamma * A

        Row player maximizing, column player minimizing

        Zero-sum game as an LP (primal --> α)

            max (0', 1) @ (x', v)
            subject to
            [-M', ones(n, 1)] @ (x', v)' <= 0
            (x', v) @ (ones(m, 1), 0) = 1
            (x', v) >= (0', -inf)
```

(continues on next page)

(continued from previous page)

```
        Zero-sum game as an LP (dual --> beta)

            min (0', 1) @ (p', u)
            subject to
            [M, -ones(m, 1)] @ (p', u)' <= 0
            (p', u) @ (ones(n, 1), 0) = 1
            (p', u) >= (0', -inf)

        Outputs:
        --------
        value: scalar
            value of the zero-sum game

        strategy: vector
            if dual = False, it is the intensity vector,
            if dual = True, it is the price vector
        """

        A, B, n, m = self.A, self.B, self.n, self.m
        M = B - y * A

        if dual == False:
            # Solve the primal LP (for details see the description)
            # (1) Define the problem for v as a maximization (linprog minimizes)
            c = np.hstack([np.zeros(m), -1])

            # (2) Add constraints :
            # ... non-negativity constraints
            bounds = tuple(m * [(0, None)] + [(None, None)])
            # ... inequality constraints
            A_iq = np.hstack([-M.T, np.ones((n, 1))])
            b_iq = np.zeros((n, 1))
            # ... normalization
            A_eq = np.hstack([np.ones(m), 0]).reshape(1, m + 1)
            b_eq = 1

            res = linprog(c, A_ub=A_iq, b_ub=b_iq, A_eq=A_eq, b_eq=b_eq,
                          bounds=bounds)

        else:
            # Solve the dual LP (for details see the description)
            # (1) Define the problem for v as a maximization (linprog minimizes)
            c = np.hstack([np.zeros(n), 1])

            # (2) Add constraints :
            # ... non-negativity constraints
            bounds = tuple(n * [(0, None)] + [(None, None)])
            # ... inequality constraints
            A_iq = np.hstack([M, -np.ones((m, 1))])
            b_iq = np.zeros((m, 1))
            # ... normalization
            A_eq = np.hstack([np.ones(n), 0]).reshape(1, n + 1)
            b_eq = 1

            res = linprog(c, A_ub=A_iq, b_ub=b_iq, A_eq=A_eq, b_eq=b_eq,
                          bounds=bounds)
```

(continues on next page)

(continued from previous page)

```
        if res.status != 0:
            print(res.message)

        # Pull out the required quantities
        value = res.x[-1]
        strategy = res.x[:-1]

        return value, strategy


    def expansion(self, tol=1e-8, maxit=1000):
        """
        The algorithm used here is described in Hamburger-Thompson-Weil
        (1967, ECTA). It is based on a simple bisection argument and utilizes
        the idea that for a given γ (= α or β), the matrix "M = B - γ * A"
        defines a two-player zero-sum game, where the optimal strategies are
        the (normalized) intensity and price vector.

        Outputs:
        --------
        alpha: scalar
            optimal expansion rate
        """

        LB, UB = self.bounds()

        for iter in range(maxit):

            γ = (LB + UB) / 2
            ZS = self.zerosum(γ=γ)
            V = ZS[0]      # value of the game with γ

            if V >= 0:
                LB = γ
            else:
                UB = γ

            if abs(UB - LB) < tol:
                γ = (UB + LB) / 2
                x = self.zerosum(γ=γ)[1]
                p = self.zerosum(γ=γ, dual=True)[1]
                break

        return γ, x, p

    def interest(self, tol=1e-8, maxit=1000):
        """
        The algorithm used here is described in Hamburger-Thompson-Weil
        (1967, ECTA). It is based on a simple bisection argument and utilizes
        the idea that for a given gamma (= alpha or beta),
        the matrix "M = B - γ * A" defines a two-player zero-sum game,
        where the optimal strategies are the (normalized) intensity and price
        vector

        Outputs:
        --------
        beta: scalar
```

(continues on next page)

(continued from previous page)

```
            optimal interest rate
        """

        LB, UB = self.bounds()

        for iter in range(maxit):
            y = (LB + UB) / 2
            ZS = self.zerosum(y=y, dual=True)
            V = ZS[0]

            if V > 0:
                LB = y
            else:
                UB = y

            if abs(UB - LB) < tol:
                y = (UB + LB) / 2
                p = self.zerosum(y=y, dual=True)[1]
                x = self.zerosum(y=y)[1]
                break

        return y, x, p
```

40.1 Notation

We use the following notation.

0 denotes a vector of zeros.

We call an n-vector **positive** and write $x \gg \mathbf{0}$ if $x_i > 0$ for all $i = 1, 2, \dots, n$.

We call a vector **non-negative** and write $x \geq \mathbf{0}$ if $x_i \geq 0$ for all $i = 1, 2, \dots, n$.

We call a vector **semi-positive** written $x > \mathbf{0}$ if $x \geq \mathbf{0}$ and $x \neq \mathbf{0}$.

For two conformable vectors x and y, $x \gg y$, $x \geq y$ and $x > y$ mean $x - y \gg \mathbf{0}$, $x - y \geq \mathbf{0}$, and $x - y > \mathbf{0}$, respectively.

We let all vectors in this lecture be column vectors; $x^\top$ denotes the transpose of x (i.e., a row vector).

Let ι_n denote a column vector composed of n ones, i.e. $\iota_n = (1, 1, \dots, 1)^\top$.

Let e^i denote a vector (of arbitrary size) containing zeros except for the i th position where it is one.

We denote matrices by capital letters. For an arbitrary matrix A, $a_{i,j}$ represents the entry in its i th row and j th column. $a_{.j}$ and $a_{i.}$ denote the j th column and i th row of A, respectively.

40.2 Model Ingredients and Assumptions

A pair (A, B) of $m \times n$ non-negative matrices defines an economy.

- m is the number of **activities** (or sectors)
- n is the number of **goods** (produced and/or consumed)
- A is called the **input matrix**; $a_{i,j}$ denotes the amount of good j consumed by activity i
- B is called the **output matrix**; $b_{i,j}$ represents the amount of good j produced by activity i

Two key assumptions restrict economy (A, B):

Assumption 40.2.1 (every good that is consumed is also produced)

$$b_{.,j} > \mathbf{0} \quad \forall j = 1, 2, \ldots, n$$

Assumption 40.2.2 (no free lunch)

$$a_{i,.} > \mathbf{0} \quad \forall i = 1, 2, \ldots, m$$

A semi-positive **intensity** m-vector x denotes levels at which activities are operated.

Therefore,

- vector $x^\top A$ gives the total amount of **goods used in production**
- vector $x^\top B$ gives **total outputs**

An economy (A, B) is said to be **productive**, if there exists a non-negative intensity vector $x \geq 0$ such that $x^\top B > x^\top A$.

The semi-positive n-vector p contains prices assigned to the n goods.

The p vector implies **cost** and **revenue** vectors

- the vector Ap tells **costs** of the vector of activities
- the vector Bp tells **revenues** from the vector of activities

Satisfaction of a property of an input-output pair (A, B) called **irreducibility** (or indecomposability) determines whether an economy can be decomposed into multiple "sub-economies".

Definition 40.2.1

For an economy (A, B), the set of goods $S \subset \{1, 2, \ldots, n\}$ is called an **independent subset** if it is possible to produce every good in S without consuming goods from outside S.

Formally, the set S is independent if $\exists T \subset \{1, 2, \ldots, m\}$ (a subset of activities) such that $a_{i,j} = 0$, $\forall i \in T$ and $j \in S^c$ and for all $j \in S$, $\exists i \in T$ for which $b_{i,j} > 0$.

The economy is **irreducible** if there are no proper independent subsets.

We study two examples, both in Chapter 9.6 of Gale [Gale, 1989]

```
# (1) Irreducible (A, B) example: α_0 = β_0
A1 = np.array([[0, 1, 0, 0],
               [1, 0, 0, 1],
               [0, 0, 1, 0]])

B1 = np.array([[1, 0, 0, 0],
               [0, 0, 2, 0],
               [0, 1, 0, 1]])

# (2) Reducible (A, B) example: β_0 < α_0
A2 = np.array([[0, 1, 0, 0, 0, 0],
               [1, 0, 1, 0, 0, 0],
               [0, 0, 0, 1, 0, 0],
               [0, 0, 1, 0, 0, 1],
               [0, 0, 0, 0, 1, 0]])

B2 = np.array([[1, 0, 0, 1, 0, 0],
               [0, 1, 0, 0, 0, 0],
               [0, 0, 1, 0, 0, 0],
               [0, 0, 0, 0, 2, 0],
               [0, 0, 0, 1, 0, 1]])
```

The following code sets up our first Neumann economy or `Neumann` instance

```
n1 = Neumann(A1, B1)
n1
```

```
Generalized von Neumann expanding model:
  - number of goods          : 4
  - number of activities     : 3

Assumptions:
  - AI:  every column of B has a positive entry    : True
  - AII: every row of A has a positive entry       : True
```

Here is a second instance of a Neumann economy

```
n2 = Neumann(A2, B2)
n2
```

```
Generalized von Neumann expanding model:
  - number of goods          : 6
  - number of activities     : 5

Assumptions:
  - AI:  every column of B has a positive entry    : True
  - AII: every row of A has a positive entry       : True
```

40.3 Dynamic Interpretation

Attach a time index t to the preceding objects, regard an economy as a dynamic system, and study sequences

$$\{(A_t, B_t)\}_{t\geq 0}, \qquad \{x_t\}_{t\geq 0}, \qquad \{p_t\}_{t\geq 0}$$

An interesting special case holds the technology process constant and investigates the dynamics of quantities and prices only.

Accordingly, in the rest of this lecture, we assume that $(A_t, B_t) = (A, B)$ for all $t \geq 0$.

A crucial element of the dynamic interpretation involves the timing of production.

We assume that production (consumption of inputs) takes place in period t, while the consequent output materializes in period $t + 1$, i.e., consumption of $x_t^\top A$ in period t results in $x_t^\top B$ amounts of output in period $t + 1$.

These timing conventions imply the following feasibility condition:

$$x_t^\top B \geq x_{t+1}^\top A \qquad \forall t \geq 1$$

which asserts that no more goods can be used today than were produced yesterday.

Accordingly, Ap_t tells the costs of production in period t and Bp_t tells revenues in period $t + 1$.

40.3.1 Balanced Growth

We follow John von Neumann in studying "balanced growth".

Let $./$ denote an elementwise division of one vector by another and let $\alpha > 0$ be a scalar.

Then **balanced growth** is a situation in which

$$x_{t+1}./x_t = \alpha, \quad \forall t \geq 0$$

With balanced growth, the law of motion of x is evidently $x_{t+1} = \alpha x_t$ and so we can rewrite the feasibility constraint as

$$x_t^\top B \geq \alpha x_t^\top A \qquad \forall t$$

In the same spirit, define $\beta \in \mathbb{R}$ as the **interest factor** per unit of time.

We assume that it is always possible to earn a gross return equal to the constant interest factor β by investing "outside the model".

Under this assumption about outside investment opportunities, a no-arbitrage condition gives rise to the following (no profit) restriction on the price sequence:

$$\beta A p_t \geq B p_t \qquad \forall t$$

This says that production cannot yield a return greater than that offered by the outside investment opportunity (here we compare values in period $t + 1$).

The balanced growth assumption allows us to drop time subscripts and conduct an analysis purely in terms of a time-invariant growth rate α and interest factor β.

40.4 Duality

Two problems are connected by a remarkable dual relationship between technological and valuation characteristics of the economy:

Definition 40.4.1

The **technological expansion problem** (TEP) for the economy (A, B) is to find a semi-positive m-vector $x > 0$ and a number $\alpha \in \mathbb{R}$ that satisfy

$$\begin{aligned} \max_{\alpha} \quad & \alpha \\ \text{s.t.} \quad & x^\top B \geq \alpha x^\top A \end{aligned}$$

Theorem 9.3 of David Gale's book [Gale, 1989] asserts that if *Assumption 40.2.1* and *Assumption 40.2.2* are both satisfied, then a maximum value of α exists and that it is positive.

The maximal value is called the *technological expansion rate* and is denoted by α_0. The associated intensity vector x_0 is the *optimal intensity vector*.

Definition 40.4.2

The economic expansion problem (EEP) for (A, B) is to find a semi-positive n-vector $p > 0$ and a number $\beta \in \mathbb{R}$ that satisfy

$$\begin{aligned} \min_{\beta} \quad & \beta \\ \text{s.t.} \quad & Bp \leq \beta Ap \end{aligned}$$

Assumption 40.2.1 and *Assumption 40.2.2* imply existence of a minimum value $\beta_0 > 0$ called the *economic expansion rate*.

The corresponding price vector p_0 is the *optimal price vector*.

Because the criterion functions in the *technological expansion* problem and the *economical expansion problem* are both linearly homogeneous, the optimality of x_0 and p_0 are defined only up to a positive scale factor.

For convenience (and to emphasize a close connection to zero-sum games), we normalize both vectors x_0 and p_0 to have unit length.

A standard duality argument (see Lemma 9.4. in (Gale, 1960) [Gale, 1989]) implies that under *Assumption 40.2.1* and *Assumption 40.2.2*, $\beta_0 \leq \alpha_0$.

But to deduce that $\beta_0 \geq \alpha_0$, *Assumption 40.2.1* and *Assumption 40.2.2* are not sufficient.

Therefore, von Neumann [von Neumann, 1937] went on to prove the following remarkable "duality" result that connects TEP and EEP.

Theorem 40.4.1 (von Neumann)

If the economy (A, B) satisfies *Assumption 40.2.1* and *Assumption 40.2.2*, then there exist (γ^*, x_0, p_0), where $\gamma^* \in [\beta_0, \alpha_0] \subset \mathbb{R}$, $x_0 > 0$ is an m-vector, $p_0 > 0$ is an n-vector, and the following arbitrage conditions hold

$$\begin{aligned} x_0^\top B &\geq \gamma^* x_0^\top A \\ Bp_0 &\leq \gamma^* Ap_0 \\ x_0^\top (B - \gamma^* A)\, p_0 &= 0 \end{aligned}$$

Proof. (Sketch)

Assumption 40.2.1 and *Assumption 40.2.2* imply that there exist (α_0, x_0) and (β_0, p_0) that solve the TEP and EEP, respectively.

If $\gamma^* > \alpha_0$, then by definition of α_0, there cannot exist a semi-positive x that satisfies $x^\top B \geq \gamma^* x^\top A$.

Similarly, if $\gamma^* < \beta_0$, there is no semi-positive p for which $Bp \leq \gamma^* Ap$. Let $\gamma^* \in [\beta_0, \alpha_0]$, then $x_0^\top B \geq \alpha_0 x_0^\top A \geq \gamma^* x_0^\top A$.

Moreover, $Bp_0 \leq \beta_0 Ap_0 \leq \gamma^* Ap_0$. These two inequalities imply $x_0 \left(B - \gamma^* A\right) p_0 = 0$.

Here the constant γ^* is both an expansion factor and an interest factor (not necessarily optimal).

We have already encountered and discussed the first two inequalities that represent feasibility and no-profit conditions.

Moreover, the equality $x_0^\top \left(B - \gamma^* A\right) p_0 = 0$ concisely expresses the requirements that if any good grows at a rate larger than γ^* (i.e., if it is *oversupplied*), then its price must be zero; and that if any activity provides negative profit, it must be unused.

Therefore, the conditions stated in *Theorem 40.4.1* encode all equilibrium conditions.

So *Theorem 40.4.1* essentially states that under *Assumption 40.2.1* and *Assumption 40.2.2* there always exists an equilibrium (γ^*, x_0, p_0) with balanced growth.

Note that *Theorem 40.4.1* is silent about uniqueness of the equilibrium. In fact, it does not rule out (trivial) cases with $x_0^\top B p_0 = 0$ so that nothing of value is produced.

To exclude such uninteresting cases, Kemeny, Morgenstern and Thompson [Kemeny *et al.*, 1956] add an extra requirement

$$x_0^\top B p_0 > 0$$

and call the associated equilibria *economic solutions*.

They show that this extra condition does not affect the existence result, while it significantly reduces the number of (relevant) solutions.

40.5 Interpretation as Two-player Zero-sum Game

To compute the equilibrium (γ^*, x_0, p_0), we follow the algorithm proposed by Hamburger, Thompson and Weil (1967), building on the key insight that an equilibrium (with balanced growth) can be solved as a particular two-player zero-sum game.

First, we introduce some notation.

Consider the $m \times n$ matrix C as a payoff matrix, with the entries representing payoffs from the **minimizing** column player to the **maximizing** row player and assume that the players can use mixed strategies. Thus,

- the row player chooses the m-vector $x > \mathbf{0}$ subject to $\iota_m^\top x = 1$
- the column player chooses the n-vector $p > \mathbf{0}$ subject to $\iota_n^\top p = 1$.

Definition 40.5.1

The $m \times n$ matrix game C has the *solution* $(x^*, p^*, V(C))$ in mixed strategies if

$$(x^*)^\top C e^j \geq V(C) \quad \forall j \in \{1, \dots, n\} \qquad \text{and} \qquad (e^i)^\top C p^* \leq V(C) \quad \forall i \in \{1, \dots, m\}$$

The number $V(C)$ is called the **value** of the game.

From the above definition, it is clear that the value $V(C)$ has two alternative interpretations:

- by playing the appropriate mixed strategy, the maximizing player can assure himself at least $V(C)$ (no matter what the column player chooses)
- by playing the appropriate mixed strategy, the minimizing player can make sure that the maximizing player will not get more than $V(C)$ (irrespective of what is the maximizing player's choice)

A famous theorem of Nash (1951) tells us that there always exists a mixed strategy Nash equilibrium for any *finite* two-player zero-sum game.

Moreover, von Neumann's Minmax Theorem [von Neumann, 1928] implies that

$$V(C) = \max_x \min_p \; x^\top C p = \min_p \max_x \; x^\top C p = (x^*)^\top C p^*$$

40.5.1 Connection with Linear Programming (LP)

Nash equilibria of a finite two-player zero-sum game solve a linear programming problem.

To see this, we introduce the following notation

- For a fixed x, let v be the value of the minimization problem: $v \equiv \min_p x^\top C p = \min_j x^\top C e^j$
- For a fixed p, let u be the value of the maximization problem: $u \equiv \max_x x^\top C p = \max_i (e^i)^\top C p$

Then the *max-min problem* (the game from the maximizing player's point of view) can be written as the *primal* LP

$$\begin{aligned} V(C) = \max \quad & v \\ \text{s.t.} \quad & v\iota_n^\top \leq x^\top C \\ & x \geq \mathbf{0} \\ & \iota_n^\top x = 1 \end{aligned}$$

while the *min-max problem* (the game from the minimizing player's point of view) is the *dual* LP

$$\begin{aligned} V(C) = \min \quad & u \\ \text{s.t.} \quad & u\iota_m \geq Cp \\ & p \geq \mathbf{0} \\ & \iota_m^\top p = 1 \end{aligned}$$

Hamburger, Thompson and Weil [Hamburger *et al.*, 1967] view the input-output pair of the economy as payoff matrices of two-player zero-sum games.

Using this interpretation, they restate *Assumption 40.2.1* and *Assumption 40.2.2* as follows

$$V(-A) < 0 \qquad \text{and} \qquad V(B) > 0$$

Proof. (Sketch)

- $\Rightarrow V(B) > 0$ implies $x_0^\top B \gg \mathbf{0}$, where x_0 is a maximizing vector. Since B is non-negative, this requires that each column of B has at least one positive entry, which is *Assumption 40.2.1*.
- $\Leftarrow$ From *Assumption 40.2.1* and the fact that $p > \mathbf{0}$, it follows that $Bp > \mathbf{0}$. This implies that the maximizing player can always choose x so that $x^\top Bp > 0$ so that it must be the case that $V(B) > 0$.

In order to (re)state *Theorem 40.4.1* in terms of a particular two-player zero-sum game, we define a matrix for $\gamma \in \mathbb{R}$

$$M(\gamma) \equiv B - \gamma A$$

For fixed γ, treating $M(\gamma)$ as a matrix game, calculating the solution of the game implies

- If $\gamma > \alpha_0$, then for all $x > 0$, there $\exists j \in \{1, \dots, n\}$, s.t. $[x^\top M(\gamma)]_j < 0$ implying that $V(M(\gamma)) < 0$.
- If $\gamma < \beta_0$, then for all $p > 0$, there $\exists i \in \{1, \dots, m\}$, s.t. $[M(\gamma)p]_i > 0$ implying that $V(M(\gamma)) > 0$.
- If $\gamma \in \{\beta_0, \alpha_0\}$, then (by *Theorem 40.4.1*) the optimal intensity and price vectors x_0 and p_0 satisfy

$$x_0^\top M(\gamma) \geq \mathbf{0}^\top \qquad \text{and} \qquad M(\gamma)p_0 \leq \mathbf{0}$$

That is, $(x_0, p_0, 0)$ is a solution of the game $M(\gamma)$ so that $V\left(M(\beta_0)\right) = V\left(M(\alpha_0)\right) = 0$.

- If $\beta_0 < \alpha_0$ and $\gamma \in (\beta_0, \alpha_0)$, then $V(M(\gamma)) = 0$.

Moreover, if x' is optimal for the maximizing player in $M(\gamma')$ for $\gamma' \in (\beta_0, \alpha_0)$ and p'' is optimal for the minimizing player in $M(\gamma'')$ where $\gamma'' \in (\beta_0, \gamma')$, then $(x', p'', 0)$ is a solution for $M(\gamma)$ $\forall \gamma \in (\gamma'', \gamma')$.

Proof. (Sketch) If x' is optimal for a maximizing player in game $M(\gamma')$, then $(x')^\top M(\gamma') \geq \mathbf{0}^\top$ and so for all $\gamma < \gamma'$.

$$(x')^\top M(\gamma) = (x')^\top M(\gamma') + (x')^\top (\gamma' - \gamma) A \geq \mathbf{0}^\top$$

hence $V(M(\gamma)) \geq 0$. If p'' is optimal for a minimizing player in game $M(\gamma'')$, then $M(\gamma'')p'' \leq \mathbf{0}$ and so for all $\gamma'' < \gamma$

$$M(\gamma)p'' = M(\gamma'') + (\gamma'' - \gamma) A p'' \leq \mathbf{0}$$

hence $V(M(\gamma)) \leq 0$.

It is clear from the above argument that β_0, α_0 are the minimal and maximal γ for which $V(M(\gamma)) = 0$.

Furthermore, Hamburger et al. [Hamburger *et al.*, 1967] show that the function $\gamma \mapsto V(M(\gamma))$ is continuous and nonincreasing in γ.

This suggests an algorithm to compute (α_0, x_0) and (β_0, p_0) for a given input-output pair (A, B).

40.5.2 Algorithm

Hamburger, Thompson and Weil [Hamburger *et al.*, 1967] propose a simple bisection algorithm to find the minimal and maximal roots (i.e. β_0 and α_0) of the function $\gamma \mapsto V(M(\gamma))$.

Step 1

First, notice that we can easily find trivial upper and lower bounds for α_0 and β_0.

- TEP requires that $x^\top(B - \alpha A) \geq \mathbf{0}^\top$ and $x > \mathbf{0}$, so if α is so large that $\max_i\{[(B - \alpha A)\iota_n]_i\} < 0$, then TEP ceases to have a solution.

Accordingly, let **UB** be the α^* that solves $\max_i\{[(B - \alpha^* A)\iota_n]_i\} = 0$.

- Similar to the upper bound, if β is so low that $\min_j\{[\iota_m^\top(B - \beta A)]_j\} > 0$, then the EEP has no solution and so we can define **LB** as the β^* that solves $\min_j\{[\iota_m^\top(B - \beta^* A)]_j\} = 0$.

The *bounds* method calculates these trivial bounds for us

```
n1.bounds()
```

```
(1.0, 2.0)
```

Step 2

Compute α_0 and β_0

- Finding α_0
  1. Fix $\gamma = \frac{UB+LB}{2}$ and compute the solution of the two-player zero-sum game associated with $M(\gamma)$. We can use either the primal or the dual LP problem.
  2. If $V(M(\gamma)) \geq 0$, then set $LB = \gamma$, otherwise let $UB = \gamma$.
  3. Iterate on 1. and 2. until $|UB - LB| < \epsilon$.
- Finding β_0
  1. Fix $\gamma = \frac{UB+LB}{2}$ and compute the solution of the two-player zero-sum game associated with $M(\gamma)$. We can use either the primal or the dual LP problem.
  2. If $V(M(\gamma)) > 0$, then set $LB = \gamma$, otherwise let $UB = \gamma$.
  3. Iterate on 1. and 2. until $|UB - LB| < \epsilon$.
- *Existence*: Since $V(M(LB)) > 0$ and $V(M(UB)) < 0$ and $V(M(\cdot))$ is a continuous, nonincreasing function, there is at least one $\gamma \in [LB, UB]$, s.t. $V(M(\gamma)) = 0$.

The *zerosum* method calculates the value and optimal strategies associated with a given γ.

```
y = 2

print(f'Value of the game with y = {y}')
print(n1.zerosum(y=y)[0])
print('Intensity vector (from the primal)')
print(n1.zerosum(y=y)[1])
print('Price vector (from the dual)')
print(n1.zerosum(y=y, dual=True)[1])
```

```
Value of the game with γ = 2
-0.24
Intensity vector (from the primal)
[0.32 0.28 0.4 ]
Price vector (from the dual)
[0.4  0.32 0.28 0.  ]
```

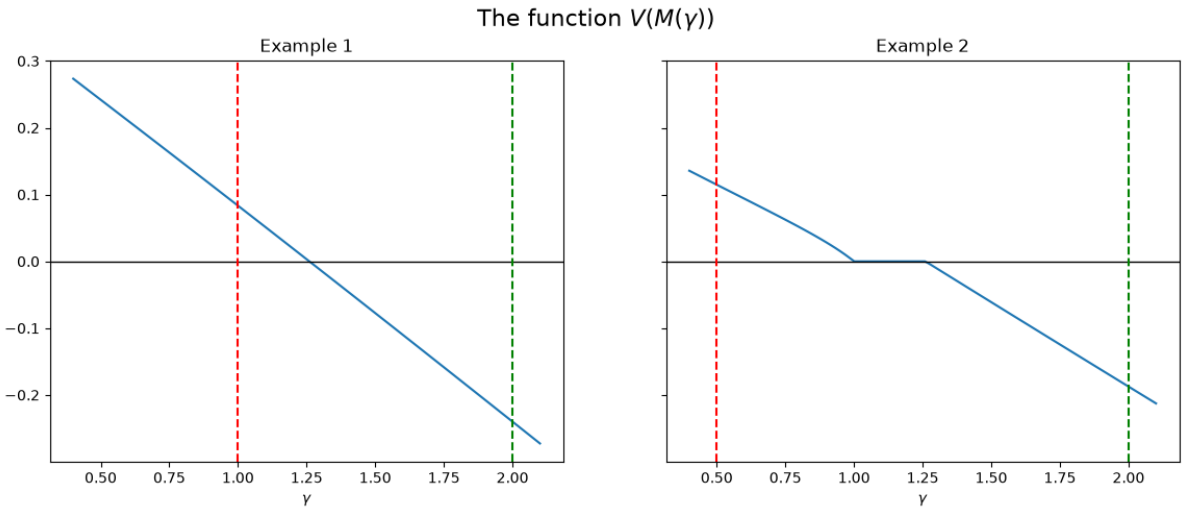
```
numb_grid = 100
γ_grid = np.linspace(0.4, 2.1, numb_grid)

value_ex1_grid = np.asarray([n1.zerosum(γ=γ_grid[i])[0]
                             for i in range(numb_grid)])
value_ex2_grid = np.asarray([n2.zerosum(γ=γ_grid[i])[0]
                             for i in range(numb_grid)])

fig, axes = plt.subplots(1, 2, figsize=(14, 5), sharey=True)
fig.suptitle(r'The function $V(M(\gamma))$', fontsize=16)

for ax, grid, N, i in zip(axes, (value_ex1_grid, value_ex2_grid),
                          (n1, n2), (1, 2)):
    ax.plot(γ_grid, grid)
    ax.set(title=f'Example {i}', xlabel=r'$\gamma$')
    ax.axhline(0, c='k', lw=1)
    ax.axvline(N.bounds()[0], c='r', ls='--', label='lower bound')
    ax.axvline(N.bounds()[1], c='g', ls='--', label='upper bound')

plt.show()
```



The *expansion* method implements the bisection algorithm for α_0 (and uses the primal LP problem for x_0)

```
α_0, x, p = n1.expansion()
print(f'α_0 = {α_0}')
print(f'x_0 = {x}')
print(f'The corresponding p from the dual = {p}')
```

```
α_0 = 1.2599210478365421
x_0 = [0.33 0.26 0.41]
The corresponding p from the dual = [0.41 0.33 0.26 0.  ]
```

The *interest* method implements the bisection algorithm for β_0 (and uses the dual LP problem for p_0)

```
β_0, x, p = n1.interest()
print(f'β_0 = {β_0}')
print(f'p_0 = {p}')
print(f'The corresponding x from the primal = {x}')
```

```
β_0 = 1.2599210478365421
p_0 = [0.41 0.33 0.26 0.  ]
The corresponding x from the primal = [0.33 0.26 0.41]
```

Of course, when γ^* is unique, it is irrelevant which one of the two methods we use – both work.

In particular, as will be shown below, in case of an irreducible (A, B) (like in Example 1), the maximal and minimal roots of $V(M(\gamma))$ necessarily coincide implying a "full duality" result, i.e. $\alpha_0 = \beta_0 = \gamma^*$ so that the expansion (and interest) rate γ^* is unique.

40.5.3 Uniqueness and Irreducibility

As an illustration, compute first the maximal and minimal roots of $V(M(\cdot))$ for our Example 2 that has a reducible input-output pair (A, B)

```
α_0, x, p = n2.expansion()
print(f'α_0 = {α_0}')
print(f'x_0 = {x}')
print(f'The corresponding p from the dual = {p}')
```

```
α_0 = 1.259921052493155
x_0 = [5.27e-10 0.00e+00 3.27e-01 2.60e-01 4.13e-01]
The corresponding p from the dual = [0.   0.21 0.33 0.26 0.21 0.  ]
```

```
β_0, x, p = n2.interest()
print(f'β_0 = {β_0}')
print(f'p_0 = {p}')
print(f'The corresponding x from the primal = {x}')
```

```
β_0 = 1.0000000009313226
p_0 = [ 5.00e-01  5.00e-01 -1.55e-09 -1.24e-09 -9.31e-10  0.00e+00]
The corresponding x from the primal = [-0.    0.    0.25  0.25  0.5 ]
```

As we can see, with a reducible (A, B), the roots found by the bisection algorithms might differ, so there might be multiple γ^* that make the value of the game with $M(\gamma^*)$ zero. (see the figure above).

Indeed, although the von Neumann theorem assures existence of the equilibrium, *Assumption 40.2.1* and *Assumption 40.2.2* are not sufficient for uniqueness. Nonetheless, Kemeny et al. (1967) show that there are at most finitely many economic solutions, meaning that there are only finitely many γ^* that satisfy $V(M(\gamma^*)) = 0$ and $x_0^\top B p_0 > 0$ and that for each such γ_i^*, there is a self-contained part of the economy (a sub-economy) that in equilibrium can expand independently with the expansion coefficient γ_i^*.

The following theorem (see Theorem 9.10. in Gale [Gale, 1989]) asserts that imposing irreducibility is sufficient for uniqueness of (γ^*, x_0, p_0).

Theorem 40.5.1

Adopt the conditions of *Theorem 40.4.1*. If the economy (A, B) is irreducible, then $\gamma^* = \alpha_0 = \beta_0$.

40.5.4 A Special Case

There is a special (A, B) that allows us to simplify the solution method significantly by invoking the powerful Perron-Frobenius theorem for non-negative matrices.

Definition 40.5.2

We call an economy *simple* if it satisfies

- $n = m$
- Each activity produces exactly one good
- Each good is produced by one and only one activity.

These assumptions imply that $B = I_n$, i.e., that B can be written as an identity matrix (possibly after reshuffling its rows and columns).

The simple model has the following special property (Theorem 9.11. in Gale [Gale, 1989]): if x_0 and $\alpha_0 > 0$ solve the TEP with (A, I_n), then

$$x_0^\top = \alpha_0 x_0^\top A \qquad \Leftrightarrow \qquad x_0^\top A = \left(\frac{1}{\alpha_0}\right) x_0^\top$$

The latter shows that $1/\alpha_0$ is a positive eigenvalue of A and x_0 is the corresponding non-negative left eigenvector.

The classic result of *Perron and Frobenius* implies that a non-negative matrix has a non-negative eigenvalue-eigenvector pair.

Moreover, if A is irreducible, then the optimal intensity vector x_0 is positive and *unique* up to multiplication by a positive scalar.

Suppose that A is reducible with k irreducible subsets $S_1, \ldots, S_k$. Let A_i be the submatrix corresponding to S_i and let α_i and β_i be the associated expansion and interest factors, respectively. Then we have

$$\alpha_0 = \max_i \{\alpha_i\} \qquad \text{and} \qquad \beta_0 = \min_i \{\beta_i\}$$

Part VI

Introduction to Dynamics

CHAPTER

FORTYONE

FINITE MARKOV CHAINS

Contents

- *Finite Markov Chains*
  - *Overview*
  - *Definitions*
  - *Simulation*
  - *Marginal Distributions*
  - *Irreducibility and Aperiodicity*
  - *Stationary Distributions*
  - *Ergodicity*
  - *Computing Expectations*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

41.1 Overview

Markov chains are one of the most useful classes of stochastic processes, being

- simple, flexible and supported by many elegant theoretical results
- valuable for building intuition about random dynamic models
- central to quantitative modeling in their own right

You will find them in many of the workhorse models of economics and finance.

In this lecture, we review some of the theory of Markov chains.

We will also introduce some of the high-quality routines for working with Markov chains available in QuantEcon.py.

Prerequisite knowledge is basic probability and linear algebra.

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import quantecon as qe
import numpy as np
from mpl_toolkits.mplot3d import Axes3D
```

41.2 Definitions

The following concepts are fundamental.

41.2.1 Stochastic Matrices

A **stochastic matrix** (or **Markov matrix**) is an $n \times n$ square matrix P such that

1. each element of P is nonnegative, and
2. each row of P sums to one

Each row of P can be regarded as a probability mass function over n possible outcomes.

It is too not difficult to check[1] that if P is a stochastic matrix, then so is the k-th power P^k for all $k \in \mathbb{N}$.

41.2.2 Markov Chains

There is a close connection between stochastic matrices and Markov chains.

To begin, let S be a finite set with n elements $\{x_1, \ldots, x_n\}$.

The set S is called the **state space** and $x_1, \ldots, x_n$ are the **state values**.

A **Markov chain** $\{X_t\}$ on S is a sequence of random variables on S that have the **Markov property**.

This means that, for any date t and any state $y \in S$,

$$\mathbb{P}\{X_{t+1} = y \,|\, X_t\} = \mathbb{P}\{X_{t+1} = y \,|\, X_t, X_{t-1}, \ldots\} \tag{41.1}$$

In other words, knowing the current state is enough to know probabilities for future states.

In particular, the dynamics of a Markov chain are fully determined by the set of values

$$P(x, y) := \mathbb{P}\{X_{t+1} = y \,|\, X_t = x\} \qquad (x, y \in S) \tag{41.2}$$

By construction,

- $P(x, y)$ is the probability of going from x to y in one unit of time (one step)
- $P(x, \cdot)$ is the conditional distribution of X_{t+1} given $X_t = x$

We can view P as a stochastic matrix where

$$P_{ij} = P(x_i, x_j) \qquad 1 \leq i, j \leq n$$

Going the other way, if we take a stochastic matrix P, we can generate a Markov chain $\{X_t\}$ as follows:

- draw X_0 from a marginal distribution ψ
- for each $t = 0, 1, \ldots$, draw X_{t+1} from $P(X_t, \cdot)$

By construction, the resulting process satisfies (41.2).

[1] Hint: First show that if P and Q are stochastic matrices then so is their product — to check the row sums, try post multiplying by a column vector of ones. Finally, argue that P^n is a stochastic matrix using induction.

41.2.3 Example 1

Consider a worker who, at any given time t, is either unemployed (state 0) or employed (state 1).

Suppose that, over a one month period,

1. An unemployed worker finds a job with probability $\alpha \in (0, 1)$.
2. An employed worker loses her job and becomes unemployed with probability $\beta \in (0, 1)$.

In terms of a Markov model, we have

- $S = \{0, 1\}$
- $P(0, 1) = \alpha$ and $P(1, 0) = \beta$

We can write out the transition probabilities in matrix form as

$$P = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix} \tag{41.3}$$

Once we have the values α and β, we can address a range of questions, such as

- What is the average duration of unemployment?
- Over the long-run, what fraction of time does a worker find herself unemployed?
- Conditional on employment, what is the probability of becoming unemployed at least once over the next 12 months?

We'll cover such applications below.

41.2.4 Example 2

From US unemployment data, Hamilton [Hamilton, 2005] estimated the stochastic matrix

$$P = \begin{pmatrix} 0.971 & 0.029 & 0 \\ 0.145 & 0.778 & 0.077 \\ 0 & 0.508 & 0.492 \end{pmatrix}$$

where

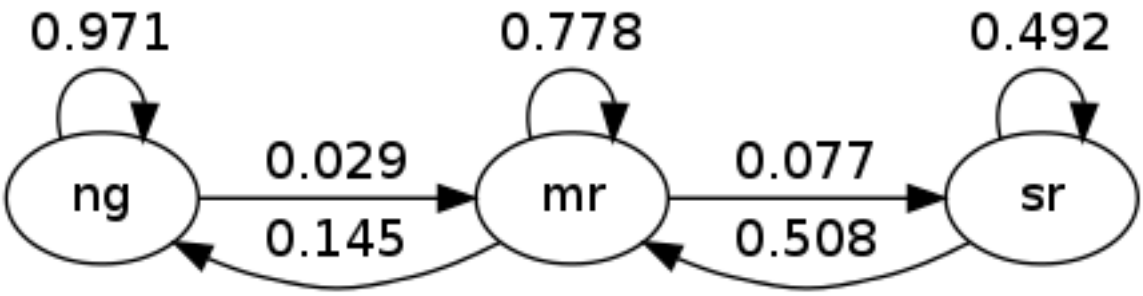
- the frequency is monthly
- the first state represents "normal growth"
- the second state represents "mild recession"
- the third state represents "severe recession"

For example, the matrix tells us that when the state is normal growth, the state will again be normal growth next month with probability 0.97.

In general, large values on the main diagonal indicate persistence in the process $\{X_t\}$.

This Markov process can also be represented as a directed graph, with edges labeled by transition probabilities

Here "ng" is normal growth, "mr" is mild recession, etc.



41.3 Simulation

One natural way to answer questions about Markov chains is to simulate them.

(To approximate the probability of event E, we can simulate many times and count the fraction of times that E occurs).

Nice functionality for simulating Markov chains exists in QuantEcon.py.

- Efficient, bundled with lots of other useful routines for handling Markov chains.

However, it's also a good exercise to roll our own routines — let's do that first and then come back to the methods in QuantEcon.py.

In these exercises, we'll take the state space to be $S = 0, \ldots, n-1$.

41.3.1 Rolling Our Own

To simulate a Markov chain, we need its stochastic matrix P and a marginal probability distribution ψ from which to draw a realization of X_0.

The Markov chain is then constructed as discussed above. To repeat:

1. At time $t = 0$, draw a realization of X_0 from ψ.
2. At each subsequent time t, draw a realization of the new state X_{t+1} from $P(X_t, \cdot)$.

To implement this simulation procedure, we need a method for generating draws from a discrete distribution.

For this task, we'll use `random.draw` from QuantEcon, which works as follows:

```
ψ = (0.3, 0.7)           # probabilities over {0, 1}
cdf = np.cumsum(ψ)       # convert into cummulative distribution
qe.random.draw(cdf, 5)   # generate 5 independent draws from ψ
```

```
array([0, 1, 0, 1, 0])
```

We'll write our code as a function that accepts the following three arguments

- A stochastic matrix `P`
- An initial state `init`
- A positive integer `sample_size` representing the length of the time series the function should return

```
def mc_sample_path(P, ψ_0=None, sample_size=1_000):

    # set up
```

(continues on next page)

(continued from previous page)

```
    P = np.asarray(P)
    X = np.empty(sample_size, dtype=int)

    # Convert each row of P into a cdf
    n = len(P)
    P_dist = [np.cumsum(P[i, :]) for i in range(n)]

    # draw initial state, defaulting to 0
    if ψ_0 is not None:
        X_0 = qe.random.draw(np.cumsum(ψ_0))
    else:
        X_0 = 0

    # simulate
    X[0] = X_0
    for t in range(sample_size - 1):
        X[t+1] = qe.random.draw(P_dist[X[t]])

    return X
```

Let's see how it works using the small matrix

```
P = [[0.4, 0.6],
     [0.2, 0.8]]
```

As we'll see later, for a long series drawn from `P`, the fraction of the sample that takes value 0 will be about 0.25.

Moreover, this is true, regardless of the initial distribution from which X_0 is drawn.

The following code illustrates this

```
X = mc_sample_path(P, ψ_0=[0.1, 0.9], sample_size=100_000)
np.mean(X == 0)
```

```
np.float64(0.25039)
```

You can try changing the initial distribution to confirm that the output is always close to 0.25, at least for the `P` matrix above.

41.3.2 Using QuantEcon's Routines

As discussed above, QuantEcon.py has routines for handling Markov chains, including simulation.

Here's an illustration using the same P as the preceding example

```
from quantecon import MarkovChain

mc = qe.MarkovChain(P)
X = mc.simulate(ts_length=1_000_000)
np.mean(X == 0)
```

```
np.float64(0.249341)
```

The QuantEcon.py routine is JIT compiled and much faster.

```
%time mc_sample_path(P, sample_size=1_000_000) # Our homemade code version
```

```
CPU times: user 2.5 s, sys: 2.79 ms, total: 2.51 s
Wall time: 2.51 s
```

```
array([0, 0, 1, ..., 1, 0, 1], shape=(1000000,))
```

```
%time mc.simulate(ts_length=1_000_000) # qe code version
```

```
CPU times: user 15.4 ms, sys: 4.02 ms, total: 19.4 ms
Wall time: 19.2 ms
```

```
array([0, 1, 1, ..., 0, 1, 1], shape=(1000000,))
```

Adding State Values and Initial Conditions

If we wish to, we can provide a specification of state values to `MarkovChain`.

These state values can be integers, floats, or even strings.

The following code illustrates

```
mc = qe.MarkovChain(P, state_values=('unemployed', 'employed'))
mc.simulate(ts_length=4, init='employed')
```

```
array(['employed', 'employed', 'employed', 'employed'], dtype='<U10')
```

```
mc.simulate(ts_length=4, init='unemployed')
```

```
array(['unemployed', 'employed', 'employed', 'employed'], dtype='<U10')
```

```
mc.simulate(ts_length=4)  # Start at randomly chosen initial state
```

```
array(['unemployed', 'employed', 'unemployed', 'employed'], dtype='<U10')
```

If we want to see indices rather than state values as outputs as we can use

```
mc.simulate_indices(ts_length=4)
```

```
array([0, 1, 1, 1])
```

41.4 Marginal Distributions

Suppose that

1. $\{X_t\}$ is a Markov chain with stochastic matrix P
2. the marginal distribution of X_t is known to be ψ_t

What then is the marginal distribution of X_{t+1}, or, more generally, of X_{t+m}?

To answer this, we let ψ_t be the marginal distribution of X_t for $t = 0, 1, 2, ...$

Our first aim is to find ψ_{t+1} given ψ_t and P.

To begin, pick any $y \in S$.

Using the law of total probability, we can decompose the probability that $X_{t+1} = y$ as follows:

$$\mathbb{P}\{X_{t+1} = y\} = \sum_{x \in S} \mathbb{P}\{X_{t+1} = y \,|\, X_t = x\} \cdot \mathbb{P}\{X_t = x\}$$

In words, to get the probability of being at y tomorrow, we account for all ways this can happen and sum their probabilities.

Rewriting this statement in terms of marginal and conditional probabilities gives

$$\psi_{t+1}(y) = \sum_{x \in S} P(x, y)\psi_t(x)$$

There are n such equations, one for each $y \in S$.

If we think of ψ_{t+1} and ψ_t as *row vectors*, these n equations are summarized by the matrix expression

$$\psi_{t+1} = \psi_t P \tag{41.4}$$

Thus, to move a marginal distribution forward one unit of time, we postmultiply by P.

By postmultiplying m times, we move a marginal distribution forward m steps into the future.

Hence, iterating on (41.4), the expression $\psi_{t+m} = \psi_t P^m$ is also valid — here P^m is the m-th power of P.

As a special case, we see that if ψ_0 is the initial distribution from which X_0 is drawn, then $\psi_0 P^m$ is the distribution of X_m.

This is very important, so let's repeat it

$$X_0 \sim \psi_0 \quad \Longrightarrow \quad X_m \sim \psi_0 P^m \tag{41.5}$$

and, more generally,

$$X_t \sim \psi_t \quad \Longrightarrow \quad X_{t+m} \sim \psi_t P^m \tag{41.6}$$

41.4.1 Multiple Step Transition Probabilities

We know that the probability of transitioning from x to y in one step is $P(x, y)$.

It turns out that the probability of transitioning from x to y in m steps is $P^m(x, y)$, the (x, y)-th element of the m-th power of P.

To see why, consider again (41.6), but now with a ψ_t that puts all probability on state x so that the transition probabilities are

- 1 in the x-th position and zero elsewhere

Inserting this into (41.6), we see that, conditional on $X_t = x$, the distribution of X_{t+m} is the x-th row of P^m.

In particular

$$\mathbb{P}\{X_{t+m} = y \,|\, X_t = x\} = P^m(x, y) = (x, y)\text{-th element of } P^m$$

41.4.2 Example: Probability of Recession

Recall the stochastic matrix P for recession and growth *considered above*.

Suppose that the current state is unknown — perhaps statistics are available only at the *end* of the current month.

We guess that the probability that the economy is in state x is $\psi(x)$.

The probability of being in recession (either mild or severe) in 6 months time is given by the inner product

$$\psi P^6 \cdot \begin{pmatrix} 0 \\ 1 \\ 1 \end{pmatrix}$$

41.4.3 Example 2: Cross-Sectional Distributions

The marginal distributions we have been studying can be viewed either as probabilities or as cross-sectional frequencies that a Law of Large Numbers leads us to anticipate for large samples.

To illustrate, recall our model of employment/unemployment dynamics for a given worker *discussed above*.

Consider a large population of workers, each of whose lifetime experience is described by the specified dynamics, with each worker's outcomes being realizations of processes that are statistically independent of all other workers' processes.

Let ψ be the current *cross-sectional* distribution over $\{0, 1\}$.

The cross-sectional distribution records fractions of workers employed and unemployed at a given moment.

- For example, $\psi(0)$ is the unemployment rate.

What will the cross-sectional distribution be in 10 periods hence?

The answer is ψP^{10}, where P is the stochastic matrix in (41.3).

This is because each worker's state evolves according to P, so ψP^{10} is a marginal distibution for a single randomly selected worker.

But when the sample is large, outcomes and probabilities are roughly equal (by an application of the Law of Large Numbers).

So for a very large (tending to infinite) population, ψP^{10} also represents fractions of workers in each state.

This is exactly the cross-sectional distribution.

41.5 Irreducibility and Aperiodicity

Irreducibility and aperiodicity are central concepts of modern Markov chain theory.

Let's see what they're about.

41.5.1 Irreducibility

Let P be a fixed stochastic matrix.

Two states x and y are said to **communicate** with each other if there exist positive integers j and k such that

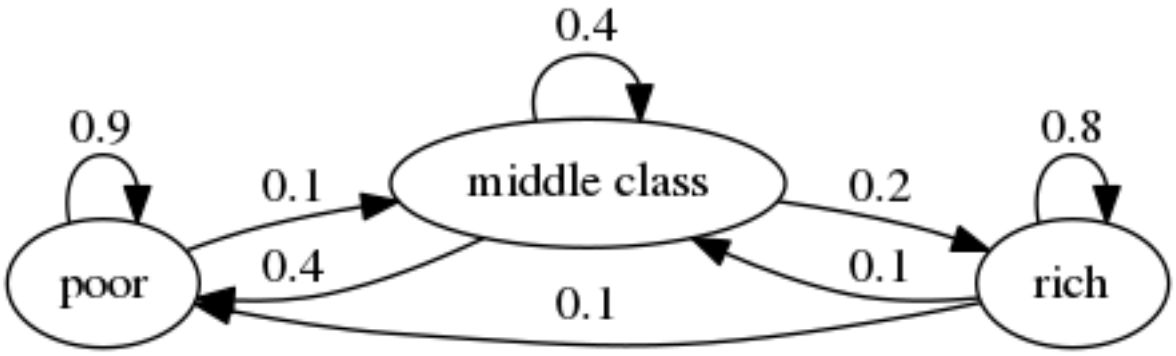
$$P^j(x, y) > 0 \quad \text{and} \quad P^k(y, x) > 0$$

In view of our discussion *above*, this means precisely that

- state x can eventually be reached from state y, and
- state y can eventually be reached from state x

The stochastic matrix P is called **irreducible** if all states communicate; that is, if x and y communicate for all (x, y) in $S \times S$.

For example, consider the following transition probabilities for wealth of a fictitious set of households



We can translate this into a stochastic matrix, putting zeros where there's no edge between nodes

$$P := \begin{pmatrix} 0.9 & 0.1 & 0 \\ 0.4 & 0.4 & 0.2 \\ 0.1 & 0.1 & 0.8 \end{pmatrix}$$

It's clear from the graph that this stochastic matrix is irreducible: we can eventually reach any state from any other state.

We can also test this using QuantEcon.py's MarkovChain class

```
P = [[0.9, 0.1, 0.0],
     [0.4, 0.4, 0.2],
     [0.1, 0.1, 0.8]]

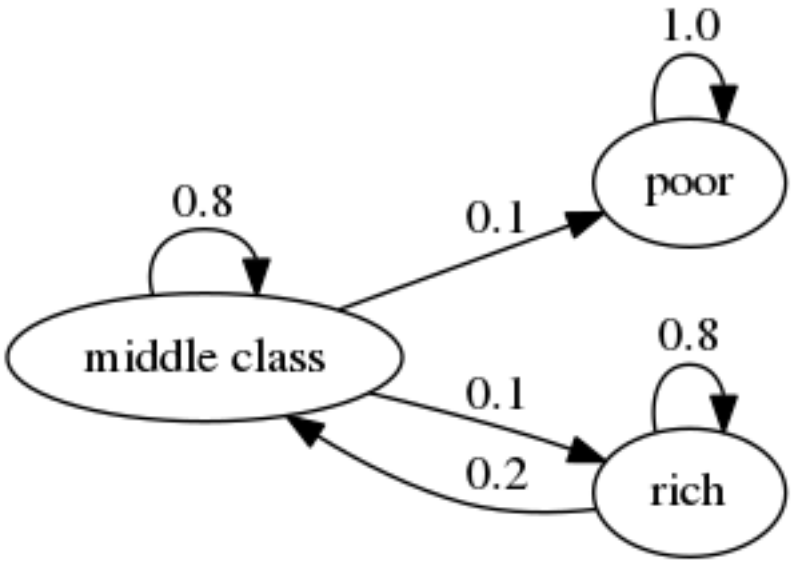
mc = qe.MarkovChain(P, ('poor', 'middle', 'rich'))
mc.is_irreducible
```

```
True
```

Here's a more pessimistic scenario in which poor people remain poor forever

This stochastic matrix is not irreducible, since, for example, rich is not accessible from poor.

Let's confirm this




```
P = [[1.0, 0.0, 0.0],
     [0.1, 0.8, 0.1],
     [0.0, 0.2, 0.8]]

mc = qe.MarkovChain(P, ('poor', 'middle', 'rich'))
mc.is_irreducible
```

```
False
```

We can also determine the "communication classes"

```
mc.communication_classes
```

```
[array(['poor'], dtype='<U6'), array(['middle', 'rich'], dtype='<U6')]
```

It might be clear to you already that irreducibility is going to be important in terms of long run outcomes.

For example, poverty is a life sentence in the second graph but not the first.

We'll come back to this a bit later.

41.5.2 Aperiodicity

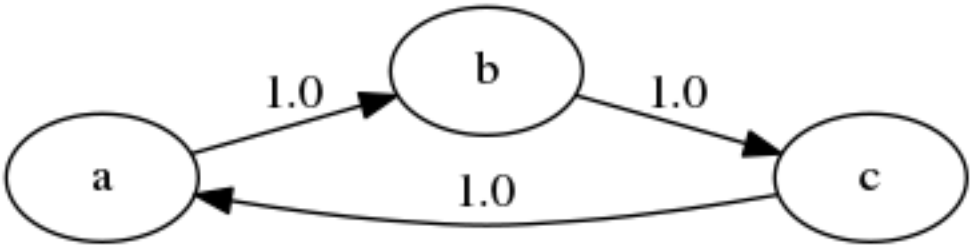
Loosely speaking, a Markov chain is called **periodic** if it cycles in a predictable way, and **aperiodic** otherwise.

Here's a trivial example with three states

The chain cycles with period 3:

```
P = [[0, 1, 0],
     [0, 0, 1],
     [1, 0, 0]]

mc = qe.MarkovChain(P)
mc.period
```




```
3
```

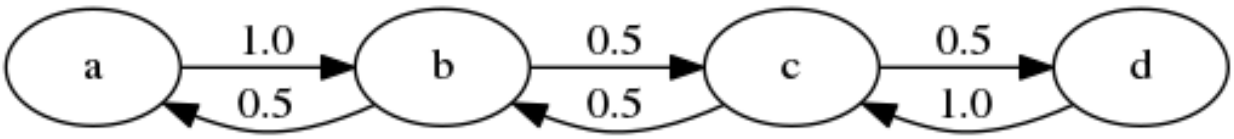
More formally, the **period** of a state x is the largest common divisor of a set of integers

$$D(x) := \{j \geq 1 : P^j(x, x) > 0\}$$

In the last example, $D(x) = \{3, 6, 9, ...\}$ for every state x, so the period is 3.

A stochastic matrix is called **aperiodic** if the period of every state is 1, and **periodic** otherwise.

For example, the stochastic matrix associated with the transition probabilities below is periodic because, for example, state a has period 2



We can confirm that the stochastic matrix is periodic with the following code

```
P = [[0.0, 1.0, 0.0, 0.0],
     [0.5, 0.0, 0.5, 0.0],
     [0.0, 0.5, 0.0, 0.5],
     [0.0, 0.0, 1.0, 0.0]]

mc = qe.MarkovChain(P)
mc.period
```

```
2
```

```
mc.is_aperiodic
```

```
False
```

41.6 Stationary Distributions

As seen in (41.4), we can shift a marginal distribution forward one unit of time via postmultiplication by P.

Some distributions are invariant under this updating process — for example,

```
P = np.array([[0.4, 0.6],
              [0.2, 0.8]])
```

(continues on next page)

(continued from previous page)

```
ψ = (0.25, 0.75)
ψ @ P
```

```
array([0.25, 0.75])
```

Such distributions are called **stationary** or **invariant**.

Formally, a marginal distribution ψ^* on S is called **stationary** for P if $\psi^* = \psi^* P$.

(This is the same notion of stationarity that we learned about in the lecture on AR(1) processes applied to a different setting.)

From this equality, we immediately get $\psi^* = \psi^* P^t$ for all t.

This tells us an important fact: If the distribution of X_0 is a stationary distribution, then X_t will have this same distribution for all t.

Hence stationary distributions have a natural interpretation as **stochastic steady states** — we'll discuss this more soon.

Mathematically, a stationary distribution is a fixed point of P when P is thought of as the map $\psi \mapsto \psi P$ from (row) vectors to (row) vectors.

Theorem. Every stochastic matrix P has at least one stationary distribution.

(We are assuming here that the state space S is finite; if not more assumptions are required)

For proof of this result, you can apply Brouwer's fixed point theorem, or see EDTC, theorem 4.3.5.

There can be many stationary distributions corresponding to a given stochastic matrix P.

- For example, if P is the identity matrix, then all marginal distributions are stationary.

To get uniqueness an invariant distribution, the transition matrix P must have the property that no nontrivial subsets of the state space are **infinitely persistent**.

A subset of the state space is infinitely persistent if other parts of the state space cannot be accessed from it.

Thus, infinite persistence of a non-trivial subset is the opposite of irreducibility.

This gives some intuition for the following fundamental theorem.

Theorem. If P is both aperiodic and irreducible, then

1. P has exactly one stationary distribution ψ^*.
2. For any initial marginal distribution ψ_0, we have $\|\psi_0 P^t - \psi^*\| \to 0$ as $t \to \infty$.

For a proof, see, for example, theorem 5.2 of [Häggström, 2002].

(Note that part 1 of the theorem only requires irreducibility, whereas part 2 requires both irreducibility and aperiodicity)

A stochastic matrix that satisfies the conditions of the theorem is sometimes called **uniformly ergodic**.

A sufficient condition for aperiodicity and irreducibility is that every element of P is strictly positive.

- Try to convince yourself of this.

41.6.1 Example

Recall our model of the employment/unemployment dynamics of a particular worker *discussed above*.

Assuming $\alpha \in (0, 1)$ and $\beta \in (0, 1)$, the uniform ergodicity condition is satisfied.

Let $\psi^* = (p, 1-p)$ be the stationary distribution, so that p corresponds to unemployment (state 0).

Using $\psi^* = \psi^* P$ and a bit of algebra yields

$$p = \frac{\beta}{\alpha + \beta}$$

This is, in some sense, a steady state probability of unemployment — more about the interpretation of this below.

Not surprisingly it tends to zero as $\beta \to 0$, and to one as $\alpha \to 0$.

41.6.2 Calculating Stationary Distributions

As discussed above, a particular Markov matrix P can have many stationary distributions.

That is, there can be many row vectors ψ such that $\psi = \psi P$.

In fact if P has two distinct stationary distributions ψ_1, ψ_2 then it has infinitely many, since in this case, as you can verify, for any $\lambda \in [0, 1]$

$$\psi_3 := \lambda \psi_1 + (1 - \lambda) \psi_2$$

is a stationary distribution for P.

If we restrict attention to the case in which only one stationary distribution exists, one way to finding it is to solve the system

$$\psi (I_n - P) = 0 \tag{41.7}$$

for ψ, where I_n is the $n \times n$ identity.

But the zero vector solves system (41.7), so we must proceed cautiously.

We want to impose the restriction that ψ is a probability distribution.

There are various ways to do this.

One option is to regard solving system (41.7) as an eigenvector problem: a vector ψ such that $\psi = \psi P$ is a left eigenvector associated with the unit eigenvalue $\lambda = 1$.

A stable and sophisticated algorithm specialized for stochastic matrices is implemented in QuantEcon.py.

This is the one we recommend:

```
P = [[0.4, 0.6],
     [0.2, 0.8]]

mc = qe.MarkovChain(P)
mc.stationary_distributions  # Show all stationary distributions
```

```
array([[0.25, 0.75]])
```

41.6.3 Convergence to Stationarity

Part 2 of the Markov chain convergence theorem *stated above* tells us that the marginal distribution of X_t converges to the stationary distribution regardless of where we begin.

This adds considerable authority to our interpretation of ψ^* as a stochastic steady state.

The convergence in the theorem is illustrated in the next figure

```
P = ((0.971, 0.029, 0.000),
     (0.145, 0.778, 0.077),
     (0.000, 0.508, 0.492))
P = np.array(P)

ψ = (0.0, 0.2, 0.8)         # Initial condition

fig = plt.figure(figsize=(8, 6))
ax = fig.add_subplot(111, projection='3d')

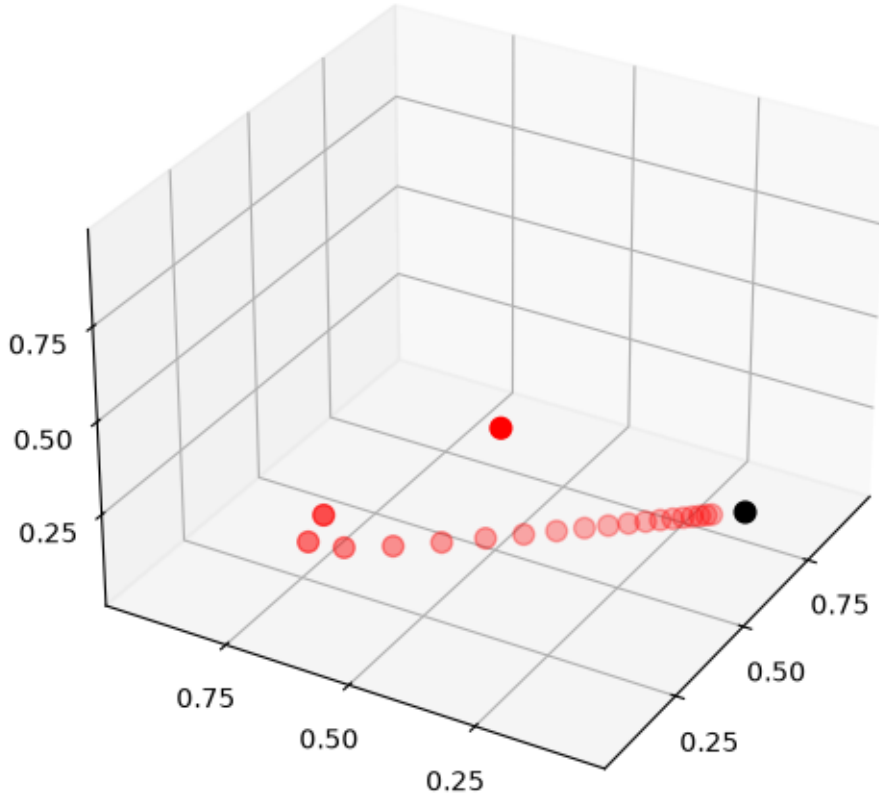
ax.set(xlim=(0, 1), ylim=(0, 1), zlim=(0, 1),
       xticks=(0.25, 0.5, 0.75),
       yticks=(0.25, 0.5, 0.75),
       zticks=(0.25, 0.5, 0.75))

x_vals, y_vals, z_vals = [], [], []
for t in range(20):
    x_vals.append(ψ[0])
    y_vals.append(ψ[1])
    z_vals.append(ψ[2])
    ψ = ψ @ P

ax.scatter(x_vals, y_vals, z_vals, c='r', s=60)
ax.view_init(30, 210)

mc = qe.MarkovChain(P)
ψ_star = mc.stationary_distributions[0]
ax.scatter(ψ_star[0], ψ_star[1], ψ_star[2], c='k', s=60)

plt.show()
```



Here

- P is the stochastic matrix for recession and growth *considered above*.
- The highest red dot is an arbitrarily chosen initial marginal probability distribution ψ, represented as a vector in $\mathbb{R}^3$.
- The other red dots are the marginal distributions ψP^t for $t = 1, 2, ...$
- The black dot is ψ^*.

You might like to try experimenting with different initial conditions.

41.7 Ergodicity

Under irreducibility, yet another important result obtains: for all $x \in S$,

$$\frac{1}{m} \sum_{t=1}^{m} \mathbf{1}\{X_t = x\} \to \psi^*(x) \quad \text{as } m \to \infty \tag{41.8}$$

Here

- $\mathbf{1}\{X_t = x\} = 1$ if $X_t = x$ and zero otherwise
- convergence is with probability one
- the result does not depend on the marginal distribution of X_0

The result tells us that the fraction of time the chain spends at state x converges to $\psi^*(x)$ as time goes to infinity.

This gives us another way to interpret the stationary distribution — provided that the convergence result in (41.8) is valid.

The convergence asserted in (41.8) is a special case of a law of large numbers result for Markov chains — see EDTC, section 4.3.4 for some additional information.

41.7.1 Example

Recall our cross-sectional interpretation of the employment/unemployment model *discussed above*.

Assume that $\alpha \in (0, 1)$ and $\beta \in (0, 1)$, so that irreducibility and aperiodicity both hold.

We saw that the stationary distribution is $(p, 1 - p)$, where

$$p = \frac{\beta}{\alpha + \beta}$$

In the cross-sectional interpretation, this is the fraction of people unemployed.

In view of our latest (ergodicity) result, it is also the fraction of time that a single worker can expect to spend unemployed.

Thus, in the long-run, cross-sectional averages for a population and time-series averages for a given person coincide.

This is one aspect of the concept of ergodicity.

41.8 Computing Expectations

We sometimes want to compute mathematical expectations of functions of X_t of the form

$$\mathbb{E}[h(X_t)] \tag{41.9}$$

and conditional expectations such as

$$\mathbb{E}[h(X_{t+k}) \mid X_t = x] \tag{41.10}$$

where

- $\{X_t\}$ is a Markov chain generated by $n \times n$ stochastic matrix P
- h is a given function, which, in terms of matrix algebra, we'll think of as the column vector

$$h = \begin{pmatrix} h(x_1) \\ \vdots \\ h(x_n) \end{pmatrix}$$

Computing the unconditional expectation (41.9) is easy.

We just sum over the marginal distribution of X_t to get

$$\mathbb{E}[h(X_t)] = \sum_{x \in S} (\psi P^t)(x) h(x)$$

Here ψ is the distribution of X_0.

Since ψ and hence ψP^t are row vectors, we can also write this as

$$\mathbb{E}[h(X_t)] = \psi P^t h$$

For the conditional expectation (41.10), we need to sum over the conditional distribution of X_{t+k} given $X_t = x$.

We already know that this is $P^k(x, \cdot)$, so

$$\mathbb{E}[h(X_{t+k}) \mid X_t = x] = (P^k h)(x) \tag{41.11}$$

The vector $P^k h$ stores the conditional expectation $\mathbb{E}[h(X_{t+k}) \mid X_t = x]$ over all x.

41.8.1 Iterated Expectations

The **law of iterated expectations** states that

$$\mathbb{E}\left[\mathbb{E}[h(X_{t+k}) \mid X_t = x]\right] = \mathbb{E}[h(X_{t+k})]$$

where the outer $\mathbb{E}$ on the left side is an unconditional distribution taken with respect to the marginal distribution ψ_t of X_t (again see equation (41.6)).

To verify the law of iterated expectations, use equation (41.11) to substitute $(P^k h)(x)$ for $E[h(X_{t+k}) \mid X_t = x]$, write

$$\mathbb{E}\left[\mathbb{E}[h(X_{t+k}) \mid X_t = x]\right] = \psi_t P^k h,$$

and note $\psi_t P^k h = \psi_{t+k} h = \mathbb{E}[h(X_{t+k})]$.

41.8.2 Expectations of Geometric Sums

Sometimes we want to compute the mathematical expectation of a geometric sum, such as $\sum_t \beta^t h(X_t)$.

In view of the preceding discussion, this is

$$\mathbb{E}\Big[\sum_{j=0}^{\infty} \beta^j h(X_{t+j}) \mid X_t = x\Big] = [(I - \beta P)^{-1} h](x)$$

where

$$(I - \beta P)^{-1} = I + \beta P + \beta^2 P^2 + \cdots$$

Premultiplication by $(I - \beta P)^{-1}$ amounts to "applying the **resolvent operator**".

41.9 Exercises

Exercise 41.9.1

According to the discussion *above*, if a worker's employment dynamics obey the stochastic matrix

$$P = \begin{pmatrix} 1 - \alpha & \alpha \\ \beta & 1 - \beta \end{pmatrix}$$

with $\alpha \in (0, 1)$ and $\beta \in (0, 1)$, then, in the long-run, the fraction of time spent unemployed will be

$$p := \frac{\beta}{\alpha + \beta}$$

In other words, if $\{X_t\}$ represents the Markov chain for employment, then $\bar{X}_m \to p$ as $m \to \infty$, where

$$\bar{X}_m := \frac{1}{m} \sum_{t=1}^{m} \mathbf{1}\{X_t = 0\}$$

This exercise asks you to illustrate convergence by computing $\bar{X}_m$ for large m and checking that it is close to p.

You will see that this statement is true regardless of the choice of initial condition or the values of α, β, provided both lie in $(0, 1)$.

Solution

We will address this exercise graphically.

The plots show the time series of $\bar{X}_m - p$ for two initial conditions.

As m gets large, both series converge to zero.

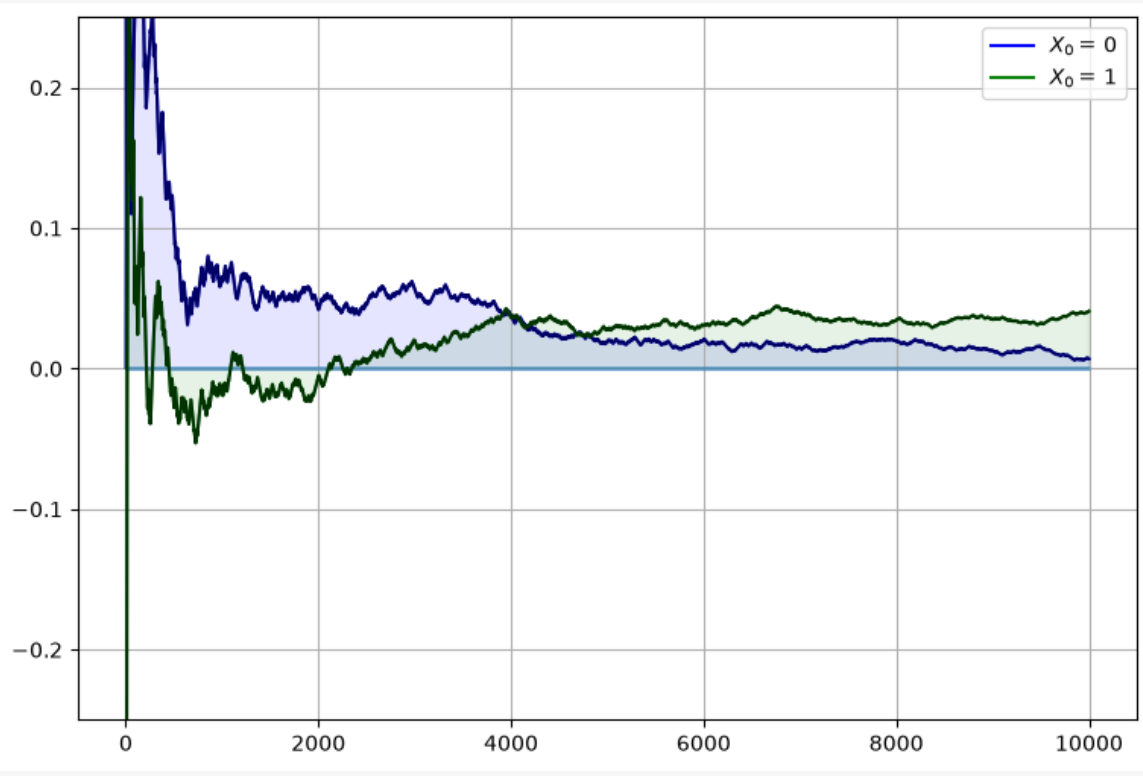
```
α = β = 0.1
N = 10000
p = β / (α + β)

P = ((1 - α,        α),               # Careful: P and p are distinct
     (    β,   1 - β))
mc = MarkovChain(P)

fig, ax = plt.subplots(figsize=(9, 6))
ax.set_ylim(-0.25, 0.25)
ax.grid()
ax.hlines(0, 0, N, lw=2, alpha=0.6)   # Horizonal line at zero

for x0, col in ((0, 'blue'), (1, 'green')):
    # Generate time series for worker that starts at x0
    X = mc.simulate(N, init=x0)
    # Compute fraction of time spent unemployed, for each n
    X_bar = (X == 0).cumsum() / (1 + np.arange(N, dtype=float))
    # Plot
    ax.fill_between(range(N), np.zeros(N), X_bar - p, color=col, alpha=0.1)
    ax.plot(X_bar - p, color=col, label=fr'$X_0 = \, {x0} $')
    # Overlay in black--make lines clearer
    ax.plot(X_bar - p, 'k-', alpha=0.6)

ax.legend(loc='upper right')
plt.show()
```



Exercise 41.9.2

A topic of interest for economics and many other disciplines is *ranking*.

Let's now consider one of the most practical and important ranking problems — the rank assigned to web pages by search engines.

(Although the problem is motivated from outside of economics, there is in fact a deep connection between search ranking systems and prices in certain competitive equilibria — see [Du *et al.*, 2013].)

To understand the issue, consider the set of results returned by a query to a web search engine.

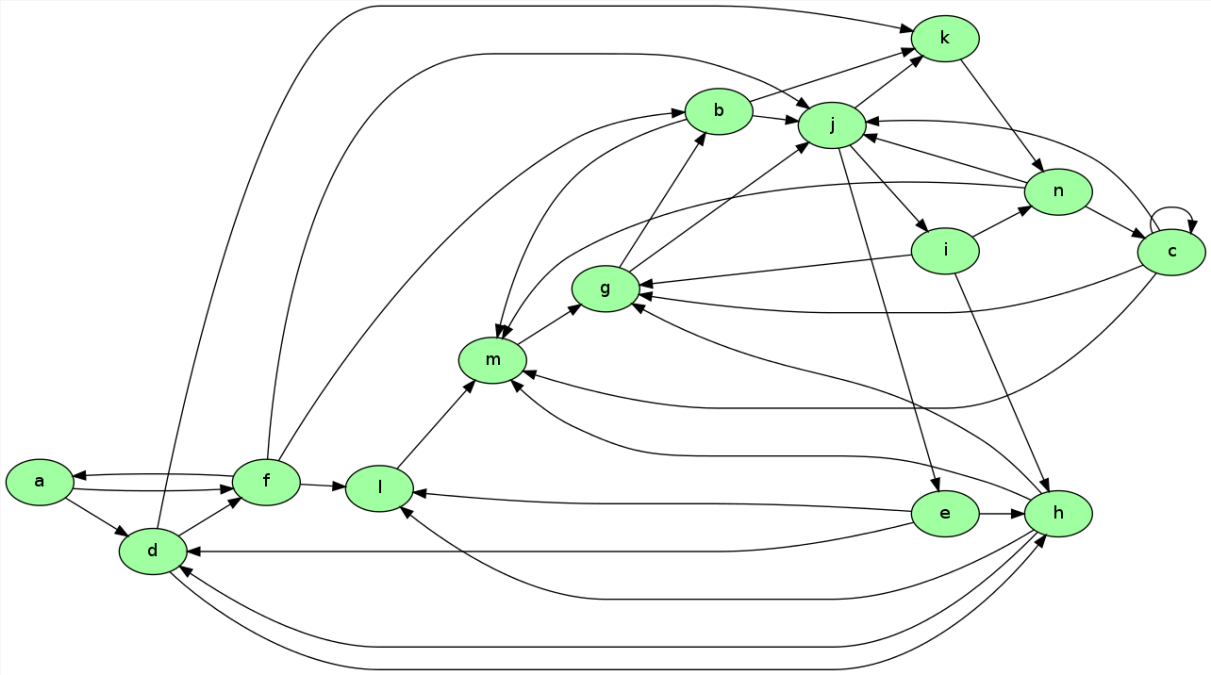
For the user, it is desirable to

1. receive a large set of accurate matches
2. have the matches returned in order, where the order corresponds to some measure of "importance"

Ranking according to a measure of importance is the problem we now consider.

The methodology developed to solve this problem by Google founders Larry Page and Sergey Brin is known as PageRank.

To illustrate the idea, consider the following diagram



Imagine that this is a miniature version of the WWW, with

- each node representing a web page
- each arrow representing the existence of a link from one page to another

Now let's think about which pages are likely to be important, in the sense of being valuable to a search engine user.

One possible criterion for the importance of a page is the number of inbound links — an indication of popularity.

By this measure, `m` and `j` are the most important pages, with 5 inbound links each.

However, what if the pages linking to `m`, say, are not themselves important?

Thinking this way, it seems appropriate to weight the inbound nodes by relative importance.

The PageRank algorithm does precisely this.

A slightly simplified presentation that captures the basic idea is as follows.

Letting j be (the integer index of) a typical page and r_j be its ranking, we set

$$r_j = \sum_{i \in L_j} \frac{r_i}{\ell_i}$$

where

- ℓ_i is the total number of outbound links from i
- L_j is the set of all pages i such that i has a link to j

This is a measure of the number of inbound links, weighted by their own ranking (and normalized by $1/\ell_i$).

There is, however, another interpretation, and it brings us back to Markov chains.

Let P be the matrix given by $P(i, j) = \mathbf{1}\{i \to j\}/\ell_i$ where $\mathbf{1}\{i \to j\} = 1$ if i has a link to j and zero otherwise.

The matrix P is a stochastic matrix provided that each page has at least one link.

With this definition of P we have

$$r_j = \sum_{i \in L_j} \frac{r_i}{\ell_i} = \sum_{\text{all } i} \mathbf{1}\{i \to j\} \frac{r_i}{\ell_i} = \sum_{\text{all } i} P(i,j) r_i$$

Writing r for the row vector of rankings, this becomes $r = rP$.

Hence r is the stationary distribution of the stochastic matrix P.

Let's think of $P(i,j)$ as the probability of "moving" from page i to page j.

The value $P(i,j)$ has the interpretation

- $P(i,j) = 1/k$ if i has k outbound links and j is one of them
- $P(i,j) = 0$ if i has no direct link to j

Thus, motion from page to page is that of a web surfer who moves from one page to another by randomly clicking on one of the links on that page.

Here "random" means that each link is selected with equal probability.

Since r is the stationary distribution of P, assuming that the uniform ergodicity condition is valid, we *can interpret* r_j as the fraction of time that a (very persistent) random surfer spends at page j.

Your exercise is to apply this ranking algorithm to the graph pictured above and return the list of pages ordered by rank.

There is a total of 14 nodes (i.e., web pages), the first named `a` and the last named `n`.

A typical line from the file has the form

```
d -> h;
```

This should be interpreted as meaning that there exists a link from `d` to `h`.

The data for this graph is shown below, and read into a file called `web_graph_data.txt` when the cell is executed.

```
%%file web_graph_data.txt
a -> d;
a -> f;
b -> j;
b -> k;
b -> m;
c -> c;
c -> g;
c -> j;
c -> m;
d -> f;
d -> h;
d -> k;
e -> d;
e -> h;
e -> l;
f -> a;
f -> b;
f -> j;
f -> l;
g -> b;
g -> j;
h -> d;
h -> g;
h -> l;
```

```
h -> m;
i -> g;
i -> h;
i -> n;
j -> e;
j -> i;
j -> k;
k -> n;
l -> m;
m -> g;
n -> c;
n -> j;
n -> m;
```

```
Writing web_graph_data.txt
```

To parse this file and extract the relevant information, you can use regular expressions.

The following code snippet provides a hint as to how you can go about this

```
import re
re.findall(r'\w', 'x +++ y ****** z')  # \w matches alphanumerics
```

```
['x', 'y', 'z']
```

```
re.findall(r'\w', 'a ^^ b &&& $$ c')
```

```
['a', 'b', 'c']
```

When you solve for the ranking, you will find that the highest ranked node is in fact `g`, while the lowest is `a`.

Solution

Here is one solution:

```
"""
Return list of pages, ordered by rank
"""
import re
from operator import itemgetter

infile = 'web_graph_data.txt'
alphabet = 'abcdefghijklmnopqrstuvwxyz'

n = 14 # Total number of web pages (nodes)

# Create a matrix Q indicating existence of links
#  * Q[i, j] = 1 if there is a link from i to j
#  * Q[i, j] = 0 otherwise
Q = np.zeros((n, n), dtype=int)
with open(infile) as f:
    edges = f.readlines()
for edge in edges:
    from_node, to_node = re.findall(r'\w', edge)
    i, j = alphabet.index(from_node), alphabet.index(to_node)
    Q[i, j] = 1
# Create the corresponding Markov matrix P
P = np.empty((n, n))
```

```
for i in range(n):
    P[i, :] = Q[i, :] / Q[i, :].sum()
mc = MarkovChain(P)
# Compute the stationary distribution r
r = mc.stationary_distributions[0]
ranked_pages = {alphabet[i] : r[i] for i in range(n)}
# Print solution, sorted from highest to lowest rank
print('Rankings\n ***')
for name, rank in sorted(ranked_pages.items(), key=itemgetter(1), reverse=1):
    print(f'{name}: {rank:.4}')
```

```
Rankings
 ***
g: 0.1607
j: 0.1594
m: 0.1195
n: 0.1088
k: 0.09106
b: 0.08326
e: 0.05312
i: 0.05312
c: 0.04834
h: 0.0456
l: 0.03202
d: 0.03056
f: 0.01164
a: 0.002911
```

Exercise 41.9.3

In numerical work, it is sometimes convenient to replace a continuous model with a discrete one.

In particular, Markov chains are routinely generated as discrete approximations to AR(1) processes of the form

$$y_{t+1} = \rho y_t + u_{t+1}$$

Here u_t is assumed to be IID and $N(0, \sigma_u^2)$.

The variance of the stationary probability distribution of $\{y_t\}$ is

$$\sigma_y^2 := \frac{\sigma_u^2}{1 - \rho^2}$$

Tauchen's method [Tauchen, 1986] is the most common method for approximating this continuous state process with a finite state Markov chain.

A routine for this already exists in QuantEcon.py but let's write our own version as an exercise.

As a first step, we choose

- n, the number of states for the discrete approximation
- m, an integer that parameterizes the width of the state space

Next, we create a state space $\{x_0, \ldots, x_{n-1}\} \subset \mathbb{R}$ and a stochastic $n \times n$ matrix P such that

- $x_0 = -m\,\sigma_y$
- $x_{n-1} = m\,\sigma_y$

- $x_{i+1} = x_i + s$ where $s = (x_{n-1} - x_0)/(n-1)$

Let F be the cumulative distribution function of the normal distribution $N(0, \sigma_u^2)$.

The values $P(x_i, x_j)$ are computed to approximate the AR(1) process — omitting the derivation, the rules are as follows:

1. If $j = 0$, then set

$$P(x_i, x_j) = P(x_i, x_0) = F(x_0 - \rho x_i + s/2)$$

2. If $j = n - 1$, then set

$$P(x_i, x_j) = P(x_i, x_{n-1}) = 1 - F(x_{n-1} - \rho x_i - s/2)$$

3. Otherwise, set

$$P(x_i, x_j) = F(x_j - \rho x_i + s/2) - F(x_j - \rho x_i - s/2)$$

The exercise is to write a function `approx_markov(rho, sigma_u, m=3, n=7)` that returns $\{x_0, \ldots, x_{n-1}\} \subset \mathbb{R}$ and $n \times n$ matrix P as described above.

- Even better, write a function that returns an instance of QuantEcon.py's MarkovChain class.

Solution

A solution from the QuantEcon.py library can be found here.

CHAPTER

FORTYTWO

INVENTORY DYNAMICS

Contents

- *Inventory Dynamics*
  - *Overview*
  - *Sample paths*
  - *Marginal distributions*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install jax
```

42.1 Overview

In this lecture, we will study the time path of inventories for firms that follow so-called s-S inventory dynamics.

Such firms

1. wait until inventory falls below some level s and then
2. order sufficient quantities to bring their inventory back up to capacity S.

These kinds of policies are common in practice and are also optimal in certain circumstances.

A review of early literature and some macroeconomic implications can be found in [Caplin, 1985].

Here, our main aim is to learn more about simulation, time series, and Markov dynamics.

While our Markov environment and many of the concepts we consider are related to those found in our *lecture on finite Markov chains*, the state space is a continuum in the current application.

Let's start with some imports

```
from functools import partial
from typing import NamedTuple
import jax
import jax.numpy as jnp
from jax import random
```

(continues on next page)

(continued from previous page)

```
import matplotlib.pyplot as plt
from sklearn.neighbors import KernelDensity
```

42.2 Sample paths

Consider a firm with inventory X_t.

The firm waits until $X_t \leq s$ and then restocks up to S units.

It faces stochastic demand $\{D_t\}$, which we assume is IID.

With notation $a^+ := \max\{a, 0\}$, inventory dynamics can be written as

$$X_{t+1} = \begin{cases} (S - D_{t+1})^+ & \text{if } X_t \leq s \\ (X_t - D_{t+1})^+ & \text{if } X_t > s \end{cases}$$

In what follows, we will assume that each D_t is lognormal, so that

$$D_t = \exp(\mu + \sigma Z_t)$$

where μ and σ are parameters and $\{Z_t\}$ is IID and standard normal.

Here's a class that stores parameters and generates time paths for inventory.

```
class Firm(NamedTuple):
    s: int    # restock trigger level
    S: int    # capacity
    μ: float  # shock location parameter
    σ: float  # shock scale parameter
```

```
@partial(jax.jit, static_argnames="sim_length")
def sim_inventory_path(firm, x_init, key, sim_length):
    """
    Simulate an inventory path of length sim_length from a single key.

    A fresh shock is generated each period by folding the period index
    into key, so callers pass one key per path rather than a whole array.

    Args:
        firm: Firm object
        x_init: Initial inventory level
        key: A single JAX random key
        sim_length: Number of periods to simulate

    Returns:
        Array of inventory levels [X_0, X_1, ..., X_{sim_length-1}]
    """

    def update(t, X):
        x = X[t - 1]
        Z = random.normal(random.fold_in(key, t))    # fresh shock for period t
        D = jnp.exp(firm.μ + firm.σ * Z)
        x_new = jnp.where(
            x <= firm.s,
```

(continues on next page)

(continued from previous page)

```
            jnp.maximum(firm.S - D, 0.0),    # restock to S, then meet demand
            jnp.maximum(x - D, 0.0),         # just meet demand
        )
        return X.at[t].set(x_new)

    X = jnp.zeros(sim_length).at[0].set(x_init)
    return jax.lax.fori_loop(1, sim_length, update, X)
```

Note

Writing `X.at[t].set(x_new)` returns a *new* array rather than mutating `X` in place, in keeping with JAX's functional style. This is not as wasteful as it looks: inside a `jax.jit`-compiled function the XLA compiler sees that the old array is no longer needed and performs the update in place, so no fresh array is allocated each period.

```
firm = Firm(s=10, S=100, μ=1.0, σ=0.5)
```

```
sim_length = 100
x_init = 50
X = sim_inventory_path(firm, x_init, random.key(21), sim_length)
```

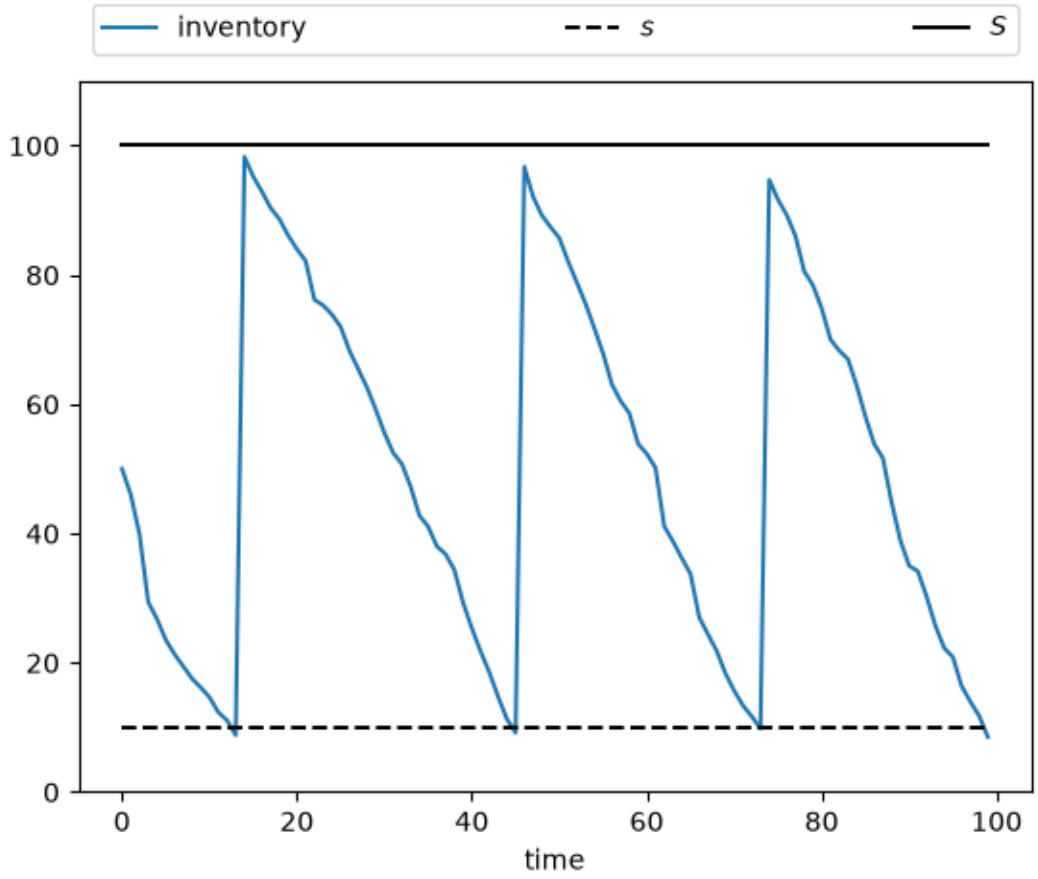
Let's run a first simulation, of a single path:

```
s, S = firm.s, firm.S

fig, ax = plt.subplots()
bbox = (0.0, 1.02, 1.0, 0.102)
legend_args = {"ncol": 3, "bbox_to_anchor": bbox, "loc": 3, "mode": "expand"}

ax.plot(X, label="inventory")
ax.plot(jnp.full(sim_length, s), "k--", label="$s$")
ax.plot(jnp.full(sim_length, S), "k-", label="$S$")
ax.set_ylim(0, S + 10)
ax.set_xlabel("time")
ax.legend(**legend_args)

plt.show()
```



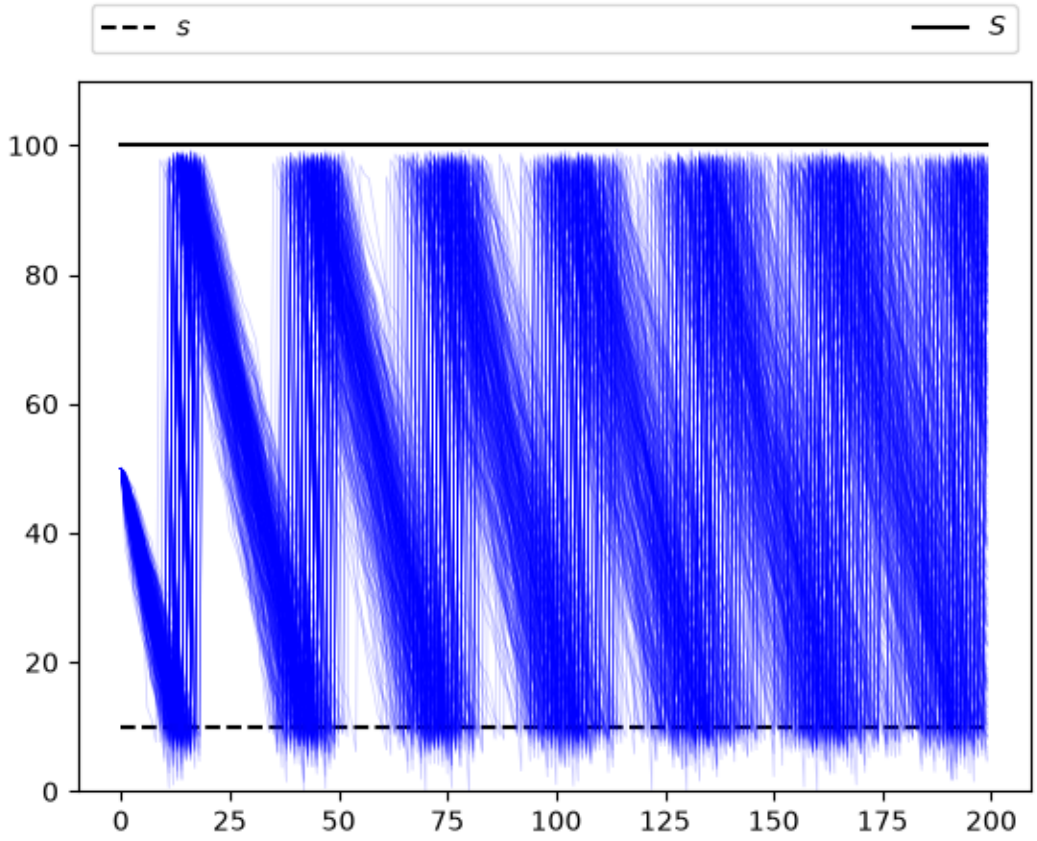
Now let's simulate multiple paths in order to build a more complete picture of the probabilities of different outcomes:

```
sim_length = 200
fig, ax = plt.subplots()

ax.plot(jnp.full(sim_length, s), "k--", label="$s$")
ax.plot(jnp.full(sim_length, S), "k-", label="$S$")
ax.set_ylim(0, S + 10)
ax.legend(**legend_args)

for i in range(400):
    X = sim_inventory_path(firm, x_init, random.key(i), sim_length)
    ax.plot(X, "b", alpha=0.2, lw=0.5)

plt.show()
```



42.3 Marginal distributions

Now let's look at the marginal distribution ψ_T of X_T for some fixed T.

We will do this by generating many draws of X_T given initial condition X_0.

With these draws of X_T we can build up a picture of its distribution ψ_T.

Here's one visualization, with $T = 50$.

```
T = 50
M = 200  # Number of draws

ymin, ymax = 0, S + 10

fig, axes = plt.subplots(1, 2, figsize=(11, 6))

for ax in axes:
    ax.grid(alpha=0.4)

ax = axes[0]

ax.set_ylim(ymin, ymax)
ax.set_ylabel("$X_t$", fontsize=16)
ax.vlines((T,), -1.5, 1.5)
```

(continues on next page)

(continued from previous page)

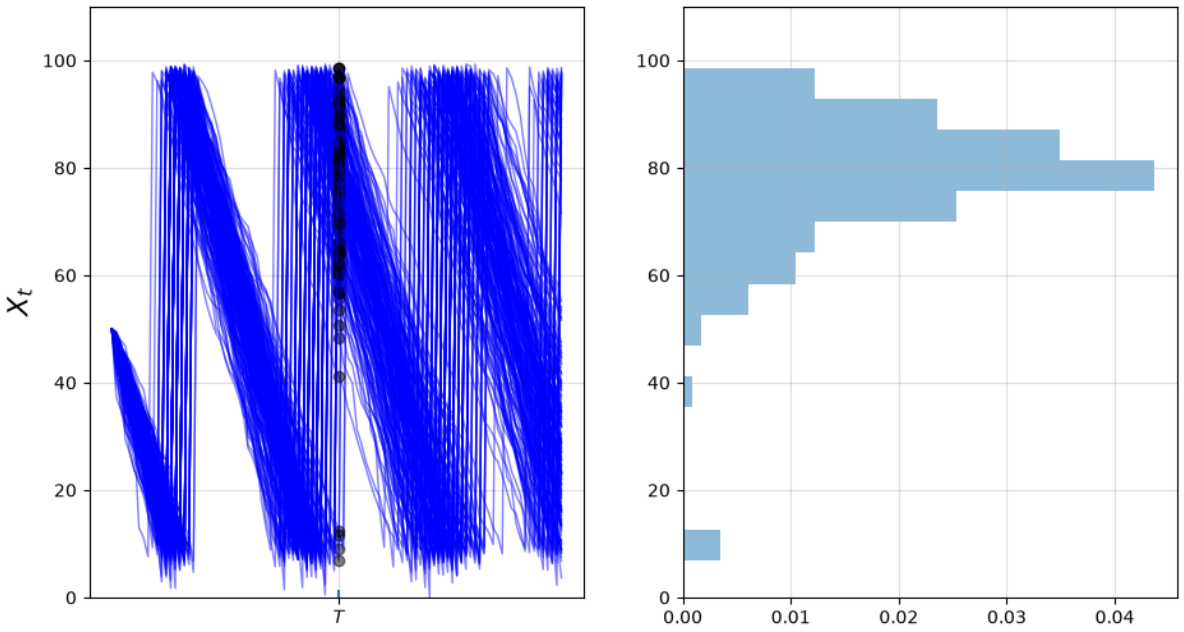
```
ax.set_xticks((T,))
ax.set_xticklabels((r"$T$",))

sample = []
for m in range(M):
    X = sim_inventory_path(firm, x_init, random.key(m), 2 * T)
    ax.plot(X, "b-", lw=1, alpha=0.5)
    ax.plot((T,), (X[T],), "ko", alpha=0.5)
    sample.append(X[T])

axes[1].set_ylim(ymin, ymax)

axes[1].hist(
    sample,
    bins=16,
    density=True,
    orientation="horizontal",
    histtype="bar",
    alpha=0.5,
)

plt.show()
```



We can build up a clearer picture by drawing more samples

```
T = 50
M = 50_000

fig, ax = plt.subplots()

# Draw all M paths at once by vectorizing over one key per path
keys = random.split(random.key(0), M)
paths = jax.vmap(
    sim_inventory_path, in_axes=(None, None, 0, None)
```

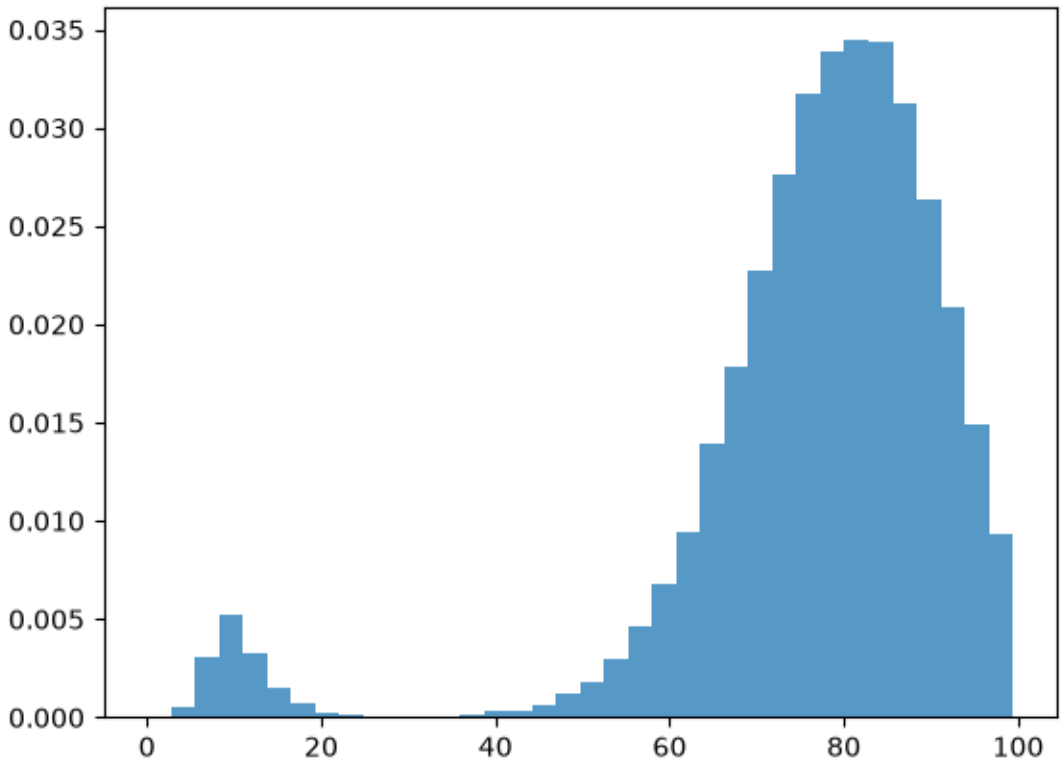
(continues on next page)

(continued from previous page)

```
)(firm, x_init, keys, T + 1)
sample = paths[:, T]

ax.hist(sample, bins=36, density=True, histtype="bar", alpha=0.75)

plt.show()
```



Note that the distribution is bimodal

- Most firms have restocked twice but a few have restocked only once (see figure with paths above).
- Firms in the second category have lower inventory.

We can also approximate the distribution using a kernel density estimator.

Kernel density estimators can be thought of as smoothed histograms.

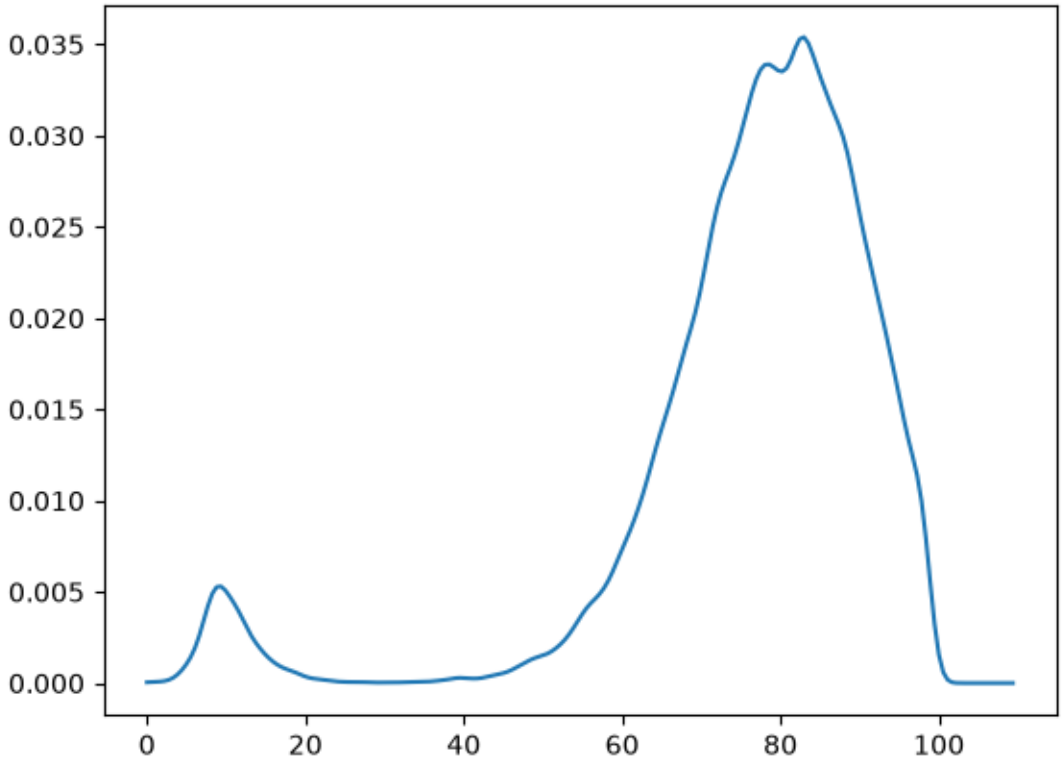
They are preferable to histograms when the distribution being estimated is likely to be smooth.

We will use a kernel density estimator from scikit-learn

```
def plot_kde(sample, ax, label=""):
    xmin, xmax = 0.9 * sample.min(), 1.1 * sample.max()
    xgrid = jnp.linspace(xmin, xmax, 200)
    kde = KernelDensity(kernel="gaussian").fit(sample[:, None])
    log_dens = kde.score_samples(xgrid[:, None])

    ax.plot(xgrid, jnp.exp(log_dens), label=label)
```

```
fig, ax = plt.subplots()
plot_kde(sample, ax)
plt.show()
```



The allocation of probability density is similar to what was shown by the histogram just above.

42.4 Exercises

Exercise 42.4.1

This model is asymptotically stationary, with a unique stationary distribution.

(See the discussion of stationarity in our lecture on AR(1) processes for background — the fundamental concepts are the same.)

In particular, the sequence of marginal distributions $\{\psi_t\}$ is converging to a unique limiting distribution that does not depend on initial conditions.

Although we will not prove this here, we can investigate it using simulation.

Your task is to generate and plot the sequence $\{\psi_t\}$ at times $t = 10, 50, 250, 500, 750$ based on the discussion above.

(The kernel density estimator is probably the best way to present each distribution.)

You should see convergence, in the sense that differences between successive distributions are getting smaller.

Try different initial conditions to verify that, in the long run, the distribution is invariant across initial conditions.

Solution

Below is one possible solution:

The computations involve a lot of CPU cycles so we have tried to write the code efficiently using `jax.jit` and `jax.vmap` to run on CPU/GPU.

```
@jax.jit
def simulate_firm_forward(firm, x_init, key, num_periods):
    """
    Simulate a single firm num_periods steps forward and return its
    final inventory level. A fresh shock is drawn each period from key.
    """

    def update(t, x):
        Z = random.normal(random.fold_in(key, t))
        D = jnp.exp(firm.μ + firm.σ * Z)
        return jnp.where(
            x <= firm.s,
            jnp.maximum(firm.S - D, 0.0),
            jnp.maximum(x - D, 0.0),
        )

    return jax.lax.fori_loop(0, num_periods, update, x_init * 1.0)


# Vectorize over firms: each has its own initial level and its own key
simulate_firms_forward = jax.vmap(
    simulate_firm_forward, in_axes=(None, 0, 0, None)
)
```

```
def shift_firms_forward(firm, current_inventory_levels, num_periods, key):
    """
    Shift multiple firms forward by num_periods using JAX vectorization.
    Returns:
        Array of new inventory levels after num_periods
    """

    # Generate one independent random key per firm
    num_firms = len(current_inventory_levels)
    firm_keys = random.split(key, num_firms)
    # Run simulation for all firms in parallel
    new_inventory_levels = simulate_firms_forward(
        firm, current_inventory_levels, firm_keys, num_periods
    )

    return new_inventory_levels
```

```
x_init = 50
num_firms = 50_000

sample_dates = 0, 10, 50, 250, 500, 750

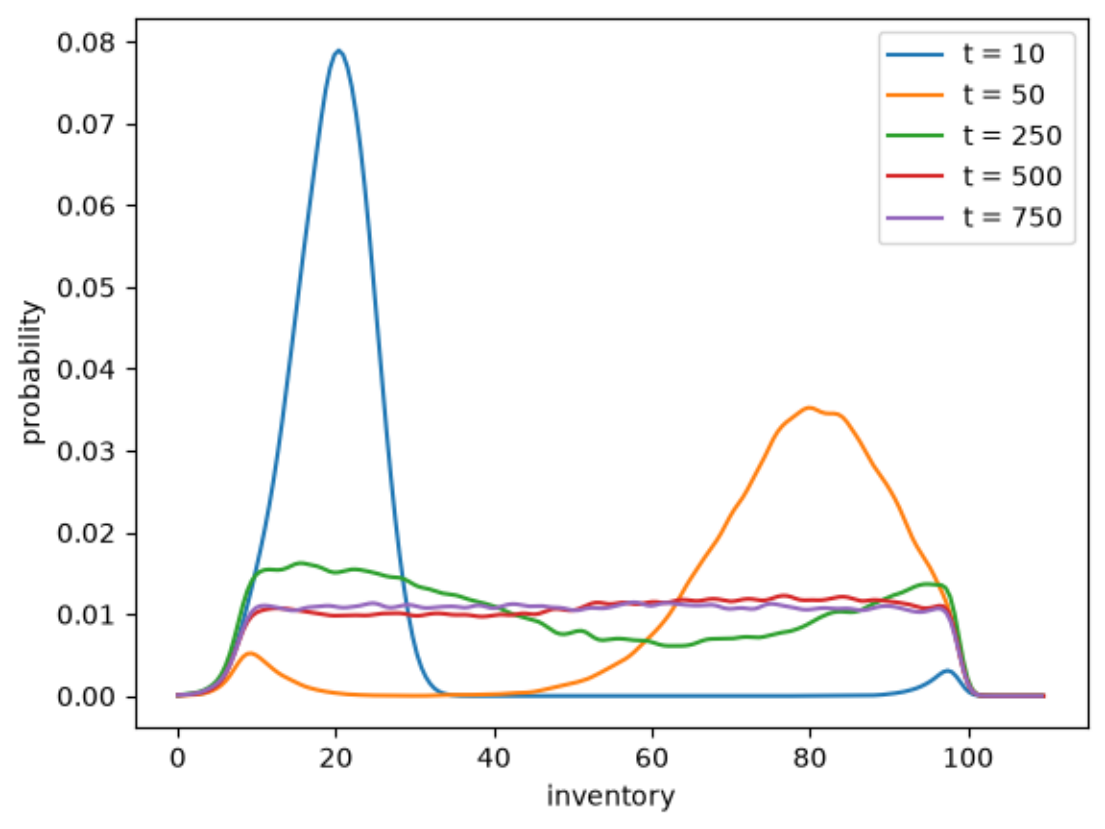
first_diffs = jnp.diff(jnp.array(sample_dates))

fig, ax = plt.subplots()

X = jnp.full(num_firms, x_init)

current_date = 0
for d in first_diffs:
    X = shift_firms_forward(firm, X, d, random.key(current_date + 1))
    current_date += d
    plot_kde(X, ax, label=f"t = {current_date}")

ax.set_xlabel("inventory")
```



Notice that by $t = 500$ or $t = 750$ the densities are barely changing.

We have reached a reasonable approximation of the stationary density.

You can convince yourself that initial conditions don't matter by testing a few of them.

For example, try rerunning the code above with all firms starting at $X_0 = 20$ or $X_0 = 80$.

Exercise 42.4.2

Using simulation, calculate the probability that firms that start with $X_0 = 70$ need to order twice or more in the first 50 periods.

You will need a large sample size to get an accurate reading.

Solution

Here is one solution.

Again, the computations are relatively intensive so we have written a specialized JAX-jitted function and using `jax.vmap` to use parallelization across firms.

Note the time the routine takes to run, as well as the output.

```
@jax.jit
def simulate_firm_restocks(firm, x_init, key, num_periods):
    """
```

```
    Simulate a single firm for num_periods and report whether it
    restocked more than once. A fresh shock is drawn each period from key.

    Returns:
        1 if the firm restocks > 1 times, 0 otherwise
    """

    def update(t, carry):
        x, restock_count = carry
        Z = random.normal(random.fold_in(key, t))
        D = jnp.exp(firm.μ + firm.σ * Z)
        restock = x <= firm.s
        x_new = jnp.where(
            restock,
            jnp.maximum(firm.S - D, 0.0),
            jnp.maximum(x - D, 0.0),
        )
        return x_new, restock_count + restock

    # Carry the inventory level and a running restock count
    _, total_restocks = jax.lax.fori_loop(
        0, num_periods, update, (x_init * 1.0, 0)
    )
    return (total_restocks > 1).astype(jnp.int32)


# Vectorize the simulation across all firms (one key each)
simulate_firms_restocks = jax.vmap(
    simulate_firm_restocks, in_axes=(None, None, 0, None)
)



def compute_freq(
    firm, x_init=70, sim_length=50, num_firms=1_000_000, key=random.key(2)
):
    """
    Compute the frequency of firms that restock 2 or more times using JAX.

    Args:
        firm: Firm dataclass
        x_init: Initial inventory level for all firms
        sim_length: Length of simulation for each firm
        num_firms: Number of firms to simulate
        key: JAX random key

    Returns:
        Fraction of firms that restock 2 or more times
    """
    # Generate one independent random key per firm
    firm_keys = random.split(key, num_firms)
    # Run simulation for all firms
    restock_indicators = simulate_firms_restocks(
        firm, x_init, firm_keys, sim_length
    )
    # Compute frequency (fraction of firms that restocked > 1 times)
    frequency = jnp.mean(restock_indicators)
    return frequency
```

```
%%time

freq = compute_freq(firm)
print(f"Frequency of at least two stock outs = {freq}")
```

```
Frequency of at least two stock outs = 0.44668900966644287
CPU times: user 1.31 s, sys: 72.2 ms, total: 1.38 s
Wall time: 749 ms
```

CHAPTER

FORTYTHREE

LINEAR STATE SPACE MODELS

Contents

- *Linear State Space Models*
  - *Overview*
  - *The Linear State Space Model*
  - *Distributions and Moments*
  - *Stationarity and Ergodicity*
  - *Noisy Observations*
  - *Prediction*
  - *Code*
  - *Exercises*

> "We may regard the present state of the universe as the effect of its past and the cause of its future" – Marquis de Laplace

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

43.1 Overview

This lecture introduces the **linear state space** dynamic system.

The linear state space system is a generalization of the scalar AR(1) process we studied before.

This model is a workhorse that carries a powerful theory of prediction.

Its many applications include:

- representing dynamics of higher-order linear systems
- predicting the position of a system j steps into the future
- predicting a geometric sum of future values of a variable like
  - non-financial income

  - dividends on a stock
  - the money supply
  - a government deficit or surplus, etc.
- key ingredient of useful models
  - Friedman's permanent income model of consumption smoothing.
  - Barro's model of smoothing total tax collections.
  - Rational expectations version of Cagan's model of hyperinflation.
  - Sargent and Wallace's "unpleasant monetarist arithmetic," etc.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
from quantecon import LinearStateSpace
from scipy.stats import norm
import random
```

43.2 The Linear State Space Model

The objects in play are:

- An $n \times 1$ vector x_t denoting the **state** at time $t = 0, 1, 2, \ldots$.
- An IID sequence of $m \times 1$ random vectors $w_t \sim N(0, I)$.
- A $k \times 1$ vector y_t of **observations** at time $t = 0, 1, 2, \ldots$.
- An $n \times n$ matrix A called the **transition matrix**.
- An $n \times m$ matrix C called the **volatility matrix**.
- A $k \times n$ matrix G sometimes called the **output matrix**.

Here is the linear state-space system

$$
\begin{aligned}
x_{t+1} &= Ax_t + Cw_{t+1} \\
y_t &= Gx_t \\
x_0 &\sim N(\mu_0, \Sigma_0)
\end{aligned}
$$

43.2.1 Primitives

The primitives of the model are

1. the matrices A, C, G
2. shock distribution, which we have specialized to $N(0, I)$
3. the distribution of the initial condition x_0, which we have set to $N(\mu_0, \Sigma_0)$

Given A, C, G and draws of x_0 and $w_1, w_2, \ldots$, the model (43.1) pins down the values of the sequences $\{x_t\}$ and $\{y_t\}$.

Even without these draws, the primitives 1–3 pin down the **probability distributions** of $\{x_t\}$ and $\{y_t\}$.

Later we'll see how to compute these distributions and their moments.

Martingale Difference Shocks

We've made the common assumption that the shocks are independent standardized normal vectors.

But some of what we say will be valid under the assumption that $\{w_{t+1}\}$ is a **martingale difference sequence**.

A martingale difference sequence is a sequence that is zero mean when conditioned on past information.

In the present case, since $\{x_t\}$ is our state sequence, this means that it satisfies

$$\mathbb{E}[w_{t+1}|x_t, x_{t-1}, ...] = 0$$

This is a weaker condition than that $\{w_t\}$ is IID with $w_{t+1} \sim N(0, I)$.

43.2.2 Examples

By appropriate choice of the primitives, a variety of dynamics can be represented in terms of the linear state space model.

The following examples help to highlight this point.

They also illustrate the wise dictum *finding the state is an art*.

Second-order Difference Equation

Let $\{y_t\}$ be a deterministic sequence that satisfies

$$y_{t+1} = \phi_0 + \phi_1 y_t + \phi_2 y_{t-1} \quad \text{s.t.} \quad y_0, y_{-1} \text{ given} \tag{43.1}$$

To map (43.1) into our state space system (43.1), we set

$$x_t = \begin{bmatrix} 1 \\ y_t \\ y_{t-1} \end{bmatrix} \qquad A = \begin{bmatrix} 1 & 0 & 0 \\ \phi_0 & \phi_1 & \phi_2 \\ 0 & 1 & 0 \end{bmatrix} \qquad C = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \qquad G = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

You can confirm that under these definitions, (43.1) and (43.1) agree.

The next figure shows the dynamics of this process when $\phi_0 = 1.1, \phi_1 = 0.8, \phi_2 = -0.8, y_0 = y_{-1} = 1$.

```
def plot_lss(A,
             C,
             G,
             n=3,
             ts_length=50):

    ar = LinearStateSpace(A, C, G, mu_0=np.ones(n))
    x, y = ar.simulate(ts_length)

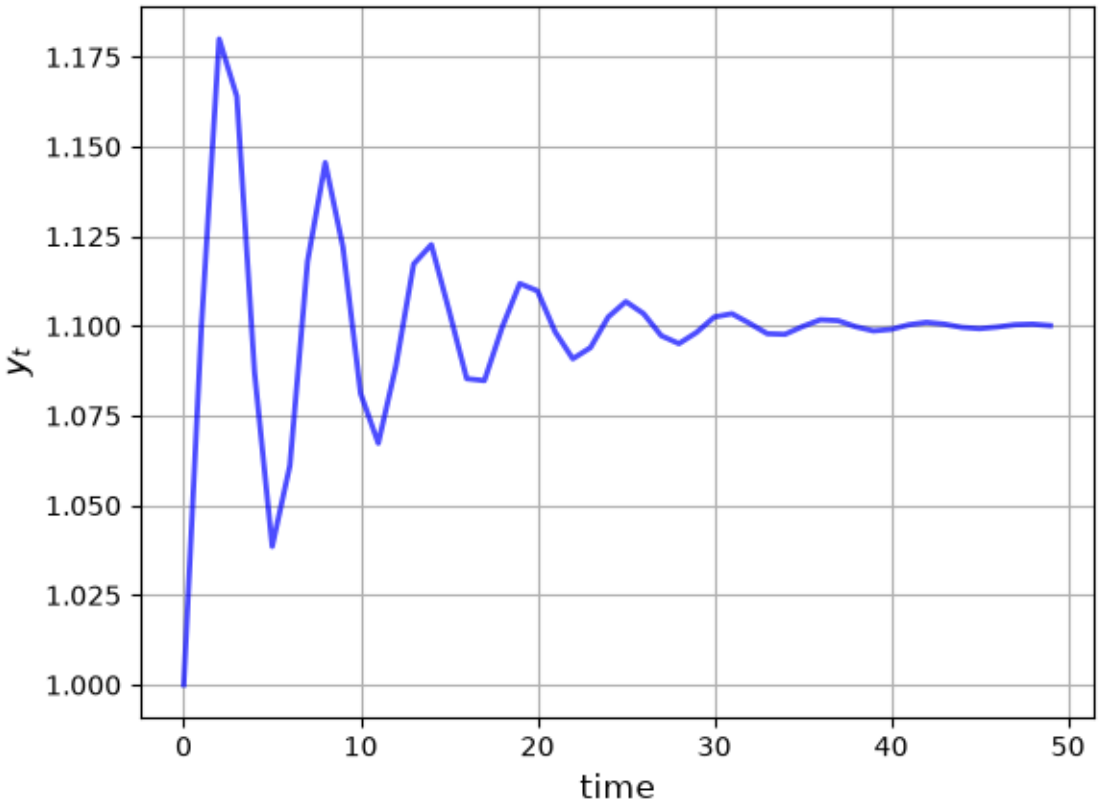
    fig, ax = plt.subplots()
    y = y.flatten()
    ax.plot(y, 'b-', lw=2, alpha=0.7)
    ax.grid()
    ax.set_xlabel('time', fontsize=12)
    ax.set_ylabel('$y_t$', fontsize=12)
    plt.show()
```

```
ϕ_0, ϕ_1, ϕ_2 = 1.1, 0.8, -0.8

A = [[1,     0,     0  ],
     [ϕ_0,   ϕ_1,   ϕ_2],
     [0,     1,     0  ]]

C = np.zeros((3, 1))
G = [0, 1, 0]

plot_lss(A, C, G)
```



Later you'll be asked to recreate this figure.

Univariate Autoregressive Processes

We can use (43.1) to represent the model

$$y_{t+1} = \phi_1 y_t + \phi_2 y_{t-1} + \phi_3 y_{t-2} + \phi_4 y_{t-3} + \sigma w_{t+1} \tag{43.2}$$

where $\{w_t\}$ is IID and standard normal.

To put this in the linear state space format we take $x_t = \begin{bmatrix} y_t & y_{t-1} & y_{t-2} & y_{t-3} \end{bmatrix}'$ and

$$A = \begin{bmatrix} \phi_1 & \phi_2 & \phi_3 & \phi_4 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \qquad C = \begin{bmatrix} \sigma \\ 0 \\ 0 \\ 0 \end{bmatrix} \qquad G = \begin{bmatrix} 1 & 0 & 0 & 0 \end{bmatrix}$$

The matrix A has the form of the **companion matrix** to the vector $\begin{bmatrix}\phi_1 & \phi_2 & \phi_3 & \phi_4\end{bmatrix}$.

The next figure shows the dynamics of this process when

$$\phi_1 = 0.5, \phi_2 = -0.2, \phi_3 = 0, \phi_4 = 0.5, \sigma = 0.2, y_0 = y_{-1} = y_{-2} = y_{-3} = 1$$

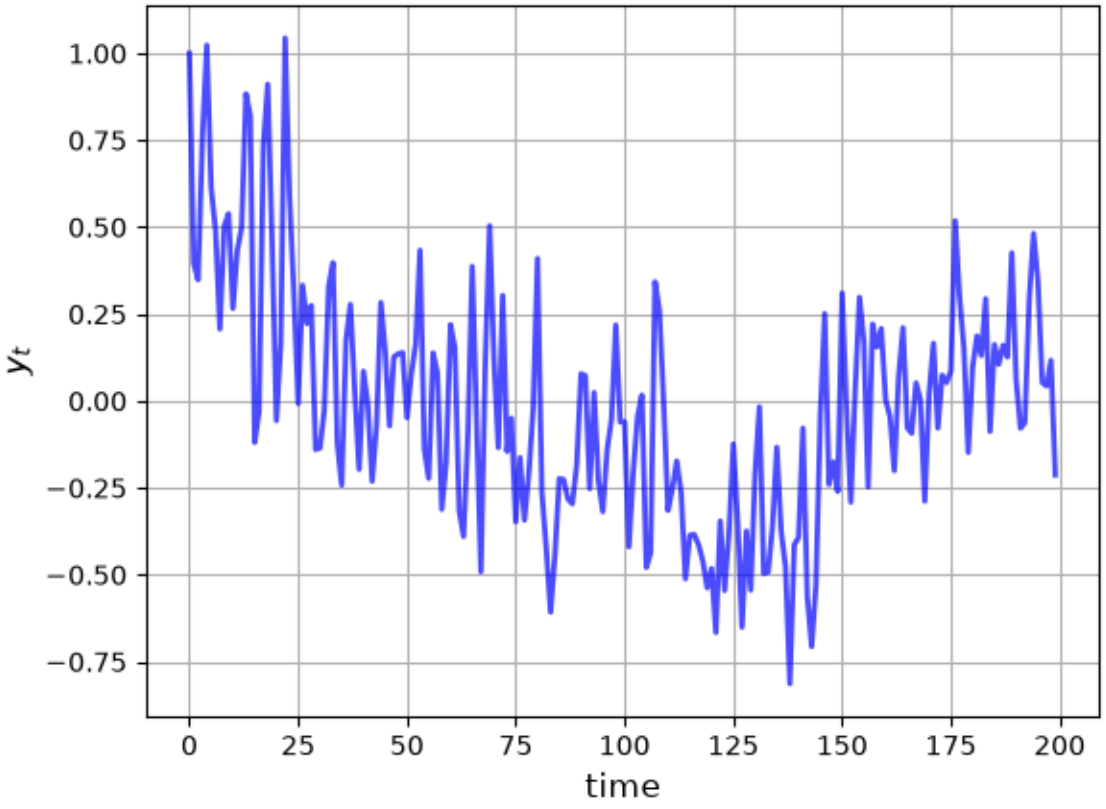
```
ϕ_1, ϕ_2, ϕ_3, ϕ_4 = 0.5, -0.2, 0, 0.5
σ = 0.2

A_1 = [[ϕ_1,   ϕ_2,   ϕ_3,   ϕ_4],
       [1,     0,     0,     0  ],
       [0,     1,     0,     0  ],
       [0,     0,     1,     0  ]]

C_1 = [[σ],
       [0],
       [0],
       [0]]

G_1 = [1, 0, 0, 0]

plot_lss(A_1, C_1, G_1, n=4, ts_length=200)
```



Vector Autoregressions

Now suppose that

- y_t is a $k \times 1$ vector
- ϕ_j is a $k \times k$ matrix and
- w_t is $k \times 1$

Then (43.2) is termed a **vector autoregression**.

To map this into (43.1), we set

$$x_t = \begin{bmatrix} y_t \\ y_{t-1} \\ y_{t-2} \\ y_{t-3} \end{bmatrix} \quad A = \begin{bmatrix} \phi_1 & \phi_2 & \phi_3 & \phi_4 \\ I & 0 & 0 & 0 \\ 0 & I & 0 & 0 \\ 0 & 0 & I & 0 \end{bmatrix} \quad C = \begin{bmatrix} \sigma \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad G = \begin{bmatrix} I & 0 & 0 & 0 \end{bmatrix}$$

where I is the $k \times k$ identity matrix and σ is a $k \times k$ matrix.

Seasonals

We can use (43.1) to represent

1. the **deterministic seasonal** $y_t = y_{t-4}$
2. the **indeterministic seasonal** $y_t = \phi_4 y_{t-4} + w_t$

In fact, both are special cases of (43.2).

With the deterministic seasonal, the transition matrix becomes

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

It is easy to check that $A^4 = I$, which implies that x_t is strictly periodic with period 4:[1]

$$x_{t+4} = x_t$$

Such an x_t process can be used to model deterministic seasonals in quarterly time series.

The *indeterministic* seasonal produces recurrent, but aperiodic, seasonal fluctuations.

Time Trends

The model $y_t = at + b$ is known as a **linear time trend**.

We can represent this model in the linear state space form by taking

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \qquad C = \begin{bmatrix} 0 \\ 0 \end{bmatrix} \qquad G = \begin{bmatrix} a & b \end{bmatrix} \tag{43.3}$$

and starting at initial condition $x_0 = \begin{bmatrix} 0 & 1 \end{bmatrix}'$.

In fact, it's possible to use the state-space system to represent polynomial trends of any order.

[1] The eigenvalues of A are $(1, -1, i, -i)$.

For instance, we can represent the model $y_t = at^2 + bt + c$ in the linear state space form by taking

$$A = \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix} \qquad C = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \qquad G = \begin{bmatrix} 2a & a+b & c \end{bmatrix}$$

and starting at initial condition $x_0 = \begin{bmatrix} 0 & 0 & 1 \end{bmatrix}'$.

It follows that

$$A^t = \begin{bmatrix} 1 & t & t(t-1)/2 \\ 0 & 1 & t \\ 0 & 0 & 1 \end{bmatrix}$$

Then $x_t' = \begin{bmatrix} t(t-1)/2 & t & 1 \end{bmatrix}$. You can now confirm that $y_t = Gx_t$ has the correct form.

43.2.3 Moving Average Representations

A nonrecursive expression for x_t as a function of $x_0, w_1, w_2, \ldots, w_t$ can be found by using (43.1) repeatedly to obtain

$$\begin{aligned} x_t &= Ax_{t-1} + Cw_t \\ &= A^2 x_{t-2} + ACw_{t-1} + Cw_t \\ &\quad \vdots \\ &= \sum_{j=0}^{t-1} A^j C w_{t-j} + A^t x_0 \end{aligned}$$

Representation (43.4) is a **moving average** representation.

It expresses $\{x_t\}$ as a linear function of

1. current and past values of the process $\{w_t\}$ and
2. the initial condition x_0

As an example of a moving average representation, let the model be

$$A = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \qquad C = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$$

You will be able to show that $A^t = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}$ and $A^j C = \begin{bmatrix} 1 & 0 \end{bmatrix}'$.

Substituting into the moving average representation (43.4), we obtain

$$x_{1t} = \sum_{j=0}^{t-1} w_{t-j} + \begin{bmatrix} 1 & t \end{bmatrix} x_0$$

where x_{1t} is the first entry of x_t.

The first term on the right is a cumulated sum of martingale differences and is therefore a martingale.

The second term is a translated linear function of time.

For this reason, x_{1t} is called a **martingale with drift**.

43.3 Distributions and Moments

43.3.1 Unconditional Moments

Using (43.1), it's easy to obtain expressions for the (unconditional) means of x_t and y_t.

We'll explain what *unconditional* and *conditional* mean soon.

Letting $\mu_t := \mathbb{E}[x_t]$ and using linearity of expectations, we find that

$$\mu_{t+1} = A\mu_t \quad \text{with} \quad \mu_0 \text{ given} \tag{43.4}$$

Here μ_0 is a primitive given in (43.1).

The variance-covariance matrix of x_t is $\Sigma_t := \mathbb{E}[(x_t - \mu_t)(x_t - \mu_t)']$.

Using $x_{t+1} - \mu_{t+1} = A(x_t - \mu_t) + Cw_{t+1}$, we can determine this matrix recursively via

$$\Sigma_{t+1} = A\Sigma_t A' + CC' \quad \text{with} \quad \Sigma_0 \text{ given} \tag{43.5}$$

As with μ_0, the matrix Σ_0 is a primitive given in (43.1).

As a matter of terminology, we will sometimes call

- μ_t the **unconditional mean** of x_t
- Σ_t the **unconditional variance-covariance matrix** of x_t

This is to distinguish μ_t and Σ_t from related objects that use conditioning information, to be defined below.

However, you should be aware that these "unconditional" moments do depend on the initial distribution $N(\mu_0, \Sigma_0)$.

Moments of the Observables

Using linearity of expectations again we have

$$\mathbb{E}[y_t] = \mathbb{E}[Gx_t] = G\mu_t \tag{43.6}$$

The variance-covariance matrix of y_t is easily shown to be

$$\text{Var}[y_t] = \text{Var}[Gx_t] = G\Sigma_t G' \tag{43.7}$$

43.3.2 Distributions

In general, knowing the mean and variance-covariance matrix of a random vector is not quite as good as knowing the full distribution.

However, there are some situations where these moments alone tell us all we need to know.

These are situations in which the mean vector and covariance matrix are all of the **parameters** that pin down the population distribution.

One such situation is when the vector in question is Gaussian (i.e., normally distributed).

This is the case here, given

1. our Gaussian assumptions on the primitives
2. the fact that normality is preserved under linear operations

In fact, it's well-known that

$$u \sim N(\bar{u}, S) \quad \text{and} \quad v = a + Bu \implies v \sim N(a + B\bar{u}, BSB') \tag{43.8}$$

In particular, given our Gaussian assumptions on the primitives and the linearity of (43.1) we can see immediately that both x_t and y_t are Gaussian for all $t \geq 0$[2].

Since x_t is Gaussian, to find the distribution, all we need to do is find its mean and variance-covariance matrix.

But in fact we've already done this, in (43.4) and (43.5).

Letting μ_t and Σ_t be as defined by these equations, we have

$$x_t \sim N(\mu_t, \Sigma_t) \tag{43.9}$$

By similar reasoning combined with (43.6) and (43.7),

$$y_t \sim N(G\mu_t, G\Sigma_t G') \tag{43.10}$$

43.3.3 Ensemble Interpretations

How should we interpret the distributions defined by (43.9)–(43.10)?

Intuitively, the probabilities in a distribution correspond to relative frequencies in a large population drawn from that distribution.

Let's apply this idea to our setting, focusing on the distribution of y_T for fixed T.

We can generate independent draws of y_T by repeatedly simulating the evolution of the system up to time T, using an independent set of shocks each time.

The next figure shows 20 simulations, producing 20 time series for $\{y_t\}$, and hence 20 draws of y_T.

The system in question is the univariate autoregressive model (43.2).

The values of y_T are represented by black dots in the left-hand figure

```
def cross_section_plot(A,
                       C,
                       G,
                       T=20,                 # Set the time
                       ymin=-0.8,
                       ymax=1.25,
                       sample_size = 20,     # 20 observations/simulations
                       n=4):                 # The number of dimensions for the initial x0

    ar = LinearStateSpace(A, C, G, mu_0=np.ones(n))

    fig, axes = plt.subplots(1, 2, figsize=(16, 5))

    for ax in axes:
        ax.grid(alpha=0.4)
        ax.set_ylim(ymin, ymax)

    ax = axes[0]
    ax.set_ylim(ymin, ymax)
    ax.set_ylabel('$y_t$', fontsize=12)
```

(continues on next page)

[2] The correct way to argue this is by induction. Suppose that x_t is Gaussian. Then (43.1) and (43.8) imply that x_{t+1} is Gaussian. Since x_0 is assumed to be Gaussian, it follows that every x_t is Gaussian. Evidently, this implies that each y_t is Gaussian.

(continued from previous page)

```
    ax.set_xlabel('time', fontsize=12)
    ax.vlines((T,), -1.5, 1.5)

    ax.set_xticks((T,))
    ax.set_xticklabels(('$T$',))

    sample = []
    for i in range(sample_size):
        rcolor = random.choice(('c', 'g', 'b', 'k'))
        x, y = ar.simulate(ts_length=T+15)
        y = y.flatten()
        ax.plot(y, color=rcolor, lw=1, alpha=0.5)
        ax.plot((T,), (y[T],), 'ko', alpha=0.5)
        sample.append(y[T])

    y = y.flatten()
    axes[1].set_ylim(ymin, ymax)
    axes[1].set_ylabel('$y_t$', fontsize=12)
    axes[1].set_xlabel('relative frequency', fontsize=12)
    axes[1].hist(sample, bins=16, density=True, orientation='horizontal', alpha=0.5)
    plt.show()
```

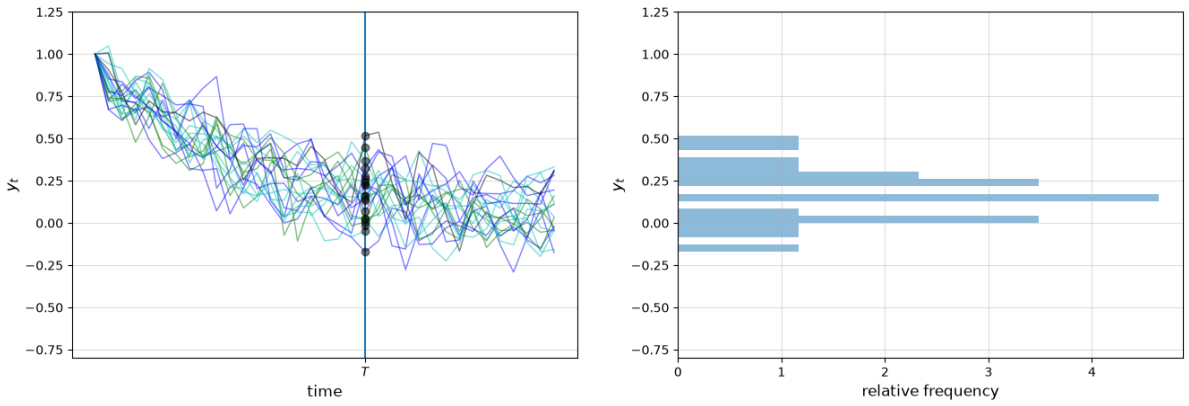
```
ϕ_1, ϕ_2, ϕ_3, ϕ_4 = 0.5, -0.2, 0, 0.5
σ = 0.1

A_2 = [[ϕ_1, ϕ_2, ϕ_3, ϕ_4],
       [1,     0,     0,     0],
       [0,     1,     0,     0],
       [0,     0,     1,     0]]

C_2 = [[σ], [0], [0], [0]]

G_2 = [1, 0, 0, 0]

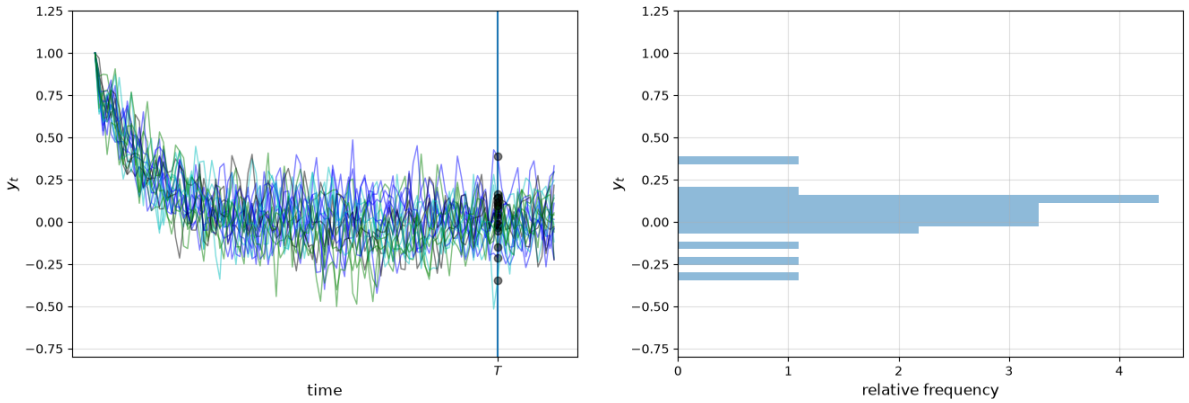
cross_section_plot(A_2, C_2, G_2)
```



In the right-hand figure, these values are converted into a rotated histogram that shows relative frequencies from our sample of 20 y_T's.

Here is another figure, this time with 100 observations

```
t = 100
cross_section_plot(A_2, C_2, G_2, T=t)
```

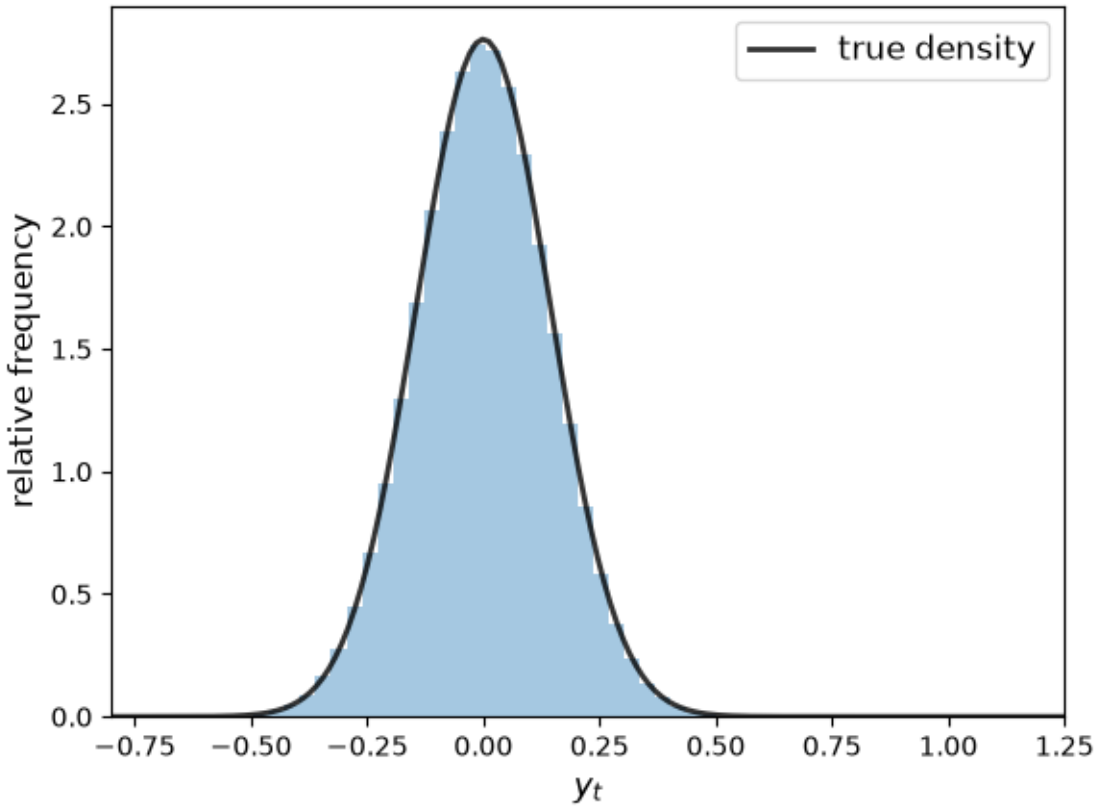


Let's now try with 500,000 observations, showing only the histogram (without rotation)

```
T = 100
ymin=-0.8
ymax=1.25
sample_size = 500_000

ar = LinearStateSpace(A_2, C_2, G_2, mu_0=np.ones(4))
fig, ax = plt.subplots()
x, y = ar.simulate(sample_size)
mu_x, mu_y, Sigma_x, Sigma_y, Sigma_yx = ar.stationary_distributions()
f_y = norm(loc=float(mu_y.item()), scale=float(np.sqrt(Sigma_y.item())))
y = y.flatten()
ygrid = np.linspace(ymin, ymax, 150)

ax.hist(y, bins=50, density=True, alpha=0.4)
ax.plot(ygrid, f_y.pdf(ygrid), 'k-', lw=2, alpha=0.8, label='true density')
ax.set_xlim(ymin, ymax)
ax.set_xlabel('$y_t$', fontsize=12)
ax.set_ylabel('relative frequency', fontsize=12)
ax.legend(fontsize=12)
plt.show()
```



The black line is the population density of y_T calculated from (43.10).

The histogram and population distribution are close, as expected.

By looking at the figures and experimenting with parameters, you will gain a feel for how the population distribution depends on the model primitives *listed above*, as intermediated by the distribution's parameters.

Ensemble Means

In the preceding figure, we approximated the population distribution of y_T by

1. generating I sample paths (i.e., time series) where I is a large number
2. recording each observation y_T^i
3. histogramming this sample

Just as the histogram approximates the population distribution, the **ensemble** or **cross-sectional average**

$$\bar{y}_T := \frac{1}{I} \sum_{i=1}^{I} y_T^i$$

approximates the expectation $\mathbb{E}[y_T] = G\mu_T$ (as implied by the law of large numbers).

Here's a simulation comparing the ensemble averages and population means at time points $t = 0, \ldots, 50$.

The parameters are the same as for the preceding figures, and the sample size is relatively small ($I = 20$).

```
I = 20
T = 50
ymin = -0.5
ymax = 1.15

ar = LinearStateSpace(A_2, C_2, G_2, mu_0=np.ones(4))

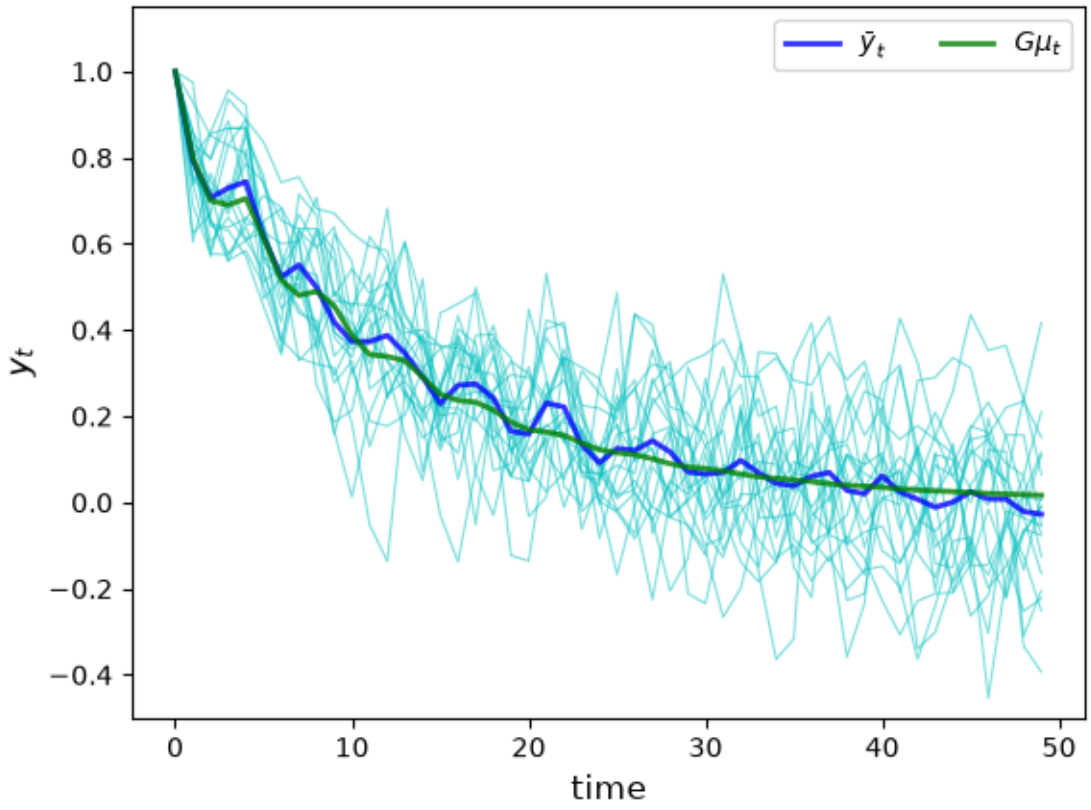
fig, ax = plt.subplots()

ensemble_mean = np.zeros(T)
for i in range(I):
    x, y = ar.simulate(ts_length=T)
    y = y.flatten()
    ax.plot(y, 'c-', lw=0.8, alpha=0.5)
    ensemble_mean = ensemble_mean + y

ensemble_mean = ensemble_mean / I
ax.plot(ensemble_mean, color='b', lw=2, alpha=0.8, label='$\\bar y_t$')
m = ar.moment_sequence()

population_means = []
for t in range(T):
    μ_x, μ_y, Σ_x, Σ_y = next(m)
    population_means.append(float(μ_y.item()))

ax.plot(population_means, color='g', lw=2, alpha=0.8, label=r'$G\mu_t$')
ax.set_ylim(ymin, ymax)
ax.set_xlabel('time', fontsize=12)
ax.set_ylabel('$y_t$', fontsize=12)
ax.legend(ncol=2)
plt.show()
```



The ensemble mean for x_t is

$$\bar{x}_T := \frac{1}{I}\sum_{i=1}^{I} x_T^i \to \mu_T \qquad (I \to \infty)$$

The limit μ_T is a "long-run average".

(By *long-run average* we mean the average for an infinite ($I = \infty$) number of sample x_T's)

Another application of the law of large numbers assures us that

$$\frac{1}{I}\sum_{i=1}^{I}(x_T^i - \bar{x}_T)(x_T^i - \bar{x}_T)' \to \Sigma_T \qquad (I \to \infty)$$

43.3.4 Joint Distributions

In the preceding discussion, we looked at the distributions of x_t and y_t in isolation.

This gives us useful information but doesn't allow us to answer questions like

- what's the probability that $x_t \geq 0$ for all t?
- what's the probability that the process $\{y_t\}$ exceeds some value a before falling below b?
- etc., etc.

Such questions concern the *joint distributions* of these sequences.

To compute the joint distribution of $x_0, x_1, \ldots, x_T$, recall that joint and conditional densities are linked by the rule

$$p(x, y) = p(y \,|\, x)p(x) \qquad \text{(joint} = \text{conditional} \times \text{marginal)}$$

From this rule we get $p(x_0, x_1) = p(x_1 \mid x_0)p(x_0)$.

The Markov property $p(x_t \mid x_{t-1}, \ldots, x_0) = p(x_t \mid x_{t-1})$ and repeated applications of the preceding rule lead us to

$$p(x_0, x_1, \ldots, x_T) = p(x_0) \prod_{t=0}^{T-1} p(x_{t+1} \mid x_t)$$

The marginal $p(x_0)$ is just the primitive $N(\mu_0, \Sigma_0)$.

In view of (43.1), the conditional densities are

$$p(x_{t+1} \mid x_t) = N(Ax_t, CC')$$

Autocovariance Functions

An important object related to the joint distribution is the **autocovariance function**

$$\Sigma_{t+j,t} := \mathbb{E}[(x_{t+j} - \mu_{t+j})(x_t - \mu_t)'] \tag{43.11}$$

Elementary calculations show that

$$\Sigma_{t+j,t} = A^j \Sigma_t \tag{43.12}$$

Notice that $\Sigma_{t+j,t}$ in general depends on both j, the gap between the two dates, and t, the earlier date.

43.4 Stationarity and Ergodicity

Stationarity and ergodicity are two properties that, when they hold, greatly aid analysis of linear state space models.

Let's start with the intuition.

43.4.1 Visualizing Stability

Let's look at some more time series from the same model that we analyzed above.

This picture shows cross-sectional distributions for y at times T, T', T''

```
def cross_plot(A,
               C,
               G,
               steady_state='False',
               T0 = 10,
               T1 = 50,
               T2 = 75,
               T4 = 100):

    ar = LinearStateSpace(A, C, G, mu_0=np.ones(4))

    if steady_state == 'True':
        μ_x, μ_y, Σ_x, Σ_y, Σ_yx = ar.stationary_distributions()
        ar_state = LinearStateSpace(A, C, G, mu_0=μ_x, Sigma_0=Σ_x)

    ymin, ymax = -0.6, 0.6
    fig, ax = plt.subplots()
```

(continues on next page)

(continued from previous page)

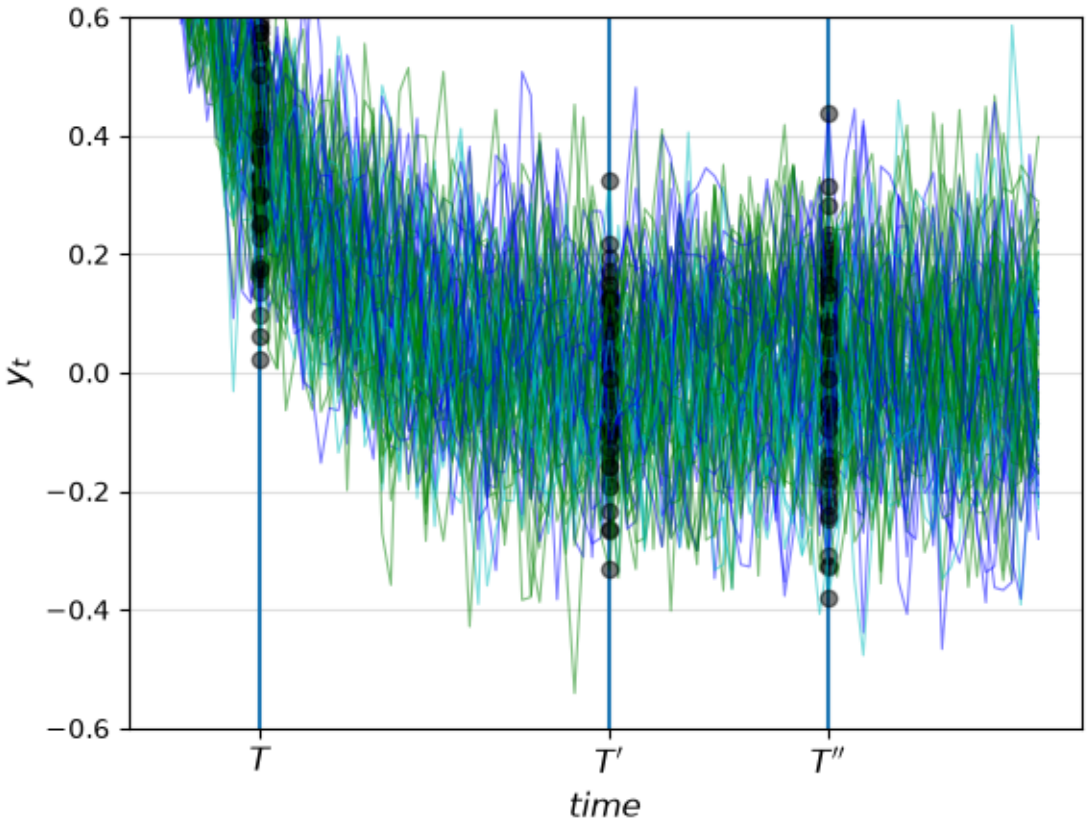
```
    ax.grid(alpha=0.4)
    ax.set_ylim(ymin, ymax)
    ax.set_ylabel('$y_t$', fontsize=12)
    ax.set_xlabel('$time$', fontsize=12)

    ax.vlines((T0, T1, T2), -1.5, 1.5)
    ax.set_xticks((T0, T1, T2))
    ax.set_xticklabels(("$T$", "$T'$", "$T''$"), fontsize=12)
    for i in range(80):
        rcolor = random.choice(('c', 'g', 'b'))

        if steady_state == 'True':
            x, y = ar_state.simulate(ts_length=T4)
        else:
            x, y = ar.simulate(ts_length=T4)

        y = y.flatten()
        ax.plot(y, color=rcolor, lw=0.8, alpha=0.5)
        ax.plot((T0, T1, T2), (y[T0], y[T1], y[T2],), 'ko', alpha=0.5)
    plt.show()
```

```
cross_plot(A_2, C_2, G_2)
```



Note how the time series "settle down" in the sense that the distributions at T' and T'' are relatively similar to each other — but unlike the distribution at T.

Apparently, the distributions of y_t converge to a fixed long-run distribution as $t \to \infty$.

When such a distribution exists it is called a **stationary distribution**.

43.4.2 Stationary Distributions

In our setting, a distribution ψ_∞ is said to be **stationary** for x_t if

$$x_t \sim \psi_\infty \quad \text{and} \quad x_{t+1} = Ax_t + Cw_{t+1} \quad \Longrightarrow \quad x_{t+1} \sim \psi_\infty$$

Since

1. in the present case, all distributions are Gaussian
2. a Gaussian distribution is pinned down by its mean and variance-covariance matrix

we can restate the definition as follows: ψ_∞ is stationary for x_t if

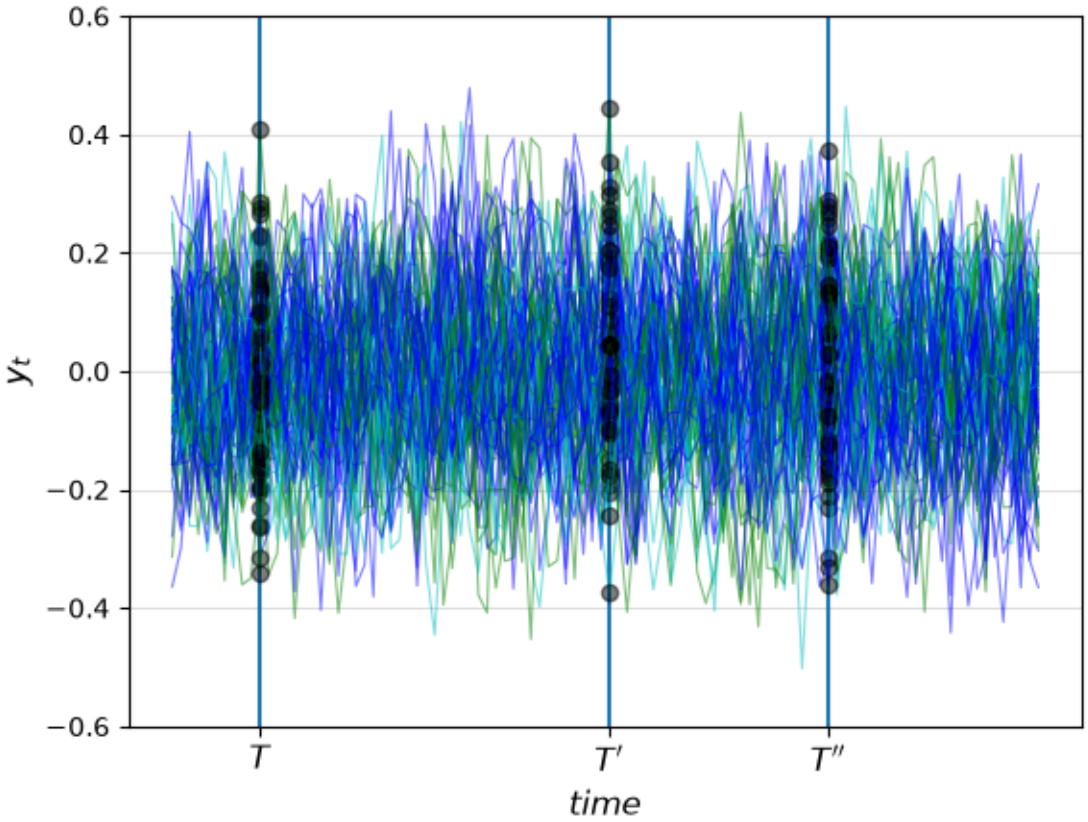
$$\psi_\infty = N(\mu_\infty, \Sigma_\infty)$$

where μ_∞ and Σ_∞ are fixed points of (43.4) and (43.5) respectively.

43.4.3 Covariance Stationary Processes

Let's see what happens to the preceding figure if we start x_0 at the stationary distribution.

```
cross_plot(A_2, C_2, G_2, steady_state='True')
```



Now the differences in the observed distributions at T, T' and T'' come entirely from random fluctuations due to the finite sample size.

By

- our choosing $x_0 \sim N(\mu_\infty, \Sigma_\infty)$
- the definitions of μ_∞ and Σ_∞ as fixed points of (43.4) and (43.5) respectively

we've ensured that

$$\mu_t = \mu_\infty \quad \text{and} \quad \Sigma_t = \Sigma_\infty \quad \text{for all } t$$

Moreover, in view of (43.12), the autocovariance function takes the form $\Sigma_{t+j,t} = A^j \Sigma_\infty$, which depends on j but not on t.

This motivates the following definition.

A process $\{x_t\}$ is said to be **covariance stationary** if

- both μ_t and Σ_t are constant in t
- $\Sigma_{t+j,t}$ depends on the time gap j but not on time t

In our setting, $\{x_t\}$ will be covariance stationary if μ_0, Σ_0, A, C assume values that imply that none of $\mu_t, \Sigma_t, \Sigma_{t+j,t}$ depends on t.

43.4.4 Conditions for Stationarity

The Globally Stable Case

The difference equation $\mu_{t+1} = A\mu_t$ is known to have *unique* fixed point $\mu_\infty = 0$ if all eigenvalues of A have moduli strictly less than unity.

That is, if `(np.absolute(np.linalg.eigvals(A)) < 1).all() == True`.

The difference equation (43.5) also has a unique fixed point in this case, and, moreover

$$\mu_t \to \mu_\infty = 0 \quad \text{and} \quad \Sigma_t \to \Sigma_\infty \quad \text{as} \quad t \to \infty$$

regardless of the initial conditions μ_0 and Σ_0.

This is the *globally stable case* — see these notes for more a theoretical treatment.

However, global stability is more than we need for stationary solutions, and often more than we want.

To illustrate, consider *our second order difference equation example*.

Here the state is $x_t = \begin{bmatrix} 1 & y_t & y_{t-1} \end{bmatrix}'$.

Because of the constant first component in the state vector, we will never have $\mu_t \to 0$.

How can we find stationary solutions that respect a constant state component?

Processes with a Constant State Component

To investigate such a process, suppose that A and C take the form

$$A = \begin{bmatrix} A_1 & a \\ 0 & 1 \end{bmatrix} \qquad C = \begin{bmatrix} C_1 \\ 0 \end{bmatrix}$$

where

- A_1 is an $(n-1) \times (n-1)$ matrix
- a is an $(n-1) \times 1$ column vector

Let $x_t = \begin{bmatrix} x_{1t}' & 1 \end{bmatrix}'$ where x_{1t} is $(n-1) \times 1$.

It follows that

$$x_{1,t+1} = A_1 x_{1t} + a + C_1 w_{t+1}$$

Let $\mu_{1t} = \mathbb{E}[x_{1t}]$ and take expectations on both sides of this expression to get

$$\mu_{1,t+1} = A_1 \mu_{1,t} + a \tag{43.13}$$

Assume now that the moduli of the eigenvalues of A_1 are all strictly less than one.

Then (43.13) has a unique stationary solution, namely,

$$\mu_{1\infty} = (I - A_1)^{-1} a$$

The stationary value of μ_t itself is then $\mu_\infty := \begin{bmatrix} \mu_{1\infty}' & 1 \end{bmatrix}'$.

The stationary values of Σ_t and $\Sigma_{t+j,t}$ satisfy

$$\begin{aligned} \Sigma_\infty &= A \Sigma_\infty A' + CC' \\ \Sigma_{t+j,t} &= A^j \Sigma_\infty \end{aligned}$$

Notice that here $\Sigma_{t+j,t}$ depends on the time gap j but not on calendar time t.

In conclusion, if

- $x_0 \sim N(\mu_\infty, \Sigma_\infty)$ and
- the moduli of the eigenvalues of A_1 are all strictly less than unity

then the $\{x_t\}$ process is covariance stationary, with constant state component.

Note

If the eigenvalues of A_1 are less than unity in modulus, then (a) starting from any initial value, the mean and variance-covariance matrix both converge to their stationary values; and (b) iterations on (43.5) converge to the fixed point of the *discrete Lyapunov equation* in the first line of (43.14).

43.4.5 Ergodicity

Let's suppose that we're working with a covariance stationary process.

In this case, we know that the ensemble mean will converge to μ_∞ as the sample size I approaches infinity.

Averages over Time

Ensemble averages across simulations are interesting theoretically, but in real life, we usually observe only a *single* realization $\{x_t, y_t\}_{t=0}^T$.

So now let's take a single realization and form the time-series averages

$$\bar{x} := \frac{1}{T}\sum_{t=1}^{T} x_t \quad \text{and} \quad \bar{y} := \frac{1}{T}\sum_{t=1}^{T} y_t$$

Do these time series averages converge to something interpretable in terms of our basic state-space representation?

The answer depends on something called *ergodicity*.

Ergodicity is the property that time series and ensemble averages coincide.

More formally, ergodicity implies that time series sample averages converge to their expectation under the stationary distribution.

In particular,

- $\frac{1}{T}\sum_{t=1}^{T} x_t \to \mu_\infty$
- $\frac{1}{T}\sum_{t=1}^{T} (x_t - \bar{x}_T)(x_t - \bar{x}_T)' \to \Sigma_\infty$
- $\frac{1}{T}\sum_{t=1}^{T} (x_{t+j} - \bar{x}_T)(x_t - \bar{x}_T)' \to A^j\Sigma_\infty$

In our linear Gaussian setting, any covariance stationary process is also ergodic.

43.5 Noisy Observations

In some settings, the observation equation $y_t = Gx_t$ is modified to include an error term.

Often this error term represents the idea that the true state can only be observed imperfectly.

To include an error term in the observation we introduce

- An IID sequence of $\ell \times 1$ random vectors $v_t \sim N(0, I)$.
- A $k \times \ell$ matrix H.

and extend the linear state-space system to

$$\begin{aligned} x_{t+1} &= Ax_t + Cw_{t+1} \\ y_t &= Gx_t + Hv_t \\ x_0 &\sim N(\mu_0, \Sigma_0) \end{aligned}$$

The sequence $\{v_t\}$ is assumed to be independent of $\{w_t\}$.

The process $\{x_t\}$ is not modified by noise in the observation equation and its moments, distributions and stability properties remain the same.

The unconditional moments of y_t from (43.6) and (43.7) now become

$$\mathbb{E}[y_t] = \mathbb{E}[Gx_t + Hv_t] = G\mu_t \tag{43.14}$$

The variance-covariance matrix of y_t is easily shown to be

$$\text{Var}[y_t] = \text{Var}[Gx_t + Hv_t] = G\Sigma_t G' + HH' \tag{43.15}$$

The distribution of y_t is therefore

$$y_t \sim N(G\mu_t, G\Sigma_t G' + HH')$$

43.6 Prediction

The theory of prediction for linear state space systems is elegant and simple.

43.6.1 Forecasting Formulas – Conditional Means

The natural way to predict variables is to use conditional distributions.

For example, the optimal forecast of x_{t+1} given information known at time t is

$$\mathbb{E}_t[x_{t+1}] := \mathbb{E}[x_{t+1} \mid x_t, x_{t-1}, \ldots, x_0] = Ax_t$$

The right-hand side follows from $x_{t+1} = Ax_t + Cw_{t+1}$ and the fact that w_{t+1} is zero mean and independent of $x_t, x_{t-1}, \ldots, x_0$.

That $\mathbb{E}_t[x_{t+1}] = \mathbb{E}[x_{t+1} \mid x_t]$ is an implication of $\{x_t\}$ having the **Markov property**.

The one-step-ahead forecast error is

$$x_{t+1} - \mathbb{E}_t[x_{t+1}] = Cw_{t+1}$$

The covariance matrix of the forecast error is

$$\mathbb{E}[(x_{t+1} - \mathbb{E}_t[x_{t+1}])(x_{t+1} - \mathbb{E}_t[x_{t+1}])'] = CC'$$

More generally, we'd like to compute the j-step ahead forecasts $\mathbb{E}_t[x_{t+j}]$ and $\mathbb{E}_t[y_{t+j}]$.

With a bit of algebra, we obtain

$$x_{t+j} = A^j x_t + A^{j-1} Cw_{t+1} + A^{j-2} Cw_{t+2} + \cdots + A^0 Cw_{t+j}$$

In view of the IID property, current and past state values provide no information about future values of the shock.

Hence $\mathbb{E}_t[w_{t+k}] = \mathbb{E}[w_{t+k}] = 0$.

It now follows from linearity of expectations that the j-step ahead forecast of x is

$$\mathbb{E}_t[x_{t+j}] = A^j x_t$$

The j-step ahead forecast of y is therefore

$$\mathbb{E}_t[y_{t+j}] = \mathbb{E}_t[Gx_{t+j} + Hv_{t+j}] = GA^j x_t$$

43.6.2 Covariance of Prediction Errors

It is useful to obtain the covariance matrix of the vector of j-step-ahead prediction errors

$$x_{t+j} - \mathbb{E}_t[x_{t+j}] = \sum_{s=0}^{j-1} A^s C w_{t-s+j} \tag{43.16}$$

Evidently,

$$V_j := \mathbb{E}_t[(x_{t+j} - \mathbb{E}_t[x_{t+j}])(x_{t+j} - \mathbb{E}_t[x_{t+j}])'] = \sum_{k=0}^{j-1} A^k C C' A^{k'} \tag{43.17}$$

V_j defined in (43.17) can be calculated recursively via $V_1 = CC'$ and

$$V_j = CC' + AV_{j-1}A', \quad j \geq 2 \tag{43.18}$$

V_j is the **conditional covariance matrix** of the errors in forecasting x_{t+j}, conditioned on time t information x_t.

Under particular conditions, V_j converges to

$$V_\infty = CC' + AV_\infty A' \tag{43.19}$$

Equation (43.19) is an example of a **discrete Lyapunov** equation in the covariance matrix V_∞.

A sufficient condition for V_j to converge is that the eigenvalues of A be strictly less than one in modulus.

Weaker sufficient conditions for convergence associate eigenvalues equaling or exceeding one in modulus with elements of C that equal 0.

43.7 Code

Our preceding simulations and calculations are based on code in the file lss.py from the QuantEcon.py package.

The code implements a class for handling linear state space models (simulations, calculating moments, etc.).

One Python construct you might not be familiar with is the use of a generator function in the method `moment_sequence()`.

Go back and read the relevant documentation if you've forgotten how generator functions work.

Examples of usage are given in the solutions to the exercises.

43.8 Exercises

Exercise 43.8.1

In several contexts, we want to compute forecasts of geometric sums of future random variables governed by the linear state-space system (43.1).

We want the following objects

- Forecast of a geometric sum of future x's, or $\mathbb{E}_t\left[\sum_{j=0}^{\infty} \beta^j x_{t+j}\right]$.
- Forecast of a geometric sum of future y's, or $\mathbb{E}_t\left[\sum_{j=0}^{\infty} \beta^j y_{t+j}\right]$.

These objects are important components of some famous and interesting dynamic models.

For example,

- if $\{y_t\}$ is a stream of dividends, then $\mathbb{E}\left[\sum_{j=0}^{\infty} \beta^j y_{t+j} | x_t\right]$ is a model of a stock price
- if $\{y_t\}$ is the money supply, then $\mathbb{E}\left[\sum_{j=0}^{\infty} \beta^j y_{t+j} | x_t\right]$ is a model of the price level

Show that:

$$\mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j x_{t+j}\right] = [I - \beta A]^{-1} x_t$$

and

$$\mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j y_{t+j}\right] = G[I - \beta A]^{-1} x_t$$

what must the modulus for every eigenvalue of A be less than?

Solution

Suppose that every eigenvalue of A has modulus strictly less than $\frac{1}{\beta}$.

It *then follows* that $I + \beta A + \beta^2 A^2 + \cdots = [I - \beta A]^{-1}$.

This leads to our formulas:

- Forecast of a geometric sum of future x's

$$\mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j x_{t+j}\right] = [I + \beta A + \beta^2 A^2 + \cdots] x_t = [I - \beta A]^{-1} x_t$$

- Forecast of a geometric sum of future y's

$$\mathbb{E}_t \left[\sum_{j=0}^{\infty} \beta^j y_{t+j}\right] = G[I + \beta A + \beta^2 A^2 + \cdots] x_t = G[I - \beta A]^{-1} x_t$$

CHAPTER

FORTYFOUR

SAMUELSON MULTIPLIER-ACCELERATOR

Contents

- *Samuelson Multiplier-Accelerator*
  - *Overview*
  - *Details*
  - *Implementation*
  - *Stochastic shocks*
  - *Government spending*
  - *Wrapping everything into a class*
  - *Using the LinearStateSpace class*
  - *Pure multiplier model*
  - *Summary*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

44.1 Overview

This lecture creates non-stochastic and stochastic versions of Paul Samuelson's celebrated multiplier accelerator model [Samuelson, 1939].

In doing so, we extend the example of the Solow model class in our second OOP lecture.

Our objectives are to

- provide a more detailed example of OOP and classes
- review a famous model
- review linear difference equations, both deterministic and stochastic

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

We'll also use the following for various tasks described below:

```
from quantecon import LinearStateSpace
import cmath
import math
import sympy
from sympy import Symbol, init_printing
from cmath import sqrt
```

44.1.1 Samuelson's model

Samuelson used a *second-order linear difference equation* to represent a model of national output based on three components:

- a *national output identity* asserting that national output or national income is the sum of consumption plus investment plus government purchases.
- a Keynesian *consumption function* asserting that consumption at time t is equal to a constant times national output at time $t-1$.
- an investment *accelerator* asserting that investment at time t equals a constant called the *accelerator coefficient* times the difference in output between period $t-1$ and $t-2$.

Consumption plus investment plus government purchases constitute *aggregate demand,* which automatically calls forth an equal amount of *aggregate supply*.

(To read about linear difference equations see here or chapter IX of [Sargent, 1987].)

Samuelson used the model to analyze how particular values of the marginal propensity to consume and the accelerator coefficient might give rise to transient **business cycles** in national output.

Possible dynamic properties include

- smooth convergence to a constant level of output
- damped business cycles that eventually converge to a constant level of output
- persistent business cycles that neither dampen nor explode

Later we present an extension that adds a random shock to the right side of the national income identity representing random fluctuations in aggregate demand.

This modification makes national output become governed by a second-order **stochastic linear difference equation** that, with appropriate parameter values, gives rise to recurrent irregular business cycles.

(To read about stochastic linear difference equations see chapter XI of [Sargent, 1987].)

44.2 Details

Let's assume that

- $\{G_t\}$ is a sequence of levels of government expenditures – we'll start by setting $G_t = G$ for all t.
- $\{C_t\}$ is a sequence of levels of aggregate consumption expenditures, a key endogenous variable in the model.
- $\{I_t\}$ is a sequence of rates of investment, another key endogenous variable.
- $\{Y_t\}$ is a sequence of levels of national income, yet another endogenous variable.
- α is the marginal propensity to consume in the Keynesian consumption function $C_t = \alpha Y_{t-1} + \gamma$.
- β is the "accelerator coefficient" in the "investment accelerator" $I_t = \beta(Y_{t-1} - Y_{t-2})$.
- $\{\epsilon_t\}$ is an IID sequence standard normal random variables.
- $\sigma \geq 0$ is a "volatility" parameter — setting $\sigma = 0$ recovers the non-stochastic case that we'll start with.

The model combines the consumption function

$$C_t = \alpha Y_{t-1} + \gamma \tag{44.1}$$

with the investment accelerator

$$I_t = \beta(Y_{t-1} - Y_{t-2}) \tag{44.2}$$

and the national income identity

$$Y_t = C_t + I_t + G_t \tag{44.3}$$

- The parameter α is peoples' **marginal propensity to consume** out of income - equation (44.1) asserts that people consume a fraction of $\alpha \in (0, 1)$ of each additional dollar of income.
- The parameter $\beta > 0$ is the investment accelerator coefficient - equation (44.2) asserts that people invest in physical capital when income is increasing and disinvest when it is decreasing.

Equations (44.1), (44.2), and (44.3) imply the following second-order linear difference equation for national income:

$$Y_t = (\alpha + \beta)Y_{t-1} - \beta Y_{t-2} + (\gamma + G_t)$$

or

$$Y_t = \rho_1 Y_{t-1} + \rho_2 Y_{t-2} + (\gamma + G_t) \tag{44.4}$$

where $\rho_1 = (\alpha + \beta)$ and $\rho_2 = -\beta$.

To complete the model, we require two *initial conditions*.

If the model is to generate time series for $t = 0, \ldots, T$, we require initial values

$$Y_{-1} = \bar{Y}_{-1}, \quad Y_{-2} = \bar{Y}_{-2}$$

We'll ordinarily set the parameters (α, β) so that starting from an arbitrary pair of initial conditions $(\bar{Y}_{-1}, \bar{Y}_{-2})$, national income Y_t converges to a constant value as t becomes large.

We are interested in studying

- the transient fluctuations in Y_t as it converges to its **steady state** level
- the *rate* at which it converges to a steady state level

The deterministic version of the model described so far — meaning that no random shocks hit aggregate demand — has only transient fluctuations.

We can convert the model to one that has persistent irregular fluctuations by adding a random shock to aggregate demand.

44.2.1 Stochastic version of the model

We create a *random* or *stochastic* version of the model by adding a random process of *shocks* or *disturbances* $\{\sigma\epsilon_t\}$ to the right side of equation (44.4), leading to the *second-order scalar linear stochastic difference equation*:

$$Y_t = (\alpha + \beta)Y_{t-1} - \beta Y_{t-2} + (\gamma + G_t) + \sigma\epsilon_t \tag{44.5}$$

44.2.2 Mathematical analysis of the model

To get started, let's set $G_t \equiv 0$, $\sigma = 0$, and $\gamma = 0$.

Then we can write equation (44.5) as

$$Y_t = \rho_1 Y_{t-1} + \rho_2 Y_{t-2}$$

or

$$Y_{t+2} - \rho_1 Y_{t+1} - \rho_2 Y_t = 0 \tag{44.6}$$

To discover the properties of the solution of (44.6), it is useful first to form the **characteristic polynomial** for (44.6):

$$z^2 - \rho_1 z - \rho_2 \tag{44.7}$$

where z is possibly a complex number.

We want to find the two **zeros** (a.k.a. **roots**) – namely λ_1, λ_2 – of the characteristic polynomial.

These are two special values of z, say $z = \lambda_1$ and $z = \lambda_2$, such that if we set z equal to one of these values in expression (44.7), the characteristic polynomial (44.7) equals zero:

$$z^2 - \rho_1 z - \rho_2 = (z - \lambda_1)(z - \lambda_2) = 0 \tag{44.8}$$

Equation (44.8) is said to *factor* the characteristic polynomial.

When the roots are complex, they will occur as a complex conjugate pair.

When the roots are complex, it is convenient to represent them in the polar form

$$\lambda_1 = re^{i\omega}, \ \lambda_2 = re^{-i\omega}$$

where r is the *amplitude* of the complex number and ω is its *angle* or *phase*.

These can also be represented as

$$\lambda_1 = r(cos(\omega) + i \sin(\omega))$$

$$\lambda_2 = r(cos(\omega) - i \sin(\omega))$$

(To read about the polar form, see here)

Given *initial conditions* Y_{-1}, Y_{-2}, we want to generate a *solution* of the difference equation (44.6).

It can be represented as

$$Y_t = \lambda_1^t c_1 + \lambda_2^t c_2$$

where c_1 and c_2 are constants that depend on the two initial conditions and on ρ_1, ρ_2.

When the roots are complex, it is useful to pursue the following calculations.

Notice that

$$\begin{aligned} Y_t &= c_1(re^{i\omega})^t + c_2(re^{-i\omega})^t \\ &= c_1 r^t e^{i\omega t} + c_2 r^t e^{-i\omega t} \\ &= c_1 r^t [\cos(\omega t) + i \sin(\omega t)] + c_2 r^t [\cos(\omega t) - i \sin(\omega t)] \\ &= (c_1 + c_2) r^t \cos(\omega t) + i(c_1 - c_2) r^t \sin(\omega t) \end{aligned}$$

The only way that Y_t can be a real number for each t is if $c_1 + c_2$ is a real number and $c_1 - c_2$ is an imaginary number.

This happens only when c_1 and c_2 are complex conjugates, in which case they can be written in the polar forms

$$c_1 = ve^{i\theta}, \quad c_2 = ve^{-i\theta}$$

So we can write

$$\begin{aligned} Y_t = &\quad ve^{i\theta} r^t e^{i\omega t} + ve^{-i\theta} r^t e^{-i\omega t} \\ = &\quad vr^t [e^{i(\omega t + \theta)} + e^{-i(\omega t + \theta)}] \\ = &\quad 2vr^t \cos(\omega t + \theta) \end{aligned}$$

where v and θ are constants that must be chosen to satisfy initial conditions for Y_{-1}, Y_{-2}.

This formula shows that when the roots are complex, Y_t displays oscillations with *period* $\breve{p} = \frac{2\pi}{\omega}$ and *damping factor* r.

We say that $\breve{p}$ is the *period* because in that amount of time the cosine wave $\cos(\omega t + \theta)$ goes through exactly one complete cycles.

(Draw a cosine function to convince yourself of this please)

Remark: Following [Samuelson, 1939], we want to choose the parameters α, β of the model so that the absolute values (of the possibly complex) roots λ_1, λ_2 of the characteristic polynomial are both strictly less than one:

$$|\lambda_j| < 1 \qquad \text{for } j = 1, 2$$

Remark: When both roots λ_1, λ_2 of the characteristic polynomial have absolute values strictly less than one, the absolute value of the larger one governs the rate of convergence to the steady state of the non stochastic version of the model.

44.2.3 Things this lecture does

We write a function to generate simulations of a $\{Y_t\}$ sequence as a function of time.

The function requires that we put in initial conditions for Y_{-1}, Y_{-2}.

The function checks that α, β are set so that λ_1, λ_2 are less than unity in absolute value (also called "modulus").

The function also tells us whether the roots are complex, and, if they are complex, returns both their real and complex parts.

If the roots are both real, the function returns their values.

We use our function written to simulate paths that are stochastic (when $\sigma > 0$).

We have written the function in a way that allows us to input $\{G_t\}$ paths of a few simple forms, e.g.,

- one time jumps in G at some time
- a permanent jump in G that occurs at some time

We proceed to use the Samuelson multiplier-accelerator model as a laboratory to make a simple OOP example.

The "state" that determines next period's Y_{t+1} is now not just the current value Y_t but also the once lagged value Y_{t-1}.

This involves a little more bookkeeping than is required in the Solow model class definition.

We use the Samuelson multiplier-accelerator model as a vehicle for teaching how we can gradually add more features to the class.

We want to have a method in the class that automatically generates a simulation, either non-stochastic ($\sigma = 0$) or stochastic ($\sigma > 0$).

We also show how to map the Samuelson model into a simple instance of the `LinearStateSpace` class described here.

We can use a `LinearStateSpace` instance to do various things that we did above with our homemade function and class.

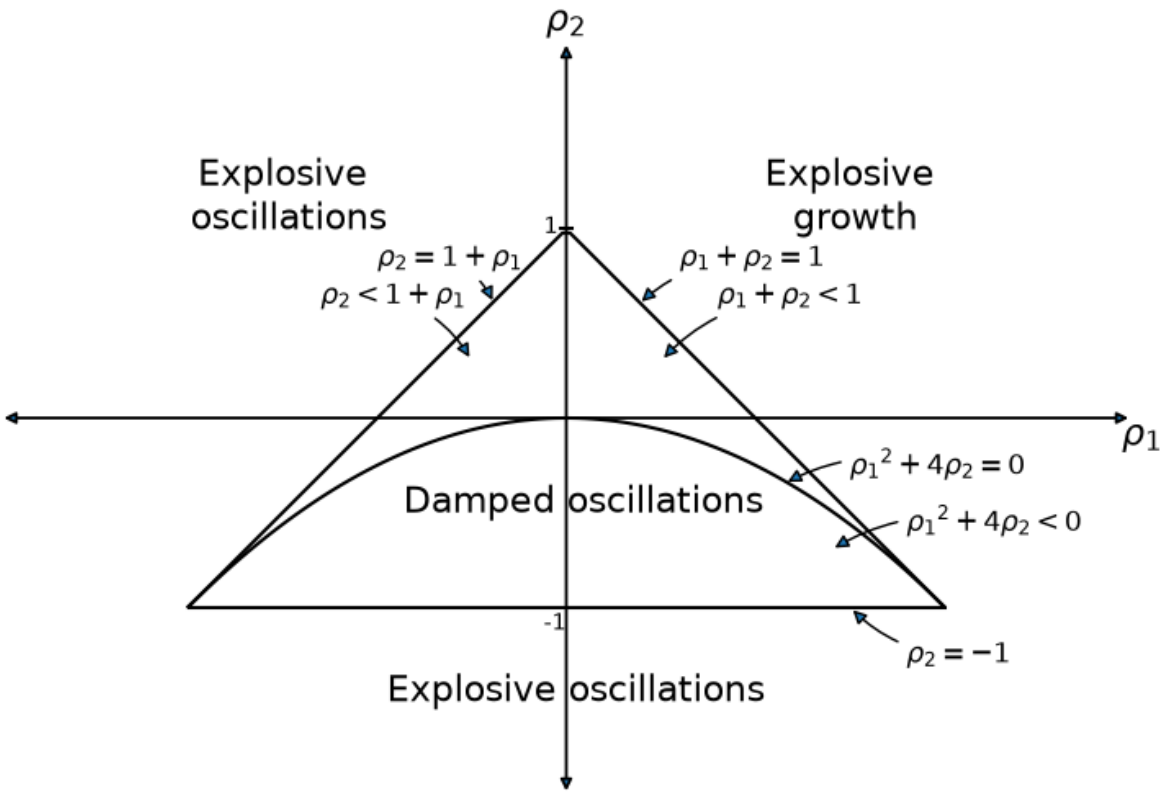
Among other things, we show by example that the eigenvalues of the matrix A that we use to form the instance of the `LinearStateSpace` class for the Samuelson model equal the roots of the characteristic polynomial (44.7) for the Samuelson multiplier accelerator model.

Here is the formula for the matrix A in the linear state space system in the case that government expenditures are a constant G:

$$A = \begin{bmatrix} 1 & 0 & 0 \\ \gamma + G & \rho_1 & \rho_2 \\ 0 & 1 & 0 \end{bmatrix}$$

44.3 Implementation

We'll start by drawing an informative graph from page 189 of [Sargent, 1987]



The graph portrays regions in which the (λ_1, λ_2) root pairs implied by the $(\rho_1 = (\alpha + \beta), \rho_2 = -\beta)$ difference equation parameter pairs in the Samuelson model are such that:

- (λ_1, λ_2) are complex with modulus less than 1 - in this case, the $\{Y_t\}$ sequence displays damped oscillations.
- (λ_1, λ_2) are both real, but one is strictly greater than 1 - this leads to explosive growth.
- (λ_1, λ_2) are both real, but one is strictly less than -1 - this leads to explosive oscillations.
- (λ_1, λ_2) are both real and both are less than 1 in absolute value - in this case, there is smooth convergence to the steady state without damped cycles.

Later we'll present the graph with a red mark showing the particular point implied by the setting of (α, β).

44.3.1 Function to describe implications of characteristic polynomial

```
def categorize_solution(ρ1, ρ2):
    """
    This function takes values of ρ1 and ρ2 and uses them
    to classify the type of solution.
    """
    discriminant = ρ1**2 + 4 * ρ2
    if ρ2 > 1 + ρ1 or ρ2 < -1:
        return "Explosive oscillations"
    elif ρ1 + ρ2 > 1:
        return "Explosive growth"
    elif discriminant < 0:
        return "Damped oscillations"
    else:
        return "Steady state convergence"

def analyze_roots(α, β, verbose=True):
    """
    Unified function to calculate roots and analyze their properties.
    """
    ρ1 = α + β
    ρ2 = -β

    # Compute characteristic polynomial roots
    roots = np.roots([1, -ρ1, -ρ2])

    # Classify solution type
    solution_type = categorize_solution(ρ1, ρ2)

    # Determine root properties
    is_complex = all(isinstance(root, complex) for root in roots)
    is_stable = all(abs(root) < 1 for root in roots)

    if verbose:
        print(f"ρ1 = {ρ1:.2f}, ρ2 = {ρ2:.2f}")
        print(f"Roots: {[f'{root:.2f}' for root in roots]}")
        print(f"Root type: {'Complex' if is_complex else 'Real'}")
        print(f"Stability: {'Stable' if is_stable else 'Unstable'}")
        print(f"Solution type: {solution_type}")

    return {
        'roots': roots,
        'rho1': ρ1,
        'rho2': ρ2,
        'is_complex': is_complex,
        'is_stable': is_stable,
```

(continues on next page)

(continued from previous page)

```
        'solution_type': solution_type
    }
```

We also write a unified simulation function that can handle both non-stochastic and stochastic versions of the model.

It allows for government spending paths of a few simple forms which we specify via a dictionary `g_params`

```
def simulate_samuelson(
    y_0, y_1, α, β, γ=10, n=100, σ=0, g_params=None, seed=0
):
    """
    Unified simulation function for Samuelson model.

    Parameters:
    g_params: dict with keys 'g', 'g_t', 'duration' for government spending
    seed: random seed for reproducible results
    """
    analysis = analyze_roots(α, β, verbose=False)
    ρ1, ρ2 = analysis['rho1'], analysis['rho2']

    # Initialize time series
    y_t = [y_0, y_1]

    # Generate shocks if stochastic
    if σ > 0:
        rng = np.random.default_rng(seed)
        ϵ = rng.normal(0, 1, n)

    # Simulate forward
    for t in range(2, n):

        # Determine government spending
        g = 0
        if g_params:
            g_val, g_t_val = g_params.get('g', 0), g_params.get('g_t', 0)
            duration = g_params.get('duration', None)
            if duration == 'permanent' and t >= g_t_val:
                g = g_val
            elif duration == 'one-off' and t == g_t_val:
                g = g_val
            elif duration is None:
                g = g_val

        # Calculate next value
        y_next = ρ1 * y_t[t-1] + ρ2 * y_t[t-2] + γ + g
        if σ > 0:
            y_next += σ * ϵ[t]

        y_t.append(y_next)

    return y_t, analysis
```

We will use this function to run simulations of the model.

But before doing that, let's test the analysis function

```
analysis = analyze_roots(α=1.3, β=0.4)
```

```
ρ1 = 1.70, ρ2 = -0.40
Roots: ['1.42', '0.28']
Root type: Real
Stability: Unstable
Solution type: Explosive growth
```

44.3.2 Function for plotting paths

A useful function for our work below is

```
def plot_y(function=None):
    """Function plots path of Y_t"""

    plt.subplots(figsize=(10, 6))
    plt.plot(function)
    plt.xlabel("$t$")
    plt.ylabel("$Y_t$")
    plt.show()
```

44.3.3 Manual or "by hand" root calculations

The following function calculates roots of the characteristic polynomial using high school algebra.

(We'll calculate the roots in other ways later using `analyze_roots`.)

The function also plots a Y_t starting from initial conditions that we set

```
def y_nonstochastic(y_0=100, y_1=80, α=0.92, β=0.5, γ=10, n=80):
    """
    This function calculates the roots of the characteristic polynomial
    by hand and returns a path of y_t starting from initial conditions
    """
    roots = []

    ρ1 = α + β
    ρ2 = -β

    print(f"ρ_1 is {ρ1:.2f}")
    print(f"ρ_2 is {ρ2:.2f}")

    discriminant = ρ1**2 + 4 * ρ2

    if discriminant == 0:
        roots.append(ρ1 / 2)
        print("Single real root: ")
        print("".join(f"{r:.2f}" for r in roots))
    elif discriminant > 0:
        roots.append((ρ1 + sqrt(discriminant).real) / 2)
        roots.append((ρ1 - sqrt(discriminant).real) / 2)
        print("Two real roots: ")
        print(" ".join(f"{r:.2f}" for r in roots))
    else:
        roots.append((ρ1 + sqrt(discriminant)) / 2)
        roots.append((ρ1 - sqrt(discriminant)) / 2)
        print("Two complex roots: ")
```

(continues on next page)

(continued from previous page)

```
        print(" ".join(f"{r.real:.2f}{r.imag:+.2f}j" for r in roots))

    if all(abs(root) < 1 for root in roots):
        print("Absolute values of roots are less than one")
    else:
        print("Absolute values of roots are not less than one")

    def transition(x, t):
        return ρ1 * x[t - 1] + ρ2 * x[t - 2] + y

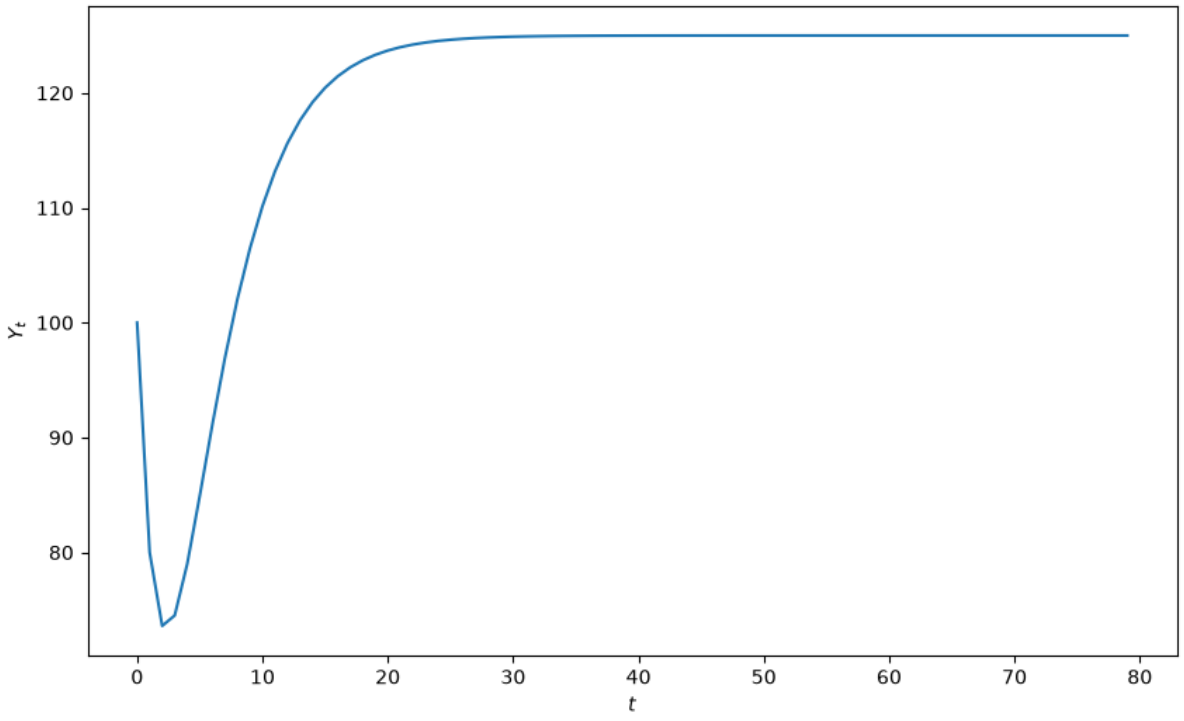
    y_t = [y_0, y_1]

    for t in range(2, n):
        y_t.append(transition(y_t, t))

    return y_t


plot_y(y_nonstochastic())
```

```
ρ_1 is 1.42
ρ_2 is -0.50
Two real roots:
0.77 0.65
Absolute values of roots are less than one
```



44.3.4 Reverse-engineering parameters to generate damped cycles

The next cell writes code that takes as inputs the modulus r and phase ϕ of a conjugate pair of complex numbers in polar form

$$\lambda_1 = r\exp(i\phi), \quad \lambda_2 = r\exp(-i\phi)$$

- The code assumes that these two complex numbers are the roots of the characteristic polynomial
- It then reverse-engineers (α, β) and (ρ_1, ρ_2), pairs that would generate those roots

```
def f(r, ϕ):
    """
    Takes modulus r and angle ϕ of complex number r exp(j ϕ)
    and creates ρ1 and ρ2 of characteristic polynomial for which
    r exp(j ϕ) and r exp(- j ϕ) are complex roots.

    Returns the multiplier coefficient $\alpha$
    and the accelerator coefficient $\beta$
    that verifies those roots.
    """
    # Create complex conjugate pair from polar coordinates
    g1 = cmath.rect(r, ϕ)
    g2 = cmath.rect(r, -ϕ)

    # Calculate corresponding ρ1, ρ2 parameters
    ρ1 = g1 + g2
    ρ2 = -g1 * g2

    # Derive α and β coefficients from ρ parameters
    β = -ρ2
    α = ρ1 - β
    return ρ1, ρ2, α, β
```

Now let's use the function in an example.

Here are the example parameters

```
r = 0.95

# Cycle period in time units
period = 10
ϕ = 2 * math.pi / period

# Apply the reverse-engineering function
ρ1, ρ2, α, β = f(r, ϕ)

print(f"α, β = {α:.2f}, {β:.2f}")
print(f"ρ1, ρ2 = {ρ1:.2f}, {ρ2:.2f}")
```

```
α, β = 0.63+0.00j, 0.90-0.00j
ρ1, ρ2 = 1.54+0.00j, -0.90+0.00j
```

The real parts of the roots are

```
print(f"ρ1 = {ρ1.real:.2f}, ρ2 = {ρ2.real:.2f}")
```

```
ρ1 = 1.54, ρ2 = -0.90
```

44.3.5 Root finding using numpy

Here we'll use numpy to compute the roots of the characteristic polynomial

```
r1, r2 = np.roots([1, -ρ1, -ρ2])

p1 = cmath.polar(r1)
p2 = cmath.polar(r2)

print(f"r, ϕ = {r:.2f}, {ϕ:.2f}")
print(f"p1, p2 = ({p1[0]:.2f}, {p1[1]:.2f}), ({p2[0]:.2f}, {p2[1]:.2f})")

print(f"α, β = {α:.2f}, {β:.2f}")
print(f"ρ1, ρ2 = {ρ1:.2f}, {ρ2:.2f}")
```

```
r, ϕ = 0.95, 0.63
p1, p2 = (0.95, 0.63), (0.95, -0.63)
α, β = 0.63+0.00j, 0.90-0.00j
ρ1, ρ2 = 1.54+0.00j, -0.90+0.00j
```

```
def y_nonstochastic(y_0=100, y_1=80, α=0.9, β=0.8, γ=10, n=80):
    """
    This function enlists numpy to calculate the roots of the characteristic
    polynomial.
    """

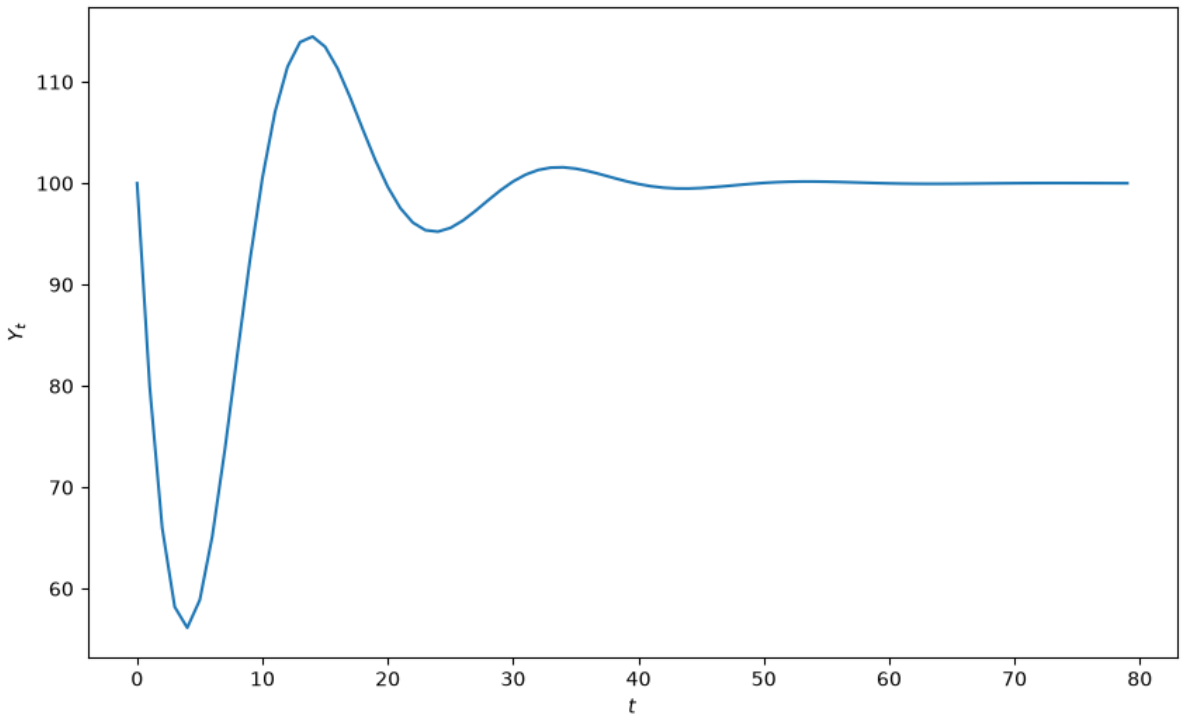
    y_series, analysis = simulate_samuelson(y_0, y_1, α, β, γ, n, 0, None, 42)

    print(f"Solution type: {analysis['solution_type']}")
    print(f"Roots are {analysis['roots']}")
    print(f"Root type: {'Complex' if analysis['is_complex'] else 'Real'}")
    print(f"Stability: {'Stable' if analysis['is_stable'] else 'Unstable'}")

    return y_series


plot_y(y_nonstochastic())
```

```
Solution type: Damped oscillations
Roots are [0.85+0.27838822j 0.85-0.27838822j]
Root type: Complex
Stability: Stable
```



44.3.6 Reverse-engineered complex roots: Example

The next cell studies the implications of reverse-engineered complex roots.

We'll generate an *undamped* cycle of period 10

```
r = 1  # Generates undamped, nonexplosive cycles

period = 10  # Length of cycle in units of time
ϕ = 2 * math.pi / period

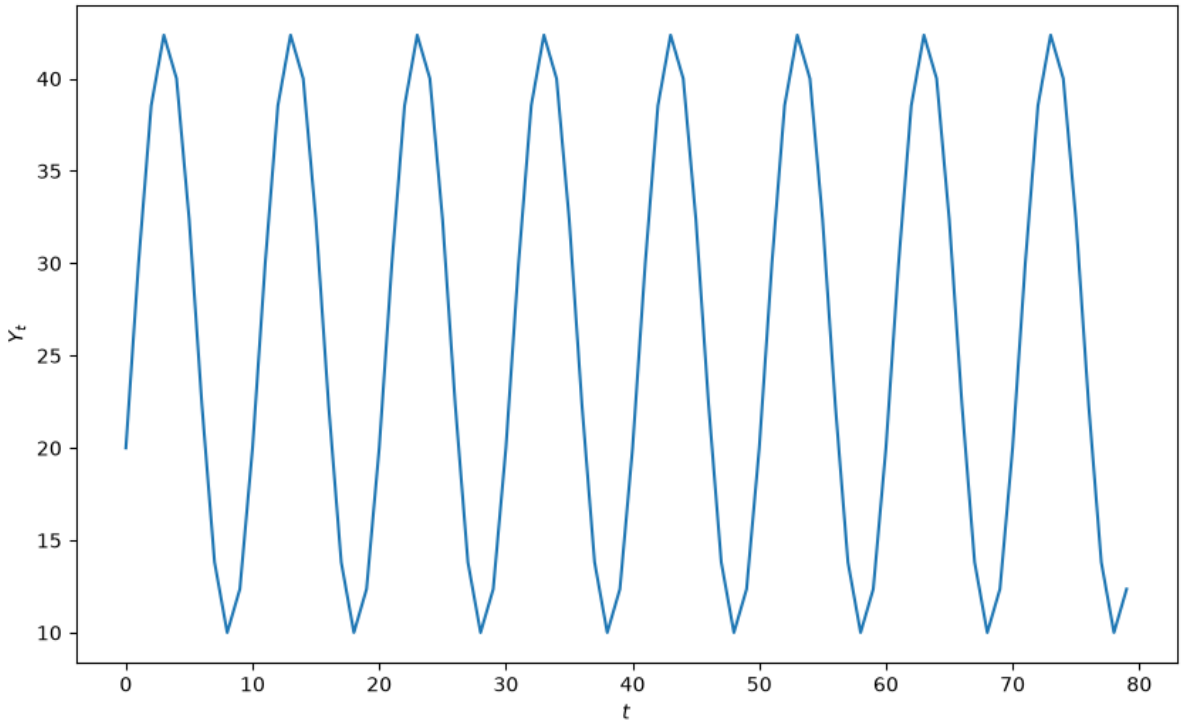
# Apply the reverse-engineering function f
ρ1, ρ2, α, β = f(r, ϕ)

# Extract real parts for numerical computation
α = α.real
β = β.real

print(f"α, β = {α:.2f}, {β:.2f}")

ytemp = y_nonstochastic(α=α, β=β, y_0=20, y_1=30)
plot_y(ytemp)
```

```
α, β = 0.62, 1.00
Solution type: Damped oscillations
Roots are [0.80901699+0.58778525j 0.80901699-0.58778525j]
Root type: Complex
Stability: Stable
```



44.3.7 Digression: Using Sympy to find roots

We can also use sympy to compute analytic formulas for the roots

```
init_printing()

r1 = Symbol("ρ_1")
r2 = Symbol("ρ_2")
z = Symbol("z")

sympy.solve(z**2 - r1 * z - r2, z)
```

$$\left[\rho_1/2 - \sqrt{\rho_1^2 + 4\rho_2}/2,\ \rho_1/2 + \sqrt{\rho_1^2 + 4\rho_2}/2\right]$$

```
α = Symbol("α")
β = Symbol("β")
r1 = α + β
r2 = -β

sympy.solve(z**2 - r1 * z - r2, z)
```

$$\left[\alpha/2 + \beta/2 - \sqrt{\alpha^2 + 2\alpha\beta + \beta^2 - 4\beta}/2,\ \alpha/2 + \beta/2 + \sqrt{\alpha^2 + 2\alpha\beta + \beta^2 - 4\beta}/2\right]$$

44.4 Stochastic shocks

Now we'll construct some code to simulate the stochastic version of the model that emerges when we add a random shock process to aggregate demand

```
def y_stochastic(y_0=0, y_1=0, α=0.8, β=0.2, γ=10, n=100, σ=5):
    """
    This function takes parameters of a stochastic version of
    the model, analyzes the roots of the characteristic
    polynomial and generates a simulation.
    """

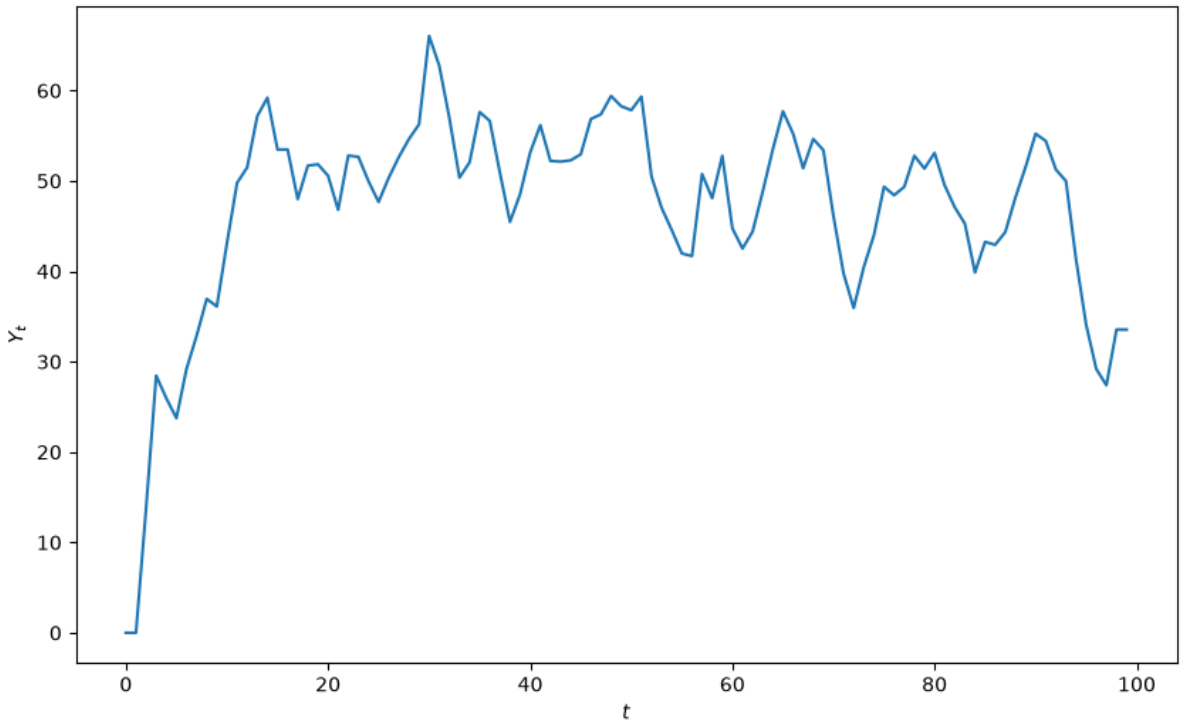
    y_series, analysis = simulate_samuelson(y_0, y_1, α, β, γ, n, σ, None, 42)

    print(f"Solution type: {analysis['solution_type']}")
    print(f"Roots are {[f'{root:.2f}' for root in analysis['roots']]}")
    print(f"Root type: {'Complex' if analysis['is_complex'] else 'Real'}")
    print(f"Stability: {'Stable' if analysis['is_stable'] else 'Unstable'}")

    return y_series


plot_y(y_stochastic())
```

```
Solution type: Steady state convergence
Roots are ['0.72', '0.28']
Root type: Real
Stability: Stable
```



Let's do a simulation in which there are shocks and the characteristic polynomial has complex roots

```
r = 0.97

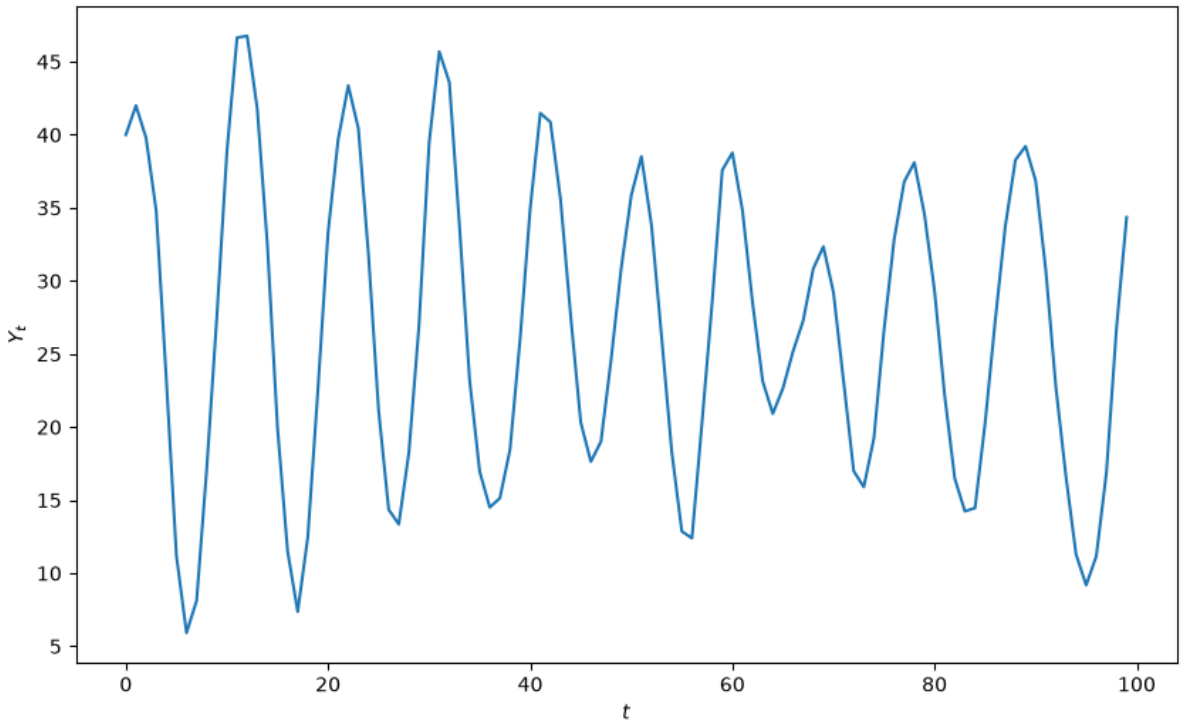
period = 10  # Length of cycle in units of time
ϕ = 2 * math.pi / period

# Apply the reverse-engineering function f
ρ1, ρ2, α, β = f(r, ϕ)

# Extract real parts for numerical computation
α = α.real
β = β.real

print(f"α, β = {α:.2f}, {β:.2f}")
plot_y(y_stochastic(y_0=40, y_1=42, α=α, β=β, σ=2, n=100))
```

```
α, β = 0.63, 0.94
Solution type: Damped oscillations
Roots are ['0.78+0.57j', '0.78-0.57j']
Root type: Complex
Stability: Stable
```



44.5 Government spending

This function computes a response to either a permanent or one-off increase in government expenditures

```
def y_stochastic_g(
    y_0=20, y_1=20, α=0.8, β=0.2, γ=10,
    n=100, σ=2, g=0, g_t=0, duration="permanent"
):
    """
    This program computes a response to a permanent increase
    in government expenditures that occurs at time 20
    """

    g_params = (
        {'g': g, 'g_t': g_t, 'duration': duration} if g != 0 else None
    )
    y_series, analysis = simulate_samuelson(
        y_0, y_1, α, β, γ, n, σ, g_params, 42
    )

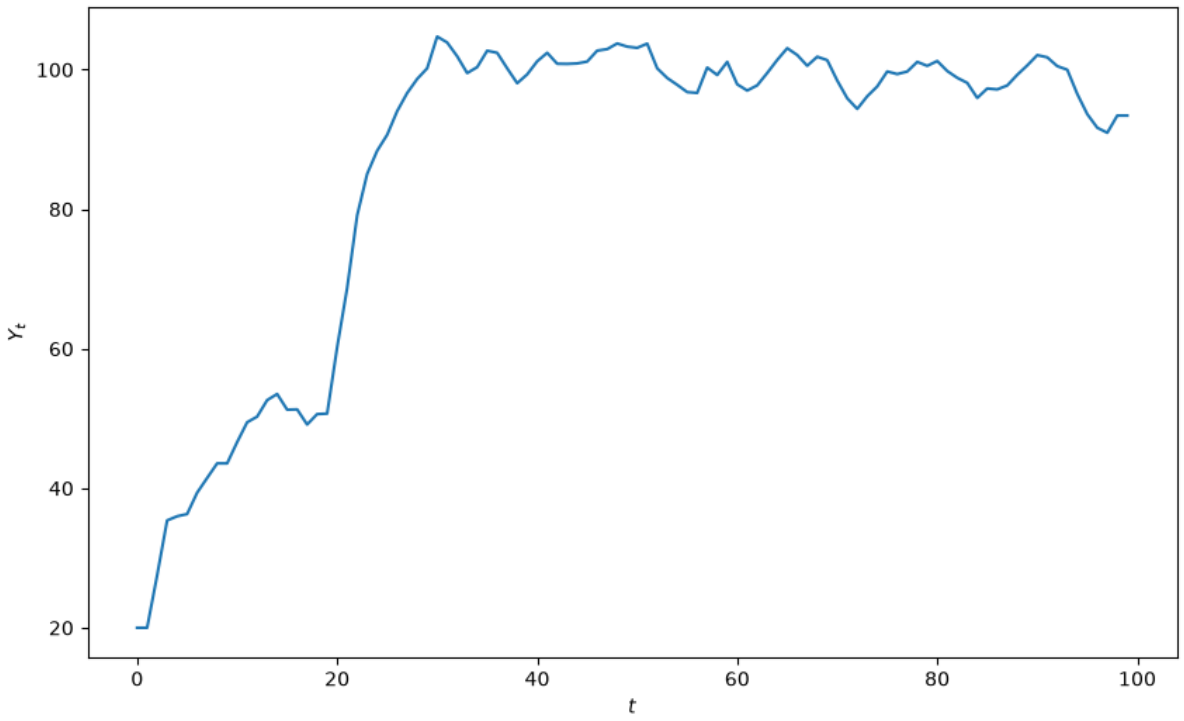
    print(f"Solution type: {analysis['solution_type']}")
    print(f"Roots: {analysis['roots']}")
    print(f"Root type: {'Complex' if analysis['is_complex'] else 'Real'}")
    print(f"Stability: {'Stable' if analysis['is_stable'] else 'Unstable'}")

    return y_series
```

A permanent government spending shock can be simulated as follows

```
plot_y(y_stochastic_g(g=10, g_t=20, duration="permanent"))
```

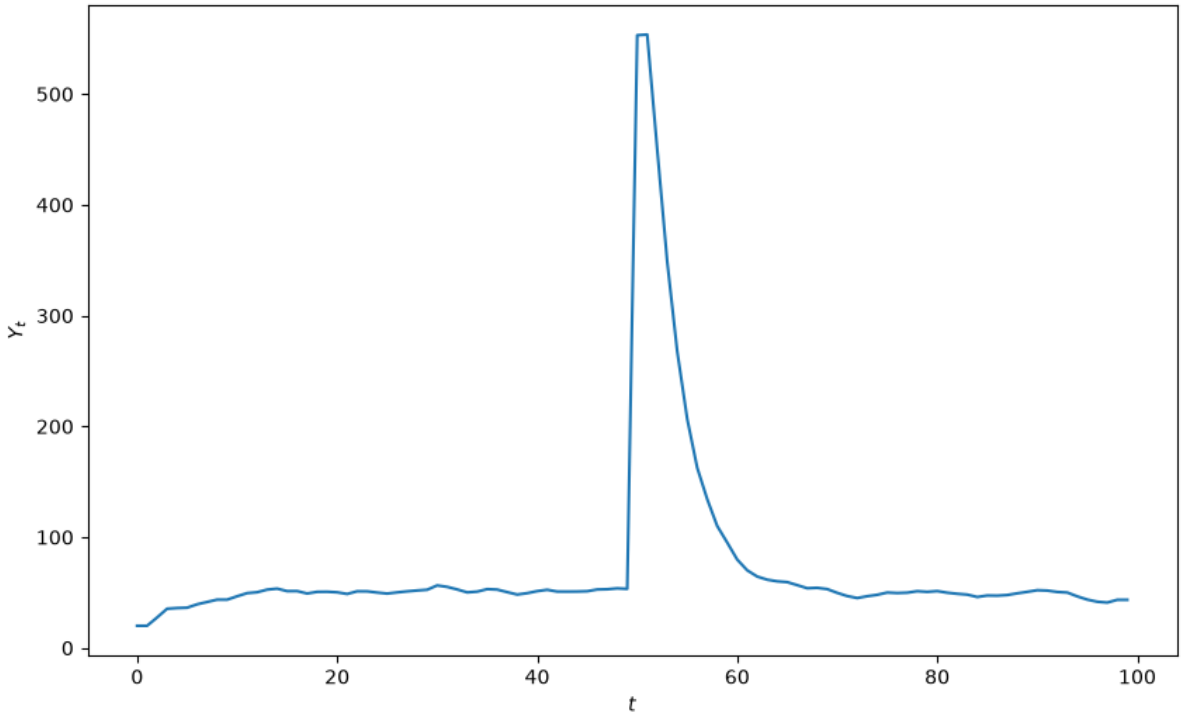
```
Solution type: Steady state convergence
Roots: [0.7236068 0.2763932]
Root type: Real
Stability: Stable
```



We can also see the response to a one time jump in government expenditures

```
plot_y(y_stochastic_g(g=500, g_t=50, duration="one-off"))
```

```
Solution type: Steady state convergence
Roots: [0.7236068 0.2763932]
Root type: Real
Stability: Stable
```



44.6 Wrapping everything into a class

Up to now, we have written functions to do the work.

Now we'll roll up our sleeves and write a Python class called `Samuelson` for the Samuelson model

```
class Samuelson:
    r"""
    This class represents the Samuelson model, otherwise known as the
    multiplier-accelerator model.
    The model combines the Keynesian multiplier
    with the accelerator theory of investment.

    The path of output is governed by a linear
    second-order difference equation

    .. math::

        Y_t = (\alpha + \beta) Y_{t-1} - \beta Y_{t-2} + (\gamma + G_t) + \sigma \
↪epsilon_t

    Parameters
    ----------
    y_0 : scalar
        Initial condition for Y_0
    y_1 : scalar
        Initial condition for Y_1
    α : scalar
        Marginal propensity to consume
    β : scalar
```

(continues on next page)

(continued from previous page)

```
        Accelerator coefficient
    n : int
        Number of iterations
    σ : scalar
        Volatility parameter. It must be greater than or equal to 0. Set
        equal to 0 for a non-stochastic model.
    g : scalar
        Government spending shock
    g_t : int
        Time at which government spending shock occurs. Must be specified
        when duration != None.
    duration : {None, 'permanent', 'one-off'}
        Specifies type of government spending shock. If none, government
        spending equal to g for all t.

    """

    def __init__(
        self, y_0=100, y_1=50,
        α=1.3, β=0.2, γ=10, n=100, σ=0, g=0, g_t=0, duration=None
    ):

        self.y_0, self.y_1, self.α, self.β = y_0, y_1, α, β
        self.n, self.g, self.g_t, self.duration = n, g, g_t, duration
        self.γ, self.σ = γ, σ

        # Use unified analysis function
        self.analysis = analyze_roots(α, β, verbose=False)
        self.ρ1, self.ρ2 = self.analysis['rho1'], self.analysis['rho2']
        self.roots = self.analysis['roots']

    def root_type(self):
        return "Complex conjugate" if self.analysis['is_complex'] else "Real"

    def root_less_than_one(self):
        return self.analysis['is_stable']

    def solution_type(self):
        return self.analysis['solution_type']

    def generate_series(self, seed=0):
        g_params = (
            {'g': self.g, 'g_t': self.g_t, 'duration': self.duration}
            if self.g != 0 else None
        )
        y_series, _ = simulate_samuelson(
            self.y_0, self.y_1, self.α, self.β, self.γ,
            self.n, self.σ, g_params, seed
        )
        return y_series

    def summary(self):
        print("Summary\n" + "-" * 50)
        print(f"Root type: {self.root_type()}")
        print(f"Solution type: {self.solution_type()}")
        print(f"Roots: {str(self.roots)}")
```

(continues on next page)

(continued from previous page)

```
        if self.root_less_than_one() == True:
            print("Absolute value of roots is less than one")
        else:
            print("Absolute value of roots is not less than one")

        if self.σ > 0:
            print("Stochastic series with σ = " + str(self.σ))
        else:
            print("Non-stochastic series")

        if self.g != 0:
            print("Government spending equal to " + str(self.g))

        if self.duration != None:
            print(
                self.duration.capitalize()
                + " government spending shock at t = "
                + str(self.g_t)
            )

    def plot(self, seed=0):
        fig, ax = plt.subplots(figsize=(10, 6))
        ax.plot(self.generate_series(seed))
        ax.set(xlabel="iteration", xlim=(0, self.n))
        ax.set_ylabel("$Y_t$", rotation=0)

        # Display model parameters on the plot
        paramstr = (
            f"$α={self.α:.2f}$ \n $β={self.β:.2f}$ \n "
            f"$\\gamma={self.γ:.2f}$ \n $\\sigma={self.σ:.2f}$ \n "
            f"$\\rho_1={self.ρ1:.2f}$ \n $\\rho_2={self.ρ2:.2f}$"
        )
        props = dict(fc="white", pad=10, alpha=0.5)
        ax.text(
            0.87,
            0.05,
            paramstr,
            transform=ax.transAxes,
            fontsize=12,
            bbox=props,
            va="bottom",
        )

        return fig

    def param_plot(self):


        fig = param_plot()
        ax = fig.gca()

        # Display eigenvalues in the legend
        for i, root in enumerate(self.roots):
            if isinstance(root, complex):

                # Handle sign formatting for complex number display
                operator = ["+", ""]
```

(continues on next page)

(continued from previous page)

```
                root_real = self.roots[i].real
                root_imag = self.roots[i].imag
                label = (
                    rf"$\lambda_{i+1} = {root_real:.2f}"
                    rf"{operator[i]} {root_imag:.2f}i$"
                )
            else:
                label = rf"$\lambda_{i+1} = {self.roots[i].real:.2f}$"

            # Add invisible point for legend entry
            ax.scatter(
                0, 0, s=0, label=label
            )

        # Mark current parameter values on the stability diagram
        ax.scatter(
            self.ρ1,
            self.ρ2,
            s=100,
            c="red",
            marker="+",
            label=r"$(\rho_1, \rho_2)$",
            zorder=5,
        )

        plt.legend(fontsize=12, loc=3)

        return fig
```

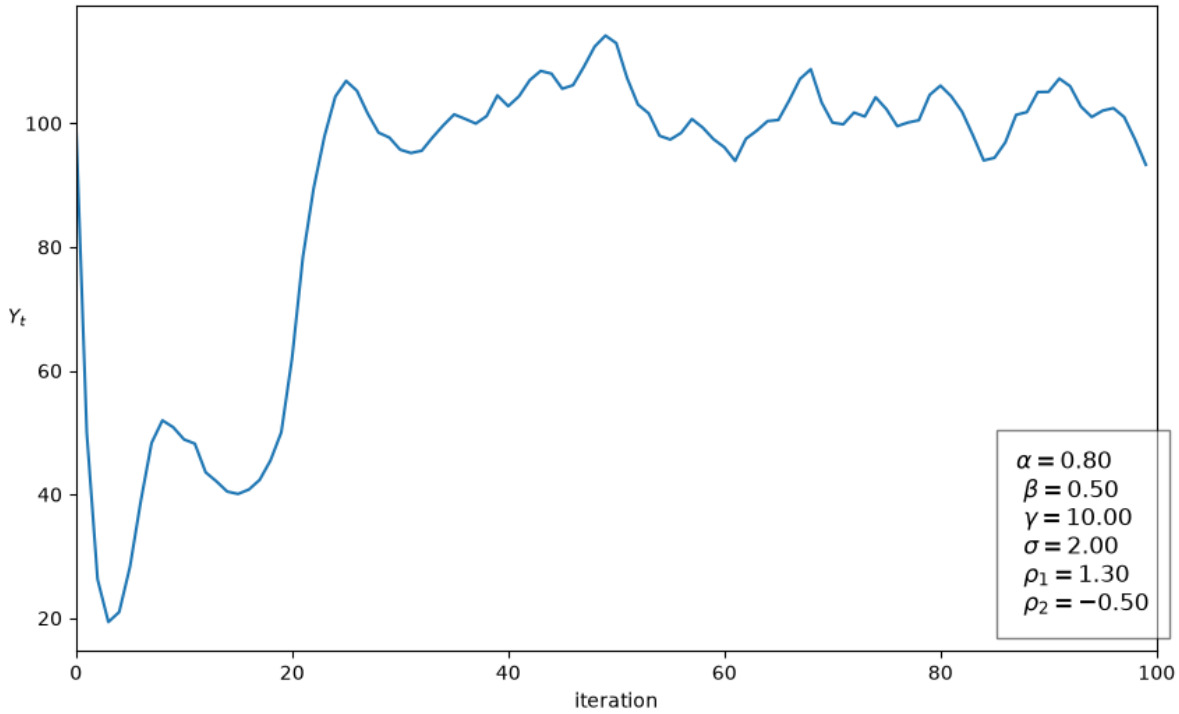
44.6.1 Illustration of Samuelson class

Now we'll put our Samuelson class to work on an example

```
sam = Samuelson(α=0.8, β=0.5, σ=2, g=10, g_t=20, duration="permanent")
sam.summary()
```

```
Summary
-----------------------------------------------
Root type: Complex conjugate
Solution type: Damped oscillations
Roots: [0.65+0.27838822j 0.65-0.27838822j]
Absolute value of roots is less than one
Stochastic series with σ = 2
Government spending equal to 10
Permanent government spending shock at t = 20
```

```
sam.plot()
plt.show()
```

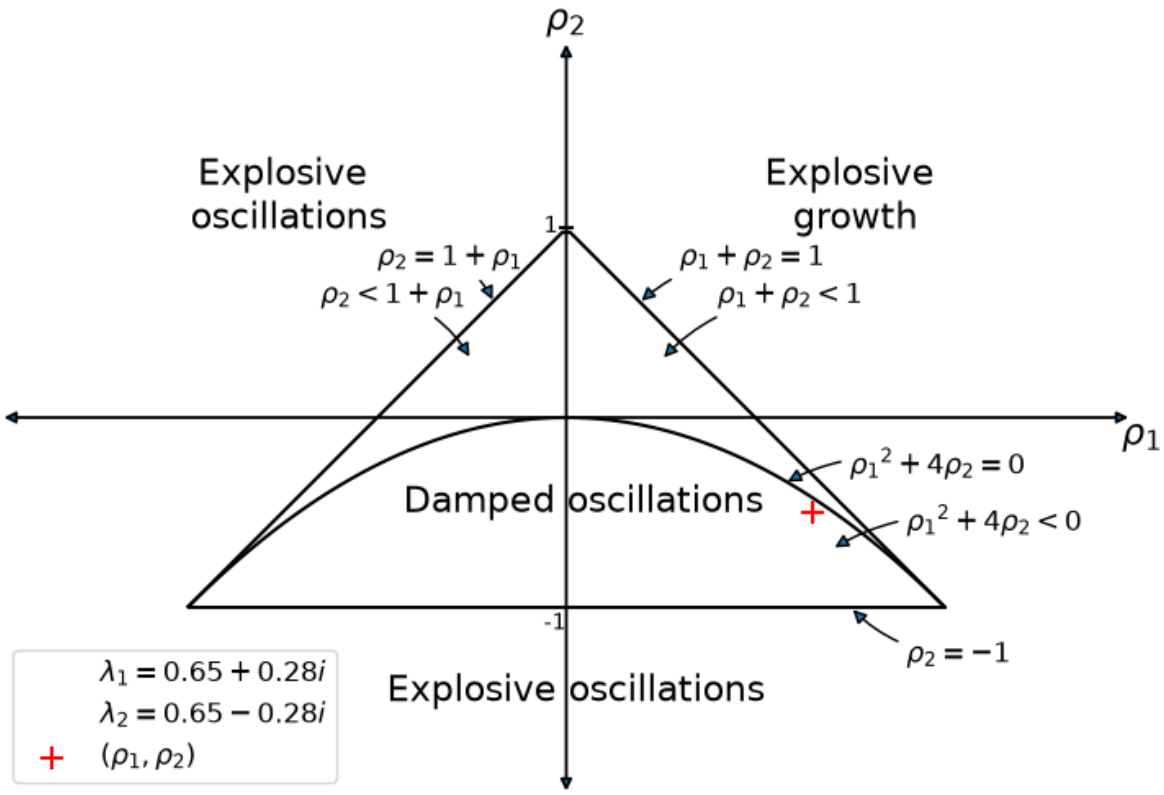


44.6.2 Using the graph

We'll use our graph to show where the roots lie and how their location is consistent with the behavior of the path just graphed.

The red + sign shows the location of the roots

```
sam.param_plot()
plt.show()
```



44.7 Using the LinearStateSpace class

It turns out that we can use the QuantEcon.py LinearStateSpace class to do much of the work that we have done from scratch above.

Here is how we map the Samuelson model into an instance of a `LinearStateSpace` class

```
α = 0.8
β = 0.9
ρ1 = α + β
ρ2 = -β
γ = 10
σ = 1
g = 10
n = 100

A = [[1, 0, 0], [γ + g, ρ1, ρ2], [0, 1, 0]]

G = [
    [γ + g, ρ1, ρ2],   # Y_{t+1}
    [γ, α, 0],         # C_{t+1}
    [0, β, -β],        # I_{t+1}
]

μ_0 = [1, 100, 50]
C = np.zeros((3, 1))
C[1] = σ  # Shock variance
```

(continues on next page)

(continued from previous page)

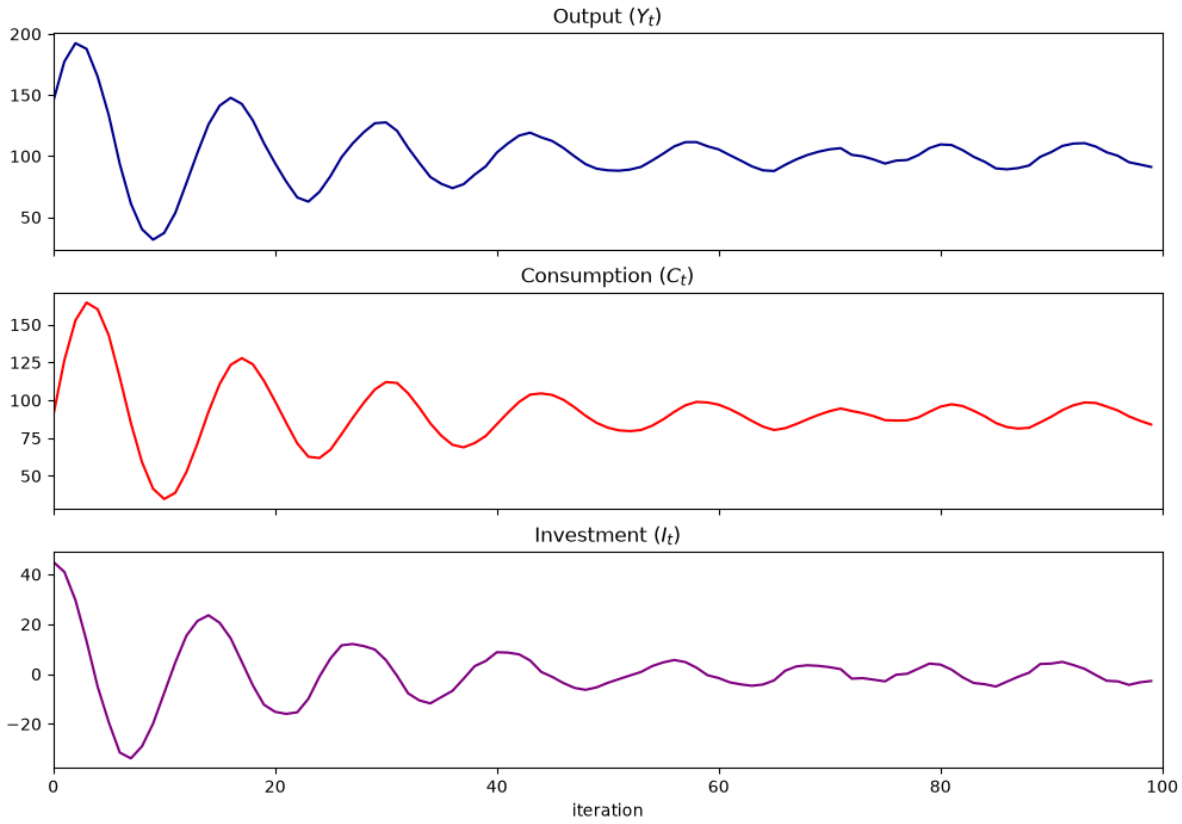
```
sam_t = LinearStateSpace(A, C, G, mu_0=μ_0)

rng = np.random.default_rng()
x, y = sam_t.simulate(ts_length=n, random_state=rng)

fig, axes = plt.subplots(3, 1, sharex=True, figsize=(12, 8))
titles = ["Output ($Y_t$)", "Consumption ($C_t$)", "Investment ($I_t$)"]
colors = ["darkblue", "red", "purple"]
for ax, series, title, color in zip(axes, y, titles, colors):
    ax.plot(series, color=color)
    ax.set(title=title, xlim=(0, n))

axes[-1].set_xlabel("iteration")

plt.show()
```



44.7.1 Other methods in the `LinearStateSpace` class

Let's plot *impulse response functions* for the instance of the Samuelson model using a method in the `LinearStateSpace` class

```
imres = sam_t.impulse_response()
imres = np.asarray(imres)
y1 = imres[:, :, 0]
y2 = imres[:, :, 1]
y1.shape
```

$$(2,\ 6,\ 1)$$

Now let's compute the zeros of the characteristic polynomial by simply calculating the eigenvalues of A

```
A = np.asarray(A)
w, v = np.linalg.eig(A)
print(np.round(w, 2))
```

```
[0.85+0.42j 0.85-0.42j 1.  +0.j  ]
```

44.7.2 Inheriting methods from `LinearStateSpace`

We could also create a subclass of `LinearStateSpace` (inheriting all its methods and attributes) to add more functions to use

```
class SamuelsonLSS(LinearStateSpace):
    """
    This subclass creates a Samuelson multiplier-accelerator model
    as a linear state space system.
    """

    def __init__(self, y_0=100, y_1=50, α=0.8, β=0.9, γ=10, σ=1, g=10):

        self.α, self.β = α, β
        self.y_0, self.y_1, self.g = y_0, y_1, g
        self.γ, self.σ = γ, σ

        # Set initial state vector
        self.initial_μ = [1, y_0, y_1]

        self.ρ1 = α + β
        self.ρ2 = -β

        # Construct state transition matrix
        self.A = [[1, 0, 0], [γ + g, self.ρ1, self.ρ2], [0, 1, 0]]

        # Construct observation matrix
        self.G = [
            [γ + g, self.ρ1, self.ρ2],  # Y_{t+1}
            [γ, α, 0],                  # C_{t+1}
            [0, β, -β],                 # I_{t+1}
        ]

        self.C = np.zeros((3, 1))
```

(continues on next page)

(continued from previous page)

```
        self.C[1] = σ  # Shock variance

        # Initialize the LinearStateSpace instance
        LinearStateSpace.__init__(
            self, self.A, self.C, self.G, mu_0=self.initial_μ
        )

    # Create unicode aliases for mu_0 and Sigma_0 in the parent class
    @property
    def μ_0(self):
        return self.mu_0

    @μ_0.setter
    def μ_0(self, value):
        self.mu_0 = value

    @property
    def Σ_0(self):
        return self.Sigma_0

    @Σ_0.setter
    def Σ_0(self, value):
        self.Sigma_0 = value

    def plot_simulation(self, ts_length=100, stationary=True, seed=0):

        # Store original distribution parameters
        temp_μ = self.μ_0
        temp_Σ = self.Σ_0

        # Use stationary distribution for simulation
        if stationary == True:
            try:
                (self.μ_x, self.μ_y, self.Σ_x, self.Σ_y, self.Σ_yx
                ) = self.stationary_distributions()
                self.μ_0 = self.μ_x
                self.Σ_0 = self.Σ_x

            # Handle case where stationary distribution doesn't exist
            except ValueError:
                print("Stationary distribution does not exist")

        rng = np.random.default_rng(seed)
        x, y = self.simulate(ts_length, random_state=rng)

        fig, axes = plt.subplots(3, 1, sharex=True, figsize=(12, 8))
        titles = ["Output ($Y_t$)",
                  "Consumption ($C_t$)",
                  "Investment ($I_t$)"]
        colors = ["darkblue", "red", "purple"]
        for ax, series, title, color in zip(axes, y, titles, colors):
            ax.plot(series, color=color)
            ax.set(title=title, xlim=(0, n))

        axes[-1].set_xlabel("iteration")
        plt.show()
```

(continues on next page)

(continued from previous page)

```
        # Restore original distribution parameters
        self.μ_0 = temp_μ
        self.Σ_0 = temp_Σ

    def plot_irf(self, j=5):

        x, y = self.impulse_response(j)

        # Reshape impulse responses for plotting
        yimf = np.array(y).flatten().reshape(j + 1, 3).T

        fig, axes = plt.subplots(3, 1, sharex=True, figsize=(12, 8))
        labels = ["$Y_t$", "$C_t$", "$I_t$"]
        colors = ["darkblue", "red", "purple"]
        for ax, series, label, color in zip(axes, yimf, labels, colors):
            ax.plot(series, color=color)
            ax.set(xlim=(0, j))
            ax.set_ylabel(label, rotation=0, fontsize=14, labelpad=10)

        axes[0].set_title("Impulse response functions")
        axes[-1].set_xlabel("iteration")
        plt.show()

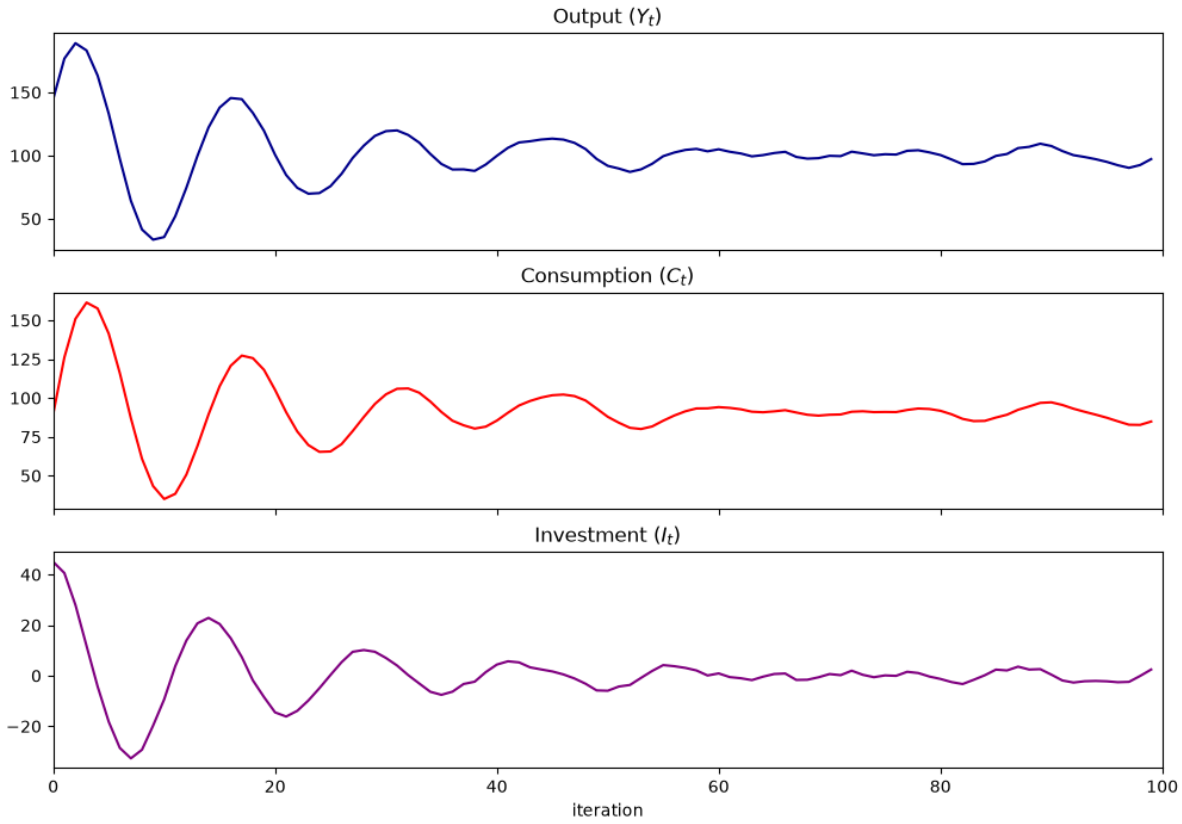
    def multipliers(self, j=5):
        x, y = self.impulse_response(j)
        return np.sum(np.array(y).flatten().reshape(j + 1, 3), axis=0)
```

44.7.3 Illustrations

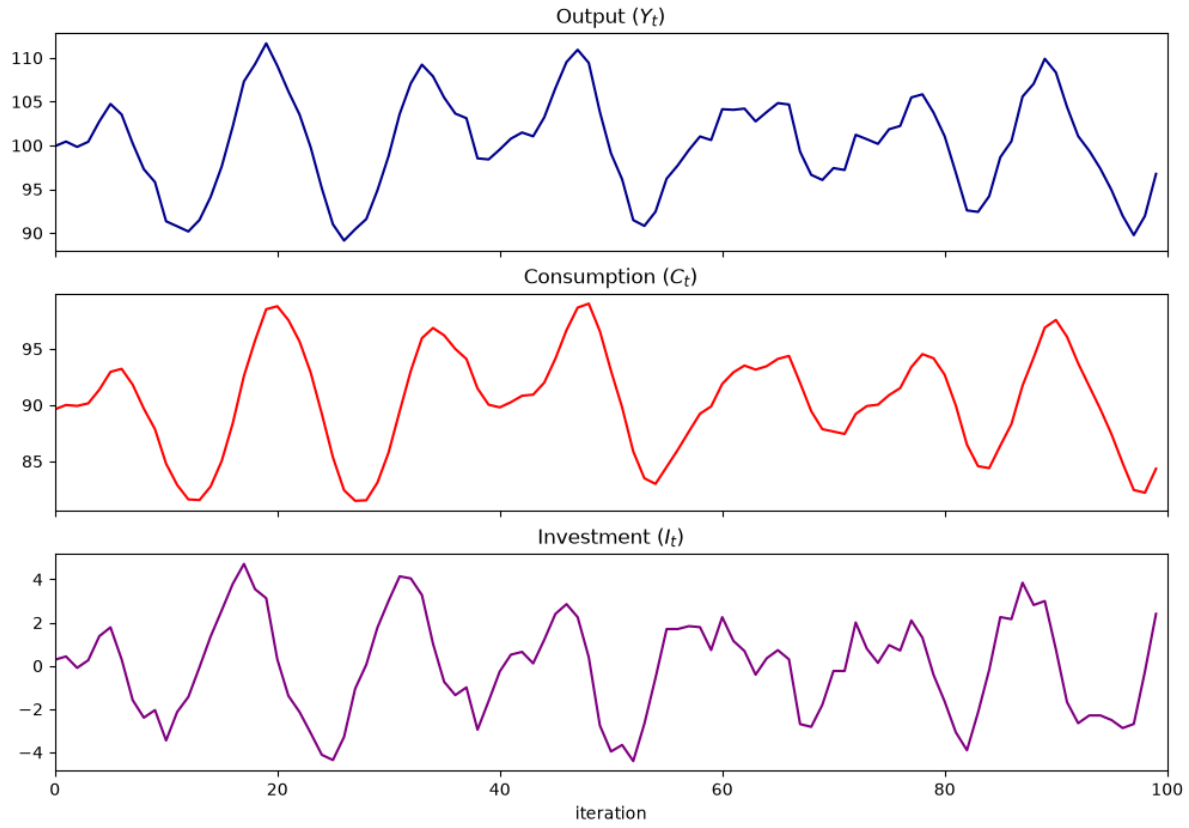
Let's show how we can use the `SamuelsonLSS`

```
samlss = SamuelsonLSS()
```

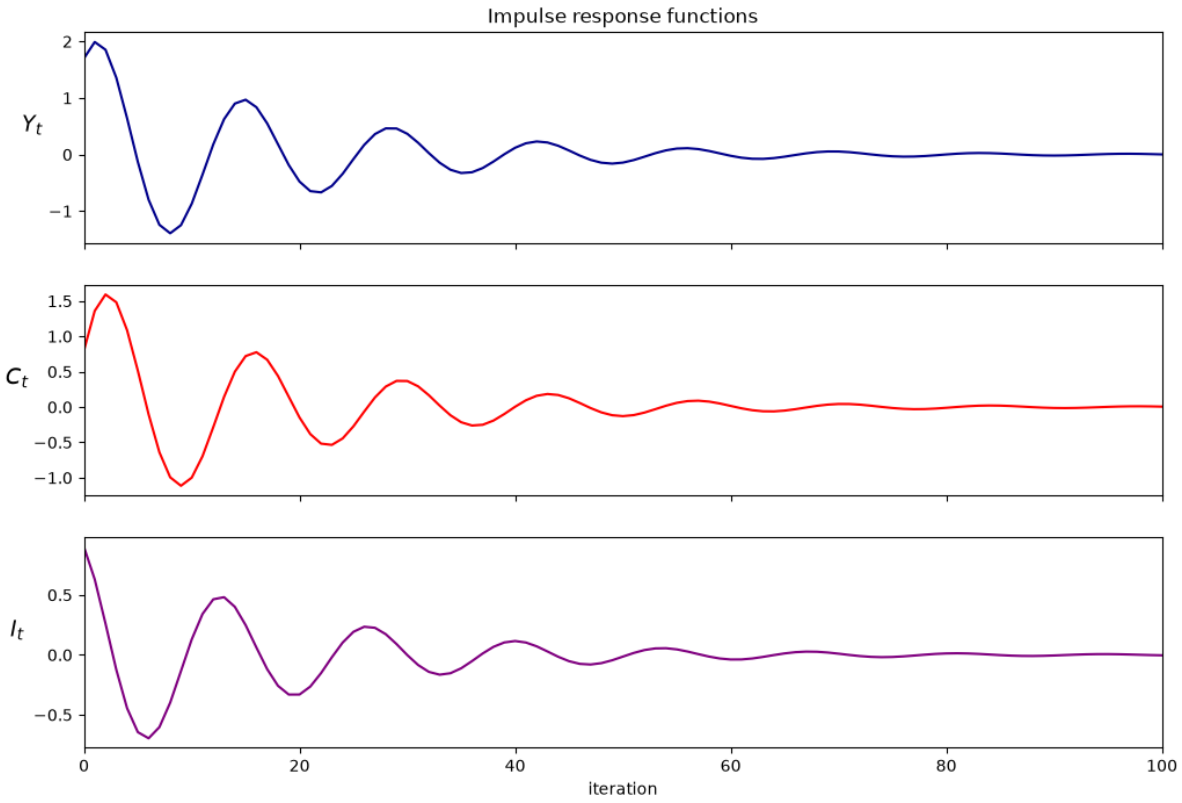
```
samlss.plot_simulation(100, stationary=False)
plt.show()
```




```
samlss.plot_simulation(100, stationary=True)
plt.show()
```




```
samlss.plot_irf(100)
plt.show()
```




```
samlss.multipliers()
```

```
array([7.414389, 6.835896, 0.578493])
```

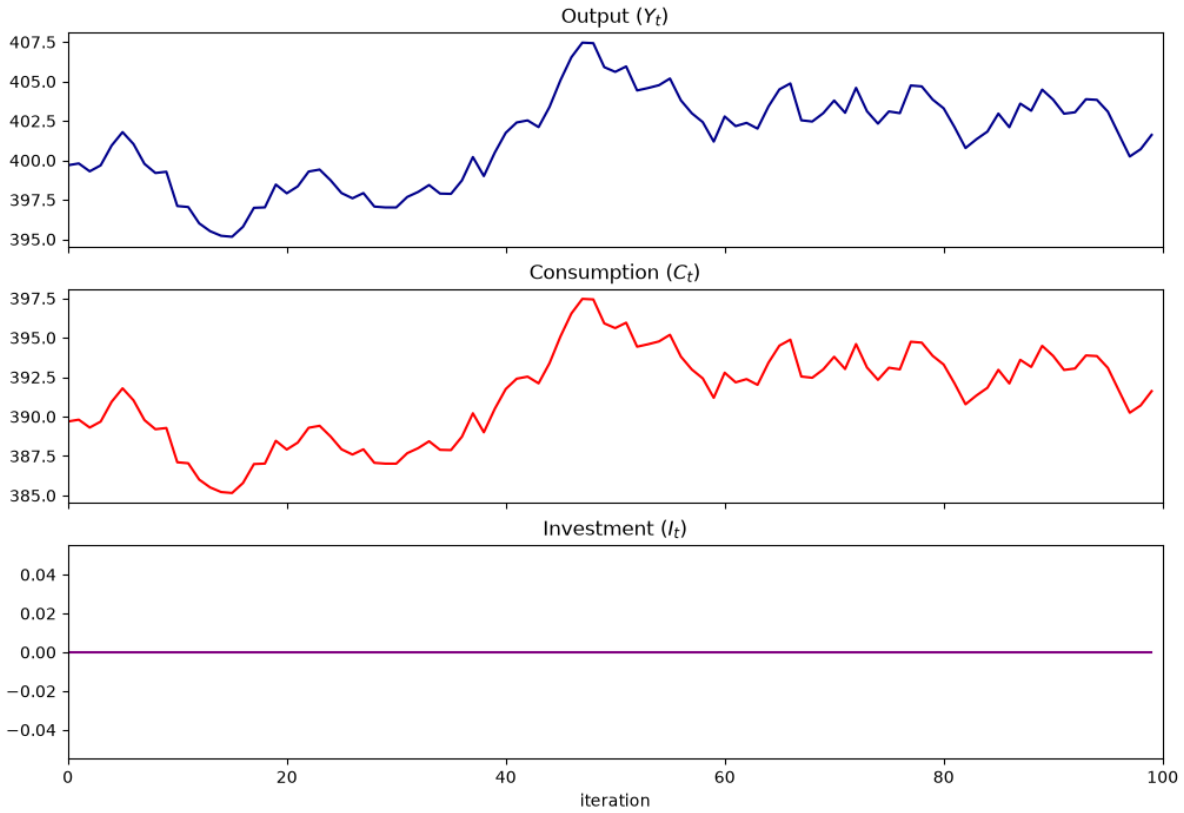
44.8 Pure multiplier model

Let's shut down the accelerator by setting $b = 0$ to get a pure multiplier model

- the absence of cycles gives an idea about why Samuelson included the accelerator

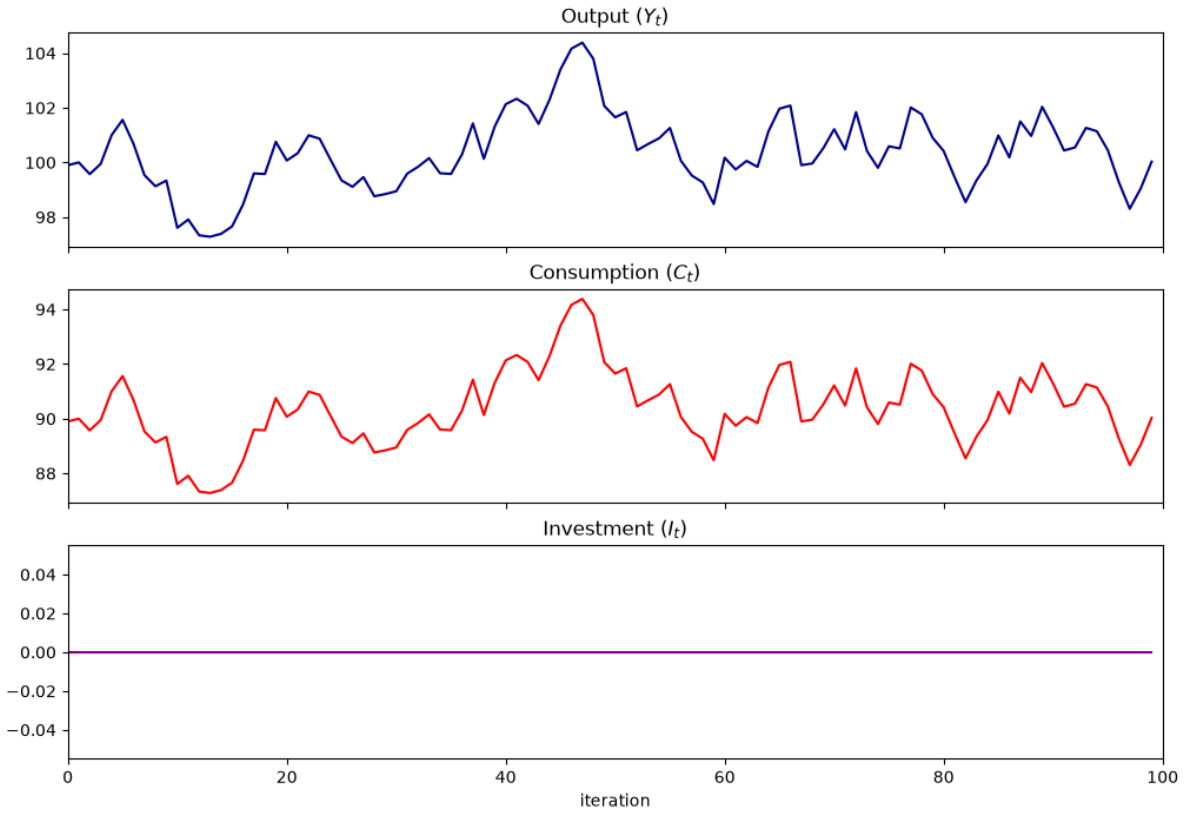
```
pure_multiplier = SamuelsonLSS(α=0.95, β=0)
```

```
pure_multiplier.plot_simulation()
```

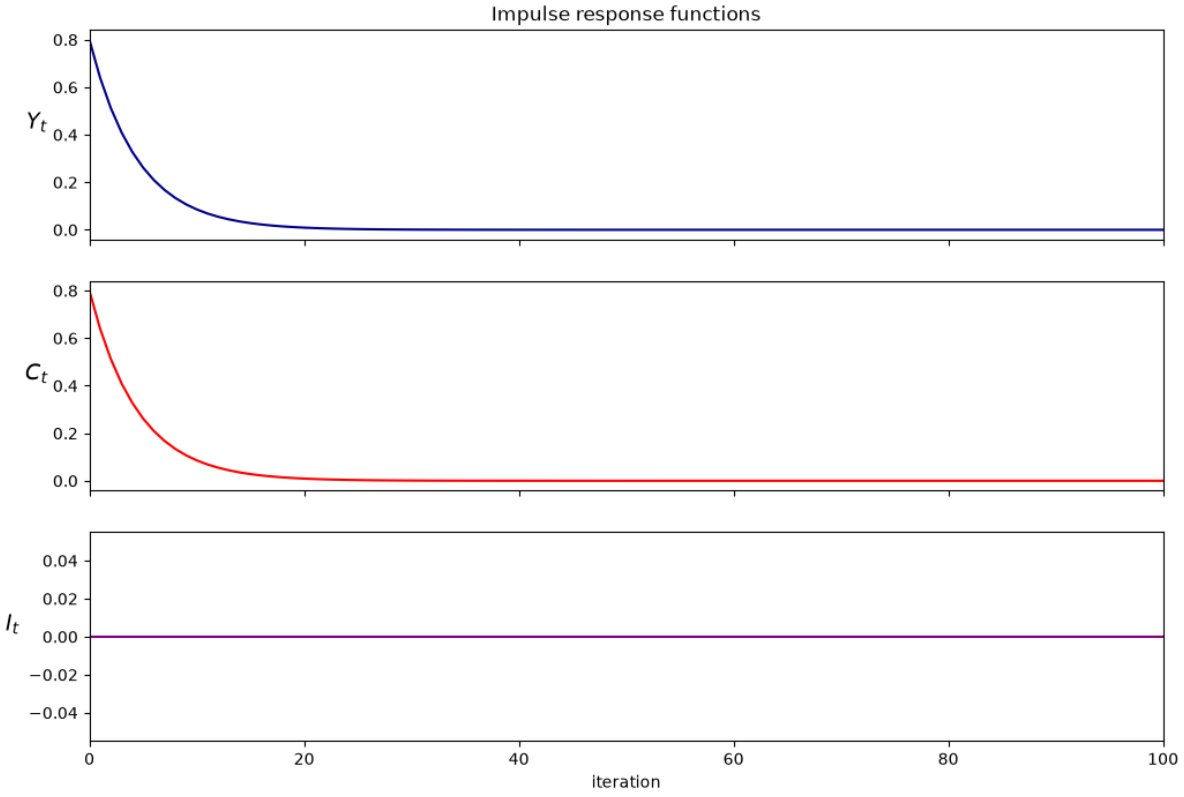



```
pure_multiplier = SamuelsonLSS(α=0.8, β=0)
```

```
pure_multiplier.plot_simulation()
```




```
pure_multiplier.plot_irf(100)
```



44.9 Summary

In this lecture, we wrote functions and classes to represent non-stochastic and stochastic versions of the Samuelson (1939) multiplier-accelerator model, described in [Samuelson, 1939].

We saw that different parameter values led to different output paths, which could either be stationary, explosive, or oscillating.

We also were able to represent the model using the QuantEcon.py LinearStateSpace class.

CHAPTER

FORTYFIVE

THE ACCELERATION PRINCIPLE AND THE NATURE OF BUSINESS CYCLES

Contents

- *The Acceleration Principle and the Nature of Business Cycles*
  - *Overview*
  - *Empirical foundation for the acceleration principle*
  - *Acceleration enables oscillations*
  - *A linear system with shocks*
  - *From autocovariances to spectra*
  - *Spectral peaks in the Hansen-Samuelson model*
  - *Real roots can produce peaks in general models*
  - *A calibrated model in the frequency domain*
  - *Summary*
  - *Exercises*

45.1 Overview

This lecture studies two classic papers by Gregory Chow:

- Chow [1968] presents empirical evidence for the acceleration principle, describes how acceleration promotes oscillations, and analyzes conditions for the emergence of spectral peaks in linear difference equation subjected to random shocks
- Chow and Levitan [1969] presents a spectral analysis of a calibrated US macroeconometric model and teaches about spectral gains, coherences, and lead–lag patterns

These papers are related to ideas in the following lectures:

- The multiplier–accelerator mechanism in *Samuelson Multiplier-Accelerator*
- Linear stochastic difference equations and autocovariances in *Linear State Space Models*
- Eigenmodes of multivariate dynamics in *VARs and DMDs*

- Fourier ideas in *Circulant Matrices* (and, for empirical estimation, the advanced lecture Estimation of Spectra)

Chow [1968] builds on earlier empirical work testing the acceleration principle on US investment data.

We start with that empirical evidence before developing the theoretical framework.

We will keep returning to three ideas:

- In deterministic models, oscillations indicate complex eigenvalues of a transition matrix.
- In stochastic models, a "cycle" shows up as a local peak in a (univariate) spectral density.
- Spectral peaks depend on eigenvalues, but also on how shocks enter and on how observables load on eigenmodes.

Let's start with some standard imports:

```
import numpy as np
import matplotlib.pyplot as plt
```

We will use the following helper functions throughout the lecture

```
def spectral_density_var1(A, V, ω_grid):
    """Spectral density matrix for VAR(1): y_t = A y_{t-1} + u_t."""
    A, V = np.asarray(A), np.asarray(V)
    n = A.shape[0]
    I = np.eye(n)
    F = np.empty((len(ω_grid), n, n), dtype=complex)
    for k, ω in enumerate(ω_grid):
        H = np.linalg.inv(I - np.exp(-1j * ω) * A)
        F[k] = (H @ V @ H.conj().T) / (2 * np.pi)
    return F

def spectrum_of_linear_combination(F, b):
    """Spectrum of x_t = b'y_t given the spectral matrix F(ω)."""
    b = np.asarray(b).reshape(-1, 1)
    return np.array([np.real((b.T @ F[k] @ b).item())
                                for k in range(F.shape[0])])

def simulate_var1(A, V, T, burn=200, seed=1234):
    r"""Simulate y_t = A y_{t-1} + u_t with u_t \sim N(0, V)."""
    rng = np.random.default_rng(seed)
    A, V = np.asarray(A), np.asarray(V)
    n = A.shape[0]
    chol = np.linalg.cholesky(V)
    y = np.zeros((T + burn, n))

    for t in range(1, T + burn):
        y[t] = A @ y[t - 1] + chol @ rng.standard_normal(n)

    return y[burn:]

def sample_autocorrelation(x, max_lag):
    """Sample autocorrelation of a 1d array from lag 0 to max_lag."""
    x = np.asarray(x)
    x = x - x.mean()
    denom = np.dot(x, x)
    acf = np.empty(max_lag + 1)
    for k in range(max_lag + 1):
        acf[k] = np.dot(x[:-k] if k else x, x[k:]) / denom
    return acf
```

45.2 Empirical foundation for the acceleration principle

Chow [1968] opens by reviewing empirical evidence for the acceleration principle from earlier macroeconometric work.

Using annual observations for 1931–40 and 1948–63, Chow tested the acceleration equation on three investment categories:

- new construction
- gross private domestic investment in producers' durable equipment combined with change in business inventories
- the last two variables separately

In each case, when the regression included both Y_t and Y_{t-1} (where Y is gross national product minus taxes net of transfers), the coefficient on Y_{t-1} was of *opposite sign* and slightly smaller in absolute value than the coefficient on Y_t.

Equivalently, when expressed in terms of ΔY_t and Y_{t-1}, the coefficient on Y_{t-1} was a small fraction of the coefficient on ΔY_t.

45.2.1 An example: automobile demand

Chow presents a clean illustration using data on net investment in automobiles from his earlier work on automobile demand.

Using annual data for 1922–41 and 1948–57, he estimates by least squares:

$$y_t^n = \underset{(0.0022)}{0.0155} Y_t \; \underset{(0.0020)}{-0.0144}\, Y_{t-1} \; \underset{(0.0056)}{-0.0239}\, p_t \; \underset{(0.0040)}{+0.0199}\, p_{t-1} + \underset{(0.101)}{0.351} y_{t-1}^n + \text{const.} \tag{45.1}$$

where:

- Y_t is real disposable personal income per capita
- p_t is a relative price index for automobiles
- y_t^n is per capita net investment in passenger automobiles
- standard errors appear in parentheses

The key observation: the coefficients on Y_{t-1} and p_{t-1} are *the negatives* of the coefficients on Y_t and p_t.

This pattern is exactly what the acceleration principle predicts.

45.2.2 From stock adjustment to acceleration

The empirical support for acceleration should not be surprising once we accept a stock-adjustment demand equation for capital:

$$s_{it} = a_i Y_t + b_i s_{i,t-1} \tag{45.2}$$

where s_{it} is the stock of capital good i.

The acceleration equation (45.1) is essentially the *first difference* of (45.2).

Net investment is the change in stock, $y_{it}^n = \Delta s_{it}$, and differencing (45.2) gives:

$$y_{it}^n = a_i \Delta Y_t + b_i y_{i,t-1}^n \tag{45.3}$$

The coefficients on Y_t and Y_{t-1} in the level form are a_i and $-a_i(1 - b_i)$ respectively.

They are opposite in sign and similar in magnitude when b_i is not too far from unity.

This connection between stock adjustment and acceleration is central to Chow's argument about why acceleration matters for business cycles.

45.3 Acceleration enables oscillations

Having established the empirical evidence for acceleration, we now examine why it matters theoretically for generating oscillations.

Chow [1968] asks a fundamental question: if we build a macro model using only standard demand equations with simple distributed lags, can the system generate sustained oscillations?

He shows that, under natural sign restrictions, the answer is no.

Stock-adjustment demand for durable goods leads to investment equations where the coefficient on Y_{t-1} is negative.

This negative coefficient captures the **acceleration effect**: investment responds not just to the level of income, but to its rate of change.

This negative coefficient is also what makes complex roots possible in the characteristic equation.

Without it, Chow proves that demand systems with only positive coefficients have real positive roots, and hence no oscillatory dynamics.

The *Samuelson Multiplier-Accelerator* lecture explores this mechanism in detail through the Hansen-Samuelson multiplier-accelerator model.

Here we briefly illustrate the effect.

Take the multiplier–accelerator law of motion:

$$Y_t = cY_{t-1} + v(Y_{t-1} - Y_{t-2}),$$

and rewrite it as a first-order system in (Y_t, Y_{t-1}).

```
def samuelson_transition(c, v):
    return np.array([[c + v, -v], [1.0, 0.0]])

# Compare weak vs strong acceleration
# Weak: c=0.8, v=0.1 gives real roots (discriminant > 0)
# Strong: c=0.6, v=0.8 gives complex roots (discriminant < 0)
cases = [("weak acceleration", 0.8, 0.1),
         ("strong acceleration", 0.6, 0.8)]
A_list = [samuelson_transition(c, v) for _, c, v in cases]

for (label, c, v), A in zip(cases, A_list):
    eig = np.linalg.eigvals(A)
    disc = (c + v)**2 - 4*v
    print(
        f"{label}: c={c}, v={v}, discriminant={disc:.2f}, eigenvalues={eig}")
```

```
weak acceleration: c=0.8, v=0.1, discriminant=0.41, eigenvalues=[0.77015621 0.
↪12984379]
strong acceleration: c=0.6, v=0.8, discriminant=-1.24, eigenvalues=[0.7+0.
↪55677644j 0.7-0.55677644j]
```

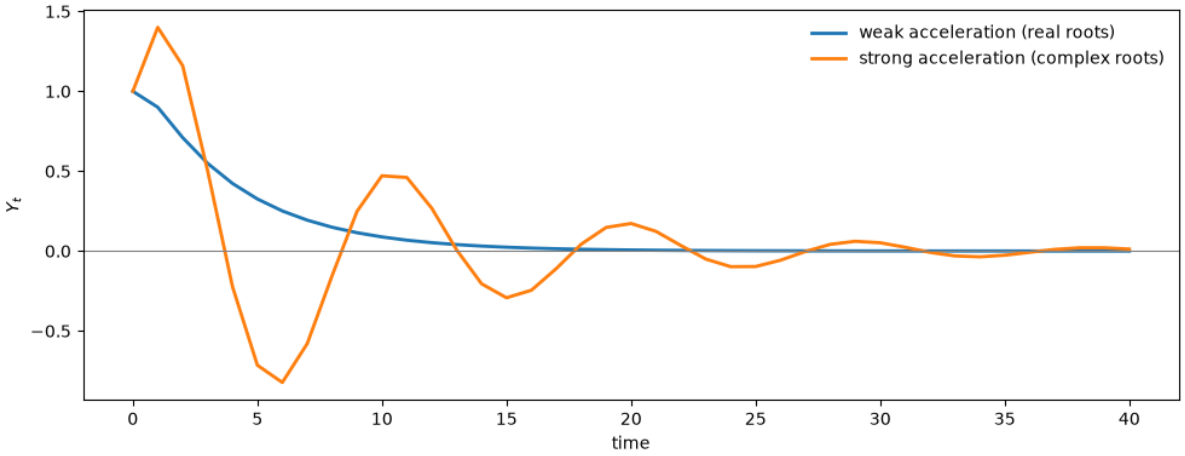
With weak acceleration ($v = 0.1$), the discriminant is positive and the roots are real.

With strong acceleration ($v = 0.8$), the discriminant is negative and the roots are complex conjugates that enable oscillatory dynamics.

Now let's see how these different eigenvalue structures affect the impulse responses to a one-time shock in Y

```
T = 40
s0 = np.array([1.0, 0.0])
irfs = []
for A in A_list:
    s = s0.copy()
    path = np.empty(T + 1)
    for t in range(T + 1):
        path[t] = s[0]
        s = A @ s
    irfs.append(path)

fig, ax = plt.subplots(figsize=(10, 4))
ax.plot(range(T + 1), irfs[0], lw=2,
                label="weak acceleration (real roots)")
ax.plot(range(T + 1), irfs[1], lw=2,
                label="strong acceleration (complex roots)")
ax.axhline(0.0, lw=0.8, color='gray')
ax.set_xlabel("time")
ax.set_ylabel(r"$Y_t$")
ax.legend(frameon=False)
plt.tight_layout()
plt.show()
```



With weak acceleration, the impulse response decays monotonically.

With strong acceleration, it oscillates.

We can ask how the eigenvalues change as we increase the accelerator v.

As we increase the accelerator v, the eigenvalues move further from the origin.

For this model, the eigenvalue modulus is $|\lambda| = \sqrt{v}$, so the stability boundary is $v = 1$.

```
v_grid = [0.2, 0.4, 0.6, 0.8, 0.95]
c = 0.6
T_irf = 40  # periods for impulse response

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

for v in v_grid:
    A = samuelson_transition(c, v)
```

(continues on next page)

(continued from previous page)

```
    eig = np.linalg.eigvals(A)

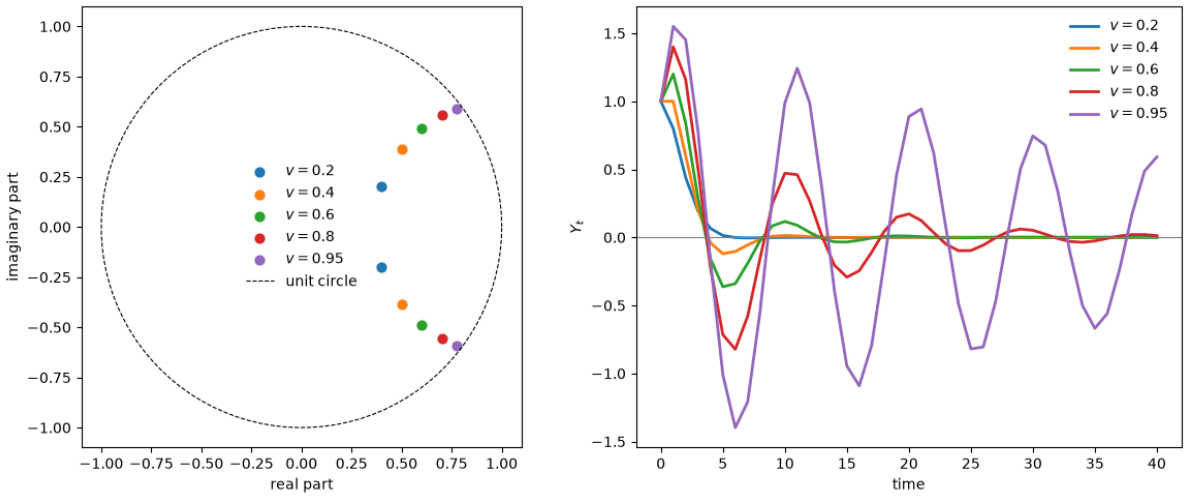
    # Eigenvalues (left panel)
    axes[0].scatter(eig.real, eig.imag, s=40, label=f'$v={v}$')

    # Impulse response (right panel)
    s = np.array([1.0, 0.0])
    irf = np.empty(T_irf + 1)
    for t in range(T_irf + 1):
        irf[t] = s[0]
        s = A @ s
    axes[1].plot(range(T_irf + 1), irf, lw=2, label=f'$v={v}$')

# Visualize the eigenvalue locations and the unit circle
θ_circle = np.linspace(0, 2*np.pi, 100)
axes[0].plot(np.cos(θ_circle), np.sin(θ_circle),
             'k--', lw=0.8, label='unit circle')
axes[0].set_xlabel('real part')
axes[0].set_ylabel('imaginary part')
axes[0].set_aspect('equal')
axes[0].legend(frameon=False)

# impulse response panel
axes[1].axhline(0, lw=0.8, color='gray')
axes[1].set_xlabel('time')
axes[1].set_ylabel(r'$Y_t$')
axes[1].legend(frameon=False)

plt.tight_layout()
plt.show()
```



As v increases, eigenvalues approach the unit circle and oscillations become more persistent.

This illustrates that acceleration creates complex eigenvalues, which are necessary for oscillatory dynamics in deterministic systems.

But what happens when we add random shocks?

An insight of Ragnar Frisch [Frisch, 1933] was that damped oscillations can be "maintained" when the system is continuously perturbed by random disturbances.

To study this formally, we need to introduce the stochastic framework.

45.4 A linear system with shocks

We analyze a first-order linear stochastic system

$$y_t = Ay_{t-1} + u_t, \qquad \mathbb{E}[u_t] = 0, \qquad \mathbb{E}[u_t u_t^\top] = V, \qquad \mathbb{E}[u_t u_{t-k}^\top] = 0, \quad k \neq 0. \tag{45.4}$$

When the eigenvalues of A are strictly inside the unit circle, the process is covariance stationary and its autocovariances exist.

In the notation of *Linear State Space Models*, this is the same stability condition that guarantees a unique solution to a discrete Lyapunov equation.

Define the lag-k autocovariance matrices

$$\Gamma_k := \mathbb{E}[y_t y_{t-k}^\top]. \tag{45.5}$$

Standard calculations (also derived in [Chow, 1968]) give the recursion

$$\Gamma_k = A\Gamma_{k-1}, \quad k \geq 1, \qquad \text{and} \qquad \Gamma_0 = A\Gamma_0 A^\top + V. \tag{45.6}$$

The second equation is the discrete Lyapunov equation for Γ_0.

Chow [1968] motivates the stochastic analysis with a quote from Ragnar Frisch:

> The examples we have discussed … show that when a [deterministic] economic system gives rise to oscillations, these will most frequently be damped. But in reality the cycles … are generally not damped. How can the maintenance of the swings be explained? … One way which I believe is particularly fruitful and promising is to study what would become of the solution of a determinate dynamic system if it were exposed to a stream of erratic shocks … Thus, by connecting the two ideas: (1) the continuous solution of a determinate dynamic system and (2) the discontinuous shocks intervening and supplying the energy that may maintain the swings–we get a theoretical setup which seems to furnish a rational interpretation of those movements which we have been accustomed to see in our statistical time data.
>
> – Ragnar Frisch (1933) [Frisch, 1933]

Chow's main insight is that oscillations in the deterministic system are *neither necessary nor sufficient* for producing "cycles" in the stochastic system.

We have to bring the stochastic element into the picture.

We will show that even when eigenvalues are real (no deterministic oscillations), the stochastic system can exhibit cyclical patterns in its autocovariances and spectral densities.

45.4.1 Autocovariances in terms of eigenvalues

Let $\lambda_1, \ldots, \lambda_p$ be the distinct, possibly complex, eigenvalues of A, and let B be the matrix whose columns are the corresponding right eigenvectors:

$$AB = BD_\lambda, \quad \text{or equivalently} \quad A = BD_\lambda B^{-1} \tag{45.7}$$

where $D_\lambda = \text{diag}(\lambda_1, \ldots, \lambda_p)$.

Define canonical variables $z_t = B^{-1} y_t$.

These satisfy the decoupled dynamics:

$$z_t = D_\lambda z_{t-1} + \varepsilon_t \tag{45.8}$$

where $\varepsilon_t = B^{-1}u_t$ has covariance matrix $W = B^{-1}V(B^{-1})^\top$.

The autocovariance matrix of the canonical variables, denoted Γ_k^*, satisfies

$$\Gamma_k^* = D_\lambda^k \Gamma_0^*, \quad k = 1, 2, 3, \ldots \tag{45.9}$$

and

$$\Gamma_0^* = \left(\frac{w_{ij}}{1 - \lambda_i \lambda_j} \right) \tag{45.10}$$

where w_{ij} are elements of W.

The autocovariance matrices of the original variables are then

$$\Gamma_k = B\Gamma_k^* B^\top = B D_\lambda^k \Gamma_0^* B^\top, \quad k = 0, 1, 2, \ldots \tag{45.11}$$

The scalar autocovariance $\gamma_{ij,k} = \mathbb{E}[y_{it} y_{j,t-k}]$ is a *linear combination* of powers of the eigenvalues:

$$\gamma_{ij,k} = \sum_m \sum_n b_{im} b_{jn} \gamma^*_{mn,0} \lambda_m^k = \sum_m d_{ij,m} \lambda_m^k \tag{45.12}$$

Compare this to the deterministic time path from initial condition y_0:

$$y_{it} = \sum_j b_{ij} z_{j0} \lambda_j^t \tag{45.13}$$

Both the autocovariance function (45.12) and the deterministic path (45.13) are linear combinations of λ_m^k (or λ_j^t).

45.4.2 Complex roots and damped oscillations

When eigenvalues come in complex conjugate pairs $\lambda = re^{\pm i\theta}$ with $r < 1$, their contribution to the autocovariance function is a **damped cosine**:

$$2sr^k \cos(\theta k + \phi) \tag{45.14}$$

for appropriate amplitude s and phase ϕ determined by the eigenvector loadings.

In the deterministic model, such complex roots generate damped oscillatory time paths.

In the stochastic model, they generate damped oscillatory autocovariance functions.

It is in this sense that deterministic oscillations could be "maintained" in the stochastic model, but as we will see, the connection between eigenvalues and spectral peaks is more subtle than this suggests.

45.5 From autocovariances to spectra

Chow's key step is to translate the autocovariance sequence $\{\Gamma_k\}$ into a frequency-domain object.

The **spectral density matrix** is the Fourier transform of Γ_k:

$$F(\omega) := \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \Gamma_k e^{-i\omega k}, \qquad \omega \in [0, \pi]. \tag{45.15}$$

For the VAR(1) system (45.4), this sum has a closed form

$$F(\omega) = \frac{1}{2\pi} \left(I - Ae^{-i\omega}\right)^{-1} V \left(I - A^\top e^{i\omega}\right)^{-1}. \tag{45.16}$$

$F(\omega)$ tells us how much variation in y_t is associated with cycles of (angular) frequency ω.

Higher frequencies correspond to rapid oscillations, meaning short cycles where the series completes many up-and-down movements per unit of time.

Lower frequencies correspond to slower oscillations, meaning long cycles that unfold over extended periods.

The corresponding cycle length (or period) is

$$T(\omega) = \frac{2\pi}{\omega}. \tag{45.17}$$

Thus, a frequency of $\omega = \pi$ corresponds to the shortest possible cycle of $T = 2$ periods, while frequencies near zero correspond to very long cycles.

When the spectral density $F(\omega)$ is concentrated at particular frequencies, it indicates that the time series exhibits pronounced cyclical behavior at those frequencies.

The advanced lecture Estimation of Spectra explains how to estimate $F(\omega)$ from data.

Here we focus on the model-implied spectrum.

We saw earlier that acceleration creates complex eigenvalues, which enable oscillatory impulse responses.

But do complex roots guarantee a spectral peak?

Are they necessary for one?

Chow provides precise answers for the Hansen-Samuelson model.

45.6 Spectral peaks in the Hansen-Samuelson model

Chow [1968] provides a detailed spectral analysis of the Hansen-Samuelson multiplier-accelerator model, deriving exact conditions for when spectral peaks occur.

The analysis reveals that in this specific model, complex roots are *necessary* for a peak, but as we will see later, this is not true in general.

45.6.1 The model as a first-order system

The second-order Hansen-Samuelson equation can be written as a first-order system:

$$\begin{bmatrix} y_{1t} \\ y_{2t} \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ 1 & 0 \end{bmatrix} \begin{bmatrix} y_{1,t-1} \\ y_{2,t-1} \end{bmatrix} + \begin{bmatrix} u_{1t} \\ 0 \end{bmatrix} \tag{45.18}$$

where $y_{2t} = y_{1,t-1}$ is simply the lagged value of y_{1t}.

This structure implies a special relationship among the autocovariances:

$$\gamma_{11,k} = \gamma_{22,k} = \gamma_{12,k-1} = \gamma_{21,k+1} \tag{45.19}$$

Using the autocovariance recursion, Chow shows that this leads to the condition

$$\gamma_{11,-1} = d_{11,1}\lambda_1^{-1} + d_{11,2}\lambda_2^{-1} = \gamma_{11,1} = d_{11,1}\lambda_1 + d_{11,2}\lambda_2 \tag{45.20}$$

which constrains the spectral density in a useful way.

45.6.2 The spectral density formula

From equations (45.12) and the scalar kernel $g_i(\omega) = (1-\lambda_i^2)/(1+\lambda_i^2-2\lambda_i\cos\omega)$, the spectral density of y_{1t} is:

$$f_{11}(\omega) = d_{11,1}g_1(\omega) + d_{11,2}g_2(\omega) \tag{45.21}$$

which can be written in the combined form:

$$f_{11}(\omega) = \frac{d_{11,1}(1-\lambda_1^2)(1+\lambda_2^2) + d_{11,2}(1-\lambda_2^2)(1+\lambda_1^2) - 2[d_{11,1}(1-\lambda_1^2)\lambda_2 + d_{11,2}(1-\lambda_2^2)\lambda_1]\cos\omega}{(1+\lambda_1^2-2\lambda_1\cos\omega)(1+\lambda_2^2-2\lambda_2\cos\omega)} \tag{45.22}$$

A key observation: due to condition (45.20), the *numerator is not a function of* $\cos\omega$.

Therefore, to find a maximum of $f_{11}(\omega)$, we need only find a minimum of the denominator.

45.6.3 Conditions for a spectral peak

The first derivative of the denominator with respect to ω is:

$$2[(1+\lambda_1^2)\lambda_2 + (1+\lambda_2^2)\lambda_1]\sin\omega - 8\lambda_1\lambda_2\cos\omega\sin\omega \tag{45.23}$$

For $0 < \omega < \pi$, we have $\sin\omega > 0$, so the derivative equals zero if and only if:

$$(1+\lambda_1^2)\lambda_2 + (1+\lambda_2^2)\lambda_1 = 4\lambda_1\lambda_2\cos\omega \tag{45.24}$$

For *complex conjugate roots* $\lambda_1 = re^{i\theta}$, $\lambda_2 = re^{-i\theta}$, substitution into (45.24) gives:

$$\cos\omega = \frac{1+r^2}{2r}\cos\theta \tag{45.25}$$

The second derivative confirms this is a maximum when $\omega < \frac{3\pi}{4}$.

The necessary condition for a valid solution is:

$$-1 < \frac{1+r^2}{2r}\cos\theta < 1 \tag{45.26}$$

We can interpret it as:

- When $r \approx 1$, the factor $(1+r^2)/2r \approx 1$, so $\omega \approx \theta$
- When r is small (e.g., 0.3 or 0.4), condition (45.26) can only be satisfied if $\cos\theta \approx 0$, meaning $\theta \approx \pi/2$ (cycles of approximately 4 periods)

If $\theta = 54^\circ$ (corresponding to cycles of 6.67 periods) and $r = 0.4$, then $(1+r^2)/2r = 1.45$, giving $\cos\omega = 1.45 \times 0.588 = 0.85$, or $\omega = 31.5^\circ$, corresponding to cycles of 11.4 periods, which is much longer than the deterministic cycle.

```
def peak_condition_factor(r):
    """Compute (1 + r^2) / (2r)"""
    return (1 + r**2) / (2 * r)

θ_deg = 54
θ = np.deg2rad(θ_deg)
r_grid = np.linspace(0.3, 0.99, 100)

# For each r, compute the implied peak frequency
ω_peak = []
for r in r_grid:
```

(continues on next page)

(continued from previous page)

```
    factor = peak_condition_factor(r)
    cos_ω = factor * np.cos(θ)
    if -1 < cos_ω < 1:
        ω_peak.append(np.arccos(cos_ω))
    else:
        ω_peak.append(np.nan)

ω_peak = np.array(ω_peak)
period_peak = 2 * np.pi / ω_peak

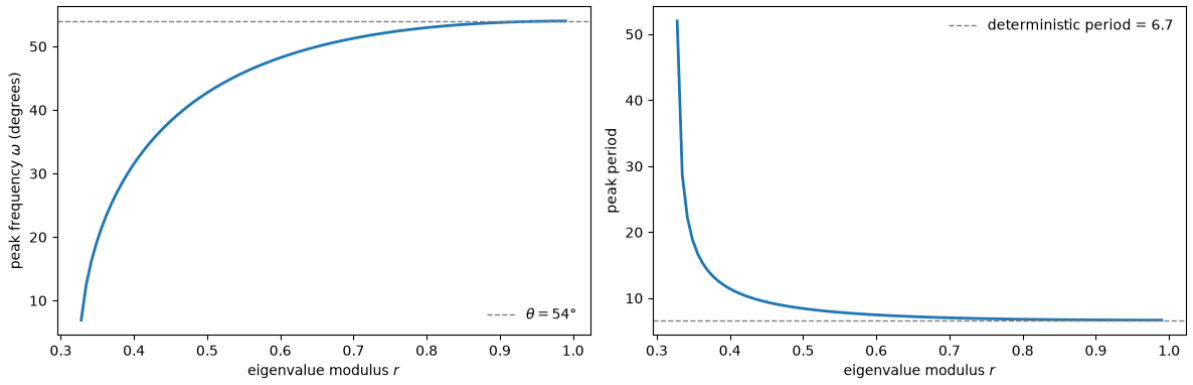
fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].plot(r_grid, np.rad2deg(ω_peak), lw=2)
axes[0].axhline(θ_deg, ls='--', lw=1.0, color='gray',
        label=rf'$\theta = {θ_deg}°$')
axes[0].set_xlabel('eigenvalue modulus $r$')
axes[0].set_ylabel(r'peak frequency $\omega$ (degrees)')
axes[0].legend(frameon=False)

axes[1].plot(r_grid, period_peak, lw=2)
axes[1].axhline(360/θ_deg, ls='--', lw=1.0, color='gray',
        label=rf'deterministic period = {360/θ_deg:.1f}')
axes[1].set_xlabel('eigenvalue modulus $r$')
axes[1].set_ylabel('peak period')
axes[1].legend(frameon=False)

plt.tight_layout()
plt.show()

r_example = 0.4
factor = peak_condition_factor(r_example)
cos_ω = factor * np.cos(θ)
ω_example = np.arccos(cos_ω)
print(f"Chow's example: r = {r_example}, θ = {θ_deg}°")
print(f"  cos(ω) = {cos_ω:.3f}")
print(f"  ω = {np.rad2deg(ω_example):.1f}°")
print(f"  Peak period = {360/np.rad2deg(ω_example):.1f}")
```




```
Chow's example: r = 0.4, θ = 54°
  cos(ω) = 0.852
  ω = 31.5°
  Peak period = 11.4
```

As $r \to 1$, the peak frequency converges to θ.

For smaller r, the peak frequency can differ substantially from the deterministic oscillation frequency.

45.6.4 Real positive roots cannot produce peaks

For *real and positive roots* $\lambda_1, \lambda_2 > 0$, the first-order condition (45.24) cannot be satisfied.

To see why, recall that a spectral peak at an interior frequency $\omega \in (0, \pi)$ requires

$$\cos\omega = \frac{(1+\lambda_1^2)\lambda_2 + (1+\lambda_2^2)\lambda_1}{4\lambda_1\lambda_2}.$$

For this to have a solution, we need the right-hand side to lie in $[-1, 1]$.

But for positive λ_1, λ_2, the numerator exceeds $4\lambda_1\lambda_2$:

$$(1+\lambda_1^2)\lambda_2 + (1+\lambda_2^2)\lambda_1 - 4\lambda_1\lambda_2 = \lambda_1(1-\lambda_2)^2 + \lambda_2(1-\lambda_1)^2. \tag{45.27}$$

The right-hand side is a sum of two non-negative terms (each is a positive number times a square).

It equals zero only if both $\lambda_1 = 1$ and $\lambda_2 = 1$, which violates the stability condition $|\lambda_i| < 1$.

For any stable system with real positive roots, this expression is strictly positive, so

$$\cos\omega = \frac{(1+\lambda_1^2)\lambda_2 + (1+\lambda_2^2)\lambda_1}{4\lambda_1\lambda_2} > 1, \tag{45.28}$$

which is impossible.

This is a key result: in the Hansen-Samuelson model, *complex roots are necessary* for a spectral peak at interior frequencies.

The following figure illustrates the difference in spectra between a case with complex roots and a case with real roots

```
ω_grid = np.linspace(1e-3, np.pi - 1e-3, 800)
V_hs = np.array([[1.0, 0.0], [0.0, 0.0]])  # shock only in first equation

# Case 1: Complex roots (c=0.6, v=0.8)
c_complex, v_complex = 0.6, 0.8
A_complex = samuelson_transition(c_complex, v_complex)
eig_complex = np.linalg.eigvals(A_complex)

# Case 2: Real roots (c=0.8, v=0.1)
c_real, v_real = 0.8, 0.1
A_real = samuelson_transition(c_real, v_real)
eig_real = np.linalg.eigvals(A_real)

print(
    f"Complex case (c={c_complex}, v={v_complex}): eigenvalues = {eig_complex}")
print(
    f"Real case (c={c_real}, v={v_real}): eigenvalues = {eig_real}")

F_complex = spectral_density_var1(A_complex, V_hs, ω_grid)
F_real = spectral_density_var1(A_real, V_hs, ω_grid)

f11_complex = np.real(F_complex[:, 0, 0])
f11_real = np.real(F_real[:, 0, 0])

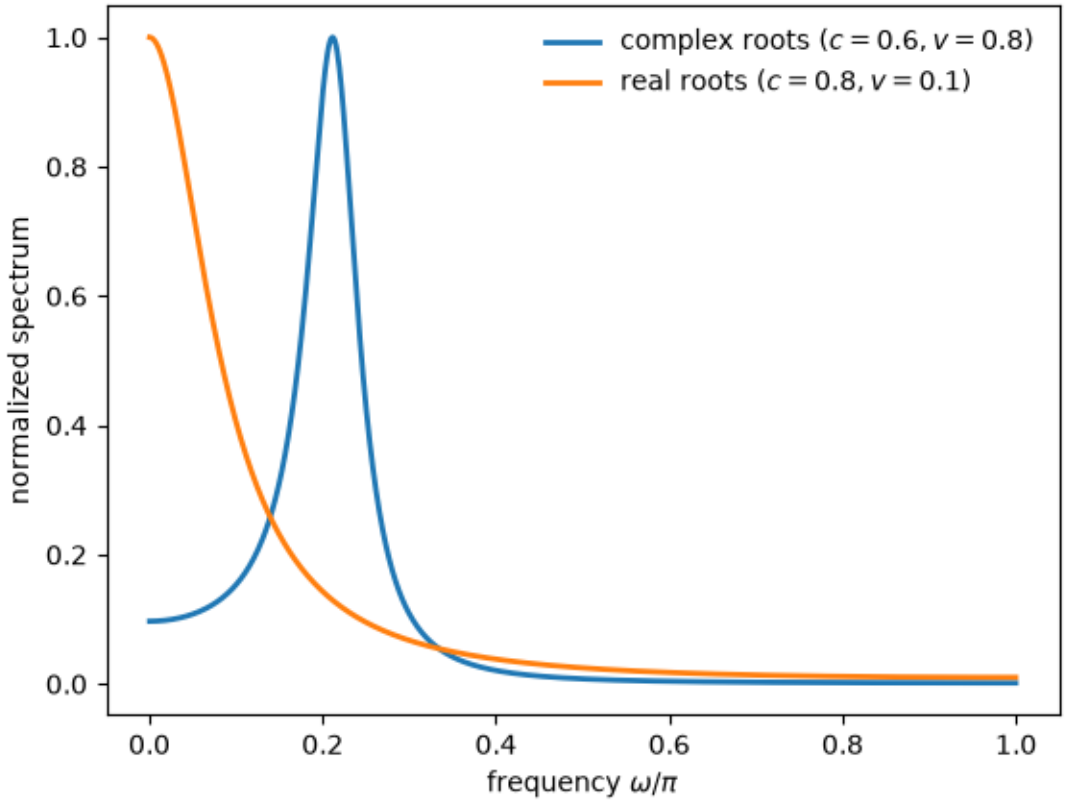
fig, ax = plt.subplots()
```

(continues on next page)

(continued from previous page)

```
ax.plot(ω_grid / np.pi, f11_complex / np.max(f11_complex), lw=2,
        label=fr'complex roots ($c={c_complex}, v={v_complex}$)')
ax.plot(ω_grid / np.pi, f11_real / np.max(f11_real), lw=2,
        label=fr'real roots ($c={c_real}, v={v_real}$)')
ax.set_xlabel(r'frequency $\omega/\pi$')
ax.set_ylabel('normalized spectrum')
ax.legend(frameon=False)
plt.show()
```

```
Complex case (c=0.6, v=0.8): eigenvalues = [0.7+0.55677644j 0.7-0.55677644j]
Real case (c=0.8, v=0.1): eigenvalues = [0.77015621 0.12984379]
```



With complex roots, the spectrum has a clear interior peak.

With real roots, the spectrum is monotonically decreasing and no interior peak is possible.

45.7 Real roots can produce peaks in general models

While real positive roots cannot produce spectral peaks in the Hansen-Samuelson model, Chow [1968] emphasizes that this is *not true in general*.

In multivariate systems, the spectral density of a linear combination of variables can have interior peaks even when all eigenvalues are real and positive.

45.7.1 Example

Chow constructs the following explicit example with two real positive eigenvalues:

$$\lambda_1 = 0.1, \quad \lambda_2 = 0.9 \tag{45.29}$$

$$w_{11} = w_{22} = 1, \quad w_{12} = 0.8 \tag{45.30}$$

$$b_{m1} = 1, \quad b_{m2} = -0.01 \tag{45.31}$$

The spectral density of the linear combination $x_t = b_m^\top y_t$ is:

$$f_{mm}(\omega) = \frac{0.9913}{1.01 - 0.2\cos\omega} - \frac{0.001570}{1.81 - 1.8\cos\omega} \tag{45.32}$$

Chow tabulates the values:

ω	0	$\pi/8$	$2\pi/8$	$3\pi/8$	$4\pi/8$	$5\pi/8$	$6\pi/8$	$7\pi/8$	π
$f_{mm}(\omega)$	1.067	1.183	1.191	1.138	1.061	0.981	0.912	0.860	0.829

The peak at ω slightly below $\pi/8$ (corresponding to periods around 11) is "quite pronounced."

In the following figure, we reproduce this table, but with Python, we can plot a finer grid to find the peak more accurately

```
λ1, λ2 = 0.1, 0.9
w11, w22, w12 = 1.0, 1.0, 0.8
bm1, bm2 = 1.0, -0.01

# Construct the system
A_chow_ex = np.diag([λ1, λ2])

# W is the canonical shock covariance; we need V = B W B^T
# For diagonal A with distinct eigenvalues, B = I, so V = W
V_chow_ex = np.array([[w11, w12], [w12, w22]])
b_chow_ex = np.array([bm1, bm2])

# Chow's formula
def chow_spectrum_formula(ω):
    term1 = 0.9913 / (1.01 - 0.2 * np.cos(ω))
    term2 = 0.001570 / (1.81 - 1.8 * np.cos(ω))
    return term1 - term2

# Compute via formula and via our general method
ω_table = np.array([0, np.pi/8, 2*np.pi/8, 3*np.pi/8, 4*np.pi/8,
                    5*np.pi/8, 6*np.pi/8, 7*np.pi/8, np.pi])
f_formula = np.array([chow_spectrum_formula(ω) for ω in ω_table])

# General method
ω_grid_fine = np.linspace(1e-4, np.pi, 1000)
F_chow_ex = spectral_density_var1(A_chow_ex, V_chow_ex, ω_grid_fine)
f_general = spectrum_of_linear_combination(F_chow_ex, b_chow_ex)

# Normalize to match Chow's table scale
scale = f_formula[0] / spectrum_of_linear_combination(
    spectral_density_var1(
        A_chow_ex, V_chow_ex, np.array([0.0])), b_chow_ex)[0]
```

(continues on next page)

(continued from previous page)

```
print("Chow's Table (equation 67):")
print("ω/π:        ", "  ".join([f"{ω/np.pi:.3f}" for ω in ω_table]))
print("f_mm(ω):    ", "  ".join([f"{f:.3f}" for f in f_formula]))

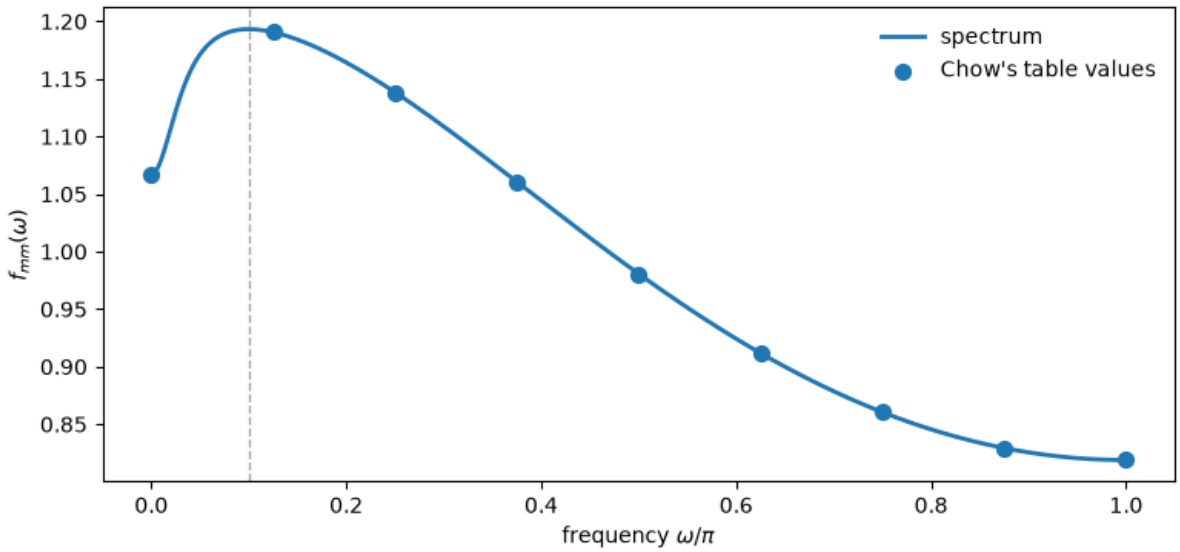
fig, ax = plt.subplots(figsize=(9, 4))
ax.plot(ω_grid_fine / np.pi, f_general * scale, lw=2,
            label='spectrum')
ax.scatter(ω_table / np.pi, f_formula, s=50, zorder=3,
            label="Chow's table values")

# Mark the peak
i_peak = np.argmax(f_general)
ω_peak = ω_grid_fine[i_peak]
ax.axvline(ω_peak / np.pi, ls='--', lw=1.0, color='gray', alpha=0.7)
ax.set_xlabel(r'frequency $\omega/\pi$')
ax.set_ylabel(r'$f_{mm}(\omega)$')
ax.legend(frameon=False)
plt.show()

print(f"\nPeak at ω/π ≈ {ω_peak/np.pi:.3f}, period ≈ {2*np.pi/ω_peak:.1f}")
```

```
Chow's Table (equation 67):
ω/π:         0.000  0.125  0.250  0.375  0.500  0.625  0.750  0.875  1.000
f_mm(ω):     1.067  1.191  1.138  1.061  0.981  0.912  0.860  0.829  0.819

Peak at ω/π ≈ 0.100, period ≈ 20.0
```



The peak appears at $\omega/\pi \approx 0.10$, which corresponds to a cycle length of approximately 20 periods, again much longer than the deterministic cycles implied by the eigenvalues.

45.7.2 The Slutsky connection

Chow connects this result to Slutsky's [Slutzky, 1937] finding that moving averages of a random series have recurrent cycles.

The VAR(1) model can be written as an infinite moving average:

$$y_t = u_t + Au_{t-1} + A^2 u_{t-2} + \cdots \tag{45.33}$$

This amounts to taking an infinite moving average of the random vectors u_t with "geometrically declining" weights $A^0, A^1, A^2, ...$

For a scalar process with $0 < \lambda < 1$, no distinct cycles can emerge.

But for a matrix A with real roots between 0 and 1, cycles *can* emerge in linear combinations of the variables.

As Chow puts it: "When neither of two (canonical) variables has distinct cycles… a linear combination can have a peak in its spectral density."

45.7.3 The general lesson

The examples above illustrate the following central points:

1. In the *Hansen-Samuelson model specifically*, complex roots are necessary for a spectral peak
2. But in *general multivariate systems*, complex roots are neither necessary nor sufficient
3. The full spectral shape depends on:
    - The eigenvalues of A
    - The shock covariance structure V
    - How the observable of interest loads on the eigenmodes (the vector b)

45.8 A calibrated model in the frequency domain

Chow and Levitan [1969] use the frequency-domain objects from Chow [1968] to study a calibrated annual macroeconometric model.

They work with five annual aggregates:

- $y_1 = C$ (consumption),
- $y_2 = I_1$ (equipment plus inventories),
- $y_3 = I_2$ (construction),
- $y_4 = R_a$ (long rate),
- $y_5 = Y_1 = C + I_1 + I_2$ (private-domestic GNP),

and add $y_6 = y_{1,t-1}$ to rewrite the original system in first-order form.

Throughout this section, frequency is measured in cycles per year, $f = \omega/2\pi \in [0, 1/2]$.

Following the paper, we normalize each spectrum to have area 1 over $[0, 1/2]$ so plots compare shape rather than scale.

Our goal is to reconstruct the transition matrix A and then compute and interpret the model-implied spectra, gains/coherences, and phase differences.

45.8.1 The cycle subsystem

The paper starts from a reduced form with exogenous inputs,

$$y_t = Ay_{t-1} + Cx_t + u_t. \tag{45.34}$$

To study cycles, they remove the deterministic component attributable to x_t and focus on the zero-mean subsystem

$$y_t = Ay_{t-1} + u_t. \tag{45.35}$$

For second moments, the only additional ingredient is the covariance matrix $V = \mathbb{E}[u_t u_t^\top]$.

Chow and Levitan compute it from structural parameters via

$$V = M^{-1}\Sigma(M^{-1})^\top \tag{45.36}$$

where Σ is the covariance of structural residuals and M is the matrix of contemporaneous structural coefficients.

Here we take A and V as given and ask what they imply for spectra and cross-spectra.

The 6×6 reduced-form shock covariance matrix V (scaled by 10^{-7}) reported by Chow and Levitan is:

$$V = \begin{bmatrix} 8.250 & 7.290 & 2.137 & 2.277 & 17.68 & 0 \\ 7.290 & 7.135 & 1.992 & 2.165 & 16.42 & 0 \\ 2.137 & 1.992 & 0.618 & 0.451 & 4.746 & 0 \\ 2.277 & 2.165 & 0.451 & 1.511 & 4.895 & 0 \\ 17.68 & 16.42 & 4.746 & 4.895 & 38.84 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{45.37}$$

The sixth row and column are zeros because y_6 is an identity (lagged y_1).

The transition matrix A has six characteristic roots:

$$\begin{aligned} &\lambda_1 = 0.9999725, \quad \lambda_2 = 0.9999064, \quad \lambda_3 = 0.4838, \\ &\lambda_4 = 0.0761 + 0.1125i, \quad \lambda_5 = 0.0761 - 0.1125i, \quad \lambda_6 = -0.00004142. \end{aligned} \tag{45.38}$$

Two roots are near unity because two structural equations are in first differences.

One root (λ_6) is theoretically zero because of the identity $y_5 = y_1 + y_2 + y_3$.

The complex conjugate pair $\lambda_{4,5}$ has modulus $|\lambda_4| = \sqrt{0.0761^2 + 0.1125^2} \approx 0.136$.

The right eigenvector matrix B (columns are eigenvectors corresponding to $\lambda_1, \ldots, \lambda_6$):

$$B = \begin{bmatrix} -0.008 & 1.143 & 0.320 & 0.283 + 0.581i & 0.283 - 0.581i & 0.000 \\ -0.000 & 0.013 & -0.586 & -2.151 + 0.742i & -2.151 - 0.742i & 2.241 \\ -0.001 & 0.078 & 0.889 & -0.215 + 0.135i & -0.215 - 0.135i & 0.270 \\ 1.024 & 0.271 & 0.069 & -0.231 + 0.163i & -0.231 - 0.163i & 0.307 \\ -0.009 & 1.235 & 0.623 & -2.082 + 1.468i & -2.082 - 1.468i & 2.766 \\ -0.008 & 1.143 & 0.662 & 4.772 + 0.714i & 4.772 - 0.714i & -4.399 \end{bmatrix}. \tag{45.39}$$

Together, V, $\{\lambda_i\}$, and B are sufficient to compute all spectral and cross-spectral densities.

45.8.2 Reconstructing A and computing $F(\omega)$

The paper reports (λ, B, V), which is enough to reconstruct $A = B\,\mathrm{diag}(\lambda_1, \ldots, \lambda_6)\,B^{-1}$ and then compute the model-implied spectral objects.

```
λ = np.array([
    0.9999725, 0.9999064, 0.4838,
    0.0761 + 0.1125j, 0.0761 - 0.1125j, -0.00004142
], dtype=complex)

B = np.array([
    [-0.008, 1.143, 0.320, 0.283+0.581j, 0.283-0.581j, 0.000],
    [-0.000, 0.013, -0.586, -2.151+0.742j, -2.151-0.742j, 2.241],
    [-0.001, 0.078, 0.889, -0.215+0.135j, -0.215-0.135j, 0.270],
    [1.024, 0.271, 0.069, -0.231+0.163j, -0.231-0.163j, 0.307],
    [-0.009, 1.235, 0.623, -2.082+1.468j, -2.082-1.468j, 2.766],
    [-0.008, 1.143, 0.662, 4.772+0.714j, 4.772-0.714j, -4.399]
], dtype=complex)

V = np.array([
    [8.250, 7.290, 2.137, 2.277, 17.68, 0],
    [7.290, 7.135, 1.992, 2.165, 16.42, 0],
    [2.137, 1.992, 0.618, 0.451, 4.746, 0],
    [2.277, 2.165, 0.451, 1.511, 4.895, 0],
    [17.68, 16.42, 4.746, 4.895, 38.84, 0],
    [0, 0, 0, 0, 0, 0]
]) * 1e-7

D_λ = np.diag(λ)
A_chow = B @ D_λ @ np.linalg.inv(B)
A_chow = np.real(A_chow)
print("eigenvalues of reconstructed A:")
print(np.linalg.eigvals(A_chow).round(6))
```

```
eigenvalues of reconstructed A:
[-4.10000e-05+0.j      7.61000e-02+0.1125j  7.61000e-02-0.1125j
  4.83800e-01+0.j      9.99906e-01+0.j      9.99973e-01+0.j    ]
```

45.8.3 Canonical coordinates

Chow and Levitan's canonical transformation uses $z_t = B^{-1} y_t$, giving dynamics $z_t = D_\lambda z_{t-1} + e_t$.

Accordingly, the canonical shock covariance is

$$W = B^{-1} V (B^{-1})^\top.$$

```
B_inv = np.linalg.inv(B)
W = B_inv @ V @ B_inv.T
print("diagonal of W:")
print(np.diag(W).round(10))
```

```
diagonal of W:
[ 8.560e-08-0.00e+00j  3.638e-07-0.00e+00j  1.300e-08-0.00e+00j
 -7.880e-08+6.26e-08j -7.880e-08-6.26e-08j -1.650e-08+0.00e+00j]
```

Chow and Levitan derive the following closed-form formula for the spectral density matrix:

$$F(\omega) = B \left[\frac{w_{ij}}{(1 - \lambda_i e^{-i\omega})(1 - \lambda_j e^{i\omega})} \right] B^\top, \tag{45.40}$$

where w_{ij} are elements of the canonical shock covariance W.

```
def spectral_density_chow(λ, B, W, ω_grid):
    """Spectral density via Chow's eigendecomposition formula."""
    p = len(λ)
    F = np.zeros((len(ω_grid), p, p), dtype=complex)
    for k, ω in enumerate(ω_grid):
        F_star = np.zeros((p, p), dtype=complex)
        for i in range(p):
            for j in range(p):
                denom = (1 - λ[i] * np.exp(-1j * ω)) \
                    * (1 - λ[j] * np.exp(1j * ω))
                F_star[i, j] = W[i, j] / denom
        F[k] = B @ F_star @ B.T
    return F / (2 * np.pi)

freq = np.linspace(1e-4, 0.5, 5000)      # cycles/year in [0, 1/2]
ω_grid = 2 * np.pi * freq                # radians in [0, π]
F_chow = spectral_density_chow(λ, B, W, ω_grid)
```

Let's plot the univariate spectra of consumption (y_1) and equipment plus inventories (y_2)

```
variable_names = ['$C$', '$I_1$', '$I_2$', '$R_a$', '$Y_1$']
freq_ticks = [1/18, 1/9, 1/6, 1/4, 1/3, 1/2]
freq_labels = [r'$\frac{1}{18}$', r'$\frac{1}{9}$', r'$\frac{1}{6}$',
               r'$\frac{1}{4}$', r'$\frac{1}{3}$', r'$\frac{1}{2}$']

def paper_frequency_axis(ax):
    ax.set_xlim([0.0, 0.5])
    ax.set_xticks(freq_ticks)
    ax.set_xticklabels(freq_labels)
    ax.set_xlabel(r'frequency $\omega/2\pi$')

# Normalized spectra (areas set to 1)
S = np.real(np.diagonal(F_chow, axis1=1, axis2=2))[:, :5]
df = np.diff(freq)
areas = np.sum(0.5 * (S[1:] + S[:-1]) * df[:, None], axis=0)
S_norm = S / areas
mask = freq >= 0.0

fig, axes = plt.subplots(1, 2, figsize=(10, 6))

# Figure I.1: consumption (log scale)
axes[0].plot(freq[mask], S_norm[mask, 0], lw=2)
axes[0].set_yscale('log')
paper_frequency_axis(axes[0])
axes[0].set_ylabel(r'normalized $f_{11}(\omega)$')

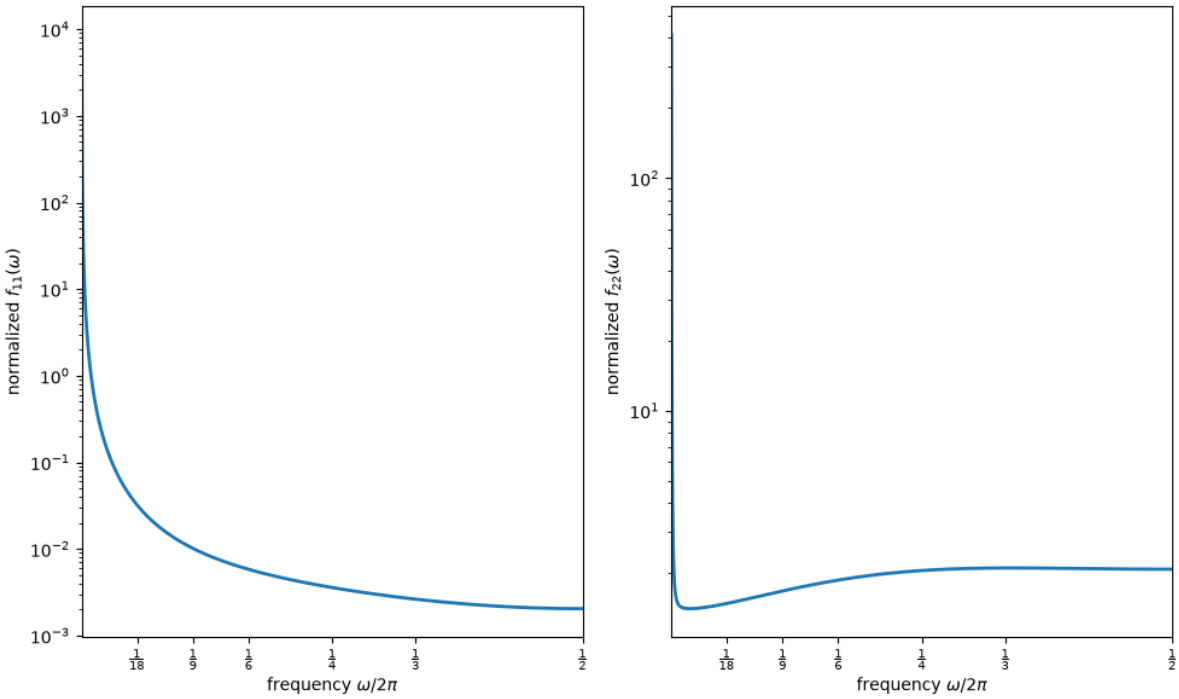
# Figure I.2: equipment + inventories (log scale)
axes[1].plot(freq[mask], S_norm[mask, 1], lw=2)
axes[1].set_yscale('log')
paper_frequency_axis(axes[1])
axes[1].set_ylabel(r'normalized $f_{22}(\omega)$')
```

(continues on next page)

(continued from previous page)

```
plt.tight_layout()
plt.show()

i_peak = np.argmax(S_norm[mask, 1])
f_peak = freq[mask][i_peak]
```



The left panel corresponds to consumption and declines monotonically with frequency.

It illustrates Granger's "typical spectral shape" for macroeconomic time series.

The right panel corresponds to equipment plus inventories and shows the clearest (but still very flat) interior-frequency bump.

Chow and Levitan associate the dominance of very low frequencies in both plots with strong persistence and long-run movements.

Very large low-frequency power can arise from eigenvalues extremely close to one, which occurs mechanically when some equations are written in first differences.

Local peaks are not automatic: complex roots may have small modulus, and multivariate interactions can generate peaks even when all roots are real.

The interior bump in the right panel corresponds to cycles of roughly three years, with the spectrum nearly flat over cycles between about two and four years.

(This discussion follows Section II of [Chow and Levitan, 1969].)

45.8.4 How variables move together across frequencies

Beyond univariate spectra, we can ask how pairs of variables covary at each frequency.

The **cross-spectrum** $f_{ij}(\omega) = c_{ij}(\omega) - i \cdot q_{ij}(\omega)$ decomposes into the cospectrum c_{ij} and the quadrature spectrum q_{ij}.

The **cross-amplitude** is $g_{ij}(\omega) = |f_{ij}(\omega)| = \sqrt{c_{ij}^2 + q_{ij}^2}$.

The **squared coherence** measures linear association at frequency ω:

$$R_{ij}^2(\omega) = \frac{|f_{ij}(\omega)|^2}{f_{ii}(\omega) f_{jj}(\omega)} \in [0, 1]. \tag{45.41}$$

Coherence measures how much of the variance of y_i at frequency ω can be "explained" by y_j at the same frequency.

High coherence means the two series move together tightly at that frequency.

The **gain** is the frequency-response coefficient when regressing y_i on y_j:

$$G_{ij}(\omega) = \frac{|f_{ij}(\omega)|}{f_{jj}(\omega)}. \tag{45.42}$$

It measures how much y_i responds to a unit change in y_j at frequency ω.

For instance, a gain of 0.9 at low frequencies means long-cycle movements in y_j translate almost one-for-one to y_i, and a gain of 0.3 at high frequencies means short-cycle movements are dampened.

The **phase** captures lead-lag relationships (in radians):

$$\Delta_{ij}(\omega) = \tan^{-1}\left(\frac{q_{ij}(\omega)}{c_{ij}(\omega)}\right). \tag{45.43}$$

```
def cross_spectral_measures(F, i, j):
    """Compute coherence, gain (y_i on y_j), and phase between variables i and j."""
    f_ij = F[:, i, j]
    f_ii, f_jj = np.real(F[:, i, i]), np.real(F[:, j, j])
    g_ij = np.abs(f_ij)
    coherence = (g_ij**2) / (f_ii * f_jj)
    gain = g_ij / f_jj
    phase = np.arctan2(-np.imag(f_ij), np.real(f_ij))
    return coherence, gain, phase
```

We now plot gain and coherence as in Figures II.1–II.4 of [Chow and Levitan, 1969].

```
gnp_idx = 4

fig, axes = plt.subplots(1, 2, figsize=(8, 6))

for idx, var_idx in enumerate([0, 1]):
    coherence, gain, phase = cross_spectral_measures(F_chow, var_idx, gnp_idx)
    ax = axes[idx]

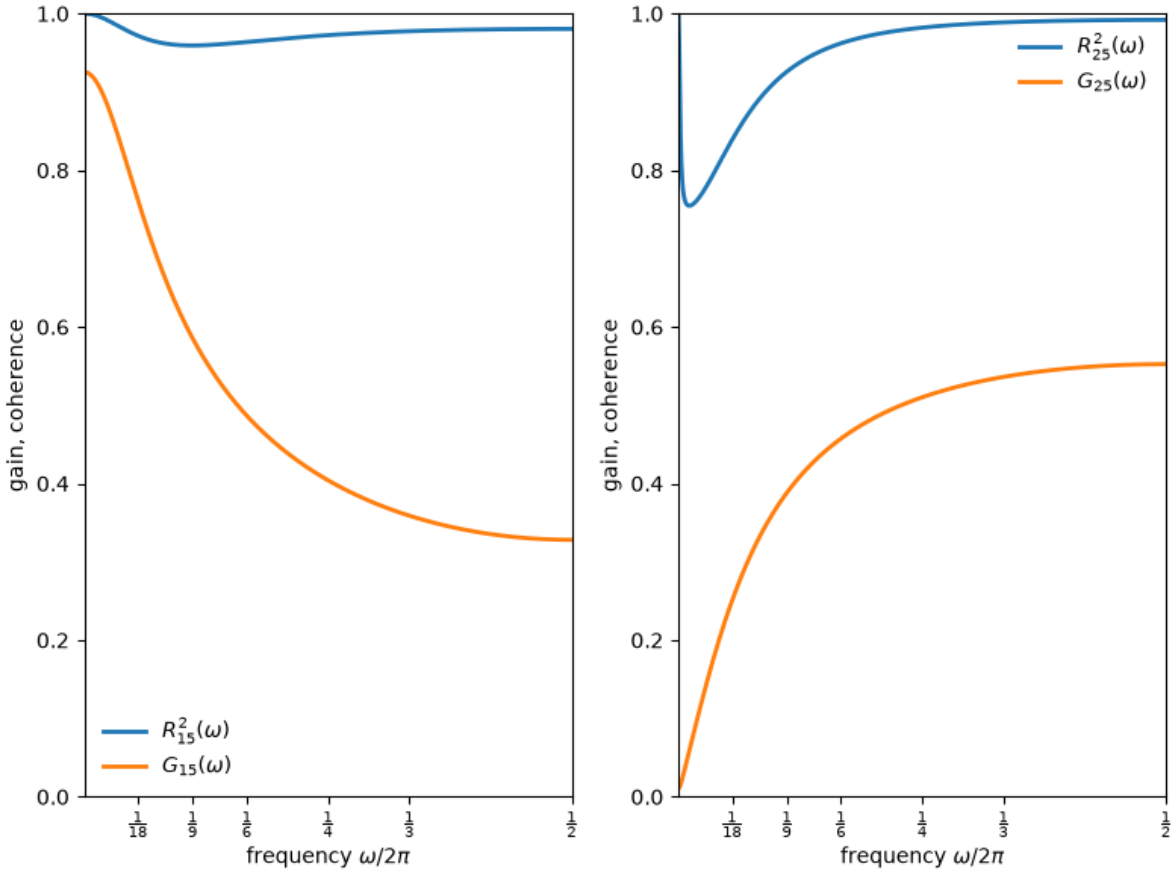
    ax.plot(freq[mask], coherence[mask],
            lw=2, label=rf'$R^2_{{{var_idx+1}5}}(\omega)$')
    ax.plot(freq[mask], gain[mask],
            lw=2, label=rf'$G_{{{var_idx+1}5}}(\omega)$')
    paper_frequency_axis(ax)
    ax.set_ylim([0, 1.0])
    ax.set_ylabel('gain, coherence')
```

(continues on next page)

(continued from previous page)

```
    ax.legend(frameon=False, loc='best')

plt.tight_layout()
plt.show()
```



The gain and coherence patterns differ across components (Figures II.1–II.2 of [Chow and Levitan, 1969]):

- Consumption vs private-domestic GNP (left panel):
  - Gain is about 0.9 at very low frequencies but falls below 0.4 for cycles shorter than four years.
  - This is evidence that short-cycle income movements translate less into consumption than long-cycle movements, consistent with permanent-income interpretations.
  - Coherence remains high throughout.
- Equipment plus inventories vs private-domestic GNP (right panel):
  - Gain *rises* with frequency, exceeding 0.5 for short cycles.
  - This is the frequency-domain signature of acceleration and volatile short-run inventory movements.

```
fig, axes = plt.subplots(1, 2, figsize=(8, 6))

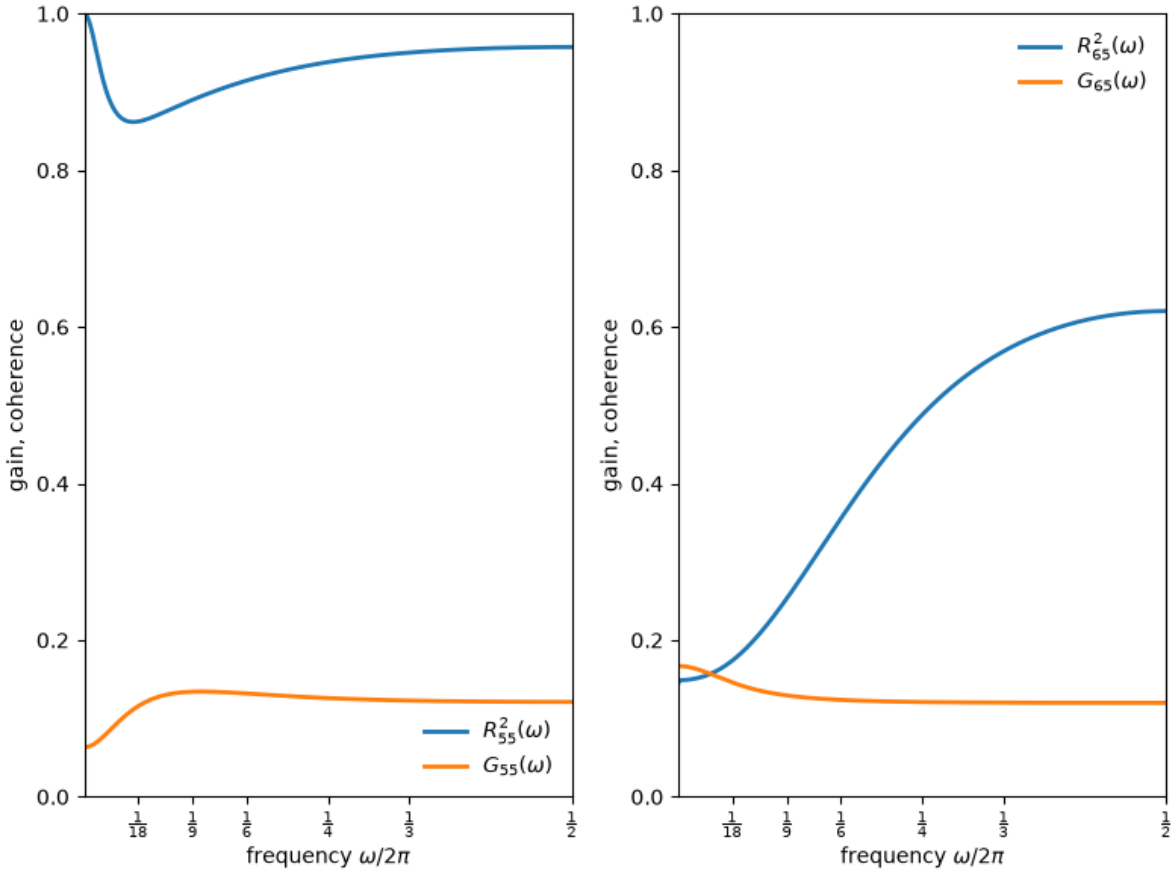
for idx, var_idx in enumerate([2, 3]):
    coherence, gain, phase = cross_spectral_measures(F_chow, var_idx, gnp_idx)
    ax = axes[idx]
```

(continues on next page)

(continued from previous page)

```
    ax.plot(freq[mask], coherence[mask],
            lw=2, label=rf'$R^2_{{{var_idx+3}5}}(\omega)$')
    ax.plot(freq[mask], gain[mask],
            lw=2, label=rf'$G_{{{var_idx+3}5}}(\omega)$')
    paper_frequency_axis(ax)
    ax.set_ylim([0, 1.0])
    ax.set_ylabel('gain, coherence')
    ax.legend(frameon=False, loc='best')

plt.tight_layout()
plt.show()
```



- New construction vs private-domestic GNP (left panel):
  - Gain peaks at medium cycle lengths (around 0.1 for short cycles).
  - Coherence for both investment series stays fairly high across frequencies.
- Long-bond yield vs private-domestic GNP (right panel):
  - Gain varies less across frequencies than real activity series.
  - Coherence with output is comparatively low at business-cycle frequencies, making it hard to explain interest-rate movements by inverting a money-demand equation.

45.8.5 Lead-lag relationships

The phase tells us which variable leads at each frequency.

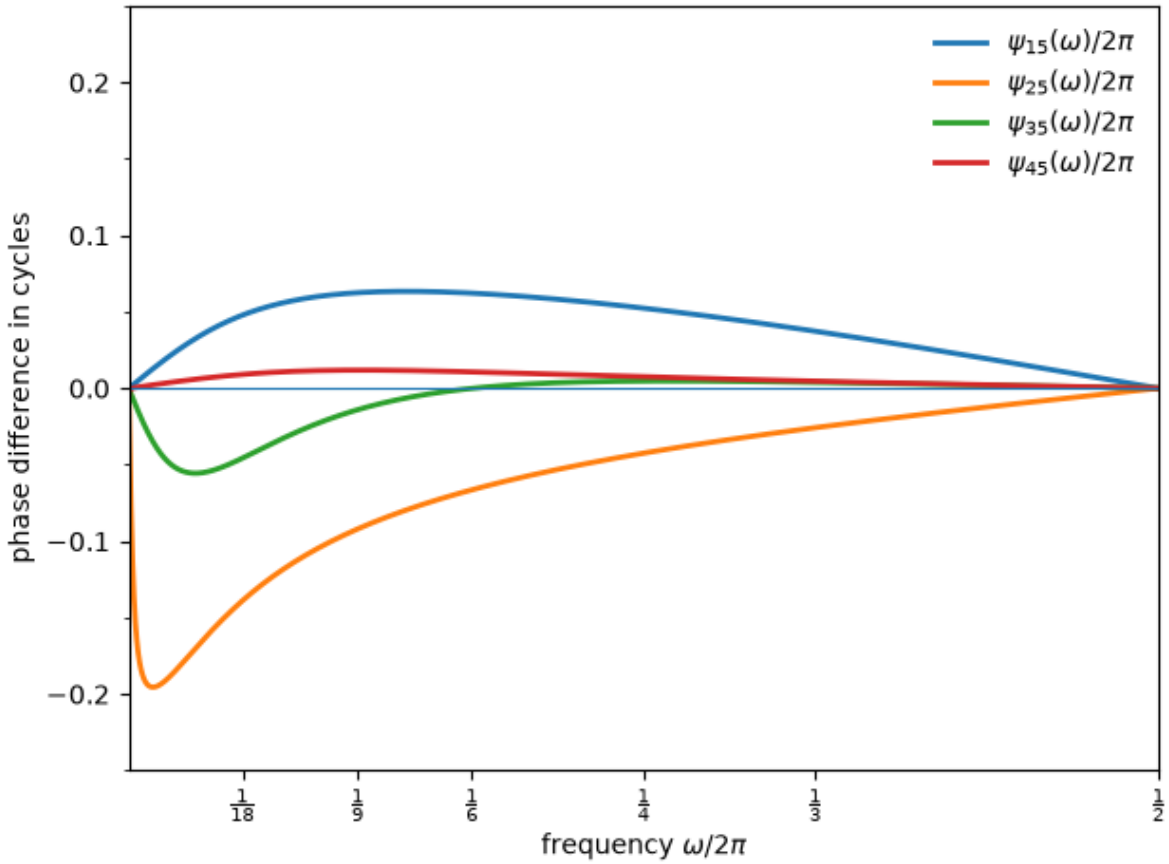
Positive phase means output leads the component; negative phase means the component leads output.

```
fig, ax = plt.subplots()

labels = [r'$\psi_{15}(\omega)/2\pi$', r'$\psi_{25}(\omega)/2\pi$',
          r'$\psi_{35}(\omega)/2\pi$', r'$\psi_{45}(\omega)/2\pi$']

for var_idx in range(4):
    coherence, gain, phase = cross_spectral_measures(F_chow, var_idx, gnp_idx)
    phase_cycles = phase / (2 * np.pi)
    ax.plot(freq[mask], phase_cycles[mask], lw=2, label=labels[var_idx])

ax.axhline(0, lw=0.8)
paper_frequency_axis(ax)
ax.set_ylabel('phase difference in cycles')
ax.set_ylim([-0.25, 0.25])
ax.set_yticks(np.arange(-0.25, 0.3, 0.05), minor=True)
ax.legend(frameon=False)
plt.tight_layout()
plt.show()
```



The phase relationships reveal that:

- Output leads consumption by a small fraction of a cycle (about 0.06 cycles at a 6-year period, 0.04 cycles at a 3-year

period).

- Equipment plus inventories tends to lead output (by about 0.07 cycles at a 6-year period, 0.03 cycles at a 3-year period).
- New construction leads at low frequencies and is close to coincident at higher frequencies.
- The bond yield lags output slightly, remaining close to coincident in timing.

These implied leads and lags are broadly consistent with turning-point timing summaries reported elsewhere, and simulations of the same model deliver similar lead–lag ordering at turning points (Figure III of [Chow and Levitan, 1969]).

45.8.6 Building blocks of spectral shape

Each eigenvalue contributes a characteristic spectral shape through the *scalar kernel*

$$g_i(\omega) = \frac{1 - |\lambda_i|^2}{|1 - \lambda_i e^{-i\omega}|^2} = \frac{1 - |\lambda_i|^2}{1 + |\lambda_i|^2 - 2\mathrm{Re}(\lambda_i)\cos\omega + 2\mathrm{Im}(\lambda_i)\sin\omega}. \tag{45.44}$$

For real λ_i, this simplifies to

$$g_i(\omega) = \frac{1 - \lambda_i^2}{1 + \lambda_i^2 - 2\lambda_i \cos\omega}.$$

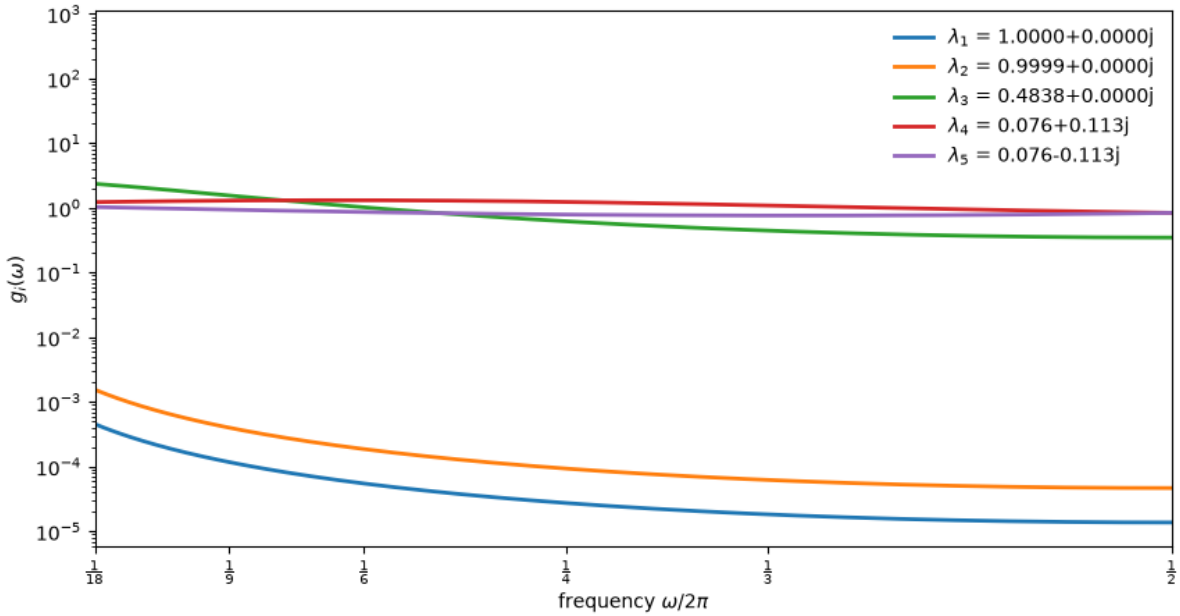
Each observable spectral density is a linear combination of these kernels (plus cross-terms).

Below, we plot the scalar kernels for each eigenvalue to see how they shape the overall spectra

```
def scalar_kernel(λ_i, ω_grid):
    """scalar spectral kernel g_i(ω)."""
    λ_i = complex(λ_i)
    mod_sq = np.abs(λ_i)**2
    return np.array(
        [(1 - mod_sq) / np.abs(1 - λ_i * np.exp(-1j * ω))**2
        for ω in ω_grid])

fig, ax = plt.subplots(figsize=(10, 5))
for i, λ_i in enumerate(λ):
    if np.abs(λ_i) > 0.01:
        g_i = scalar_kernel(λ_i, ω_grid)
        label = f'$\\lambda_{i+1}$ = {λ_i:.4f}' \
        if np.isreal(λ_i) else f'$\\lambda_{i+1}$ = {λ_i:.3f}'
        ax.semilogy(freq, g_i, label=label, lw=2)

ax.set_xlabel(r'frequency $\omega/2\pi$')
ax.set_ylabel('$g_i(\\omega)$')
ax.set_xlim([1/18, 0.5])
ax.set_xticks(freq_ticks)
ax.set_xticklabels(freq_labels)
ax.legend(frameon=False)
plt.show()
```



The figure reveals how eigenvalue magnitude shapes spectral contributions:

- *Near-unit eigenvalues* ($\lambda_1, \lambda_2 \approx 1$) produce kernels sharply peaked at low frequencies as these drive the strong low-frequency power seen in the spectra above.
- *The moderate eigenvalue* ($\lambda_3 \approx 0.48$) contributes a flatter component that spreads power more evenly across frequencies.
- *The complex pair* ($\lambda_{4,5}$) has such small modulus ($|\lambda_{4,5}| \approx 0.136$) that its kernel is nearly flat, which is too weak to generate a pronounced interior peak.

This decomposition explains why the spectra look the way they do: the near-unit eigenvalues dominate, concentrating variance at very low frequencies.

The complex pair, despite enabling oscillatory dynamics in principle, has insufficient modulus to produce a visible spectral peak.

45.9 Summary

Chow [1968] draws several conclusions that remain relevant for understanding business cycles.

The acceleration principle receives strong empirical support: the negative coefficient on lagged output in investment equations is a robust finding across datasets.

The relationship between eigenvalues and spectral peaks is more subtle than it first appears:

- Complex roots guarantee oscillatory autocovariances, but they are neither necessary nor sufficient for a pronounced spectral peak.
- In the Hansen–Samuelson model specifically, complex roots *are* necessary for a peak.
- But in general multivariate systems, even real roots can produce peaks through the interaction of shocks and eigenvector loadings.

Chow and Levitan [1969] demonstrate what these objects look like in a calibrated system: strong low-frequency power from near-unit eigenvalues, frequency-dependent gains and coherences, and lead–lag relations that vary with cycle length.

Their results are consistent with Granger's "typical spectral shape" for economic time series.

That is a monotonically decreasing function of frequency, driven by the near-unit eigenvalues that arise when some equations are specified in first differences.

Understanding whether this shape reflects the true data-generating process requires analyzing the spectral densities implied by structural econometric models.

45.10 Exercises

Exercise 45.10.1

Plot impulse responses and spectra side-by-side for several values of the accelerator v in the Hansen-Samuelson model, showing how acceleration strength affects both the time-domain and frequency-domain signatures.

Use the same v values as in the main text: $v \in \{0.2, 0.4, 0.6, 0.8, 0.95\}$ with $c = 0.6$.

Solution

Here is one solution:

```
v_grid_ex1 = [0.2, 0.4, 0.6, 0.8, 0.95]
c_ex1 = 0.6
freq_ex1 = np.linspace(1e-4, 0.5, 2000)
ω_grid_ex1 = 2 * np.pi * freq_ex1
V_ex1 = np.array([[1.0, 0.0], [0.0, 0.0]])
T_irf_ex1 = 40

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

for v in v_grid_ex1:
    A = samuelson_transition(c_ex1, v)

    # impulse response (left panel)
    s = np.array([1.0, 0.0])
    irf = np.empty(T_irf_ex1 + 1)
    for t in range(T_irf_ex1 + 1):
        irf[t] = s[0]
        s = A @ s
    axes[0].plot(range(T_irf_ex1 + 1), irf, lw=2, label=f'$v={v}$')

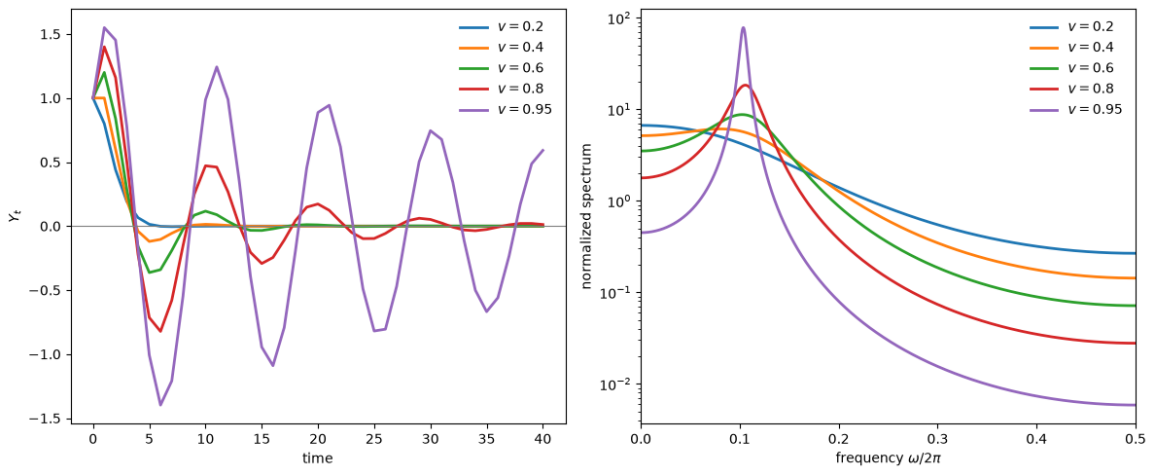
    # spectrum (right panel)
    F = spectral_density_var1(A, V_ex1, ω_grid_ex1)
    f11 = np.real(F[:, 0, 0])
    df = np.diff(freq_ex1)
    area = np.sum(0.5 * (f11[1:] + f11[:-1]) * df)
    f11_norm = f11 / area
    axes[1].plot(freq_ex1, f11_norm, lw=2, label=f'$v={v}$')

axes[0].axhline(0, lw=0.8, color='gray')
axes[0].set_xlabel('time')
axes[0].set_ylabel(r'$Y_t$')
axes[0].legend(frameon=False)

axes[1].set_xlabel(r'frequency $\omega/2\pi$')
axes[1].set_ylabel('normalized spectrum')
```

```
axes[1].set_xlim([0, 0.5])
axes[1].set_yscale('log')
axes[1].legend(frameon=False)

plt.tight_layout()
plt.show()
```



As v increases, eigenvalues approach the unit circle: oscillations become more persistent in the time domain (left), and the spectral peak becomes sharper in the frequency domain (right).

Complex roots produce a pronounced peak at interior frequencies–the spectral signature of business cycles.

Exercise 45.10.2

Verify spectral peak condition (45.25) numerically for the Hansen-Samuelson model.

1. For a range of eigenvalue moduli $r \in [0.3, 0.99]$ with fixed $\theta = 60°$, compute:
   - The theoretical peak frequency from formula: $\cos\omega = \frac{1+r^2}{2r}\cos\theta$
   - The actual peak frequency by numerically maximizing the spectral density
2. Plot both on the same graph and verify they match.
3. Identify the range of r for which no valid peak exists (when the condition (45.26) is violated).

Solution

Here is one solution:

```
θ_ex = np.pi / 3  # 60 degrees
r_grid = np.linspace(0.3, 0.99, 50)
ω_grid_ex = np.linspace(1e-3, np.pi - 1e-3, 1000)
V_hs_ex = np.array([[1.0, 0.0], [0.0, 0.0]])

ω_theory = []
ω_numerical = []

for r in r_grid:
    # Theoretical peak
```

```
    factor = (1 + r**2) / (2 * r)
    cos_ω = factor * np.cos(θ_ex)
    if -1 < cos_ω < 1:
        ω_theory.append(np.arccos(cos_ω))
    else:
        ω_theory.append(np.nan)

    # Numerical peak from spectral density
    # Construct Hansen-Samuelson with eigenvalues r*exp(+-iθ)
    # This corresponds to c + v = 2r*cos(θ), v = r^2
    v = r**2
    c = 2 * r * np.cos(θ_ex) - v
    A_ex = samuelson_transition(c, v)
    F_ex = spectral_density_var1(A_ex, V_hs_ex, ω_grid_ex)
    f11 = np.real(F_ex[:, 0, 0])
    i_max = np.argmax(f11)

    # Only count as a peak if it's not at the boundary
    if 5 < i_max < len(ω_grid_ex) - 5:
        ω_numerical.append(ω_grid_ex[i_max])
    else:
        ω_numerical.append(np.nan)

ω_theory = np.array(ω_theory)
ω_numerical = np.array(ω_numerical)

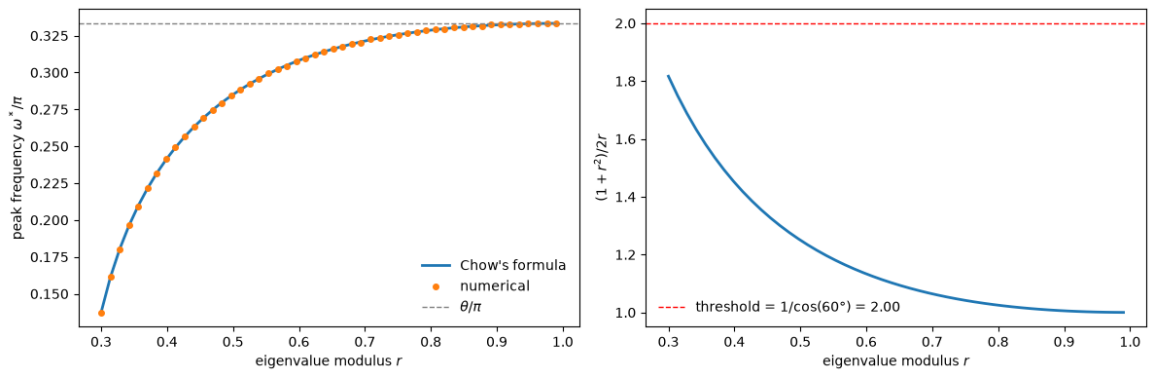
fig, axes = plt.subplots(1, 2, figsize=(12, 4))

# Plot peak frequencies
axes[0].plot(r_grid, ω_theory / np.pi, lw=2, label="Chow's formula")
axes[0].plot(r_grid, ω_numerical / np.pi, 'o', markersize=4, label='numerical')
axes[0].axhline(θ_ex / np.pi, ls='--', lw=1.0, color='gray', label=r'$\theta/\pi$')
axes[0].set_xlabel('eigenvalue modulus $r$')
axes[0].set_ylabel(r'peak frequency $\omega^*/\pi$')
axes[0].legend(frameon=False)

# Plot the factor (1+r^2)/2r to show when peaks are valid
axes[1].plot(r_grid, (1 + r_grid**2) / (2 * r_grid), lw=2)
axes[1].axhline(1 / np.cos(θ_ex), ls='--', lw=1.0, color='red',
                label=f'threshold = 1/cos({np.rad2deg(θ_ex):.0f}°) = {1/np.cos(θ_
 ↪ex):.2f}')
axes[1].set_xlabel('eigenvalue modulus $r$')
axes[1].set_ylabel(r'$(1+r^2)/2r$')
axes[1].legend(frameon=False)

plt.tight_layout()
plt.show()

# Find threshold r below which no peak exists
valid_mask = ~np.isnan(ω_theory)
if valid_mask.any():
    r_threshold = r_grid[valid_mask][0]
    print(f"Peak exists for r >= {r_threshold:.2f}")
```




```
Peak exists for r >= 0.30
```

The theoretical and numerical peak frequencies match closely.

As $r \to 1$, the peak frequency converges to θ.

For smaller r, the factor $(1 + r^2)/2r$ exceeds the threshold, and no valid peak exists.

Exercise 45.10.3

In the "real roots but a peak" example, hold A fixed and vary the shock correlation (the off-diagonal entry of V) between 0 and 0.99.

When does the interior-frequency peak appear, and how does its location change?

Solution

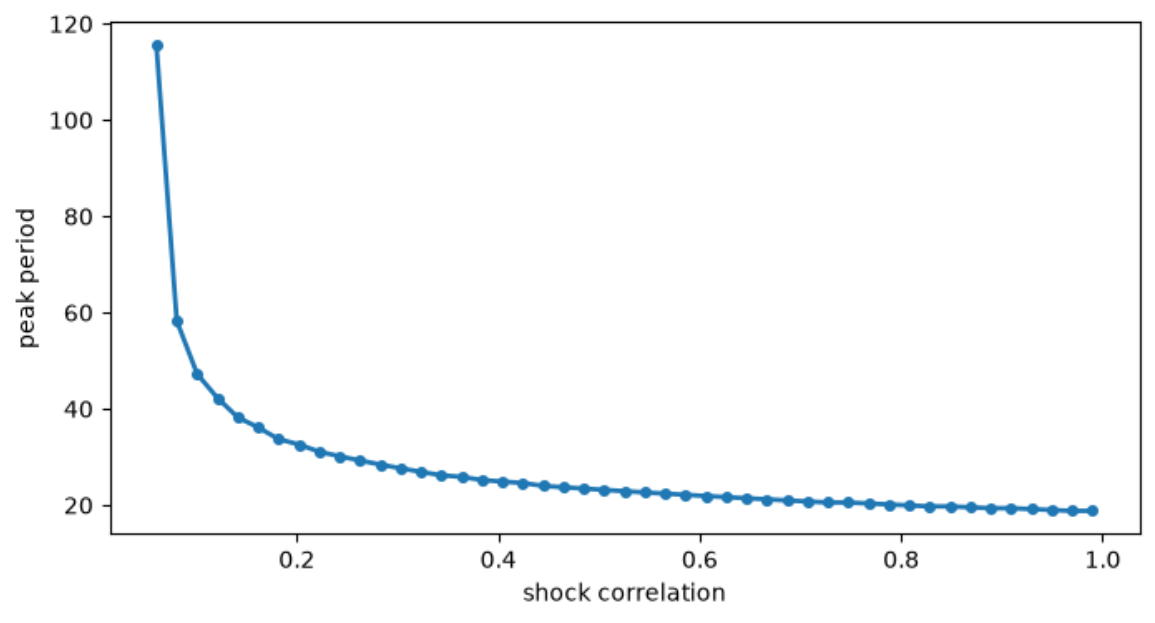
Here is one solution:

```
A_ex3 = np.diag([0.1, 0.9])
b_ex3 = np.array([1.0, -0.01])
corr_grid = np.linspace(0, 0.99, 50)
peak_periods = []
for corr in corr_grid:
    V_ex3 = np.array([[1.0, corr], [corr, 1.0]])
    F_ex3 = spectral_density_var1(A_ex3, V_ex3, ω_grid_ex)
    f_x = spectrum_of_linear_combination(F_ex3, b_ex3)
    i_max = np.argmax(f_x)
    if 5 < i_max < len(ω_grid_ex) - 5:
        peak_periods.append(2 * np.pi / ω_grid_ex[i_max])
    else:
        peak_periods.append(np.nan)

fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(corr_grid, peak_periods, marker='o', lw=2, markersize=4)
ax.set_xlabel('shock correlation')
ax.set_ylabel('peak period')
plt.show()

threshold_idx = np.where(~np.isnan(peak_periods))[0]
if len(threshold_idx) > 0:
    print(
```

```
    f"interior peak when correlation >= {corr_grid[threshold_idx[0]]:.2f}")
```




```
interior peak when correlation >= 0.06
```

The interior peak appears only when the shock correlation exceeds a threshold.

This illustrates that spectral peaks depend on the full system structure, not just eigenvalues.

Exercise 45.10.4

Using the calibrated Chow-Levitan parameters, compute the autocovariance matrices $\Gamma_0, \Gamma_1, \ldots, \Gamma_{10}$ using:

1. The recursion $\Gamma_k = A\Gamma_{k-1}$ with Γ_0 from the Lyapunov equation.
2. Chow's eigendecomposition formula $\Gamma_k = BD_\lambda^k \Gamma_0^* B^\top$ where Γ_0^* is the canonical covariance.

Verify that both methods give the same result.

Solution

Here is one solution:

```
from scipy.linalg import solve_discrete_lyapunov

Γ_0_lyap = solve_discrete_lyapunov(A_chow, V)
Γ_recursion = [Γ_0_lyap]
for k in range(1, 11):
    Γ_recursion.append(A_chow @ Γ_recursion[-1])

p = len(λ)
Γ_0_star = np.zeros((p, p), dtype=complex)
for i in range(p):
    for j in range(p):
        Γ_0_star[i, j] = W[i, j] / (1 - λ[i] * λ[j])
```

```
Γ_eigen = []
for k in range(11):
    D_k = np.diag(λ**k)
    Γ_eigen.append(np.real(B @ D_k @ Γ_0_star @ B.T))

print("Comparison of Γ_5 (first 3x3 block):")
print("\nRecursion method:")
print(np.real(Γ_recursion[5][:3, :3]).round(10))
print("\nEigendecomposition method:")
print(Γ_eigen[5][:3, :3].round(10))
print("\nMax absolute difference:",
        np.max(np.abs(np.real(Γ_recursion[5]) - Γ_eigen[5])))
```

```
Comparison of Γ_5 (first 3x3 block):

Recursion method:
[[2.5417901e-03 2.9310700e-05 1.7370210e-04]
 [2.8886900e-05 3.3220000e-07 1.9737000e-06]
 [1.7356020e-04 2.0028000e-06 1.1861700e-05]]

Eigendecomposition method:
[[2.5417901e-03 2.9310700e-05 1.7370210e-04]
 [2.8886900e-05 3.3220000e-07 1.9737000e-06]
 [1.7356020e-04 2.0028000e-06 1.1861700e-05]]

Max absolute difference: 9.500386588534582e-14
```

Both methods produce essentially identical results, up to numerical precision.

Exercise 45.10.5

Modify the Chow-Levitan model by changing λ_3 from 0.4838 to 0.95.

1. Recompute the spectral densities.
2. How does this change affect the spectral shape for each variable?
3. What economic interpretation might correspond to this parameter change?

Solution

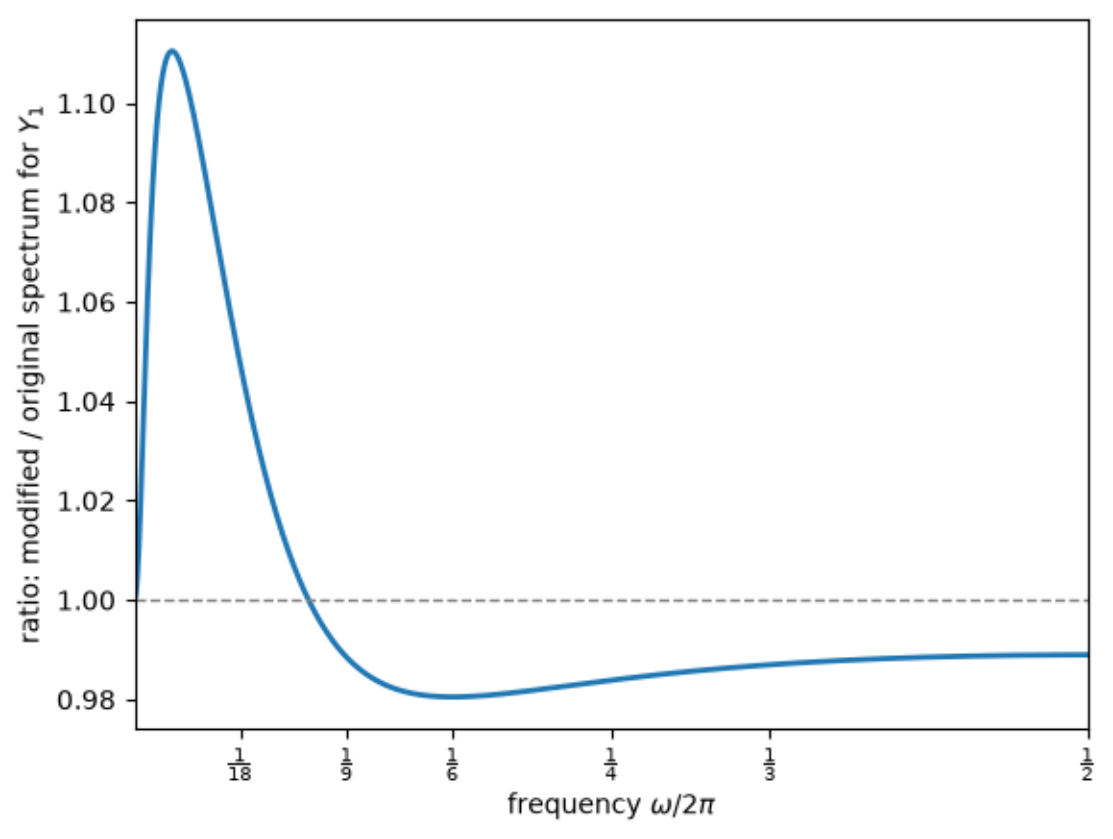
Here is one solution:

```
# Modify λ_3 and reconstruct the transition matrix
λ_modified = λ.copy()
λ_modified[2] = 0.95
D_λ_mod = np.diag(λ_modified)
A_mod = np.real(B @ D_λ_mod @ np.linalg.inv(B))

# Compute spectra using the VAR(1) formula with original V
F_mod = spectral_density_var1(A_mod, V, ω_grid)
F_orig = spectral_density_var1(A_chow, V, ω_grid)

# Plot ratio of spectra for output (Y_1)
f_orig = np.real(F_orig[:, 4, 4])
f_mod = np.real(F_mod[:, 4, 4])
```

```
fig, ax = plt.subplots()
ax.plot(freq, f_mod / f_orig, lw=2)
ax.axhline(1.0, ls='--', lw=1, color='gray')
paper_frequency_axis(ax)
ax.set_ylabel(r"ratio: modified / original spectrum for $Y_1$")
plt.show()
```



The near-unit eigenvalues ($\lambda_1, \lambda_2 \approx 0.9999$) dominate the output spectrum so heavily that changing λ_3 from 0.48 to 0.95 produces only a small relative effect.

The ratio plot reveals the change: the modified spectrum has slightly more power at low-to-medium frequencies and slightly less at high frequencies.

Economically, increasing λ_3 adds persistence to the mode it governs.

CHAPTER

FORTYSIX

KESTEN PROCESSES AND FIRM DYNAMICS

Contents

- *Kesten Processes and Firm Dynamics*
  - *Overview*
  - *Kesten processes*
  - *Heavy tails*
  - *Application: firm dynamics*
  - *Exercises*

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install --upgrade quantecon yfinance jax
```

46.1 Overview

Previously in AR(1) Processes, we learned about linear scalar-valued stochastic processes (AR(1) models).

Now we generalize these linear models slightly by allowing the multiplicative coefficient to be stochastic.

Such processes are known as Kesten processes after German–American mathematician Harry Kesten (1931–2019).

Although simple to write down, Kesten processes are interesting for at least two reasons:

1. A number of significant economic processes are or can be described as Kesten processes.

2. Kesten processes generate interesting dynamics, including, in some cases, heavy-tailed cross-sectional distributions.

We will discuss these issues as we go along.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
import yfinance as yf
```

Additional technical background related to this lecture can be found in the monograph by [Buraczewski *et al.*, 2016].

46.2 Kesten processes

A **Kesten process** is a stochastic process of the form

$$X_{t+1} = a_{t+1} X_t + \eta_{t+1} \tag{46.1}$$

where $\{a_t\}_{t \geq 1}$ and $\{\eta_t\}_{t \geq 1}$ are IID sequences.

We are interested in the dynamics of $\{X_t\}_{t \geq 0}$ when X_0 is given.

We will focus on the nonnegative scalar case, where X_t takes values in $\mathbb{R}_+$.

In particular, we will assume that

- the initial condition X_0 is nonnegative,
- $\{a_t\}_{t \geq 1}$ is a nonnegative IID stochastic process and
- $\{\eta_t\}_{t \geq 1}$ is another nonnegative IID stochastic process, independent of the first.

46.2.1 Example: GARCH volatility

The GARCH model is common in financial applications, where time series such as asset returns exhibit time varying volatility.

For example, consider the following plot of daily returns on the Nasdaq Composite Index for the period 1st January 2006 to 1st November 2019.

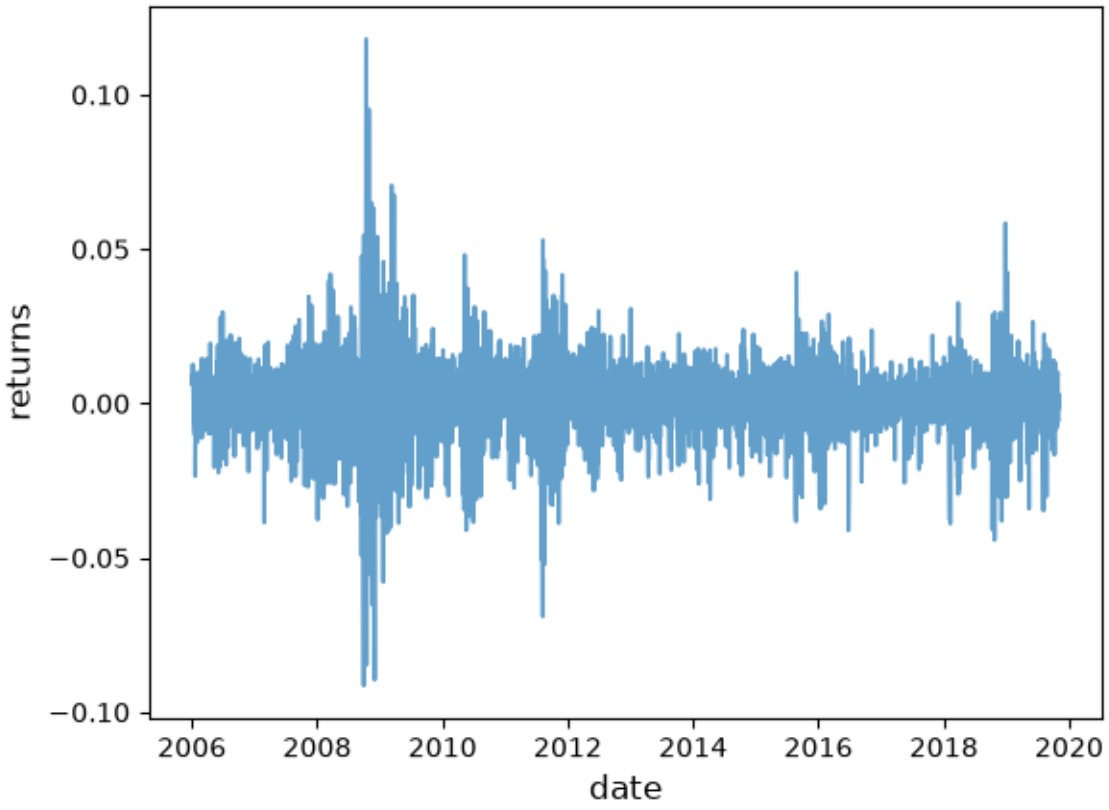
```
s = yf.download(
    "^IXIC", "2006-1-1", "2019-11-1", auto_adjust=False
)["Adj Close"]

r = s.pct_change()

fig, ax = plt.subplots()
ax.plot(r, alpha=0.7)
ax.set_ylabel("returns", fontsize=12)
ax.set_xlabel("date", fontsize=12)

plt.show()
```

```
[**********100%**********]  1 of 1 completed
```



Notice how the series exhibits bursts of volatility (high variance) and then settles down again.

GARCH models can replicate this feature.

The GARCH(1, 1) volatility process takes the form

$$\sigma_{t+1}^2 = \alpha_0 + \sigma_t^2(\alpha_1 \xi_{t+1}^2 + \beta) \tag{46.2}$$

where $\{\xi_t\}$ is IID with $\mathbb{E}\xi_t^2 = 1$ and all parameters are positive.

Returns on a given asset are then modeled as

$$r_t = \sigma_t \zeta_t \tag{46.3}$$

where $\{\zeta_t\}$ is again IID and independent of $\{\xi_t\}$.

The volatility sequence $\{\sigma_t^2\}$, which drives the dynamics of returns, is a Kesten process.

46.2.2 Example: wealth dynamics

Suppose that a given household saves a fixed fraction s of its current wealth in every period.

The household earns labor income y_t at the start of time t.

Wealth then evolves according to

$$w_{t+1} = R_{t+1} s w_t + y_{t+1} \tag{46.4}$$

where $\{R_t\}$ is the gross rate of return on assets.

If $\{R_t\}$ and $\{y_t\}$ are both IID, then (46.4) is a Kesten process.

46.2.3 Stationarity

In earlier lectures, such as the one on AR(1) Processes, we introduced the notion of a stationary distribution.

In the present context, we can define a stationary distribution as follows:

The distribution F^* on $\mathbb{R}$ is called **stationary** for the Kesten process (46.1) if

$$X_t \sim F^* \quad \Longrightarrow \quad a_{t+1} X_t + \eta_{t+1} \sim F^* \tag{46.5}$$

In other words, if the current state X_t has distribution F^*, then so does the next period state X_{t+1}.

We can write this alternatively as

$$F^*(y) = \int \mathbb{P}\{a_{t+1} x + \eta_{t+1} \leq y\} F^*(dx) \quad \text{for all } y \geq 0. \tag{46.6}$$

The left hand side is the distribution of the next period state when the current state is drawn from F^*.

The equality in (46.6) states that this distribution is unchanged.

46.2.4 Cross-sectional interpretation

There is an important cross-sectional interpretation of stationary distributions, discussed previously but worth repeating here.

Suppose, for example, that we are interested in the wealth distribution — that is, the current distribution of wealth across households in a given country.

Suppose further that

- the wealth of each household evolves independently according to (46.4),
- F^* is a stationary distribution for this stochastic process and
- there are many households.

Then F^* is a steady state for the cross-sectional wealth distribution in this country.

In other words, if F^* is the current wealth distribution then it will remain so in subsequent periods, *ceteris paribus*.

To see this, suppose that F^* is the current wealth distribution.

What is the fraction of households with wealth less than y next period?

To obtain this, we sum the probability that wealth is less than y tomorrow, given that current wealth is w, weighted by the fraction of households with wealth w.

Noting that the fraction of households with wealth in interval dw is $F^*(dw)$, we get

$$\int \mathbb{P}\{R_{t+1} sw + y_{t+1} \leq y\} F^*(dw)$$

By the definition of stationarity and the assumption that F^* is stationary for the wealth process, this is just $F^*(y)$.

Hence the fraction of households with wealth in $[0, y]$ is the same next period as it is this period.

Since y was chosen arbitrarily, the distribution is unchanged.

46.2.5 Conditions for stationarity

The Kesten process $X_{t+1} = a_{t+1}X_t + \eta_{t+1}$ does not always have a stationary distribution.

For example, if $a_t \equiv \eta_t \equiv 1$ for all t, then $X_t = X_0 + t$, which diverges to infinity.

To prevent this kind of divergence, we require that $\{a_t\}$ is strictly less than 1 most of the time.

In particular, if

$$\mathbb{E} \ln a_t < 0 \quad \text{and} \quad \mathbb{E}\eta_t < \infty \tag{46.7}$$

then a unique stationary distribution exists on $\mathbb{R}_+$.

- See, for example, theorem 2.1.3 of [Buraczewski *et al.*, 2016], which provides slightly weaker conditions.

As one application of this result, we see that the wealth process (46.4) will have a unique stationary distribution whenever labor income has finite mean and $\mathbb{E} \ln R_t + \ln s < 0$.

46.3 Heavy tails

Under certain conditions, the stationary distribution of a Kesten process has a Pareto tail.

(See our lecture Heavy-Tailed Distributions for background.)

This fact is significant for economics because of the prevalence of Pareto-tailed distributions.

46.3.1 The Kesten–Goldie theorem

To state the conditions under which the stationary distribution of a Kesten process has a Pareto tail, we first recall that a random variable is called **nonarithmetic** if its distribution is not concentrated on $\{\dots, -2t, -t, 0, t, 2t, \dots\}$ for any $t \geq 0$.

For example, any random variable with a density is nonarithmetic.

The famous Kesten–Goldie Theorem (see, e.g., [Buraczewski *et al.*, 2016], theorem 2.4.4) states that if

1. the stationarity conditions in (46.7) hold,
2. the random variable a_t is positive with probability one and nonarithmetic,
3. $\mathbb{P}\{a_t x + \eta_t = x\} < 1$ for all $x \in \mathbb{R}_+$ and
4. there exists a positive constant α such that

$$\mathbb{E}a_t^\alpha = 1, \quad \mathbb{E}\eta_t^\alpha < \infty, \quad \text{and} \quad \mathbb{E}[a_t^{\alpha+1}] < \infty$$

then the stationary distribution of the Kesten process has a Pareto tail with tail index α.

More precisely, if F^* is the unique stationary distribution and $X^* \sim F^*$, then

$$\lim_{x \to \infty} x^\alpha \mathbb{P}\{X^* > x\} = c$$

for some positive constant c.

46.3.2 Intuition

Later we will illustrate the Kesten–Goldie Theorem using rank-size plots.

Prior to doing so, we can give the following intuition for the conditions.

Two important conditions are that $\mathbb{E} \ln a_t < 0$, so the model is stationary, and $\mathbb{E} a_t^\alpha = 1$ for some $\alpha > 0$.

The first condition implies that the distribution of a_t has a large amount of probability mass below 1.

The second condition implies that the distribution of a_t has at least some probability mass at or above 1.

The first condition gives us existence of the stationary condition.

The second condition means that the current state can be expanded by a_t.

If this occurs for several concurrent periods, the effects compound each other, since a_t is multiplicative.

This leads to spikes in the time series, which fill out the extreme right hand tail of the distribution.

The spikes in the time series are visible in the following simulation, which generates 10 paths when a_t and b_t are lognormal.

```
μ = -0.5
σ = 1.0

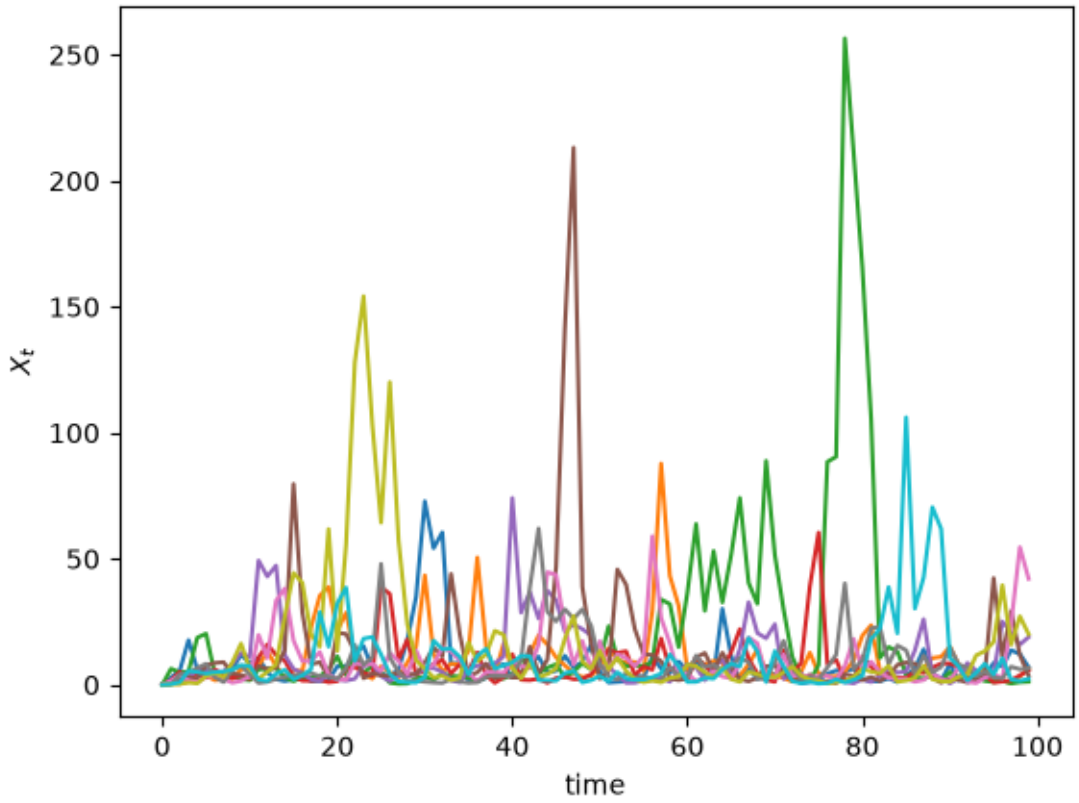
def kesten_ts(rng, ts_length=100):
    x = np.zeros(ts_length)
    for t in range(ts_length - 1):
        a = np.exp(μ + σ * rng.standard_normal())
        b = np.exp(rng.standard_normal())
        x[t+1] = a * x[t] + b
    return x

fig, ax = plt.subplots()

num_paths = 10
rng = np.random.default_rng(12)

for i in range(num_paths):
    ax.plot(kesten_ts(rng))

ax.set(xlabel="time", ylabel="$X_t$")
plt.show()
```



46.4 Application: firm dynamics

As noted in our lecture Heavy-Tailed Distributions, for common measures of firm size such as revenue or employment, the US firm size distribution exhibits a Pareto tail (see, e.g., [Axtell, 2001], [Gabaix, 2016]).

Let us try to explain this rather striking fact using the Kesten–Goldie Theorem.

46.4.1 Gibrat's law

It was postulated many years ago by Robert Gibrat [Gibrat, 1931] that firm size evolves according to a simple rule whereby size next period is proportional to current size.

This is now known as Gibrat's law of proportional growth.

We can express this idea by stating that a suitably defined measure s_t of firm size obeys

$$\frac{s_{t+1}}{s_t} = a_{t+1} \tag{46.8}$$

for some positive IID sequence $\{a_t\}$.

One implication of Gibrat's law is that the growth rate of individual firms does not depend on their size.

However, over the last few decades, research contradicting Gibrat's law has accumulated in the literature.

For example, it is commonly found that, on average,

1. small firms grow faster than large firms (see, e.g., [Evans, 1987] and [Hall, 1987]) and

2. the growth rate of small firms is more volatile than that of large firms [Dunne *et al.*, 1989].

On the other hand, Gibrat's law is generally found to be a reasonable approximation for large firms [Evans, 1987].

We can accommodate these empirical findings by modifying (46.8) to

$$s_{t+1} = a_{t+1} s_t + b_{t+1} \tag{46.9}$$

where $\{a_t\}$ and $\{b_t\}$ are both IID and independent of each other.

In the exercises you are asked to show that (46.9) is more consistent with the empirical findings presented above than Gibrat's law in (46.8).

46.4.2 Heavy tails

So what has this to do with Pareto tails?

The answer is that (46.9) is a Kesten process.

If the conditions of the Kesten–Goldie Theorem are satisfied, then the firm size distribution is predicted to have heavy tails — which is exactly what we see in the data.

In the exercises below we explore this idea further, generalizing the firm size dynamics and examining the corresponding rank-size plots.

We also try to illustrate why the Pareto tail finding is significant for quantitative analysis.

46.5 Exercises

Exercise 46.5.1

Simulate and plot 15 years of daily returns (consider each year as having 250 working days) using the GARCH(1, 1) process in (46.2)–(46.3).

Take ξ_t and ζ_t to be independent and standard normal.

Set $\alpha_0 = 0.00001, \alpha_1 = 0.1, \beta = 0.9$ and $\sigma_0 = 0$.

Compare visually with the Nasdaq Composite Index returns *shown above*.

While the time path differs, you should see bursts of high volatility.

Solution

Here is one solution:

```
α_0 = 1e-5
α_1 = 0.1
β = 0.9

years = 15
days = years * 250

def garch_ts(rng, ts_length=days):
    σ2 = 0
    r = np.zeros(ts_length)
```

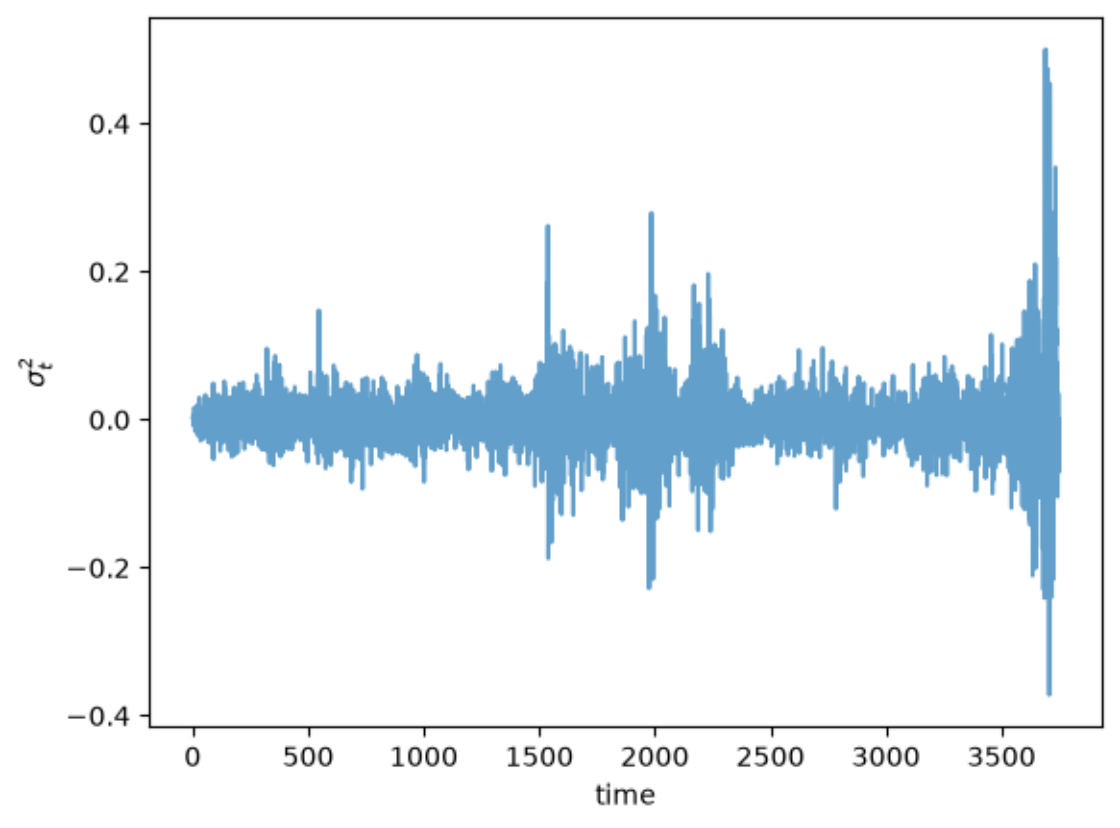
```
    for t in range(ts_length - 1):
        ξ = rng.standard_normal()
        σ2 = α_0 + σ2 * (α_1 * ξ**2 + β)
        r[t] = np.sqrt(σ2) * rng.standard_normal()
    return r

fig, ax = plt.subplots()

rng = np.random.default_rng(12)

ax.plot(garch_ts(rng), alpha=0.7)

ax.set(xlabel="time", ylabel="$\\sigma_t^2$")
plt.show()
```



Exercise 46.5.2

In our discussion of firm dynamics, it was claimed that (46.9) is more consistent with the empirical literature than Gibrat's law in (46.8).

(The empirical literature was reviewed immediately above (46.9).)

In what sense is this true (or false)?

Solution

The empirical findings are that

1. small firms grow faster than large firms and
2. the growth rate of small firms is more volatile than that of large firms.

Also, Gibrat's law is generally found to be a reasonable approximation for large firms than for small firms

The claim is that the dynamics in (46.9) are more consistent with points 1-2 than Gibrat's law.

To see why, we rewrite (46.9) in terms of growth dynamics:

$$\frac{s_{t+1}}{s_t} = a_{t+1} + \frac{b_{t+1}}{s_t} \tag{46.10}$$

Taking $s_t = s$ as given, the mean and variance of firm growth are

$$\mathbb{E}a + \frac{\mathbb{E}b}{s} \quad \text{and} \quad \mathbb{V}a + \frac{\mathbb{V}b}{s^2}$$

Both of these decline with firm size s, consistent with the data.

Moreover, the law of motion (46.10) clearly approaches Gibrat's law (46.8) as s_t gets large.

Exercise 46.5.3

Consider an arbitrary Kesten process as given in (46.1).

Suppose that $\{a_t\}$ is lognormal with parameters (μ, σ).

In other words, each a_t has the same distribution as $\exp(\mu + \sigma Z)$ when Z is standard normal.

Suppose further that $\mathbb{E}\eta_t^r < \infty$ for every $r > 0$, as would be the case if, say, η_t is also lognormal.

Show that the conditions of the Kesten–Goldie theorem are satisfied if and only if $\mu < 0$.

Obtain the value of α that makes the Kesten–Goldie conditions hold.

Solution

Since a_t has a density it is nonarithmetic.

Since a_t has the same density as $a = \exp(\mu + \sigma Z)$ when Z is standard normal, we have

$$\mathbb{E} \ln a_t = \mathbb{E}(\mu + \sigma Z) = \mu,$$

and since η_t has finite moments of all orders, the stationarity condition holds if and only if $\mu < 0$.

Given the properties of the lognormal distribution (which has finite moments of all orders), the only other condition in doubt is existence of a positive constant α such that $\mathbb{E} a_t^\alpha = 1$.

This is equivalent to the statement

$$\exp\left(\alpha\mu + \frac{\alpha^2\sigma^2}{2}\right) = 1.$$

Solving for α gives $\alpha = -2\mu/\sigma^2$.

Exercise 46.5.4

One unrealistic aspect of the firm dynamics specified in (46.9) is that it ignores entry and exit.

In any given period and in any given market, we observe significant numbers of firms entering and exiting the market.

Empirical discussion of this can be found in a famous paper by Hugo Hopenhayn [Hopenhayn, 1992].

In the same paper, Hopenhayn builds a model of entry and exit that incorporates profit maximization by firms and market clearing quantities, wages and prices.

In his model, a stationary equilibrium occurs when the number of entrants equals the number of exiting firms.

In this setting, firm dynamics can be expressed as

$$s_{t+1} = e_{t+1} \mathbb{1}\{s_t < \bar{s}\} + (a_{t+1} s_t + b_{t+1}) \mathbb{1}\{s_t \geq \bar{s}\} \tag{46.11}$$

Here

- the state variable s_t represents productivity (which is a proxy for output and hence firm size),
- the IID sequence $\{e_t\}$ is thought of as a productivity draw for a new entrant and
- the variable $\bar{s}$ is a threshold value that we take as given, although it is determined endogenously in Hopenhayn's model.

The idea behind (46.11) is that firms stay in the market as long as their productivity s_t remains at or above $\bar{s}$.

- In this case, their productivity updates according to (46.9).

Firms choose to exit when their productivity s_t falls below $\bar{s}$.

- In this case, they are replaced by a new firm with productivity e_{t+1}.

What can we say about dynamics?

Although (46.11) is not a Kesten process, it does update in the same way as a Kesten process when s_t is large.

So perhaps its stationary distribution still has Pareto tails?

Your task is to investigate this question via simulation and rank-size plots.

The approach will be to

1. generate M draws of s_T when M and T are large and
2. plot the largest 1,000 of the resulting draws in a rank-size plot.

(The distribution of s_T will be close to the stationary distribution when T is large.)

In the simulation, assume that

- each of a_t, b_t and e_t is lognormal,
- the parameters are

```
μ_a = -0.5        # location parameter for a
σ_a = 0.1         # scale parameter for a
μ_b = 0.0         # location parameter for b
σ_b = 0.5         # scale parameter for b
μ_e = 0.0         # location parameter for e
σ_e = 0.5         # scale parameter for e
s_bar = 1.0       # threshold
T = 500           # sampling date
M = 1_000_000     # number of firms
s_init = 1.0      # initial condition for each firm
```

Solution

Here's one solution.

We need to simulate a large cross-section of firms, so this is a good use case for JAX.

We store the firm parameters in a `NamedTuple` and write the update rule (46.11) for a single firm as a pure function.

```
import jax
import jax.numpy as jnp
from functools import partial
from typing import NamedTuple


class Firm(NamedTuple):
    μ_a: float
    σ_a: float
    μ_b: float
    σ_b: float
    μ_e: float
    σ_e: float
    s_bar: float


def update(s, key, firm):
    "Update a single firm's size s, given a fresh random key."
    a_key, b_key, e_key = jax.random.split(key, 3)
    a = jnp.exp(firm.μ_a + firm.σ_a * jax.random.normal(a_key))
    b = jnp.exp(firm.μ_b + firm.σ_b * jax.random.normal(b_key))
    e = jnp.exp(firm.μ_e + firm.σ_e * jax.random.normal(e_key))
    return jnp.where(s < firm.s_bar, e, a * s + b)
```

We simulate one firm for `T` periods with `lax.scan`, then use `vmap` to run all `M` firms in parallel.

```
@partial(jax.jit, static_argnames=("T", "M"))
def generate_draws(firm, s_init=1.0, T=500, M=1_000_000, seed=0):
    keys = jax.random.split(jax.random.PRNGKey(seed), M)

    def sim_firm(key):
        step_keys = jax.random.split(key, T)
        s_final, _ = jax.lax.scan(
            lambda s, k: (update(s, k, firm), None), s_init, step_keys
        )
        return s_final

    return jax.vmap(sim_firm)(keys)
```

Now we build the firm from the parameters above and generate the draws.

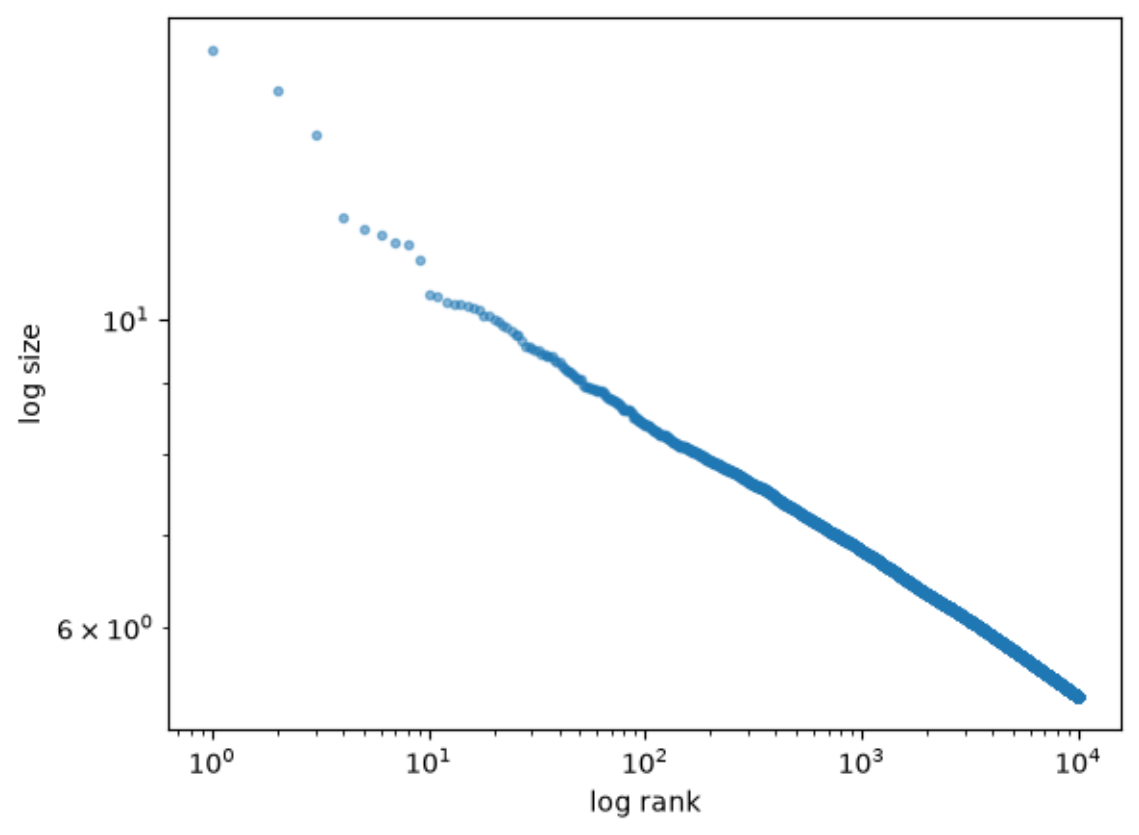
```
firm = Firm(μ_a, σ_a, μ_b, σ_b, μ_e, σ_e, s_bar)
data = generate_draws(firm, s_init=s_init, T=T, M=M)
```

Now we produce the rank-size plot:

```
fig, ax = plt.subplots()

rank_data, size_data = qe.rank_size(data, c=0.01)
ax.loglog(rank_data, size_data, "o", markersize=3.0, alpha=0.5)
ax.set_xlabel("log rank")
ax.set_ylabel("log size")

plt.show()
```



The plot produces a straight line, consistent with a Pareto tail.

CHAPTER

FORTYSEVEN

WEALTH DISTRIBUTION DYNAMICS

Contents

- *Wealth Distribution Dynamics*
  - *Overview*
  - *Lorenz Curves and the Gini Coefficient*
  - *A Model of Wealth Dynamics*
  - *Implementation*
  - *Applications*
  - *Exercises*

See also

A version of this lecture using JAX is available here

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

47.1 Overview

This notebook gives an introduction to wealth distribution dynamics, with a focus on

- modeling and computing the wealth distribution via simulation,
- measures of inequality such as the Lorenz curve and Gini coefficient, and
- how inequality is affected by the properties of wage income and returns on assets.

One interesting property of the wealth distribution we discuss is Pareto tails.

The wealth distribution in many countries exhibits a Pareto tail

- See this lecture for a definition.
- For a review of the empirical evidence, see, for example, [Benhabib and Bisin, 2018].

This is consistent with high concentration of wealth amongst the richest households.

It also gives us a way to quantify such concentration, in terms of the tail index.

One question of interest is whether or not we can replicate Pareto tails from a relatively simple model.

47.1.1 A Note on Assumptions

The evolution of wealth for any given household depends on their savings behavior.

Modeling such behavior will form an important part of this lecture series.

However, in this particular lecture, we will be content with rather ad hoc (but plausible) savings rules.

We do this to more easily explore the implications of different specifications of income dynamics and investment returns.

At the same time, all of the techniques discussed here can be plugged into models that use optimization to obtain savings rules.

We will use the following imports.

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
from numba import jit, float64, prange
from numba.experimental import jitclass
```

47.2 Lorenz Curves and the Gini Coefficient

Before we investigate wealth dynamics, we briefly review some measures of inequality.

47.2.1 Lorenz Curves

One popular graphical measure of inequality is the Lorenz curve.

The package QuantEcon.py, already imported above, contains a function to compute Lorenz curves.

To illustrate, suppose that

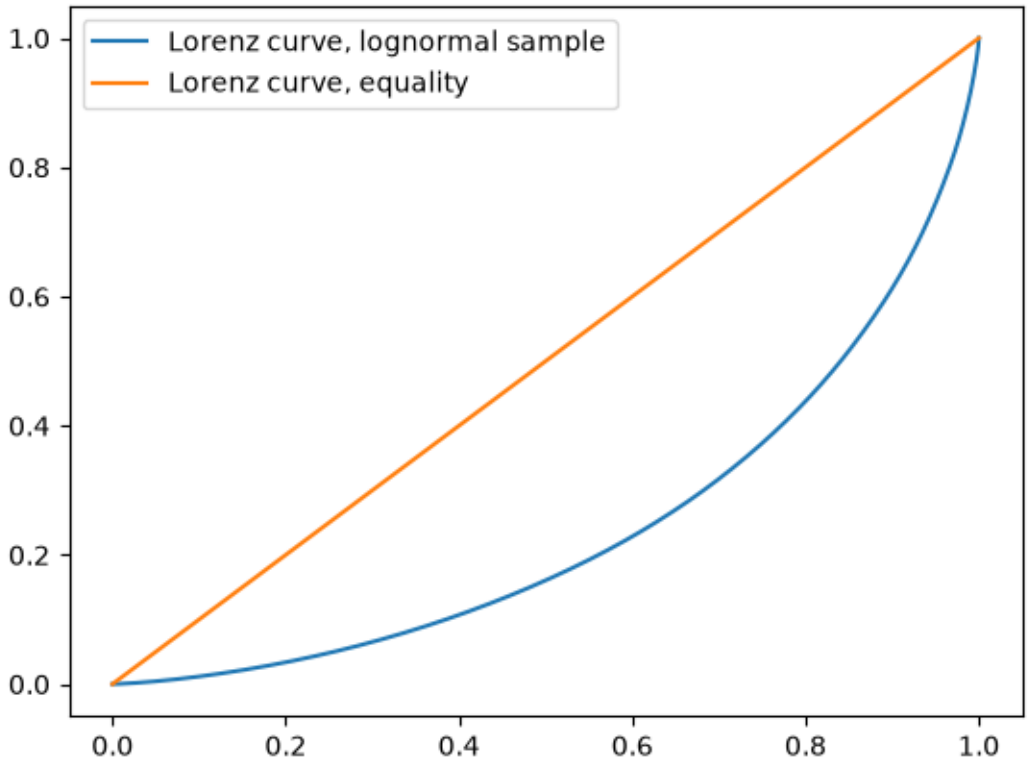
```
rng = np.random.default_rng()
n = 10_000                       # size of sample
w = np.exp(rng.standard_normal(n))   # lognormal draws
```

is data representing the wealth of 10,000 households.

We can compute and plot the Lorenz curve as follows:

```
f_vals, l_vals = qe.lorenz_curve(w)

fig, ax = plt.subplots()
ax.plot(f_vals, l_vals, label='Lorenz curve, lognormal sample')
ax.plot(f_vals, f_vals, label='Lorenz curve, equality')
ax.legend()
plt.show()
```



This curve can be understood as follows: if point (x, y) lies on the curve, it means that, collectively, the bottom $(100x)\%$ of the population holds $(100y)\%$ of the wealth.

The "equality" line is the 45 degree line (which might not be exactly 45 degrees in the figure, depending on the aspect ratio).

A sample that produces this line exhibits perfect equality.

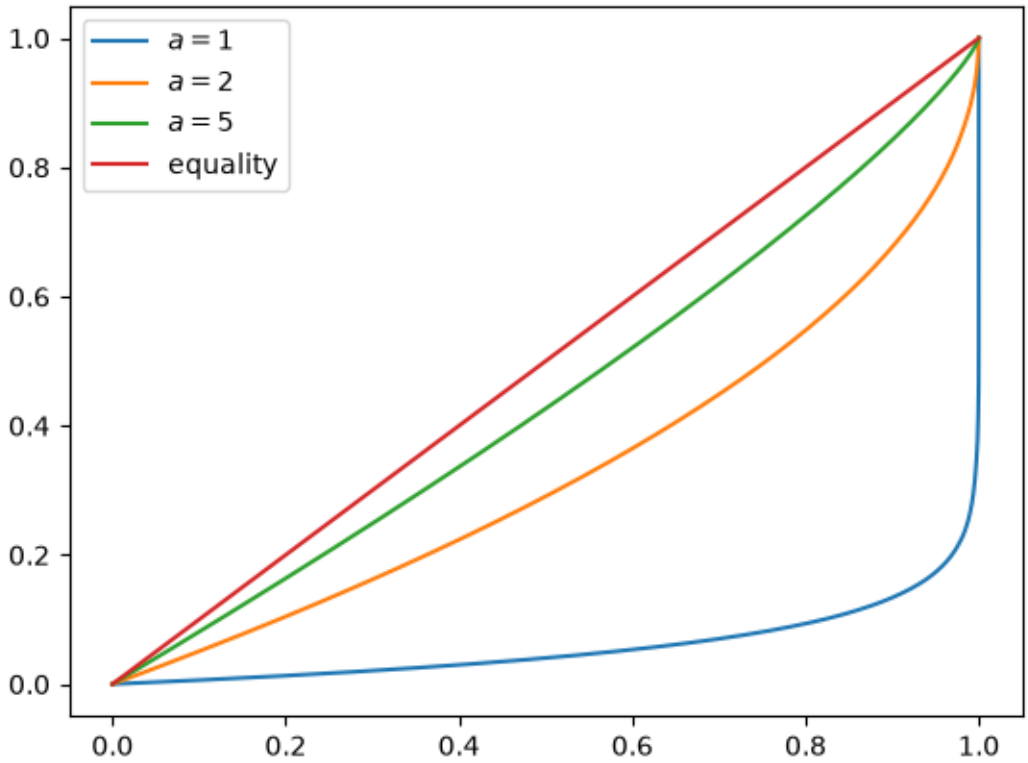
The other line in the figure is the Lorenz curve for the lognormal sample, which deviates significantly from perfect equality.

For example, the bottom 80% of the population holds around 40% of total wealth.

Here is another example, which shows how the Lorenz curve shifts as the underlying distribution changes.

We generate 10,000 observations using the Pareto distribution with a range of parameters, and then compute the Lorenz curve corresponding to each set of observations.

```
a_vals = (1, 2, 5)              # Pareto tail index
n = 10_000                      # size of each sample
fig, ax = plt.subplots()
for a in a_vals:
    u = rng.uniform(size=n)
    y = u**(-1/a)               # distributed as Pareto with tail index a
    f_vals, l_vals = qe.lorenz_curve(y)
    ax.plot(f_vals, l_vals, label=f'$a = {a}$')
ax.plot(f_vals, f_vals, label='equality')
ax.legend()
plt.show()
```



You can see that, as the tail parameter of the Pareto distribution increases, inequality decreases.

This is to be expected, because a higher tail index implies less weight in the tail of the Pareto distribution.

47.2.2 The Gini Coefficient

The definition and interpretation of the Gini coefficient can be found on the corresponding Wikipedia page.

A value of 0 indicates perfect equality (corresponding the case where the Lorenz curve matches the 45 degree line) and a value of 1 indicates complete inequality (all wealth held by the richest household).

The QuantEcon.py library contains a function to calculate the Gini coefficient.

We can test it on the Weibull distribution with parameter a, where the Gini coefficient is known to be

$$G = 1 - 2^{-1/a}$$

Let's see if the Gini coefficient computed from a simulated sample matches this at each fixed value of a.

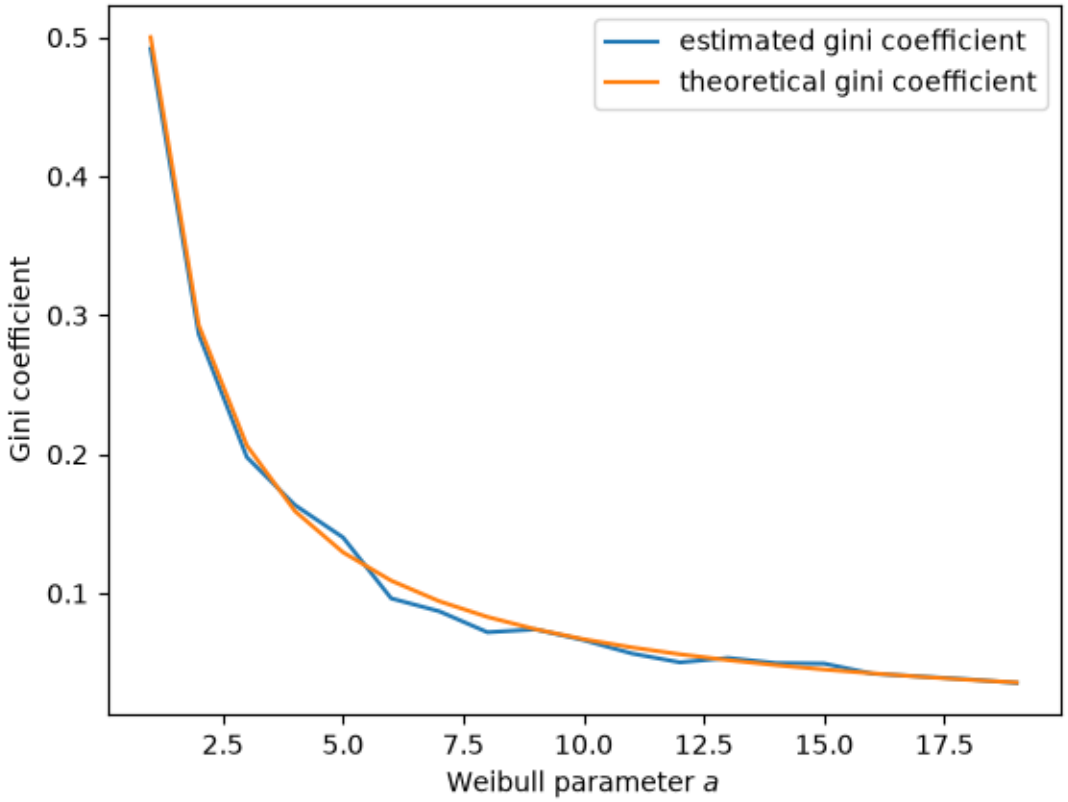
```
a_vals = range(1, 20)
ginis = []
ginis_theoretical = []
n = 100

fig, ax = plt.subplots()
for a in a_vals:
    y = rng.weibull(a, size=n)
    ginis.append(qe.gini_coefficient(y))
    ginis_theoretical.append(1 - 2**(-1/a))
ax.plot(a_vals, ginis, label='estimated gini coefficient')
```

(continues on next page)

(continued from previous page)

```
ax.plot(a_vals, ginis_theoretical, label='theoretical gini coefficient')
ax.legend()
ax.set_xlabel("Weibull parameter $a$")
ax.set_ylabel("Gini coefficient")
plt.show()
```



The simulation shows that the fit is good.

47.3 A Model of Wealth Dynamics

Having discussed inequality measures, let us now turn to wealth dynamics.

The model we will study is

$$w_{t+1} = (1 + r_{t+1})s(w_t) + y_{t+1} \tag{47.1}$$

where

- w_t is wealth at time t for a given household,
- r_t is the rate of return of financial assets,
- y_t is current non-financial (e.g., labor) income and
- $s(w_t)$ is current wealth net of consumption

Letting $\{z_t\}$ be a correlated state process of the form

$$z_{t+1} = az_t + b + \sigma_z \epsilon_{t+1}$$

we'll assume that

$$R_t := 1 + r_t = c_r \exp(z_t) + \exp(\mu_r + \sigma_r \xi_t)$$

and

$$y_t = c_y \exp(z_t) + \exp(\mu_y + \sigma_y \zeta_t)$$

Here $\{(\epsilon_t, \xi_t, \zeta_t)\}$ is IID and standard normal in $\mathbb{R}^3$.

The value of c_r should be close to zero, since rates of return on assets do not exhibit large trends.

When we simulate a population of households, we will assume all shocks are idiosyncratic (i.e., specific to individual households and independent across them).

Regarding the savings function s, our default model will be

$$s(w) = s_0 w \cdot \mathbb{1}\{w \geq \hat{w}\} \tag{47.2}$$

where s_0 is a positive constant.

Thus, for $w < \hat{w}$, the household saves nothing. For $w \geq \bar{w}$, the household saves a fraction s_0 of their wealth.

We are using something akin to a fixed savings rate model, while acknowledging that low wealth households tend to save very little.

47.4 Implementation

Here's some type information to help Numba.

```
wealth_dynamics_data = [
    ('w_hat',  float64),    # savings parameter
    ('s_0',    float64),    # savings parameter
    ('c_y',    float64),    # labor income parameter
    ('μ_y',    float64),    # labor income paraemter
    ('σ_y',    float64),    # labor income parameter
    ('c_r',    float64),    # rate of return parameter
    ('μ_r',    float64),    # rate of return parameter
    ('σ_r',    float64),    # rate of return parameter
    ('a',      float64),    # aggregate shock parameter
    ('b',      float64),    # aggregate shock parameter
    ('σ_z',    float64),    # aggregate shock parameter
    ('z_mean', float64),    # mean of z process
    ('z_var', float64),     # variance of z process
    ('y_mean', float64),    # mean of y process
    ('R_mean', float64)     # mean of R process
]
```

Here's a class that stores instance data and implements methods that update the aggregate state and household wealth.

```
@jitclass(wealth_dynamics_data)
class WealthDynamics:

    def __init__(self,
                 w_hat=1.0,
                 s_0=0.75,
                 c_y=1.0,
```

(continues on next page)

(continued from previous page)

```
                μ_y=1.0,
                σ_y=0.2,
                c_r=0.05,
                μ_r=0.1,
                σ_r=0.5,
                a=0.5,
                b=0.0,
                σ_z=0.1):

        self.w_hat, self.s_0 = w_hat, s_0
        self.c_y, self.μ_y, self.σ_y = c_y, μ_y, σ_y
        self.c_r, self.μ_r, self.σ_r = c_r, μ_r, σ_r
        self.a, self.b, self.σ_z = a, b, σ_z

        # Record stationary moments
        self.z_mean = b / (1 - a)
        self.z_var = σ_z**2 / (1 - a**2)
        exp_z_mean = np.exp(self.z_mean + self.z_var / 2)
        self.R_mean = c_r * exp_z_mean + np.exp(μ_r + σ_r**2 / 2)
        self.y_mean = c_y * exp_z_mean + np.exp(μ_y + σ_y**2 / 2)

        # Test a stability condition that ensures wealth does not diverge
        # to infinity.
        α = self.R_mean * self.s_0
        if α >= 1:
            raise ValueError("Stability condition failed.")

    def parameters(self):
        """
        Collect and return parameters.
        """
        parameters = (self.w_hat, self.s_0,
                      self.c_y, self.μ_y, self.σ_y,
                      self.c_r, self.μ_r, self.σ_r,
                      self.a, self.b, self.σ_z)
        return parameters

    def update_states(self, w, z):
        """
        Update one period, given current wealth w and persistent
        state z.
        """

        # Simplify names
        params = self.parameters()
        w_hat, s_0, c_y, μ_y, σ_y, c_r, μ_r, σ_r, a, b, σ_z = params
        zp = a * z + b + σ_z * np.random.randn()

        # Update wealth
        y = c_y * np.exp(zp) + np.exp(μ_y + σ_y * np.random.randn())
        wp = y
        if w >= w_hat:
            R = c_r * np.exp(zp) + np.exp(μ_r + σ_r * np.random.randn())
            wp += R * s_0 * w
        return wp, zp
```

Here's function to simulate the time series of wealth for in individual households.

```
@jit
def wealth_time_series(wdy, w_0, n):
    """
    Generate a single time series of length n for wealth given
    initial value w_0.

    The initial persistent state z_0 for each household is drawn from
    the stationary distribution of the AR(1) process.

        * wdy: an instance of WealthDynamics
        * w_0: scalar
        * n: int


    """
    z = wdy.z_mean + np.sqrt(wdy.z_var) * np.random.randn()
    w = np.empty(n)
    w[0] = w_0
    for t in range(n-1):
        w[t+1], z = wdy.update_states(w[t], z)
    return w
```

Now here's function to simulate a cross section of households forward in time.

Note the use of parallelization to speed up computation.

```
@jit(parallel=True)
def update_cross_section(wdy, w_distribution, shift_length=500):
    """
    Shifts a cross-section of household forward in time

    * wdy: an instance of WealthDynamics
    * w_distribution: array_like, represents current cross-section

    Takes a current distribution of wealth values as w_distribution
    and updates each w_t in w_distribution to w_{t+j}, where
    j = shift_length.

    Returns the new distribution.

    """
    new_distribution = np.empty_like(w_distribution)

    # Update each household
    for i in prange(len(new_distribution)):
        z = wdy.z_mean + np.sqrt(wdy.z_var) * np.random.randn()
        w = w_distribution[i]
        for t in range(shift_length-1):
            w, z = wdy.update_states(w, z)
        new_distribution[i] = w
    return new_distribution
```

Parallelization is very effective in the function above because the time path of each household can be calculated independently once the path for the aggregate state is known.

47.5 Applications

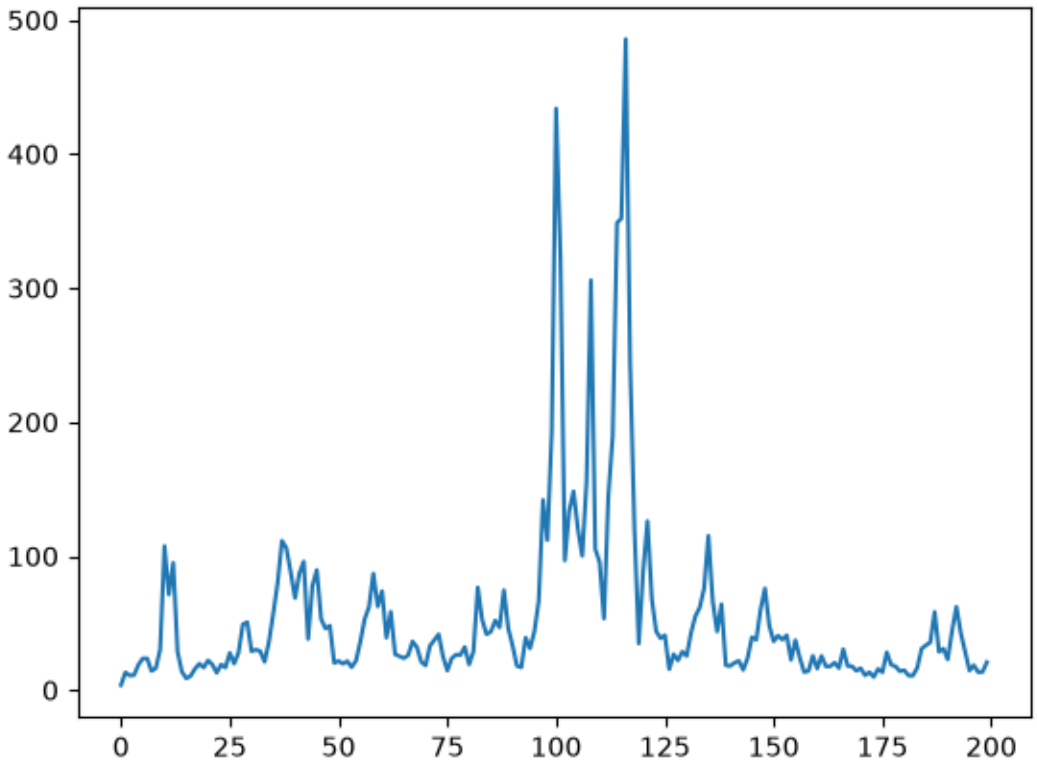
Let's try simulating the model at different parameter values and investigate the implications for the wealth distribution.

47.5.1 Time Series

Let's look at the wealth dynamics of an individual household.

```
wdy = WealthDynamics()
ts_length = 200
w = wealth_time_series(wdy, wdy.y_mean, ts_length)
```

```
fig, ax = plt.subplots()
ax.plot(w)
plt.show()
```



Notice the large spikes in wealth over time.

Such spikes are similar to what we observed in time series when *we studied Kesten processes*.

47.5.2 Inequality Measures

Let's look at how inequality varies with returns on financial assets.

The next function generates a cross section and then computes the Lorenz curve and Gini coefficient.

```
def generate_lorenz_and_gini(wdy, num_households=100_000, T=500):
    """
    Generate the Lorenz curve data and gini coefficient corresponding to a
    WealthDynamics mode by simulating num_households forward to time T.
    """
    ψ_0 = np.full(num_households, wdy.y_mean)
    z_0 = wdy.z_mean

    ψ_star = update_cross_section(wdy, ψ_0, shift_length=T)
    return qe.gini_coefficient(ψ_star), qe.lorenz_curve(ψ_star)
```

Now we investigate how the Lorenz curves associated with the wealth distribution change as return to savings varies.

The code below plots Lorenz curves for three different values of μ_r.

If you are running this yourself, note that it will take one or two minutes to execute.

This is unavoidable because we are executing a CPU intensive task.

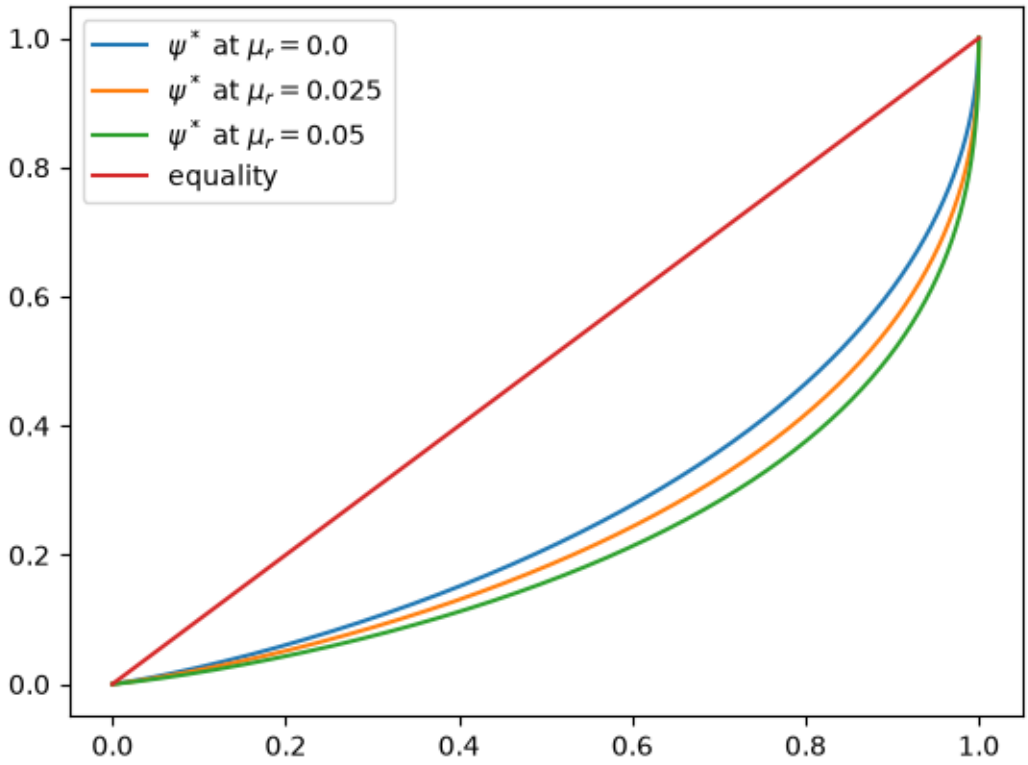
In fact the code, which is JIT compiled and parallelized, runs extremely fast relative to the number of computations.

```
%%time

fig, ax = plt.subplots()
μ_r_vals = (0.0, 0.025, 0.05)
gini_vals = []

for μ_r in μ_r_vals:
    wdy = WealthDynamics(μ_r=μ_r)
    gv, (f_vals, l_vals) = generate_lorenz_and_gini(wdy)
    ax.plot(f_vals, l_vals, label=fr'$\psi^*$ at $\mu_r = {μ_r:0.2}$')
    gini_vals.append(gv)

ax.plot(f_vals, f_vals, label='equality')
ax.legend(loc="upper left")
plt.show()
```




```
CPU times: user 39.5 s, sys: 35.8 ms, total: 39.6 s
Wall time: 5.66 s
```

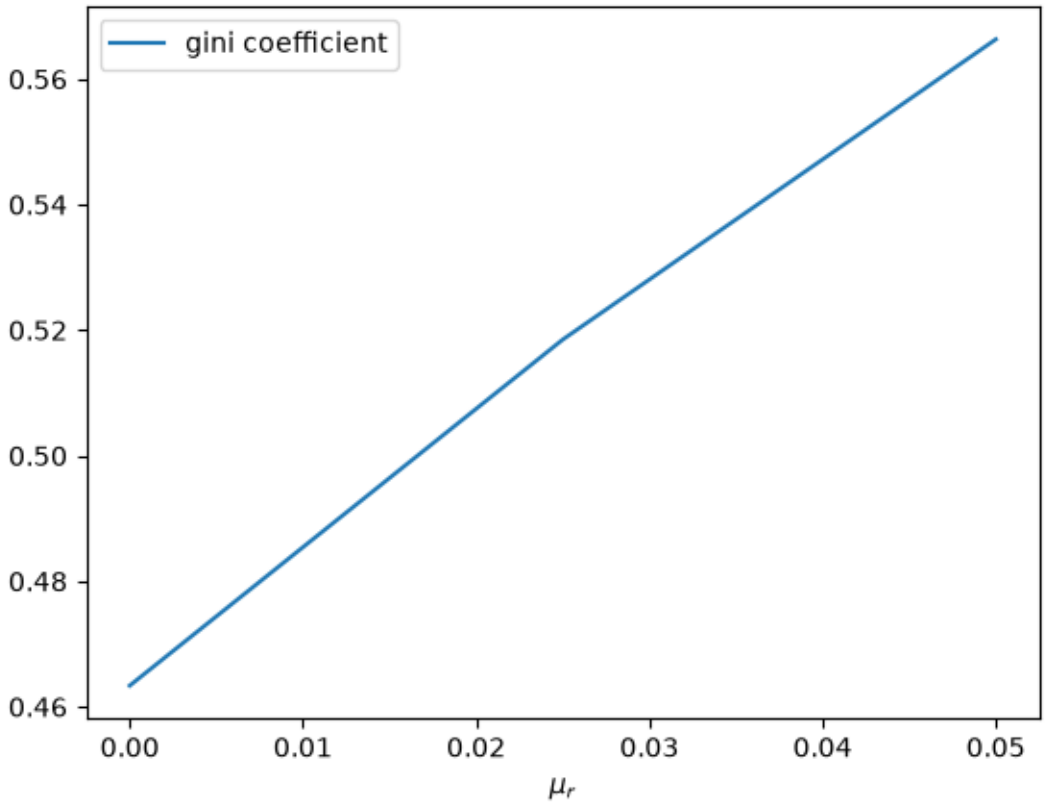
The Lorenz curve shifts downwards as returns on financial income rise, indicating a rise in inequality.

We will look at this again via the Gini coefficient immediately below, but first consider the following image of our system resources when the code above is executing:

Since the code is both efficiently JIT compiled and fully parallelized, it's close to impossible to make this sequence of tasks run faster without changing hardware.

Now let's check the Gini coefficient.

```
fig, ax = plt.subplots()
ax.plot(μ_r_vals, gini_vals, label='gini coefficient')
ax.set_xlabel(r"$\mu_r$")
ax.legend()
plt.show()
```



Once again, we see that inequality increases as returns on financial income rise.

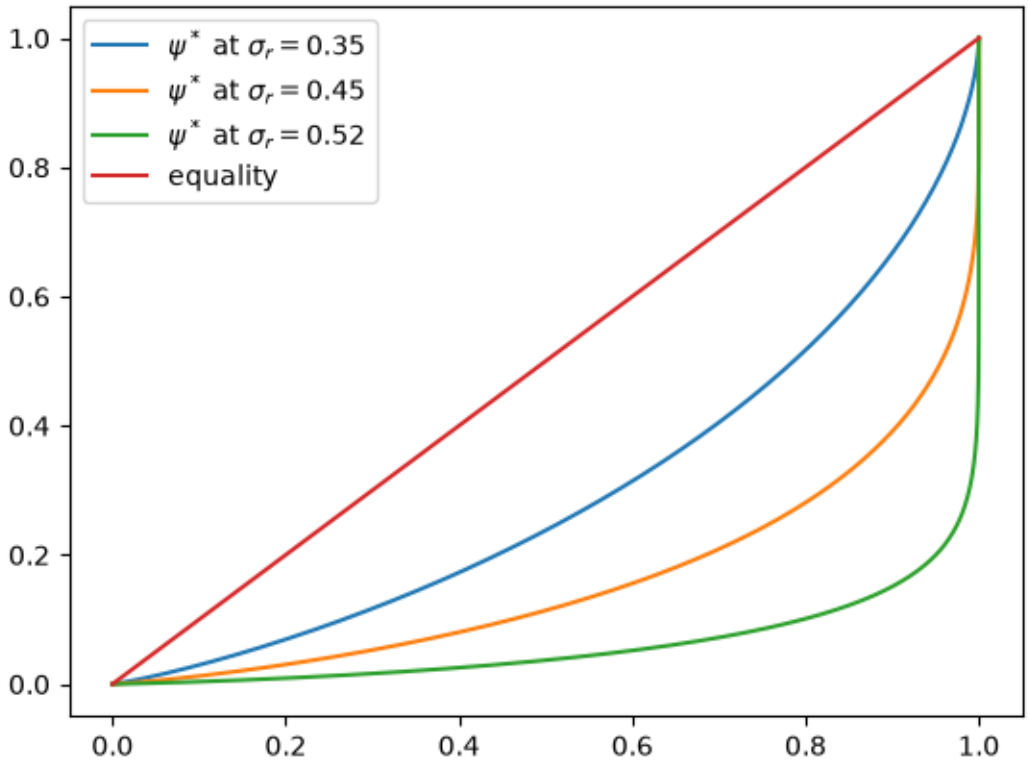
Let's finish this section by investigating what happens when we change the volatility term σ_r in financial returns.

```
%%time

fig, ax = plt.subplots()
σ_r_vals = (0.35, 0.45, 0.52)
gini_vals = []

for σ_r in σ_r_vals:
    wdy = WealthDynamics(σ_r=σ_r)
    gv, (f_vals, l_vals) = generate_lorenz_and_gini(wdy)
    ax.plot(f_vals, l_vals, label=fr'$\psi^*$ at $\sigma_r = {σ_r:0.2}$')
    gini_vals.append(gv)

ax.plot(f_vals, f_vals, label='equality')
ax.legend(loc="upper left")
plt.show()
```




```
CPU times: user 39 s, sys: 30 ms, total: 39 s
Wall time: 5.09 s
```

We see that greater volatility has the effect of increasing inequality in this model.

47.6 Exercises

Exercise 47.6.1

For a wealth or income distribution with Pareto tail, a higher tail index suggests lower inequality.

Indeed, it is possible to prove that the Gini coefficient of the Pareto distribution with tail index a is $1/(2a-1)$.

To the extent that you can, confirm this by simulation.

In particular, generate a plot of the Gini coefficient against the tail index using both the theoretical value just given and the value computed from a sample via `qe.gini_coefficient`.

For the values of the tail index, use `a_vals = np.linspace(1, 10, 25)`.

Use sample of size 1,000 for each a and the sampling method for generating Pareto draws employed in the discussion of Lorenz curves for the Pareto distribution.

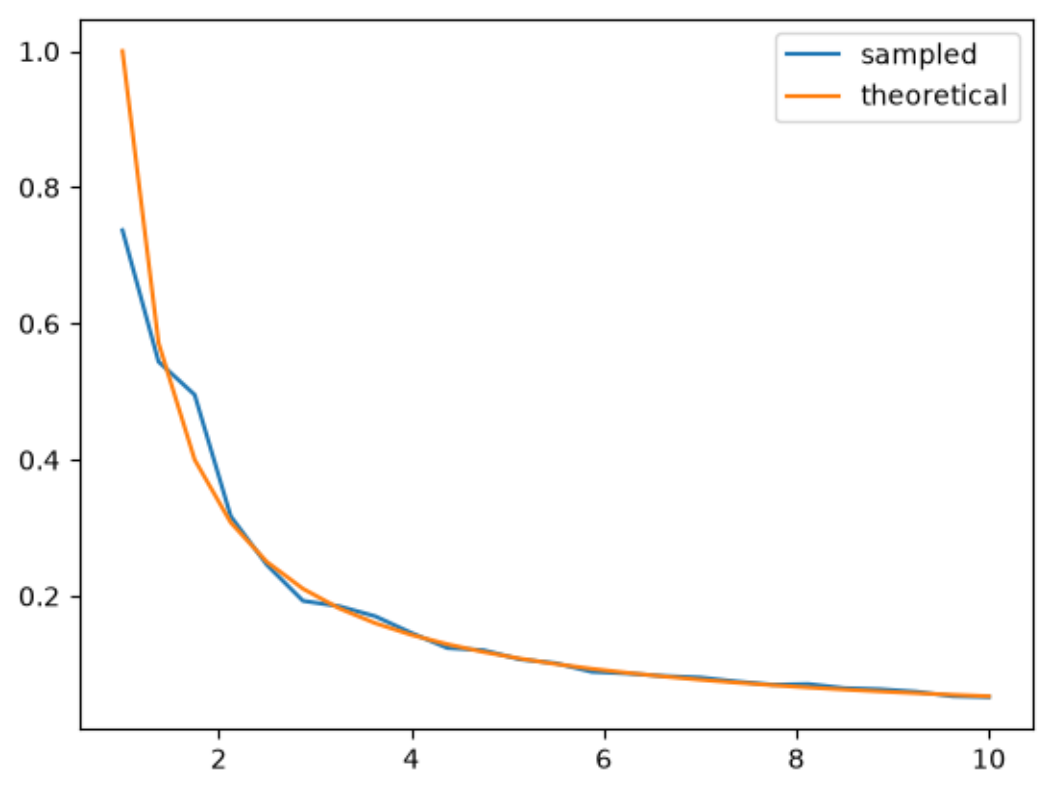
To the extent that you can, interpret the monotone relationship between the Gini index and a.

Solution

Here is one solution, which produces a good match between theory and simulation.

```
rng = np.random.default_rng()
a_vals = np.linspace(1, 10, 25)  # Pareto tail index
ginis = np.empty_like(a_vals)

n = 1000                         # size of each sample
fig, ax = plt.subplots()
for i, a in enumerate(a_vals):
    y = rng.uniform(size=n)**(-1/a)
    ginis[i] = qe.gini_coefficient(y)
ax.plot(a_vals, ginis, label='sampled')
ax.plot(a_vals, 1/(2*a_vals - 1), label='theoretical')
ax.legend()
plt.show()
```



In general, for a Pareto distribution, a higher tail index implies less weight in the right hand tail.

This means less extreme values for wealth and hence more equality.

More equality translates to a lower Gini index.

Exercise 47.6.2

The wealth process (47.1) is similar to a *Kesten process*.

This is because, according to (47.2), savings is constant for all wealth levels above $\hat{w}$.

When savings is constant, the wealth process has the same quasi-linear structure as a Kesten process, with multiplicative and additive shocks.

The Kesten–Goldie theorem tells us that Kesten processes have Pareto tails under a range of parameterizations.

The theorem does not directly apply here, since savings is not always constant and since the multiplicative and additive terms in (47.1) are not IID.

At the same time, given the similarities, perhaps Pareto tails will arise.

To test this, run a simulation that generates a cross-section of wealth and generate a rank-size plot.

If you like, you can use the function `rank_size` from the `quantecon` library (documentation here).

In viewing the plot, remember that Pareto tails generate a straight line. Is this what you see?

For sample size and initial conditions, use

```
num_households = 250_000
T = 500                                          # shift forward T periods
ψ_0 = np.full(num_households, wdy.y_mean)    # initial distribution
z_0 = wdy.z_mean
```

Solution

First let's generate the distribution:

```
num_households = 250_000
T = 500  # how far to shift forward in time
wdy = WealthDynamics()
ψ_0 = np.full(num_households, wdy.y_mean)
z_0 = wdy.z_mean

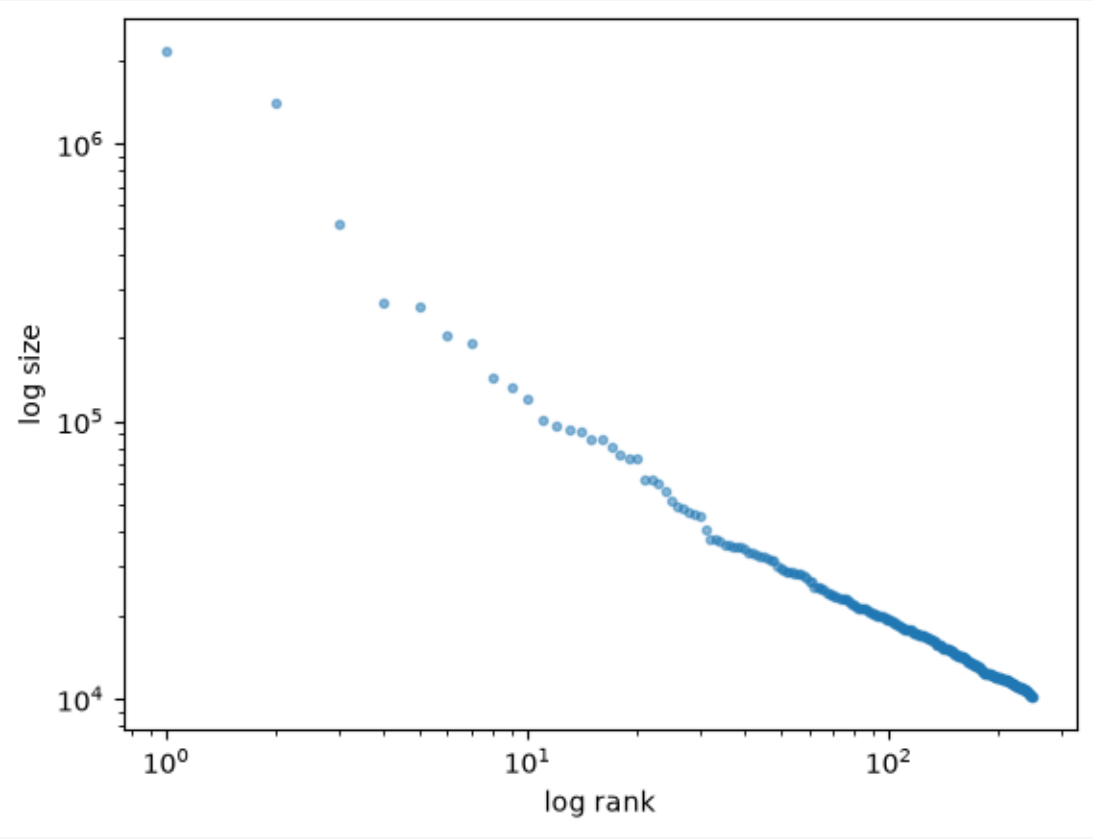
ψ_star = update_cross_section(wdy, ψ_0, shift_length=T)
```

Now let's see the rank-size plot:

```
fig, ax = plt.subplots()

rank_data, size_data = qe.rank_size(ψ_star, c=0.001)
ax.loglog(rank_data, size_data, 'o', markersize=3.0, alpha=0.5)
ax.set_xlabel("log rank")
ax.set_ylabel("log size")

plt.show()
```


10^6
10^5
10^4
log size
10^0
10^1
10^2
log rank

CHAPTER

FORTYEIGHT

A FIRST LOOK AT THE KALMAN FILTER

Contents

- *A First Look at the Kalman Filter*
  - *Overview*
  - *The basic idea*
  - *Convergence*
  - *Implementation*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

48.1 Overview

This lecture provides a simple and intuitive introduction to the Kalman filter

It is aimed at readers who either

- have heard of the Kalman filter but don't know how it works, or
- know the Kalman filter equations, but don't know where they come from

Subsequent lectures use the same recursive logic in more applied and more econometric settings.

See *Another Look at the Kalman Filter* for an economic application in which a firm infers a worker's hidden human capital and effort.

See *The Kalman Filter and Vector Autoregressions* for a derivation of the innovations representation and its connection to vector autoregressions.

Priors, Escapes, and Learning Cycles gives a government that tracks drifting regression coefficients with this filter, and *Drifts and Volatilities* gives an econometrician who does the same for a vector autoregression.

For additional (more advanced) reading on the Kalman filter, see

- [Ljungqvist and Sargent, 2018], section 2.7
- [Anderson and Moore, 2005]

The second reference presents a comprehensive treatment of the Kalman filter.

Required knowledge: Familiarity with matrix manipulations, multivariate normal distributions, covariance matrices, etc.

We'll need the following imports:

```
import matplotlib.pyplot as plt
from scipy import linalg
import numpy as np
from quantecon import Kalman, LinearStateSpace
from scipy.stats import norm, multivariate_normal
from scipy.integrate import quad
from scipy.linalg import eigvals
```

48.2 The basic idea

The Kalman filter has many applications in economics, but for now let's pretend that we are rocket scientists.

A missile has been launched from a hostile country and our mission is to track it.

Let $X_t \in \mathbb{R}^2$ denote the current location of the missile—a pair indicating latitude-longitude coordinates on a map.

At the present moment in time, the location X_t is unknown, but we do have some beliefs about it.

We could certainly produce a point prediction.

For example, it could mark a point on the globe somewhere in northern Mongolia.

But the fact is that we are uncertain.

And the President wants to know: what is the probability that the missile is within 500km of Manhattan?

A point prediction doesn't address that question.

Hence it's best if we can express our current understanding via a bivariate probability density p.

- Now $\int_E p(x)dx$ indicates the probability that the missile is in region E.

We will call p our **prior** for the random variable X.

To keep things tractable, we assume for now that our prior is Gaussian.

In particular, we take

$$p = N(\mu, \Sigma) \tag{48.1}$$

where μ is the (vector) mean of the distribution—a natural point prediction—and Σ is a 2×2 covariance matrix.

In our simulations, we will suppose that

$$\mu = \begin{pmatrix} 0.2 \\ -0.2 \end{pmatrix}, \qquad \Sigma = \begin{pmatrix} 0.4 & 0.3 \\ 0.3 & 0.45 \end{pmatrix} \tag{48.2}$$

This density p is shown below as a contour map, with the center of the red ellipse being equal to μ.

```
# Set up the Gaussian prior density p
Σ = np.array([[0.4, 0.3],
              [0.3, 0.45]])
μ = np.array([[0.2],
              [-0.2]])
# Define the matrices G and R from the measurement equation Y = G X + v
```

(continues on next page)

(continued from previous page)

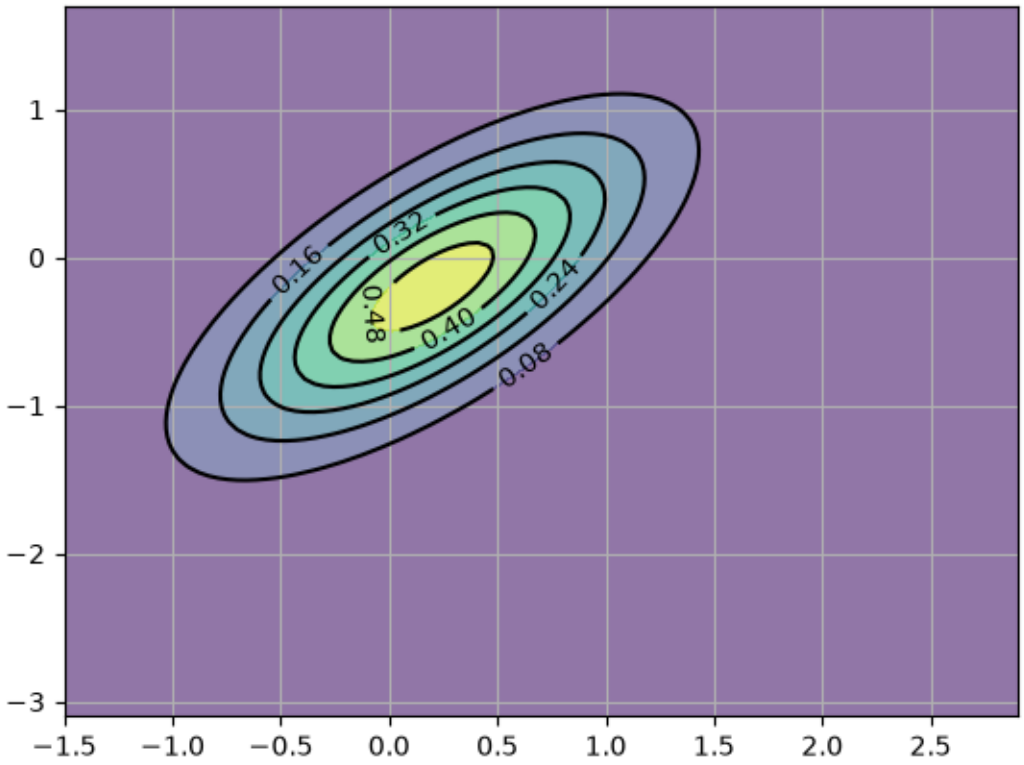
```
G = np.array([[1, 0],
              [0, 1]])
R = 0.5 * Σ
# The matrices A and Q
A = np.array([[1.2, 0],
              [0, -0.2]])
Q = 0.3 * Σ
# The observed value of y
y = np.array([[2.3],
              [-1.9]])

# Set up grid for plotting
x_grid = np.linspace(-1.5, 2.9, 100)
y_grid = np.linspace(-3.1, 1.7, 100)
X, Y = np.meshgrid(x_grid, y_grid)

def gen_gaussian_plot_vals(μ, C):
    "Z values for plotting the bivariate Gaussian N(μ, C)"
    pos = np.dstack((X, Y))
    return multivariate_normal(μ.ravel(), C).pdf(pos)

# Plot the figure

fig, ax = plt.subplots()
ax.grid()
Z = gen_gaussian_plot_vals(μ, Σ)
ax.contourf(X, Y, Z, 6, alpha=0.6, cmap="viridis")
cs = ax.contour(X, Y, Z, 6, colors="black")
ax.clabel(cs, inline=1, fontsize=10)
plt.show()
```



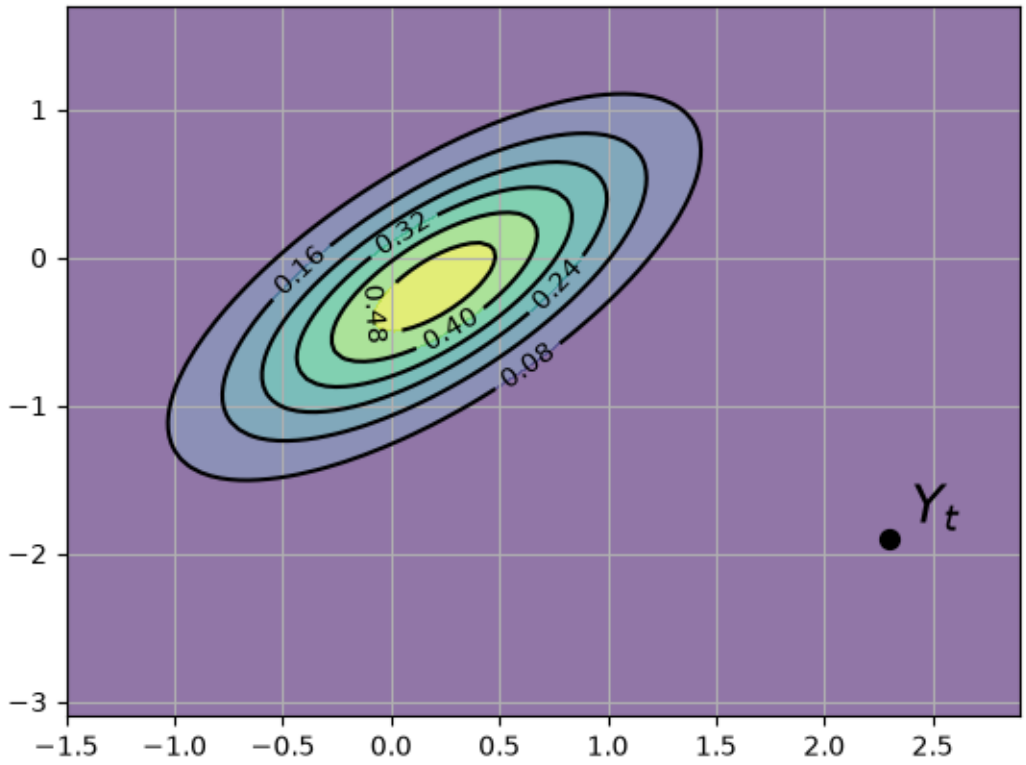
48.2.1 The filtering step

We are now presented with some good news and some bad news.

The good news is that the missile has been located by our sensors, which report that the current location is $Y_t = (2.3, -1.9)$.

The next figure shows the original prior p and the new reported signal Y_t

```
fig, ax = plt.subplots()
ax.grid()
Z = gen_gaussian_plot_vals(μ, Σ)
ax.contourf(X, Y, Z, 6, alpha=0.6, cmap="viridis")
cs = ax.contour(X, Y, Z, 6, colors="black")
ax.clabel(cs, inline=1, fontsize=10)
y_1, y_2 = y[0].item(), y[1].item()
ax.scatter(y_1, y_2, marker="o", s=50, color="black", zorder=3)
ax.text(y_1 + 0.1, y_2 + 0.1, "$Y_t$", fontsize=20, color="black")
plt.show()
```



The bad news is that our sensors are imprecise.

The sensor report is a noisy signal distorted by measurement error.

In particular, we should interpret the output of our sensor not as $Y_t = X_t$, but rather as

$$Y_t = GX_t + v_t, \quad \text{where} \quad v_t \sim N(0, R) \tag{48.3}$$

Here G and R are 2×2 matrices, with R being symmetric and positive definite.

We assume that

- G and R are known
- the noise term v_t is unobservable and independent of X_t

How then should we combine our prior $X_t \sim N(\mu, \Sigma)$ and this new information Y_t to improve our understanding of the location of the missile?

As you may have guessed, the answer is to use Bayes' theorem.

It tells us how to update the prior density $p(x)$ for X_t to the posterior density $p(x \mid y)$ after observing Y_t:

$$p(x \mid Y_t) = \frac{p(Y_t \mid x)\, p(x)}{p(Y_t)}$$

where $p(Y_t) = \int p(Y_t \mid x)\, p(x) dx$.

In solving for $p(x \mid Y_t)$, we observe that

- $p(x)$ is the prior density $N(\mu, \Sigma)$.
- $p(Y_t \mid x)$ is the conditional density of Y_t given $X_t = x$.
- In view of (48.3), this conditional density is $N(Gx, R)$.

Due to our linear Gaussian framework, the updated density turns out to be Gaussian as well.

In particular, the solution is known to be

$$p(x \mid Y_t) = N(\mu^F, \Sigma^F)$$

where

$$\mu^F := \mu + \Sigma G^\top (G \Sigma G^\top + R)^{-1} (y - G\mu) \tag{48.4}$$

and

$$\Sigma^F := \Sigma - \Sigma G^\top (G \Sigma G^\top + R)^{-1} G \Sigma \tag{48.5}$$

Note

A proof can be found in [Bishop, 2006].

To get from his expressions to the ones used above, you will also need to apply the Woodbury matrix identity.

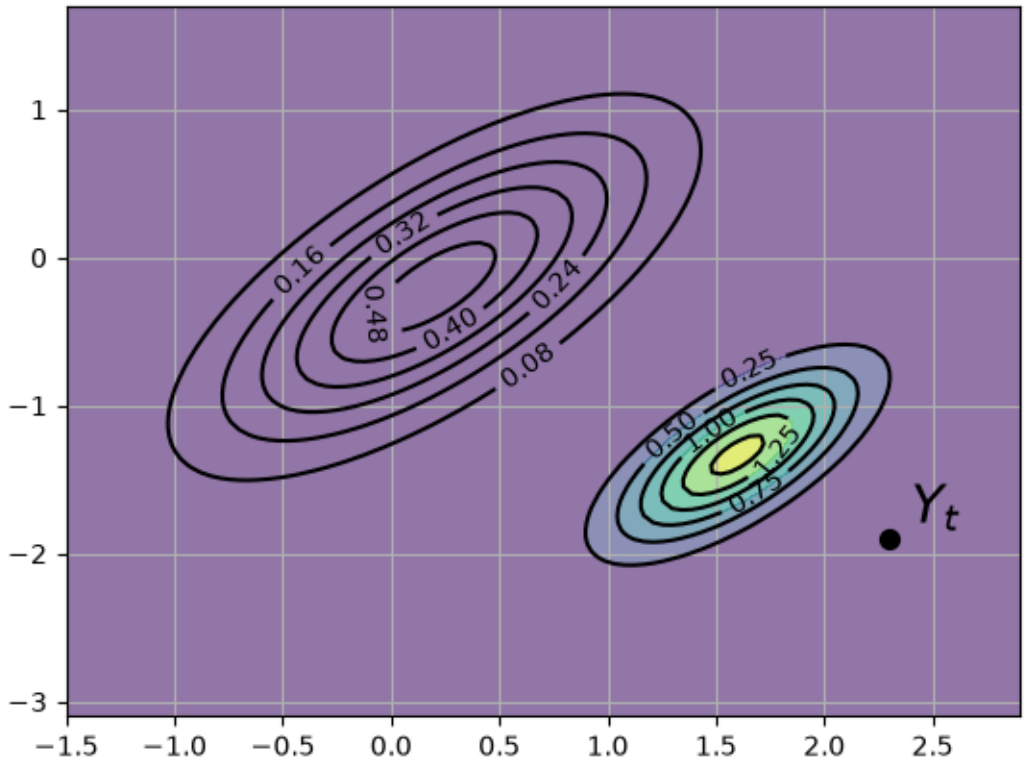
Here $\Sigma G^\top (G \Sigma G^\top + R)^{-1}$ is the matrix of population regression coefficients of the hidden state deviation $X_t - \mu$ on the *signal surprise* $Y_t - G\mu$.

This new density $p(x \mid Y_t) = N(\mu^F, \Sigma^F)$ is shown in the next figure via contour lines and the color map.

The original density is left in as contour lines for comparison

```
fig, ax = plt.subplots()
ax.grid()

Z = gen_gaussian_plot_vals(μ, Σ)
cs1 = ax.contour(X, Y, Z, 6, colors="black")
ax.clabel(cs1, inline=1, fontsize=10)
M = Σ @ G.T @ linalg.inv(G @ Σ @ G.T + R)
μ_F = μ + M @ (y - G @ μ)
Σ_F = Σ - M @ G @ Σ
new_Z = gen_gaussian_plot_vals(μ_F, Σ_F)
cs2 = ax.contour(X, Y, new_Z, 6, colors="black")
ax.clabel(cs2, inline=1, fontsize=10)
ax.contourf(X, Y, new_Z, 6, alpha=0.6, cmap="viridis")
y_1, y_2 = y[0].item(), y[1].item()
ax.scatter(y_1, y_2, marker="o", s=50, color="black", zorder=3)
ax.text(y_1 + 0.1, y_2 + 0.1, "$Y_t$", fontsize=20, color="black")
plt.show()
```



Our new density twists the prior $p(x)$ in a direction determined by the new information $Y_t - G\mu$.

In generating the figure, we set G to the identity matrix and $R = 0.5\Sigma$ for Σ defined in (48.2).

48.2.2 The forecast step

What have we achieved so far?

We have obtained probabilities for the current location of the state (missile) given prior and current information.

This is called "filtering" rather than forecasting because we are filtering out noise rather than looking into the future.

The posterior $p(x \mid Y_t) = N(\mu^F, \Sigma^F)$ is called the **filtering distribution** for X_t after observing Y_t

But now let's suppose that we are given another task: to predict the location of the missile after one unit of time (whatever that may be) has elapsed.

To do this we need a model of how the state evolves.

Let's suppose that we have one, and that it's linear and Gaussian.

In particular,

$$X_{t+1} = AX_t + W_{t+1}, \quad \text{where} \quad W_t \sim N(0, Q) \tag{48.6}$$

Our aim is to combine this law of motion and our current filtering distribution $N(\mu^F, \Sigma^F)$ to come up with a new **predictive** distribution for the location in one unit of time.

In view of (48.6), all we have to do is introduce a random vector $X^F \sim N(\mu^F, \Sigma^F)$ and work out the distribution of $AX^F + W$ where W is independent of X^F and has distribution $N(0, Q)$.

Since linear combinations of Gaussians are Gaussian, $AX^F + W$ is Gaussian.

Standard calculations and the expressions in (48.4)–(48.5) tell us that

$$\begin{aligned}\mathbb{E}[AX^F + W] &= A\mathbb{E}[X^F] + \mathbb{E}[W] \\ &= A\mu^F \\ &= A\mu + A\Sigma G^\top (G\Sigma G^\top + R)^{-1}(Y_t - G\mu)\end{aligned}$$

and

$$\begin{aligned}\text{Var}[AX^F + W] &= A\,\text{Var}[X^F]A^\top + Q \\ &= A\Sigma^F A^\top + Q \\ &= A\Sigma A^\top + Q - A\Sigma G^\top (G\Sigma G^\top + R)^{-1} G\Sigma A^\top\end{aligned}$$

The matrix $A\Sigma G^\top (G\Sigma G^\top + R)^{-1}$ is often written as K_Σ and called the **Kalman gain**.

- The subscript Σ has been added to remind us that K_Σ depends on Σ, but not Y_t or μ.

Using this notation, we can summarize our results as follows.

Our updated prediction is the density $N(\mu_{\text{new}}, \Sigma_{\text{new}})$ where

$$\begin{aligned}\mu_{\text{new}} &:= A\mu + K_\Sigma (y - G\mu) \\ \Sigma_{\text{new}} &:= A\Sigma A^\top - K_\Sigma G\Sigma A^\top + Q\end{aligned}$$

- The density $p_{\text{new}}(x) = N(\mu_{\text{new}}, \Sigma_{\text{new}})$ is called the **predictive distribution**

The predictive distribution is the new density shown in the following figure, where the update has used parameters.

$$A = \begin{pmatrix} 1.2 & 0.0 \\ 0.0 & -0.2 \end{pmatrix}, \qquad Q = 0.3\Sigma$$

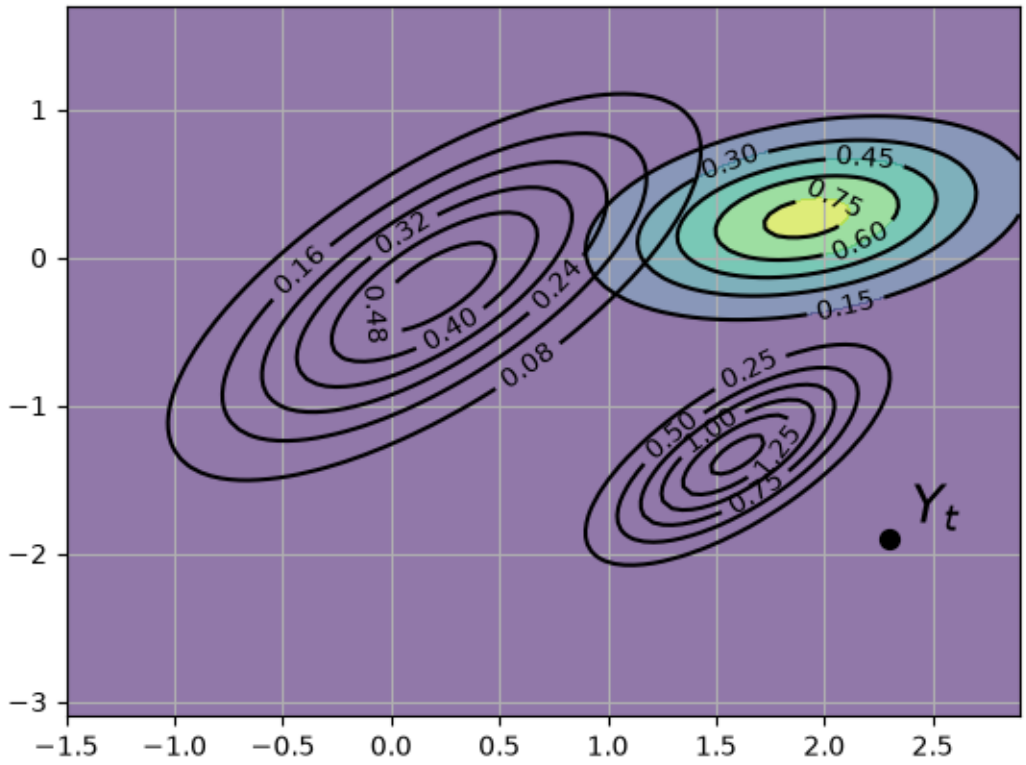
```
fig, ax = plt.subplots()
ax.grid()

# Density 1
Z = gen_gaussian_plot_vals(μ, Σ)
cs1 = ax.contour(X, Y, Z, 6, colors="black")
ax.clabel(cs1, inline=1, fontsize=10)

# Density 2
M = Σ @ G.T @ linalg.inv(G @ Σ @ G.T + R)
μ_F = μ + M @ (y - G @ μ)
Σ_F = Σ - M @ G @ Σ
Z_F = gen_gaussian_plot_vals(μ_F, Σ_F)
cs2 = ax.contour(X, Y, Z_F, 6, colors="black")
ax.clabel(cs2, inline=1, fontsize=10)

# Density 3
new_μ = A @ μ_F
new_Σ = A @ Σ_F @ A.T + Q
new_Z = gen_gaussian_plot_vals(new_μ, new_Σ)
cs3 = ax.contour(X, Y, new_Z, 6, colors="black")
ax.clabel(cs3, inline=1, fontsize=10)
ax.contourf(X, Y, new_Z, 6, alpha=0.6, cmap="viridis")
y_1, y_2 = y[0].item(), y[1].item()
ax.scatter(y_1, y_2, marker="o", s=50, color="black", zorder=3)
ax.text(y_1 + 0.1, y_2 + 0.1, "$Y_t$", fontsize=20, color="black")

plt.show()
```



48.2.3 The recursive procedure

Let's look back at what we've done.

We started the current period with a prior density $p_t(x)$ for the hidden state X_t.

We then observed the signal Y_t and updated the prior density to the filtering density $p_t(x \mid Y_t)$.

Finally, we used the law of motion (48.6) for $\{X_t\}$ to update to the predictive density $p_{t+1}(x)$ for X_{t+1}.

If we now step into the next period, we are ready to go round again, taking $p_{t+1}(x)$ as the current prior density and reading in the new observation Y_{t+1}.

Using this time-indexed notation, the full recursive procedure is:

1. Start the current period with prior density $p_t(x) = N(\mu_t, \Sigma_t)$ for X_t.
2. Observe current signal $Y_t = y_t$.
3. Compute the filtering density $p_t(x \mid y_t) = N(\mu_t^F, \Sigma_t^F)$ from $p_t(x)$ and y_t, applying Bayes rule and the conditional distribution (48.3).
4. Compute the predictive density $p_{t+1}(x) = N(\mu_{t+1}, \Sigma_{t+1})$ for X_{t+1} from the filtering density and (48.6).
5. Increment t by one and go to step 1.

Repeating (48.7), the dynamics for μ_t and Σ_t are as follows

$$\mu_{t+1} = A\mu_t + K_{\Sigma_t}(y_t - G\mu_t)$$
$$\Sigma_{t+1} = A\Sigma_t A^\top - K_{\Sigma_t} G \Sigma_t A^\top + Q$$

These are the standard dynamic equations for the Kalman filter (see, for example, [Ljungqvist and Sargent, 2018], page 58).

Note

Here μ_t is the filter's prediction of the hidden state X_t.

In much of the Kalman filter literature it is written $\hat{x}_t$, emphasizing that it is an estimate of X_t.

48.3 Convergence

The matrix Σ_t is a measure of the uncertainty of our prediction μ_t of X_t.

Apart from special cases, this uncertainty will never be fully resolved, regardless of how much time elapses.

One reason is that our prediction μ_t is made based on information available at $t-1$, not t.

Even if we knew the precise realized value $X_{t-1} = x_{t-1}$ (which we don't), the transition equation (48.6) implies that $X_t = Ax_{t-1} + W_t$.

Since the shock W_t is not observable at $t-1$, any time $t-1$ prediction of X_t will incur some error (unless W_t is degenerate).

However, it is certainly possible that Σ_t converges to a constant matrix as $t \to \infty$.

To study this topic, let's expand the second equation in (48.7):

$$\Sigma_{t+1} = A\Sigma_t A^\top - A\Sigma_t G^\top (G\Sigma_t G^\top + R)^{-1} G\Sigma_t A^\top + Q \tag{48.7}$$

This is a nonlinear difference equation in Σ_t.

A fixed point of (48.7) is a constant matrix Σ such that

$$\Sigma = A\Sigma A^\top - A\Sigma G^\top (G\Sigma G^\top + R)^{-1} G\Sigma A^\top + Q \tag{48.8}$$

Equation (48.7) is known as a discrete-time Riccati difference equation.

Equation (48.8) is known as a discrete-time algebraic Riccati equation.

Conditions under which a fixed point exists and the sequence $\{\Sigma_t\}$ converges to it are discussed in [Anderson *et al.*, 1996] and [Anderson and Moore, 2005], chapter 4.

A sufficient (but not necessary) condition is that all the eigenvalues λ_i of A satisfy $|\lambda_i| < 1$.

See, for example, [Anderson and Moore, 2005], p. 77.

(This strong condition assures that the unconditional distribution of X_t converges as $t \to \infty$.)

In this case, for any symmetric nonnegative definite initial choice of Σ_0, the sequence $\{\Sigma_t\}$ in (48.7) converges to a nonnegative symmetric matrix Σ that solves (48.8).

48.4 Implementation

The class `Kalman` from the QuantEcon.py package implements the Kalman filter

- Instance data consists of:
  - the moments (μ_t, Σ_t) of the current prior, stored as the attributes `x_hat` and `Sigma` (the mean μ_t is named `x_hat` because it is also written $\hat{x}_t$ in much of the literature).
  - An instance of the LinearStateSpace class from QuantEcon.py.

The latter represents a linear state space model of the form

$$X_{t+1} = AX_t + Cw_{t+1}$$
$$Y_t = GX_t + Hv_t$$

where X_t and Y_t denote random variables, and the shocks w_t and v_t are IID standard normals.

To connect this with the notation of this lecture we set

$$Q := CC^\top \quad \text{and} \quad R := HH^\top$$

- The class `Kalman` from the QuantEcon.py package has a number of methods, some that we will wait to use until we study more advanced applications in subsequent lectures.
- Methods pertinent for this lecture are:
  - `prior_to_filtered`, which updates (μ_t, Σ_t) to (μ_t^F, Σ_t^F)
  - `filtered_to_forecast`, which updates the filtering distribution to the predictive distribution – which becomes the new prior $(\mu_{t+1}, \Sigma_{t+1})$
  - `update`, which combines the last two methods
  - a `stationary_values`, which computes the solution to (48.8) and the corresponding (stationary) Kalman gain

You can view the program on GitHub.

48.5 Exercises

Exercise 48.5.1

Consider the following simple application of the Kalman filter, loosely based on [Ljungqvist and Sargent, 2018], section 2.9.2.

Suppose that

- all variables are scalars
- the hidden state $\{X_t\}$ is in fact constant, equal to some $\theta \in \mathbb{R}$ unknown to the modeler

State dynamics are therefore given by (48.6) with $A = 1$, $Q = 0$ and $X_0 = \theta$.

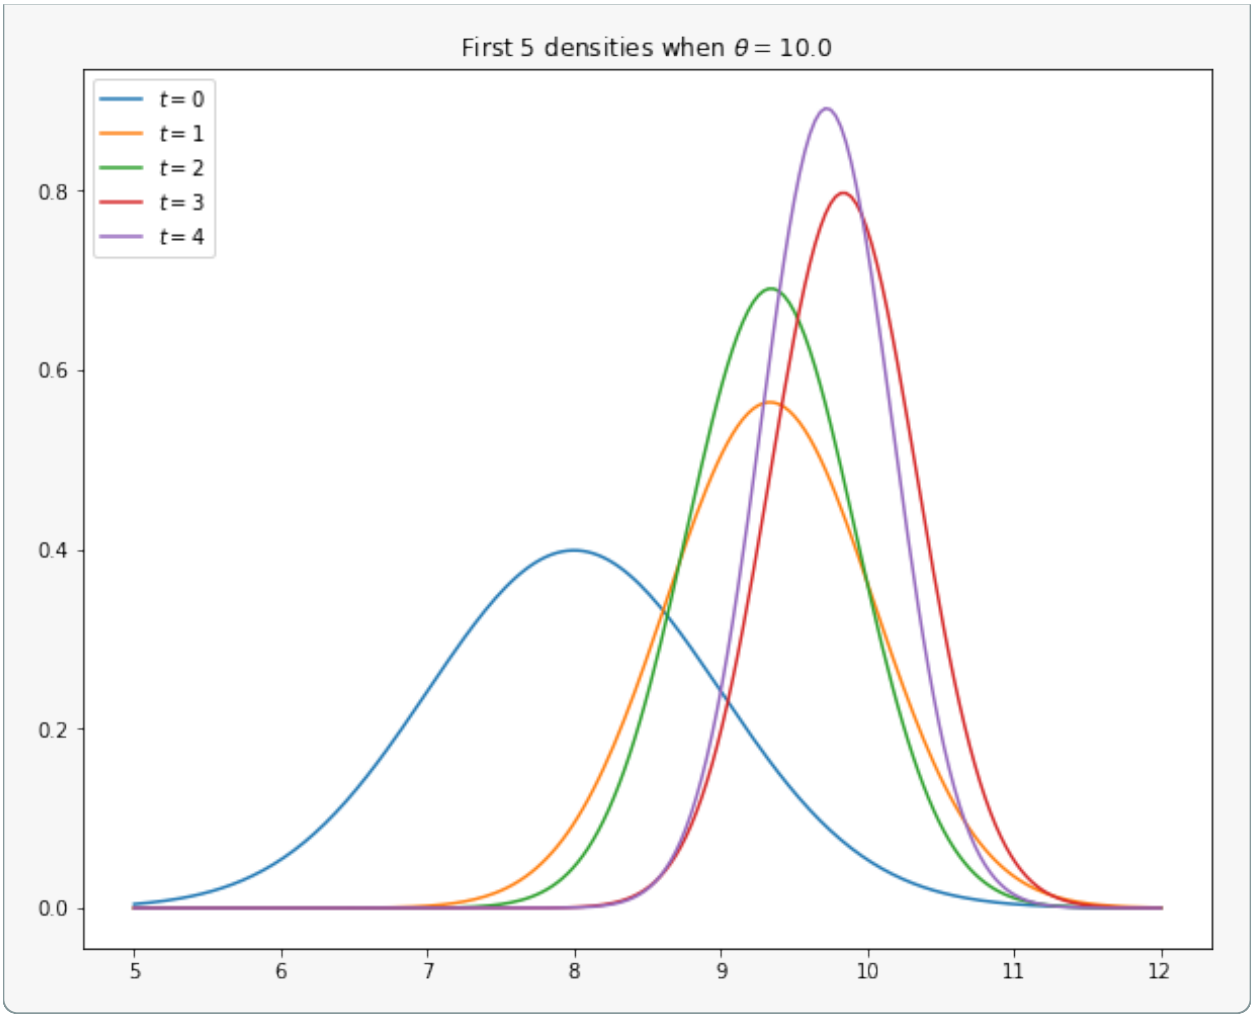
The measurement equation is $Y_t = \theta + v_t$ where v_t is $N(0, 1)$ and IID.

The task of this exercise is to simulate the model and, using the code from `kalman.py`, plot the first five predictive densities $p_t(x) = N(\mu_t, \Sigma_t)$ for X_t.

As shown in [Ljungqvist and Sargent, 2018], sections 2.9.1–2.9.2, these distributions asymptotically put all mass on the unknown value θ.

In the simulation, take $\theta = 10$, $\mu_0 = 8$ and $\Sigma_0 = 1$.

Your figure should – modulo randomness – look something like this



Solution

Here is one solution:

```
# Parameters
θ = 10  # Constant value of state X_t
A, C, G, H = 1, 0, 1, 1
ss = LinearStateSpace(A, C, G, H, mu_0=θ)

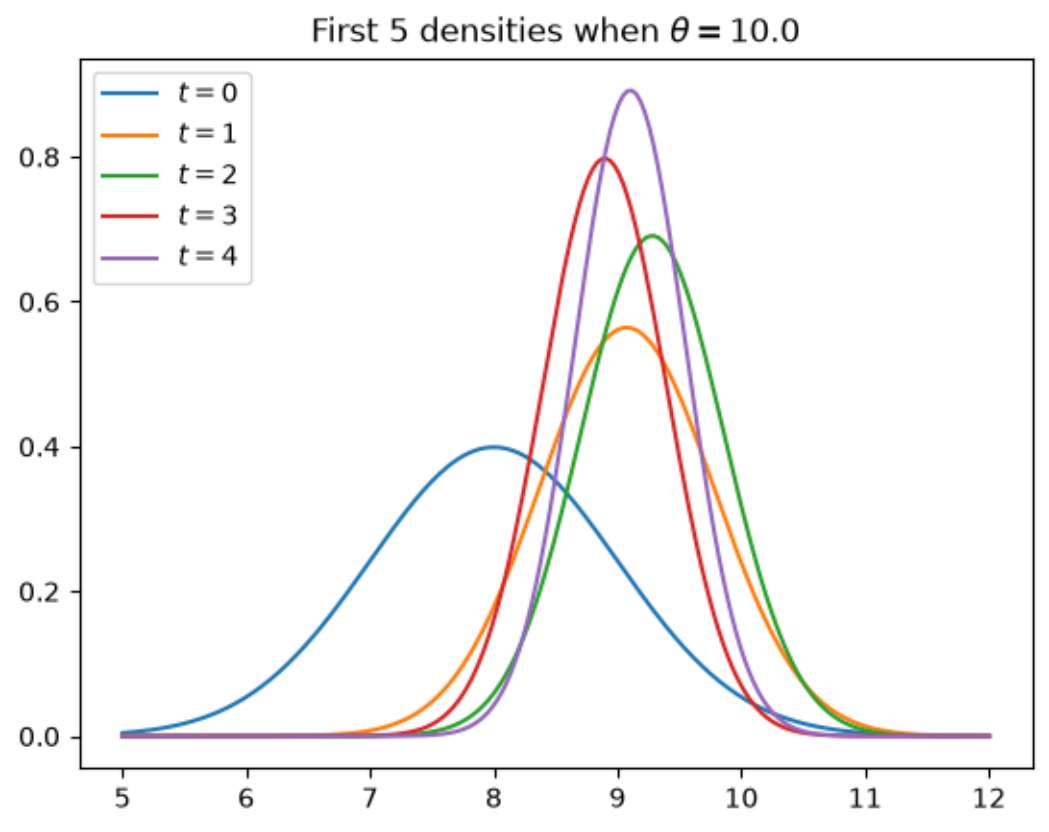
# Set prior, initialize kalman filter
μ_0, Σ_0 = 8, 1
kalman = Kalman(ss, μ_0, Σ_0)

# Draw observations of y from state space model
N = 5
x, y = ss.simulate(N)
y = y.flatten()

# Set up plot
fig, ax = plt.subplots()
xgrid = np.linspace(θ - 5, θ + 2, 200)
```

```
for i in range(N):
    # Record the current predicted mean and variance
    m, v = kalman.x_hat.item(), kalman.Sigma.item()
    # Plot, update filter
    ax.plot(xgrid, norm.pdf(xgrid, loc=m, scale=np.sqrt(v)), label=f'$t={i}$')
    kalman.update(y[i])

ax.set_title(f'First {N} densities when $\\theta = {θ:.1f}$')
ax.legend(loc='upper left')
plt.show()
```



Exercise 48.5.2

The preceding figure gives some support to the idea that probability mass converges to θ.

To get a better idea, choose a small $\epsilon > 0$ and calculate the error

$$z_t := 1 - \int_{\theta-\epsilon}^{\theta+\epsilon} p_t(x)dx$$

for $t = 0, 1, 2, \ldots, T$.

Plot z_t against t, setting $\epsilon = 0.1$ and $T = 600$.

Solution

Here is one solution:

```
ϵ = 0.1
θ = 10  # Constant value of state X_t
A, C, G, H = 1, 0, 1, 1
ss = LinearStateSpace(A, C, G, H, mu_0=θ)

μ_0, Σ_0 = 8, 1
kalman = Kalman(ss, μ_0, Σ_0)

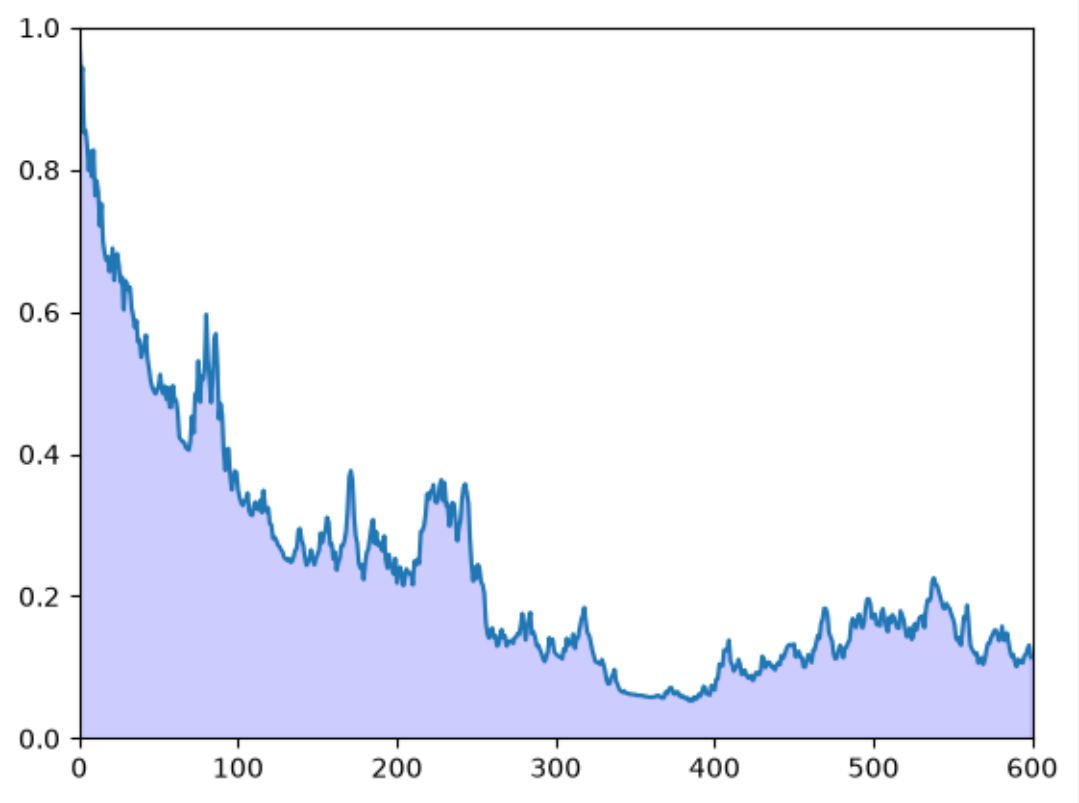
T = 600
z = np.empty(T)
x, y = ss.simulate(T)
y = y.flatten()

for t in range(T):
    # Record the current predicted mean and variance and plot their densities
    m, v = kalman.x_hat.item(), kalman.Sigma.item()

    f = lambda x: norm.pdf(x, loc=m, scale=np.sqrt(v))
    integral, error = quad(f, θ - ϵ, θ + ϵ)
    z[t] = 1 - integral

    kalman.update(y[t])

fig, ax = plt.subplots()
ax.set_ylim(0, 1)
ax.set_xlim(0, T)
ax.plot(range(T), z)
ax.fill_between(range(T), np.zeros(T), z, color="blue", alpha=0.2)
plt.show()
```



Exercise 48.5.3

As discussed *above*, if the shock sequence $\{W_t\}$ is not degenerate, then it is not in general possible to predict X_t without error at time $t-1$ (and this would be the case even if we could observe X_{t-1}).

Let's now compare the prediction μ_t made by the Kalman filter against a competitor who **is** allowed to observe X_{t-1}.

This competitor will use the conditional expectation $\mathbb{E}[X_t \mid X_{t-1}]$, which in this case is AX_{t-1}.

The conditional expectation is known to be the optimal prediction method in terms of minimizing mean squared error.

(More precisely, the minimizer of $\mathbb{E} \|X_t - g(X_{t-1})\|^2$ with respect to g is $g^*(X_{t-1}) := \mathbb{E}[X_t \mid X_{t-1}]$)

Thus we are comparing the Kalman filter against a competitor who has more information (in the sense of being able to observe the latent state) and behaves optimally in terms of minimizing squared error.

Our horse race will be assessed in terms of realized squared error.

In particular, your task is to generate a graph plotting simulated realizations of both $\|X_t - AX_{t-1}\|^2$ and $\|X_t - \mu_t\|^2$ against t for $t = 1, \ldots, 49$.

In the code below, `x[:, t]` is the realized value of X_t along the simulated path.

For the parameters, set $G = I$, $R = 0.5I$ and $Q = 0.3I$, where I is the 2×2 identity.

Set

$$A = \begin{pmatrix} 0.5 & 0.4 \\ 0.6 & 0.3 \end{pmatrix}$$

To initialize the prior density, set

$$\Sigma_0 = \begin{pmatrix} 0.9 & 0.3 \\ 0.3 & 0.9 \end{pmatrix}$$

and $\mu_0 = (8, 8)$.

Finally, set the realized initial state to $x_0 = (0, 0)$.

Solution

Here is one solution:

```
# Define A, C, G, H
G = np.identity(2)
H = np.sqrt(0.5) * np.identity(2)

A = [[0.5, 0.4],
     [0.6, 0.3]]
C = np.sqrt(0.3) * np.identity(2)

# Set up state space mode, initial value x_0 set to zero
ss = LinearStateSpace(A, C, G, H, mu_0 = np.zeros(2))

# Define the prior density
Σ = [[0.9, 0.3],
     [0.3, 0.9]]
Σ = np.array(Σ)
μ = np.array([8, 8])

# Initialize the Kalman filter
kn = Kalman(ss, μ, Σ)

# Print eigenvalues of A
print("Eigenvalues of A:")
print(eigvals(A))

# Print stationary Σ
S, K = kn.stationary_values()
print("Stationary prediction error variance:")
print(S)

# Generate the plot
T = 50
x, y = ss.simulate(T)

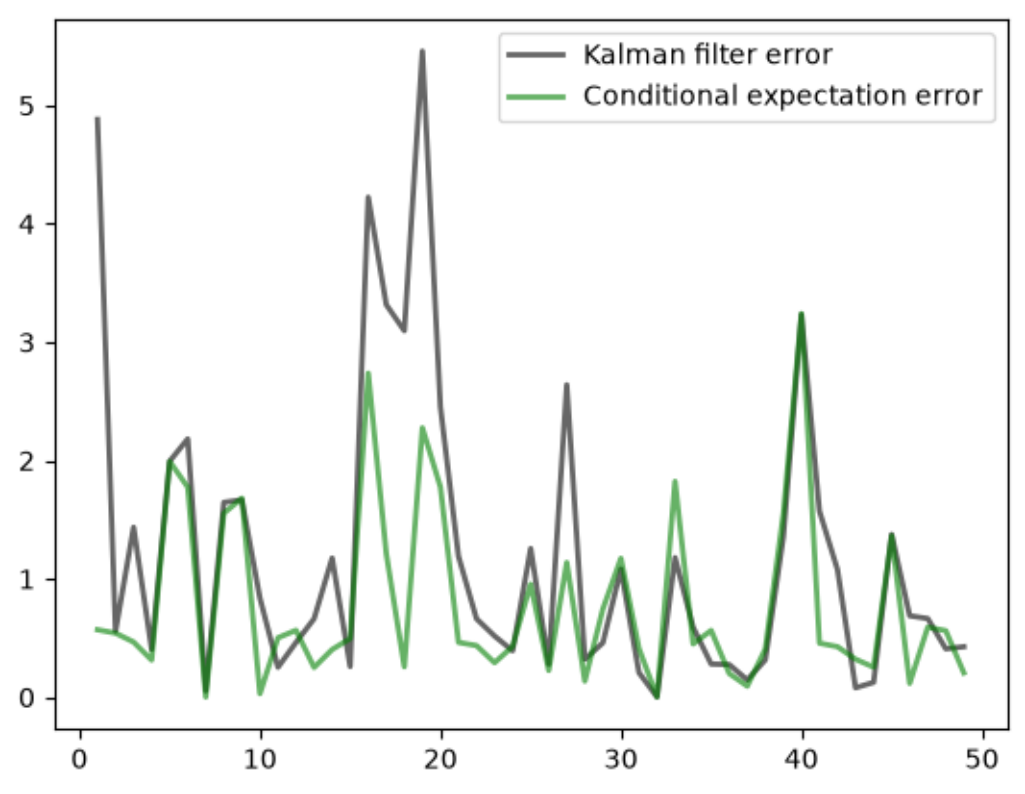
e1 = np.empty(T-1)
e2 = np.empty(T-1)

for t in range(1, T):
    kn.update(y[:, t-1])
    diff1 = x[:, t] - kn.x_hat.flatten()
    diff2 = x[:, t] - A @ x[:, t-1]
    e1[t-1] = diff1 @ diff1
    e2[t-1] = diff2 @ diff2

fig, ax = plt.subplots()
```

```
ax.plot(range(1, T), e1, 'k-', lw=2, alpha=0.6,
        label='Kalman filter error')
ax.plot(range(1, T), e2, 'g-', lw=2, alpha=0.6,
        label='Conditional expectation error')
ax.legend()
plt.show()
```

```
Eigenvalues of A:
[ 0.9+0.j -0.1+0.j]
Stationary prediction error variance:
[[0.40329108 0.1050718 ]
 [0.1050718  0.41061709]]
```



Observe how, after an initial learning period, the Kalman filter performs quite well, even relative to the competitor who predicts optimally with knowledge of the latent state.

Exercise 48.5.4

Try varying the coefficient 0.3 in $Q = 0.3I$ up and down.

Observe how the diagonal values in the stationary solution Σ (see (48.8)) increase and decrease in line with this coefficient.

The interpretation is that more randomness in the law of motion for X_t causes more (permanent) uncertainty in prediction.

CHAPTER

FORTYNINE

ANOTHER LOOK AT THE KALMAN FILTER

Contents

- *Another Look at the Kalman Filter*
  - *A worker's output*
  - *A firm's wage-setting policy*
  - *A state-space representation*
  - *An innovations representation*
  - *Some computational experiments*
  - *Future extensions*

In *A First Look at the Kalman Filter*, we used a Kalman filter to estimate locations of a rocket.

In this lecture, we'll use the Kalman filter to infer a worker's human capital and the effort that the worker devotes to accumulating human capital, neither of which the firm observes directly.

This lecture is an application of the filtering and forecasting recursion introduced in *A First Look at the Kalman Filter*.

The firm learns about those things only by observing a history of the output that the worker generates for the firm, and from understanding how that output depends on the worker's human capital and how human capital evolves as a function of the worker's effort.

We'll posit a rule that expresses how much the firm pays the worker each period as a function of the firm's information each period.

The lecture *The Kalman Filter and Vector Autoregressions* uses the same recursion to construct innovations, likelihoods, and vector autoregressions.

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

To conduct simulations, we bring in these imports, as in *A First Look at the Kalman Filter*.

```
import matplotlib.pyplot as plt
import numpy as np
from quantecon import Kalman, LinearStateSpace
from typing import NamedTuple
from scipy.stats import multivariate_normal
import matplotlib as mpl
```

(continues on next page)

(continued from previous page)

```
mpl.rcParams['text.usetex'] = True
mpl.rcParams['text.latex.preamble'] = r'\usepackage{amsmath,amsfonts}'
```

49.1 A worker's output

A representative worker is permanently employed at a firm.

A worker's output is described by the following dynamic process:

$$
\begin{aligned}
h_{t+1} &= \alpha h_t + \beta u_t + c\epsilon_{t+1}, \quad \epsilon_{t+1} \sim N(0,1) \\
u_{t+1} &= u_t \\
y_t &= g h_t + v_t, \quad v_t \sim N(0,R)
\end{aligned} \tag{49.1}
$$

Here

- h_t is the logarithm of human capital at time t
- u_t is the logarithm of the worker's effort at accumulating human capital at t
- y_t is the logarithm of the worker's output at time t
- ϵ_{t+1} is an IID standard normal shock to human capital
- $h_0 \sim N(\hat{h}_0, \sigma^2_{h,0})$
- $u_0 \sim N(\hat{u}_0, \sigma^2_{u,0})$

Parameters of the model are $\alpha, \beta, c, R, g, \hat{h}_0, \hat{u}_0, \sigma_{h,0}, \sigma_{u,0}$ where $\sigma_{h,0}$ and $\sigma_{u,0}$ are standard deviations of the firm's initial beliefs about h_0 and u_0.

We assume that h_0, u_0, $\{\epsilon_t\}$, and $\{v_t\}$ are mutually independent.

At time 0, a firm has hired the worker.

The worker is permanently attached to the firm and so works for the same firm at all dates $t = 0, 1, 2, ...$

At the beginning of time 0, the firm observes neither the worker's innate initial human capital h_0 nor its hard-wired permanent effort level u_0.

The firm believes that u_0 for a particular worker is drawn from a Gaussian probability distribution, and so is described by $u_0 \sim N(\hat{u}_0, \sigma^2_{u,0})$.

The h_t part of a worker's "type" moves over time, while the equation $u_{t+1} = u_t$ implies $u_t = u_0$ for all t.

Thus, from the firm's point of view, effort is a fixed, unobserved component of the worker's type that must be inferred from output observations.

At the beginning of time $t \geq 1$, before setting the wage w_t, the firm has observed the history $y^{t-1} = [y_{t-1}, y_{t-2}, \ldots, y_0]$.

The firm does not observe the worker's "type" (h_0, u_0).

After production in period t, the firm observes the worker's output y_t, which it then uses to update its beliefs going into period $t+1$.

49.2 A firm's wage-setting policy

At time $t \geq 1$, before observing current output y_t, the firm sets the worker's log wage using the past output history y^{t-1}:

$$w_t = g\mathbb{E}[h_t | y^{t-1}], \quad t \geq 1$$

and at time 0 pays the worker a log wage equal to the unconditional mean of y_0:

$$w_0 = g\hat{h}_0$$

In using this payment rule, the firm is taking into account that the worker's log output today is partly due to the random component v_t that comes entirely from luck, and that is assumed to be independent of h_t and u_t.

49.3 A state-space representation

Write system *(49.1.1)* in the state-space form

$$\begin{bmatrix} h_{t+1} \\ u_{t+1} \end{bmatrix} = \begin{bmatrix} \alpha & \beta \\ 0 & 1 \end{bmatrix} \begin{bmatrix} h_t \\ u_t \end{bmatrix} + \begin{bmatrix} c \\ 0 \end{bmatrix} \epsilon_{t+1}$$
$$y_t = \begin{bmatrix} g & 0 \end{bmatrix} \begin{bmatrix} h_t \\ u_t \end{bmatrix} + v_t$$

which is equivalent with

$$\begin{aligned} x_{t+1} &= Ax_t + C\epsilon_{t+1} \\ y_t &= Gx_t + v_t \\ x_0 &\sim N(\hat{x}_0, \Sigma_0) \end{aligned} \tag{49.2}$$

where

$$x_t = \begin{bmatrix} h_t \\ u_t \end{bmatrix}, \quad \hat{x}_0 = \begin{bmatrix} \hat{h}_0 \\ \hat{u}_0 \end{bmatrix}, \quad \Sigma_0 = \begin{bmatrix} \sigma_{h,0}^2 & 0 \\ 0 & \sigma_{u,0}^2 \end{bmatrix}$$

To compute the firm's wage setting policy, we first create a `NamedTuple` to store the parameters of the model

```
class WorkerModel(NamedTuple):
    A: np.ndarray
    C: np.ndarray
    G: np.ndarray
    R: float
    xhat_0: np.ndarray
    Σ_0: np.ndarray


def create_worker(α=.8, β=.2, c=.2,
                  R=.5, g=1.0, hhat_0=4, uhat_0=4,
                  σ_h=2, σ_u=2):

    A = np.array([[α, β],
                  [0, 1]])
    C = np.array([[c],
                  [0]])
    G = np.array([g, 0])
```

(continues on next page)

(continued from previous page)

```
    # Define initial state and covariance matrix
    xhat_0 = np.array([[hhat_0],
                       [uhat_0]])

    # σ_h and σ_u are standard deviations, so Σ_0 holds their squares
    Σ_0 = np.array([[σ_h**2, 0],
                    [0, σ_u**2]])

    return WorkerModel(A=A, C=C, G=G, R=R, xhat_0=xhat_0, Σ_0=Σ_0)
```

Please note how the `WorkerModel` NamedTuple creates all of the objects required to compute an associated state-space representation (49.2).

This is handy, because in order to simulate a history $\{y_t, h_t\}$ for a worker, we'll want to form state space system for him/her by using the `LinearStateSpace` class.

```
# Define A, C, G, R, xhat_0, Σ_0
worker = create_worker()
A, C, G, R = worker.A, worker.C, worker.G, worker.R
xhat_0, Σ_0 = worker.xhat_0, worker.Σ_0

# Create a LinearStateSpace object
ss = LinearStateSpace(A, C, G, np.sqrt(R),
        mu_0=xhat_0, Sigma_0=np.zeros((2,2)))

T = 100
x, y = ss.simulate(T)
y = y.flatten()

h_0, u_0 = x[0, 0], x[1, 0]
```

We set `Sigma_0=np.zeros((2,2))` so that the simulation fixes a particular worker's initial state (h_0, u_0), while the firm still enters period 0 with the non-degenerate prior beliefs $\hat{x}_0$ and Σ_0 that drive its Kalman filter.

Next, to compute the firm's policy for setting the log wage based on the information it has about the worker, we use the Kalman filter described in this QuantEcon lecture *A First Look at the Kalman Filter*.

In particular, we want to compute all of the objects in an "innovation representation".

49.4 An innovations representation

We have all the objects in hand required to form an innovations representation for the output process $\{y_t\}_{t=0}^{T-1}$ for a worker.

Let's code that up now.

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + K_t a_t \\ y_t &= G\hat{x}_t + a_t \end{aligned}$$

where $\hat{x}_t = \mathbb{E}[x_t | y^{t-1}]$ is the firm's forecast of the state formed before y_t is observed, and K_t is the Kalman gain matrix at time t.

Here $a_t = y_t - G\hat{x}_t$ is the **innovation** at time t, the firm's one-step-ahead error in forecasting output y_t from the history y^{t-1}.

Because $\hat{x}_t$ is conditioned on y^{t-1} rather than y^t, the gain K_t folds together the filtering update that uses the current observation y_t and the one-step-ahead prediction that advances the state to $t+1$.

Writing $\Sigma_t = \mathbb{E}[(x_t - \hat{x}_t)(x_t - \hat{x}_t)' | y^{t-1}]$ for the conditional covariance of the state, the gain is

$$K_t = A\Sigma_t G'(G\Sigma_t G' + R)^{-1} = AL_t,$$

where $L_t = \Sigma_t G'(G\Sigma_t G' + R)^{-1}$ is the filtering gain that updates the firm's beliefs about x_t once y_t is observed.

We accomplish this in the following code that uses the `Kalman` class.

```
kalman = Kalman(ss, xhat_0, Σ_0)
Σ_t = np.zeros((*Σ_0.shape, T))
y_hat_t = np.zeros(T)
x_hat_t = np.zeros((2, T))

for t in range(T):
    # Record the firm's belief about x_t given y^{t-1}, before seeing y_t
    x_hat, Σ = kalman.x_hat, kalman.Sigma
    Σ_t[:, :, t] = Σ
    x_hat_t[:, t] = x_hat.reshape(-1)
    y_hat_t[t] = (worker.G @ x_hat).item()

    # Then incorporate the observation y_t and advance the filter to t+1
    kalman.update(y[t])

u_hat_t = x_hat_t[1, :]
```

For this fixed worker initial state, we plot $\mathbb{E}[y_t | y^{t-1}] = G\hat{x}_t$ where $\hat{x}_t = \mathbb{E}[x_t | y^{t-1}]$.

We also plot $\mathbb{E}[u_0 | y^{t-1}]$, which is the firm's inference about a worker's hard-wired "work ethic" u_0, conditioned on information y^{t-1} that it has about him or her coming into period t.

We can watch how the firm updates its inference $\mathbb{E}[u_0 | y^{t-1}]$ about the worker's work ethic as more output observations arrive.

```
fig, ax = plt.subplots(1, 2)

ax[0].plot(y_hat_t, label=r'$\mathbb{E}[y_t| y^{t-1}]$')
ax[0].set_xlabel('time')
ax[0].set_ylabel(r'$\mathbb{E}[y_t| y^{t-1}]$')
ax[0].set_title(r'$\mathbb{E}[y_t| y^{t-1}]$ over time')
ax[0].legend()

ax[1].plot(u_hat_t, label=r'$\mathbb{E}[u_0|y^{t-1}]$')
ax[1].axhline(y=u_0, color='grey',
              linestyle='dashed', label=fr'$u_0={u_0:.2f}$')
ax[1].set_xlabel('time')
ax[1].set_ylabel(r'$\mathbb{E}[u_0|y^{t-1}]$')
ax[1].set_title('Inferred work ethic over time')
ax[1].legend()

fig.tight_layout()
plt.show()
```

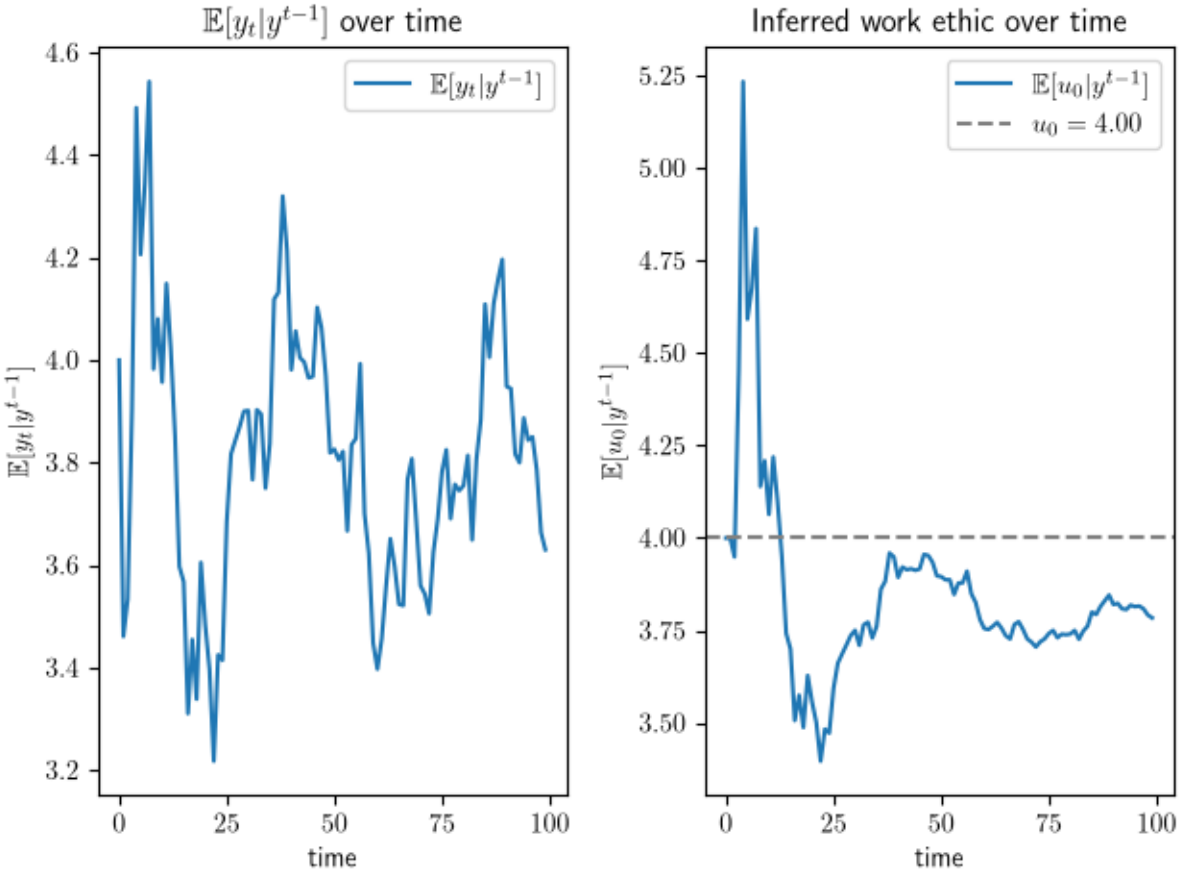


Fig. 49.1: Firm's output forecast and inferred work ethic over time

49.5 Some computational experiments

Let's look at Σ_0 and Σ_{T-1} in order to see how much the firm learns about the hidden state during the horizon we have set.

```
print(Σ_t[:, :, 0])
```

```
[[4. 0.]
 [0. 4.]]
```

```
print(Σ_t[:, :, -1])
```

```
[[0.09403859 0.00997018]
 [0.00997018 0.01592704]]
```

Evidently, the conditional variances become smaller over time.

It is enlightening to portray how the firm's conditional beliefs evolve by plotting contours of the conditional bivariate normal density of x_t given y^{t-1} at various t's.

```
# Create a grid of points for contour plotting
h_range = np.linspace(x_hat_t[0, :].min()-0.5*Σ_t[0, 0, 1],
                      x_hat_t[0, :].max()+0.5*Σ_t[0, 0, 1], 100)
u_range = np.linspace(x_hat_t[1, :].min()-0.5*Σ_t[1, 1, 1],
                      x_hat_t[1, :].max()+0.5*Σ_t[1, 1, 1], 100)
h, u = np.meshgrid(h_range, u_range)

# Create a figure with subplots for each time step
fig, axs = plt.subplots(1, 3, figsize=(12, 7))

# Iterate through each time step
for i, t in enumerate(np.linspace(0, T-1, 3, dtype=int)):
    # Create a multivariate normal distribution with x_hat and Σ at time step t
    μ = x_hat_t[:, t]
    cov = Σ_t[:, :, t]
    mvn = multivariate_normal(mean=μ, cov=cov)

    # Evaluate the multivariate normal PDF on the grid
    pdf_values = mvn.pdf(np.dstack((h, u)))

    # Create a contour plot for the PDF
    con = axs[i].contour(h, u, pdf_values, cmap='viridis')
    axs[i].clabel(con, inline=1, fontsize=10)
    axs[i].set_title(f'time step {t}')
    axs[i].set_xlabel(r'$h_{{{}}}$'.format(str(t)))
    axs[i].set_ylabel(r'$u_{{{}}}$'.format(str(t)))

    cov_latex = (
        r'$\Sigma_{{{}}}= \begin{{bmatrix}} {:.2f} & {:.2f} \\ '
        r'{:.2f} & {:.2f} \end{{bmatrix}}$'
    ).format(t, cov[0, 0], cov[0, 1], cov[1, 0], cov[1, 1])
    axs[i].text(0.33, -0.15, cov_latex, transform=axs[i].transAxes)


plt.tight_layout()
plt.show()
```

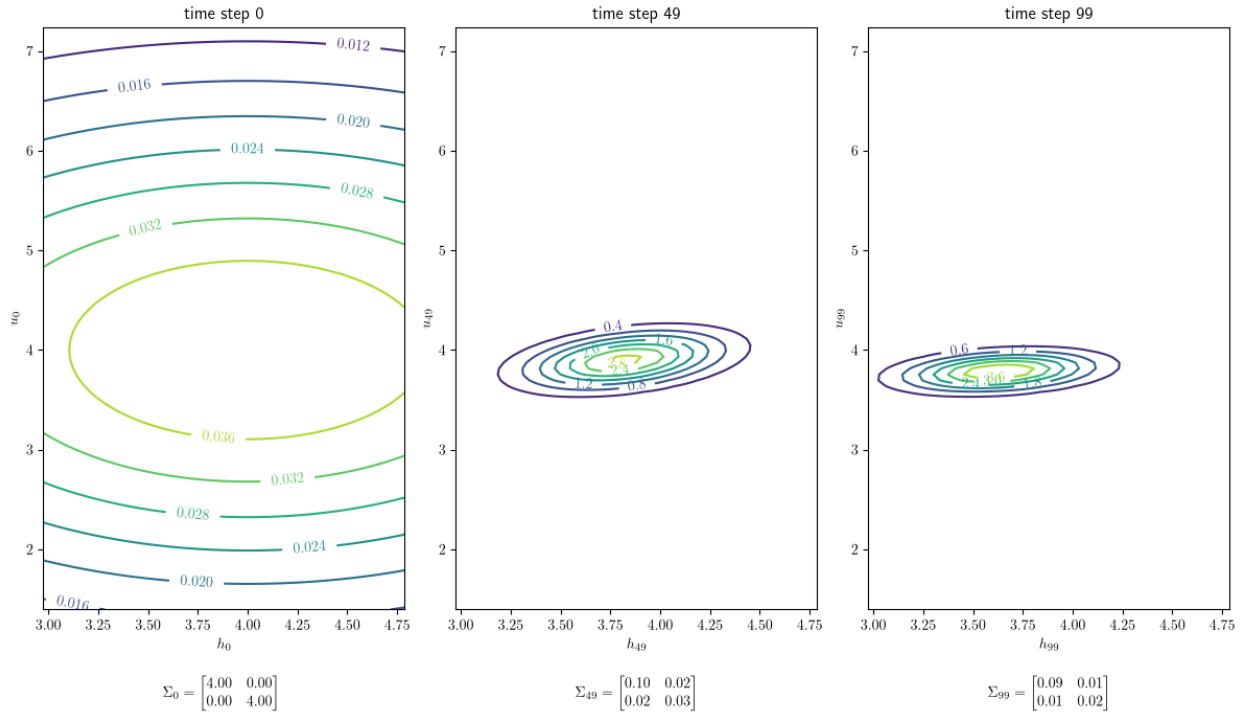


Fig. 49.2: Density contours for the firm's belief about x_t at three dates

Note how the accumulation of evidence y^{t-1} affects the shape of the density contours as sample size t grows.

Now let's use our code to set the hidden state x_0 to a particular vector in order to watch how a firm learns starting from some x_0 we are interested in.

For example, let's say $h_0 = 0$ and $u_0 = 4$.

Here is one way to do this.

```
# For example, we might want h_0 = 0 and u_0 = 4
μ_0 = np.array([[0.0],
                [4.0]])

# Create a LinearStateSpace object with Sigma_0 as a matrix of zeros
ss_example = LinearStateSpace(A, C, G, np.sqrt(R), mu_0=μ_0,
                              # This line forces exact h_0=0 and u_0=4
                              Sigma_0=np.zeros((2, 2))
                             )

T = 100
x, y = ss_example.simulate(T)
y = y.flatten()

# Now h_0=0 and u_0=4
h_0, u_0 = x[0, 0], x[1, 0]
print('h_0 =', h_0)
print('u_0 =', u_0)
```

```
h_0 = 0.0
u_0 = 4.0
```

Another way to accomplish the same goal is to use the following code.

```
# If we want to set the initial
# h_0 = hhat_0 = 0.0 and u_0 = uhat_0 = 4.0:
worker_example = create_worker(hhat_0=0.0, uhat_0=4.0)

# The firm's prior stays at the original xhat_0 and Σ_0
ss_example = LinearStateSpace(A, C, G, np.sqrt(R),
                              # This line takes h_0=hhat_0 and u_0=uhat_0
                              mu_0=worker_example.xhat_0,
                              # This line forces exact h_0=hhat_0 and u_0=uhat_0
                              Sigma_0=np.zeros((2, 2))
                             )

T = 100
x, y = ss_example.simulate(T)
y = y.flatten()

# Now h_0 and u_0 will be exactly hhat_0 and uhat_0
h_0, u_0 = x[0, 0], x[1, 0]
print('h_0 =', h_0)
print('u_0 =', u_0)
```

```
h_0 = 0.0
u_0 = 4.0
```

For this worker, let's generate a plot like the one above.

```
# The firm filters output from ss_example using its prior beliefs
# xhat_0 and Σ_0
kalman = Kalman(ss_example, xhat_0, Σ_0)
Σ_t = []
y_hat_t = np.zeros(T)
u_hat_t = np.zeros(T)

# Then we iteratively update the Kalman filter class using
# observation y based on the linear state model above:
for t in range(T):
    # Record the firm's belief about x_t given y^{t-1}, before seeing y_t
    x_hat, Σ = kalman.x_hat, kalman.Sigma
    Σ_t.append(Σ)
    y_hat_t[t] = (G @ x_hat).item()
    u_hat_t[t] = x_hat[1].item()

    # Then incorporate the observation y_t and advance the filter to t+1
    kalman.update(y[t])


# Generate plots for y_hat_t and u_hat_t
fig, ax = plt.subplots(1, 2)

ax[0].plot(y_hat_t, label=r'$\mathbb{E}[y_t| y^{t-1}]$')
ax[0].set_xlabel('time')
ax[0].set_ylabel(r'$\mathbb{E}[y_t| y^{t-1}]$')
```

(continues on next page)

(continued from previous page)

```
ax[0].set_title(r'$\mathbb{E}[y_t| y^{t-1}]$ over time')
ax[0].legend()

ax[1].plot(u_hat_t, label=r'$\mathbb{E}[u_0|y^{t-1}]$')
ax[1].axhline(y=u_0, color='grey',
              linestyle='dashed', label=fr'$u_0={u_0:.2f}$')
ax[1].set_xlabel('time')
ax[1].set_ylabel(r'$\mathbb{E}[u_0|y^{t-1}]$')
ax[1].set_title('Inferred work ethic over time')
ax[1].legend()

fig.tight_layout()
plt.show()
```

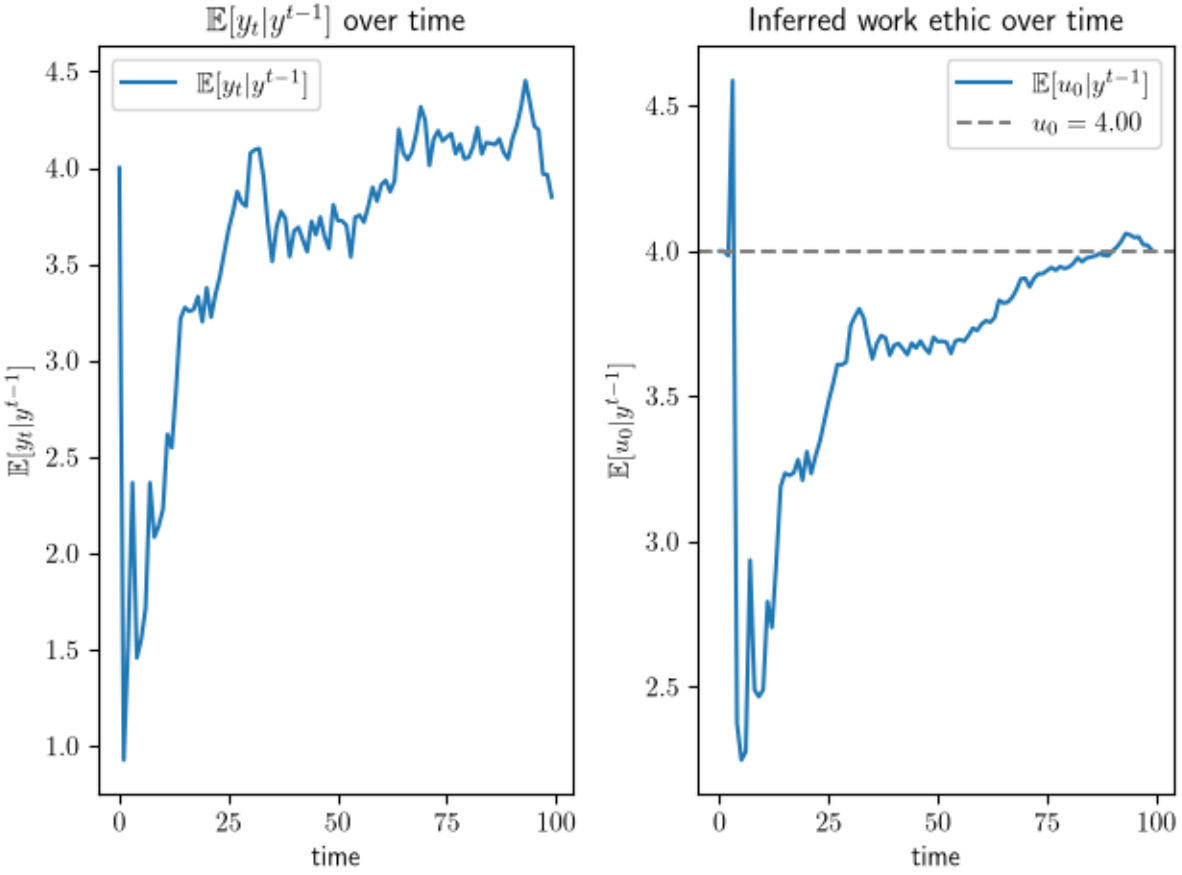


Fig. 49.3: Output forecast and inferred work ethic for the worker with $u_0 = 4$

More generally, we can change some or all of the parameters defining a worker in our `create_worker` factory function.

Here is an example.

```
# We can set these parameters when creating a worker, just like classes!
hard_working_worker = create_worker(α=.4, β=.8,
                        hhat_0=7.0, uhat_0=100, σ_h=2.5, σ_u=3.2)
```

(continues on next page)

(continued from previous page)

```
print(hard_working_worker)
```

```
WorkerModel(A=array([[0.4, 0.8],
       [0. , 1. ]]), C=array([[0.2],
       [0. ]]), G=array([1., 0.]), R=0.5, xhat_0=array([[  7.],
       [100.]]), Σ_0=array([[ 6.25,  0.  ],
       [ 0.  , 10.24]]))
```

We can also simulate the system for $T = 100$ periods for different workers.

When effort feeds into output through human capital, so that $g \neq 0$ and $\beta \neq 0$, the firm's uncertainty about u_0 falls over time and the inferred work ethic converges toward the true u_0.

If $\beta = 0$, effort never affects h_t, and if $g = 0$, output carries no information about h_t, so in either case the firm cannot learn u_0 from output alone.

This shows that, under these observability conditions, the filter gradually teaches the firm about the worker's effort.

For three workers, we first plot the gap between the firm's inferred work ethic and the true u_0 over time.

```
num_workers = 3
T = 100
fig, ax = plt.subplots(figsize=(7, 7))

for i in range(num_workers):
    worker = create_worker(uhat_0=4+2*i)
    simulate_workers(worker, T, ax, name=fr'$\hat u_0 = {4+2*i}$',
                     random_state=2 + i)
ax.set_ylim(ymin=-2, ymax=2)
ax.legend()
plt.show()
```

In this simulation the firm's inferred work ethic moves toward the true u_0.

Under the correctly specified observable linear-Gaussian model, the posterior mean of u_0 is consistent, so this gap shrinks as the output history grows.

By setting `diff=False`, we instead plot the level of each worker's inferred work ethic $\mathbb{E}[u_0|y^{t-1}]$, shown together with a dashed line marking the true u_0.

```
fig, ax = plt.subplots(figsize=(7, 7))

uhat_0s = [2, -2, 1]
αs = [0.2, 0.3, 0.5]
βs = [0.2, 0.9, 0.3]

for i, (uhat_0, α, β) in enumerate(zip(uhat_0s, αs, βs)):
    worker = create_worker(uhat_0=uhat_0, α=α, β=β)
    simulate_workers(worker, T, ax, diff=False,
                     name=r'$u_{{{}, 0}}$'.format(i),
                     random_state=3 + i)

ax.legend(bbox_to_anchor=(1, 0.5))
plt.show()
```

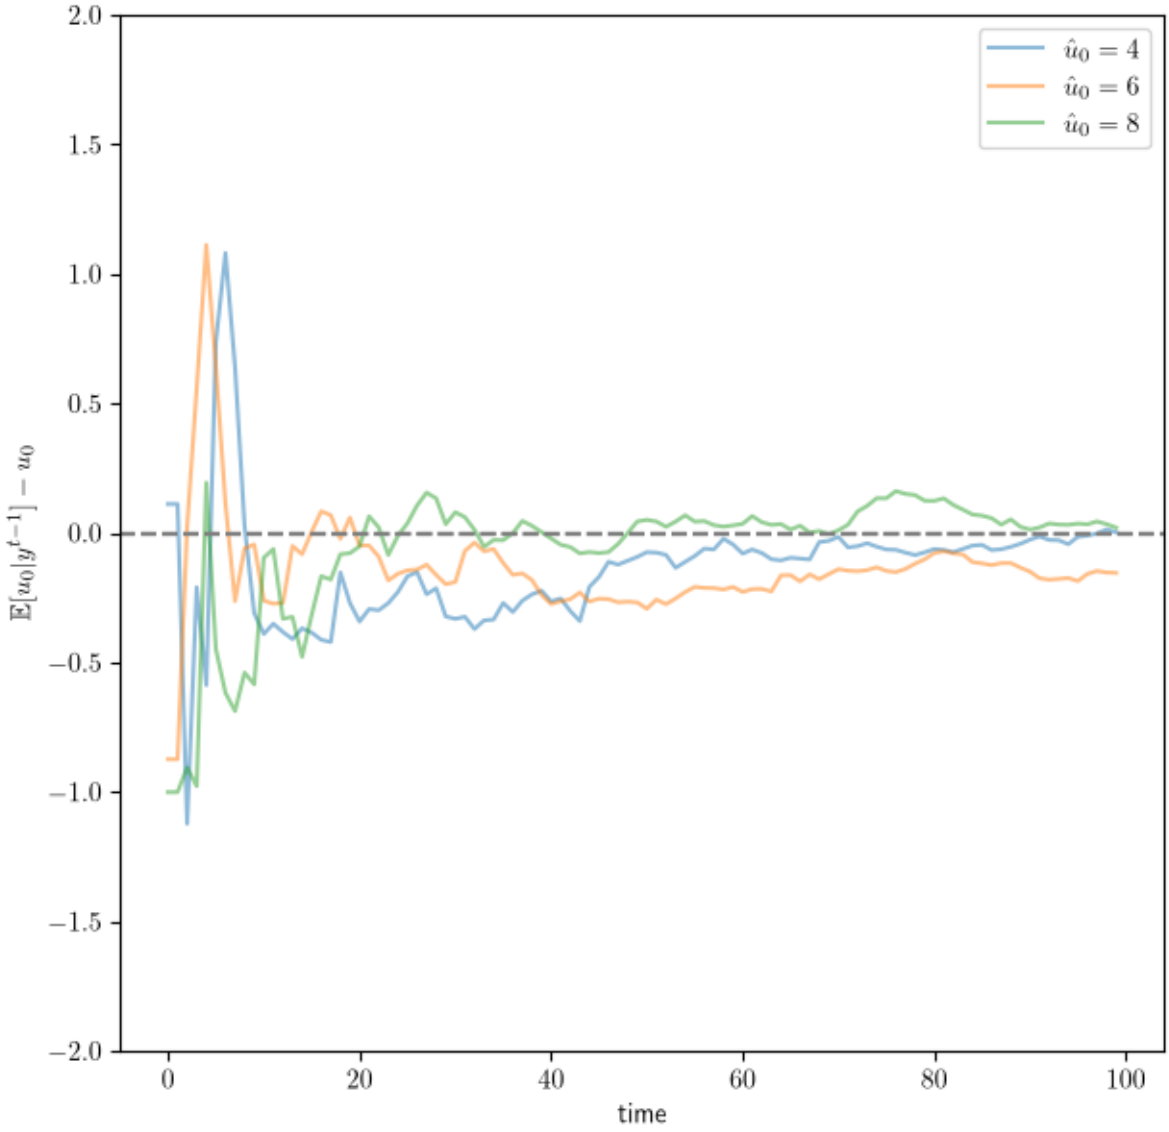


Fig. 49.4: Difference between inferred and true work ethic over time

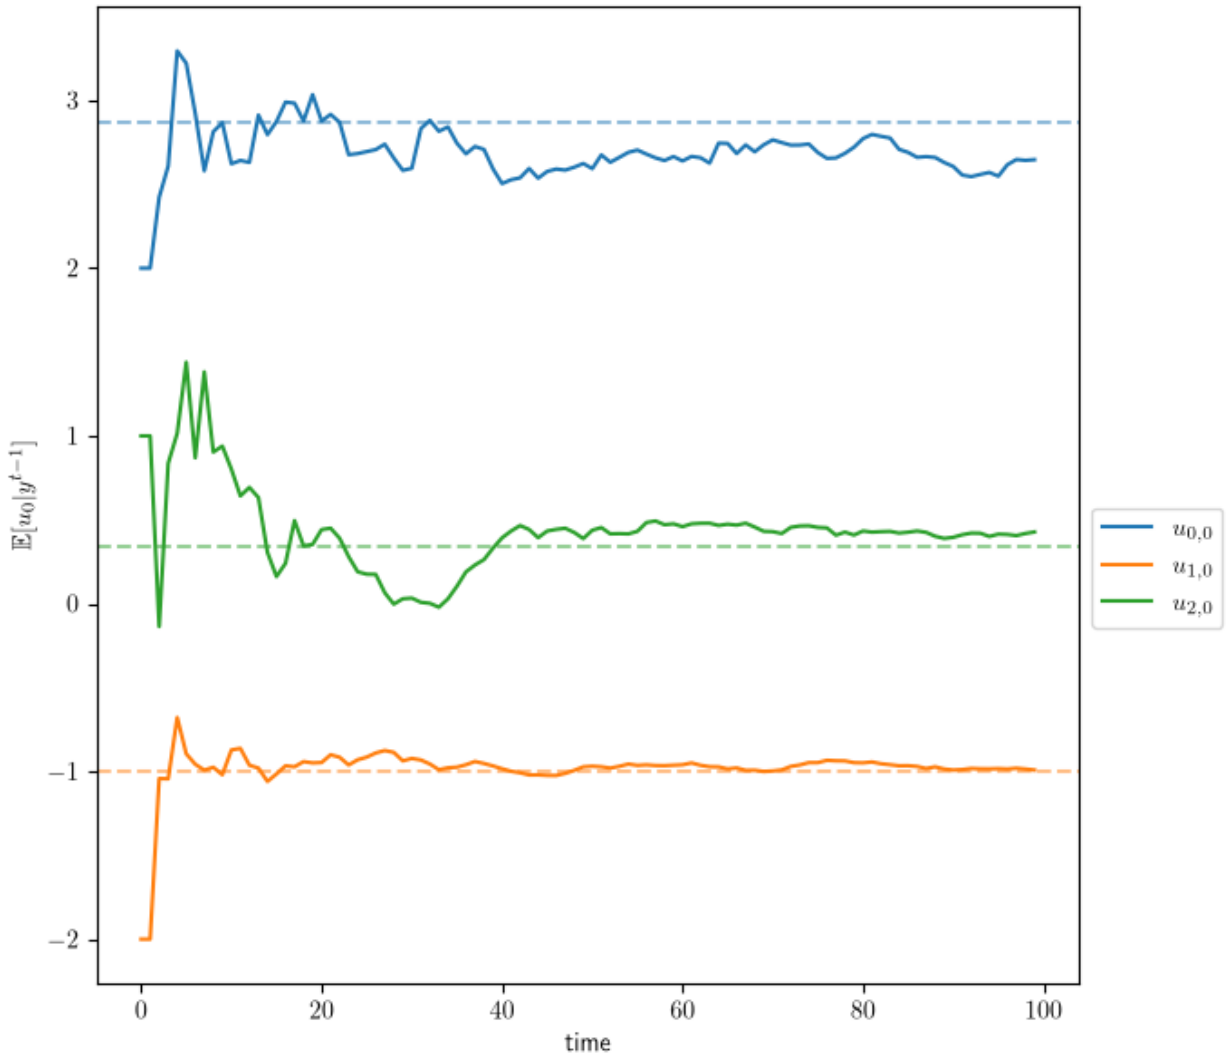


Fig. 49.5: Inferred work ethic over time for three workers

These three workers differ in α and β as well as in $\hat{u}_0$, and the speed of learning differs sharply across them.

The worker with the largest β, here $u_{1,0}$ with $\beta = 0.9$, settles onto its dashed true value almost immediately, while the worker with the smallest β, here $u_{0,0}$ with $\beta = 0.2$, converges slower.

The reason is that effort affects output only through human capital, so in these stable examples with $|\alpha| < 1$ its steady-state effect on output is governed by $g\beta/(1-\alpha)$, and a small β leaves the firm with too little signal to pin down u_0 over this horizon.

The speed of learning also reflects the measurement noise R, the shock scale c, and the firm's prior variances.

We can also give every worker the same true initial state, here $h_0 = 2$ and $u_0 = 1$, by passing a fixed `μ_sim_0` and a zero `Σ_sim_0` to `simulate_workers`.

```
fig, ax = plt.subplots(figsize=(7, 7))

μ_sim_0 = np.array([[2.0],
                    [1.0]])
Σ_sim_0 = np.zeros((2,2))

uhat_0s = [2, -2, 1]
αs = [0.2, 0.3, 0.5]
βs = [0.2, 0.9, 0.3]

for i, (uhat_0, α, β) in enumerate(zip(uhat_0s, αs, βs)):
    worker = create_worker(uhat_0=uhat_0, α=α, β=β)
    simulate_workers(worker, T, ax, μ_sim_0=μ_sim_0, Σ_sim_0=Σ_sim_0,
                     diff=False, name=r'$u_{{{}, 0}}$'.format(i))

# This controls the boundary of plots
ax.set_ylim(ymin=-3, ymax=3)
ax.legend(bbox_to_anchor=(1, 0.5))
plt.show()
```

Even though the firm begins from different prior means $\hat{u}_0$, all three workers share the true work ethic $u_0 = 1$, and the inferred paths converge to that common dashed line at speeds that again reflect each worker's β.

Finally, we track a single worker type under two different true effort levels, comparing a hard-working worker with $u_0 = 100$ against a normal worker with $u_0 = 30$.

```
T = 50
fig, ax = plt.subplots(figsize=(7, 7))

μ_sim_0_1 = np.array([[1],
                      [100]])
μ_sim_0_2 = np.array([[1],
                      [30]])
Σ_sim_0 = np.zeros((2, 2))

worker = create_worker(uhat_0=1, α=0.5, β=0.3)
simulate_workers(worker, T, ax, μ_sim_0=μ_sim_0_1, Σ_sim_0=Σ_sim_0,
                 diff=False, name=r'Hard-working worker')
simulate_workers(worker, T, ax, μ_sim_0=μ_sim_0_2, Σ_sim_0=Σ_sim_0,
                 diff=False, name=r'Normal worker')
ax.legend(bbox_to_anchor=(1, 0.5))
plt.show()
```

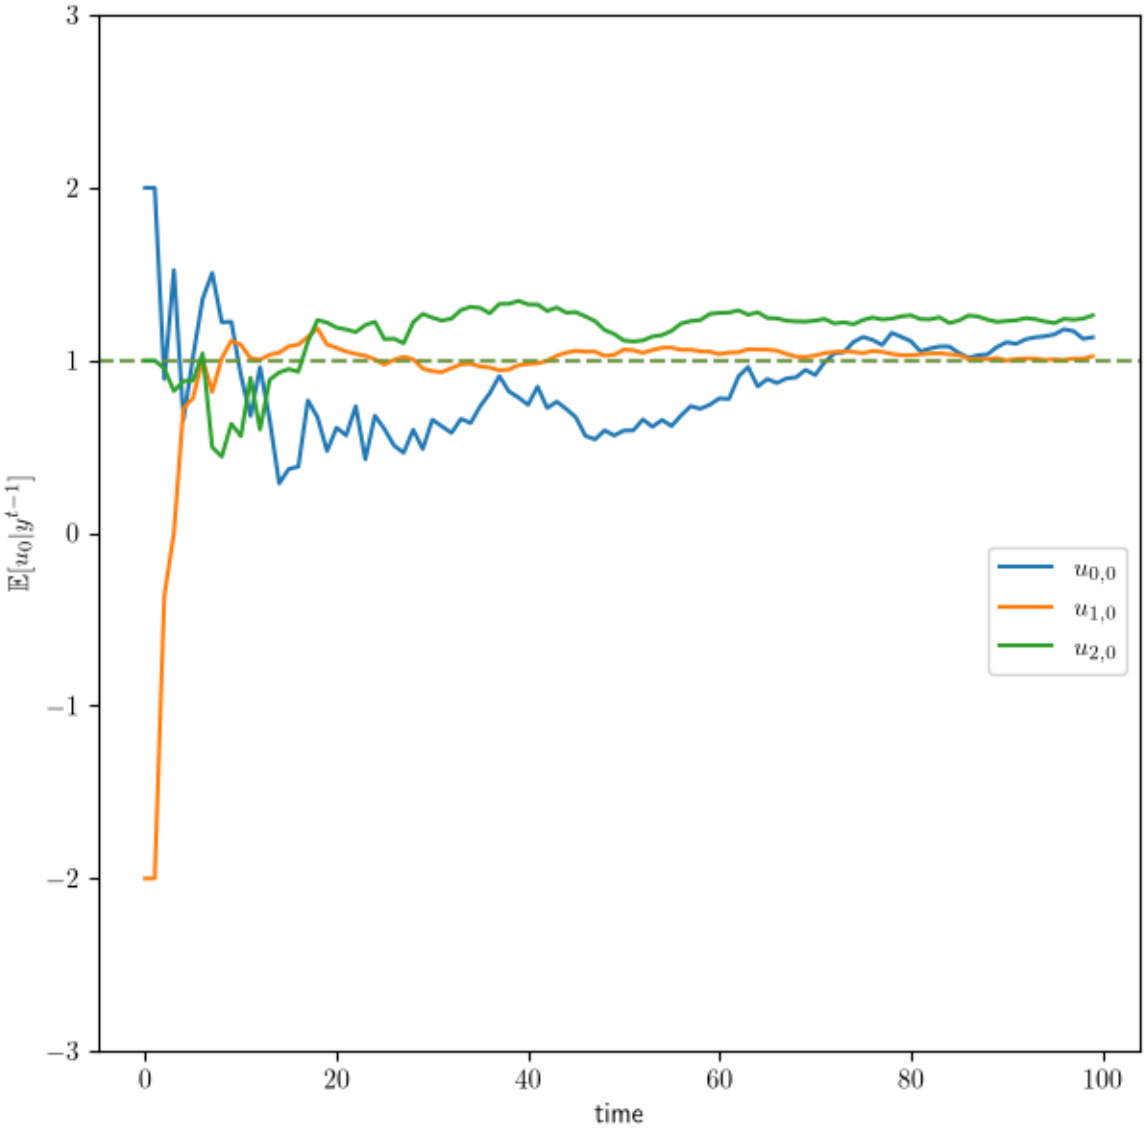


Fig. 49.6: Inferred work ethic when every worker starts at $h_0 = 2$ and $u_0 = 1$

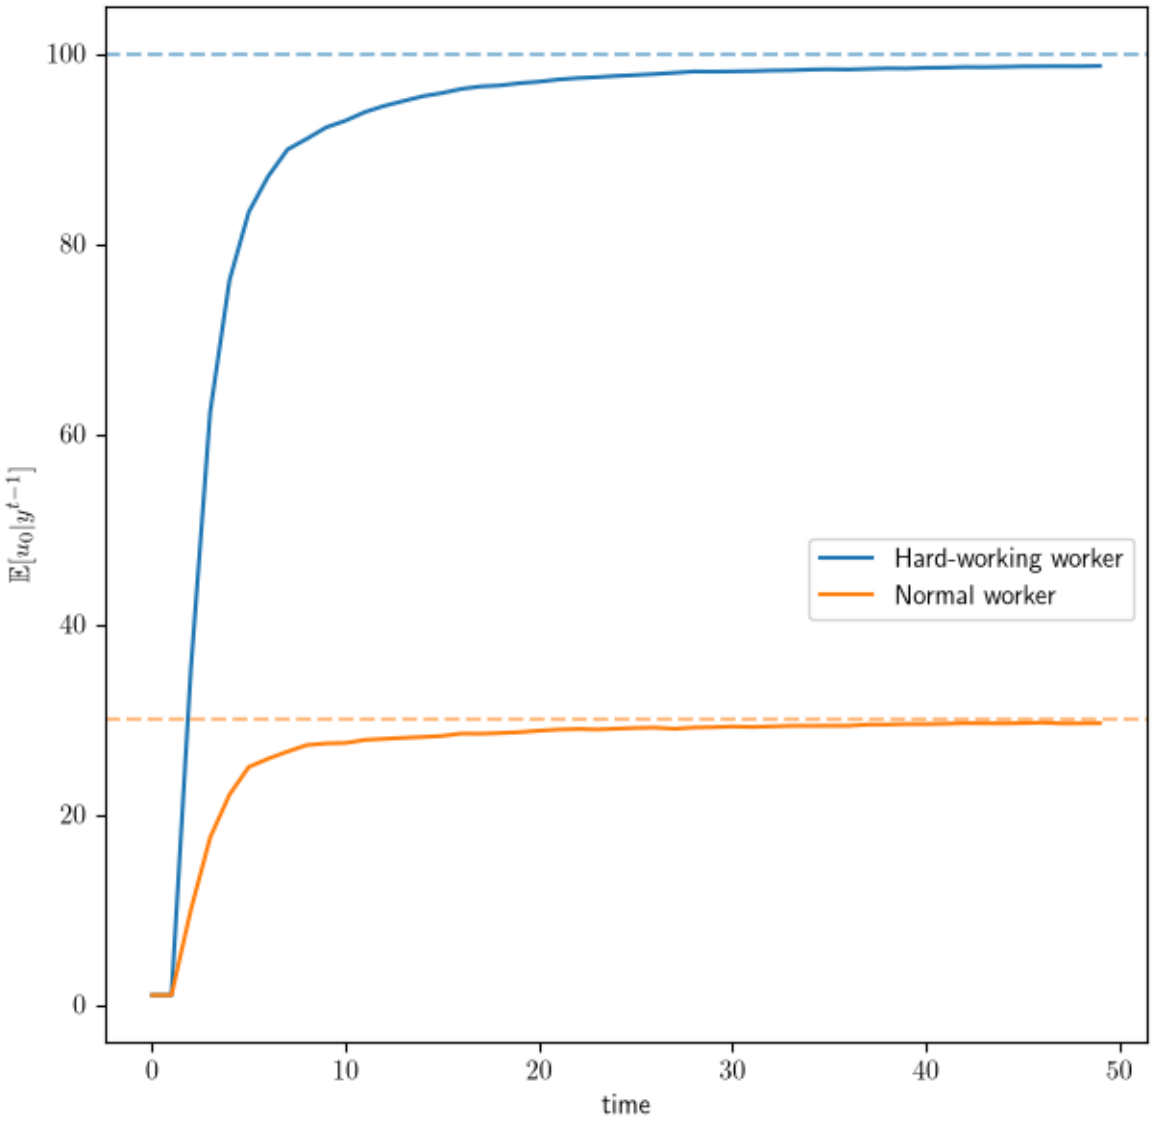


Fig. 49.7: A hard-working worker and a less hard-working worker

Both inferred paths start from the firm's common prior $\widehat{u}_0 = 1$ and climb toward their different true values, showing that the filter corrects the gap between prior and truth as evidence accumulates.

49.6 Future extensions

We can do lots of enlightening experiments by creating new types of workers and letting the firm learn about their hidden (to the firm) states by observing just their output histories.

CHAPTER

FIFTY

THE KALMAN FILTER AND VECTOR AUTOREGRESSIONS

Contents

- *The Kalman Filter and Vector Autoregressions*
  - *Overview*
  - *The state space system*
  - *The Kalman filter*
  - *The Gram-Schmidt process*
  - *Hidden Markov model*
  - *Estimation*
  - *Vector autoregressions and the Kalman filter*
  - *Spectral factorization identity*
  - *Python implementation*
  - *Where this leads*
  - *Summary*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

50.1 Overview

This lecture derives the **Kalman filter** for a linear Gaussian state space system and then uses it to construct **vector autoregressions (VARs)**.

It builds on *A First Look at the Kalman Filter*, where the filter is introduced through filtering and forecasting distributions.

It complements *Another Look at the Kalman Filter*, where the same recursion is used in a worker-firm learning problem.

This lecture shifts attention from filtering a hidden state to representing and estimating the observable process generated by that hidden state.

Our approach rests on repeated applications of the **population linear least squares** projection formula, the insight that computing a conditional expectation of a jointly normal random vector is the same as running a population OLS regression.

The lecture covers:

- deriving the Kalman filter recursions from first principles
- the **matrix Riccati difference equation** governing conditional covariance matrices
- the **innovations representation** and the **Gram-Schmidt** whitening property
- the structure of a **hidden Markov model**
- the **likelihood function** for a state space system and its role in maximum likelihood and Bayesian estimation
- how the time-invariant Kalman filter generates a **vector autoregression**
- why the Kalman filter is an essential tool for *interpreting VARs* estimated from economic data

A sequel, *Vector Autoregressions for Subsets of Variables*, applies this machinery to a question that arises constantly in practice: what happens to a VAR when an econometrician observes only some of the variables that appear in it.

50.2 The state space system

The Kalman filter applies to the **state space system** for $t \geq 0$:

$$\begin{aligned} x_{t+1} &= Ax_t + Cw_{t+1} \\ y_t &= Gx_t + v_t \end{aligned} \tag{50.1}$$

where

- x_t is an $n \times 1$ **state vector** (hidden, unobserved)
- y_t is an $m \times 1$ vector of **signals** on the hidden state (observed)
- w_{t+1} is a $p \times 1$ IID sequence of normal random variables with mean 0 and identity covariance matrix
- v_t is an IID sequence of normal random variables with mean zero and covariance matrix R
- w_t and v_s are orthogonal at all pairs of dates

The coefficient matrices have the following dimensions: A is $n \times n$, C is $n \times p$, G is $m \times n$, and R is $m \times m$.

The initial state satisfies

$$x_0 \sim N(\hat{x}_0, \Sigma_0) \tag{50.2}$$

We observe $y_t, \ldots, y_0$ but *not* $x_t, \ldots, x_0$ at time t, and we know all first and second moments implied by (50.1) and (50.2).

50.3 The Kalman filter

50.3.1 Starting distribution

Working forward in time starting at $t = 0$, before observing y_0, the specification (50.1)-(50.2) implies that the marginal distribution of y_0 is

$$y_0 \sim N(G\hat{x}_0,\ G\Sigma_0 G^\top + R) \tag{50.3}$$

For $t \geq 0$ let $y^t = [y_t, y_{t-1}, \ldots, y_0]$.

We want a convenient recursive representation of the conditional distribution of y_t given history y^{t-1}.

The Kalman filter attains this by constructing recursive formulas for $\hat{x}_t$ and Σ_t such that the distribution of y_t conditional on y^{t-1} generalizes (50.3) to

$$y_t \sim N(G\hat{x}_t,\ G\Sigma_t G^\top + R) \tag{50.4}$$

for $t \geq 1$, where the distribution of x_t conditional on y^{t-1} is $N(\hat{x}_t, \Sigma_t)$.

The objects $\hat{x}_t$ and Σ_t characterize the **population regression**

$$\hat{x}_t = \mathbb{E}[x_t \mid y_{t-1}, \dots, y_0]$$

and the **conditional covariance matrix**

$$\Sigma_t = \mathbb{E}[(x_t - \hat{x}_t)(x_t - \hat{x}_t)^\top \mid y_{t-1}, \dots, y_0]$$

50.3.2 Derivation

At each date, our approach is to *regress what we do not know on what we know*.

Note

Because our assumptions imply that $\{x_t, y_t\}_{t=0}^{\infty}$ is a jointly normal stochastic process, linear least squares regressions equal conditional mathematical expectations.

Each step below is an application of Bayes' law.

Under the weaker assumption that all means and covariances exist without joint normality, the same calculations yield "wide-sense conditional expectations" that coincide with true conditional expectations only when those conditional expectations are linear.

We arrive at $t = 0$ knowing $\hat{x}_0$ and Σ_0.

The information about x_0 in y_0 that is new relative to $(\hat{x}_0, \Sigma_0)$ is the **innovation**

$$a_0 \equiv y_0 - G\hat{x}_0$$

Let L_0 be the population regression coefficient of the hidden-state error $x_0 - \hat{x}_0$ on the signal surprise $y_0 - G\hat{x}_0$.

The conditional mean $\mathbb{E}[x_0 \mid y_0] = \hat{x}_0 + L_0(y_0 - G\hat{x}_0)$ satisfies the population regression formula

$$x_0 - \hat{x}_0 = L_0(y_0 - G\hat{x}_0) + \eta \tag{50.5}$$

where η is the least squares residual.

Orthogonality of η to $(y_0 - G\hat{x}_0)$ pins down L_0 via the normal equations

$$\mathbb{E}(x_0 - \hat{x}_0)(y_0 - G\hat{x}_0)^\top = L_0\, \mathbb{E}(y_0 - G\hat{x}_0)(y_0 - G\hat{x}_0)^\top$$

Evaluating the moment matrices and solving for L_0 gives

$$L_0 = \Sigma_0 G^\top (G\Sigma_0 G^\top + R)^{-1} \tag{50.6}$$

Thus L_0 updates the estimate of x_0, while $K_0 = AL_0$ updates the forecast of x_1.

To forecast x_1, note that

$$x_1 = A\hat{x}_0 + A(x_0 - \hat{x}_0) + Cw_1 \tag{50.7}$$

Applying (50.5) gives $\mathbb{E}[x_1 \mid y_0] = A\hat{x}_0 + AL_0(y_0 - G\hat{x}_0)$, which we write as

$$\hat{x}_1 = A\hat{x}_0 + K_0(y_0 - G\hat{x}_0) \tag{50.8}$$

where the **Kalman gain** at time 0 is

$$K_0 = A\Sigma_0 G^\top (G\Sigma_0 G^\top + R)^{-1} \tag{50.9}$$

Subtracting (50.8) from (50.7) yields

$$x_1 - \hat{x}_1 = A(x_0 - \hat{x}_0) + Cw_1 - K_0(y_0 - G\hat{x}_0) \tag{50.10}$$

Using (50.10) and $y_0 = Gx_0 + v_0$ to evaluate $\Sigma_1 \equiv \mathbb{E}[(x_1 - \hat{x}_1)(x_1 - \hat{x}_1)^\top \mid y_0]$ gives

$$\Sigma_1 = (A - K_0 G)\Sigma_0 (A - K_0 G)^\top + CC^\top + K_0 R K_0^\top \tag{50.11}$$

Thus $f(x_1 \mid y_0) \sim N(\hat{x}_1, \Sigma_1)$.

Collecting the time-0 equations:

$$\begin{aligned} a_0 &= y_0 - G\hat{x}_0 \\ K_0 &= A\Sigma_0 G^\top (G\Sigma_0 G^\top + R)^{-1} \\ \hat{x}_1 &= A\hat{x}_0 + K_0 a_0 \\ \Sigma_1 &= CC^\top + K_0 R K_0^\top + (A - K_0 G)\Sigma_0 (A - K_0 G)^\top \end{aligned} \tag{50.12}$$

System (50.12) maps a mean-covariance pair $(\hat{x}_0, \Sigma_0)$ into a new pair $(\hat{x}_1, \Sigma_1)$, with auxiliary outputs (a_0, K_0).

Recognizing that "we are in the same situation at the start of period 1 as at the start of period 0" activates a recursion, the **Kalman filter**.

50.3.3 The Kalman filter recursions

Iterating system (50.12) yields the Kalman filter for $t \geq 0$:

$$\begin{aligned} a_t &= y_t - G\hat{x}_t \\ K_t &= A\Sigma_t G^\top (G\Sigma_t G^\top + R)^{-1} \\ \hat{x}_{t+1} &= A\hat{x}_t + K_t a_t \\ \Sigma_{t+1} &= CC^\top + K_t R K_t^\top + (A - K_t G)\Sigma_t (A - K_t G)^\top \end{aligned} \tag{50.13}$$

Here K_t is the **Kalman gain** at time t.

50.3.4 The matrix Riccati equation

Substituting the expression for K_t from the second line of (50.13) into the fourth line gives an equivalent update formula:

$$\Sigma_{t+1} = A\Sigma_t A^\top + CC^\top - A\Sigma_t G^\top (G\Sigma_t G^\top + R)^{-1} G\Sigma_t A^\top \tag{50.14}$$

Equation (50.14) is the **matrix Riccati difference equation**.

It governs the sequence of conditional covariance matrices $\{\Sigma_t\}_{t=0}^{\infty}$ without reference to the observations $\{y_t\}$.

50.4 The Gram-Schmidt process

The random vector

$$a_t = y_t - \mathbb{E}[y_t \mid y_{t-1}, \ldots, y_0]$$

is the **innovation** of y_t with respect to y^{t-1}, the part of y_t that cannot be predicted from past observations.

Note that $\mathbb{E} a_t a_t^\top = G \Sigma_t G^\top + R$, the matrix whose inverse appears in the Kalman gain formula (50.13).

A direct calculation using $a_t = G(x_t - \hat{x}_t) + v_t$ shows that $\mathbb{E} a_t a_{t-1}^\top = 0$ and, more generally, $\mathbb{E}[a_t \mid a_{t-1}, \ldots, a_0] = 0$.

Note

An alternative argument from first principles: let $H(y^t)$ denote the closed linear span of y^t.

Since $a_{t+1} = y_{t+1} - \mathbb{E}[y_{t+1} \mid y^t]$ is a least-squares error, $a_{t+1} \perp H(y^t)$, and in particular $a_{t+1} \perp a_t$.

Thus $\{a_t\}$ is a white-noise process of innovations to $\{y_t\}$.

Sometimes (50.13) is called a **whitening filter**: it takes the signal process $\{y_t\}$ as input and produces the white-noise innovation process $\{a_t\}$ as output.

With $H(a^t)$ defined analogously, $H(a^t) = H(y^t)$, and $[a_t, \ldots, a_0]$ is an orthogonal basis for that common space.

Rather than computing $\mathbb{E}[x_t \mid y_{t-1}, \ldots, y_0]$ via one large regression, the Kalman filter performs a sequence of small regressions on successive orthogonal components of the basis $[a_{t-1}, \ldots, a_0]$, an instance of the **Gram-Schmidt procedure**.

50.5 Hidden Markov model

System (50.1)-(50.2) is an example of a **hidden Markov model**.

The observed process $\{y_t\}_{t=0}^\infty$ is *not* Markov, but the hidden process $\{x_t\}_{t=0}^\infty$ *is* Markov.

So is the process of means and covariances $\{(\hat{x}_t, \Sigma_t)\}$, which are sufficient statistics for the distribution of x_t conditional on $[y_{t-1}, \ldots, y_0]$.

50.6 Estimation

50.6.1 The innovations representation

The **innovations representation** emerging from the Kalman filter is

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + K_t a_t \\ y_t &= G\hat{x}_t + a_t \end{aligned} \tag{50.15}$$

where $\hat{x}_t = \mathbb{E}[x_t \mid y^{t-1}]$ for $t \geq 1$ and $\mathbb{E}[a_t a_t^\top \mid y^{t-1}] = G\Sigma_t G^\top + R \equiv \Omega_t$.

For $t \geq 1$, $\mathbb{E}[y_t \mid y^{t-1}] = G\hat{x}_t$ and the conditional distribution of y_t given y^{t-1} is $N(G\hat{x}_t, \Omega_t)$.

The objects $(G\hat{x}_t, \Omega_t)$ emerging from the Kalman filter recursions therefore completely characterize this conditional distribution.

50.6.2 The likelihood function

We can factor the likelihood of a sample $(y_T, y_{T-1}, \dots, y_0)$ as

$$f(y_T, \dots, y_0) = f(y_T \mid y^{T-1})\, f(y_{T-1} \mid y^{T-2}) \cdots f(y_1 \mid y_0)\, f(y_0) \tag{50.16}$$

The log conditional density of the $m \times 1$ vector y_t is

$$\log f(y_t \mid y^{t-1}) = -\frac{m}{2}\log(2\pi) - \frac{1}{2}\log\det(\Omega_t) - \frac{1}{2}\, a_t^\top \Omega_t^{-1} a_t \tag{50.17}$$

Using (50.17) and (50.13) together, we can evaluate the likelihood (50.16) recursively for any parameter vector θ that underlies the matrices A, G, C, R.

Such calculations are at the heart of efficient strategies for computing **maximum likelihood estimators** of free parameters.

50.6.3 Bayesian inference

The likelihood function is also central to **Bayesian inference**.

Where θ is the parameter vector, y_0^T the data, and $\tilde{p}(\theta)$ a prior density over θ before seeing y_0^T, Bayes' law gives the **posterior**

$$\tilde{p}(\theta \mid y_0^T) = \frac{f(y_0^T \mid \theta)\, \tilde{p}(\theta)}{\int f(y_0^T \mid \theta)\, \tilde{p}(\theta)\, d\theta}$$

The denominator is the marginal joint density $f(y_0^T)$.

50.7 Vector autoregressions and the Kalman filter

50.7.1 Convergence to a steady state

Under conditions discussed by Anderson *et al.* [1996], iterations on the Riccati equation (50.14) converge to a **time-invariant** matrix Σ from any positive semi-definite starting value Σ_0.

A time-invariant fixed point $\Sigma_t = \Sigma$ of (50.14) is the covariance matrix of x_t around

$$\mathbb{E}[x_t \mid \{y_s\}_{s \le t-1}]$$

where the conditioning extends over the **semi-infinite** past $s \le t-1$.

50.7.2 A time-invariant VAR

If the fixed point Σ exists and we initialize the filter at $\Sigma_0 = \Sigma$, the innovations representation (50.15) becomes time-invariant:

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + K a_t \\ y_t &= G\hat{x}_t + a_t \end{aligned} \tag{50.18}$$

where $\mathbb{E} a_t a_t^\top = G\Sigma G^\top + R$ and the **steady-state Kalman gain** is $K = A\Sigma G^\top (G\Sigma G^\top + R)^{-1}$.

From (50.18) we obtain $\hat{x}_{t+1} = (A - KG)\hat{x}_t + K y_t$.

If the eigenvalues of $A - KG$ are bounded in modulus strictly below unity, we can solve this equation forward to get

$$\hat{x}_{t+1} = \sum_{j=0}^{\infty} (A - KG)^j K \, y_{t-j} \tag{50.19}$$

Substituting (50.19) into the observation equation of (50.18) gives the **vector autoregression**

$$y_t = G \sum_{j=0}^{\infty} (A - KG)^j K \, y_{t-j-1} + a_t \tag{50.20}$$

where by construction

$$\mathbb{E}\left[a_t \, y_{t-j-1}^\top\right] = 0 \quad \forall \, j \geq 0 \tag{50.21}$$

The orthogonality conditions (50.21) identify (50.20) as a vector autoregression.

Letting L denote the lag operator, so that $Lx_t = x_{t-1}$, the **moving average representation** deduced from (50.18) is

$$y_t = \left[I + G(I - AL)^{-1} KL\right] a_t = \left[I + G \sum_{j=0}^{\infty} A^j K L^{j+1}\right] a_t$$

50.7.3 Interpreting VARs

Equilibria of economic models (or their linear or log-linear approximations) typically take the form of state space system (50.1).

This hidden Markov model disturbs the state x_t by the $p \times 1$ shock vector w_{t+1} and perturbs the $m \times 1$ vector of observables y_t by the $m \times 1$ measurement error v_t.

An economic theory typically makes w_{t+1} and v_t directly interpretable as shocks to preferences, technologies, endowments, or information sets.

The state space system (50.1) represents $\{y_t\}$ in terms of these **interpretable shocks**.

However, in the typical situation these shocks *cannot* be recovered directly from the y_t's, even when A, G, C, R are known.

The innovations representation (50.18) represents the *same* stochastic process $\{y_t\}$ in terms of the $m \times 1$ vector a_t of innovations that would be recovered by running an infinite-order population vector autoregression.

Its role in mapping the original representation (50.1) to the VAR (50.20) makes the Kalman filter an indispensable tool for *interpreting vector autoregressions*.

50.8 Spectral factorization identity

Because the original state space system (50.1) and the innovations representation (50.18) describe the *same* stochastic process $\{y_t\}$, they imply two distinct formulas for the **spectral density matrix** of $\{y_t\}$.

Equating those formulas yields the *spectral factorization identity*.

50.8.1 Two representations of the spectral density

Consider first the original state space system.

Writing the first line of (50.1) as $x_t = (zI - A)^{-1} C w_{t+1}$ (using the z-transform convention $z^{-1} x_t = x_{t-1}$), the covariance generating function of $\{x_t\}$ is

$$S_x(z) = (zI - A)^{-1} C C^\top (z^{-1} I - A^\top)^{-1}.$$

Since v_t is orthogonal to x_t, the spectral density of $\{y_t\}$ is

$$S_y(z) = G(zI - A)^{-1} C C^\top (z^{-1} I - A^\top)^{-1} G^\top + R. \tag{50.22}$$

Now consider the innovations representation.

The time-invariant innovations representation (50.18) gives $y_t = [G(zI - A)^{-1} K + I]\, a_t$.

Since a_t is white noise with covariance matrix $G \Sigma G^\top + R$, the spectral density is also

$$S_y(z) = \big[G(zI - A)^{-1} K + I\big](G \Sigma G^\top + R)\big[K^\top (z^{-1} I - A^\top)^{-1} G^\top + I\big]. \tag{50.23}$$

50.8.2 The spectral factorization identity

Equating (50.22) and (50.23) gives the **spectral factorization identity**:

$$G(zI - A)^{-1} C C^\top (z^{-1} I - A^\top)^{-1} G^\top + R = \big[G(zI - A)^{-1} K + I\big](G \Sigma G^\top + R)\big[K^\top (z^{-1} I - A^\top)^{-1} G^\top + I\big]. \tag{50.24}$$

The left side expresses $S_y(z)$ in terms of the **structural shocks** (w_{t+1}, v_t) and the matrices (A, C, G, R).

The right side expresses the same object as a spectral factor built from the innovations a_t and the steady-state Kalman gain K.

50.8.3 Wold and autoregressive representations

Starting from the innovations representation (50.18), we can obtain both a Wold moving average representation and an autoregressive representation.

For the Wold representation, iterate the state equation in (50.18) to express the current observation in terms of current and past innovations.

With L denoting the lag operator, (50.18) implies

$$y_t = \big[I + G(I - AL)^{-1} K L\big]\, a_t \tag{50.25}$$

This is the Wold moving-average representation of $\{y_t\}$ in terms of its one-step forecast errors.

For the autoregressive representation, invert the moving-average operator in (50.25) and solve for a_t in terms of current and past observations.

Using the identity

$$\big[I + G(I - AL)^{-1} K L\big]^{-1} = I - G\,[I - (A - KG)L]^{-1}\, KL$$

gives

$$y_t = G\big[I - (A - KG)L\big]^{-1} K\, y_{t-1} + a_t = \sum_{j=1}^{\infty} G(A - KG)^{j-1} K\, y_{t-j} + a_t, \tag{50.26}$$

which is the vector autoregression already stated in (50.20).

The key analytical fact is that, under mild stability conditions, the zeros of $\det[G(zI - A)^{-1}K + I]$ all lie *inside* the unit circle.

This ensures that the moving-average operator in (50.25) has a causal one-sided inverse.

Hence a_t lies in the closed linear span of current and past observations y^t, so a_t is the population forecast error in the VAR.

50.9 Python implementation

We now illustrate the theory using the `quantecon` library, which provides `LinearStateSpace` and `Kalman` classes that implement everything derived above.

We use the following imports:

```
import numpy as np
import matplotlib.pyplot as plt
import quantecon as qe
```

50.9.1 A scalar hidden AR(1) model

Consider a scalar hidden AR(1) state observed with measurement noise:

$$x_{t+1} = \rho\, x_t + \sigma_w\, w_{t+1}$$
$$y_t = x_t + \sigma_v\, v_t$$

with $w_t, v_t \sim N(0,1)$ IID.

Here ρ is the persistence parameter, while σ_w and σ_v are shock standard deviations.

The `LinearStateSpace` class parameterizes measurement noise by a matrix H such that $R = HH^\top$.

```
# Model parameters
ρ = 0.9
σ_w = 0.5
σ_v = 1.0

# State-space matrices
A = np.array([[ρ]])
C = np.array([[σ_w]])
G = np.array([[1.0]])
R = np.array([[σ_v**2]])

# Build a LinearStateSpace and a Kalman filter object
H = np.array([[σ_v]])   # measurement noise factor: R = H @ H.T
lss = qe.LinearStateSpace(
  A, C, G, H, mu_0=np.zeros(1), Sigma_0=np.eye(1) * 10.0)
kf = qe.Kalman(lss)
kf.set_state(np.zeros(1), np.eye(1) * 10.0)  # diffuse prior
```

We first simulate a sample path of the true hidden state and noisy observations.

```
T = 200
x_path, y_path = lss.simulate(ts_length=T, random_state=42)

# Shapes: x_path is (n, T), y_path is (m, T)
x_true = x_path[0, :]
y_obs = y_path[0, :]
```

We then run the Kalman filter manually, step by step.

The `Kalman.update` method performs a *complete* cycle, moving the prior to the filtering distribution and then on to next period's prior.

So $(\hat{x}_t, \Sigma_t)$ as defined in (50.13) are the values held by the object *before* `update` is called, and we record them there.

```
x_hats = np.zeros(T)
Sigmas = np.zeros(T)
innovations = np.zeros(T)

for t in range(T):
    x_hats[t] = kf.x_hat.item()       # x_hat_t = E[x_t | y^{t-1}]
    Sigmas[t] = kf.Sigma.item()       # Sigma_t
    innovations[t] = y_obs[t] - (G @ kf.x_hat).item()
    kf.update(y_obs[t:t+1])           # one full filter cycle
```

Getting this ordering right matters.

Recording `kf.x_hat` *after* the call would store the one-step-ahead forecast $\hat{x}_{t+1}$, and differencing it from y_t would produce a series that is not the innovation and does not have variance $G\Sigma G^\top + R$.

```
fig, axes = plt.subplots(3, 1, figsize=(10, 8), sharex=True)

t_range = np.arange(T)

axes[0].plot(t_range, x_true, lw=2,
        label='true state $x_t$')
axes[0].plot(t_range, x_hats, lw=2,
        linestyle='--', label=r'$\hat{x}_t$ (Kalman)')
axes[0].plot(t_range, y_obs, alpha=0.35, lw=2, label='observation $y_t$')
axes[0].set_title('state and observation')
axes[0].legend(fontsize=9)

axes[1].plot(t_range, Sigmas, color='C1', lw=2,
             label=r'conditional variance $\Sigma_t$')
axes[1].axhline(kf.Sigma_infinity[0, 0], ls='--', color='k',
                label=r'steady-state $\Sigma_\infty$')
axes[1].set_title('conditional variance')
axes[1].legend(fontsize=9)

axes[2].plot(t_range, innovations, color='C2', lw=2, alpha=0.7,
             label=r'innovation $a_t = y_t - G\hat{x}_t$')
axes[2].set_title('innovation')
axes[2].set_xlabel('time $t$')
axes[2].legend(fontsize=9)
fig.tight_layout()
plt.show()
```

After a short adjustment, the Kalman estimate follows the hidden state much more closely than the raw noisy observations.

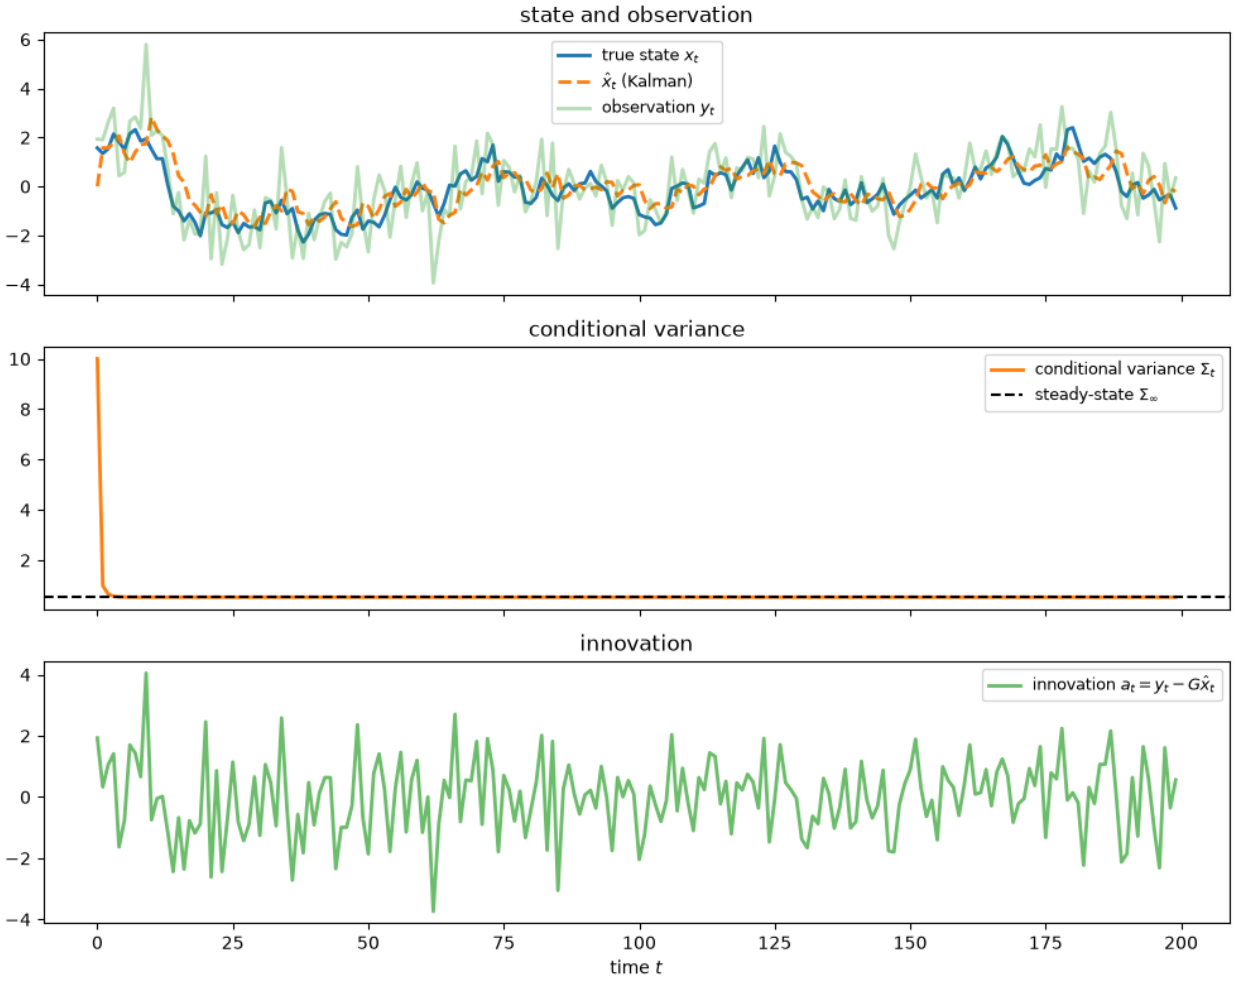


Fig. 50.1: Scalar Kalman filtering

The conditional variance drops quickly from the diffuse prior and then settles onto its steady-state value.

The innovation series fluctuates around zero, as a one-step forecast error should.

Its sample standard deviation should be close to $\sqrt{G\Sigma_\infty G^\top + R}$.

```
print(f"sample sd of innovations   = {innovations.std():.4f}")
print(f"steady-state sd            = "
      f"{np.sqrt(kf.Sigma_infinity[0, 0] + R[0, 0]):.4f}")
print(f"first-order autocorrelation = "
      f"{np.corrcoef(innovations[1:], innovations[:-1])[0, 1]:.4f}")
```

```
sample sd of innovations   = 1.2297
steady-state sd            = 1.2373
first-order autocorrelation = -0.0455
```

The autocorrelation is near zero, confirming that the filter has whitened the observed series.

50.9.2 Convergence of the Riccati equation

The `Kalman` class computes the steady-state covariance Σ_∞ by solving the discrete algebraic Riccati equation directly.

```
Sigma_inf, K_inf = kf.stationary_values()

print(f"Steady-state covariance  Σ_inf = {Sigma_inf[0, 0]:.6f}")
print(f"Kalman filter converged to Σ_t = {Sigmas[-1]:.6f}")
print(f"Steady-state Kalman gain K  = {K_inf[0, 0]:.6f}")

A_minus_KG = A - K_inf @ G
eigval = np.linalg.eigvals(A_minus_KG)[0]
print(f"\nEigenvalue of (A - KG)      = {eigval:.6f}")
print(f"Stable VAR: {np.abs(eigval) < 1}")
```

```
Steady-state covariance  Σ_inf = 0.530899
Kalman filter converged to Σ_t = 0.530899
Steady-state Kalman gain K  = 0.312110

Eigenvalue of (A - KG)      = 0.587890
Stable VAR: True
```

50.9.3 The VAR representation

Using (50.20), the coefficients in the infinite-order VAR representation are $G(A-KG)^j K$ for $j = 0, 1, 2, \ldots$

We retrieve them via `stationary_coefficients`:

```
J = 30
var_coeffs = kf.stationary_coefficients(J, coeff_type='var')

# Lag j+1 coefficient matrices
lags = np.arange(1, J + 1)
coeff_values = np.array([var_coeffs[j][0, 0] for j in range(J)])

fig, ax = plt.subplots()
ax.stem(lags, coeff_values, basefmt=' ')
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel('lag $j$')
ax.set_ylabel(r'VAR coefficient $G(A{-}KG)^{j-1}K$')
fig.tight_layout()
plt.show()
```

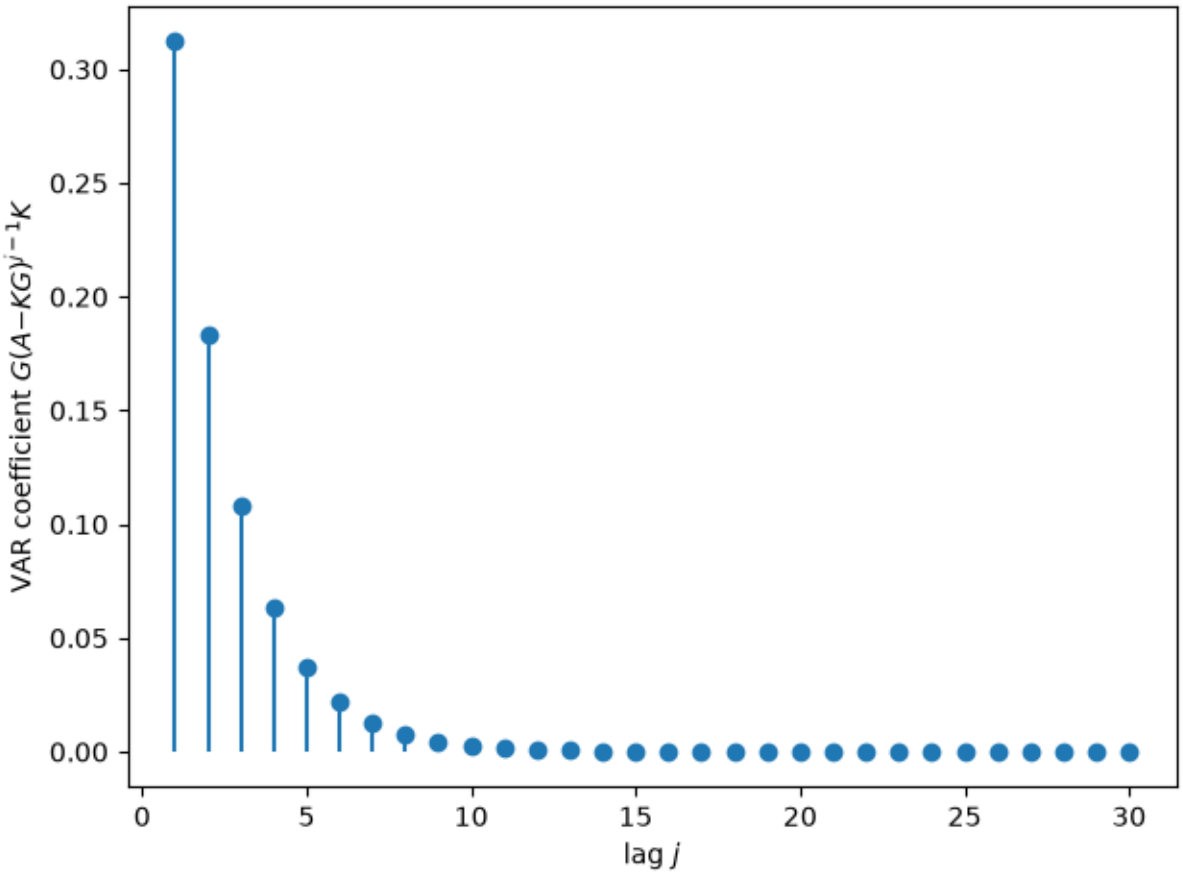


Fig. 50.2: VAR coefficients from the innovations representation

Most of the autoregressive weight is concentrated in the first few lags.

The later coefficients are nearly zero, so a short finite-lag VAR captures most of the infinite-order representation in this example.

50.9.4 Likelihood evaluation

We use (50.17) to evaluate the log-likelihood of the simulated sample.

```
def log_likelihood(A, C, G, R, y_data, x_hat_0, Sigma_0):
    """Evaluate the log-likelihood using the Kalman filter recursions."""
    H_ = np.linalg.cholesky(R)    # R = H_ @ H_.T
    lss_ = qe.LinearStateSpace(A, C, G, H_, mu_0=x_hat_0, Sigma_0=Sigma_0)
    kf_ = qe.Kalman(lss_)
    kf_.set_state(x_hat_0, Sigma_0)

    T_, m_ = y_data.shape
    loglik = 0.0

    for t in range(T_):
        x_h = kf_.x_hat
        Sig = kf_.Sigma
        Omega = G @ Sig @ G.T + R        # innovation covariance
        a_t = y_data[t] - (G @ x_h).flatten()

        sign, logdet = np.linalg.slogdet(Omega)
        loglik += -0.5 * (m_ * np.log(2 * np.pi) + logdet
                          + float(a_t @ np.linalg.solve(Omega, a_t)))
        kf_.update(y_data[t])

    return loglik


y_data_col = y_obs.reshape(-1, 1)
ll = log_likelihood(A, C, G, R,
                    y_data_col,
                    np.zeros(1), np.eye(1) * 10.0)
print(f"Log-likelihood of sample: {ll:.4f}")
```

```
Log-likelihood of sample: -325.2335
```

50.10 Where this leads

The Kalman filter maps a state space system into the VAR that a population regression of y_t on its own past would recover.

That map is most interesting when the state x_t contains variables that the econometrician does *not* see.

The sequel *Vector Autoregressions for Subsets of Variables* takes the leading case of this: Y_t follows a finite-order VAR, and the econometrician observes only a subvector $y_t = S_y Y_t$.

There we build general code that, given the VAR for Y_t and the selector matrix S_y, returns the VAR and moving average representations for y_t and expresses the innovations in the small system as a distributed lag of the innovations in the large one.

50.11 Summary

The Kalman filter recursively updates beliefs about a hidden state by combining a prior forecast with the new information contained in the current observation.

The Riccati equation tracks how the filter's conditional covariance evolves, and its steady state makes the filter time invariant.

At that steady state, the innovations representation writes the observable process in terms of one-step forecast errors that are white noise by construction.

Solving the innovations representation backward gives an infinite-order VAR, while solving it forward gives the Wold moving-average representation.

Which variables an econometrician observes determines the Kalman gain and therefore the VAR and Wold representations, even when the underlying state dynamics are held fixed.

Vector Autoregressions for Subsets of Variables pursues that observation systematically.

50.12 Exercises

Exercise 50.12.1

Consider the scalar AR(1) state space system used above with $\rho = 0.9$, $\sigma_w = 0.5$, $\sigma_v = 1.0$.

Derive an algebraic expression for the **steady-state** conditional variance Σ_∞ by solving the scalar Riccati equation (50.14) at its fixed point $\Sigma_{t+1} = \Sigma_t = \Sigma$.

Show that Σ satisfies a quadratic equation, find its positive root, and verify numerically that your formula matches `kf.Sigma_infinity`.

Solution

Here is one solution:

Setting $\Sigma_{t+1} = \Sigma_t = \Sigma$ in the scalar version of (50.14) with $A = \rho$, $CC^\top = \sigma_w^2$, $GG^\top = 1$, $R = \sigma_v^2$:

$$\Sigma = \rho^2 \Sigma + \sigma_w^2 - \frac{\rho^2 \Sigma^2}{\Sigma + \sigma_v^2}$$

Multiplying through by $\Sigma + \sigma_v^2$ and rearranging:

$$\Sigma^2 + \left[\sigma_v^2(1-\rho^2) - \sigma_w^2\right]\Sigma - \sigma_v^2 \sigma_w^2 = 0$$

Taking the positive root of this quadratic:

$$\Sigma_\infty = \frac{\sigma_w^2 - \sigma_v^2(1-\rho^2) + \sqrt{\left[\sigma_v^2(1-\rho^2) - \sigma_w^2\right]^2 + 4\sigma_v^2\sigma_w^2}}{2}$$

```
ρ_, σ_w_, σ_v_ = 0.9, 0.5, 1.0

b = σ_v_**2 * (1 - ρ_**2) - σ_w_**2
discriminant = b**2 + 4 * σ_v_**2 * σ_w_**2
Sigma_formula = (-b + np.sqrt(discriminant)) / 2

A_ = np.array([[ρ_]])
C_ = np.array([[σ_w_]])
G_ = np.array([[1.0]])
R_ = np.array([[σ_v_**2]])
H_ = np.array([[σ_v_]])   # R_ = H_ @ H_.T
lss_ = qe.LinearStateSpace(A_, C_, G_, H_, mu_0=np.zeros(1), Sigma_0=np.eye(1))
kf_ = qe.Kalman(lss_)

print(f"Analytical Σ_inf   = {Sigma_formula:.8f}")
print(f"Numerical  Σ_inf   = {kf_.Sigma_infinity[0, 0]:.8f}")
```

```
Analytical Σ_inf   = 0.53089919
Numerical  Σ_inf   = 0.53089919
```

Exercise 50.12.2

This exercise considers a two-dimensional state with a one-dimensional observation:

$$A = \begin{pmatrix} 0.9 & 0.1 \\ 0 & 0.8 \end{pmatrix}, \quad C = \begin{pmatrix} 0.4 \\ 0.1 \end{pmatrix}, \quad G = \begin{pmatrix} 1 & 0 \end{pmatrix}, \quad R = [0.5]$$

1. Simulate $T = 500$ observations from this system starting from a diffuse prior.
2. Run the Kalman filter and plot both components of $\hat{x}_t$ against the true hidden state path.
3. Compute and report the steady-state covariance Σ_∞ and Kalman gain K_∞.
4. Check that the eigenvalues of $A - K_\infty G$ lie strictly inside the unit circle, confirming that the VAR representation (50.20) is stable.

Solution

Here is one solution:

```
A2 = np.array([[0.9, 0.1],
               [0.0, 0.8]])
C2 = np.array([[0.4],
               [0.1]])
G2 = np.array([[1.0, 0.0]])
R2 = np.array([[0.5]])
H2 = np.array([[np.sqrt(0.5)]])   # R2 = H2 @ H2.T

lss2 = qe.LinearStateSpace(A2, C2, G2, H2,
                           mu_0=np.zeros(2),
                           Sigma_0=np.eye(2) * 5.0)
kf2 = qe.Kalman(lss2)
kf2.set_state(np.zeros(2), np.eye(2) * 5.0)

T2 = 500
```

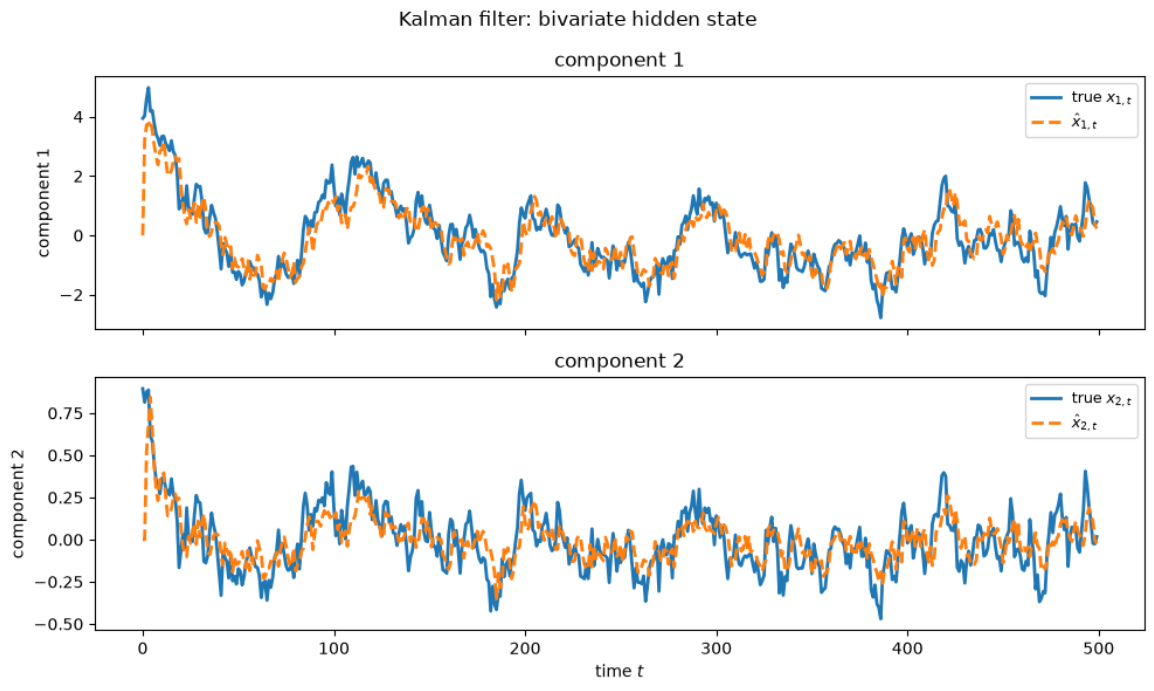
```
x2_path, y2_path = lss2.simulate(ts_length=T2, random_state=0)

x_hats2 = np.zeros((T2, 2))
for t in range(T2):
    x_hats2[t] = kf2.x_hat.ravel()      # record x_hat_t before updating
    kf2.update(y2_path[:, t])

fig, axes = plt.subplots(2, 1, figsize=(10, 6), sharex=True)
for i, ax in enumerate(axes):
    ax.plot(x2_path[i, :], lw=2, label=f'true $x_{{{i+1},t}}$')
    ax.plot(x_hats2[:, i], lw=2, ls='--', label=rf'$\hat{{x}}_{{{i+1},t}}$')
    ax.set_title(f'component {i+1}')
    ax.legend(fontsize=9)
    ax.set_ylabel(f'component {i+1}')
axes[1].set_xlabel('time $t$')
fig.suptitle('Kalman filter: bivariate hidden state')
fig.tight_layout()
plt.show()

# Steady-state values
Sigma2_inf, K2_inf = kf2.stationary_values()
print("Steady-state covariance Σ_inf:")
print(np.round(Sigma2_inf, 5))
print("\nSteady-state Kalman gain K_inf:")
print(np.round(K2_inf, 5))

# Eigenvalues of A - K_inf G
AKG2 = A2 - K2_inf @ G2
eigvals2 = np.linalg.eigvals(AKG2)
print(f"\nEigenvalues of A - K_inf G: {np.round(eigvals2, 5)}")
print(f"Stable VAR: {np.all(np.abs(eigvals2) < 1)}")
```



```
Steady-state covariance Σ_inf:
[[0.32854 0.07219]
 [0.07219 0.0166 ]]

Steady-state Kalman gain K_inf:
[[0.36559]
 [0.06971]]

Eigenvalues of A - K_inf G: [0.56394 0.77047]
Stable VAR: True
```

After the diffuse-prior transient, both Kalman estimates move closely with their corresponding hidden state paths.

The second component is not observed directly, so its tracking comes from the state dynamics and its connection to the observed signal.

The printed eigenvalues then check the separate stability condition for the VAR representation.

Exercise 50.12.3

This exercise studies likelihood and parameter estimation using the scalar model from the main text with true parameters $(\rho, \sigma_w, \sigma_v) = (0.9, 0.5, 1.0)$:

1. Simulate $T = 300$ observations.
2. Write a function that evaluates the **log-likelihood** as a function of $\rho \in (0, 1)$, holding $\sigma_w = 0.5$ and $\sigma_v = 1.0$ fixed, and plot the log-likelihood against ρ for a grid of values.
3. Locate the maximum numerically and check that it is close to the true value $\rho = 0.9$.

Solution

Here is one solution:

```
# True parameters
ρ_true, sw_true, sv_true = 0.9, 0.5, 1.0

A_t = np.array([[ρ_true]])
C_t = np.array([[sw_true]])
G_t = np.array([[1.0]])
R_t = np.array([[sv_true**2]])
H_t = np.array([[sv_true]])   # R_t = H_t @ H_t.T

lss_t = qe.LinearStateSpace(A_t, C_t, G_t, H_t,
                            mu_0=np.zeros(1), Sigma_0=np.eye(1))
_, y_sim = lss_t.simulate(ts_length=300, random_state=7)
y_sim = y_sim.T          # shape (300, 1)

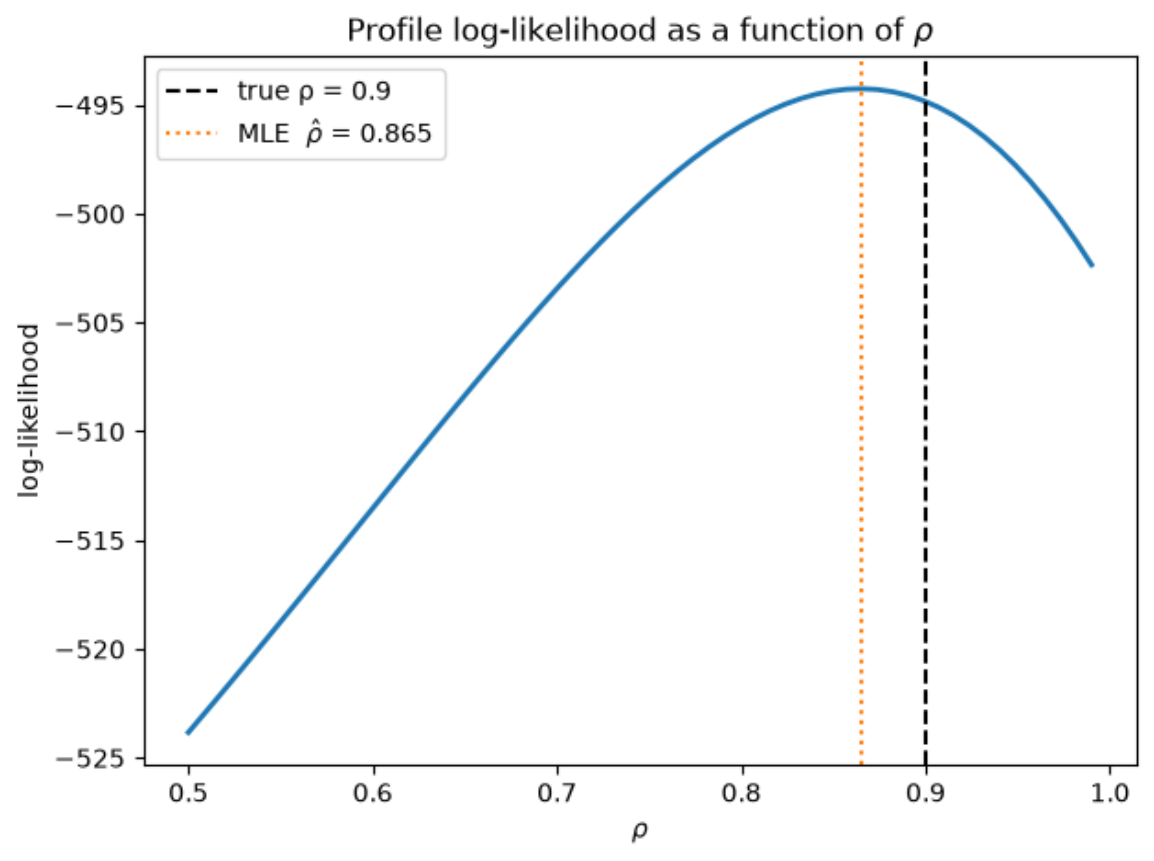
def ll_rho(ρ_val):
    A_ = np.array([[ρ_val]])
    C_ = np.array([[sw_true]])
    G_ = np.array([[1.0]])
    R_ = np.array([[sv_true**2]])
    return log_likelihood(A_, C_, G_, R_, y_sim,
                          np.zeros(1), np.eye(1) * 10.0)
```

```
ρ_grid = np.linspace(0.5, 0.99, 60)
ll_vals = np.array([ll_rho(r) for r in ρ_grid])

ρ_mle = ρ_grid[np.argmax(ll_vals)]

fig, ax = plt.subplots()
ax.plot(ρ_grid, ll_vals, lw=2)
ax.axvline(ρ_true, color='k',   ls='--', label=f'true ρ = {ρ_true}')
ax.axvline(ρ_mle, color='C1', ls=':', label=f'MLE  $\\hat{{\\rho}}$ = {ρ_mle:.3f}')
ax.set_xlabel(r'$\rho$')
ax.set_ylabel('log-likelihood')
ax.set_title('Profile log-likelihood as a function of $\\rho$')
ax.legend()
fig.tight_layout()
plt.show()

print(f"True ρ = {ρ_true},  MLE ρ_hat = {ρ_mle:.4f}")
```




```
True ρ = 0.9,  MLE ρ_hat = 0.8654
```

The likelihood curve is single-peaked and reaches its maximum close to the true value $\rho = 0.9$.

The small gap between the grid maximizer and the true value comes from finite-sample randomness and the use of a discrete grid.

CHAPTER

FIFTYONE

VECTOR AUTOREGRESSIONS FOR SUBSETS OF VARIABLES

Contents

- *Vector Autoregressions for Subsets of Variables*
  - *Overview*
  - *The large system*
  - *The small system*
  - *Four representations*
  - *Two cases where nothing is lost*
  - *Code*
  - *Example 1: a bivariate VAR(2)*
  - *Example 2: an omitted interest rate*
  - *Summary*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

51.1 Overview

An economic model delivers a vector autoregression for a list of variables Y_t.

An econometrician often observes only *some* of them.

This lecture answers three questions about that situation.

Given an mth order VAR for an $n \times 1$ vector Y_t and a selector matrix S_y that extracts an $n_y \times 1$ subvector $y_t = S_y Y_t$:

1. What vector autoregression does y_t obey?
2. What is its moving average representation?
3. How are the innovations in the small system related to the innovations in the large one?

The answers all come from the *Kalman filter*.

The state is the history of Y_t that a finite-order VAR requires, and the observation is the subvector y_t.

Because y_t is a subvector of the state and not a noisy signal about it, this is a state space system with *no* measurement error.

The main results are

- y_t obeys an **infinite-order** VAR whose coefficients we compute exactly,
- the innovation a_t of the small system is a **one-sided distributed lag** of the innovations ε_t of the large system,
- that distributed lag has a wide coefficient matrix at every lag, so ε_t *cannot* be recovered from the history of y_t,
- the forecast error variance of the small system exceeds that of the large one by a quantity we compute.

Two special cases where the small VAR stays finite-order are identified and verified.

This lecture generalizes the worked example that used to close *The Kalman Filter and Vector Autoregressions*.

Let's start with imports.

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe

plt.rcParams['figure.figsize'] = (10, 5)
np.set_printoptions(precision=4, suppress=True)
```

51.2 The large system

Let Y_t be $n \times 1$ and suppose it obeys the mth order vector autoregression

$$Y_{t+1} = A_1 Y_t + A_2 Y_{t-1} + \cdots + A_m Y_{t-m+1} + \varepsilon_{t+1}, \qquad \mathbb{E}\,\varepsilon_t \varepsilon_t^\top = V, \tag{51.1}$$

where $\{\varepsilon_t\}$ is a serially uncorrelated sequence with $\mathbb{E}[\varepsilon_{t+1} \mid Y_t, Y_{t-1}, ...] = 0$.

Stack m lags into the $nm \times 1$ state vector

$$X_t = \begin{pmatrix} Y_t \\ Y_{t-1} \\ \vdots \\ Y_{t-m+1} \end{pmatrix}.$$

Then (51.1) becomes the first-order **companion form**

$$X_{t+1} = A X_t + C \varepsilon_{t+1}, \qquad A = \begin{pmatrix} A_1 & A_2 & \cdots & A_{m-1} & A_m \\ I & 0 & \cdots & 0 & 0 \\ 0 & I & \cdots & 0 & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & 0 & \cdots & I & 0 \end{pmatrix}, \qquad C = \begin{pmatrix} I_n \\ 0 \\ \vdots \\ 0 \end{pmatrix}. \tag{51.2}$$

We write A without a subscript for the companion matrix and $A_1, \ldots, A_m$ with subscripts for the VAR coefficient matrices.

Let $J = \begin{pmatrix} I_n & 0 & \cdots & 0 \end{pmatrix}$ be the $n \times nm$ matrix that reads Y_t off the state, so that $Y_t = J X_t$ and $C = J^\top$.

We assume all eigenvalues of A are strictly inside the unit circle, so $\{Y_t\}$ is covariance stationary.

51.3 The small system

The econometrician observes

$$y_t = S_y Y_t = G X_t, \qquad G = S_y J, \tag{51.3}$$

where S_y is $n_y \times n$ with $n_y < n$.

Usually S_y picks out n_y coordinates of Y_t, but nothing below requires that, so linear combinations are allowed too.

Equations (51.2) and (51.3) are a state space system of the form studied in *The Kalman Filter and Vector Autoregressions*, with shock loading C, observation matrix G, and measurement error covariance

$$R = 0.$$

The observation is *exact*, but it is a strict subvector of the state, so the econometrician still faces a filtering problem.

Let Σ solve the steady-state Riccati equation

$$\Sigma = A\Sigma A^\top + CVC^\top - A\Sigma G^\top \left(G\Sigma G^\top\right)^{-1} G\Sigma A^\top \tag{51.4}$$

with associated Kalman gain

$$K = A\Sigma G^\top \left(G\Sigma G^\top\right)^{-1}. \tag{51.5}$$

Here Σ is the covariance matrix of $X_t - \mathbb{E}[X_t \mid y^{t-1}]$.

Note

Because X_t contains lags of Y_t whose S_y components the econometrician has already seen exactly, Σ is singular.

That is harmless.

What the Kalman gain (51.5) requires is that the *innovation* covariance $G\Sigma G^\top$ be nonsingular, which holds whenever no linear combination of y_t is perfectly predictable from y^{t-1}.

The innovation in the small system is

$$a_t = y_t - \mathbb{E}[y_t \mid y^{t-1}] = G(X_t - \hat{X}_t), \qquad \Omega \equiv \mathbb{E}\, a_t a_t^\top = G\Sigma G^\top. \tag{51.6}$$

51.4 Four representations

51.4.1 The Wold representation

The steady-state innovations representation derived in *The Kalman Filter and Vector Autoregressions* is

$$\hat{X}_{t+1} = A\hat{X}_t + K a_t, \qquad y_t = G\hat{X}_t + a_t. \tag{51.7}$$

Solving (51.7) forward gives the moving average representation of y_t in terms of its own innovations,

$$y_t = \sum_{h=0}^{\infty} \Psi_h \, a_{t-h}, \qquad \Psi_0 = I_{n_y}, \qquad \Psi_h = GA^{h-1}K \quad (h \geq 1). \tag{51.8}$$

51.4.2 The vector autoregression

Solving (51.7) backward instead gives

$$y_t = \sum_{j=1}^{\infty} B_j \, y_{t-j} + a_t, \qquad B_j = G(A - KG)^{j-1} K. \tag{51.9}$$

This is an infinite-order VAR, convergent because the eigenvalues of $A - KG$ lie inside the unit circle.

51.4.3 Innovations of the small system in terms of the large one

This is the question that motivates the lecture.

Let $e_t = X_t - \hat{X}_t$ be the filtering error.

Subtracting the Kalman recursion $\hat{X}_{t+1} = A\hat{X}_t + Ka_t$ from the state equation (51.2) and using $a_t = Ge_t$ gives

$$e_{t+1} = (A - KG)e_t + C\varepsilon_{t+1}. \tag{51.10}$$

Solving (51.10) backward and premultiplying by G yields the answer.

Proposition 51.4.1

The innovations of the small system are the one-sided distributed lag

$$a_t = \sum_{j=0}^{\infty} \Gamma_j \, \varepsilon_{t-j}, \qquad \Gamma_j = G(A - KG)^j C, \tag{51.11}$$

of the innovations of the large system, with leading coefficient

$$\Gamma_0 = GC = S_y J J^\top = S_y.$$

Since each Γ_j is $n_y \times n$ with $n_y < n$, the map from $\{\varepsilon_t\}$ to $\{a_t\}$ has no inverse.

Knowing the entire history of y_t is not enough to recover ε_t.

51.4.4 The structural moving average

For comparison, iterating (51.2) gives y_t directly in terms of the large system's innovations,

$$y_t = \sum_{j=0}^{\infty} \Phi_j \, \varepsilon_{t-j}, \qquad \Phi_j = GA^j C. \tag{51.12}$$

51.4.5 A forecast error that contains past shocks

Separating the $j = 0$ term in (51.11) and using $\Gamma_0 = S_y$ gives

$$a_t = \underbrace{S_y \varepsilon_t}_{\text{full-information forecast error}} + \underbrace{\sum_{j=1}^{\infty} \Gamma_j \, \varepsilon_{t-j}}_{\text{shocks realized before } t}. \tag{51.13}$$

The first term is what it appears to be.

Because $\mathbb{E}[\varepsilon_t \mid Y^{t-1}] = 0$, we have $y_t - \mathbb{E}[y_t \mid Y^{t-1}] = S_y \varepsilon_t$, so $S_y \varepsilon_t$ is the error made in forecasting y_t by someone who observes the *entire* history of Y.

The second term deserves a pause, because at first sight it looks impossible.

By construction a_t is a forecast error, orthogonal to everything known at $t-1$.

Yet (51.13) says that a_t loads on $\varepsilon_{t-1}, \varepsilon_{t-2}, \ldots$, shocks that had already been realized by then.

The two facts are consistent, and both are true:

$$\mathbb{E}\, a_t\, \varepsilon_{t-j}^\top = \Gamma_j V \neq 0 \quad (j \geq 1), \qquad \text{while} \qquad \mathbb{E}\, a_t\, y_{t-k}^\top = 0 \quad (k \geq 1). \tag{51.14}$$

The resolution is that past shocks are known to the *large* system's econometrician, not to the small one.

The small econometrician's information set is $H(y^{t-1})$, the closed linear span of $y_{t-1}, y_{t-2}, \ldots$, and ε_{t-j} does not lie in it.

Let P_{t-1} denote projection onto $H(y^{t-1})$.

Applying P_{t-1} to (51.13), using $P_{t-1} a_t = 0$ and $P_{t-1} \varepsilon_t = 0$, delivers the identity

$$\sum_{j=1}^{\infty} \Gamma_j\, P_{t-1} \varepsilon_{t-j} = 0. \tag{51.15}$$

So the distributed lag in (51.13) loads only on the parts of past shocks that the small econometrician has *not yet* learned,

$$a_t = S_y \varepsilon_t + \sum_{j=1}^{\infty} \Gamma_j (\varepsilon_{t-j} - P_{t-1} \varepsilon_{t-j}). \tag{51.16}$$

A shock that happened three quarters ago can still be news today, if the only series you watch has not finished revealing it.

That is the whole content of the filtering problem, and it is why Ω exceeds $S_y V S_y^\top$.

51.4.6 What the coefficients Γ_j are

Substituting the moving average (51.12) into the autoregression (51.9) and matching the coefficient on ε_{t-k} gives a second formula for the same objects,

$$\Gamma_k = \Phi_k - \sum_{j=1}^{k} B_j\, \Phi_{k-j}. \tag{51.17}$$

So Γ_k measures the failure of the small system's own autoregression to reproduce the large system's k-lag response.

The case $k = 1$ is worth writing out.

Since $\Phi_0 = S_y$ and $\Phi_1 = GAC = S_y A_1$,

$$\Gamma_1 = S_y A_1 - B_1 S_y. \tag{51.18}$$

Suppose S_y selects coordinates, and partition $Y_t = (y_t^\top, \tilde{y}_t^\top)^\top$ as before.

Then $B_1 S_y$ has zeros in the columns belonging to the dropped variables, so (51.18) reads

$$\Gamma_1 = \begin{pmatrix} A_1^{yy} - B_1 & A_1^{y\tilde{y}} \end{pmatrix}. \tag{51.19}$$

The loading of a_t on last period's *omitted* shocks is exactly $A_1^{y\tilde{y}}$, the block through which the omitted variables enter the retained equations.

Equation (51.19) also previews *Proposition 51.5.2*: block exogeneity sets $A_1^{y\tilde{y}} = 0$, which forces $B_1 = A_1^{yy}$ and hence $\Gamma_1 = 0$.

51.4.7 A factorization identity

Representations (51.8), (51.11), and (51.12) describe the same process, so with $\Psi(z) = \sum_h \Psi_h z^h$ and similarly for Γ and Φ,

$$\Phi(z) = \Psi(z)\,\Gamma(z). \tag{51.20}$$

Equation (51.20) says the structural moving average operator factors into the Wold operator of the small system times the innovation map.

It is a sharp numerical check on everything above, and we use it as one.

Matching variances at each lag also gives

$$\Omega = \sum_{j=0}^{\infty} \Gamma_j V \Gamma_j^\top = S_y V S_y^\top + \sum_{j=1}^{\infty} \Gamma_j V \Gamma_j^\top. \tag{51.21}$$

Every term in the second sum is positive semidefinite, so

$$\Omega \succeq S_y V S_y^\top.$$

The small system's one-step forecast error variance is never smaller than the corresponding block of the large system's, and (51.21) says exactly how much is lost.

51.5 Two cases where nothing is lost

Proposition 51.5.1

If $S_y = I_n$, then $\Gamma_0 = I_n$ and $\Gamma_j = 0$ for $j \geq 1$, so $a_t = \varepsilon_t$, and (51.9) collapses to the original mth order VAR (51.1).

Nothing is hidden, so the Wold innovations *are* the structural innovations.

The second case is more interesting.

Partition $Y_t = (y_t^\top, \tilde{y}_t^\top)^\top$ and correspondingly

$$A_k = \begin{pmatrix} A_k^{yy} & A_k^{y\tilde{y}} \\ A_k^{\tilde{y}y} & A_k^{\tilde{y}\tilde{y}} \end{pmatrix}, \qquad k = 1, \dots, m.$$

Proposition 51.5.2

Suppose y_t is **block exogenous**, meaning $A_k^{y\tilde{y}} = 0$ for all k, so that no lag of the omitted variables appears in the equations for y.

Then $\Gamma_j = 0$ for $j \geq 1$, $a_t = S_y \varepsilon_t$, $\Omega = S_y V S_y^\top$, and y_t obeys the finite mth order VAR

$$y_t = \sum_{k=1}^{m} A_k^{yy}\, y_{t-k} + a_t.$$

Proof. Block exogeneity makes the y rows of (51.1) read $y_{t+1} = \sum_k A_k^{yy} y_{t+1-k} + S_y \varepsilon_{t+1}$.

The right side involves only lags of y, and $S_y \varepsilon_{t+1}$ is orthogonal to the whole history Y^t and hence to y^t.

So this *is* the projection of y_{t+1} on y^t, which identifies it as the Wold representation.

Note what *Proposition 51.5.2* does *not* require: V need not be block diagonal.

Contemporaneous correlation between ε^y and $\varepsilon^{\tilde{y}}$ is fine.

What matters for whether an omitted variable damages a VAR is Granger causality, not contemporaneous correlation.

51.6 Code

The class below packages everything.

```
class VARSubsystem:
    """
    A VAR for Y and the implied representations for a subvector y = S_y Y.

        Y[t+1] = A_1 Y[t] + ... + A_m Y[t-m+1] + eps[t+1],  E eps eps' = V
        y[t]   = S_y Y[t]

    Parameters
    ----------
    A_list : list of (n, n) arrays, the VAR coefficient matrices A_1, ..., A_m
    V      : (n, n) array, covariance matrix of eps
    S_y    : (n_y, n) selector matrix
    """

    def __init__(self, A_list, V, S_y):
        self.A_list = [np.atleast_2d(np.asarray(a, dtype=float)) for a in A_list]
        self.V = np.atleast_2d(np.asarray(V, dtype=float))
        self.S_y = np.atleast_2d(np.asarray(S_y, dtype=float))
        n, m = self.A_list[0].shape[0], len(self.A_list)
        self.n, self.m, self.n_y = n, m, self.S_y.shape[0]

        # companion form
        self.A = np.zeros((n * m, n * m))
        self.A[:n] = np.hstack(self.A_list)
        if m > 1:
            self.A[n:, :n * (m - 1)] = np.eye(n * (m - 1))
        self.J = np.zeros((n, n * m))
        self.J[:, :n] = np.eye(n)
        self.C = self.J.T
        self.G = self.S_y @ self.J
        self.Q = self.C @ self.V @ self.C.T
        self._Sigma = self._K = None

    def companion_eigenvalues(self):
        return np.linalg.eigvals(self.A)

    def stationary_filter(self):
        """Steady-state (Sigma, K) from the Riccati equation with R = 0."""
```

(continues on next page)

(continued from previous page)

```
        if self._Sigma is None:
            A, G = self.A, self.G
            R = np.zeros((self.n_y, self.n_y))
            Sigma = qe.solve_discrete_riccati(A.T, G.T, self.Q, R)
            Omega = G @ Sigma @ G.T
            self._Sigma = Sigma
            self._K = A @ Sigma @ G.T @ np.linalg.inv(Omega)
        return self._Sigma, self._K

    def innovation_cov(self):
        """Omega = E a a', the one-step forecast error covariance of y."""
        Sigma, _ = self.stationary_filter()
        return self.G @ Sigma @ self.G.T

    def wold(self, h_max=20):
        """Psi[h] in y[t] = sum_h Psi[h] a[t-h]; Psi[0] = I."""
        _, K = self.stationary_filter()
        Psi = np.empty((h_max + 1, self.n_y, self.n_y))
        Psi[0], P = np.eye(self.n_y), np.eye(self.A.shape[0])
        for h in range(1, h_max + 1):
            Psi[h] = self.G @ P @ K
            P = P @ self.A
        return Psi

    def var_coefficients(self, h_max=20):
        """B[j-1] in y[t] = sum_j B[j] y[t-j] + a[t], for j = 1, ..., h_max."""
        _, K = self.stationary_filter()
        M = self.A - K @ self.G
        B, P = np.empty((h_max, self.n_y, self.n_y)), np.eye(self.A.shape[0])
        for j in range(h_max):
            B[j] = self.G @ P @ K
            P = P @ M
        return B

    def innovation_map(self, h_max=20):
        """Gamma[j] in a[t] = sum_j Gamma[j] eps[t-j]."""
        _, K = self.stationary_filter()
        M = self.A - K @ self.G
        Gamma, P = np.empty((h_max + 1, self.n_y, self.n)), np.eye(self.A.shape[0])
        for j in range(h_max + 1):
            Gamma[j] = self.G @ P @ self.C
            P = P @ M
        return Gamma

    def structural_ma(self, h_max=20):
        """Phi[j] in y[t] = sum_j Phi[j] eps[t-j]."""
        Phi, P = np.empty((h_max + 1, self.n_y, self.n)), np.eye(self.A.shape[0])
        for j in range(h_max + 1):
            Phi[j] = self.G @ P @ self.C
            P = P @ self.A
        return Phi

    def simulate(self, T, seed=0, burn=200):
        """Simulate Y and the innovations eps that generated it."""
        rng = np.random.default_rng(seed)
        L = np.linalg.cholesky(self.V)
        eps = rng.standard_normal((T + burn, self.n)) @ L.T
```

(continues on next page)

(continued from previous page)

```
        X, Y = np.zeros(self.n * self.m), np.zeros((T + burn, self.n))
        for t in range(T + burn):
            X = self.A @ X + self.C @ eps[t]
            Y[t] = self.J @ X
        return Y[burn:], eps[burn:]

    def filter_innovations(self, y_path):
        """Recover a[t] from observed y by running the steady-state filter."""
        _, K = self.stationary_filter()
        x_hat = np.zeros(self.A.shape[0])
        a = np.empty((len(y_path), self.n_y))
        for t in range(len(y_path)):
            a[t] = y_path[t] - self.G @ x_hat
            x_hat = self.A @ x_hat + K @ a[t]
        return a


def convolve(Psi, Gamma, h_max):
    """(Psi * Gamma)[h] = sum_{k=0}^{h} Psi[k] Gamma[h-k]."""
    out = np.zeros((h_max + 1, Psi.shape[1], Gamma.shape[2]))
    for h in range(h_max + 1):
        for k in range(h + 1):
            out[h] += Psi[k] @ Gamma[h - k]
    return out
```

A single routine collects the diagnostics we want to see for every example.

```
def report(model, h_max=30, label=''):
    """Print the identities that every subsystem must satisfy."""
    Sigma, K = model.stationary_filter()
    A, G, V = model.A, model.G, model.V
    resid = Sigma - (A @ Sigma @ A.T + model.Q
                     - A @ Sigma @ G.T @ np.linalg.inv(G @ Sigma @ G.T)
                       @ G @ Sigma @ A.T)
    Psi = model.wold(h_max)
    Gamma = model.innovation_map(h_max)
    Phi = model.structural_ma(h_max)
    Omega, Vy = model.innovation_cov(), model.S_y @ V @ model.S_y.T

    print(f'--- {label} (n = {model.n}, m = {model.m}, n_y = {model.n_y})')
    print(f'  max |eig| of companion A      {np.max(abs(model.companion_
↪eigenvalues())):.6f}')
    print(f'  max |eig| of A - KG           {np.max(abs(np.linalg.eigvals(A - K @␣
↪G))):.6f}')
    print(f'  Riccati residual              {np.abs(resid).max():.2e}')
    print(f'  |Gamma[0] - S_y|              {np.abs(Gamma[0] - model.S_y).max():.2e}')
    print(f'  |Phi - Psi * Gamma|           '
          f'{np.abs(convolve(Psi, Gamma, h_max) - Phi).max():.2e}')
    Gamma_long = model.innovation_map(300)      # the sum in (SS) is infinite
    print(f'  |Omega - sum Gamma V Gamma\'|  '
          f'{np.abs(Omega - sum(Gamma_long[j] @ V @ Gamma_long[j].T for j in␣
↪range(301))).max():.2e}')
    print(f'  max |Gamma[j]|, j >= 1        {np.abs(Gamma[1:]).max():.3e}')
    print(f'  det Omega / det S_y V S_y\'    '
          f'{np.linalg.det(Omega) / np.linalg.det(Vy):.4f}')
    return Psi, Gamma, Phi
```

51.7 Example 1: a bivariate VAR(2)

Two observable series r_t and z_t obey the VAR(2)

$$\begin{pmatrix} r_{t+1} \\ z_{t+1} \end{pmatrix} = A_1 \begin{pmatrix} r_t \\ z_t \end{pmatrix} + A_2 \begin{pmatrix} r_{t-1} \\ z_{t-1} \end{pmatrix} + \varepsilon_{t+1},$$

with

$$A_1 = \begin{pmatrix} 0.80 & 0.75 \\ 0 & 0.75 \end{pmatrix}, \qquad A_2 = \begin{pmatrix} 0.05 & -0.72 \\ 0 & 0.20 \end{pmatrix}, \qquad V = I_2.$$

Note that z is block exogenous: its equation contains no lags of r.

But r is *not*, since z enters the r equation with both lags.

So dropping z should matter, while dropping r should not.

```
A1 = np.array([[0.80,  0.75],
               [0.00,  0.75]])
A2 = np.array([[0.05, -0.72],
               [0.00,  0.20]])
V2 = np.eye(2)

S_both = np.eye(2)                    # observe (r, z)
S_r    = np.array([[1.0, 0.0]])       # observe r only
S_z    = np.array([[0.0, 1.0]])       # observe z only

mod_both = VARSubsystem([A1, A2], V2, S_both)
mod_r    = VARSubsystem([A1, A2], V2, S_r)
mod_z    = VARSubsystem([A1, A2], V2, S_z)

Psi_both, Gam_both, _ = report(mod_both, label='observe r and z')
print()
Psi_r, Gam_r, _ = report(mod_r, label='observe r only')
print()
Psi_z, Gam_z, _ = report(mod_z, label='observe z only')
```

```
--- observe r and z (n = 2, m = 2, n_y = 2)
  max |eig| of companion A      0.958631
  max |eig| of A - KG           0.000000
  Riccati residual              0.00e+00
  |Gamma[0] - S_y|              0.00e+00
  |Phi - Psi * Gamma|           0.00e+00
  |Omega - sum Gamma V Gamma'|  0.00e+00
  max |Gamma[j]|, j >= 1        0.000e+00
  det Omega / det S_y V S_y'    1.0000

--- observe r only (n = 2, m = 2, n_y = 1)
  max |eig| of companion A      0.958631
  max |eig| of A - KG           0.959007
  Riccati residual              1.78e-15
  |Gamma[0] - S_y|              0.00e+00
  |Phi - Psi * Gamma|           2.22e-16
  |Omega - sum Gamma V Gamma'|  2.00e-15
  max |Gamma[j]|, j >= 1        7.500e-01
  det Omega / det S_y V S_y'    1.5786
```

(continues on next page)

(continued from previous page)

```
--- observe z only (n = 2, m = 2, n_y = 1)
  max |eig| of companion A      0.958631
  max |eig| of A - KG           0.858258
  Riccati residual              4.44e-16
  |Gamma[0] - S_y|              0.00e+00
  |Phi - Psi * Gamma|           0.00e+00
  |Omega - sum Gamma V Gamma'|  0.00e+00
  max |Gamma[j]|, j >= 1        0.000e+00
  det Omega / det S_y V S_y'    1.0000
```

Every identity holds to machine precision.

The three cases differ exactly as the propositions predict.

Observing both variables gives $\Gamma_j = 0$ for $j \geq 1$, so $a_t = \varepsilon_t$, as *Proposition 51.5.1* requires.

Observing only z, which is block exogenous, also gives $\Gamma_j = 0$ for $j \geq 1$, so $a_t = \varepsilon_{z,t}$, as *Proposition 51.5.2* requires.

Observing only r is different.

Here $\Gamma_j \neq 0$ for $j \geq 1$, and the ratio of forecast error variances reports how much the r-only econometrician loses.

```
print('observe r only:')
print(f'  Omega           = {mod_r.innovation_cov()[0, 0]:.4f}')
print(f'  S_y V S_y\'      = {(S_r @ V2 @ S_r.T)[0, 0]:.4f}')
print('\n  Gamma[j] for j = 0, ..., 5 (rows: response of a to eps_r, eps_z)')
print(Gam_r[:6, 0, :])
```

```
observe r only:
  Omega           = 1.5786
  S_y V S_y'      = 1.0000

  Gamma[j] for j = 0, ..., 5 (rows: response of a to eps_r, eps_z)
[[ 1.      0.    ]
 [ 0.0769  0.75  ]
 [-0.0097 -0.0998]
 [ 0.0017  0.0125]
 [ 0.0002 -0.0023]
 [ 0.0004 -0.0003]]
```

The forecast error variance is over 50 percent larger than V_{11}.

The extra variance is entirely attributable to past ε_z shocks that the econometrician sees only through their effect on r.

51.7.1 The VAR for the subsystem

```
B_r = mod_r.var_coefficients(12)

fig, ax = plt.subplots()
ax.stem(np.arange(1, 13), B_r[:, 0, 0], basefmt=' ')
ax.axhline(0, color='k', lw=0.6)
ax.set_xlabel('lag $j$')
ax.set_ylabel('$B_j$')
ax.set_title(r'Population VAR coefficients for $r_t$ when only $r$ is observed')
fig.tight_layout()
plt.show()
```

(continues on next page)

(continued from previous page)

```
print('B_1, ..., B_6:', np.round(B_r[:6, 0, 0], 5))
print('\nfor comparison, A_1[0,0] and A_2[0,0]:', A1[0, 0], A2[0, 0])
```

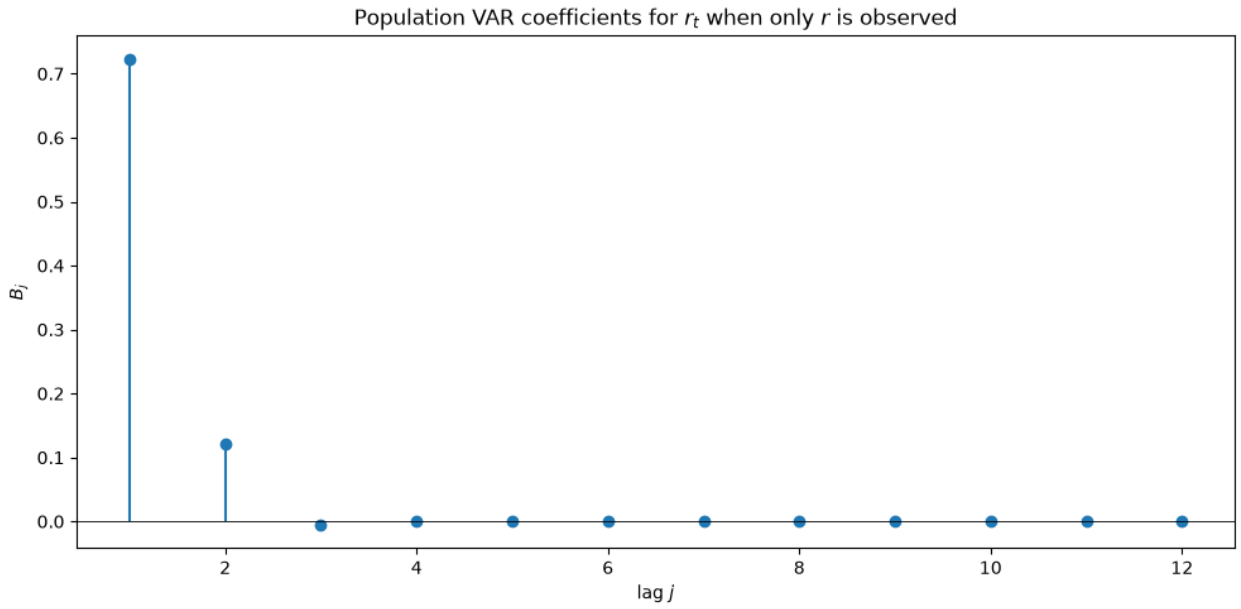


Fig. 51.1: VAR coefficients for the subsystem

```
B_1, ..., B_6: [ 0.7231  0.1212 -0.0056  0.0007 -0.0001 -0.    ]

for comparison, A_1[0,0] and A_2[0,0]: 0.8 0.05
```

The infinite-order VAR for r alone is dominated by two lags, but neither coefficient equals the corresponding coefficient in the bivariate system.

Dropping z does not simply delete the z columns of the VAR; it changes what is left.

51.7.2 Wold impulse responses

```
H = 25
h = np.arange(H + 1)
Psi_both = mod_both.wold(H)

fig, axes = plt.subplots(2, 2, figsize=(10, 6), sharex=True)
names, shocks = [r'$r_t$', r'$z_t$'], [r'$a_{r,t}$', r'$a_{z,t}$']
for i in range(2):
    for j in range(2):
        axes[i, j].plot(h, Psi_both[:, i, j], lw=2)
        axes[i, j].axhline(0, color='k', lw=0.6, ls='--')
        axes[i, j].set_title(f'{names[i]} to {shocks[j]}', fontsize=10)
        if i == 1:
            axes[i, j].set_xlabel('horizon $h$')
fig.suptitle('Wold responses, both variables observed')
fig.tight_layout()
plt.show()
```

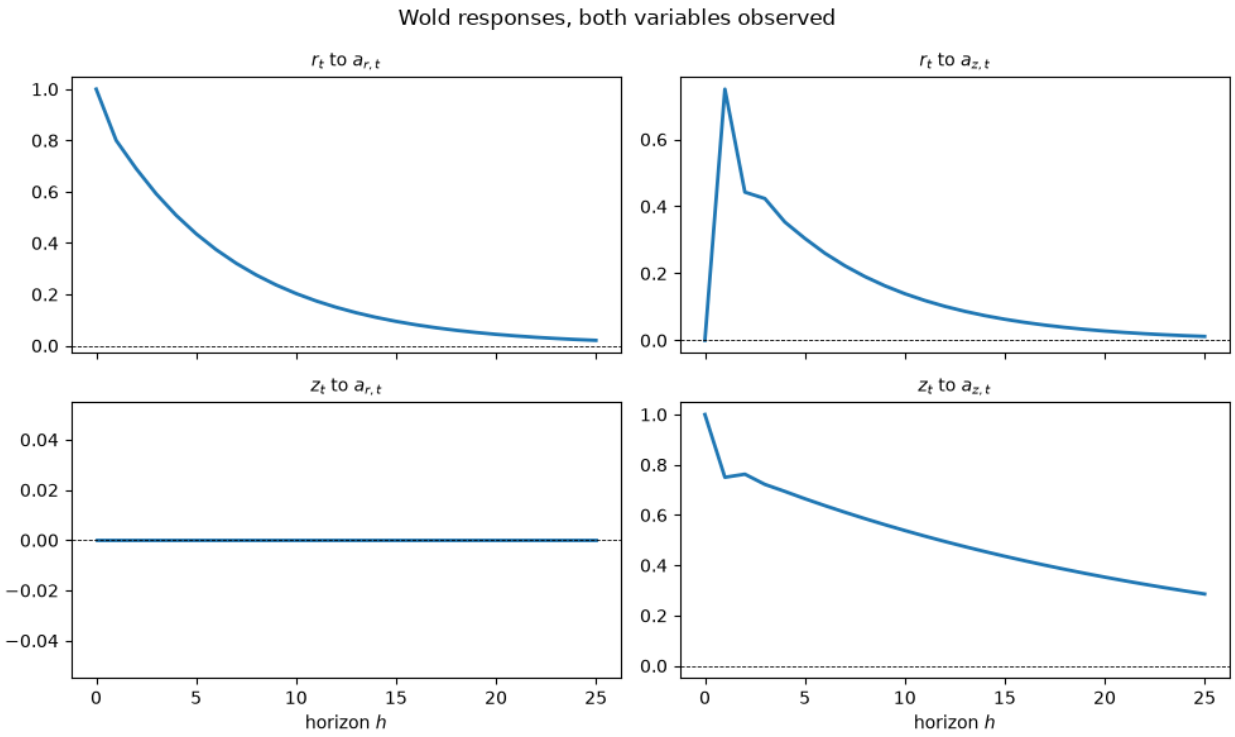


Fig. 51.2: Wold responses when both variables are observed

Because $a_t = \varepsilon_t$ here, these Wold responses coincide with the structural responses of the bivariate VAR.

```
Psi_r = mod_r.wold(H)
Phi_r = mod_r.structural_ma(H)

fig, ax = plt.subplots()
ax.plot(h, Psi_r[:, 0, 0], lw=2, label=r'$\Psi_h$: $r_t$ to its own innovation $a_t$')
ax.plot(h, Phi_r[:, 0, 0], lw=2, ls='--',
        label=r'$\Phi_h$: $r_t$ to $\varepsilon_{r,t}$')
ax.plot(h, Phi_r[:, 0, 1], lw=2, ls=':',
        label=r'$\Phi_h$: $r_t$ to $\varepsilon_{z,t}$')
ax.axhline(0, color='k', lw=0.6)
ax.set_xlabel('horizon $h$')
ax.set_ylabel('response')
ax.set_title(r'Wold versus structural responses of $r_t$')
ax.legend()
fig.tight_layout()
plt.show()
```

The Wold response to a_t is not the response to either structural shock.

It is a blend, and (51.20) says exactly how the blending works.

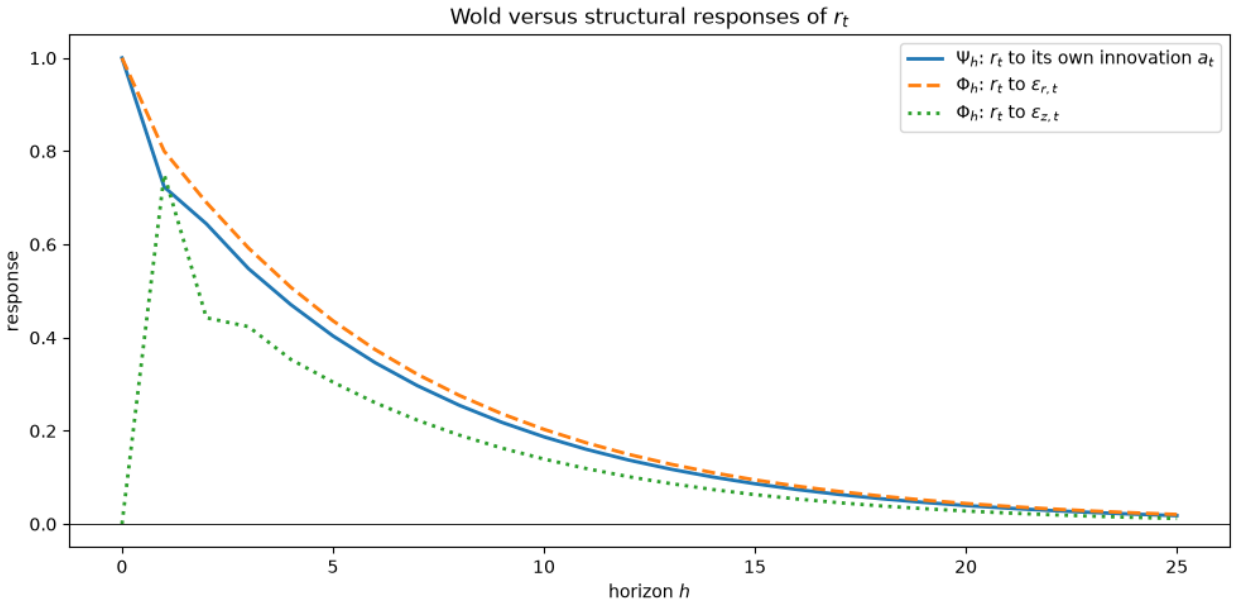


Fig. 51.3: Wold response when only r is observed

51.7.3 Checking the innovation map by simulation

Proposition *Proposition 51.4.1* is a statement about population objects.

We check it on a long simulated sample by comparing the innovations that the Kalman filter extracts from the observed r series with the distributed lag $\sum_j \Gamma_j \varepsilon_{t-j}$ built from the shocks that generated it.

```
T = 100_000
Y_sim, eps_sim = mod_r.simulate(T, seed=1)
y_sim = Y_sim @ S_r.T

a_filtered = mod_r.filter_innovations(y_sim)

J_lag = 400          # the distributed lag is infinite, so truncate generously
Gam_long = mod_r.innovation_map(J_lag)
a_theory = np.zeros_like(a_filtered)
for j in range(J_lag + 1):
    a_theory[j:] += eps_sim[:T - j] @ Gam_long[j].T

burn = J_lag + 1
print(f'correlation of the two series   '
      f'{np.corrcoef(a_filtered[burn:, 0], a_theory[burn:, 0])[0, 1]:.8f}')
print(f'max absolute difference         '
      f'{np.abs(a_filtered[burn:] - a_theory[burn:]).max():.2e}')
print(f'sample variance of a            {a_filtered[burn:, 0].var():.4f}')
print(f'population Omega                {mod_r.innovation_cov()[0, 0]:.4f}')
print(f'sample corr(a_t, a_(t-1))       '
      f'{np.corrcoef(a_filtered[burn + 1:, 0], a_filtered[burn:-1, 0])[0, 1]:.4f}')
```

```
correlation of the two series   1.00000000
max absolute difference         6.45e-10
sample variance of a            1.5706
population Omega                1.5786
```

(continues on next page)

(continued from previous page)

```
sample corr(a_t, a_(t-1))       -0.0048
```

The two constructions of a_t agree to the accuracy of the truncated distributed lag, the sample variance of a_t matches Ω, and a_t is serially uncorrelated.

51.7.4 The two orthogonality facts

Now we check (51.14) directly, by running two regressions on the simulated data.

The first regresses a_t on current and lagged *structural* shocks, which the small econometrician cannot see.

The second regresses a_t on lagged *observations*, which are all that the small econometrician can see.

```
P_lags = 4
Z_eps = np.column_stack([eps_sim[P_lags - l:T - l] for l in range(P_lags + 1)])
target = a_filtered[P_lags:, 0]
b_eps = np.linalg.lstsq(Z_eps, target, rcond=None)[0].reshape(P_lags + 1, 2)
fit = Z_eps @ b_eps.ravel()
r2_eps = 1 - ((target - fit) ** 2).sum() / target.var() / len(target)

print('OLS of a_t on eps_t, ..., eps_{t-4}:')
print(b_eps)
print('population Gamma_0, ..., Gamma_4:')
print(Gam_long[:P_lags + 1, 0, :])
print(f'max discrepancy {np.abs(b_eps - Gam_long[:P_lags + 1, 0, :]).max():.2e}')
print(f'R^2 = {r2_eps:.6f}')
```

```
OLS of a_t on eps_t, ..., eps_{t-4}:
[[ 1.     -0.    ]
 [ 0.0769  0.75  ]
 [-0.0097 -0.0998]
 [ 0.0017  0.0125]
 [ 0.0002 -0.0023]]
population Gamma_0, ..., Gamma_4:
[[ 1.      0.    ]
 [ 0.0769  0.75  ]
 [-0.0097 -0.0998]
 [ 0.0017  0.0125]
 [ 0.0002 -0.0023]]
max discrepancy 1.64e-05
R^2 = 0.999996
```

```
Q_lags = 8
Z_y = np.column_stack([y_sim[Q_lags - l - 1:T - l - 1, 0] for l in range(Q_lags)])
tgt = a_filtered[Q_lags:, 0]
b_y = np.linalg.lstsq(Z_y, tgt, rcond=None)[0]
resid = tgt - Z_y @ b_y
r2_y = 1 - (resid ** 2).sum() / tgt.var() / len(tgt)

print('OLS of a_t on y_{t-1}, ..., y_{t-8}:')
print(np.round(b_y, 5))
print(f'R^2 = {r2_y:.6f}')
```

```
OLS of a_t on y_{t-1}, ..., y_{t-8}:
[-0.0046  0.0044 -0.0015  0.001   0.0018 -0.0063  0.0055 -0.0024]
R^2 = 0.000056
```

The first regression recovers the Γ_j and fits almost perfectly, the small shortfall coming only from truncating the distributed lag at four lags.

So a_t really is built out of shocks stretching back before t.

The second explains essentially nothing, confirming that a_t is nonetheless orthogonal to the small econometrician's information set.

We can see why by asking how much of each past shock the small econometrician has managed to learn.

```
print('R^2 from projecting a structural shock on y_{t-1}, ..., y_{t-8}')
for lag in [0, 1, 2, 3]:
    r2s = []
    for k in range(2):
        shock = eps_sim[Q_lags - lag:T - lag, k]
        c = np.linalg.lstsq(Z_y, shock, rcond=None)[0]
        e = shock - Z_y @ c
        r2s.append(1 - (e ** 2).sum() / shock.var() / len(shock))
    print(f'   eps_(t-{lag}):   eps_r {r2s[0]:6.4f}    eps_z {r2s[1]:6.4f}')
```

```
R^2 from projecting a structural shock on y_{t-1}, ..., y_{t-8}
   eps_(t-0):   eps_r 0.0000    eps_z 0.0001
   eps_(t-1):   eps_r 0.6314    eps_z 0.0001
   eps_(t-2):   eps_r 0.6351    eps_z 0.3561
   eps_(t-3):   eps_r 0.6352    eps_z 0.3624
```

The current shock ε_t is entirely unpredictable from y^{t-1}, as it must be.

Shocks from two or more periods back are substantially learned.

The interesting row is ε_{t-1}: its r component is largely known, while its z component is *completely* unknown, because z_{t-1} reaches r only with a one-period lag and so has not yet shown up anywhere in y^{t-1}.

That is why (51.19) gives a_t a loading on $\varepsilon_{z,t-1}$ equal to the full structural coefficient $A_1^{y\tilde{y}} = 0.75$, while its loading on $\varepsilon_{r,t-1}$ is only the much smaller residual $A_1^{yy} - B_1$.

```
B1 = mod_r.var_coefficients(1)[0]
print(f'Gamma_1                   = {Gam_long[1, 0, :]}')
print(f'[A1[0,0] - B_1,  A1[0,1]] = '
      f'[{A1[0, 0] - B1[0, 0]:.4f}, {A1[0, 1]:.4f}]')

Phi_r = mod_r.structural_ma(6)
B_r6 = mod_r.var_coefficients(6)
recursion = np.array([Phi_r[k] - sum(B_r6[j - 1] @ Phi_r[k - j]
                                     for j in range(1, k + 1))
                      for k in range(7)])
print(f'\nmax |Gamma_k - (Phi_k - sum_j B_j Phi_(k-j))| = '
      f'{np.abs(recursion - Gam_long[:7]).max():.2e}')
```

```
Gamma_1                   = [0.0769 0.75  ]
[A1[0,0] - B_1,  A1[0,1]] = [0.0769, 0.7500]

max |Gamma_k - (Phi_k - sum_j B_j Phi_(k-j))| = 1.25e-16
```

51.8 Example 2: an omitted interest rate

Now a trivariate VAR(1) in output growth g_t, inflation π_t, and an interest rate i_t, from which the econometrician drops i_t.

We contrast two coefficient matrices that differ *only* in whether the interest rate feeds back onto g and π.

$$A_1^{\text{exog}} = \begin{pmatrix} 0.60 & 0.10 & 0.00 \\ 0.15 & 0.55 & 0.00 \\ 0.30 & 0.40 & 0.70 \end{pmatrix}, \qquad A_1^{\text{fb}} = \begin{pmatrix} 0.60 & 0.10 & -0.35 \\ 0.15 & 0.55 & 0.25 \\ 0.30 & 0.40 & 0.70 \end{pmatrix}.$$

The shock covariance matrix V is the same in both, and it is *not* diagonal, so the interest rate innovation is contemporaneously correlated with the other two.

```
V3 = np.array([[0.36, 0.05, 0.02],
               [0.05, 0.25, 0.06],
               [0.02, 0.06, 0.16]])
S_gpi = np.array([[1.0, 0.0, 0.0],
                  [0.0, 1.0, 0.0]])

A_exog = np.array([[0.60, 0.10, 0.00],
                   [0.15, 0.55, 0.00],
                   [0.30, 0.40, 0.70]])
A_fb = np.array([[0.60, 0.10, -0.35],
                 [0.15, 0.55,  0.25],
                 [0.30, 0.40,  0.70]])

mod_exog = VARSubsystem([A_exog], V3, S_gpi)
mod_fb = VARSubsystem([A_fb], V3, S_gpi)

_, Gam_exog, _ = report(mod_exog, label='i is block exogenous')
print()
_, Gam_fb, _ = report(mod_fb, label='i feeds back')
```

```
--- i is block exogenous (n = 3, m = 1, n_y = 2)
  max |eig| of companion A      0.700000
  max |eig| of A - KG           0.700000
  Riccati residual              5.55e-17
  |Gamma[0] - S_y|              0.00e+00
  |Phi - Psi * Gamma|           2.78e-17
  |Omega - sum Gamma V Gamma'|  0.00e+00
  max |Gamma[j]|, j >= 1        1.110e-16
  det Omega / det S_y V S_y'    1.0000

--- i feeds back (n = 3, m = 1, n_y = 2)
  max |eig| of companion A      0.792688
  max |eig| of A - KG           0.558938
  Riccati residual              1.11e-16
  |Gamma[0] - S_y|              0.00e+00
  |Phi - Psi * Gamma|           1.11e-16
  |Omega - sum Gamma V Gamma'|  1.11e-16
  max |Gamma[j]|, j >= 1        3.500e-01
  det Omega / det S_y V S_y'    1.1614
```

The block exogenous case behaves exactly as *Proposition 51.5.2* says it must, despite the correlated shocks.

The feedback case does not.

```
print('block exogenous: B_1 versus the (g, pi) block of A_1')
print(mod_exog.var_coefficients(3)[0])
print(A_exog[:2, :2])
print(f'  max |B_j| for j >= 2: {np.abs(mod_exog.var_coefficients(12)[1:]).max():.2e}
 ↪')

print('\nfeedback: B_1 versus the (g, pi) block of A_1')
print(mod_fb.var_coefficients(3)[0])
print(A_fb[:2, :2])
print(f'  max |B_j| for j >= 2: {np.abs(mod_fb.var_coefficients(12)[1:]).max():.4f}')
```

```
block exogenous: B_1 versus the (g, pi) block of A_1
[[0.6  0.1 ]
 [0.15 0.55]]
[[0.6  0.1 ]
 [0.15 0.55]]
  max |B_j| for j >= 2: 6.66e-17

feedback: B_1 versus the (g, pi) block of A_1
[[ 0.643  -0.0373]
 [ 0.1193  0.648 ]]
[[0.6  0.1 ]
 [0.15 0.55]]
  max |B_j| for j >= 2: 0.1455
```

With block exogeneity the subsystem VAR is *exactly* the corresponding block of the large VAR, and it stops at one lag.

With feedback the one-lag coefficients are distorted and higher-order terms appear.

The effect on the coefficient of lagged inflation in the output growth equation is worth noticing: a positive number in the large system becomes a negative one in the subsystem.

```
J_max = 10
labels = [r'$\varepsilon_g$', r'$\varepsilon_\pi$', r'$\varepsilon_i$']
fig, axes = plt.subplots(2, 2, figsize=(11, 6), sharex=True)
for col, (Gam, ttl) in enumerate([(Gam_exog, 'block exogenous'),
                                  (Gam_fb, 'feedback')]):
    for row, obs in enumerate([r'$a_g$', r'$a_\pi$']):
        ax = axes[row, col]
        for k in range(3):
            ax.plot(np.arange(J_max + 1), Gam[:J_max + 1, row, k],
                    marker='o', ms=3, lw=1.5, label=labels[k])
        ax.axhline(0, color='k', lw=0.6)
        ax.set_title(f'{obs}, {ttl}', fontsize=10)
        if row == 1:
            ax.set_xlabel('lag $j$')
        if row == 0 and col == 0:
            ax.legend(fontsize=8)
axes[0, 0].set_ylabel(r'$\Gamma_j$')
axes[1, 0].set_ylabel(r'$\Gamma_j$')
fig.suptitle(r'Coefficients $\Gamma_j$ in $a_t = \sum_j \Gamma_j \varepsilon_{t-j}$')
fig.tight_layout()
plt.show()
```

In the left column only the $j = 0$ coefficients are nonzero, and they equal the rows of S_y.

In the right column the omitted interest rate shock ε_i leaks into the observed innovations at every lag.

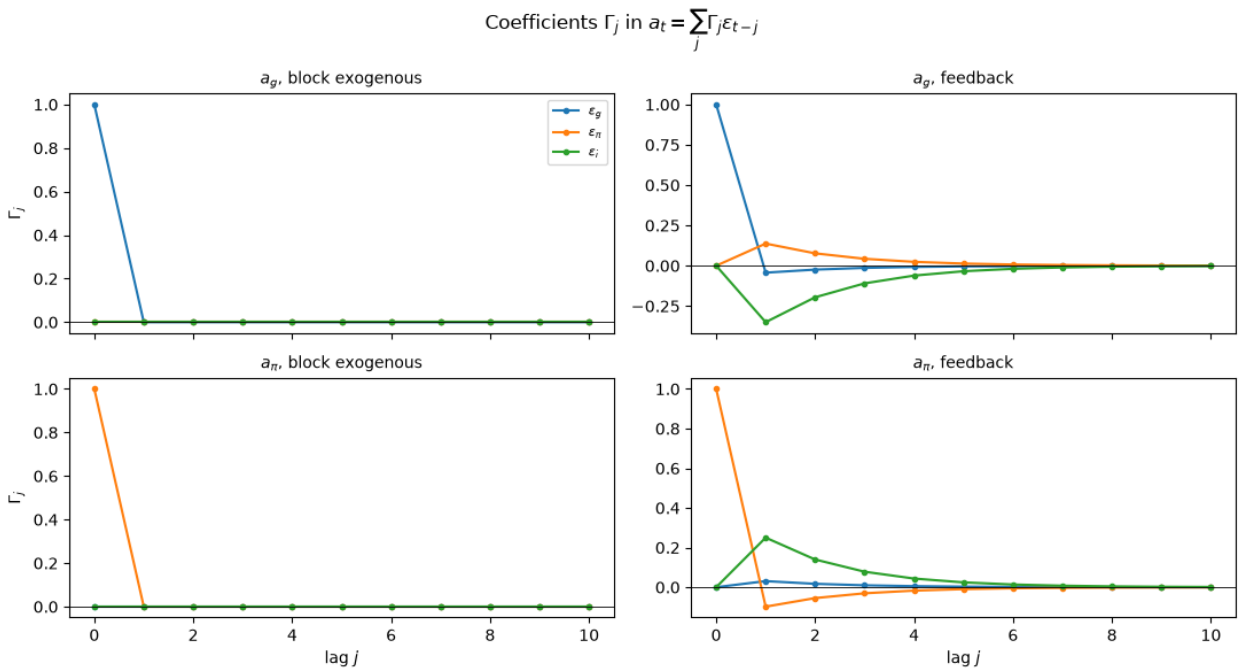


Fig. 51.4: Innovation map coefficients with and without feedback

51.9 Summary

A finite-order VAR for Y_t implies, for any subvector $y_t = S_y Y_t$, an infinite-order VAR whose coefficients $B_j = G(A - KG)^{j-1}K$ come from the steady-state Kalman filter for the companion system.

The innovations of the small system are a one-sided distributed lag $a_t = \sum_j \Gamma_j \varepsilon_{t-j}$ of the innovations of the large system, with $\Gamma_j = G(A - KG)^j C$ and $\Gamma_0 = S_y$.

Because Γ_j is wide, that map cannot be inverted, which is a precise statement of the informational deficiency of a subsystem VAR.

The price is measured by $\Omega - S_y V S_y^\top = \sum_{j \geq 1} \Gamma_j V \Gamma_j^\top$.

The price is zero when everything is observed, and also when the retained block is block exogenous, in which case the subsystem VAR is exactly the corresponding block of the original one.

51.10 Exercises

Exercise 51.10.1

Take the trivariate system of Example 2 and put the feedback of the interest rate onto (g, π) under your control by writing

$$A_1(\theta) = A_1^{\text{exog}} + \theta \begin{pmatrix} 0 & 0 & -0.35 \\ 0 & 0 & 0.25 \\ 0 & 0 & 0 \end{pmatrix}.$$

For θ on a grid from 0 to 1.5, plot

1. $\det \Omega / \det(S_y V S_y^\top)$, the information lost by dropping i_t,
2. $\max_{j \geq 1} |\Gamma_j|$, the size of the leakage of past shocks into the observed innovations.

Explain the shape you find at $\theta = 0$.

Solution

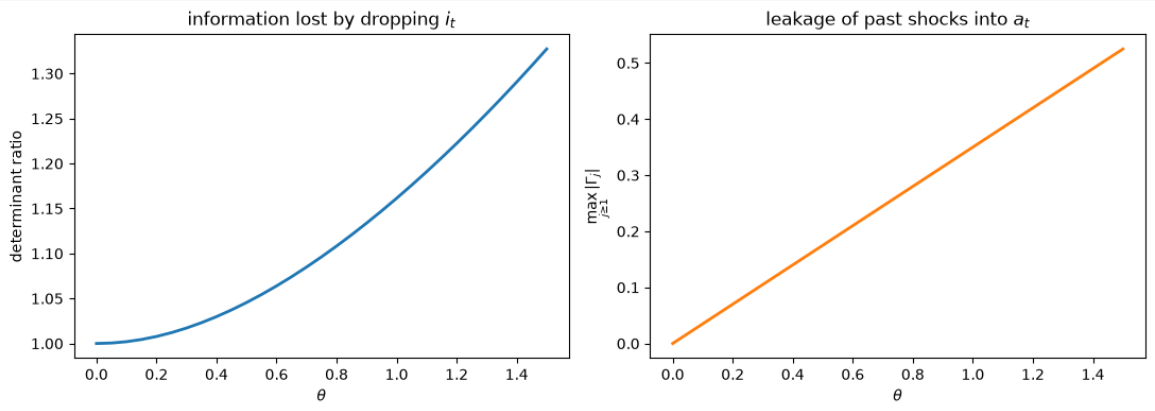
Here is one solution:

```
E = np.zeros((3, 3))
E[0, 2], E[1, 2] = -0.35, 0.25

thetas = np.linspace(0, 1.5, 31)
det_ratio, leak = [], []
for th in thetas:
    mod = VARSubsystem([A_exog + th * E], V3, S_gpi)
    Om = mod.innovation_cov()
    det_ratio.append(np.linalg.det(Om)
                     / np.linalg.det(S_gpi @ V3 @ S_gpi.T))
    leak.append(np.abs(mod.innovation_map(40)[1:]).max())

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
axes[0].plot(thetas, det_ratio, lw=2)
axes[0].set(xlabel=r'$\theta$', ylabel='determinant ratio',
            title='information lost by dropping $i_t$')
axes[1].plot(thetas, leak, lw=2, color='C1')
axes[1].set(xlabel=r'$\theta$', ylabel=r'$\max_{j \geq 1} |\Gamma_j|$',
            title='leakage of past shocks into $a_t$')
for ax in axes:
    ax.axhline(ax.get_ylim()[0], color='k', lw=0.6)
fig.tight_layout()
plt.show()

print(f'at theta = 0:   ratio = {det_ratio[0]:.6f}, leakage = {leak[0]:.2e}')
print(f'at theta = 1:   ratio = {det_ratio[20]:.4f}, leakage = {leak[20]:.4f}')
```




```
at theta = 0:   ratio = 1.000000, leakage = 1.11e-16
at theta = 1:   ratio = 1.1614, leakage = 0.3500
```

Both curves start at their minima, exactly zero leakage and a determinant ratio of exactly one.

That is *Proposition 51.5.2*: at $\theta = 0$ the interest rate does not Granger cause (g, π), so dropping it costs nothing even

though its innovation is contemporaneously correlated with the others.

Both measures rise as the feedback strengthens.

Exercise 51.10.2

Return to the bivariate system, but replace A_1 and A_2 by the single matrix

$$A_1 = \begin{pmatrix} 0.5 & 0.6 \\ 0 & \rho_z \end{pmatrix}, \qquad V = I_2,$$

so that ρ_z controls the persistence of the omitted variable z.

Observing r only, report for $\rho_z \in \{0.2, 0.5, 0.75, 0.9, 0.95, 0.99\}$

1. the forecast error variance Ω,
2. the number of lags needed before $\sum_{j>p} |B_j| < 10^{-3}$.

What does a persistent omitted variable do to the VAR that an econometrician should fit?

Solution

Here is one solution:

```
print(' rho_z    Omega    lags needed')
for rho_z in [0.2, 0.5, 0.75, 0.9, 0.95, 0.99]:
    mod = VARSubsystem([np.array([[0.5, 0.6], [0.0, rho_z]])],
                       np.eye(2), np.array([[1.0, 0.0]]))
    B = mod.var_coefficients(400)
    tails = [np.abs(B[p:]).sum() for p in range(400)]
    p_need = next(p for p in range(400) if tails[p] < 1e-3)
    print(f' {rho_z:5.2f}   {mod.innovation_cov()[0, 0]:6.4f}      {p_need:3d}')
```

```
rho_z    Omega    lags needed
 0.20   1.3708        3
 0.50   1.4359        5
 0.75   1.5625        5
 0.90   1.6910        7
 0.95   1.7454        8
 0.99   1.7937        8
```

Both columns rise with ρ_z.

A persistent omitted variable both inflates the forecast error variance and lengthens the autoregression, because the econometrician must reach further back to extract the same information about z from the history of r.

An empirical VAR with too few lags will therefore be worst exactly where the missing variable is most persistent.

Exercise 51.10.3

The coefficients B_j are population objects.

Simulate $T = 200{,}000$ observations of the bivariate system of Example 1, retain only r_t, and fit finite-order autoregressions of orders $p = 1, 2, 4, 8$ by ordinary least squares.

Compare the estimates with $B_1, \ldots, B_p$ and the residual variance with Ω.

Which order is too short, and what does fitting too short a lag length do to the first coefficient?

Solution

Here is one solution:

```
Y_big, _ = mod_r.simulate(200_000, seed=3)
r_series = Y_big[:, 0]
B_pop = mod_r.var_coefficients(8)[:, 0, 0]

for p in [1, 2, 4, 8]:
    X = np.column_stack([r_series[p - 1 - l:len(r_series) - 1 - l]
                         for l in range(p)])
    zz = r_series[p:]
    b_hat = np.linalg.lstsq(X, zz, rcond=None)[0]
    resid = zz - X @ b_hat
    print(f'p = {p}')
    print(f'   OLS        {np.round(b_hat, 4)}')
    print(f'   population {np.round(B_pop[:p], 4)}')
    print(f'   residual variance {resid.var():.4f}   Omega {mod_r.innovation_
↪cov()[0, 0]:.4f}')
```

```
p = 1
   OLS        [0.8179]
   population [0.7231]
   residual variance 1.6037   Omega 1.5786
p = 2
   OLS        [0.7201 0.1196]
   population [0.7231 0.1212]
   residual variance 1.5807   Omega 1.5786
p = 4
   OLS        [ 0.7207  0.1238 -0.0036 -0.0025]
   population [ 0.7231  0.1212 -0.0056  0.0007]
   residual variance 1.5807   Omega 1.5786
p = 8
   OLS        [ 0.7207  0.1238 -0.0035 -0.0022 -0.0019 -0.0011  0.003   0.0003]
   population [ 0.7231  0.1212 -0.0056  0.0007 -0.0001 -0.     -0.     -0.    ]
   residual variance 1.5806   Omega 1.5786
```

An AR(1) is too short.

Its single coefficient is pulled well above B_1, because it has to stand in for the omitted second lag, and its residual variance exceeds Ω.

From $p = 2$ onward the estimates track the population coefficients and the residual variance settles on Ω, which matches the finding above that B_j is negligible beyond the second lag in this example.

CHAPTER

FIFTYTWO

ORGANIZATION CAPITAL

52.1 Overview

This lecture describes a theory of **organization capital** proposed by Prescott and Visscher [1980].

Prescott and Visscher define organization capital as information that a firm accumulates about its employees, teams, and production processes.

This information is an *asset* to the firm because it affects the production possibility set and is produced jointly with output.

Costs of adjusting the stock of organization capital constrain the firm's growth rate, providing an explanation for

1. why firm growth rates are independent of firm size (Gibrat's Law)
2. why adjustment costs for rapid growth arise endogenously rather than being assumed

The paper offers three examples of organization capital:

- *Personnel information*: knowledge about the match between workers and tasks
- *Team information*: knowledge about how well groups of workers mesh
- *Firm-specific human capital*: skills of employees enhanced by on-the-job training

In each case, the investment possibilities lead firms to grow at a common rate, yielding constant returns to scale together with increasing costs of rapid size adjustment.

Note

The theory is related to ideas of Coase [1937] and Williamson [1975] about the nature of the firm.

Prescott and Visscher stress the firm's role as a storehouse of information and argue that incentives within the firm are created for efficient accumulation and use of that information.

Let's start with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.stats import norm
from scipy.optimize import brentq
```

52.2 The basic idea

The firm is a storehouse of information.

Within the firm, incentives are created for the efficient accumulation and use of that information.

Prescott and Visscher exploit this concept to explain certain facts about firm growth and size distribution.

The key insight: the process by which information is accumulated naturally leads to

1. *constant returns to scale*, and
2. *increasing costs to rapid firm size adjustment*

Constant returns to scale explain the absence of an observed unique optimum firm size (see Stigler [1958]).

Without costs of adjustment, the pattern of investment by firms in the face of a change in market demand would exhibit discontinuities we do not observe.

Further, without a cost penalty to rapid growth, the first firm to discover a previously untapped market would preempt competition by usurping all profitable investments as they appear, thus implying monopoly more prevalent than it is.

52.3 Personnel information as organization capital

The first example of organization capital is information about the match between workers and tasks.

52.3.1 Setup

Workers have different sets of skills and talents.

A variable θ measures the aptitude of a worker for a particular kind of work.

- Workers with high θ have comparative advantage in tasks requiring repeated attention to detail
- Workers with low θ have comparative advantage in work requiring broadly defined duties

The population distribution of θ is normal with mean zero and precision (inverse of variance) π:

$$\theta \sim N(0, 1/\pi)$$

When a worker is hired from the labor pool, neither the worker nor the employer knows θ.

Both know only the population distribution.

52.3.2 Three tasks

If q units of output are produced, assume:

- $\varphi_1 q$ workers are assigned to *task 1* (screening)
- φq workers are assigned to *task 2*
- the remaining workers are assigned to *task 3*

where $\varphi_1 + 2\varphi = 1$.

> **Note**
>
> The fixed coefficients technology requires a constant ratio between the number of personnel in jobs 2 and 3 and the number assigned to job 1.

For task 1, the screening task, per unit cost of production is *invariant* to the θ-values of the individuals assigned.

However, the larger a worker's θ, the larger is his product in task 2 relative to his product in task 3.

Consequently:

- a worker with a highly positive θ is much better suited for task 2
- a worker with a highly negative θ is much better suited for task 3

52.3.3 Bayesian learning

Performance in tasks 2 or 3 cannot be observed at the individual level.

But information about a worker's θ-value can be obtained from observing performance in task 1, the screening task.

The expert supervising the apprentice determines a value of z each period:

$$z_{it} = \theta_i + \epsilon_{it} \tag{52.1}$$

where $\epsilon_{it} \sim N(0, 1)$ are independently distributed over both workers i and periods t.

After n observations on a worker in the screening job, the *posterior distribution* of θ is normal with

posterior mean:

$$m = \frac{1}{\pi + n} \sum_{k=1}^{n} z_k \tag{52.2}$$

posterior precision:

$$h = \pi + n \tag{52.3}$$

Knowledge of an individual is thus completely characterized by the pair (m, h).

```
def bayesian_update(z_observations, prior_precision):
    """
    Compute posterior mean and precision after observing signals.
    """
    n = len(z_observations)
    h = prior_precision + n
    m = np.sum(z_observations) / h
    return m, h
```

Let's visualize how the posterior evolves as we observe a worker whose true $\theta = 0.8$:

```
rng = np.random.default_rng(0)

θ_true = 0.8
π = 1.0

T = 20
```

(continues on next page)

(continued from previous page)

```
ε = rng.standard_normal(T)
z_signals = θ_true + ε

posterior_means = []
posterior_stds = []

for n in range(1, T + 1):
    m, h = bayesian_update(z_signals[:n], π)
    posterior_means.append(m)
    posterior_stds.append(1 / np.sqrt(h))

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

ax = axes[0]
ax.plot(range(1, T + 1), posterior_means, '-o', markersize=4, lw=2,
        label='Posterior mean $m$')
ax.axhline(θ_true, color='r', linestyle='--',
           label=fr'True $\theta = {θ_true}$')
ax.set_xlabel('number of observations $n$')
ax.set_ylabel('posterior mean $m$')
ax.legend()

ax = axes[1]
ax.plot(range(1, T + 1), posterior_stds, '-o', markersize=4, lw=2,
        label=r'Posterior std $1/\sqrt{h}$')
ax.set_xlabel('number of observations $n$')
ax.set_ylabel('posterior standard deviation')
ax.legend()

plt.tight_layout()
plt.show()
```

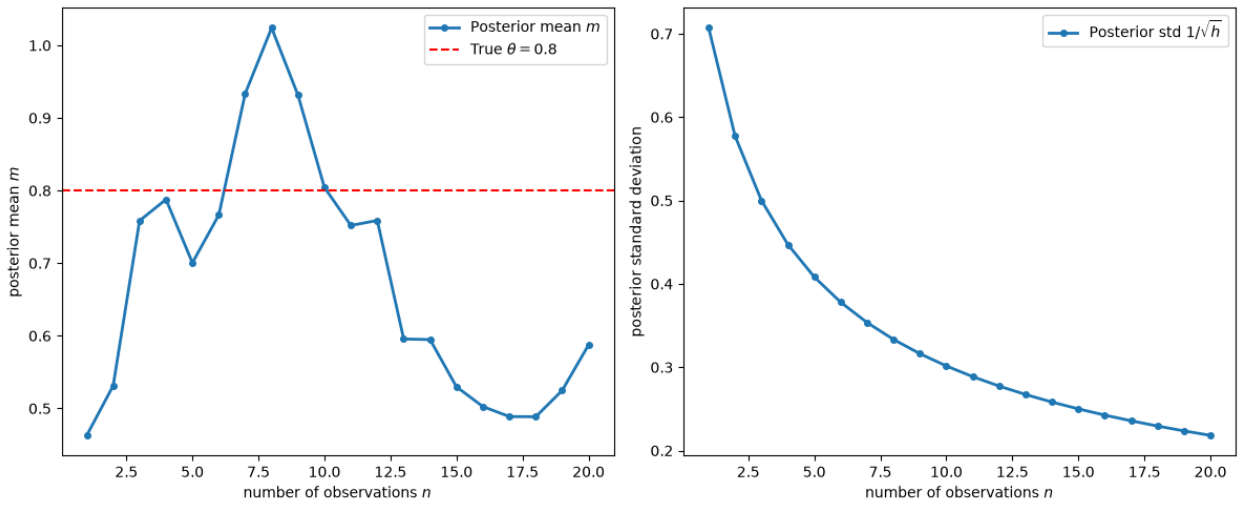


Fig. 52.1: Posterior mean convergence and uncertainty

As the number of screening observations n increases, the posterior mean converges to the true θ, and the posterior uncertainty shrinks at rate $1/\sqrt{n}$.

52.3.4 Per unit costs of production

Under the nonsequential assignment rule, employees with the greatest seniority are assigned to jobs 2 and 3, while newer employees remain in the screening task.

Workers with $m > 0$ are assigned to task 2, and those with $m \leq 0$ to task 3.

Per unit costs of production, assuming this assignment after n screening periods, are:

$$c(n) = c_1 + c_2 + c_3 - E\{\theta \mid m > 0\} + E\{\theta \mid m \leq 0\} \tag{52.4}$$

Because m is normally distributed, evaluation of the conditional expectation in (52.4) yields per unit costs as a function of n:

$$c(n) = c - 0.7978 \frac{n}{\pi(\pi + n)} \tag{52.5}$$

where $c = c_1 + c_2 + c_3$ and $0.7978 = 2 \int_0^\infty \frac{t}{\sqrt{2\pi}} e^{-t^2/2} dt$.

Note

The constant $0.7978 \approx \sqrt{2/\pi}$ is the mean of the standard half-normal distribution.

It arises from computing $E[\theta \mid m > 0] - E[\theta \mid m \leq 0]$ for a normal distribution.

The function $c(n)$ decreases at a *decreasing rate* in n.

More screening observations reduce costs but with diminishing returns.

```
def cost_per_unit(n_vals, π, c_bar=1.0):
    """
    Per unit cost as a function of screening periods n.
    """
    n_vals = np.asarray(n_vals, dtype=float)
    return c_bar - 0.7978 * n_vals / (π * (π + n_vals))


fig, ax = plt.subplots(figsize=(10, 6))

n_vals = np.linspace(0.1, 50, 200)

for π in [0.5, 1.0, 2.0, 5.0]:
    costs = cost_per_unit(n_vals, π)
    ax.plot(n_vals, costs, lw=2, label=fr'$\pi = {π}$')

ax.set_xlabel('screening periods $n$')
ax.set_ylabel('per unit cost $c(n)$')
ax.legend()
ax.set_xlim(0, 50)
plt.tight_layout()
plt.show()
```

The figure shows that:

- costs decrease with more screening time n
- the decrease is at a declining rate (diminishing returns to screening)

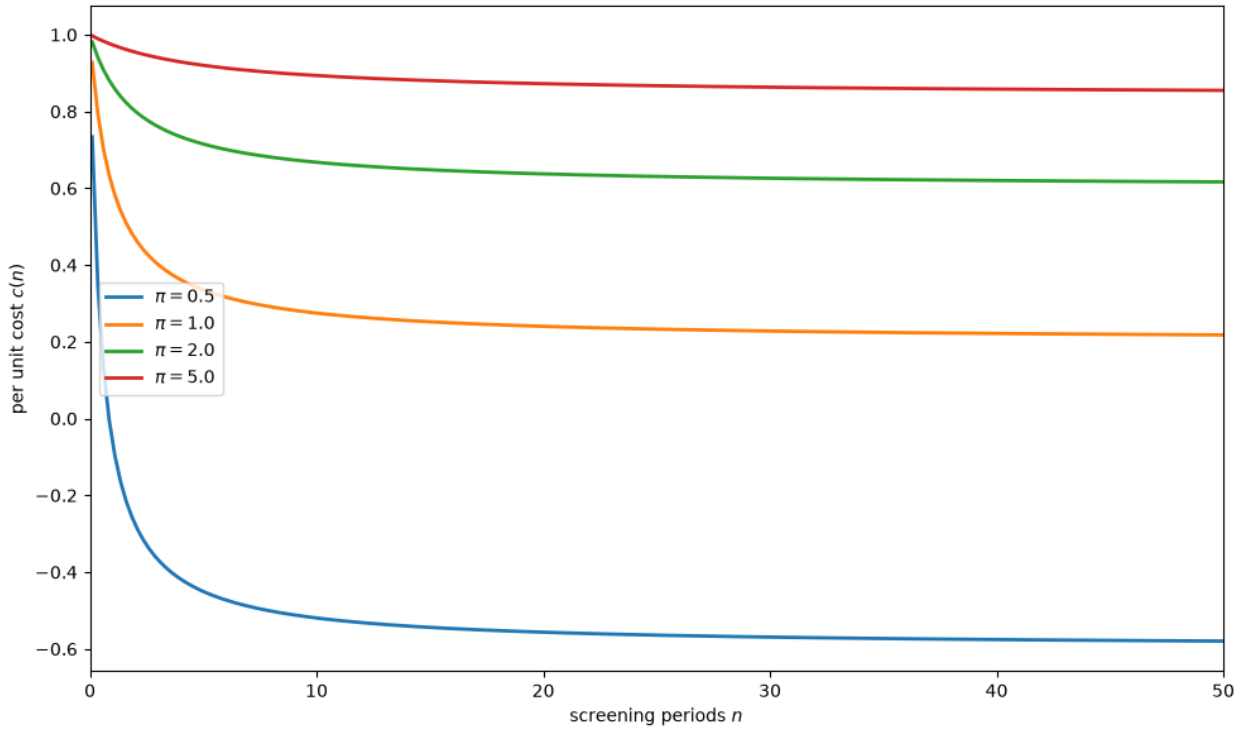


Fig. 52.2: Per unit costs by screening time

- for smaller prior precision π (more initial uncertainty about worker types), the gains from screening are larger

This diminishing-returns structure is the source of the *increasing costs of rapid adjustment*.

52.3.5 Growth rate and screening time

The greater the growth rate, the smaller must be n — the time spent in the screening task before assignment to job 2 or 3.

If γ is the growth rate of output and ρ is the quit rate, and y_i is the current number of vintage i employees, then

$$(1+\gamma)y_{i+1} = (1-\rho)y_i$$

Letting $\xi = (1-\rho)/(1+\gamma)$, from the above $y_i = \xi^i y_0$.

For the fixed coefficients technology, the fraction of present personnel with vintage greater than n must equal $2\varphi/(\varphi_1 + 2\varphi)$, which gives:

$$\xi^{n+1} = \frac{2\varphi}{\varphi_1 + 2\varphi} \tag{52.6}$$

Solving for n as a function of γ:

$$n(\gamma) = \frac{\log(2\varphi) - \log(\varphi_1 + 2\varphi)}{\log(1-\rho) - \log(1+\gamma)} - 1 \quad \text{for } \gamma > -\rho \tag{52.7}$$

```
def screening_time(γ, ρ, φ1, φ):
    """
```

(continues on next page)

(continued from previous page)

```
    Screening time n as a function of growth rate γ.
    """
    γ = np.asarray(γ, dtype=float)
    numerator = np.log(2 * ψ) - np.log(ψ1 + 2 * ψ)
    denominator = np.log(1 - ρ) - np.log(1 + γ)
    return numerator / denominator - 1


ρ = 0.1
ψ1 = 0.5
ψ = 0.25

γ_vals = np.linspace(-0.05, 0.30, 200)

valid = γ_vals > -ρ
γ_valid = γ_vals[valid]
n_vals = screening_time(γ_valid, ρ, ψ1, ψ)
mask = n_vals > 0
γ_plot = γ_valid[mask]
n_plot = n_vals[mask]

fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(γ_plot, n_plot, lw=2)
ax.set_xlabel(r'growth rate $\gamma$')
ax.set_ylabel(r'screening periods $n(\gamma)$')
ax.set_xlim(γ_plot[0], γ_plot[-1])
plt.tight_layout()
plt.show()
```

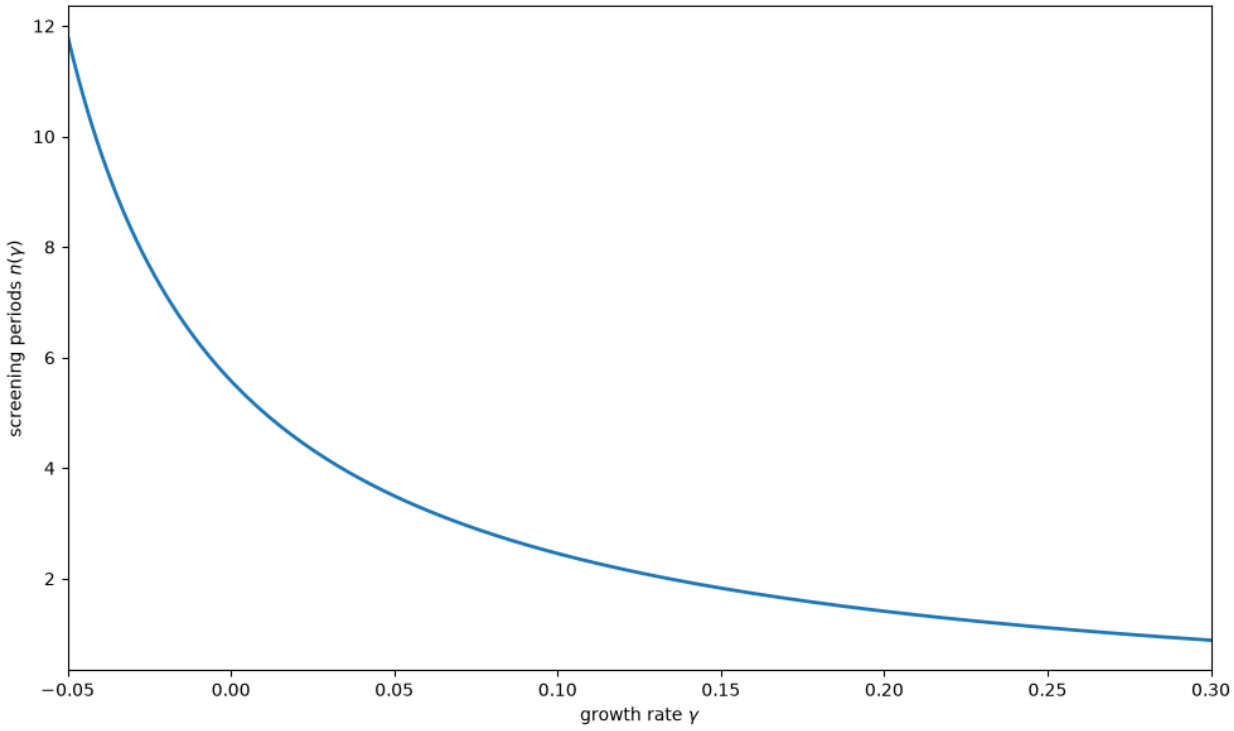


Fig. 52.3: Screening time vs. growth rate

The figure shows the key trade-off: *faster growth forces shorter screening periods*.

When growth is rapid, new workers must be promoted from the screening task to productive tasks more quickly, so less information is gathered about each worker before assignment.

52.3.6 Combined effect: growth rate and per unit costs

Composing the functions $c(n)$ and $n(\gamma)$ reveals how per unit costs depend on the growth rate:

```
fig, ax = plt.subplots(figsize=(10, 6))

π = 1.0
c_bar = 1.0

n_of_γ = screening_time(γ_plot, ρ, ψ1, ψ)
costs_of_γ = cost_per_unit(n_of_γ, π, c_bar)

ax.plot(γ_plot, costs_of_γ, lw=2)
ax.set_xlabel(r'growth rate $\gamma$')
ax.set_ylabel(r'per unit cost $c(n(\gamma))$')
ax.set_xlim(γ_plot[0], γ_plot[-1])
plt.tight_layout()
plt.show()
```

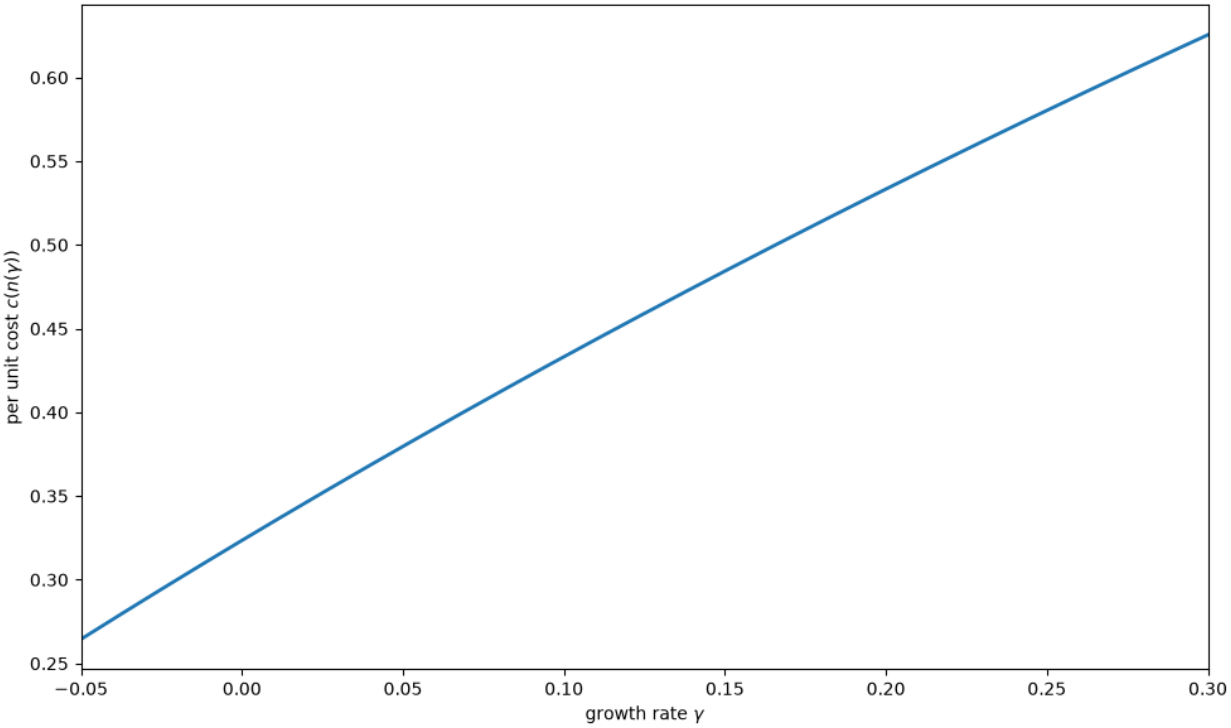


Fig. 52.4: Per unit costs vs. growth rate

This establishes the key result: *increasing costs of rapid adjustment arise endogenously* from the trade-off between screening and growth.

The faster the firm grows, the less time it has to screen workers, the poorer the match between workers and tasks, and the higher the per unit production costs.

52.4 Industry equilibrium

Firm growth rates are independent of firm size in this model because the mathematical structure of the technology constraint is the same as that considered in Lucas Jr [1967], except that the stock of organization capital is a vector rather than a scalar.

The technology set facing price-taking firms is a **convex cone**: there are constant returns to scale.

Constant returns and internal adjustment costs, along with some costs of transferring capital between firms, yield an optimum rate of firm growth *independent of the firm's size* — this is Gibrat's Law.

The bounded, downward-sloping, inverse industry demand function is

$$P_t = p(Q_t, u_t)$$

where Q_t is the sum of output over all firms and u_t is a demand shock subject to a stationary Markov process.

Prescott and Visscher show that a competitive equilibrium exists using the framework of Lucas and Prescott [1971].

The discounted consumer surplus to be maximized is

$$\sum_{t=0}^{\infty} \beta^t \left\{ \int_0^{Q_t} p(y, u_t) dy - Bw - Q_t \sum_i (A_{i2t} + A_{i3t}) c(i) / \sum_i (A_{i2t} + A_{i3t}) \right\} \tag{52.8}$$

where A_{i2t}, A_{i3t}, and B are obtained by summing a_{i2t}, a_{i3t}, and b, respectively, over all firms in the industry.

52.4.1 Key property: growth rates independent of size

If two firms have organization capital vectors $\underline{k}$ that are proportional at a point in time, they will be proportional in all future periods.

That is, *growth rates are independent of firm size*.

```
def simulate_firm_growth(T, γ, ρ, q0, seed=42):
    """
    Simulate firm output growth with stochastic shocks.
    """
    rng = np.random.default_rng(seed)
    output = np.zeros(T)
    output[0] = q0
    for t in range(1, T):
        shock = rng.normal(0, 0.02)
        output[t] = output[t-1] * (1 + γ + shock)
    return output


T = 50
γ_eq = 0.05
ρ = 0.1
```

(continues on next page)

(continued from previous page)

```
fig, axes = plt.subplots(1, 2, figsize=(14, 5))

ax = axes[0]
for q0, label in [(10, 'Small firm'), (50, 'Medium firm'),
                  (200, 'Large firm')]:
    output = simulate_firm_growth(T, y_eq, ρ, q0,
                                  seed=int(q0))
    ax.plot(range(T), output, lw=2, label=f'{label} ($q_0={q0}$)')
ax.set_xlabel('period')
ax.set_ylabel('output $q_t$')
ax.legend()

ax = axes[1]
for q0, label in [(10, 'Small firm'), (50, 'Medium firm'),
                  (200, 'Large firm')]:
    output = simulate_firm_growth(T, y_eq, ρ, q0,
                                  seed=int(q0))
    ax.plot(range(T), np.log(output), lw=2,
            label=f'{label} ($q_0={q0}$)')
ax.set_xlabel('period')
ax.set_ylabel(r'$\log(q_t)$')
ax.legend()

plt.tight_layout()
plt.show()
```

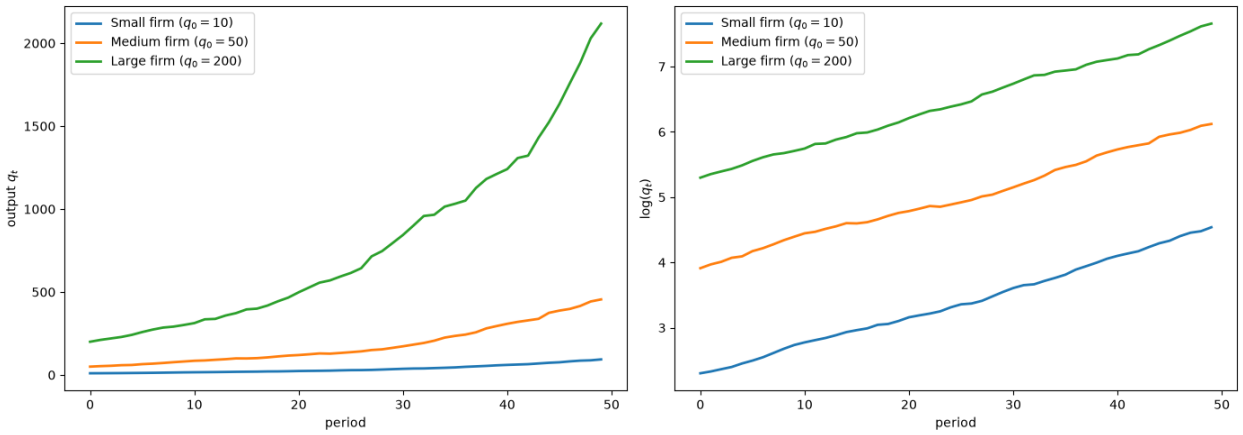


Fig. 52.5: Firm output levels and growth rates

The right panel shows that all firms grow at the same rate regardless of initial size — the log output paths are parallel.

This is **Gibrat's Law**: growth rates are independent of firm size.

52.5 Bayesian screening simulation

Let's simulate the full screening and assignment process for a single firm.

We draw workers from the population, observe their signals in the screening task, and then assign them to the appropriate productive task based on the posterior mean.

```
def simulate_screening(n_workers, n_screen, π, seed=123):
    """
    Simulate screening and assignment of workers.
    """
    rng = np.random.default_rng(seed)

    θ = rng.normal(0, 1/np.sqrt(π), n_workers)
    signals = (θ[:, None]
               + rng.normal(0, 1, (n_workers, n_screen)))
    posterior_means = signals.sum(axis=1) / (π + n_screen)

    assignment = np.where(posterior_means > 0, 2, 3)
    correct_assignment = np.where(θ > 0, 2, 3)
    misassignment_rate = np.mean(assignment != correct_assignment)

    return {
        'theta': θ,
        'posterior_means': posterior_means,
        'assignment': assignment,
        'correct_assignment': correct_assignment,
        'misassignment_rate': misassignment_rate
    }


π = 1.0
n_workers = 5000
screening_periods = [1, 3, 5, 10, 20, 50]

fig, axes = plt.subplots(2, 3, figsize=(15, 10))
axes = axes.flatten()

misassignment_rates = []

for idx, n_screen in enumerate(screening_periods):
    results = simulate_screening(n_workers, n_screen, π)
    misassignment_rates.append(results['misassignment_rate'])

    ax = axes[idx]
    θ = results['theta']
    m = results['posterior_means']

    correct = results['assignment'] == results['correct_assignment']
    ax.scatter(θ[correct], m[correct], alpha=0.1, s=5,
               color='blue', label='Correct')
    ax.scatter(θ[~correct], m[~correct], alpha=0.3, s=5,
               color='red', label='Misassigned')
    ax.axhline(0, color='k', linewidth=0.5)
    ax.axvline(0, color='k', linewidth=0.5)
    mis = results['misassignment_rate']
    ax.set_xlabel(r'true $\theta$')
```

(continues on next page)

(continued from previous page)

```
    ax.set_ylabel('posterior mean $m$')
    if idx == 0:
        ax.legend(markerscale=5, loc='upper left')

plt.tight_layout()
plt.show()
```

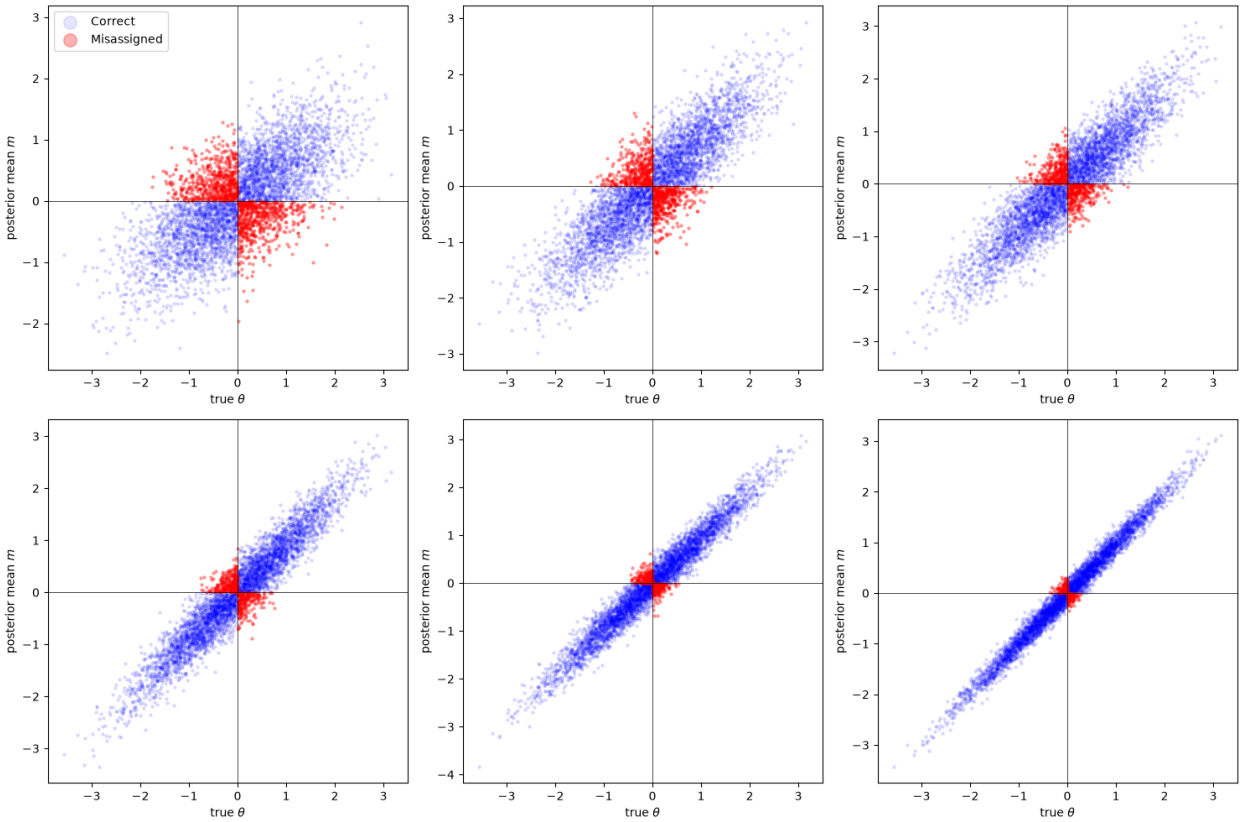


Fig. 52.6: Screening and worker assignment accuracy

Red dots are workers who are *misassigned* — placed in the wrong productive task because the posterior mean had the wrong sign relative to their true θ.

As n increases:

- The posterior mean m becomes more strongly correlated with θ
- Misassignment rates fall

```
fig, ax = plt.subplots(figsize=(10, 6))

n_range = np.arange(1, 51)
mis_rates = []
for n_screen in n_range:
    results = simulate_screening(n_workers, n_screen, π)
    mis_rates.append(results['misassignment_rate'])
```

(continues on next page)

(continued from previous page)

```
ax.plot(n_range, mis_rates, '-o', markersize=3, lw=2)
ax.set_xlabel('screening periods $n$')
ax.set_ylabel('misassignment rate')
plt.tight_layout()
plt.show()
```

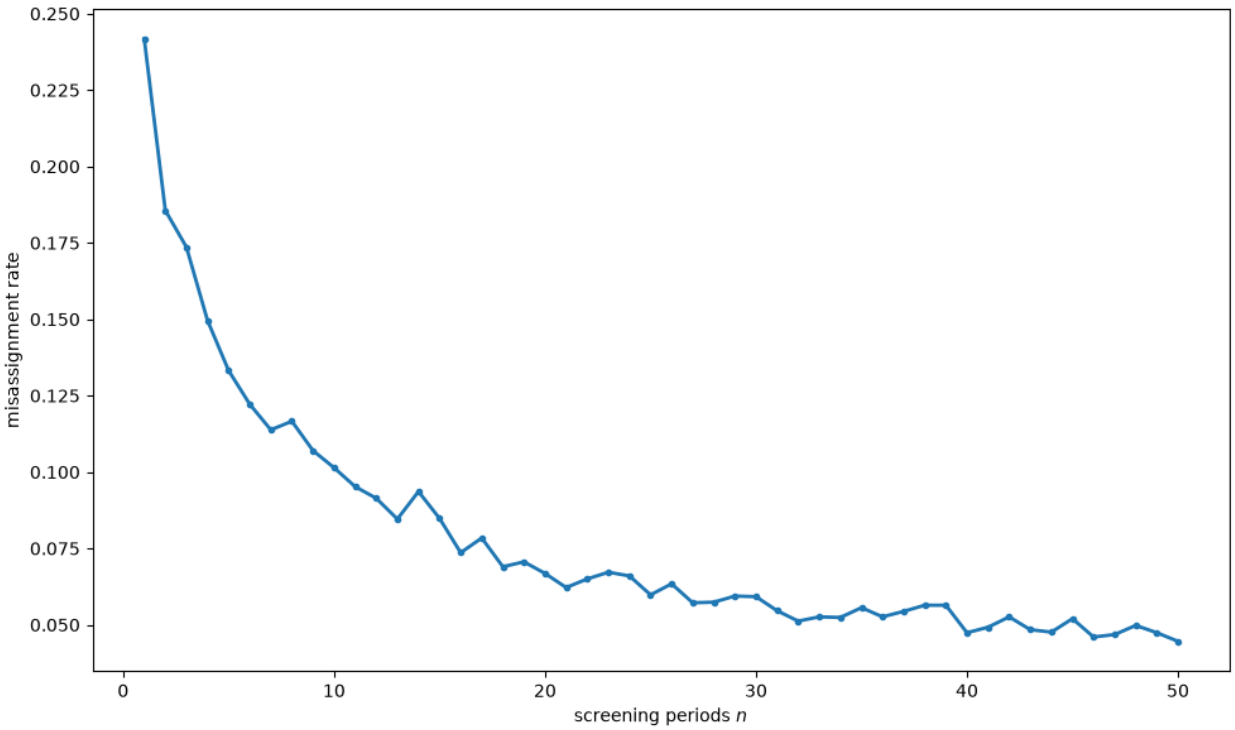


Fig. 52.7: Misassignment rate by screening time

This confirms the theoretical prediction: the cost savings from better assignment exhibit *diminishing returns* in the screening time n.

52.6 Team information

Personnel information need not be valuable only because it facilitates the matching of workers to tasks.

Another equally valuable use of personnel information is in the *matching of workers to workers*.

What is important to performance in many activities within the firm is not just the aptitude of an individual assigned to a task, but also how well the characteristics of the individual mesh with those of others performing related duties.

52.6.1 Structure

Suppose workers are grouped into teams, and team i assigned to a screening task has an observed productivity indicator

$$z_{it} = \theta_i + \epsilon_{it}$$

where:

- θ_i is a deterministic component directly related to how well team workers are paired
- $\epsilon_{it} \sim N(0,1)$ are i.i.d. stochastic components

The θ from all possible teams are approximately independently and normally distributed $N(\mu, 1/\pi)$.

After n observations on team i, the posterior distribution on θ_i is normal with

$$m = \mu + \frac{1}{\pi + n}\sum_{k=1}^{n}(z_k - \mu)$$

and precision $h = \pi + n$.

If dissolution of a team also dissolves the accrued information, the team information model has the *same mathematical structure* as the personnel information model.

```
def simulate_team_screening(n_teams, n_screen, π, μ=0.5,
                            seed=456):
    """
    Simulate team screening with Bayesian updating.
    """
    rng = np.random.default_rng(seed)

    θ = rng.normal(μ, 1/np.sqrt(π), n_teams)
    signals = (θ[:, None]
               + rng.normal(0, 1, (n_teams, n_screen)))
    z_bar = signals.mean(axis=1)
    post_means = μ + n_screen * (z_bar - μ) / (π + n_screen)
    post_prec = π + n_screen

    return {
        'theta': θ,
        'posterior_means': post_means,
        'posterior_precision': post_prec
    }


fig, axes = plt.subplots(1, 3, figsize=(15, 5))

for idx, n_screen in enumerate([1, 5, 20]):
    results = simulate_team_screening(500, n_screen, π=1.0, μ=0.5)

    ax = axes[idx]
    ax.scatter(results['theta'], results['posterior_means'],
               alpha=0.4, s=10)
    lims = [-1.5, 2.5]
    ax.plot(lims, lims, 'r--', alpha=0.5, lw=2, label='45° line')
    ax.set_xlabel(r'true team quality $\theta$')
    ax.set_ylabel('posterior mean $m$')
    ax.set_xlim(lims)
    ax.set_ylim(lims)
```

(continues on next page)

(continued from previous page)

```
    ax.legend()
    ax.set_aspect('equal')

plt.tight_layout()
plt.show()
```

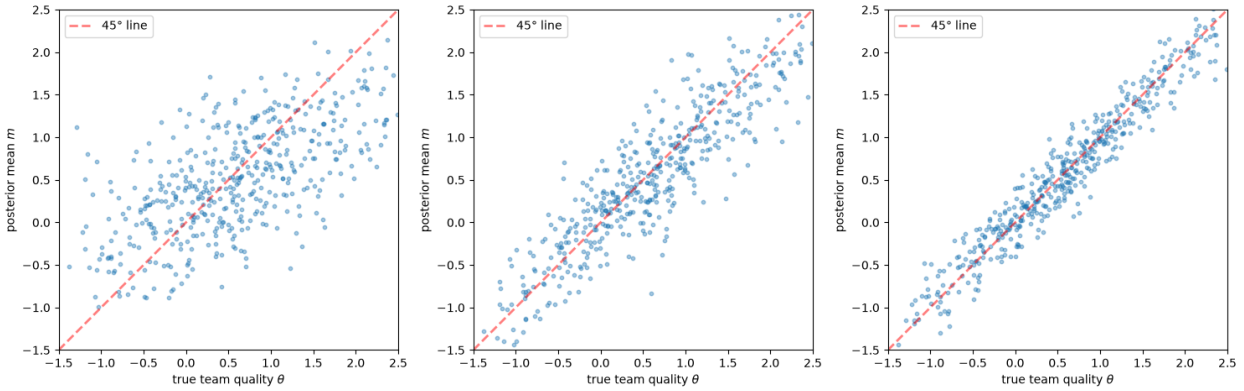


Fig. 52.8: Team quality estimates by screening periods

As with individual screening, more observations improve the precision of team quality estimates.

Rapid growth forces fewer observations before team assignments must be finalized, leading to higher costs.

52.7 Firm-specific human capital

The third example: organization capital consists of the **human capital** of the firm's employees.

The capacity of the organization to function effectively as a production unit is determined largely by the level and meshing of the skills of the employees.

Note

The case for the human capital of employees being part of the capital stock of the firm is well established (see Becker [1975]).

Productivity in the future depends on levels of human capital in the future, but to acquire human capital for the future, a sacrifice in real resources is required in the present.

The key features are:

- Output and skill enhancement are **joint products** resulting from the combination of labor inputs possessing different skill levels
- Experienced and inexperienced workers are combined in one of several available technical processes to generate the firm's product, and in the process, the overall competence of the work force is improved
- The transformation frontier between current output and future human capital is *concave* and linearly homogeneous

This gives the technology set the structure of a closed convex cone with a vertex at the origin — sufficient for optimal proportional growth by firms.

52.7.1 Concave transformation frontier

```
def transformation_frontier(q, α=0.7):
    """
    Concave transformation frontier between output and human capital.
    """
    q = np.asarray(q, dtype=float)
    return (1 - q**α)**(1/α)


fig, ax = plt.subplots(figsize=(8, 8))

q_vals = np.linspace(0, 1, 200)

for α in [0.5, 0.7, 1.0, 1.5]:
    hk = transformation_frontier(q_vals, α)
    ax.plot(q_vals, hk,
            label=fr'$\alpha = {α}$', lw=2)

ax.set_xlabel('current output $q$ (fraction of capacity)')
ax.set_ylabel('future human capital increment $\\Delta h$')
ax.legend()
ax.set_xlim(0, 1.05)
ax.set_ylim(0, 1.05)
ax.set_aspect('equal')
plt.tight_layout()
plt.show()
```

The concavity of the transformation frontier means that moving from an extremely unbalanced bundle of production and learning activity to a more balanced bundle entails little sacrifice.

But a workday consisting primarily of learning also has diminishing returns, creating the cost of rapid adjustment.

52.8 Costs of transferring organization capital

If there were no cost to transferring organization capital from one firm to another, the model would not place constraints on the firm's growth rate.

Firms could then merge, divest, or pirate each other's personnel without a cost penalty and thus produce a pattern of growth not restricted by the model.

Organization capital is *not* costlessly moved, however:

1. *Moving is disruptive*: relocating from one locale to another is disruptive to both employee and family
2. *Information is firm-specific*: the information set that makes a person productive in one organization may not make that person as productive in another, even if both firms produce identical output
    - Facility with a computer system at one firm
    - Knowing whom to ask when problems arise
    - Rapport with buyers or sellers

These are types of organization capital in one firm that *cannot be transferred costlessly* to another.

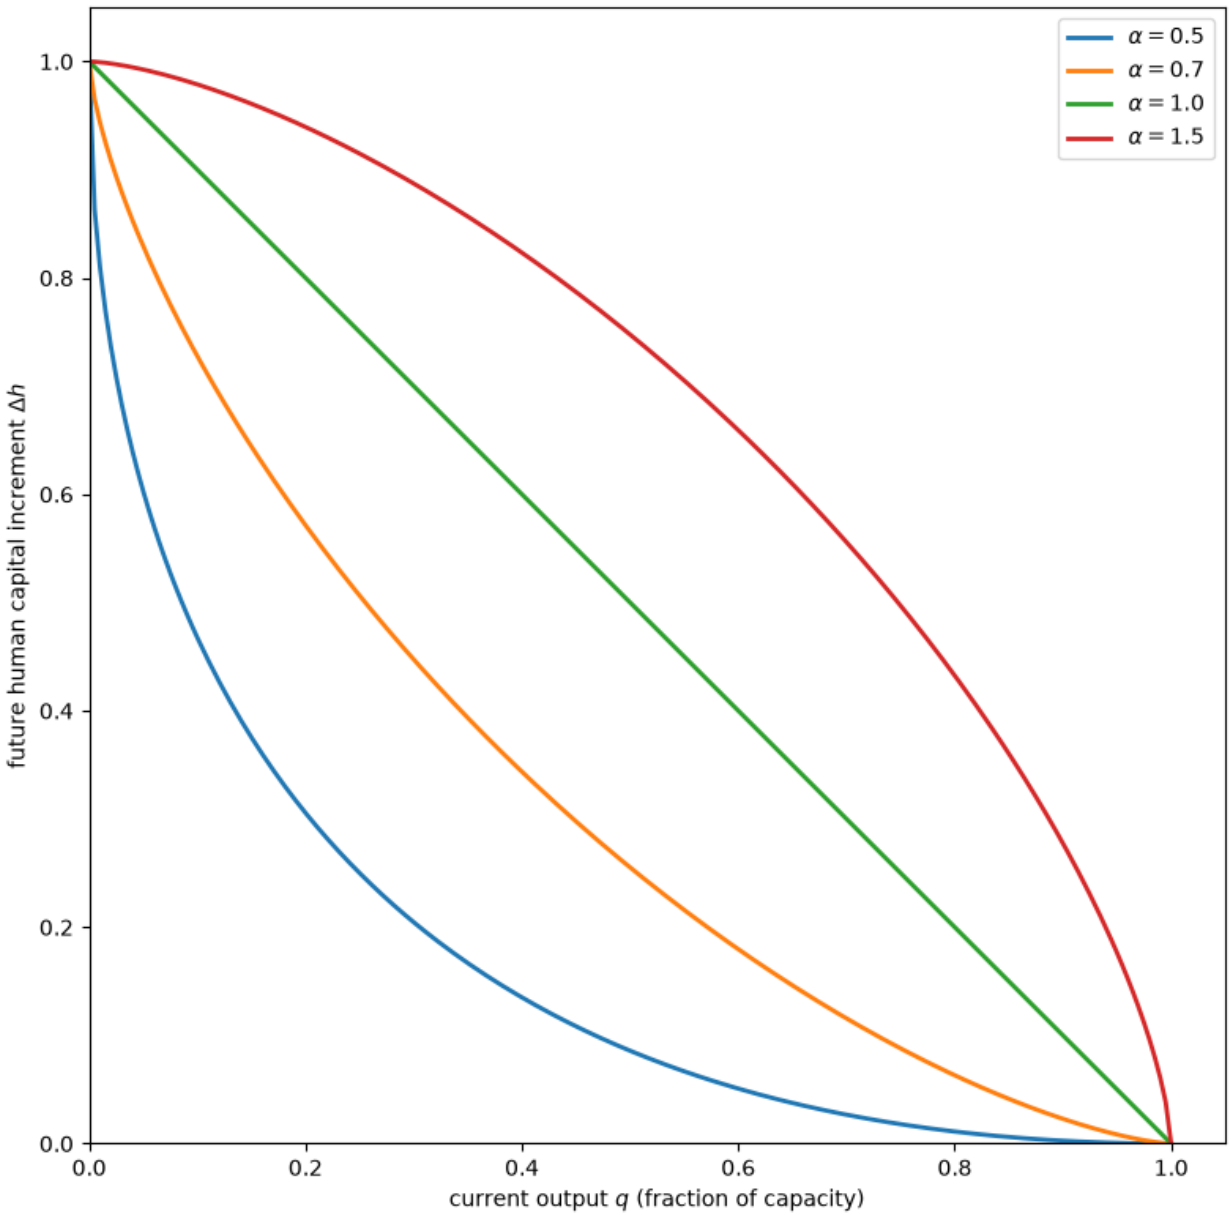


Fig. 52.9: Concave transformation frontier

52.9 Summary and implications

The Prescott-Visscher model provides a unified framework in which:

- The firm exists as an entity because it is an efficient structure for accumulating, storing, and using information
- *Constant returns to scale* arise because once the best combinations of worker types are discovered, nothing prevents the firm from replicating those combinations with proportional gains in product
- *Increasing adjustment costs* arise endogenously from the trade-off between current production and investment in organization capital
- *Gibrat's Law* — growth rates independent of firm size — is a natural implication
- Large firms should have growth rates that display *less variance* than small firms because large firms are essentially portfolios of smaller production units

```
def simulate_growth_rate_distribution(n_firms, n_subunits, γ,
                                      σ, T=100, seed=789):
    """
    Simulate growth rate distributions for firms of different sizes.
    """
    rng = np.random.default_rng(seed)
    subunit_growth = rng.normal(γ, σ,
                                (n_firms, n_subunits, T))
    firm_growth = subunit_growth.mean(axis=1)
    return firm_growth.mean(axis=1)


fig, ax = plt.subplots(figsize=(10, 6))

sizes = {'Small (1 unit)': 1,
         'Medium (5 units)': 5,
         'Large (20 units)': 20}

γ = 0.05
σ = 0.10

for label, n_sub in sizes.items():
    rates = simulate_growth_rate_distribution(
        2000, n_sub, γ, σ)
    ax.hist(rates, bins=50, alpha=0.5, density=True,
            label=f'{label}: std={rates.std():.4f}')

ax.set_xlabel('average growth rate')
ax.set_ylabel('density')
ax.legend()
ax.axvline(γ, color='k', linestyle='--',
           label=r'$\gamma$', alpha=0.5)
plt.tight_layout()
plt.show()
```

The figure shows that although all firms have the *same mean growth rate* (Gibrat's Law), large firms display *less variance* in realized growth rates because they are effectively portfolios of independent subunits.

This is consistent with the empirical findings of Mansfield [1962] and Hymer and Pashigian [1962].

The essence of the Prescott-Visscher theory is that the nature of the firm is tied to *organization capital*.

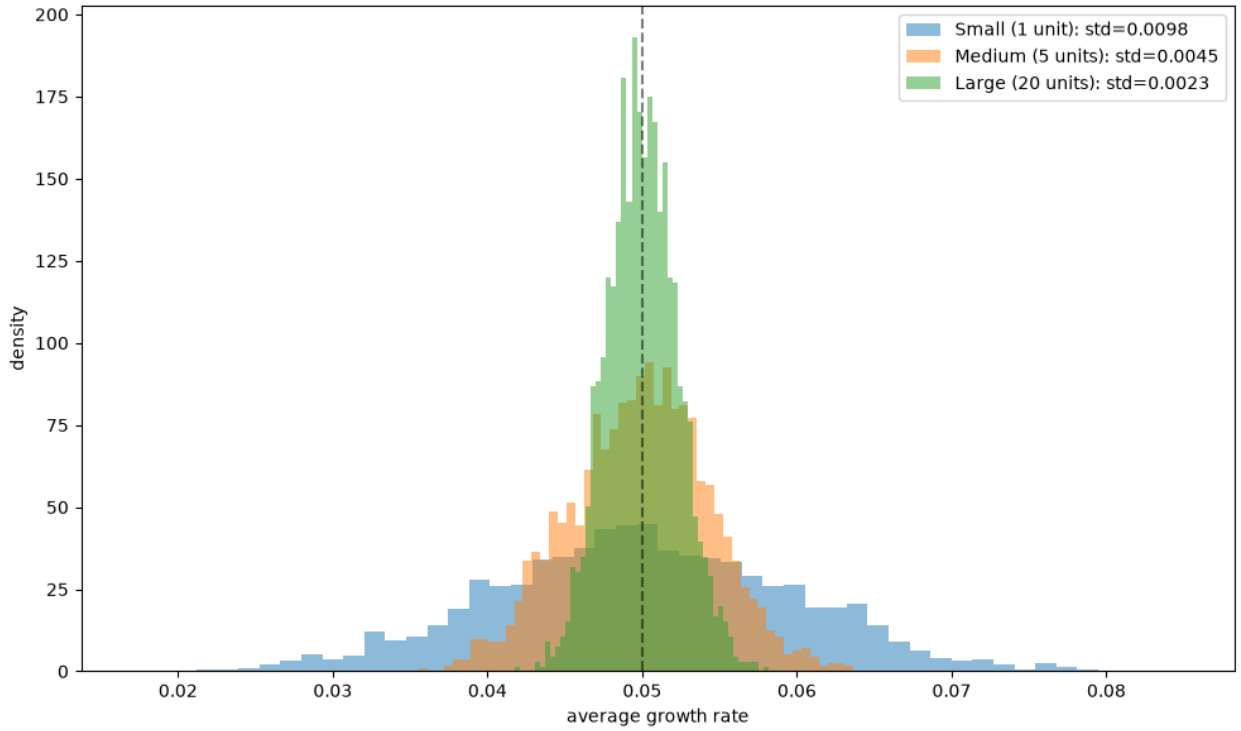


Fig. 52.10: Growth rate distributions by firm size

What distinguishes the firm from other relationships is that it is a structure within which agents have the incentive to acquire and reveal information in a manner that is less costly than in possible alternative institutions.

CHAPTER

FIFTYTHREE

TWO MODELS OF MEASUREMENTS AND THE INVESTMENT ACCELERATOR

Contents

- *Two Models of Measurements and the Investment Accelerator*
  - *Overview*
  - *The economic model*
  - *Measurement errors*
  - *A classical model of measurements initially collected by an agency*
  - *A model of optimal estimates reported by an agency*
  - *Simulation*
  - *Summary*
  - *Exercises*

53.1 Overview

"Rational expectations econometrics" aims to interpret economic time series in terms of objects that are meaningful to economists, namely, parameters describing preferences, technologies, information sets, endowments, and equilibrium concepts.

When fully worked out, rational expectations models typically deliver a well-defined mapping from these economically interpretable parameters to the moments of the time series determined by the model.

If accurate observations on these time series are available, one can use that mapping to implement parameter estimation methods based either on the likelihood function or on the method of moments.

Note

This is why econometric estimation is often called an "inverse" problem, while simulating a model for given parameter values is called a "direct" problem. The direct problem refers to the mapping we have just described, while the inverse problem involves somehow applying an "inverse" of that mapping to a data set that is treated as if it were one draw from the joint probability distribution described by the mapping.

However, if only error-ridden data exist for the variables of interest, then more steps are needed to extract parameter estimates.

In effect, we require a model of the data reporting agency, one that is workable enough that we can determine the mapping induced jointly by the dynamic economic model and the measurement process to the probability law for the measured data.

The model chosen for the data collection agency is an aspect of an econometric specification that can make big differences in inferences about the economic structure.

Sargent [1989] describes two alternative models of data generation in a *permanent income* economy in which the investment accelerator, the mechanism studied in these two quantecon lectures – *Samuelson Multiplier-Accelerator* and *The Acceleration Principle and the Nature of Business Cycles* – shapes business cycle fluctuations.

- In Model 1, the data collecting agency simply reports the error-ridden data that it collects.
- In Model 2, the data collection agency first collects error-ridden data that satisfy a classical errors-in-variables model, then filters the data, and reports the filtered objects.

Although the two models have the same "deep parameters," they produce quite different sets of restrictions on the data.

In this lecture we follow Sargent [1989] and study how these alternative measurement schemes affect empirical implications.

We start with imports and helper functions to be used throughout this lecture to generate LaTeX output

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from scipy import linalg
from IPython.display import Latex

np.set_printoptions(precision=3, suppress=True)

def df_to_latex_matrix(df, label=''):
    """Convert DataFrame to LaTeX matrix."""
    lines = [r'\begin{bmatrix}']

    for idx, row in df.iterrows():
        row_str = ' & '.join(
          [f'{v:.4f}' if isinstance(v, (int, float))
            else str(v) for v in row]) + r' \\'
        lines.append(row_str)

    lines.append(r'\end{bmatrix}')

    if label:
        return '$' + label + ' = ' + '\n'.join(lines) + '$'
    else:
        return '$' + '\n'.join(lines) + '$'

def df_to_latex_array(df):
    """Convert DataFrame to LaTeX array."""
    n_rows, n_cols = df.shape

    # Build column format (centered columns)
    col_format = 'c' * (n_cols + 1)  # +1 for index

    # Start array
    lines = [r'\begin{array}{' + col_format + '}']
```

(continues on next page)

(continued from previous page)

```
    # Header row
    header = ' & '.join([''] + [str(c) for c in df.columns]) + r' \\'
    lines.append(header)
    lines.append(r'\hline')

    # Data rows
    for idx, row in df.iterrows():
        row_str = str(idx) + ' & ' + ' & '.join(
          [f'{v:.3f}' if isinstance(v, (int, float)) else str(v)
          for v in row]) + r' \\'
        lines.append(row_str)

    lines.append(r'\end{array}')

    return '$' + '\n'.join(lines) + '$'
```

53.2 The economic model

The data are generated by a linear-quadratic version of a stochastic optimal growth model that is an instance of models described in this quantecon lecture: *The Permanent Income Model*.

A social planner chooses a stochastic process for $\{c_t, k_{t+1}\}_{t=0}^{\infty}$ that maximizes

$$E \sum_{t=0}^{\infty} \beta^t \left(u_0 + u_1 c_t - \frac{u_2}{2} c_t^2 \right) \tag{53.1}$$

subject to the restrictions imposed by the technology

$$c_t + k_{t+1} = f k_t + \theta_t, \qquad \beta f^2 > 1. \tag{53.2}$$

Here c_t is consumption, k_t is the capital stock, f is the gross rate of return on capital, and θ_t is an endowment or technology shock following

$$a(L)\, \theta_t = \varepsilon_t, \tag{53.3}$$

where L is the backward shift (or 'lag') operator, ε_t is a scalar white noise, and $a(z) = 1 - a_1 z - a_2 z^2 - \cdots - a_r z^r$ has all its zeroes outside the unit circle.

53.2.1 Optimal decision rule

The optimal decision rule for c_t is

$$c_t = \frac{-\alpha}{f-1} + \left(1 - \frac{1}{\beta f^2}\right) \left[\frac{L - f^{-1} a(f^{-1})^{-1} a(L)}{L - f^{-1}}\, \theta_t + f k_t \right], \qquad k_{t+1} = f k_t + \theta_t - c_t, \tag{53.4}$$

where $\alpha = u_1 [1 - (\beta f)^{-1}] / u_2$.

Equations (53.3) and (53.4) exhibit the cross-equation restrictions characteristic of rational expectations models: the parameters $a_1, \ldots, a_r$ of the shock process also appear in the decision rule for consumption.

53.2.2 Net income and the accelerator

Define net output or national income as

$$y_{nt} = (f-1)k_t + \theta_t. \tag{53.5}$$

Note that (53.2) and (53.5) imply $(k_{t+1} - k_t) + c_t = y_{nt}$.

To obtain both a version of Friedman [1956]'s geometric distributed lag consumption function and a distributed lag accelerator, we impose two assumptions:

1. $a(L) = 1$, so that θ_t is white noise.
2. $\beta f = 1$, so the rate of return on capital equals the rate of time preference.

Assumption 1 is crucial for the strict form of the accelerator.

Relaxing it to allow serially correlated θ_t preserves an accelerator in a broad sense but loses the sharp geometric-lag form of (53.8).

Adding a second shock breaks the one-index structure entirely and can generate nontrivial Granger causality even without measurement error.

Assumption 2 is less important, affecting only various constants.

Under both assumptions, $\alpha = 0$ and (53.4) simplifies to

$$c_t = (1 - f^{-1})\,\theta_t + (f-1)\,k_t, \tag{53.6}$$

so that the capital stock obeys $k_{t+1} = k_t + f^{-1}\theta_t$.

When (53.6), (53.5), and (53.2) are combined, the optimal plan satisfies

$$c_t = \left(\frac{1-\beta}{1-\beta L}\right) y_{nt}, \tag{53.7}$$

$$k_{t+1} - k_t = f^{-1}\left(\frac{1-L}{1-\beta L}\right) y_{nt}, \tag{53.8}$$

$$y_{nt} = \theta_t + (1-\beta)(\theta_{t-1} + \theta_{t-2} + \cdots). \tag{53.9}$$

Equation (53.7) is Friedman's consumption model: consumption is a geometric distributed lag of income, with the decay coefficient β equal to the discount factor.

Equation (53.8) is the distributed lag accelerator: investment is a geometric distributed lag of the first difference of income.

This is the same mechanism that Chow [1968] documented empirically (see *The Acceleration Principle and the Nature of Business Cycles*).

Equation (53.9) states that the first difference of net income, $y_{nt} - y_{n,t-1} = \theta_t - \beta\theta_{t-1}$, is a first-order moving average process with innovation equal to the innovation of the endowment shock θ_t.

As Muth [1960] showed, such a process is optimally forecast via a geometric distributed lag or "adaptive expectations" scheme.

Consumption itself is a martingale: (53.6) and $k_t - k_{t-1} = f^{-1}\theta_{t-1}$ imply $c_t - c_{t-1} = (1-\beta)\theta_t$, the random walk property of consumption emphasized by Hall [1978].

So consumption, the capital stock, and net income all have unit roots, while investment $k_{t+1} - k_t = f^{-1}\theta_t$ is serially uncorrelated.

53.2.3 The accelerator puzzle

When all variables are measured accurately and are driven by the single shock θ_t, the spectral density matrix of the covariance-stationary vector $(c_t - c_{t-1},\ k_{t+1} - k_t,\ y_{nt} - y_{n,t-1})$ has rank one at all frequencies.

Each variable is an invertible one-sided distributed lag of the same white noise, so no variable Granger-causes any other.

Empirically, however, measures of output Granger-cause investment but not vice versa.

Sargent [1989] shows that measurement error can resolve this puzzle.

To illustrate, suppose first that output y_{nt} is measured perfectly while consumption and capital are each polluted by serially correlated measurement errors v_{ct} and v_{kt} orthogonal to θ_t.

Let $\bar{c}_t$ and $\bar{k}_{t+1} - \bar{k}_t$ denote the measured series. Then

$$\bar{c}_t = \left(\frac{1-\beta}{1-\beta L}\right) y_{nt} + v_{ct}, \tag{53.10}$$

$$\bar{k}_{t+1} - \bar{k}_t = \beta \left(\frac{1-L}{1-\beta L}\right) y_{nt} + (v_{k,t+1} - v_{kt}), \tag{53.11}$$

where y_{nt} is given by (53.9).

In this case income Granger-causes consumption and investment but is not Granger-caused by them.

When each measured series is corrupted by measurement error, every measured variable will in general Granger-cause every other.

The strength of this Granger causality, as measured by decompositions of j-step-ahead prediction error variances, depends on the relative variances of the measurement errors.

In this case, each observed series mixes the common signal θ_t with idiosyncratic measurement noise.

A series whose measurement error is small **relative to the variation in the true series** tracks θ_t more closely, so its innovations contain more information about future values of the other series.

Accordingly, in a forecast-error-variance decomposition, shocks to better-measured series account for a larger share of other variables' j-step-ahead prediction errors.

In a one-common-index model like this one (θ_t is the common index), better-measured variables extend more Granger causality to less well measured series than vice versa.

This asymmetry drives the numerical results we observe soon.

53.2.4 State-space formulation

Let's map the economic model and the measurement process into a linear state-space framework.

Set $f = 1.05$ and $\theta_t \sim \mathcal{N}(0, 1)$.

Define the state and observation vectors

$$x_t = \begin{bmatrix} k_t \\ \theta_t \end{bmatrix}, \qquad z_t = \begin{bmatrix} y_{nt} \\ c_t \\ \Delta k_t \end{bmatrix},$$

where $\Delta k_t \equiv k_{t+1} - k_t$ denotes investment at time t.

The error-free data are described by the state-space system

$$\begin{aligned} x_{t+1} &= Ax_t + w_t, \\ z_t &= Cx_t, \end{aligned} \tag{53.12}$$

where $w_t = \begin{bmatrix} 0 \\ \theta_{t+1} \end{bmatrix}$ has covariance $Ew_t w_t^\top = Q$ and the matrices are

$$A = \begin{bmatrix} 1 & f^{-1} \\ 0 & 0 \end{bmatrix}, \qquad C = \begin{bmatrix} f-1 & 1 \\ f-1 & 1-f^{-1} \\ 0 & f^{-1} \end{bmatrix}, \qquad Q = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}.$$

Here we date the shock w_t that arrives between t and $t+1$ by t; this matches the timing convention for the measurement errors below.

Q is singular because there is only one source of randomness θ_t; the capital stock k_t evolves deterministically given θ_t.

Notice that the rows of C satisfy $C_c + C_{\Delta k} = C_y$, which is the national income identity $c_t + \Delta k_t = y_{nt}$.

```
# Baseline structural matrices for the true economy
f = 1.05
β = 1 / f

A = np.array([
    [1.0, 1.0 / f],
    [0.0, 0.0]
])

C = np.array([
    [f - 1.0, 1.0],
    [f - 1.0, 1.0 - 1.0 / f],
    [0.0, 1.0 / f]
])

Q = np.array([
    [0.0, 0.0],
    [0.0, 1.0]
])
```

53.2.5 True impulse responses

Before introducing measurement error, we compute the impulse response of the error-free variables to a unit shock $\theta_0 = 1$.

This benchmark clarifies what changes when we later switch from error-free variables to variables reported by the statistical agency.

The response shows the investment accelerator clearly: the full impact on net income y_n occurs at lag 0, while consumption adjusts by only $1 - f^{-1} \approx 0.048$ and investment absorbs the remainder.

From lag 1 onward the economy is in its new steady state, with permanently higher capital, income, and consumption

```
def table2_irf(A, C, n_lags=6):
    x = np.array([0.0, 1.0])  # k_0 = 0, theta_0 = 1
    rows = []
    for j in range(n_lags):
        y_n, c, d_k = C @ x
        rows.append([y_n, c, d_k])
        x = A @ x
    return pd.DataFrame(rows, columns=[r'y_n', r'c', r'\Delta k'],
                        index=pd.Index(range(n_lags), name='lag'))

table2 = table2_irf(A, C, n_lags=6)
display(Latex(df_to_latex_array(table2)))
```

	y_n	c	Δk
0	1.000	0.048	0.952
1	0.048	0.048	0.000
2	0.048	0.048	0.000
3	0.048	0.048	0.000
4	0.048	0.048	0.000
5	0.048	0.048	0.000

53.3 Measurement errors

Let's add the measurement layer that generates reported data.

The econometrician does not observe z_t directly but instead sees $\bar{z}_t = z_t + v_t$, where v_t is a vector of measurement errors.

Measurement errors follow an AR(1) process

$$v_{t+1} = Dv_t + \eta_t, \tag{53.13}$$

where η_t is a vector white noise with $E\eta_t\eta_t^\top = \Sigma_\eta$ and $Ew_t v_s^\top = 0$ for all t, s.

The parameters are

$$D = \text{diag}(\rho_y, \rho_c, \rho_{\Delta k}) = \text{diag}(0.6, 0.7, 0.3), \qquad \Sigma_\eta = \text{diag}(\sigma_\eta^2), \quad \sigma_\eta = (0.05, 0.035, 0.65),$$

so the unconditional covariance of v_t is

$$R = \text{diag}\left(\frac{\sigma_{\eta,i}^2}{1-\rho_i^2}\right).$$

The innovation standard deviations σ_η are smallest for consumption (0.035), next for income (0.05), and largest for investment (0.65).

What matters for Granger-causality asymmetries, however, is the size of each measurement error **relative to the signal** in the corresponding true series.

The innovations to the true series are θ_t for income (standard deviation 1), $(1-\beta)\theta_t$ for consumption (standard deviation 0.048), and $\beta\theta_t$ for investment (standard deviation 0.95).

Relative to these, income is very well measured (a ratio of about 0.05), while consumption and investment are both poorly measured (ratios of about 0.7).

```
ρ = np.array([0.6, 0.7, 0.3])
D = np.diag(ρ)

# Innovation std. devs of η_t
σ_η = np.array([0.05, 0.035, 0.65])
Σ_η = np.diag(σ_η**2)

# Unconditional covariance of measurement errors v_t
R = np.diag((σ_η / np.sqrt(1.0 - ρ**2))**2)

print(f"f = {f},  β = 1/f = {β:.6f}")
print()
display(Latex(df_to_latex_matrix(pd.DataFrame(A), 'A')))
display(Latex(df_to_latex_matrix(pd.DataFrame(C), 'C')))
display(Latex(df_to_latex_matrix(pd.DataFrame(D), 'D')))
```

```
f = 1.05,  β = 1/f = 0.952381
```

$$A = \begin{bmatrix} 1.0000 & 0.9524 \\ 0.0000 & 0.0000 \end{bmatrix}$$

$$C = \begin{bmatrix} 0.0500 & 1.0000 \\ 0.0500 & 0.0476 \\ 0.0000 & 0.9524 \end{bmatrix}$$

$$D = \begin{bmatrix} 0.6000 & 0.0000 & 0.0000 \\ 0.0000 & 0.7000 & 0.0000 \\ 0.0000 & 0.0000 & 0.3000 \end{bmatrix}$$

We will analyze the two reporting schemes separately, but first we need a solver for the steady-state Kalman gain and error covariances.

The function below iterates on the Riccati equation until convergence, returning the Kalman gain K, the state covariance S, and the innovation covariance V

```
def steady_state_kalman(A, C_obs, Q, R, W=None, tol=1e-13, max_iter=200_000):
    """
    Solve steady-state Kalman equations for
        x_{t+1} = A x_t + w_t
        y_t     = C_obs x_t + v_t
    with cov(w)=Q, cov(v)=R, cov(w,v)=W.
    """
    n = A.shape[0]
    m = C_obs.shape[0]
    if W is None:
        W = np.zeros((n, m))

    S = Q.copy()
    for _ in range(max_iter):
        V = C_obs @ S @ C_obs.T + R
        K = (A @ S @ C_obs.T + W) @ np.linalg.inv(V)
        S_new = Q + A @ S @ A.T - K @ V @ K.T

        if np.max(np.abs(S_new - S)) < tol:
            S = S_new
            break
        S = S_new

    V = C_obs @ S @ C_obs.T + R
    K = (A @ S @ C_obs.T + W) @ np.linalg.inv(V)
    return K, S, V
```

With structural matrices and tools we need in place, we now follow Sargent [1989]'s two reporting schemes in sequence.

53.4 A classical model of measurements initially collected by an agency

A data collecting agency observes a noise-corrupted version of z_t, namely

$$\bar{z}_t = Cx_t + v_t. \tag{53.14}$$

We refer to this as *Model 1*: the agency collects noisy data and reports them without filtering.

To represent the second moments of the $\bar{z}_t$ process, it is convenient to obtain its population vector autoregression.

The error vector in the vector autoregression is the innovation to $\bar{z}_t$ and can be taken to be the white noise in a Wold moving average representation, which can be obtained by "inverting" the autoregressive representation.

The population vector autoregression, and how it depends on the parameters of the state-space system and the measurement error process, carries insights about how to interpret estimated vector autoregressions for $\bar{z}_t$.

Constructing the vector autoregression is also useful as an intermediate step in computing the likelihood of a sample of $\bar{z}_t$'s as a function of the free parameters $\{A, C, D, Q, \Sigma_\eta\}$.

The particular method that will be used to construct the vector autoregressive representation also proves useful as an intermediate step in constructing a model of an optimal reporting agency.

We use recursive (Kalman filtering) methods to obtain the vector autoregression for $\bar{z}_t$.

53.4.1 Quasi-differencing

Because the measurement errors v_t are serially correlated, the standard Kalman filter with white-noise measurement error cannot be applied directly to $\bar{z}_t = Cx_t + v_t$.

An alternative is to augment the state vector with the measurement-error AR components (see Appendix B of Sargent [1989]).

Here we take the quasi-differencing route described in Sargent [1989], which reduces the system to one with serially uncorrelated observation noise.

Define

$$\tilde{z}_t = \bar{z}_{t+1} - D\bar{z}_t, \qquad \bar{\nu}_t = Cw_t + \eta_t, \qquad \bar{C} = CA - DC. \tag{53.15}$$

Then the state-space system (53.12), the measurement error process (53.13), and the observation equation (53.14) imply the state-space system

$$\begin{aligned} x_{t+1} &= Ax_t + w_t, \\ \tilde{z}_t &= \bar{C}\, x_t + \bar{\nu}_t, \end{aligned} \tag{53.16}$$

where $(w_t, \bar{\nu}_t)$ is a white noise process with

$$E \begin{bmatrix} w_t \\ \bar{\nu}_t \end{bmatrix} \begin{bmatrix} w_t \\ \bar{\nu}_t \end{bmatrix}^\top = \begin{bmatrix} Q & W_1 \\ W_1^\top & R_1 \end{bmatrix}, \qquad R_1 = CQC^\top + \Sigma_\eta, \quad W_1 = QC^\top. \tag{53.17}$$

Notice that R_1 involves the covariance matrix Σ_η of the **innovations** η_t to the measurement errors, not the unconditional covariance matrix R of the measurement errors themselves: quasi-differencing removes the serially correlated part of v_t.

System (53.16) with covariances (53.17) is characterized by the five matrices $[A, \bar{C}, Q, R_1, W_1]$.

53.4.2 Innovations representation

Associated with (53.16) and (53.17) is the **innovations representation** for $\tilde{z}_t$,

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + K_1 u_t, \\ \tilde{z}_t &= \bar{C}\hat{x}_t + u_t, \end{aligned} \tag{53.18}$$

where

$$\begin{aligned} \hat{x}_t &= E[x_t \mid \tilde{z}_{t-1}, \tilde{z}_{t-2}, \ldots, \hat{x}_0] = E[x_t \mid \bar{z}_t, \bar{z}_{t-1}, \ldots], \\ u_t &= \tilde{z}_t - E[\tilde{z}_t \mid \tilde{z}_{t-1}, \tilde{z}_{t-2}, \ldots] = \bar{z}_{t+1} - E[\bar{z}_{t+1} \mid \bar{z}_t, \bar{z}_{t-1}, \ldots], \end{aligned} \tag{53.19}$$

$[K_1, S_1]$ are computed from the steady-state Kalman filter applied to $[A, \bar{C}, Q, R_1, W_1]$, and

$$S_1 = E[(x_t - \hat{x}_t)(x_t - \hat{x}_t)^\top]. \tag{53.20}$$

From (53.19), u_t is the innovation process for the $\bar{z}_t$ process.

53.4.3 Wold representation

System (53.18) and definition (53.15) can be used to obtain a Wold vector moving average representation for the $\bar{z}_t$ process:

$$\bar{z}_{t+1} = (I - DL)^{-1}[\bar{C}(I - AL)^{-1}K_1 L + I]u_t, \tag{53.21}$$

where L is the lag operator.

From (53.16) and (53.18) the innovation covariance is

$$V_1 = E\, u_t u_t^\top = \bar{C}\, S_1\, \bar{C}^\top + R_1. \tag{53.22}$$

Below we compute K_1, S_1, and V_1 numerically

```
C_bar = C @ A - D @ C
R1 = C @ Q @ C.T + Σ_η
W1 = Q @ C.T

K1, S1, V1 = steady_state_kalman(A, C_bar, Q, R1, W1)
```

53.4.4 Computing coefficients in a Wold moving average representation

To compute the moving average coefficients in (53.21) numerically, define the augmented state

$$r_t = \begin{bmatrix} \hat{x}_{t-1} \\ \bar{z}_{t-1} \end{bmatrix},$$

with dynamics

$$r_{t+1} = F_1 r_t + G_1 u_{t-1}, \qquad \bar{z}_t = H_1 r_t + u_{t-1},$$

where

$$F_1 = \begin{bmatrix} A & 0 \\ \bar{C} & D \end{bmatrix}, \quad G_1 = \begin{bmatrix} K_1 \\ I \end{bmatrix}, \quad H_1 = [\bar{C} \;\; D].$$

(Recall that u_{t-1} is the innovation to $\bar{z}_t$.)

The moving average coefficients are then $\psi_0 = I$ and $\psi_j = H_1 F_1^{j-1} G_1$ for $j \geq 1$.

```
F1 = np.block([
    [A, np.zeros((2, 3))],
    [C_bar, D]
])
G1 = np.vstack([K1, np.eye(3)])
H1 = np.hstack([C_bar, D])


def measured_wold_coeffs(F, G, H, n_terms=25):
    psi = [np.eye(3)]
    Fpow = np.eye(F.shape[0])
    for _ in range(1, n_terms):
        psi.append(H @ Fpow @ G)
        Fpow = Fpow @ F
    return psi


def fev_contributions(psi, V, n_horizons=20):
    """
    Returns contrib[var, shock, h-1] = contribution at horizon h.
    """
    P = linalg.cholesky(V, lower=True)
    out = np.zeros((3, 3, n_horizons))
    for h in range(1, n_horizons + 1):
        acc = np.zeros((3, 3))
        for j in range(h):
            T = psi[j] @ P
            acc += T**2
        out[:, :, h - 1] = acc
    return out


psi1 = measured_wold_coeffs(F1, G1, H1, n_terms=40)
resp1 = np.array(
      [psi1[j] @ linalg.cholesky(V1, lower=True) for j in range(14)])
decomp1 = fev_contributions(psi1, V1, n_horizons=20)
```

53.4.5 Gaussian likelihood

The Gaussian log-likelihood function for a sample $\{\bar{z}_t,\ t = 0, \dots, T\}$, conditioned on an initial state estimate $\hat{x}_0$, can be represented as

$$\mathcal{L}^* = -\tfrac{3}{2} T \ln 2\pi - \tfrac{1}{2} T \ln |V_1| - \tfrac{1}{2} \sum_{t=0}^{T-1} u_t^\top V_1^{-1} u_t, \tag{53.23}$$

where u_t is a function of $\{\bar{z}_t\}$ defined by (53.24) below and the 3 in the first term is the number of measured series.

To use (53.18) to compute $\{u_t\}$, it is useful to represent it as

$$\begin{aligned} \hat{x}_{t+1} &= (A - K_1 \bar{C})\, \hat{x}_t + K_1 \tilde{z}_t, \\ u_t &= -\bar{C}\, \hat{x}_t + \tilde{z}_t, \end{aligned} \tag{53.24}$$

where $\tilde{z}_t = \bar{z}_{t+1} - D\bar{z}_t$ is the quasi-differenced observation.

Given $\hat{x}_0$, equation (53.24) can be used recursively to compute a $\{u_t\}$ process.

Equations (53.23) and (53.24) give the likelihood function of a sample of error-corrupted data $\{\bar{z}_t\}$.

53.4.6 Forecast-error-variance decomposition

To measure the relative importance of each innovation, we decompose the j-step-ahead forecast-error variance of each measured variable.

Write $\bar{z}_{t+j} - E_t \bar{z}_{t+j} = \sum_{i=0}^{j-1} \psi_i u_{t+j-i}$, where, to lighten notation, here u_{t+j-i} denotes the innovation to $\bar{z}_{t+j-i}$.

Let P be the lower-triangular Cholesky factor of V_1 so that the orthogonalized innovations are $e_t = P^{-1} u_t$.

Then the contribution of orthogonalized innovation k to the j-step-ahead variance of variable m is $\sum_{i=0}^{j-1} (\psi_i P)_{mk}^2$.

We order the variables as y_n, c, Δk.

With this ordering, the first orthogonalized innovation is the innovation in y_n, the second is the part of the innovation in c that is orthogonal to the innovation in y_n, and the third is the part of the innovation in Δk that is orthogonal to the first two.

The table below shows the cumulative contribution of each orthogonalized innovation to the forecast-error variance of y_n, c, and Δk at horizons 1 through 20.

Each panel fixes one orthogonalized innovation and reports its cumulative contribution to each variable's forecast-error variance.

Rows are forecast horizons and columns are forecasted variables.

```
horizons = np.arange(1, 21)
labels = [r'y_n', r'c', r'\Delta k']

def fev_table(decomp, shock_idx, horizons):
    return pd.DataFrame(
        np.round(decomp[:, shock_idx, :].T, 4),
        columns=labels,
        index=pd.Index(horizons, name='j')
    )
```

```
shock_titles = [r'\text{A. Innovation in } y_n',
                r'\text{B. Innovation in } c',
                r'\text{C. Innovation in } \Delta k']

parts = []
for i, title in enumerate(shock_titles):
    arr = df_to_latex_array(fev_table(decomp1, i, horizons)).strip('$')
    parts.append(
      r'\begin{array}{c} ' + title + r' \\ ' + arr + r' \end{array}')

display(Latex('$' + r' \quad '.join(parts) + '$'))
```

A. Innovation in y_n

	y_n	c	Δk
1	1.004	0.002	0.905
2	1.006	0.004	0.905
3	1.008	0.007	0.905
4	1.011	0.009	0.905
5	1.013	0.011	0.905
6	1.015	0.014	0.905
7	1.018	0.016	0.905
8	1.020	0.018	0.905
9	1.022	0.020	0.905
10	1.024	0.023	0.905
11	1.027	0.025	0.905
12	1.029	0.027	0.905
13	1.031	0.029	0.905
14	1.034	0.032	0.905
15	1.036	0.034	0.905
16	1.038	0.036	0.905
17	1.040	0.038	0.905
18	1.043	0.041	0.905
19	1.045	0.043	0.905
20	1.047	0.045	0.905

B. Innovation in c

	y_n	c	Δk
1	0.000	0.001	0.000
2	0.000	0.002	0.000
3	0.000	0.002	0.000
4	0.000	0.002	0.000
5	0.000	0.003	0.000
6	0.000	0.003	0.000
7	0.000	0.003	0.000
8	0.000	0.003	0.000
9	0.000	0.003	0.000
10	0.000	0.003	0.000
11	0.000	0.003	0.000
12	0.000	0.003	0.000
13	0.000	0.003	0.000
14	0.000	0.003	0.000
15	0.000	0.003	0.000
16	0.000	0.003	0.000
17	0.000	0.003	0.000
18	0.000	0.003	0.000
19	0.000	0.003	0.000
20	0.000	0.003	0.000

C. Innovation in Δk

	y_n	c	Δk
1	0.000	0.000	0.425
2	0.000	0.000	0.463
3	0.000	0.000	0.466
4	0.000	0.000	0.467
5	0.000	0.000	0.467
6	0.000	0.000	0.467
7	0.000	0.000	0.467
8	0.000	0.000	0.467
9	0.000	0.000	0.467
10	0.000	0.000	0.467
11	0.000	0.000	0.467
12	0.000	0.000	0.467
13	0.000	0.000	0.467
14	0.000	0.000	0.467
15	0.000	0.000	0.467
16	0.000	0.000	0.467
17	0.000	0.000	0.467
18	0.000	0.000	0.467
19	0.000	0.000	0.467
20	0.000	0.000	0.467

The income innovation accounts for substantial proportions of forecast-error variance in all three variables, while the consumption and investment innovations contribute mainly to their own variances.

This is a **Granger causality** pattern: income appears to Granger-cause consumption and investment, but not vice versa.

This matches the paper's message that, in a one-common-index model, the relatively best measured series has the strongest predictive content.

Because income is ordered first, the decomposition attributes to the income innovation all of the variance that it shares with the other innovations.

Below we confirm the Granger-causality pattern with a calculation that does not depend on an ordering.

Let's look at the covariance matrix of the innovations

```
print('Covariance matrix of innovations:')
df_v1 = pd.DataFrame(np.round(V1, 4), index=labels, columns=labels)
display(Latex(df_to_latex_matrix(df_v1)))
```

```
Covariance matrix of innovations:
```

$$\begin{bmatrix} 1.0037 & 0.0476 & 0.9530 \\ 0.0476 & 0.0035 & 0.0453 \\ 0.9530 & 0.0453 & 1.3298 \end{bmatrix}$$

The covariance matrix of the innovations is not diagonal, but the eigenvalues are well separated as shown below

```
print('Eigenvalues of covariance matrix:')
print(np.sort(np.linalg.eigvalsh(V1))[::-1].round(4))
```

```
Eigenvalues of covariance matrix:
[2.136 0.2   0.001]
```

The first eigenvalue is much larger than the others, consistent with the presence of a dominant common shock θ_t

53.4.7 Granger causality without an ordering

The population vector autoregression for $\bar{z}_t$ gives a direct view of Granger causality.

Iterating on (53.24) gives

$$\tilde{z}_t = \sum_{j=1}^{\infty} \bar{C}(A - K_1\bar{C})^{j-1} K_1 \tilde{z}_{t-j} + u_t.$$

Let $B_j = \bar{C}(A - K_1\bar{C})^{j-1}K_1$.

Substituting $\tilde{z}_t = \bar{z}_{t+1} - D\bar{z}_t$ and collecting terms gives the vector autoregression

$$\bar{z}_{t+1} = \sum_{j=1}^{\infty} \Pi_j \bar{z}_{t+1-j} + u_t, \qquad \Pi_1 = D + B_1, \quad \Pi_j = B_j - B_{j-1}D \ \ (j \geq 2).$$

Variable i fails to Granger-cause variable m if and only if the (m, i) entries of all the Π_j are zero.

Here are the first two coefficient matrices; rows are equations and columns are lagged variables.

```
M_filter = A - K1 @ C_bar
print("eigenvalues of A - K1 C_bar:", np.linalg.eigvals(M_filter).round(4))

B_coef = [C_bar @ np.linalg.matrix_power(M_filter, j) @ K1 for j in range(3)]
Π = [D + B_coef[0]] + [B_coef[j] - B_coef[j-1] @ D for j in range(1, 3)]

eq_labels = [r'\text{eq. } ' + lab for lab in labels]
for j in range(2):
    df = pd.DataFrame(Π[j], index=eq_labels, columns=labels)
    display(Latex(r'$\Pi_' + str(j + 1) + ' = '
                  + df_to_latex_array(df).strip('$') + '$'))
```

```
eigenvalues of A - K1 C_bar: [0.932 0.594]
```

$$\Pi_1 = \begin{array}{l|ccc} & y_n & c & \Delta k \\ \hline \text{eq. } y_n & 0.056 & -0.048 & -0.004 \\ \text{eq. } c & 0.013 & 0.720 & 0.000 \\ \text{eq. } \Delta k & -0.281 & -0.037 & 0.298 \end{array}$$

$$\Pi_2 = \begin{array}{l|ccc} & y_n & c & \Delta k \\ \hline \text{eq. } y_n & 0.047 & 0.035 & -0.001 \\ \text{eq. } c & 0.012 & 0.006 & 0.000 \\ \text{eq. } \Delta k & 0.016 & 0.013 & -0.000 \end{array}$$

The eigenvalues of $A - K_1\bar{C}$ are inside the unit circle, so the vector autoregression converges.

In the investment equation, the coefficient on lagged income is large (about -0.28).

In the income equation, the coefficient on lagged investment is tiny (about -0.004).

Lagged income helps to predict investment, but lagged investment barely helps to predict income, which is the Granger-causality pattern that Sargent [1989] set out to explain.

Exercise 53.8.2 computes an ordering-free measure of the strength of Granger causality due to Geweke [1982].

53.4.8 Wold impulse responses

Impulse responses in the Wold representation are reported using orthogonalized innovations (Cholesky factorization of V_1 with ordering y_n, c, Δk).

Under this method, lag-0 responses reflect both contemporaneous covariance and the Cholesky ordering.

We first define a helper function to format the response coefficients as a LaTeX array

```
lags = np.arange(14)

def wold_response_table(resp, shock_idx, lags):
    return pd.DataFrame(
        np.round(resp[:, :, shock_idx], 4),
        columns=labels,
        index=pd.Index(lags, name='j')
    )
```

Now we report the impulse responses to each orthogonalized innovation in a single table with three panels

```
wold_titles = [r'\text{A. Response to } y_n \text{ innovation}',
               r'\text{B. Response to } c \text{ innovation}',
               r'\text{C. Response to } \Delta k \text{ innovation}']

parts = []
for i, title in enumerate(wold_titles):
    arr = df_to_latex_array(wold_response_table(resp1, i, lags)).strip('$')
    parts.append(
      r'\begin{array}{c} ' + title + r' \\ ' + arr + r' \end{array}')

display(Latex('$' + r' \quad '.join(parts) + '$'))
```

A. Response to y_n innovation

	y_n	c	Δk
0	1.002	0.048	0.951
1	0.050	0.048	0.000
2	0.049	0.048	0.000
3	0.048	0.048	0.000
4	0.048	0.048	0.000
5	0.048	0.048	0.000
6	0.048	0.048	0.000
7	0.048	0.048	0.000
8	0.048	0.048	0.000
9	0.048	0.048	0.000
10	0.048	0.048	0.000
11	0.048	0.048	0.000
12	0.048	0.048	0.000
13	0.048	0.048	0.000

B. Response to c innovation

	y_n	c	Δk
0	0.000	0.035	0.003
1	−0.002	0.025	−0.000
2	−0.000	0.019	−0.000
3	0.001	0.014	−0.000
4	0.002	0.010	−0.000
5	0.002	0.008	−0.000
6	0.002	0.006	−0.000
7	0.002	0.005	−0.000
8	0.002	0.004	−0.000
9	0.002	0.004	−0.000
10	0.002	0.003	−0.000
11	0.002	0.003	−0.000
12	0.002	0.003	−0.000
13	0.002	0.003	−0.000

C. Response to Δk innovation

	y_n	c	Δk
0	0.000	0.000	0.652
1	−0.003	0.000	0.194
2	−0.002	0.000	0.058
3	−0.001	0.000	0.018
4	−0.001	0.000	0.005
5	−0.000	0.000	0.002
6	−0.000	0.000	0.001
7	−0.000	0.000	0.000
8	0.000	0.000	0.000
9	0.000	0.000	0.000
10	0.000	0.000	0.000
11	0.000	0.000	0.000
12	0.000	0.000	0.000
13	0.000	0.000	0.000

At impact, the first orthogonalized innovation loads on all three measured variables.

At subsequent lags the income innovation generates persistent responses in all three variables because, being the best-measured series, its innovation is dominated by the true permanent shock θ_t.

The consumption and investment innovations produce responses that decay at rates close to the autoregressive coefficients of their respective measurement errors ($\rho_c = 0.7$, $\rho_{\Delta k} = 0.3$), with little spillover to other variables.

53.5 A model of optimal estimates reported by an agency

Suppose that instead of reporting the error-corrupted data $\bar{z}_t$, the data collecting agency reports linear least-squares projections of the true data on a history of the error-corrupted data.

This model provides a possible way of interpreting two features of the data-reporting process.

- *seasonal adjustment*: if the components of v_t have strong seasonals, the optimal filter will assume a shape that can be interpreted partly in terms of a seasonal adjustment filter, one that is one-sided in current and past $\bar{z}_t$'s.
- *data revisions*: if z_t contains current and lagged values of some variable of interest, then the model simultaneously determines "preliminary," "revised," and "final" estimates as successive conditional expectations based on progressively longer histories of error-ridden observations.

To make this operational, we impute to the reporting agency a model of the joint process generating the true data and the measurement errors.

We assume that the reporting agency has "rational expectations": it knows the economic and measurement structure leading to (53.16)–(53.17).

To prepare its estimates, the reporting agency itself computes the Kalman filter to obtain the innovations representation (53.18).

Rather than reporting the error-corrupted data $\bar{z}_t$, the agency reports $\hat{z}_t = G\hat{x}_t$, where G is a "selection matrix," possibly equal to C, for the data reported by the agency.

The reported data are $\hat{z}_t = G\hat{x}_t = E[Gx_t \mid \bar{z}_t, \bar{z}_{t-1}, \ldots, \hat{x}_0]$.

The state-space representation for the reported data is then

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + K_1 u_t, \\ \hat{z}_t &= G\hat{x}_t, \end{aligned} \tag{53.25}$$

where the first line of (53.25) is from the innovations representation (53.18).

Note that u_t is the innovation to $\bar{z}_{t+1}$ and is *not* the innovation to $\hat{z}_t$.

To obtain a Wold representation for $\hat{z}_t$ and the likelihood function for a sample of $\hat{z}_t$ requires that we obtain an innovations representation for (53.25).

53.5.1 Innovations representation for filtered data

To add a little generality to (53.25) we amend it to the system

$$\begin{aligned} \hat{x}_{t+1} &= A\hat{x}_t + K_1 u_t, \\ \hat{z}_t &= G\hat{x}_t + \zeta_t, \end{aligned} \tag{53.26}$$

where ζ_t is a type 2 white-noise measurement error process ("typos") with presumably very small covariance matrix R_2.

The covariance matrix of the joint noise is

$$E \begin{bmatrix} K_1 u_t \\ \zeta_t \end{bmatrix} \begin{bmatrix} K_1 u_t \\ \zeta_t \end{bmatrix}^\top = \begin{bmatrix} Q_2 & 0 \\ 0 & R_2 \end{bmatrix}, \tag{53.27}$$

where $Q_2 = K_1 V_1 K_1^\top$.

If R_2 is singular, it is necessary to adjust the Kalman filtering formulas by using transformations that induce a "reduced order observer."

In practice, we approximate a zero R_2 matrix with the matrix ϵI for a small $\epsilon > 0$ to keep the Kalman filter numerically well-conditioned.

For system (53.26) and (53.27), an innovations representation is

$$\begin{aligned} \breve{x}_{t+1} &= A\breve{x}_t + K_2 a_t, \\ \hat{z}_t &= G\breve{x}_t + a_t, \end{aligned} \tag{53.28}$$

where

$$\begin{aligned} a_t &= \hat{z}_t - E[\hat{z}_t \mid \hat{z}_{t-1}, \hat{z}_{t-2}, ...], \\ \breve{x}_t &= E[\hat{x}_t \mid \hat{z}_{t-1}, \hat{z}_{t-2}, \dots, \breve{x}_0], \\ S_2 &= E[(\hat{x}_t - \breve{x}_t)(\hat{x}_t - \breve{x}_t)^\top], \\ [K_2, S_2] &= \text{kalmanfilter}(A, G, Q_2, R_2, 0). \end{aligned} \tag{53.29}$$

Thus $\{a_t\}$ is the innovation process for the reported data $\hat{z}_t$, with innovation covariance

$$V_2 = E\, a_t a_t^\top = G\, S_2\, G^\top + R_2. \tag{53.30}$$

53.5.2 Wold representation

A Wold moving average representation for $\hat{z}_t$ is found from (53.28) to be

$$\hat{z}_t = [G(I - AL)^{-1} K_2 L + I] a_t, \tag{53.31}$$

with coefficients $\psi_0 = I$ and $\psi_j = GA^{j-1}K_2$ for $j \geq 1$.

Note that this is simpler than the Model 1 Wold representation (53.21) because there is no quasi-differencing to undo.

53.5.3 Gaussian likelihood

When a method analogous to Model 1 is used, a Gaussian log-likelihood for $\hat{z}_t$ can be computed by first computing an $\{a_t\}$ sequence from observations on $\hat{z}_t$ by using

$$\begin{aligned} \breve{x}_{t+1} &= (A - K_2 G)\, \breve{x}_t + K_2 \hat{z}_t, \\ a_t &= -G\, \breve{x}_t + \hat{z}_t. \end{aligned} \tag{53.32}$$

The likelihood function for a sample of T observations $\{\hat{z}_t\}$ is then

$$\mathcal{L}^{**} = -\tfrac{3}{2} T \ln 2\pi - \tfrac{1}{2} T \ln |V_2| - \tfrac{1}{2} \sum_{t=0}^{T-1} a_t^\top V_2^{-1} a_t. \tag{53.33}$$

Note that relative to computing the likelihood function (53.23) for the error-corrupted data, computing the likelihood function for the optimally filtered data requires more calculations.

Both likelihood functions require that the Kalman filter (53.19) be computed, while the likelihood function for the filtered data requires that the Kalman filter (53.29) also be computed.

In effect, in order to interpret and use the filtered data reported by the agency, it is necessary to retrace the steps that the agency used to synthesize those data.

The Kalman filter (53.19) is supposed to be formed by the agency.

The agency need not use Kalman filter (53.29) because it does not need the Wold representation for the filtered data.

In our parameterization $G = C$.

Because the rows of C satisfy the national income identity, the reported data satisfy $\hat{c}_t + \Delta\hat{k}_t = \hat{y}_{nt}$ exactly.

So when $R_2 = 0$, the innovation covariance matrix V_2 is singular, and the likelihood (53.33) of all three reported series is degenerate; Exercise 53.8.3 explores this "stochastic singularity."

```
Q2 = K1 @ V1 @ K1.T
ε = 1e-6

K2, S2, V2 = steady_state_kalman(A, C, Q2, ε * np.eye(3))


def filtered_wold_coeffs(A, C, K, n_terms=25):
    psi = [np.eye(3)]
    Apow = np.eye(2)
    for _ in range(1, n_terms):
        psi.append(C @ Apow @ K)
        Apow = Apow @ A
    return psi


psi2 = filtered_wold_coeffs(A, C, K2, n_terms=40)
resp2 = np.array(
  [psi2[j] @ linalg.cholesky(V2, lower=True) for j in range(14)])
decomp2 = fev_contributions(psi2, V2, n_horizons=20)
```

53.5.4 Forecast-error-variance decomposition

Because the filtered data are nearly noiseless, the innovation covariance V_2 is close to singular with one dominant eigenvalue.

This means the filtered economy is driven by essentially one shock, just like the true economy

```
parts = []
for i, title in enumerate(shock_titles):
    arr = df_to_latex_array(fev_table(decomp2, i, horizons)).strip('$')
    parts.append(
      r'\begin{array}{c} ' + title + r' \\ ' + arr + r' \end{array}')

display(Latex('$' + r' \quad '.join(parts) + '$'))
```

A. Innovation in y_n

	y_n	c	Δk
1	0.997	0.002	0.904
2	0.999	0.004	0.904
3	1.001	0.007	0.904
4	1.003	0.009	0.904
5	1.006	0.011	0.904
6	1.008	0.014	0.904
7	1.010	0.016	0.904
8	1.012	0.018	0.904
9	1.015	0.020	0.904
10	1.017	0.023	0.904
11	1.019	0.025	0.904
12	1.021	0.027	0.904
13	1.024	0.029	0.904
14	1.026	0.032	0.904
15	1.028	0.034	0.904
16	1.030	0.036	0.904
17	1.033	0.038	0.904
18	1.035	0.041	0.904
19	1.037	0.043	0.904
20	1.040	0.045	0.904

B. Innovation in c

	y_n	c	Δk
1	0.000	0.000	0.000
2	0.000	0.000	0.000
3	0.000	0.000	0.000
4	0.000	0.000	0.000
5	0.000	0.000	0.000
6	0.000	0.000	0.000
7	0.000	0.000	0.000
8	0.000	0.000	0.000
9	0.000	0.000	0.000
10	0.000	0.000	0.000
11	0.000	0.000	0.000
12	0.000	0.000	0.000
13	0.000	0.000	0.000
14	0.000	0.000	0.000
15	0.000	0.000	0.000
16	0.000	0.000	0.000
17	0.000	0.000	0.000
18	0.000	0.000	0.000
19	0.000	0.000	0.000
20	0.000	0.000	0.000

C. Innovation in Δk

	y_n	c	Δk
1	0.000	0.000	0.000
2	0.000	0.000	0.000
3	0.000	0.000	0.000
4	0.000	0.000	0.000
5	0.000	0.000	0.000
6	0.000	0.000	0.000
7	0.000	0.000	0.000
8	0.000	0.000	0.000
9	0.000	0.000	0.000
10	0.000	0.000	0.000
11	0.000	0.000	0.000
12	0.000	0.000	0.000
13	0.000	0.000	0.000
14	0.000	0.000	0.000
15	0.000	0.000	0.000
16	0.000	0.000	0.000
17	0.000	0.000	0.000
18	0.000	0.000	0.000
19	0.000	0.000	0.000
20	0.000	0.000	0.000

In Model 2, the first innovation accounts for virtually all forecast-error variance, just as in the true economy where the single structural shock θ_t drives everything.

The second and third innovations contribute negligibly.

This confirms that filtering strips away the measurement noise that created the appearance of multiple independent sources of variation in Model 1.

The covariance matrix and eigenvalues of the Model 2 innovations are

```
print('Covariance matrix of innovations:')
df_v2 = pd.DataFrame(np.round(V2, 4), index=labels, columns=labels)
display(Latex(df_to_latex_matrix(df_v2)))
```

```
Covariance matrix of innovations:
```

$$\begin{bmatrix} 0.9965 & 0.0475 & 0.9490 \\ 0.0475 & 0.0023 & 0.0452 \\ 0.9490 & 0.0452 & 0.9038 \end{bmatrix}$$

```
print('Eigenvalues of covariance matrix:')
print(np.array2string(np.sort(np.linalg.eigvalsh(V2))[::-1],
                      formatter={'float': '{:.2e}'.format}))
```

```
Eigenvalues of covariance matrix:
[1.90e+00 8.54e-06 1.00e-06]
```

One eigenvalue is large.

A second eigenvalue is positive but tiny: it reflects revisions to the agency's estimate of the capital stock, which enter the reported data only through the small coefficient $f - 1$.

The third eigenvalue equals ϵ and would be exactly zero if $R_2 = 0$, because of the national income identity.

As Sargent [1989] emphasizes, the two models of measurement produce quite different inferences about the economy's dynamics despite sharing identical underlying parameters.

53.5.5 Wold impulse responses

We again use orthogonalized Wold representation impulse responses with a Cholesky decomposition of V_2 ordered as y_n, c, Δk.

```
parts = []
for i, title in enumerate(wold_titles):
    arr = df_to_latex_array(
      wold_response_table(resp2, i, lags)).strip('$')
    parts.append(
      r'\begin{array}{c} ' + title + r' \\ ' + arr + r' \end{array}')

display(Latex('$' + r' \quad '.join(parts) + '$'))
```

A. Response to y_n innovation

	y_n	c	Δk
0	0.998	0.048	0.951
1	0.048	0.048	0.000
2	0.048	0.048	0.000
3	0.048	0.048	0.000
4	0.048	0.048	0.000
5	0.048	0.048	0.000
6	0.048	0.048	0.000
7	0.048	0.048	0.000
8	0.048	0.048	0.000
9	0.048	0.048	0.000
10	0.048	0.048	0.000
11	0.048	0.048	0.000
12	0.048	0.048	0.000
13	0.048	0.048	0.000

B. Response to c innovation

	y_n	c	Δk
0	0.000	0.002	−0.002
1	0.002	0.002	0.000
2	0.002	0.002	0.000
3	0.002	0.002	0.000
4	0.002	0.002	0.000
5	0.002	0.002	0.000
6	0.002	0.002	0.000
7	0.002	0.002	0.000
8	0.002	0.002	0.000
9	0.002	0.002	0.000
10	0.002	0.002	0.000
11	0.002	0.002	0.000
12	0.002	0.002	0.000
13	0.002	0.002	0.000

C. Response to Δk innovation

	y_n	c	Δk
0	0.000	0.000	0.002
1	−0.001	−0.001	0.000
2	−0.001	−0.001	0.000
3	−0.001	−0.001	0.000
4	−0.001	−0.001	0.000
5	−0.001	−0.001	0.000
6	−0.001	−0.001	0.000
7	−0.001	−0.001	0.000
8	−0.001	−0.001	0.000
9	−0.001	−0.001	0.000
10	−0.001	−0.001	0.000
11	−0.001	−0.001	0.000
12	−0.001	−0.001	0.000
13	−0.001	−0.001	0.000

The income innovation in Model 2 produces responses that closely approximate the true impulse response function from the structural shock θ_t.

Readers can compare the left table with the table in the *True impulse responses* section above.

The numbers are essentially the same.

The consumption and investment innovations produce responses that are orders of magnitude smaller, confirming that the filtered data are driven by essentially one shock.

Unlike Model 1, the filtered data from Model 2 *cannot* reproduce the apparent Granger causality pattern that the accelerator literature has documented empirically.

Hence, at the population level, the two measurement models imply different empirical stories even though they share the same structural economy.

- In Model 1 (raw data), measurement noise creates multiple innovations and an apparent Granger-causality pattern.
- In Model 2 (filtered data), innovations collapse back to essentially one dominant shock, mirroring the true one-index economy.

Let's verify these implications in a finite sample simulation.

53.6 Simulation

The tables above characterize population moments of the two models.

Let's simulate 80 periods of true, measured, and filtered data to compare population implications with finite-sample behavior.

First, we define a function to simulate the true economy, generate measured data with AR(1) measurement errors, and apply the Model 1 Kalman filter to produce filtered estimates

```
def simulate_series(seed=7909, T=80, k0=10.0):
    """
    Simulate true, measured, and filtered series.
    """
    rng = np.random.default_rng(seed)

    # True state/observables
    θ = rng.normal(0.0, 1.0, size=T)
    k = np.empty(T + 1)
    k[0] = k0

    y = np.empty(T)
    c = np.empty(T)
    dk = np.empty(T)

    for t in range(T):
        x_t = np.array([k[t], θ[t]])
        y[t], c[t], dk[t] = C @ x_t
        k[t + 1] = k[t] + (1.0 / f) * θ[t]

    # Measured data with AR(1) errors
    v_prev = np.zeros(3)
    v = np.empty((T, 3))
    for t in range(T):
        η_t = rng.multivariate_normal(np.zeros(3), Σ_η)
        v_prev = D @ v_prev + η_t
        v[t] = v_prev

    z_meas = np.column_stack([y, c, dk]) + v

    # Filtered data via Model 1 transformed filter
    xhat_prev = np.array([k0, 0.0])
    z_prev = C @ xhat_prev    # model-implied value of the unobserved z̄_{-1}
    z_filt = np.empty((T, 3))
    k_filt = np.empty(T)

    for t in range(T):
        z_bar_t = z_meas[t] - D @ z_prev
        u_t = z_bar_t - C_bar @ xhat_prev
        xhat_t = A @ xhat_prev + K1 @ u_t

        z_filt[t] = C @ xhat_t
        k_filt[t] = xhat_t[0]
```

(continues on next page)

(continued from previous page)

```
        xhat_prev = xhat_t
        z_prev = z_meas[t]

    out = {
        "y_true": y, "c_true": c, "dk_true": dk, "k_true": k[:-1],
        "y_meas": z_meas[:, 0], "c_meas": z_meas[:, 1],
        "dk_meas": z_meas[:, 2],
        "y_filt": z_filt[:, 0], "c_filt": z_filt[:, 1],
        "dk_filt": z_filt[:, 2], "k_filt": k_filt
    }
    return out


sim = simulate_series(seed=7909, T=80, k0=10.0)
```

We use the following helper function to plot the true series against either the measured or filtered series

```
def plot_true_vs_other(t, true_series, other_series,
                                    other_label, ylabel=""):
    fig, ax = plt.subplots(figsize=(8, 4))
    ax.plot(t, true_series, lw=2, label="true")
    ax.plot(t, other_series, lw=2, ls="--", label=other_label)
    ax.set_xlabel("time")
    ax.set_ylabel(ylabel)
    ax.legend()
    plt.tight_layout()
    plt.show()


t = np.arange(1, 81)
```

Let's first compare the true series with the measured series to see how measurement errors distort the data

```
plot_true_vs_other(t, sim["c_true"], sim["c_meas"],
                                    "measured", ylabel="consumption")
```

```
plot_true_vs_other(t, sim["dk_true"], sim["dk_meas"],
                                    "measured", ylabel="investment")
```

```
plot_true_vs_other(t, sim["y_true"], sim["y_meas"],
                                    "measured", ylabel="income")
```

Investment is distorted the most because its measurement error has the largest innovation variance ($\sigma_\eta = 0.65$), while income is distorted the least ($\sigma_\eta = 0.05$).

For the filtered series, we expect the Kalman filter to recover the true series more closely by stripping away measurement noise

```
plot_true_vs_other(t, sim["c_true"], sim["c_filt"],
                                    "filtered", ylabel="consumption")
```

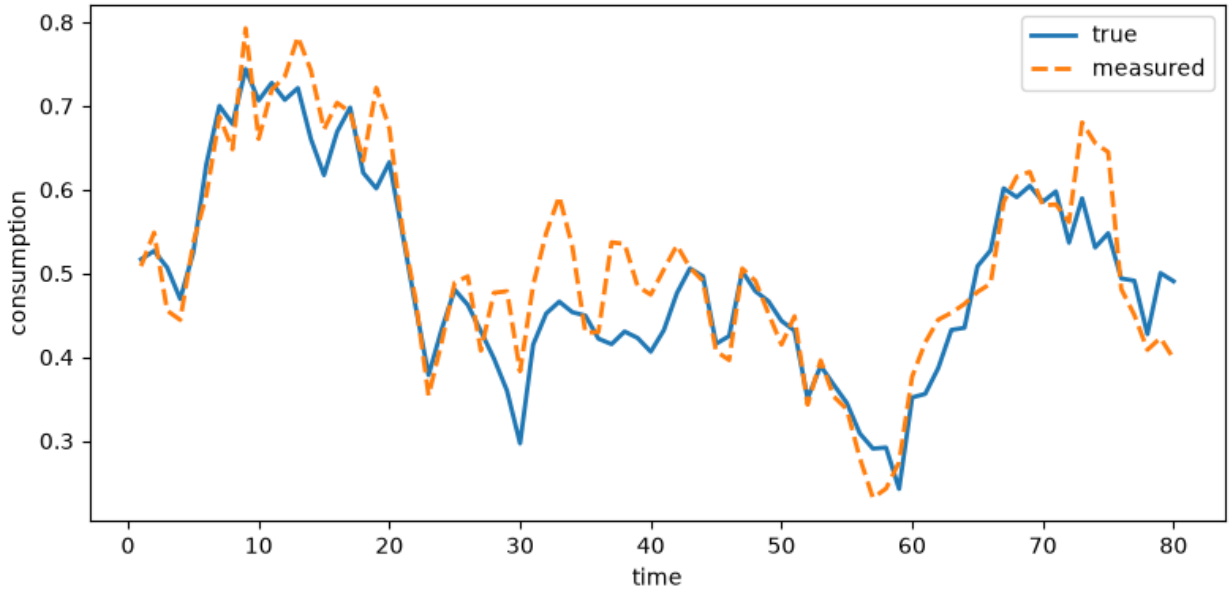


Fig. 53.1: True and measured consumption

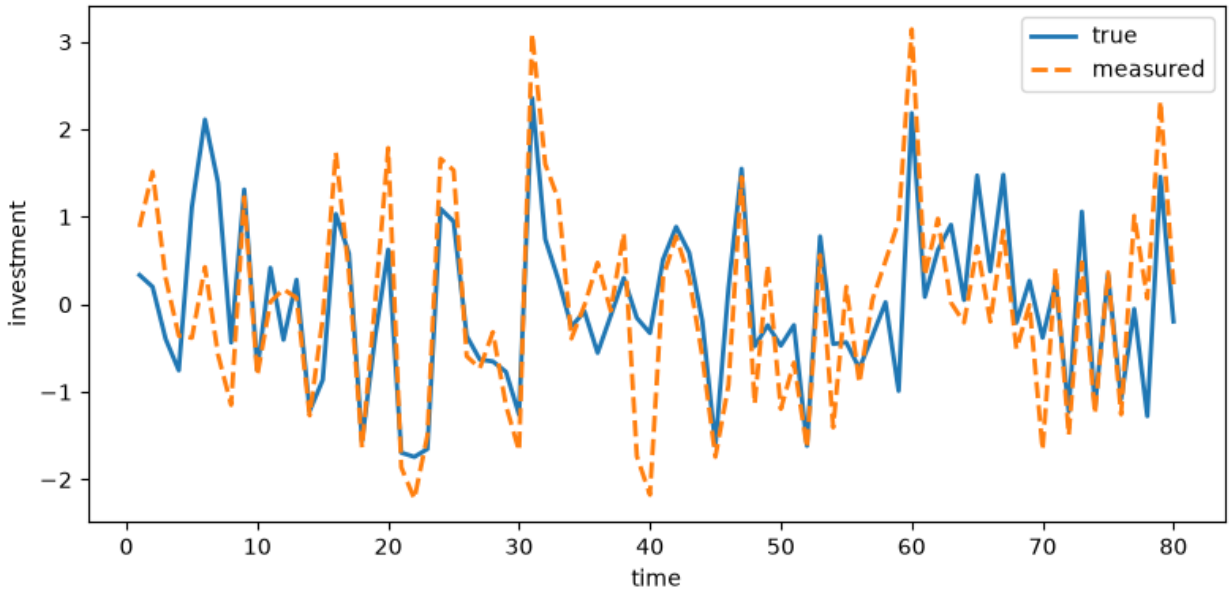


Fig. 53.2: True and measured investment

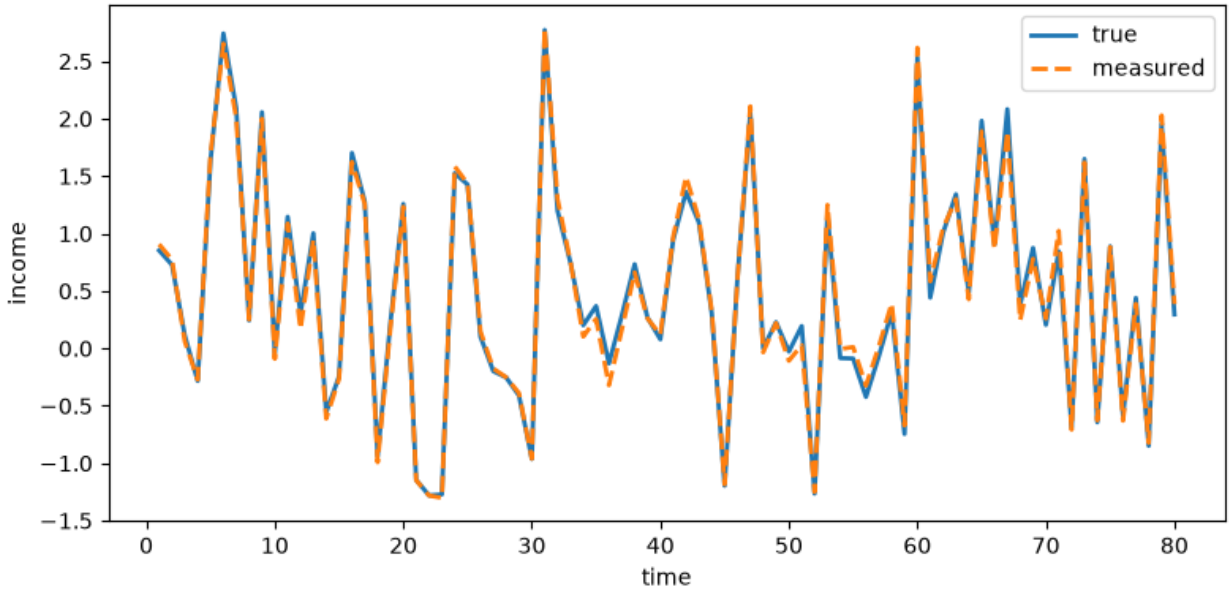


Fig. 53.3: True and measured income

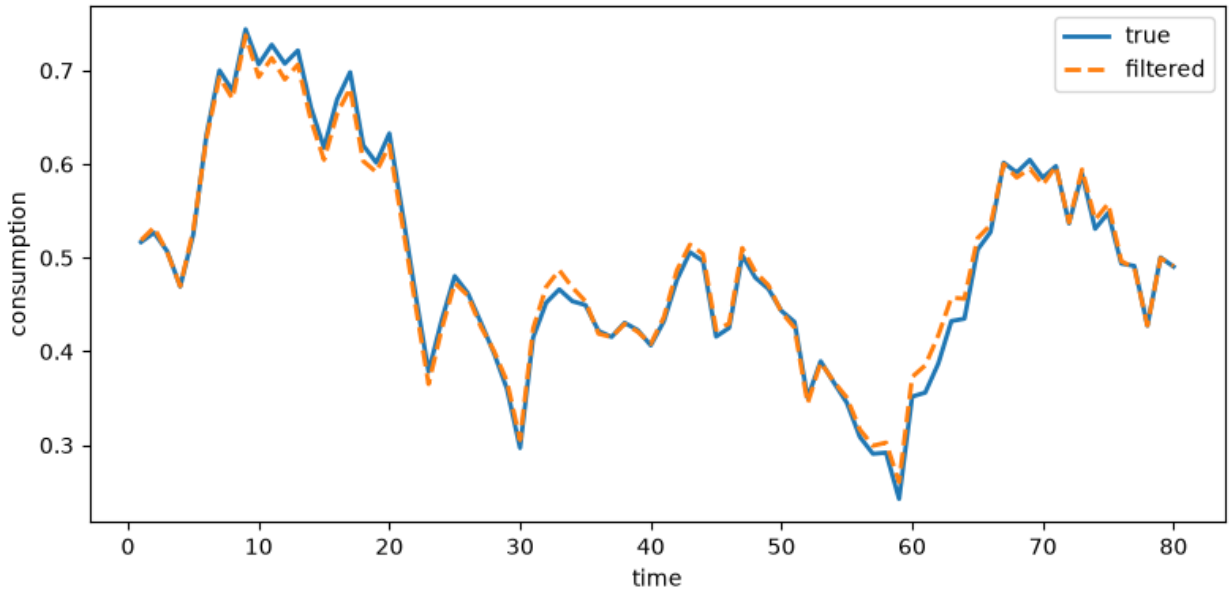


Fig. 53.4: True and filtered consumption

```
plot_true_vs_other(t, sim["dk_true"], sim["dk_filt"],
                                    "filtered", ylabel="investment")
```

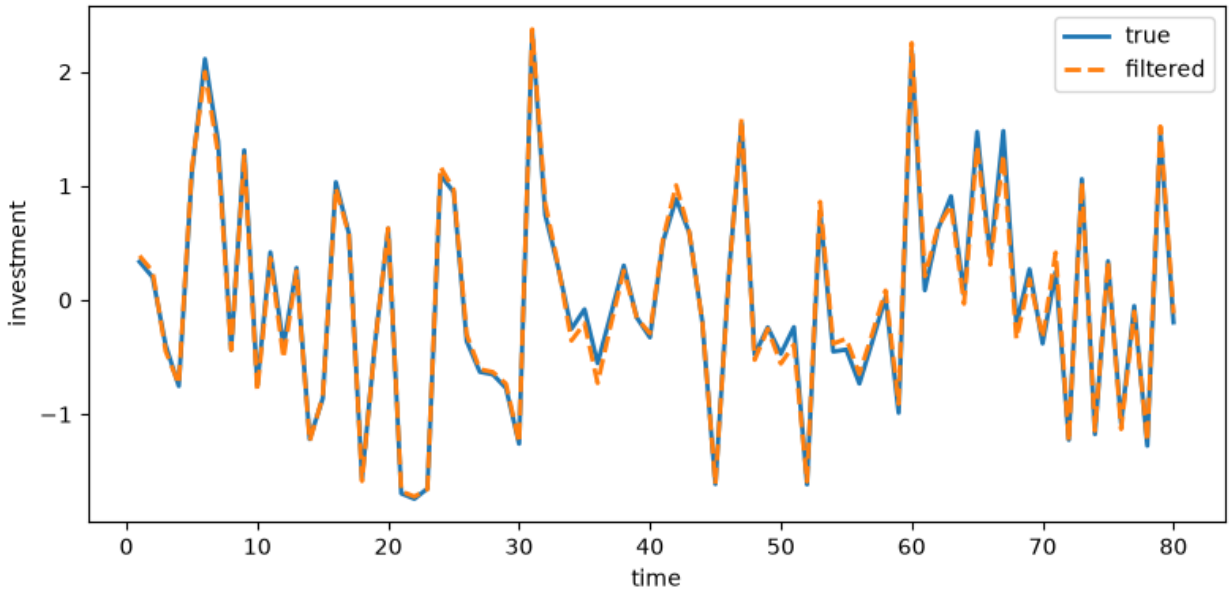


Fig. 53.5: True and filtered investment

```
plot_true_vs_other(t, sim["y_true"], sim["y_filt"],
                                    "filtered", ylabel="income")
```

```
plot_true_vs_other(t, sim["k_true"], sim["k_filt"],
                                    "filtered", ylabel="capital stock")
```

Indeed, Kalman-filtered estimates from Model 1 remove much of the measurement noise and track the truth closely.

The improvement is largest for investment, the series whose measurement error is largest; income was already well measured.

In the true model the national income identity $c_t + \Delta k_t = y_{n,t}$ holds exactly.

Independent measurement errors break this accounting identity in the measured data.

Because the filtered series are all of the form $C\hat{x}_t$, and the rows of C satisfy the identity, the filtered data satisfy it exactly.

The following figure confirms this by showing the residual $c_t + \Delta k_t - y_{n,t}$ for both measured and filtered data

```
fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(14, 4))

ax1.plot(t, sim["c_meas"] + sim["dk_meas"] - sim["y_meas"], lw=2)
ax1.axhline(0, color='black', lw=0.8, ls='--', alpha=0.5)
ax1.set_xlabel("time")
ax1.set_ylabel("measured residual")
```

(continues on next page)

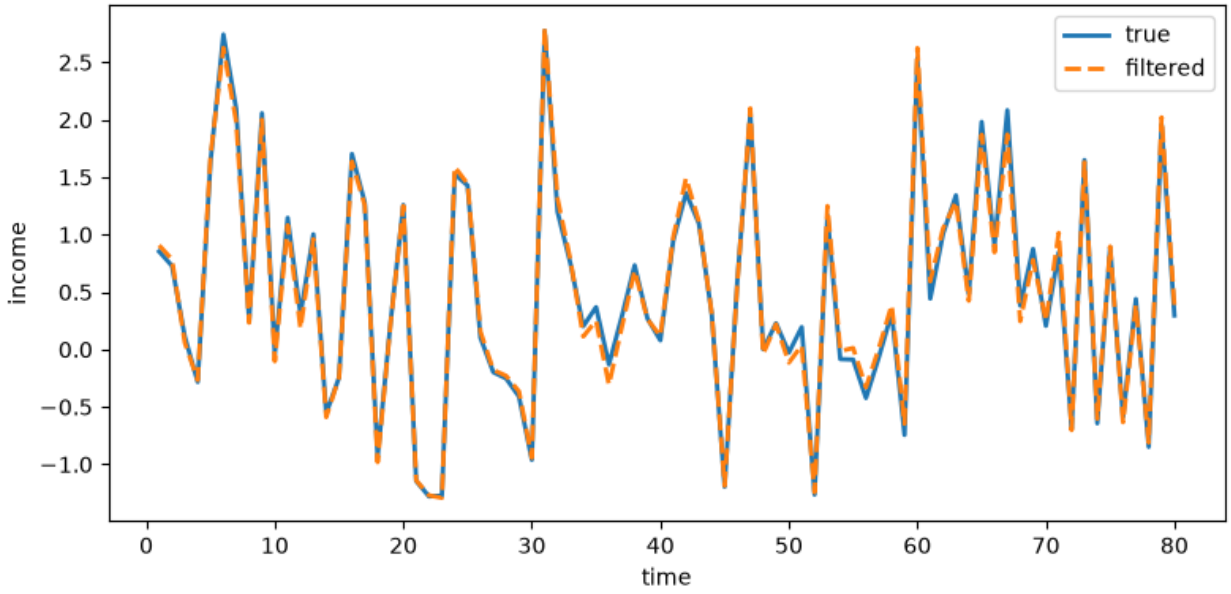


Fig. 53.6: True and filtered income

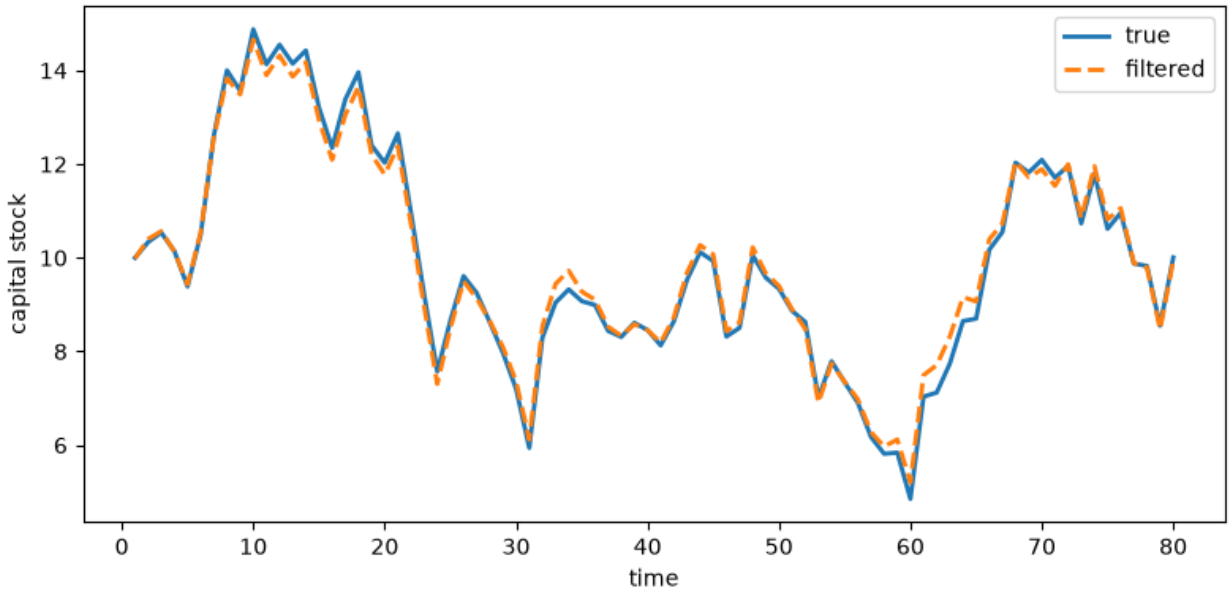


Fig. 53.7: True and filtered capital stock

(continued from previous page)

```
ax2.plot(t, sim["c_filt"] + sim["dk_filt"] - sim["y_filt"], lw=2)
ax2.axhline(0, color='black', lw=0.8, ls='--', alpha=0.5)
ax2.set_xlabel("time")
ax2.set_ylabel("filtered residual")

plt.tight_layout()
plt.show()
```

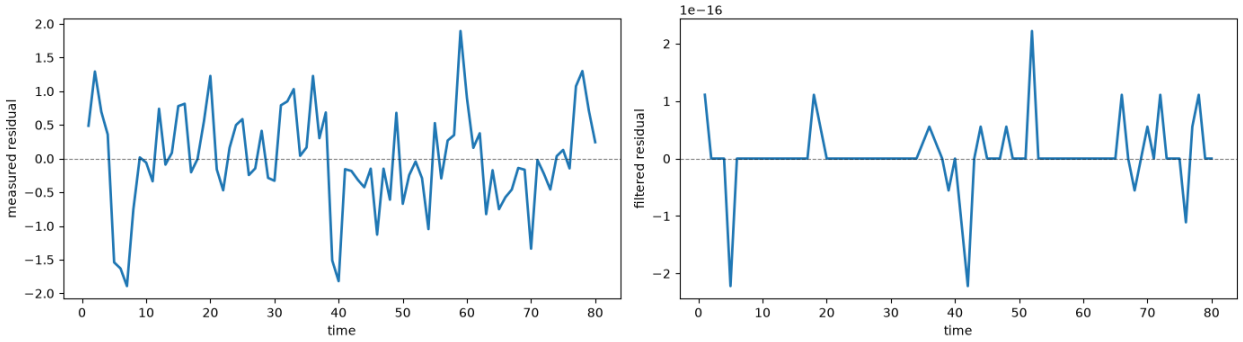


Fig. 53.8: National income identity residual

As we have predicted, the residual for the measured data is large and volatile, while the residual for the filtered data is zero up to rounding error.

53.7 Summary

Sargent [1989] shows how measurement error alters an econometrician's view of a permanent income economy driven by the investment accelerator.

The Wold representations and variance decompositions of Model 1 (raw measurements) and Model 2 (filtered measurements) differ substantially, even though the underlying economy is the same.

Measurement error can reshape inferences about which shocks drive which variables.

Model 1 reproduces the **Granger causality** pattern documented in the empirical accelerator literature: income appears to Granger-cause consumption and investment, a result Sargent [1989] attributes to measurement error and signal extraction in raw reported data.

What matters for this pattern is how large each series' measurement error is relative to the variation in the true series.

Model 2, working with filtered data, attributes nearly all variance to the single structural shock θ_t and *cannot* reproduce the Granger causality pattern.

The *Kalman filter* effectively strips measurement noise from the data, so the filtered series track the truth closely.

Raw measurement error breaks the national income accounting identity, while the filtered data satisfy it exactly, which makes the three reported series stochastically singular.

53.8 Exercises

Exercise 53.8.1

Under the assumptions $a(L) = 1$ and $\beta f = 1$, use the decision rule (53.6) and the law of motion $k_{t+1} = k_t + f^{-1}\theta_t$ to show that

1. $c_t - c_{t-1} = (1-\beta)\theta_t$, so that consumption is a martingale
2. $c_t - \beta c_{t-1} = (1-\beta)y_{nt}$, which is Friedman's consumption function (53.7)
3. $y_{nt} - y_{n,t-1} = \theta_t - \beta\theta_{t-1}$, which is equation (53.9)
4. $(k_{t+1} - k_t) - \beta(k_t - k_{t-1}) = \beta(y_{nt} - y_{n,t-1})$, which is the accelerator (53.8)

Then verify all four identities numerically using the true series in `sim`, recovering θ_t from $\theta_t = y_{nt} - (f-1)k_t$.

Solution

From $k_t - k_{t-1} = \beta\theta_{t-1}$ and (53.6),

$$c_t - c_{t-1} = (f-1)\beta\theta_{t-1} + (1-\beta)(\theta_t - \theta_{t-1}) = (1-\beta)\theta_t$$

because $(f-1)\beta = 1-\beta$.

Similarly, $y_{nt} - y_{n,t-1} = (f-1)\beta\theta_{t-1} + \theta_t - \theta_{t-1} = \theta_t - \beta\theta_{t-1}$.

For the second claim, substitute $k_t = k_{t-1} + \beta\theta_{t-1}$ into (53.6) and (53.5):

$$\begin{aligned} c_t - \beta c_{t-1} &= (f-1)(1-\beta)k_{t-1} + (1-\beta)\theta_t + [(f-1)\beta - \beta(1-\beta)]\,\theta_{t-1} \\ (1-\beta)y_{nt} &= (1-\beta)(f-1)k_{t-1} + (1-\beta)\theta_t + (1-\beta)(f-1)\beta\,\theta_{t-1} \end{aligned}$$

and use $(f-1)\beta = 1-\beta$ to see that the two right sides are equal.

Finally, $k_{t+1} - k_t = \beta\theta_t$ implies $(1-\beta L)(k_{t+1} - k_t) = \beta(\theta_t - \beta\theta_{t-1}) = \beta(y_{nt} - y_{n,t-1})$.

```
y, c, dk, k = sim["y_true"], sim["c_true"], sim["dk_true"], sim["k_true"]
θ = y - (f - 1) * k

print("martingale:  ", np.allclose(np.diff(c), (1 - β) * θ[1:]))
print("Friedman:    ", np.allclose(c[1:] - β * c[:-1], (1 - β) * y[1:]))
print("MA(1):       ", np.allclose(np.diff(y), θ[1:] - β * θ[:-1]))
print("accelerator: ", np.allclose(dk[1:] - β * dk[:-1], β * np.diff(y)))
```

```
martingale:   True
Friedman:     True
MA(1):        True
accelerator:  True
```

Exercise 53.8.2

Forecast-error-variance decompositions based on a Cholesky factorization depend on how the variables are ordered.

Geweke [1982] proposed a measure of Granger causality that does not depend on an ordering.

For measured series i and j, let σ_j^2 be the one-step-ahead prediction error variance of $\bar{z}_{j,t+1}$ given its own past only, and let $\sigma_{j|i}^2$ be its prediction error variance given the past of both $\bar{z}_i$ and $\bar{z}_j$.

Then $F_{i\to j} = \log(\sigma_j^2/\sigma_{j|i}^2) \geq 0$, with equality if and only if $\bar{z}_i$ does not Granger-cause $\bar{z}_j$.

Because the measurement errors are mutually independent AR(1) processes, both variances can be computed by applying `steady_state_kalman` to the quasi-differenced system (53.16) using only the rows of $\bar{C}$, R_1, and the columns of W_1 that correspond to the included series.

1. Compute $F_{y\to\Delta k}$, $F_{\Delta k\to y}$, $F_{y\to c}$, and $F_{c\to y}$ for the baseline measurement-error parameters.
2. Recompute them after swapping the measurement-error innovation standard deviations of income and investment, so that $\sigma_\eta = (0.65, 0.035, 0.05)$.
3. Recompute them with $\sigma_\eta = (0.3, 0.3, 0.3)$ and explain why equal measurement error variances do **not** produce symmetric Granger causality.

Solution

Here is one solution.

```
def geweke(σ_η, ρ=np.array([0.6, 0.7, 0.3]), pair=(0, 2)):
    """
    Return (F_{i->j}, F_{j->i}) for the measured series pair = (i, j).
    """
    D = np.diag(ρ)
    Σ = np.diag(np.asarray(σ_η)**2)
    C_bar = C @ A - D @ C
    R1 = C @ Q @ C.T + Σ
    W1 = Q @ C.T

    def pev(idx):
        idx = list(idx)
        return steady_state_kalman(A, C_bar[idx], Q,
                                   R1[np.ix_(idx, idx)], W1[:, idx])[2]

    i, j = pair
    V_i, V_j, V_ij = pev([i])[0, 0], pev([j])[0, 0], pev([i, j])
    return np.log(V_j / V_ij[1, 1]), np.log(V_i / V_ij[0, 0])


cases = {"baseline": (0.05, 0.035, 0.65),
         "swapped y and Δk": (0.65, 0.035, 0.05),
         "equal": (0.3, 0.3, 0.3)}

for name, σ in cases.items():
    F_yk, F_ky = geweke(σ, pair=(0, 2))
    F_yc, F_cy = geweke(σ, pair=(0, 1))
    print(f"{name:18s} y→Δk {F_yk:.4f}  Δk→y {F_ky:.4f}  "
          f"y→c {F_yc:.4f}  c→y {F_cy:.4f}")
```

```
baseline           y→Δk 0.0199  Δk→y 0.0000  y→c 0.0452  c→y 0.0001
swapped y and Δk   y→Δk 0.0000  Δk→y 0.0777  y→c 0.0045  c→y 0.0345
equal              y→Δk 0.0035  Δk→y 0.0171  y→c 0.0305  c→y 0.0009
```

In the baseline, income Granger-causes both investment and consumption, while neither Granger-causes income.

Swapping the measurement errors of income and investment reverses the direction of Granger causality between

them.

With equal measurement error variances, investment Granger-causes income more than income Granger-causes investment, and income strongly Granger-causes consumption.

What matters is not the variance of the measurement error but its size **relative to the signal** in each series.

The innovation to true investment $\beta\theta_t$ has standard deviation 0.95, while the innovation to true consumption $(1-\beta)\theta_t$ has standard deviation only 0.048.

So a measurement error with innovation standard deviation 0.3 obscures investment much less than it swamps consumption.

In the baseline, consumption's measurement error innovation (0.035) is small in absolute terms but not relative to the tiny innovation in true consumption, which is why income Granger-causes consumption.

Exercise 53.8.3

In Model 2, the agency reports $\hat{z}_t = C\hat{x}_t$.

1. Show that the reported data satisfy the national income identity $\hat{c}_t + \Delta\hat{k}_t - \hat{y}_{nt} = 0$ **exactly**, for any estimate $\hat{x}_t$.
2. Explain why this implies that the innovation covariance matrix V_2 of the reported data is singular when $R_2 = 0$.
3. Compute the eigenvalues of V_2 for $R_2 = \epsilon I$ with $\epsilon = 10^{-2}, 10^{-4}, 10^{-6}$ and interpret them.
4. What does this "stochastic singularity" imply for an econometrician who wants to compute the likelihood (53.33) from all three reported series?

Solution

The rows of C satisfy $C_c + C_{\Delta k} - C_y = [(f-1) + 0 - (f-1),\ (1-f^{-1}) + f^{-1} - 1] = [0, 0]$.

So $\hat{c}_t + \Delta\hat{k}_t - \hat{y}_{nt} = (C_c + C_{\Delta k} - C_y)\hat{x}_t = 0$ for every $\hat{x}_t$.

Because the vector $b = [-1, 1, 1]$ satisfies $b\,C\hat{x}_t = 0$ for all t, it also satisfies $b\,a_t = 0$ when $R_2 = 0$, so $bV_2b^\top = 0$ and V_2 is singular.

```
print("C_c + C_Δk − C_y =", C[1] + C[2] − C[0])
print("max |identity residual| in reported data:",
      np.max(np.abs(sim['c_filt'] + sim['dk_filt'] − sim['y_filt'])))

for ε_R in (1e-2, 1e-4, 1e-6):
    _, _, V2_ε = steady_state_kalman(A, C, Q2, ε_R * np.eye(3))
    print(f"ε = {ε_R:.0e}:  eigenvalues of V2 = "
          f"{np.array2string(np.sort(np.linalg.eigvalsh(V2_ε))[::-1], formatter={
↪'float': '{:.2e}'.format})}")
```

```
C_c + C_Δk − C_y = [0. 0.]
max |identity residual| in reported data: 2.220446049250313e-16
ε = 1e-02:  eigenvalues of V2 = [1.91e+00 1.03e-02 1.00e-02]
ε = 1e-04:  eigenvalues of V2 = [1.90e+00 1.27e-04 1.00e-04]
ε = 1e-06:  eigenvalues of V2 = [1.90e+00 8.54e-06 1.00e-06]
```

One eigenvalue of V_2 is large: it reflects news about the permanent shock θ_t.

A second eigenvalue is tiny but positive: it reflects revisions to the agency's estimate of the capital stock, which enter the reported data only through the small coefficient $f - 1$.

The third eigenvalue is approximately ϵ and vanishes as $\epsilon \to 0$ because of the identity.

An econometrician who tries to compute the Gaussian likelihood of all three reported series with $R_2 = 0$ faces a degenerate distribution: $\ln |V_2| = -\infty$.

In practice, one must either drop one of the three series, or add "type 2" reporting errors with $R_2 > 0$ as in (53.26).

In the raw data of Model 1, by contrast, independent measurement errors break the identity, so the three measured series have a nonsingular innovation covariance matrix.

Part VII

Search

CHAPTER

FIFTYFOUR

JOB SEARCH I: THE MCCALL SEARCH MODEL

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Job Search I: The McCall Search Model*
  - *Overview*
  - *The McCall Model*
  - *Computing the Optimal Policy: Take 1*
  - *Computing an Optimal Policy: Take 2*
  - *Continuous Offer Distribution*
  - *Volatility*
  - *Exercises*

> "Questioning a McCall worker is like having a conversation with an out-of-work friend: 'Maybe you are setting your sights too high', or 'Why did you quit your old job before you had a new one lined up?' This is real social science: an attempt to model, to understand, human behavior by visualizing the situation people find themselves in, the options they face and the pros and cons as they themselves see them." – Robert E. Lucas, Jr.

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

54.1 Overview

The McCall search model [McCall, 1970] helped transform economists' way of thinking about labor markets.

To clarify notions such as "involuntary" unemployment, McCall modeled the decision problem of an unemployed worker in terms of factors including

- current and likely future wages
- impatience
- unemployment compensation

To solve the decision problem McCall used dynamic programming.

Here we set up McCall's model and use dynamic programming to analyze it.

As we'll see, McCall's model is not only interesting in its own right but also an excellent vehicle for learning dynamic programming.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
import numba
import jax
import jax.numpy as jnp
from typing import NamedTuple
from functools import partial
import quantecon as qe
from quantecon.distributions import BetaBinomial
```

54.2 The McCall Model

An unemployed agent receives in each period a job offer at wage W_t.

In this lecture, we adopt the following simple environment:

- The offer sequence $\{W_t\}_{t \geq 0}$ is IID, with $q(w)$ being the probability of observing wage w in finite set $\mathbb{W}$.
- The agent observes W_t at the start of t.
- The agent knows that $\{W_t\}$ is IID with common distribution q and can use this when computing expectations.

(In later lectures, we will relax these assumptions.)

At time t, our agent has two choices:

1. Accept the offer and work permanently at constant wage W_t.
2. Reject the offer, receive unemployment compensation c, and reconsider next period.

The agent is infinitely lived and aims to maximize the expected discounted sum of earnings

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t y_t \tag{54.1}$$

The constant β lies in $(0, 1)$ and is called a **discount factor**.

The smaller is β, the more the agent discounts future earnings relative to current earnings.

The variable y_t is income, equal to

- his/her wage W_t when employed
- unemployment compensation c when unemployed

54.2.1 A Trade-Off

The worker faces a trade-off:

- Waiting too long for a good offer is costly, since the future is discounted.
- Accepting too early is costly, since better offers might arrive in the future.

To decide the optimal wait time in the face of this trade-off, we use dynamic programming.

Dynamic programming can be thought of as a two-step procedure that

1. first assigns values to "states" and
2. then deduces optimal actions given those values

We'll go through these steps in turn.

54.2.2 The Value Function

In order to optimally trade-off current and future rewards, we need to think about two things:

1. the current payoffs we get from different choices
2. the different states that those choices will lead to in next period

To weigh these two aspects of the decision problem, we need to assign *values* to states.

To this end, let $v^*(w)$ be the total lifetime value accruing to an unemployed worker who enters the current period unemployed when the wage is $w \in \mathbb{W}$.

(In particular, the agent has wage offer w in hand and can accept or reject it.)

More precisely, $v^*(w)$ denotes the total sum of expected discounted earnings when an agent always behaves in an optimal way. points in time.

Of course $v^*(w)$ is not trivial to calculate because we don't yet know what decisions are optimal and what aren't!

If we don't know what opimal choices are, it feels imposible to calculate $v^*(w)$.

But let's put this aside for now and think of v^* as a function that assigns to each possible wage w the maximal lifetime value $v^*(w)$ that can be obtained with that offer in hand.

A crucial observation is that this function v^* must satisfy

$$v^*(w) = \max \left\{ \frac{w}{1-\beta}, \, c + \beta \sum_{w' \in \mathbb{W}} v^*(w') q(w') \right\} \tag{54.2}$$

for every possible w in $\mathbb{W}$.

This is a version of the **Bellman equation**, which is ubiquitous in economic dynamics and other fields involving planning over time.

The intuition behind it is as follows:

- the first term inside the max operation is the lifetime payoff from accepting current offer, since such a worker works forever at w and values this income stream as

$$\frac{w}{1-\beta} = w + \beta w + \beta^2 w + \cdots$$

- the second term inside the max operation is the continuation value, which is the lifetime payoff from rejecting the current offer and then behaving optimally in all subsequent periods

If we optimize and pick the best of these two options, we obtain maximal lifetime value from today, given current offer w.

But this is precisely $v^*(w)$, which is the left-hand side of (54.2).

Putting this all together, we see that (54.2) is valid for all w.

54.2.3 The Optimal Policy

We still don't know how to compute v^* (although (54.2) gives us hints we'll return to below).

But suppose for now that we do know v^*.

Once we have this function in hand we can easily make optimal choices (i.e., make the right choice between accept and reject given any w).

All we have to do is select the maximal choice on the right-hand side of (54.2).

In other words, we make the best choice between stopping and continuing, given the information provided to us by v^*.

The optimal action is best thought of as a **policy**, which is, in general, a map from states to actions.

Given any w, we can read off the corresponding best choice (accept or reject) by picking the max on the right-hand side of (54.2).

Thus, we have a map from $\mathbb{W}$ to $\{0, 1\}$, with 1 meaning accept and 0 meaning reject.

We can write the policy as follows

$$\sigma(w) := \mathbf{1}\left\{\frac{w}{1-\beta} \geq c + \beta \sum_{w' \in \mathbb{W}} v^*(w')q(w')\right\}$$

Here $\mathbf{1}\{P\} = 1$ if statement P is true and equals 0 otherwise.

We can also write this as

$$\sigma(w) := \mathbf{1}\{w \geq \bar{w}\}$$

where

$$\bar{w} := (1-\beta)\left\{c + \beta \sum_{w'} v^*(w')q(w')\right\} \tag{54.3}$$

Here $\bar{w}$ (called the **reservation wage**) is a constant depending on β, c and the wage distribution.

The agent should accept if and only if the current wage offer exceeds the reservation wage.

In view of (54.3), we can compute this reservation wage if we can compute the value function.

54.3 Computing the Optimal Policy: Take 1

To put the above ideas into action, we need to compute the value function at each $w \in \mathbb{W}$.

To simplify notation, let's set

$$\mathbb{W} := \{w_1, \ldots, w_n\} \quad \text{and} \quad v^*(i) := v^*(w_i)$$

The value function is then represented by the vector $v^* = (v^*(i))_{i=1}^n$.

In view of (54.2), this vector satisfies the nonlinear system of equations

$$v^*(i) = \max\left\{\frac{w(i)}{1-\beta},\, c + \beta \sum_{j=1}^n v^*(j) q(j)\right\} \quad \text{for } i = 1, \ldots, n \tag{54.4}$$

54.3.1 The Algorithm

To compute this vector, we use successive approximations:

Step 1: pick an arbitrary initial guess $v \in \mathbb{R}^n$.

Step 2: compute a new vector $v' \in \mathbb{R}^n$ via

$$v'(i) = \max\left\{\frac{w(i)}{1-\beta},\, c + \beta \sum_{j=1}^n v(j) q(j)\right\} \quad \text{for } i = 1, \ldots, n \tag{54.5}$$

Step 3: calculate a measure of a discrepancy between v and v', such as $\max_i |v(i) - v'(i)|$.

Step 4: if the deviation is larger than some fixed tolerance, set $v = v'$ and go to step 2, else continue.

Step 5: return v.

For a small tolerance, the returned function v is a close approximation to the value function v^*.

The theory below elaborates on this point.

54.3.2 Fixed Point Theory

What's the mathematics behind these ideas?

First, one defines a mapping T from $\mathbb{R}^n$ to itself via

$$(Tv)(i) = \max\left\{\frac{w(i)}{1-\beta},\, c + \beta \sum_{j=1}^n v(j) q(j)\right\} \quad \text{for } i = 1, \ldots, n \tag{54.6}$$

(A new vector Tv is obtained from given vector v by evaluating the r.h.s. at each i.)

The element v_k in the sequence $\{v_k\}$ of successive approximations corresponds to $T^k v$.

- This is T applied k times, starting at the initial guess v

One can show that the conditions of the Banach fixed point theorem are satisfied by T on $\mathbb{R}^n$.

One implication is that T has a unique fixed point in $\mathbb{R}^n$.

- That is, a unique vector $\bar{v}$ such that $T\bar{v} = \bar{v}$.

Moreover, it's immediate from the definition of T that this fixed point is v^*.

A second implication of the Banach contraction mapping theorem is that $\{T^k v\}$ converges to the fixed point v^* regardless of v.

54.3.3 Implementation

Our default for q, the wage offer distribution, will be Beta-binomial.

```
n, a, b = 50, 200, 100                          # default parameters
q_default = jnp.array(BetaBinomial(n, a, b).pdf())
```

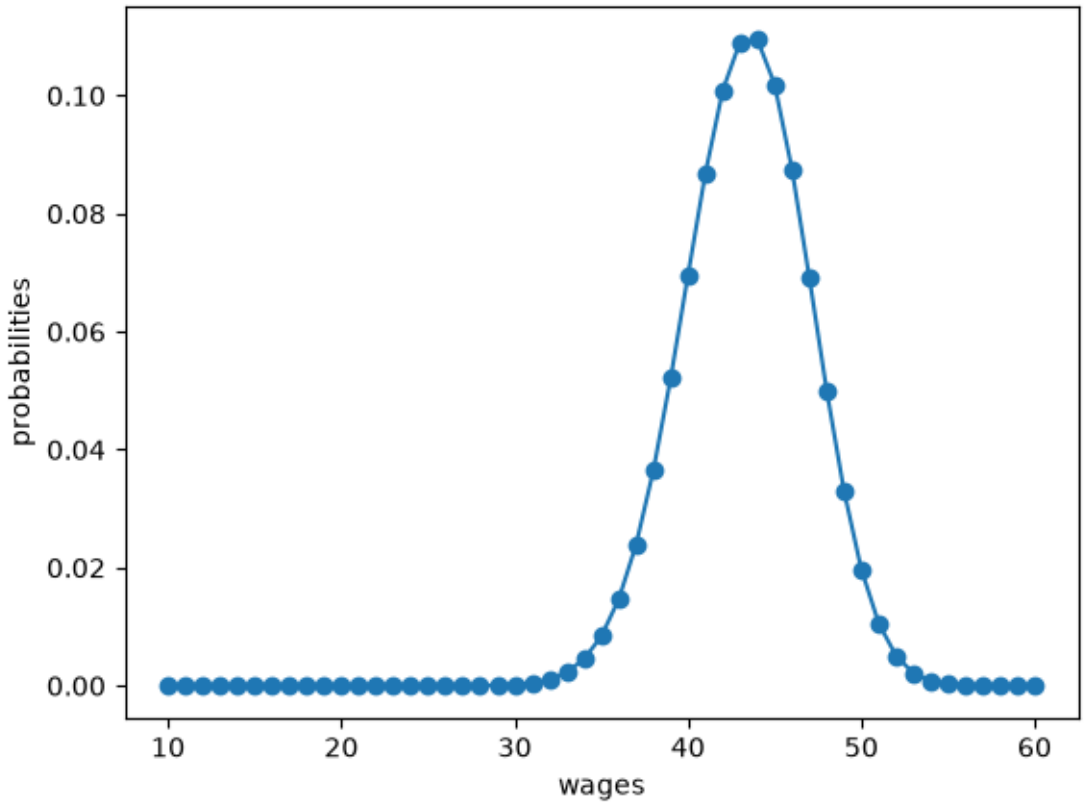
Our default set of values for wages will be

```
w_min, w_max = 10, 60
w_default = jnp.linspace(w_min, w_max, n+1)
```

Here's a plot of the probabilities of different wage outcomes:

```
fig, ax = plt.subplots()
ax.plot(w_default, q_default, '-o', label='$q(w(i))$')
ax.set_xlabel('wages')
ax.set_ylabel('probabilities')

plt.show()
```



We will use JAX to write our code.

We'll use `NamedTuple` for our model class to maintain immutability, which works well with JAX's functional programming paradigm.

Here's a class that stores the model parameters with default values.

```
class McCallModel(NamedTuple):
    c: float = 25                # unemployment compensation
    β: float = 0.99              # discount factor
    w: jnp.ndarray = w_default   # array of wage values, w[i] = wage at state i
    q: jnp.ndarray = q_default   # array of probabilities
```

We implement the Bellman operator T from (54.6), which we can write in terms of array operations as

$$Tv = \max \left\{ \frac{w}{1-\beta}, c + \beta \sum_{j=1}^{n} v(j) q(j) \right\} \tag{54.7}$$

(The first term inside the max is an array and the second is just a number – here we mean that the max comparison against this number is done element-by-element for all elements in the array.)

We can code T up as follows.

```
def T(model: McCallModel, v: jnp.ndarray):
    c, β, w, q = model
    accept = w / (1 - β)
    reject = c + β * v @ q
    return jnp.maximum(accept, reject)
```

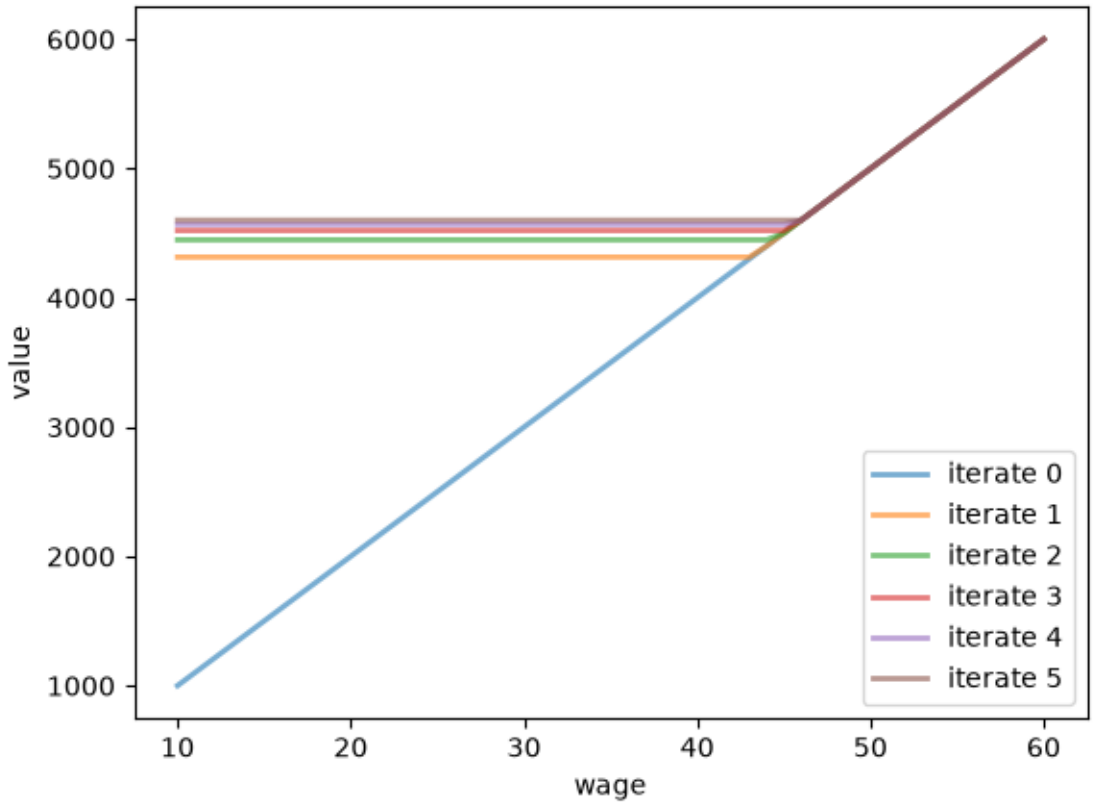
Based on these defaults, let's try plotting the first few approximate value functions in the sequence $\{T^k v\}$.

We will start from guess v given by $v(i) = w(i)/(1-\beta)$, which is the value of accepting at every given wage.

```
model = McCallModel()
c, β, w, q = model
v = w / (1 - β)    # Initial condition
fig, ax = plt.subplots()

num_plots = 6
for i in range(num_plots):
    ax.plot(w, v, '-', alpha=0.6, lw=2, label=f"iterate {i}")
    v = T(model, v)

ax.legend(loc='lower right')
ax.set_xlabel('wage')
ax.set_ylabel('value')
plt.show()
```



You can see that convergence is occurring: successive iterates are getting closer together.

Here's a more serious iteration effort to compute the limit, which continues until measured deviation between successive iterates is below `tol`.

Once we obtain a good approximation to the limit, we will use it to calculate the reservation wage.

```
def compute_reservation_wage(
        model: McCallModel,    # instance containing default parameters
        v_init: jnp.ndarray,   # initial condition for iteration
        tol: float=1e-6,       # error tolerance
        max_iter: int=500,     # maximum number of iterations for loop
    ):
    "Computes the reservation wage in the McCall job search model."
    c, β, w, q = model
    i = 0
    error = tol + 1
    v = v_init

    while i < max_iter and error > tol:
        v_next = T(model, v)
        error = jnp.max(jnp.abs(v_next - v))
        v = v_next
        i += 1

    w_bar = (1 - β) * (c + β * v @ q)
    return v, w_bar
```

The cell computes the reservation wage at the default parameters

```
model = McCallModel()
c, β, w, q = model
v_init = w / (1 - β)  # initial guess
v, w_bar = compute_reservation_wage(model, v_init)
print(w_bar)
```

```
47.316475
```

54.3.4 Comparative Statics

Now that we know how to compute the reservation wage, let's see how it varies with parameters.

Here we compare the reservation wage at two values of β.

The reservation wages will be plotted alongside the wage offer distribution, so that we can get a sense of what fraction of offers will be accepted.

```
fig, ax = plt.subplots()

# Get the default color cycle
prop_cycle = plt.rcParams['axes.prop_cycle']
colors = prop_cycle.by_key()['color']

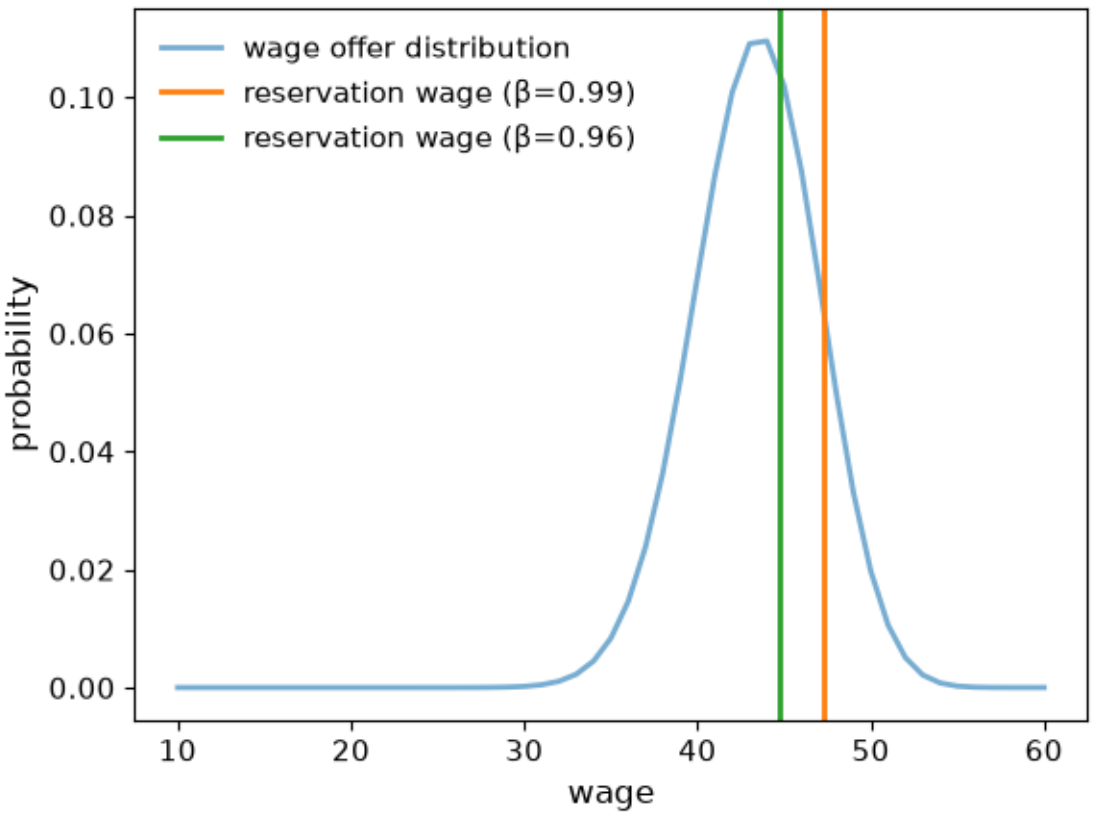
# Plot the wage offer distribution
ax.plot(w, q, '-', alpha=0.6, lw=2,
        label='wage offer distribution',
        color=colors[0])

# Compute reservation wage with default beta
model_default = McCallModel()
c, β, w, q = model_default
v_init = w / (1 - β)
v_default, res_wage_default = compute_reservation_wage(
    model_default, v_init
)

# Compute reservation wage with lower beta
β_new = 0.96
model_low_beta = McCallModel(β=β_new)
c, β_low, w, q = model_low_beta
v_init_low = w / (1 - β_low)
v_low, res_wage_low = compute_reservation_wage(
    model_low_beta, v_init_low
)

# Plot vertical lines for reservation wages
ax.axvline(x=res_wage_default, color=colors[1], lw=2,
           label=f'reservation wage (β={β})')
ax.axvline(x=res_wage_low, color=colors[2], lw=2,
           label=f'reservation wage (β={β_new})')

ax.set_xlabel('wage', fontsize=12)
ax.set_ylabel('probability', fontsize=12)
ax.tick_params(axis='both', which='major', labelsize=11)
ax.legend(loc='upper left', frameon=False, fontsize=11)
plt.show()
```



We see that the reservation wage is higher when β is higher.

This is not surprising, since higher β is associated with more patience.

Now let's look more systematically at what happens when we change β and c.

As a first step, given that we'll use it many times, let's create a more efficient, jit-complied version of the function that computes the reservation wage:

```
@jax.jit
def compute_res_wage_jitted(
        model: McCallModel,    # instance containing default parameters
        v_init: jnp.ndarray,   # initial condition for iteration
        tol: float=1e-6,       # error tolerance
        max_iter: int=500,     # maximum number of iterations for loop
    ):
    c, β, w, q = model
    i = 0
    error = tol + 1
    initial_state = v_init, i, error

    def cond(loop_state):
        v, i, error = loop_state
        return jnp.logical_and(i < max_iter, error > tol)

    def update(loop_state):
        v, i, error = loop_state
        v_next = T(model, v)
        error = jnp.max(jnp.abs(v_next - v))
        i += 1
```

(continues on next page)

(continued from previous page)

```
        new_loop_state = v_next, i, error
        return new_loop_state

    final_state = jax.lax.while_loop(cond, update, initial_state)
    v, i, error = final_state

    w_bar = (1 - β) * (c + β * v @ q)
    return v, w_bar
```

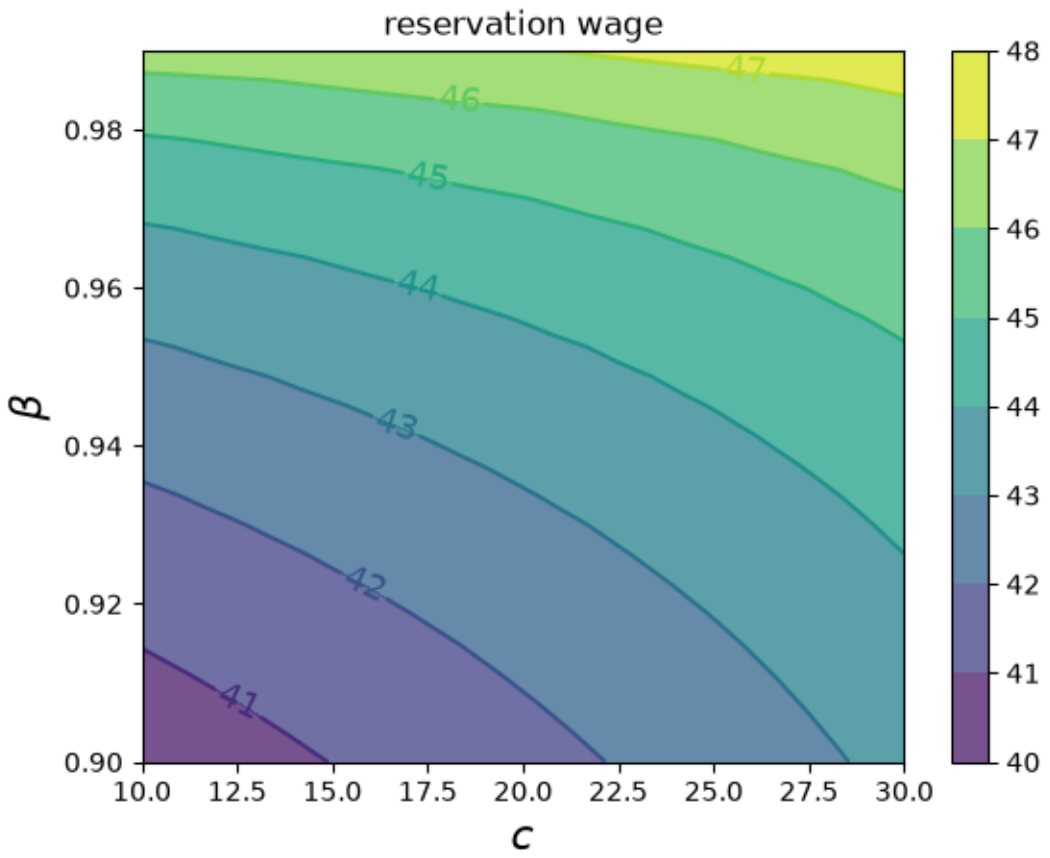
Now we compute the reservation wage at each c, β pair.

```
grid_size = 25
c_vals = jnp.linspace(10.0, 30.0, grid_size)
β_vals = jnp.linspace(0.9, 0.99, grid_size)

res_wage_matrix = np.empty((grid_size, grid_size))
model = McCallModel()
v_init = model.w / (1 - model.β)

for i, c in enumerate(c_vals):
    for j, β in enumerate(β_vals):
        model = McCallModel(c=c, β=β)
        v, w_bar = compute_res_wage_jitted(model, v_init)
        v_init = v
        res_wage_matrix[i, j] = w_bar

fig, ax = plt.subplots()
cs1 = ax.contourf(c_vals, β_vals, res_wage_matrix.T, alpha=0.75)
ctr1 = ax.contour(c_vals, β_vals, res_wage_matrix.T)
plt.clabel(ctr1, inline=1, fontsize=13)
plt.colorbar(cs1, ax=ax)
ax.set_title("reservation wage")
ax.set_xlabel("$c$", fontsize=16)
ax.set_ylabel("$β$", fontsize=16)
ax.ticklabel_format(useOffset=False)
plt.show()
```



As expected, the reservation wage increases with both patience and unemployment compensation.

54.4 Computing an Optimal Policy: Take 2

The approach to dynamic programming just described is standard and broadly applicable.

But for our McCall search model there's also an easier way that circumvents the need to compute the value function.

Let h denote the continuation value:

$$h = c + \beta \sum_{w'} v^*(w') q(w') \tag{54.8}$$

The Bellman equation can now be written as

$$v^*(w') = \max \left\{ \frac{w'}{1-\beta}, \, h \right\} \tag{54.9}$$

Now let's derive a nonlinear equation for h alone.

Starting from (54.9), we multiply both sides by $q(w')$ to get

$$v^*(w') q(w') = \max \left\{ \frac{w'}{1-\beta}, h \right\} q(w')$$

Next, we sum both sides over $w' \in \mathbb{W}$:

$$\sum_{w' \in \mathbb{W}} v^*(w') q(w') = \sum_{w' \in \mathbb{W}} \max \left\{ \frac{w'}{1-\beta}, h \right\} q(w')$$

Now multiply both sides by β:

$$\beta \sum_{w' \in \mathbb{W}} v^*(w') q(w') = \beta \sum_{w' \in \mathbb{W}} \max \left\{ \frac{w'}{1-\beta}, h \right\} q(w')$$

Add c to both sides:

$$c + \beta \sum_{w' \in \mathbb{W}} v^*(w') q(w') = c + \beta \sum_{w' \in \mathbb{W}} \max \left\{ \frac{w'}{1-\beta}, h \right\} q(w')$$

Finally, using the definition of h from (54.8), the left-hand side is just h, giving us

$$h = c + \beta \sum_{w' \in \mathbb{W}} \max \left\{ \frac{w'}{1-\beta}, h \right\} q(w') \tag{54.10}$$

This is a nonlinear equation in the single scalar h that we can solve for h.

As before, we will use successive approximations:

Step 1: pick an initial guess h.

Step 2: compute the update h' via

$$h' = c + \beta \sum_{w' \in \mathbb{W}} \max \left\{ \frac{w'}{1-\beta}, h \right\} q(w') \tag{54.11}$$

Step 3: calculate the deviation $|h - h'|$.

Step 4: if the deviation is larger than some fixed tolerance, set $h = h'$ and go to step 2, else return h.

One can again use the Banach contraction mapping theorem to show that this process always converges.

The big difference here, however, is that we're iterating on a scalar h, rather than an n-vector, $v(i), i = 1, \ldots, n$.

Here's an implementation:

```
def compute_reservation_wage_two(
        model: McCallModel,    # instance containing default parameters
        tol: float=1e-5,       # error tolerance
        max_iter: int=500,     # maximum number of iterations for loop
    ):
    c, β, w, q = model
    h = (w @ q) / (1 - β)  # initial condition
    i = 0
    error = tol + 1
    initial_loop_state = i, h, error

    def cond(loop_state):
        i, h, error = loop_state
        return jnp.logical_and(i < max_iter, error > tol)

    def update(loop_state):
        i, h, error = loop_state
        s = jnp.maximum(w / (1 - β), h)
        h_next = c + β * (s @ q)
        error = jnp.abs(h_next - h)
        i_next = i + 1
        new_loop_state = i_next, h_next, error
        return new_loop_state
```

(continues on next page)

(continued from previous page)

```
    final_state = jax.lax.while_loop(cond, update, initial_loop_state)
    i, h, error = final_state

    # Compute and return the reservation wage
    return (1 - β) * h
```

You can use this code to solve the exercise below.

54.5 Continuous Offer Distribution

The discrete wage offer distribution used above is convenient for theory and computation, but many realistic distributions are continuous (i.e., have a density).

Fortunately, the theory changes little in our simple model when we shift to a continuous offer distribution.

Recall that h in (54.8) denotes the value of not accepting a job in this period but then behaving optimally in all subsequent periods.

To shift to a continuous offer distribution, we can replace (54.8) by

$$h = c + \beta \int v^*(s') q(s') ds'. \tag{54.12}$$

Equation (54.10) becomes

$$h = c + \beta \int \max \left\{ \frac{w(s')}{1-\beta}, h \right\} q(s') ds' \tag{54.13}$$

The aim is to solve this nonlinear equation by iteration, and from it obtain the reservation wage.

54.5.1 Implementation with Lognormal Wages

Let's implement this for the case where

- the state sequence $\{s_t\}$ is IID and standard normal and
- the wage function is $w(s) = \exp(\mu + \sigma s)$.

This gives us a lognormal wage distribution.

We use Monte Carlo integration to evaluate the integral, averaging over a large number of wage draws.

For default parameters, we use `c=25, β=0.99, σ=0.5, μ=2.5`.

```
class McCallModelContinuous(NamedTuple):
    c: float              # unemployment compensation
    β: float              # discount factor
    σ: float              # scale parameter in lognormal distribution
    μ: float              # location parameter in lognormal distribution
    w_draws: jnp.ndarray  # draws of wages for Monte Carlo


def create_mccall_continuous(
        c=25, β=0.99, σ=0.5, μ=2.5, mc_size=1000, seed=1234
    ):
    key = jax.random.key(seed)
```

(continues on next page)

(continued from previous page)

```
    s = jax.random.normal(key, (mc_size,))
    w_draws = jnp.exp(μ + σ * s)
    return McCallModelContinuous(c, β, σ, μ, w_draws)


@jax.jit
def compute_reservation_wage_continuous(model, max_iter=500, tol=1e-5):
    c, β, σ, μ, w_draws = model

    h = jnp.mean(w_draws) / (1 - β)  # initial guess

    def update(state):
        h, i, error = state
        integral = jnp.mean(jnp.maximum(w_draws / (1 - β), h))
        h_next = c + β * integral
        error = jnp.abs(h_next - h)
        return h_next, i + 1, error

    def cond(state):
        h, i, error = state
        return jnp.logical_and(i < max_iter, error > tol)

    initial_state = (h, 0, tol + 1)
    final_state = jax.lax.while_loop(cond, update, initial_state)
    h_final, _, _ = final_state

    # Now compute the reservation wage
    return (1 - β) * h_final
```

Now let's investigate how the reservation wage changes with c and β using a contour plot.

```
grid_size = 25
c_vals = jnp.linspace(10.0, 30.0, grid_size)
β_vals = jnp.linspace(0.9, 0.99, grid_size)

def compute_R_element(c, β):
    model = create_mccall_continuous(c=c, β=β)
    return compute_reservation_wage_continuous(model)

# First, vectorize over β (holding c fixed)
compute_R_over_β = jax.vmap(compute_R_element, in_axes=(None, 0))

# Next, vectorize over c (applying the above function to each c)
compute_R_vectorized = jax.vmap(compute_R_over_β, in_axes=(0, None))

# Apply to compute the full grid
R = compute_R_vectorized(c_vals, β_vals)
```

```
fig, ax = plt.subplots()

cs1 = ax.contourf(c_vals, β_vals, R.T, alpha=0.75)
ctr1 = ax.contour(c_vals, β_vals, R.T)

plt.clabel(ctr1, inline=1, fontsize=13)
plt.colorbar(cs1, ax=ax)
```

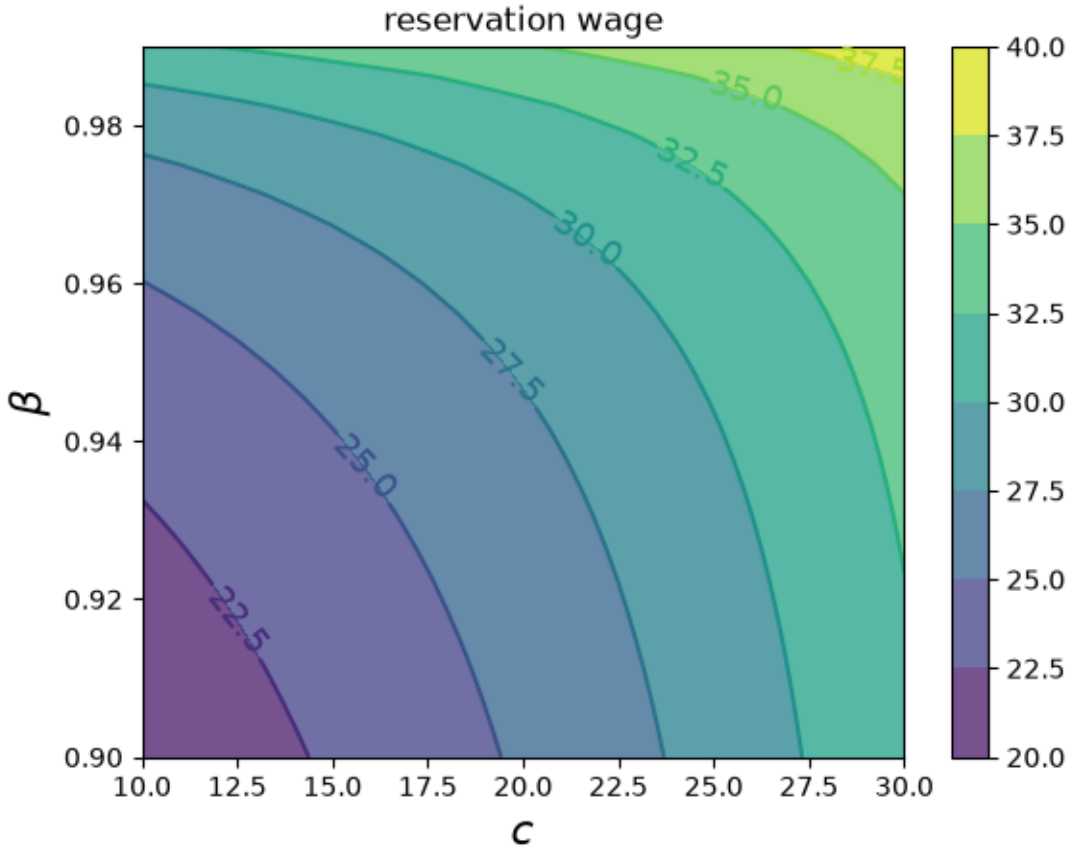
(continues on next page)

(continued from previous page)

```
ax.set_title("reservation wage")
ax.set_xlabel("$c$", fontsize=16)
ax.set_ylabel("$β$", fontsize=16)

ax.ticklabel_format(useOffset=False)

plt.show()
```



As with the discrete case, the reservation wage increases with both patience and unemployment compensation.

54.6 Volatility

An interesting feature of the McCall model is that increased volatility in wage offers tends to increase the reservation wage.

The intuition is that volatility is attractive to the worker because they can enjoy the upside (high wage offers) while rejecting the downside (low wage offers).

Hence, with more volatility, workers are more willing to continue searching rather than accept a given offer, which means the reservation wage rises.

To illustrate this phenomenon, we use a mean-preserving spread of the wage distribution.

In particular, we vary the scale parameter σ in the lognormal wage distribution $w(s) = \exp(\mu + \sigma s)$ while adjusting μ to keep the mean constant.

Recall that for a lognormal distribution with parameters μ and σ, the mean is $\exp(\mu + \sigma^2/2)$.

To keep the mean constant at some value m, we need:

$$\mu = \ln(m) - \frac{\sigma^2}{2}$$

Let's implement this and compute the reservation wage for different values of σ:

```
# Fix the mean wage
mean_wage = 20.0

# Create a range of volatility values
σ_vals = jnp.linspace(0.1, 1.0, 25)

# Given σ, compute μ to maintain constant mean
def compute_μ_for_mean(σ, mean_wage):
    return jnp.log(mean_wage) - (σ**2) / 2

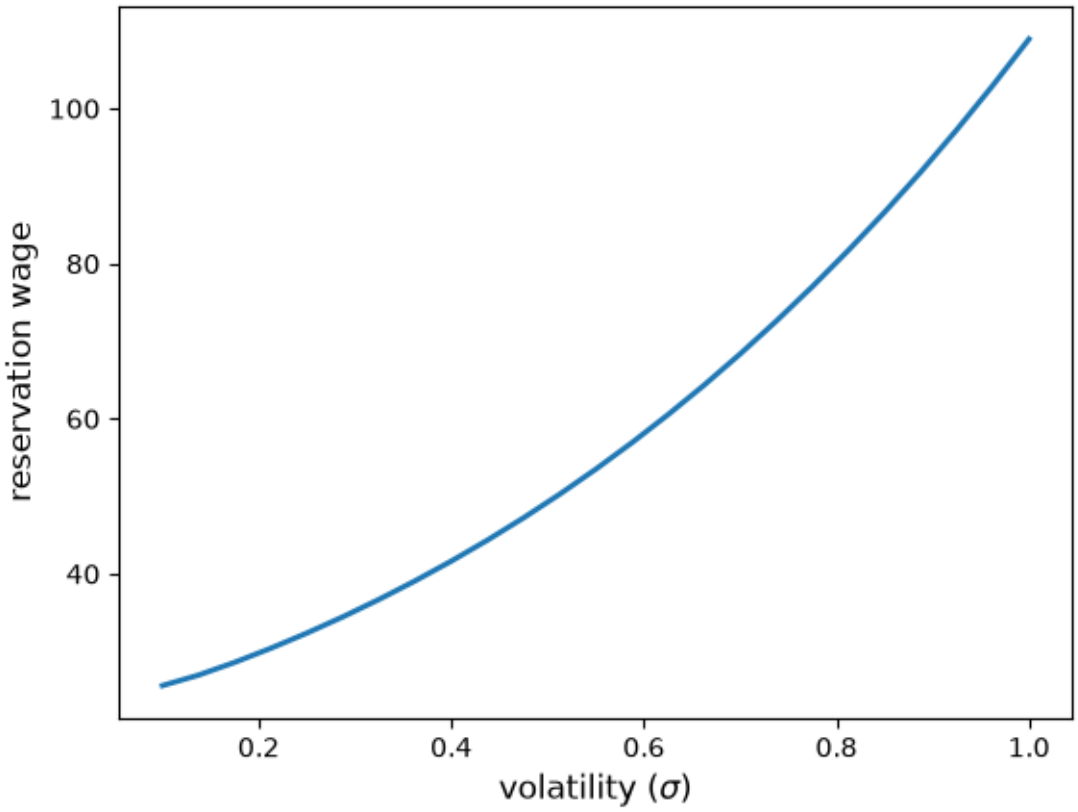
# Compute reservation wage for each volatility level
res_wages_volatility = []

for σ in σ_vals:
    μ = compute_μ_for_mean(σ, mean_wage)
    model = create_mccall_continuous(σ=float(σ), μ=float(μ))
    w_bar = compute_reservation_wage_continuous(model)
    res_wages_volatility.append(w_bar)

res_wages_volatility = jnp.array(res_wages_volatility)
```

Now let's plot the reservation wage as a function of volatility:

```
fig, ax = plt.subplots()
ax.plot(σ_vals, res_wages_volatility, linewidth=2)
ax.set_xlabel(r'volatility ($\sigma$)', fontsize=12)
ax.set_ylabel('reservation wage', fontsize=12)
plt.show()
```



As expected, the reservation wage is increasing in σ.

54.6.1 Lifetime Value and Volatility

We've seen that the reservation wage increases with volatility.

It's also the case that maximal lifetime value increases with volatility.

Higher volatility provides more upside potential, while at the same time workers can protect themselves against downside risk by rejecting low offers.

This option value translates into higher expected lifetime utility.

To demonstrate this, we will:

1. Compute the reservation wage for each volatility level
2. Calculate the expected discounted value of the lifetime income stream associated with that reservation wage, using Monte Carlo.

The simulation works as follows:

1. Compute the present discounted value of one lifetime earnings path, from a given wage path.
2. Average over a large number of such calculations to approximate expected discounted value.

We truncate each path at $T = 100$, which provides sufficient resolution for our purposes.

```
@jax.jit
def simulate_lifetime_value(key, model, w_bar, n_periods=100):
    """
```

(continues on next page)

(continued from previous page)

```
    Simulate one realization of the wage path and compute lifetime value.

    Parameters:
    -----------
    key : jax.random.key
        Random key for JAX
    model : McCallModelContinuous
        The model containing parameters
    w_bar : float
        The reservation wage
    n_periods : int
        Number of periods to simulate

    Returns:
    --------
    lifetime_value : float
        Discounted sum of income over n_periods
    """
    c, β, σ, μ, w_draws = model

    # Draw all wage offers upfront
    key, subkey = jax.random.split(key)
    s_vals = jax.random.normal(subkey, (n_periods,))
    wage_offers = jnp.exp(μ + σ * s_vals)

    # Determine which offers are acceptable
    accept = wage_offers >= w_bar

    # Track employment status: employed from first acceptance onward
    employed = jnp.cumsum(accept) > 0

    # Get the accepted wage (first wage where accept is True)
    first_accept_idx = jnp.argmax(accept)
    accepted_wage = wage_offers[first_accept_idx]

    # Earnings at each period: accepted_wage if employed, c if unemployed
    earnings = jnp.where(employed, accepted_wage, c)

    # Compute discounted sum
    periods = jnp.arange(n_periods)
    discount_factors = β ** periods
    lifetime_value = jnp.sum(discount_factors * earnings)

    return lifetime_value


@jax.jit
def compute_mean_lifetime_value(model, w_bar, num_reps=10000, seed=1234):
    """
    Compute mean lifetime value across many simulations.

    """
    key = jax.random.key(seed)
    keys = jax.random.split(key, num_reps)

    # Vectorize the simulation across all replications
    simulate_fn = jax.vmap(simulate_lifetime_value, in_axes=(0, None, None))
```

(continues on next page)

(continued from previous page)

```
    lifetime_values = simulate_fn(keys, model, w_bar)
    return jnp.mean(lifetime_values)
```

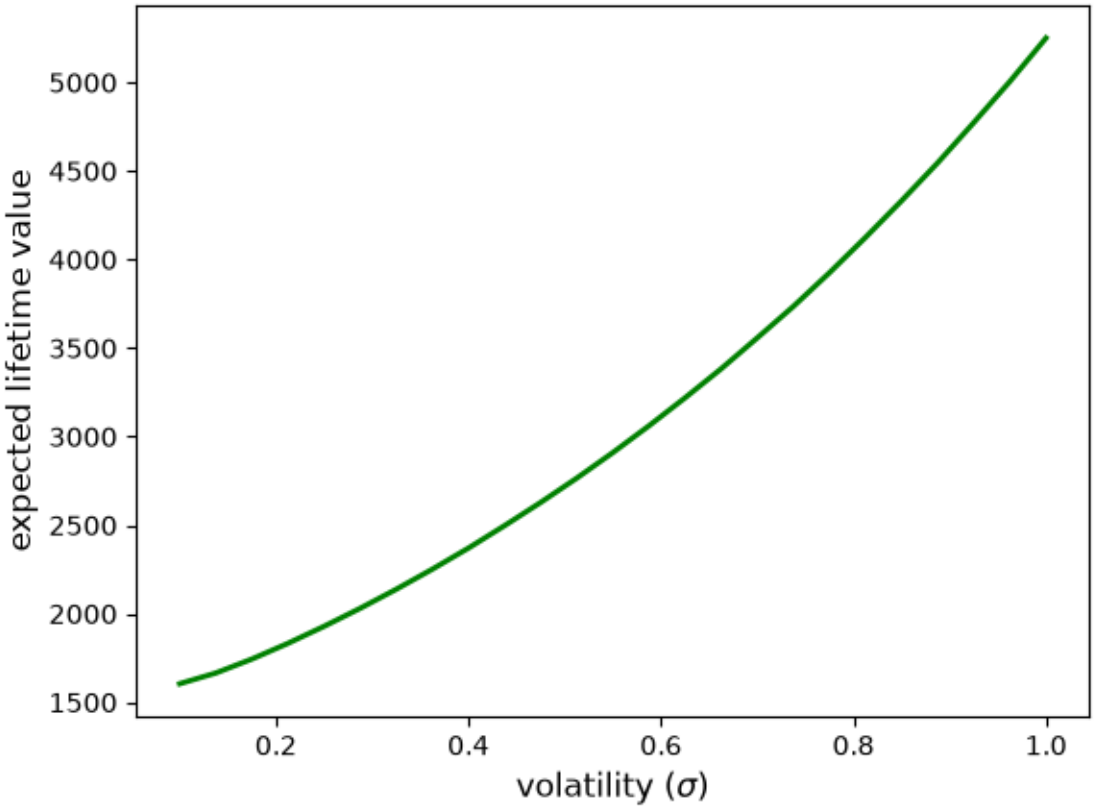
Now let's compute the expected lifetime value for each volatility level:

```
# Use the same volatility range and mean wage
σ_vals = jnp.linspace(0.1, 1.0, 25)
mean_wage = 20.0

lifetime_vals = []
for σ in σ_vals:
    μ = compute_μ_for_mean(σ, mean_wage)
    model = create_mccall_continuous(σ=σ, μ=μ)
    w_bar = compute_reservation_wage_continuous(model)
    lv = compute_mean_lifetime_value(model, w_bar)
    lifetime_vals.append(lv)
```

Let's visualize the expected lifetime value as a function of volatility:

```
fig, ax = plt.subplots()
ax.plot(σ_vals, lifetime_vals, linewidth=2, color='green')
ax.set_xlabel(r'volatility ($\sigma$)', fontsize=12)
ax.set_ylabel('expected lifetime value', fontsize=12)
plt.show()
```



The plot confirms that despite workers setting higher reservation wages when facing more volatile wage offers (as shown above), they achieve higher expected lifetime values due to the option value of search.

54.7 Exercises

Exercise 54.7.1

Compute the average duration of unemployment when $\beta = 0.99$ and c takes the following values

```
c_vals = np.linspace(10, 40, 4)
```

That is, start the agent off as unemployed, compute their reservation wage given the parameters, and then simulate to see how long it takes to accept.

Repeat a large number of times and take the average.

Plot mean unemployment duration as a function of c in `c_vals`.

Try to explain what you see.

Solution

Here's a solution using the continuous wage offer distribution with JAX.

```
def compute_stopping_time_continuous(w_bar, key, model):
    """
    Compute stopping time by drawing wages from the continuous distribution
    until one exceeds `w_bar`.

    Parameters:
    -----------
    w_bar : float
        The reservation wage
    key : jax.random.key
        Random key for JAX
    model : McCallModelContinuous
        The model containing wage draws

    Returns:
    --------
    t_final : int
        The stopping time (number of periods until acceptance)
    """
    c, β, σ, μ, w_draws = model

    def update(loop_state):
        t, key, accept = loop_state
        key, subkey = jax.random.split(key)
        # Draw a standard normal and transform to wage
        s = jax.random.normal(subkey)
        w = jnp.exp(μ + σ * s)
        accept = w >= w_bar
        t = t + 1
        return t, key, accept

    def cond(loop_state):
        _, _, accept = loop_state
        return jnp.logical_not(accept)

    initial_loop_state = (0, key, False)
```

```
    t_final, _, _ = jax.lax.while_loop(cond, update, initial_loop_state)
    return t_final


def compute_mean_stopping_time_continuous(w_bar, model, num_reps=100000,␣
↪seed=1234):
    """
    Generate a mean stopping time over `num_reps` repetitions.

    Parameters:
    -----------
    w_bar : float
        The reservation wage
    model : McCallModelContinuous
        The model containing parameters
    num_reps : int
        Number of simulation replications
    seed : int
        Random seed

    Returns:
    --------
    mean_time : float
        Average stopping time across all replications
    """
    # Generate a key for each MC replication
    key = jax.random.key(seed)
    keys = jax.random.split(key, num_reps)

    # Vectorize compute_stopping_time_continuous and evaluate across keys
    compute_fn = jax.vmap(compute_stopping_time_continuous, in_axes=(None, 0,␣
↪None))
    obs = compute_fn(w_bar, keys, model)

    # Return mean stopping time
    return jnp.mean(obs)


# Compute mean stopping time for different values of c
c_vals = jnp.linspace(10, 40, 4)

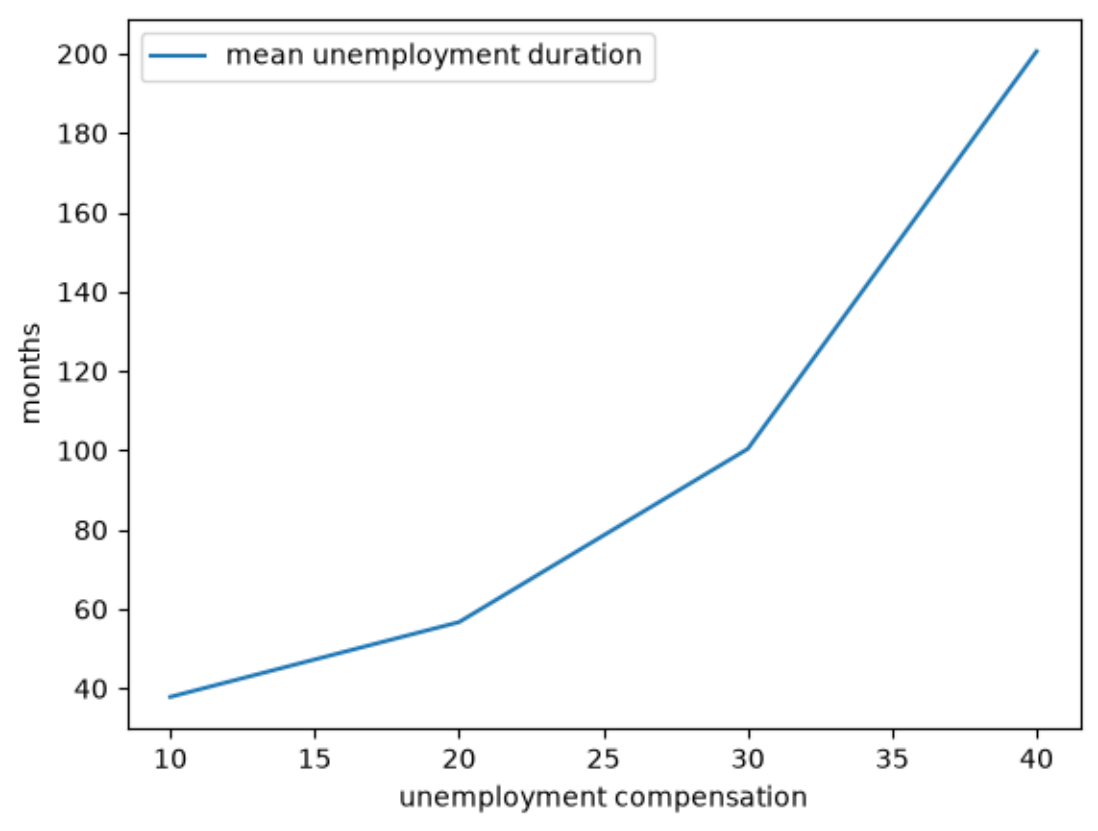
@jax.jit
def compute_stop_time_for_c_continuous(c):
    """Compute mean stopping time for a given compensation value c."""
    model = create_mccall_continuous(c=c)
    w_bar = compute_reservation_wage_continuous(model)
    return compute_mean_stopping_time_continuous(w_bar, model)

# Vectorize across all c values
compute_stop_time_vectorized = jax.vmap(compute_stop_time_for_c_continuous)
stop_times = compute_stop_time_vectorized(c_vals)

fig, ax = plt.subplots()

ax.plot(c_vals, stop_times, label="mean unemployment duration")
ax.set(xlabel="unemployment compensation", ylabel="months")
ax.legend()
```

```
plt.show()
```


mean unemployment duration
200
180
160
140
120
100
80
60
40
months
10
15
20
25
30
35
40
unemployment compensation

CHAPTER

FIFTYFIVE

JOB SEARCH II: SEARCH AND SEPARATION

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Job Search II: Search and Separation*
  - *Overview*
  - *The model*
  - *Solving the model*
  - *Code*
  - *A simplifying transformation*
  - *Implementation*
  - *Impact of parameters*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax myst-nb
```

55.1 Overview

Previously *we looked* at the McCall job search model [McCall, 1970] as a way of understanding unemployment and worker decisions.

One unrealistic feature of that version of the model was that every job is permanent.

In this lecture, we extend the model by introducing job separation.

Once separation enters the picture, the agent comes to view

- the loss of a job as a capital loss, and
- a spell of unemployment as an *investment* in searching for an acceptable job

The other minor addition is that a utility function will be included to make worker preferences slightly more sophisticated.

We'll need the following imports

```
import matplotlib.pyplot as plt
import numpy as np
import jax
import jax.numpy as jnp
from typing import NamedTuple
from quantecon.distributions import BetaBinomial
from myst_nb import glue
```

55.2 The model

The model is similar to the *baseline McCall job search model*.

It concerns the life of an infinitely lived worker and

- the opportunities he or she (let's say he to save one character) has to work at different wages
- exogenous events that destroy his current job
- his decision making process while unemployed

The worker can be in one of two states: employed or unemployed.

He wants to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(y_t) \tag{55.1}$$

At this stage the only difference from the *baseline model* is that we've added some flexibility to preferences by introducing a utility function u.

It satisfies $u' > 0$ and $u'' < 0$.

Wage offers $\{W_t\}$ are IID with common distribution q.

The set of possible wage values is denoted by $\mathbb{W}$.

55.2.1 Timing and decisions

At the start of each period, the agent can be either

- unemployed or
- employed at some existing wage level w.

If currently employed at wage w, the worker

1. receives utility $u(w)$ from their current wage and
2. is fired with some (small) probability α, becoming unemployed next period.

If currently unemployed, the worker receives random wage offer W_t and either accepts or rejects.

If he accepts, then he begins work immediately at wage W_t.

If he rejects, then he receives unemployment compensation c.

The process then repeats.

Note

We do not allow for job search while employed—this topic is taken up in a *later lecture*.

55.3 Solving the model

We drop time subscripts in what follows and primes denote next period values.

Let

- $v_e(w)$ be maximum lifetime value for a worker who enters the current period employed with wage w
- $v_u(w)$ be maximum lifetime for a worker who who enters the current period unemployed and receives wage offer w.

Here, **maximum lifetime value** means the value of (55.1) when the worker makes optimal decisions at all future points in time.

As we now show, obtaining these functions is key to solving the model.

55.3.1 The Bellman equations

We recall that, in *the original job search model*, the value function (the value of being unemployed with a given wage offer) satisfied a Bellman equation.

Here this function again satisfies a Bellman equation that looks very similar.

$$v_u(w) = \max \left\{ v_e(w),\ u(c) + \beta \sum_{w' \in \mathbb{W}} v_u(w') q(w') \right\} \tag{55.2}$$

The difference is that the value of accepting is $v_e(w)$ rather than $w/(1-\beta)$.

We have to make this change because jobs are not permanent.

Accepting transitions the worker to employment and hence yields reward $v_e(w)$, which we discuss below.

Rejecting leads to unemployment compensation and unemployment tomorrow.

Equation (55.2) expresses the value of being unemployed with offer w in hand as a maximum over the value of two options: accept or reject the current offer.

The function v_e also satisfies a Bellman equation:

$$v_e(w) = u(w) + \beta \left[(1-\alpha) v_e(w) + \alpha \sum_{w' \in \mathbb{W}} v_u(w') q(w') \right] \tag{55.3}$$

Note

This equation differs from a traditional Bellman equation because there is no max.

There is no max because an employed agent has no choices.

Nonetheless, in keeping with most of the literature, we also refer to it as a Bellman equation.

Equation (55.3) expresses the value of being employed at wage w in terms of

- current reward $u(w)$ plus
- discounted expected reward tomorrow, given the α probability of being fired

As we will see, equations (55.3) and (55.2) provide enough information to solve for both v_e and v_u.

Once we have them in hand, we will be able to make optimal choices.

55.3.2 The reservation wage

Let

$$h := u(c) + \beta \sum_{w' \in \mathbb{W}} v_u(w') q(w') \tag{55.4}$$

This is the **continuation value** for an unemployed agent – the value of rejecting the current offer and then making optimal choices.

From (55.2), we see that an unemployed agent accepts current offer w if $v_e(w) \geq h$.

This means precisely that the value of accepting is higher than the value of rejecting.

The function v_e is increasing in w, since an employed agent is never made worse off by a higher current wage.

Hence, we can express the optimal choice as accepting wage offer w if and only if $w \geq \bar{w}$, where the **reservation wage** $\bar{w}$ is the first wage level $w \in \mathbb{W}$ such that

$$v_e(w) \geq h$$

55.4 Code

Let's now implement a solution method based on the two Bellman equations (55.2) and (55.3).

55.4.1 Set up

The default utility function is a CRRA utility function

```
def u(x, γ):
    return (x**(1 - γ) - 1) / (1 - γ)
```

Also, here's a default wage distribution, based around the BetaBinomial distribution:

```
n = 60                                  # n possible outcomes for w
w_default = jnp.linspace(10, 20, n)     # wages between 10 and 20
a, b = 600, 400                         # shape parameters
dist = BetaBinomial(n-1, a, b)          # distribution
q_default = jnp.array(dist.pdf())       # probabilities as a JAX array
```

Here's our model class for the McCall model with separation.

```
class Model(NamedTuple):
    α: float = 0.2              # job separation rate
    β: float = 0.98             # discount factor
    γ: float = 2.0              # utility parameter (CRRA)
    c: float = 6.0              # unemployment compensation
    w: jnp.ndarray = w_default  # wage outcome space
    q: jnp.ndarray = q_default  # probabilities over wage offers
```

55.4.2 Operators

We'll use a similar iterative approach to solving the Bellman equations that we adopted in the *first job search lecture*.

As a first step, to iterate on the Bellman equations, we define two operators, one for each value function.

These operators take the current value functions as inputs and return updated versions.

```
def T_u(model, v_u, v_e):
    """
    Apply the unemployment Bellman update rule and return new guess of v_u.

    """
    α, β, γ, c, w, q = model
    h = u(c, γ) + β * (v_u @ q)
    v_u_new = jnp.maximum(v_e, h)
    return v_u_new
```

```
def T_e(model, v_u, v_e):
    """
    Apply the employment Bellman update rule and return new guess of v_e.

    """
    α, β, γ, c, w, q = model
    v_e_new = u(w, γ) + β * ((1 - α) * v_e + α * (v_u @ q))
    return v_e_new
```

55.4.3 Iteration

Now we write an iteration routine, which updates the pair of arrays v_u, v_e until convergence.

More precisely, we iterate until successive realizations are closer together than some small tolerance level.

```
def solve_full_model(
        model,
        tol: float = 1e-6,
        max_iter: int = 1_000,
    ):
    """
    Solves for both value functions v_u and v_e iteratively.

    """
    α, β, γ, c, w, q = model
    i = 0
    error = tol + 1
    v_e = v_u = w / (1 - β)

    while i < max_iter and error > tol:
        v_u_next = T_u(model, v_u, v_e)
        v_e_next = T_e(model, v_u, v_e)
        error_u = jnp.max(jnp.abs(v_u_next - v_u))
        error_e = jnp.max(jnp.abs(v_e_next - v_e))
        error = jnp.max(jnp.array([error_u, error_e]))
        v_u = v_u_next
        v_e = v_e_next
        i += 1

    return v_u, v_e
```

55.4.4 Computing the reservation wage

Now that we can solve for both value functions, let's investigate the reservation wage.

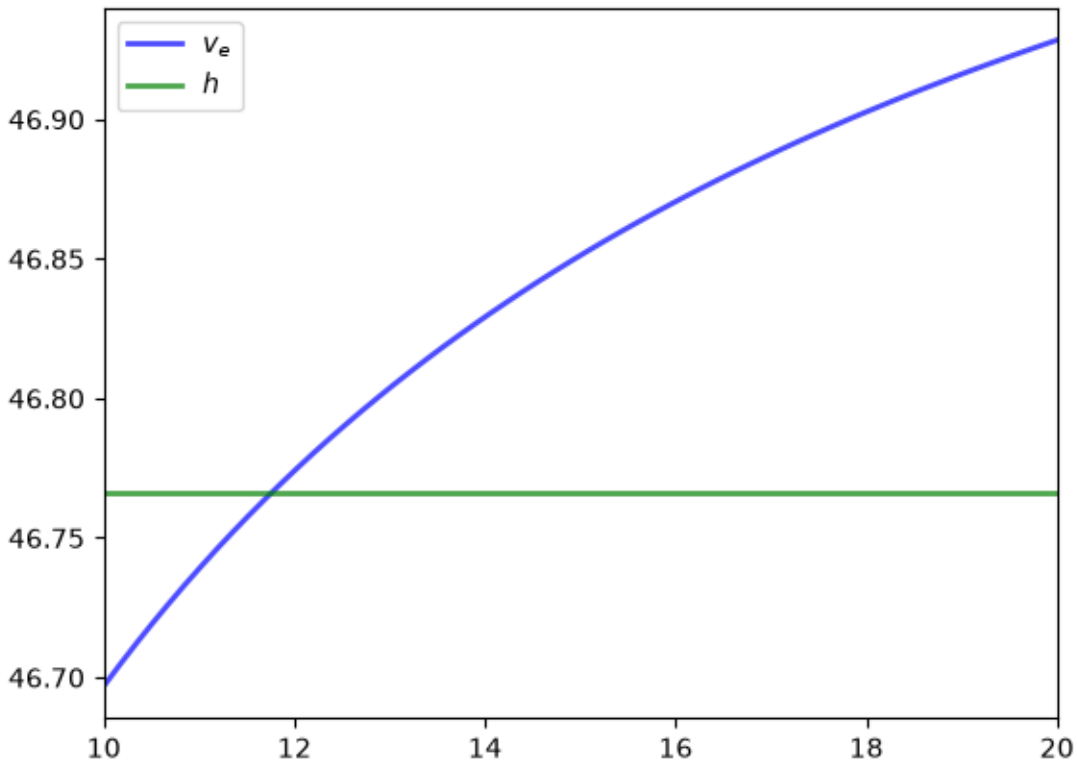
Recall from above that the reservation wage $\bar{w}$ is the first $w \in \mathbb{W}$ satisfying $v_e(w) \geq h$, where h is the continuation value defined in (55.4).

Let's compare v_e and h to see what they look like.

We'll use the default parameterizations found in the code above.

```
model = Model()
α, β, γ, c, w, q = model
v_u, v_e = solve_full_model(model)
h = u(c, γ) + β * (v_u @ q)

fig, ax = plt.subplots()
ax.plot(w, v_e, 'b-', lw=2, alpha=0.7, label='$v_e$')
ax.plot(w, [h] * len(w), 'g-', lw=2, alpha=0.7, label='$h$')
ax.set_xlim(min(w), max(w))
ax.legend()
plt.show()
```



The value v_e is increasing because higher w generates a higher wage flow conditional on staying employed.

The reservation wage is the w where these lines meet.

Let's compute this reservation wage explicitly:

```
def compute_reservation_wage_full(model):
    """
    Computes the reservation wage using the full model solution.
    """
    α, β, γ, c, w, q = model
    v_u, v_e = solve_full_model(model)
    h = u(c, γ) + β * (v_u @ q)
    # Find the first w such that v_e(w) >= h, or +inf if none exist
    accept = v_e >= h
    i = jnp.argmax(accept)  # returns first accept index
    w_bar = jnp.where(jnp.any(accept), w[i], jnp.inf)
    return w_bar

w_bar_full = compute_reservation_wage_full(model)
print(f"Reservation wage (full model): {w_bar_full:.4f}")
```

```
Reservation wage (full model): 11.8644
```

This value seems close to where the two lines meet.

55.5 A simplifying transformation

The approach above works, but iterating over two vector-valued functions is computationally expensive.

With some mathematics and some brain power, we can form a solution method that is far more efficient.

(This process will be analogous to our *second pass* at the plain vanilla McCall model, where we reduced the Bellman equation to an equation in an unknown scalar value, rather than an unknown vector.)

First, we use the continuation value h, as defined in (55.4), to write (55.2) as

$$v_u(w) = \max\{v_e(w),\, h\}$$

Taking the expectation of both sides and then discounting, this becomes

$$\beta \sum_{w'} v_u(w') q(w') = \beta \sum_{w'} \max\{v_e(w'),\, h\}\, q(w')$$

Adding $u(c)$ to both sides and using (55.4) again gives

$$h = u(c) + \beta \sum_{w'} \max\{v_e(w'),\, h\}\, q(w') \tag{55.5}$$

This is a nice scalar equation in the continuation value, which is already useful.

But we can go further, but eliminating v_e from the above equation.

55.5.1 Simplifying to a single equation

As a first step, we rearrange the expression defining h (see (55.4)) to obtain

$$\sum_{w'} v_u(w') q(w') = \frac{h - u(c)}{\beta}$$

Using this, the Bellman equation for v_e, as given in (55.3), can now be rewritten as

$$v_e(w) = u(w) + \beta \left[(1-\alpha) v_e(w) + \alpha \frac{h - u(c)}{\beta} \right] \tag{55.6}$$

Our next step is to solve (55.6) for v_e as a function of h.

Rearranging (55.6) gives

$$v_e(w) = u(w) + \beta(1-\alpha) v_e(w) + \alpha(h - u(c))$$

or

$$v_e(w) - \beta(1-\alpha) v_e(w) = u(w) + \alpha(h - u(c))$$

Solving for $v_e(w)$ gives

$$v_e(w) = \frac{u(w) + \alpha(h - u(c))}{1 - \beta(1-\alpha)} \tag{55.7}$$

Substituting this into (55.5) yields

$$h = u(c) + \beta \sum_{w' \in \mathbb{W}} \max\left\{ \frac{u(w') + \alpha(h - u(c))}{1 - \beta(1-\alpha)},\, h \right\} q(w') \tag{55.8}$$

Finally we have a single scalar equation in h!

If we can solve this for h, we can easily recover v_e using (55.7).

Then we have enough information to compute the reservation wage.

55.5.2 Solving the Bellman equations

To solve (55.8), we use the iteration rule

$$h_{n+1} = u(c) + \beta \sum_{w' \in \mathbb{W}} \max \left\{ \frac{u(w') + \alpha(h_n - u(c))}{1 - \beta(1-\alpha)}, h_n \right\} q(w') \tag{55.9}$$

starting from some initial condition h_0.

(It is possible to prove that (55.9) converges via the Banach contraction mapping theorem.)

55.6 Implementation

To implement iteration on h, we provide a function that provides one update, from h_n to h_{n+1}

```
def update_h(model, h):
    " One update of the scalar h. "
    α, β, γ, c, w, q = model
    v_e = compute_v_e(model, h)
    h_new = u(c, γ) + β * (jnp.maximum(v_e, h) @ q)
    return h_new
```

Also, we provide a function to compute v_e from (55.7).

```
def compute_v_e(model, h):
    " Compute v_e from h using the closed-form expression. "
    α, β, γ, c, w, q = model
    return (u(w, γ) + α * (h - u(c, γ))) / (1 - β * (1 - α))
```

This function will be applied once convergence is achieved.

Now we can write our model solver.

```
@jax.jit
def solve_model(model, tol=1e-5, max_iter=2000):
    " Iterates to convergence on the Bellman equations. "

    def cond(loop_state):
        h, i, error = loop_state
        return jnp.logical_and(error > tol, i < max_iter)

    def update(loop_state):
        h, i, error = loop_state
        h_new = update_h(model, h)
        error_new = jnp.abs(h_new - h)
        return h_new, i + 1, error_new

    # Initialize
    h_init = u(model.c, model.γ) / (1 - model.β)
    i_init = 0
    error_init = tol + 1
    init_state = (h_init, i_init, error_init)

    final_state = jax.lax.while_loop(cond, update, init_state)
    h_final, _, _ = final_state
```

(continues on next page)

(continued from previous page)

```
    # Compute v_e from the converged h
    v_e_final = compute_v_e(model, h_final)

    return v_e_final, h_final
```

Finally, here's a function `compute_reservation_wage` that uses all the logic above, taking an instance of `Model` and returning the associated reservation wage.

```
def compute_reservation_wage(model):
    """
    Computes the reservation wage of an instance of the McCall model
    by finding the smallest w such that v_e(w) >= h.

    """
    # Find the first i such that v_e(w_i) >= h and return w[i]
    # If no such w exists, then w_bar is set to np.inf
    v_e, h = solve_model(model)
    accept = v_e >= h
    i = jnp.argmax(accept)    # take first accept index
    w_bar = jnp.where(jnp.any(accept), model.w[i], jnp.inf)
    return w_bar
```

Let's verify that this simplified approach gives the same answer as the full model:

```
w_bar_simplified = compute_reservation_wage(model)
print(f"Reservation wage (simplified): {w_bar_simplified:.4f}")
print(f"Reservation wage (full model): {w_bar_full:.4f}")
print(f"Difference: {abs(w_bar_simplified - w_bar_full):.6f}")
```

```
Reservation wage (simplified): 11.8644
Reservation wage (full model): 11.8644
Difference: 0.000000
```

As we can see, both methods produce essentially the same reservation wage.

However, the simplified method is far more efficient.

Next we will investigate how the reservation wage varies with parameters.

55.7 Impact of parameters

In each instance below, we'll show you a figure and then ask you to reproduce it in the exercises.

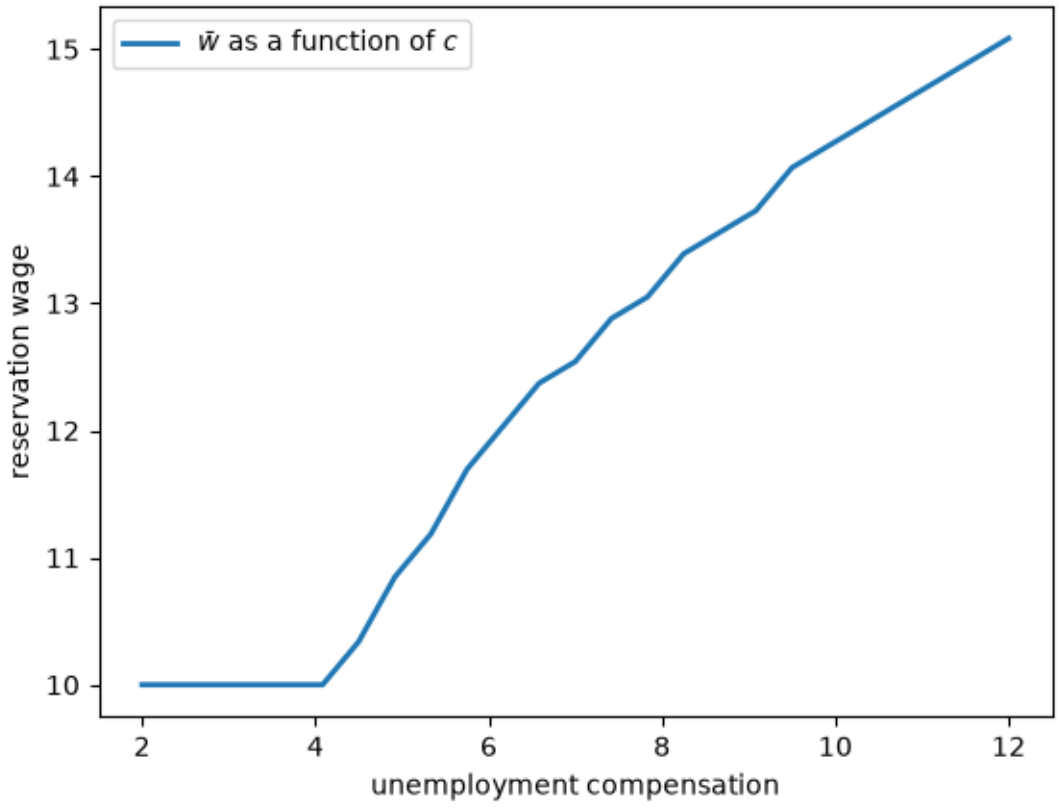
55.7.1 The reservation wage and unemployment compensation

First, let's look at how $\bar{w}$ varies with unemployment compensation.

In the figure below, we use the default parameters in the `Model` class, apart from c (which takes the values given on the horizontal axis)

As expected, higher unemployment compensation causes the worker to hold out for higher wages.

In effect, the cost of continuing job search is reduced.



55.7.2 The reservation wage and discounting

Next, let's investigate how $\bar{w}$ varies with the discount factor.

The next figure plots the reservation wage associated with different values of β

Again, the results are intuitive: More patient workers will hold out for higher wages.

55.7.3 The reservation wage and job destruction

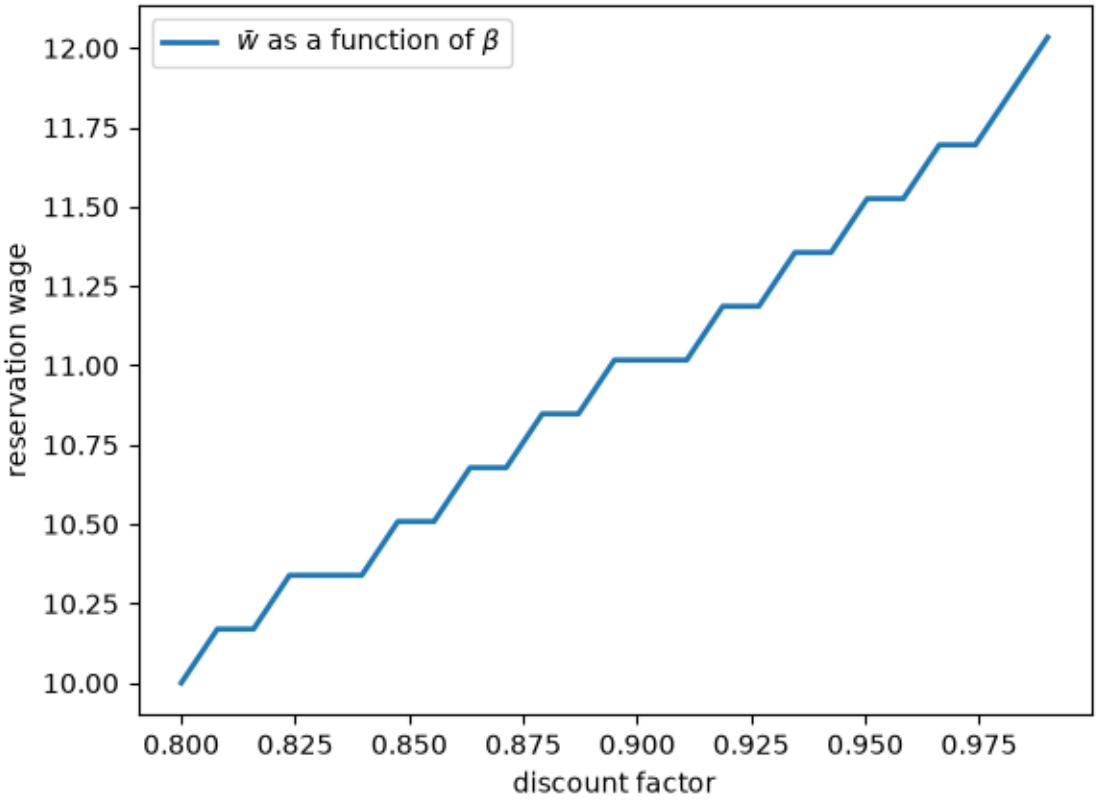
Finally, let's look at how $\bar{w}$ varies with the job separation rate α.

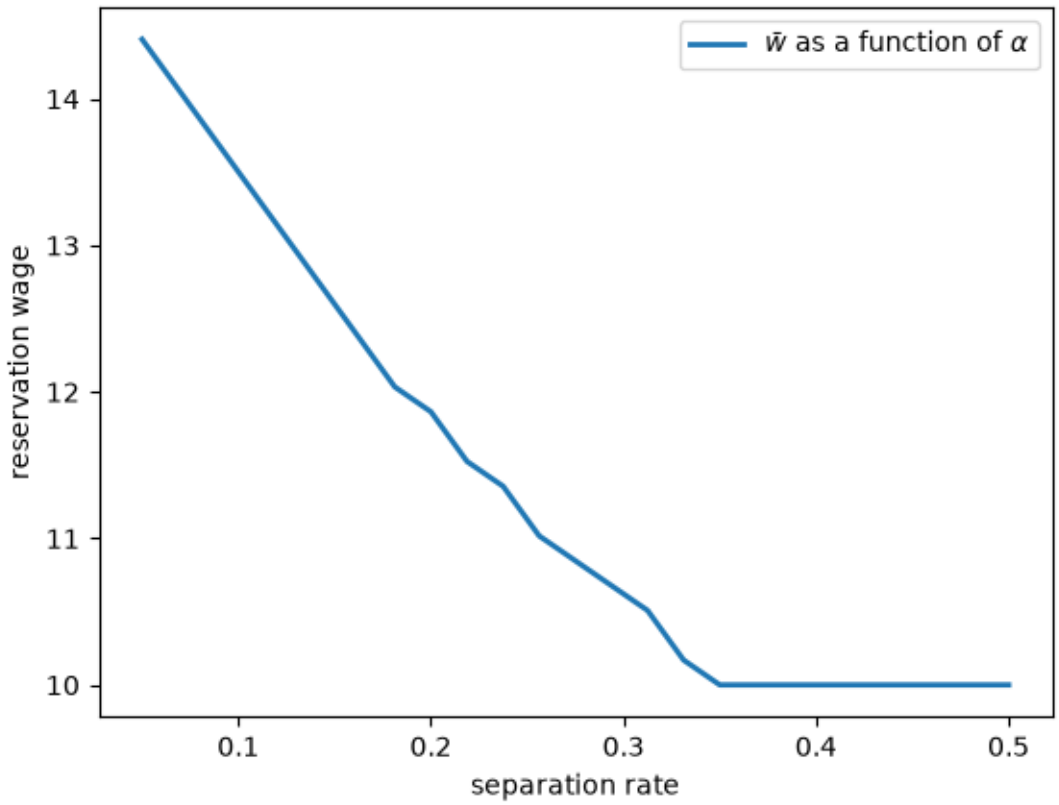
Higher α translates to a greater chance that a worker will face termination in each period once employed.

Once more, the results are in line with our intuition.

If the separation rate is high, then the benefit of holding out for a higher wage falls.

Hence the reservation wage is lower.


$\bar{w}$ as a function of β
12.00
11.75
11.50
11.25
11.00
10.75
10.50
10.25
10.00
reservation wage
0.800
0.825
0.850
0.875
0.900
0.925
0.950
0.975
discount factor


$\bar{w}$ as a function of α
14
13
12
11
10
reservation wage
0.1
0.2
0.3
0.4
0.5
separation rate

55.8 Exercises

Exercise 55.8.1

Reproduce all the reservation wage figures shown above.

Regarding the values on the horizontal axis, use

```
grid_size = 25
c_vals = jnp.linspace(2, 12, grid_size)          # unemployment compensation
β_vals = jnp.linspace(0.8, 0.99, grid_size)      # discount factors
α_vals = jnp.linspace(0.05, 0.5, grid_size)      # separation rate
```

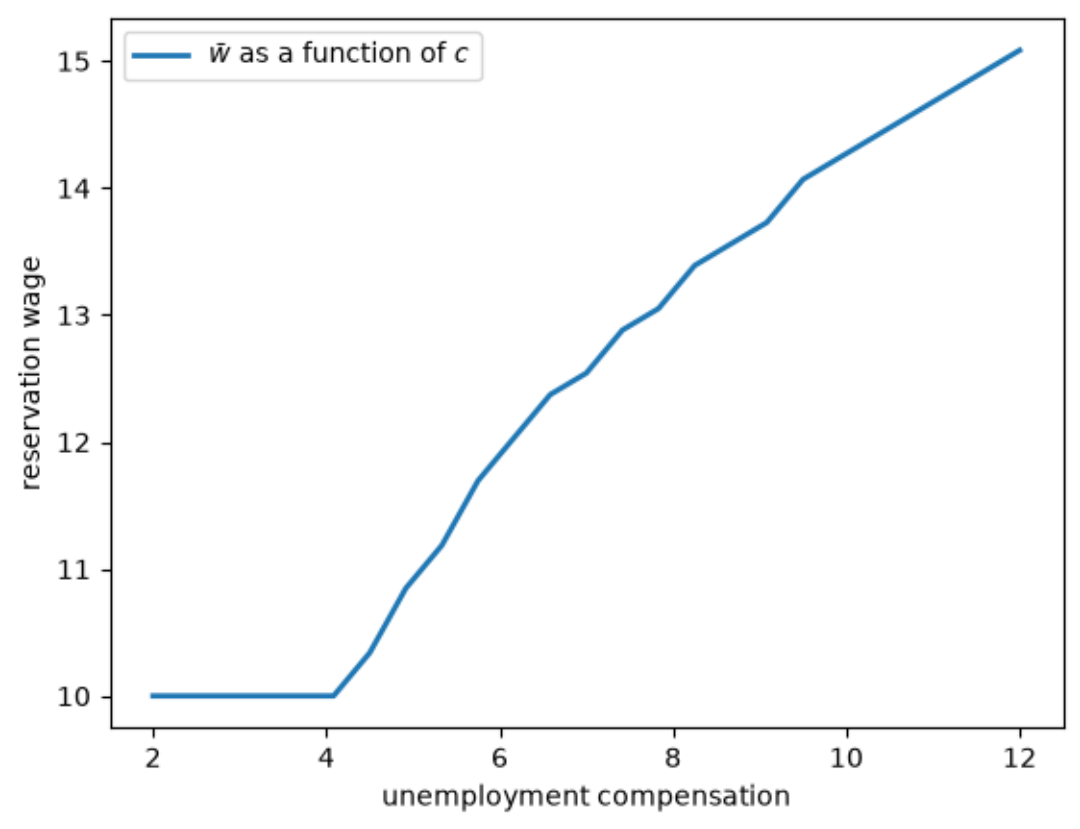
Solution

Here's the first figure.

```
def compute_res_wage_given_c(c):
    model = Model(c=c)
    w_bar = compute_reservation_wage(model)
    return w_bar

w_bar_vals = jax.vmap(compute_res_wage_given_c)(c_vals)

fig, ax = plt.subplots()
ax.set(xlabel='unemployment compensation', ylabel='reservation wage')
ax.plot(c_vals, w_bar_vals, lw=2, label=r'$\bar w$ as a function of $c$')
ax.legend()
glue("mccall_resw_c", fig, display=False)
plt.show()
```

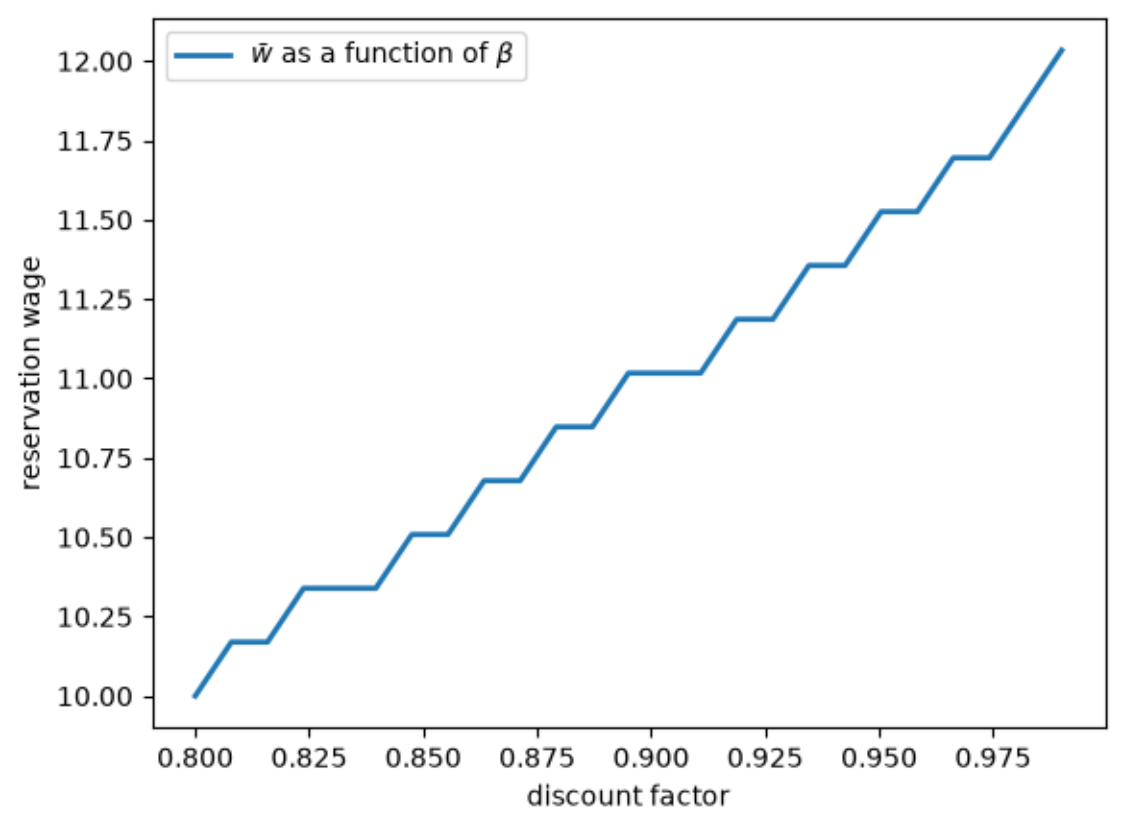


Here's the second one.

```
def compute_res_wage_given_beta(β):
    model = Model(β=β)
    w_bar = compute_reservation_wage(model)
    return w_bar

w_bar_vals = jax.vmap(compute_res_wage_given_beta)(β_vals)

fig, ax = plt.subplots()
ax.set(xlabel='discount factor', ylabel='reservation wage')
ax.plot(β_vals, w_bar_vals, lw=2, label=r'$\bar w$ as a function of $\beta$')
ax.legend()
glue("mccall_resw_beta", fig, display=False)
plt.show()
```

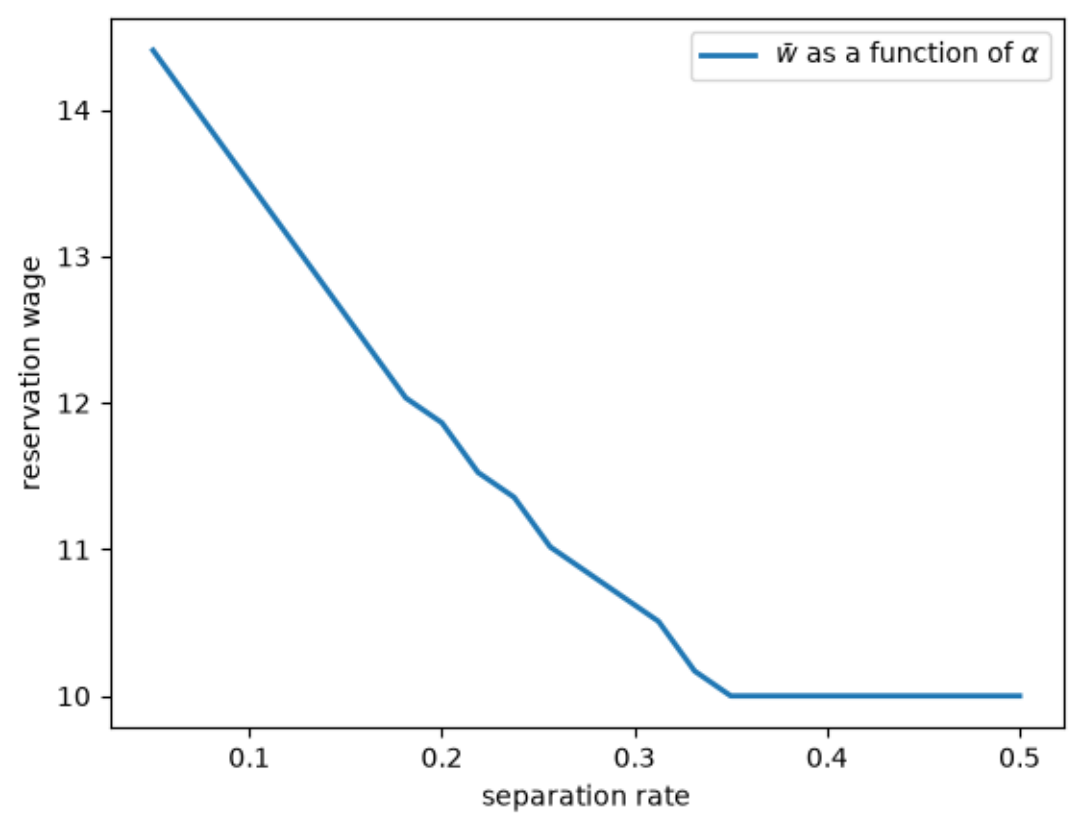


Here's the third.

```
def compute_res_wage_given_alpha(α):
    model = Model(α=α)
    w_bar = compute_reservation_wage(model)
    return w_bar

w_bar_vals = jax.vmap(compute_res_wage_given_alpha)(α_vals)

fig, ax = plt.subplots()
ax.set(xlabel='separation rate', ylabel='reservation wage')
ax.plot(α_vals, w_bar_vals, lw=2, label=r'$\bar w$ as a function of $\alpha$')
ax.legend()
glue("mccall_resw_alpha", fig, display=False)
plt.show()
```


$\bar{w}$ as a function of α
reservation wage
separation rate
14
13
12
11
10
0.1
0.2
0.3
0.4
0.5

CHAPTER

FIFTYSIX

JOB SEARCH III: SEARCH WITH SEPARATION AND MARKOV WAGES

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Job Search III: Search with Separation and Markov Wages*
  - *Model setup*
  - *Code*
  - *Improving efficiency*
  - *Sensitivity analysis*
  - *Employment simulation*
  - *Ergodic property*
  - *Cross-sectional analysis*
  - *Lower unemployment compensation (c=0.5)*
  - *Exercises*

This lecture builds on the job search model with separation presented in the *previous lecture*.

The key difference is that wage offers now follow a *Markov chain* rather than being IID.

This modification adds persistence to the wage offer process, meaning that today's wage offer provides information about tomorrow's offer.

This feature makes the model more realistic, as labor market conditions tend to exhibit serial correlation over time.

In addition to what's in Anaconda, this lecture will need the following libraries

```
!pip install quantecon jax
```

We use the following imports:

```
from quantecon.markov import tauchen
import jax.numpy as jnp
import jax
from jax import lax
from typing import NamedTuple
import matplotlib.pyplot as plt
from functools import partial
```

56.1 Model setup

The setting is as follows:

- Each unemployed agent receives a wage offer w from a finite set $\mathbb{W}$
- Wage offers follow a Markov chain with transition matrix P
- Jobs terminate with probability α each period (separation rate)
- Unemployed workers receive compensation c per period
- Future payoffs are discounted by factor $\beta \in (0, 1)$

56.1.1 Decision problem

When unemployed and receiving wage offer w, the agent chooses between:

1. Accept offer w: Become employed at wage w
2. Reject offer: Remain unemployed, receive c, get new offer next period

The wage updates are as follows:

- If an unemployed agent rejects offer w, then their next offer is drawn from $P(w, \cdot)$
- If an employed agent loses a job in which they were paid wage w, then their next offer is drawn from $P(w, \cdot)$

56.1.2 The wage offer process

To construct the wage offer process we start with an AR1 process.

$$X_{t+1} = \rho X_t + \nu Z_{t+1}$$

where $\{Z_t\}$ is IID and standard normal.

Below we will always choose $\rho \in (0, 1)$.

This means that the wage process will be positively correlated: the higher the current wage offer, the more likely we are to get a high offer tomorrow.

To go from the AR1 process to the wage offer process, we set $W_t = \exp(X_t)$.

Actually, in practice, we approximate this wage process as follows:

- discretize the AR1 process using *Tauchen's method* and

- take the exponential of the resulting wage offer values.

56.1.3 Value functions

We let

- $v_u(w)$ be the value of being unemployed when current wage offer is w
- $v_e(w)$ be the value of being employed at wage w

The Bellman equations are obvious modifications of the *IID case*.

The only change is that expectations for next period are computed using the transition matrix P conditioned on current wage w, instead of being drawn independently from q.

The unemployed worker's value function satisfies the Bellman equation

$$v_u(w) = \max \left\{ v_e(w), u(c) + \beta \sum_{w'} v_u(w') P(w, w') \right\}$$

The employed worker's value function satisfies the Bellman equation

$$v_e(w) = u(w) + \beta \left[\alpha \sum_{w'} v_u(w') P(w, w') + (1 - \alpha) v_e(w) \right]$$

As a matter of notation, given a function h assigning values to wages, it is common to set

$$(Ph)(w) = \sum_{w'} h(w') P(w, w')$$

(To understand this expression, think of P as a matrix, h as a column vector, and w as a row index.)

With this notation, the Bellman equations become

$$v_u(w) = \max\{v_e(w), u(c) + \beta (P v_u)(w)\}$$

and

$$v_e(w) = u(w) + \beta \left[\alpha (P v_u)(w) + (1 - \alpha) v_e(w)\right]$$

56.1.4 Optimal policy

Once we have the solutions v_e and v_u to these Bellman equations, we can compute the optimal policy: accept at current wage offer w if

$$v_e(w) \geq u(c) + \beta (P v_u)(w)$$

The optimal policy turns out to be a reservation wage strategy: accept all wages above some threshold.

56.2 Code

Let's now implement the model.

56.2.1 Set up

The default utility function is a CRRA utility function

```
def u(x, γ):
    return (x**(1 - γ) - 1) / (1 - γ)
```

Let's set up a `Model` class to store information needed to solve the model.

We include `P_cumsum`, the row-wise cumulative sum of the transition matrix, to optimize simulation – the details are explained below.

```
class Model(NamedTuple):
    n: int
    w_vals: jnp.ndarray
    P: jnp.ndarray
    P_cumsum: jnp.ndarray
    β: float
    c: float
    α: float
    γ: float
```

The next function holds default values and creates a `Model` instance:

```
def create_js_with_sep_model(
        n: int = 200,           # wage grid size
        ρ: float = 0.9,         # wage persistence
        ν: float = 0.2,         # wage volatility
        β: float = 0.96,        # discount factor
        α: float = 0.05,        # separation rate
        c: float = 1.0,         # unemployment compensation
        γ: float = 1.5          # utility parameter
    ) -> Model:
    """
    Creates an instance of the job search model with separation.

    """
    mc = tauchen(n, ρ, ν)
    w_vals, P = jnp.exp(jnp.array(mc.state_values)), jnp.array(mc.P)
    P_cumsum = jnp.cumsum(P, axis=1)
    return Model(n, w_vals, P, P_cumsum, β, c, α, γ)
```

56.2.2 Solution: first pass

Let's put together a (not very efficient) routine for calculating the reservation wage.

(We will think carefully about efficiency below.)

It works by starting with guesses for v_e and v_u and iterating to convergence.

Here are Bellman operators that update v_u and v_e respectively.

```
def T_u(model, v_u, v_e):
    """
    Apply the unemployment Bellman update rule and return new guess of v_u.

    """
```

(continues on next page)

(continued from previous page)

```
    n, w_vals, P, P_cumsum, β, c, α, γ = model
    h = u(c, γ) + β * P @ v_u
    v_u_new = jnp.maximum(v_e, h)
    return v_u_new
```

```
def T_e(model, v_u, v_e):
    """
    Apply the employment Bellman update rule and return new guess of v_e.

    """
    n, w_vals, P, P_cumsum, β, c, α, γ = model
    v_e_new = u(w_vals, γ) + β * ((1 - α) * v_e + α * P @ v_u)
    return v_e_new
```

Here's a routine to iterate to convergence and then compute the reservation wage.

```
def solve_model_first_pass(
        model: Model,           # instance containing default parameters
        v_u_init: jnp.ndarray,  # initial condition for v_u
        v_e_init: jnp.ndarray,  # initial condition for v_e
        tol: float=1e-6,        # error tolerance
        max_iter: int=1_000,    # maximum number of iterations for loop
    ):
    n, w_vals, P, P_cumsum, β, c, α, γ = model
    i = 0
    error = tol + 1
    v_u = v_u_init
    v_e = v_e_init

    while i < max_iter and error > tol:
        v_u_next = T_u(model, v_u, v_e)
        v_e_next = T_e(model, v_u, v_e)
        error_u = jnp.max(jnp.abs(v_u_next - v_u))
        error_e = jnp.max(jnp.abs(v_e_next - v_e))
        error = jnp.maximum(error_u, error_e)
        v_u = v_u_next
        v_e = v_e_next
        i += 1

    # Compute accept and reject values
    continuation_values = u(c, γ) + β * P @ v_u

    # Find where acceptance becomes optimal
    accept_indices = v_e >= continuation_values
    first_accept_idx = jnp.argmax(accept_indices)  # index of first True

    # If no acceptance (all False), return infinity
    # Otherwise return the wage at the first acceptance index
    w_bar = jnp.where(
        jnp.any(accept_indices), w_vals[first_accept_idx], jnp.inf
    )
    return v_u, v_e, w_bar
```

56.2.3 Road test

Let's solve the model:

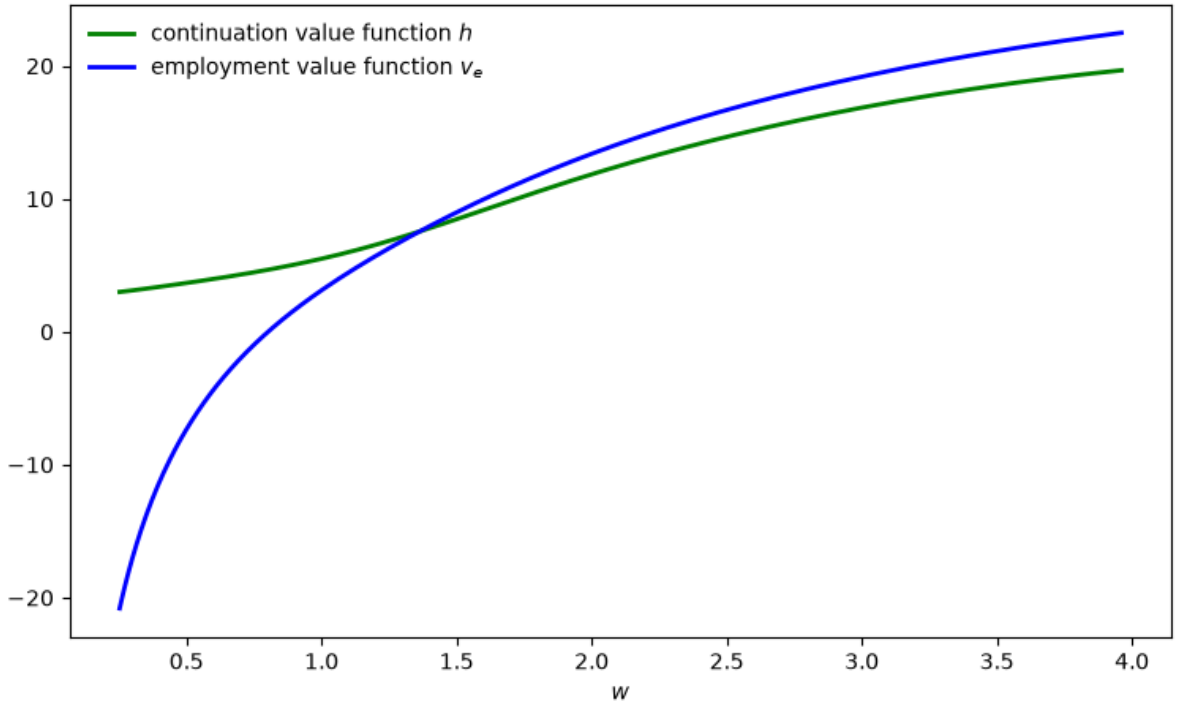
```
model = create_js_with_sep_model()
n, w_vals, P, P_cumsum, β, c, α, γ = model
v_u_init = jnp.zeros(n)
v_e_init = jnp.zeros(n)
v_u, v_e, w_bar_first = solve_model_first_pass(model, v_u_init, v_e_init)
```

Next we compute the continuation values.

```
h = u(c, γ) + β * P @ v_u
```

Let's plot our results.

```
fig, ax = plt.subplots(figsize=(9, 5.2))
ax.plot(w_vals, h, 'g-', linewidth=2,
        label="continuation value function $h$")
ax.plot(w_vals, v_e, 'b-', linewidth=2,
        label="employment value function $v_e$")
ax.legend(frameon=False)
ax.set_xlabel(r"$w$")
plt.show()
```



The reservation wage is at the intersection of v_e, and the continuation value function, which is the value of rejecting.

56.3 Improving efficiency

The solution method described above works fine but we can do much better.

First, we use the employed worker's Bellman equation to express v_e in terms of Pv_u

$$v_e(w) = \frac{1}{1-\beta(1-\alpha)} \cdot (u(w) + \alpha\beta(Pv_u)(w))$$

Next we substitute into the unemployed agent's Bellman equation to get

$$v_u(w) = \max\left\{\frac{1}{1-\beta(1-\alpha)} \cdot (u(w) + \alpha\beta(Pv_u)(w)), u(c) + \beta(Pv_u)(w)\right\}$$

Then we use value function iteration to solve for v_u.

With v_u in hand, we can recover v_e through the equations above and then compute the reservation wage.

Here's the new Bellman operator for the unemployed worker's value function:

```
def T(v: jnp.ndarray, model: Model) -> jnp.ndarray:
    """
    The Bellman operator for v_u.

    """
    n, w_vals, P, P_cumsum, β, c, α, γ = model
    d = 1 / (1 - β * (1 - α))
    v_e = d * (u(w_vals, γ) + α * β * P @ v)
    h = u(c, γ) + β * P @ v
    return jnp.maximum(v_e, h)
```

Here's a routine for value function iteration.

```
@jax.jit
def vfi(
        model: Model,
        tolerance: float = 1e-6,    # Error tolerance
        max_iter: int = 100_000,    # Max iteration bound
    ):

    v_init = jnp.zeros(model.w_vals.shape)

    def cond(loop_state):
        v, error, i = loop_state
        return (error > tolerance) & (i <= max_iter)

    def update(loop_state):
        v, error, i = loop_state
        v_new = T(v, model)
        error = jnp.max(jnp.abs(v_new - v))
        new_loop_state = v_new, error, i + 1
        return new_loop_state

    initial_state = (v_init, tolerance + 1, 1)
    final_loop_state = lax.while_loop(cond, update, initial_state)
    v_final, error, i = final_loop_state

    return v_final
```

Here is a routine that computes the reservation wage from the value function.

```
@jax.jit
def get_reservation_wage(v: jnp.ndarray, model: Model) -> float:
    """
    Calculate the reservation wage from the unemployed agents
    value function v := v_u.

    The reservation wage is the lowest wage w where accepting (v_e(w))
    is at least as good as rejecting (u(c) + β(Pv_u)(w)).

    """
    n, w_vals, P, P_cumsum, β, c, α, γ = model

    # Compute accept and reject values
    d = 1 / (1 - β * (1 - α))
    v_e = d * (u(w_vals, γ) + α * β * P @ v)
    continuation_values = u(c, γ) + β * P @ v

    # Find where acceptance becomes optimal
    accept_indices = v_e >= continuation_values
    first_accept_idx = jnp.argmax(accept_indices)  # index of first True

    # If no acceptance (all False), return infinity
    # Otherwise return the wage at the first acceptance index
    return jnp.where(jnp.any(accept_indices), w_vals[first_accept_idx], jnp.inf)
```

Let's solve the model using our new method:

```
model = create_js_with_sep_model()
n, w_vals, P, P_cumsum, β, c, α, γ = model
v_u = vfi(model)
w_bar = get_reservation_wage(v_u, model)
```

Let's verify that both methods produce the same reservation wage:

```
print(f"Reservation wage (first method):  {w_bar_first:.6f}")
print(f"Reservation wage (second method): {w_bar:.6f}")
print(f"Difference: {abs(w_bar - w_bar_first):.2e}")
```

```
Reservation wage (first method):  1.365155
Reservation wage (second method): 1.365155
Difference: 0.00e+00
```

Next we compute some related quantities for plotting.

```
d = 1 / (1 - β * (1 - α))
v_e = d * (u(w_vals, γ) + α * β * P @ v_u)
h = u(c, γ) + β * P @ v_u
```

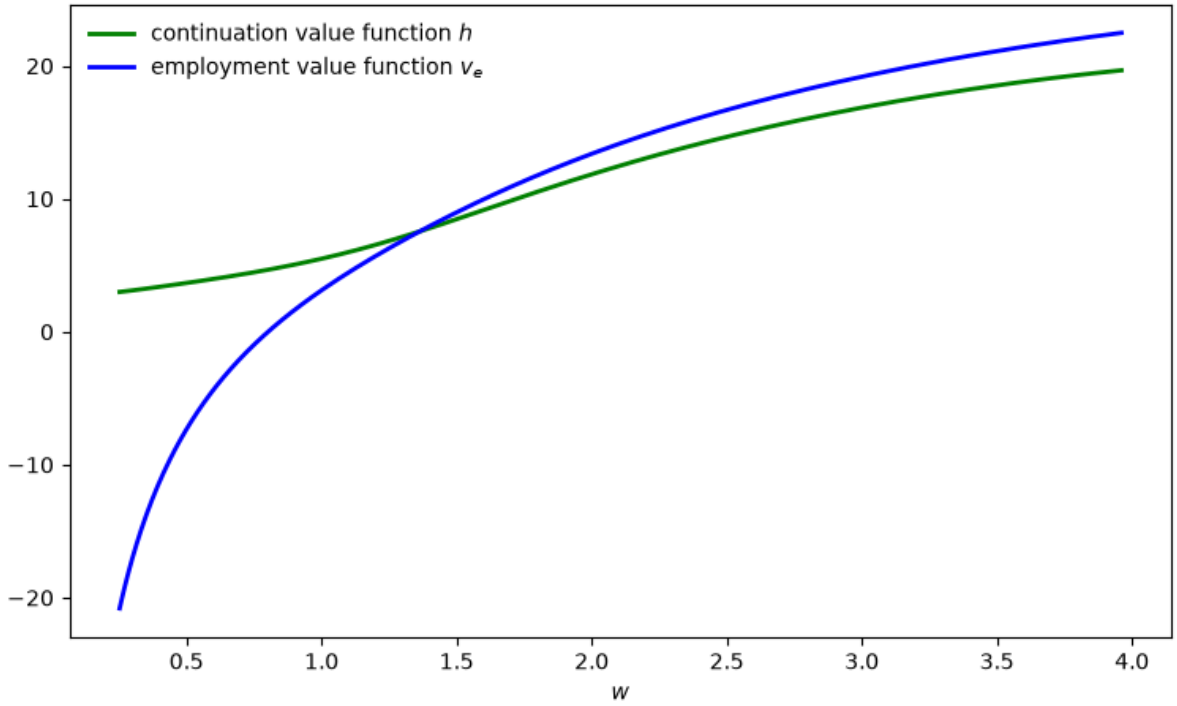
Let's plot our results.

```
fig, ax = plt.subplots(figsize=(9, 5.2))
ax.plot(w_vals, h, 'g-', linewidth=2,
        label="continuation value function $h$")
ax.plot(w_vals, v_e, 'b-', linewidth=2,
        label="employment value function $v_e$")
ax.legend(frameon=False)
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel(r"$w$")
plt.show()
```



The result is the same as before but we only iterate on one array — and also our JAX code is more efficient.

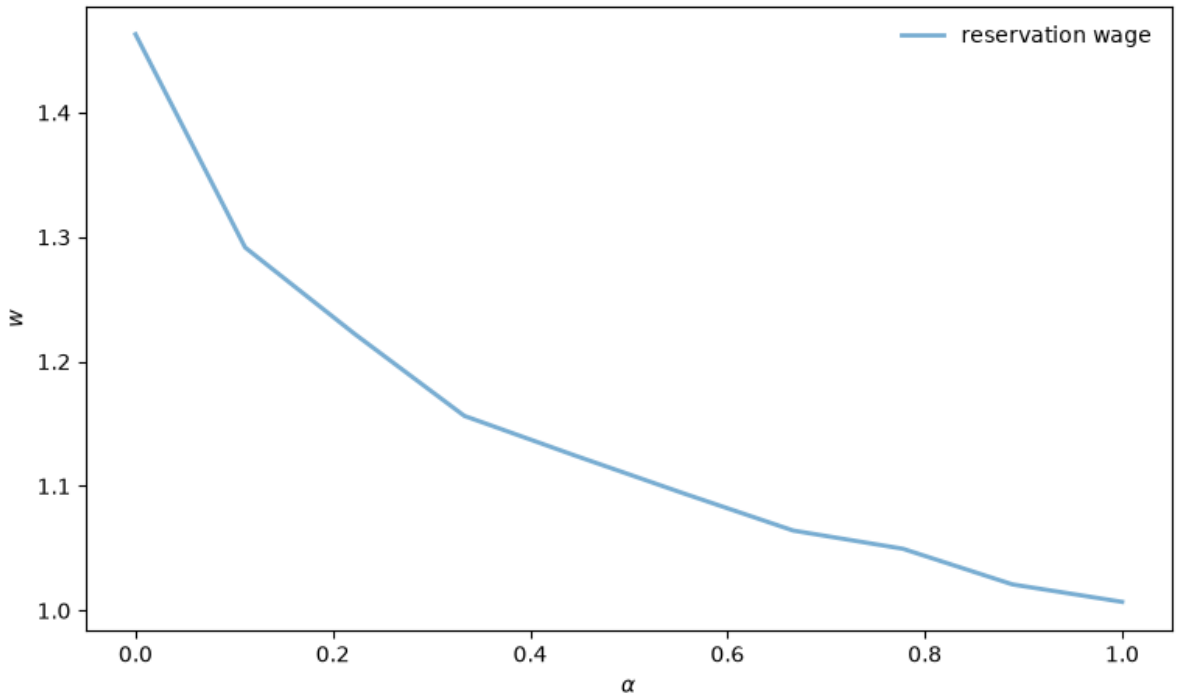
56.4 Sensitivity analysis

Let's examine how reservation wages change with the separation rate.

```
α_vals: jnp.ndarray = jnp.linspace(0.0, 1.0, 10)

w_bar_vec = []
for α in α_vals:
    model = create_js_with_sep_model(α=α)
    v_u = vfi(model)
    w_bar = get_reservation_wage(v_u, model)
    w_bar_vec.append(w_bar)

fig, ax = plt.subplots(figsize=(9, 5.2))
ax.plot(
    α_vals, w_bar_vec, linewidth=2, alpha=0.6, label="reservation wage"
)
ax.legend(frameon=False)
ax.set_xlabel(r"$\alpha$")
ax.set_ylabel(r"$w$")
plt.show()
```



Can you provide an intuitive economic story behind the outcome that you see in this figure?

56.5 Employment simulation

Now let's simulate the employment dynamics of a single agent under the optimal policy.

Note that, when simulating the Markov chain for wage offers, we need to draw from the distribution in each row of P many times.

To do this, we use the inverse transform method: draw a uniform random variable and find where it falls in the cumulative distribution.

This is implemented via `jnp.searchsorted` on the precomputed cumulative sum `P_cumsum`, which is much faster than recomputing the cumulative sum each time.

The function `update_agent` advances the agent's state by one period.

The agent's state is a pair (S_t, W_t), where S_t is employment status (0 if unemployed, 1 if employed) and W_t is

- their current wage offer, if unemployed, or
- their current wage, if employed.

```
def update_agent(key, status, wage_idx, model, w_bar):
    """
    Updates an agent's employment status and current wage.

    Parameters:
    - key: JAX random key
    - status: Current employment status (0 or 1)
    - wage_idx: Current wage, recorded as an array index
    - model: Model instance
    - w_bar: Reservation wage
```

(continues on next page)

(continued from previous page)

```
    """
    n, w_vals, P, P_cumsum, β, c, α, γ = model

    key1, key2 = jax.random.split(key)
    # Use precomputed cumulative sum for efficient sampling
    # via the inverse transform method.
    new_wage_idx = jnp.searchsorted(
        P_cumsum[wage_idx, :], jax.random.uniform(key1)
    )
    separation_occurs = jax.random.uniform(key2) < α
    # Accept if current wage meets or exceeds reservation wage
    accepts = w_vals[wage_idx] >= w_bar

    # If employed: status = 1 if no separation, 0 if separation
    # If unemployed: status = 1 if accepts, 0 if rejects
    next_status = jnp.where(
        status,
        1 - separation_occurs.astype(jnp.int32),  # employed path
        accepts.astype(jnp.int32)                 # unemployed path
    )

    # If employed: wage = current if no separation, new if separation
    # If unemployed: wage = current if accepts, new if rejects
    next_wage = jnp.where(
        status,
        jnp.where(separation_occurs, new_wage_idx, wage_idx),  # employed path
        jnp.where(accepts, wage_idx, new_wage_idx)             # unemployed path
    )

    return next_status, next_wage
```

Here's a function to simulate the employment path of a single agent.

```
def simulate_employment_path(
        model: Model,      # Model details
        w_bar: float,     # Reservation wage
        T: int = 2_000,    # Simulation length
        seed: int = 42     # Set seed for simulation
    ):
    """
    Simulate employment path for T periods starting from unemployment.

    """
    key = jax.random.key(seed)
    # Unpack model
    n, w_vals, P, P_cumsum, β, c, α, γ = model

    # Initial conditions
    status = 0
    wage_idx = 0

    wage_path = []
    status_path = []

    for t in range(T):
        wage_path.append(w_vals[wage_idx])
```

(continues on next page)

(continued from previous page)

```
        status_path.append(status)

        key, subkey = jax.random.split(key)
        status, wage_idx = update_agent(
            subkey, status, wage_idx, model, w_bar
        )

    return jnp.array(wage_path), jnp.array(status_path)
```

Let's create a comprehensive plot of the employment simulation:

```
model = create_js_with_sep_model()

# Calculate reservation wage for plotting
v_u = vfi(model)
w_bar = get_reservation_wage(v_u, model)

wage_path, employment_status = simulate_employment_path(model, w_bar)

fig, (ax1, ax2, ax3) = plt.subplots(3, 1, figsize=(8, 6))

# Plot employment status
ax1.plot(employment_status, 'b-', alpha=0.7, linewidth=1)
ax1.fill_between(
    range(len(employment_status)), employment_status, alpha=0.3, color='blue'
)
ax1.set_ylabel('employment status')
ax1.set_title('Employment path (0=unemployed, 1=employed)')
ax1.set_xticks((0, 1))
ax1.set_ylim(-0.1, 1.1)

# Plot wage path with employment status coloring
ax2.plot(wage_path, 'b-', alpha=0.7, linewidth=1)
ax2.axhline(y=w_bar, color='black', linestyle='--', alpha=0.8,
           label=f'Reservation wage: {w_bar:.2f}')
ax2.set_xlabel('time')
ax2.set_ylabel('wage')
ax2.set_title('Wage path (actual and offers)')
ax2.legend()

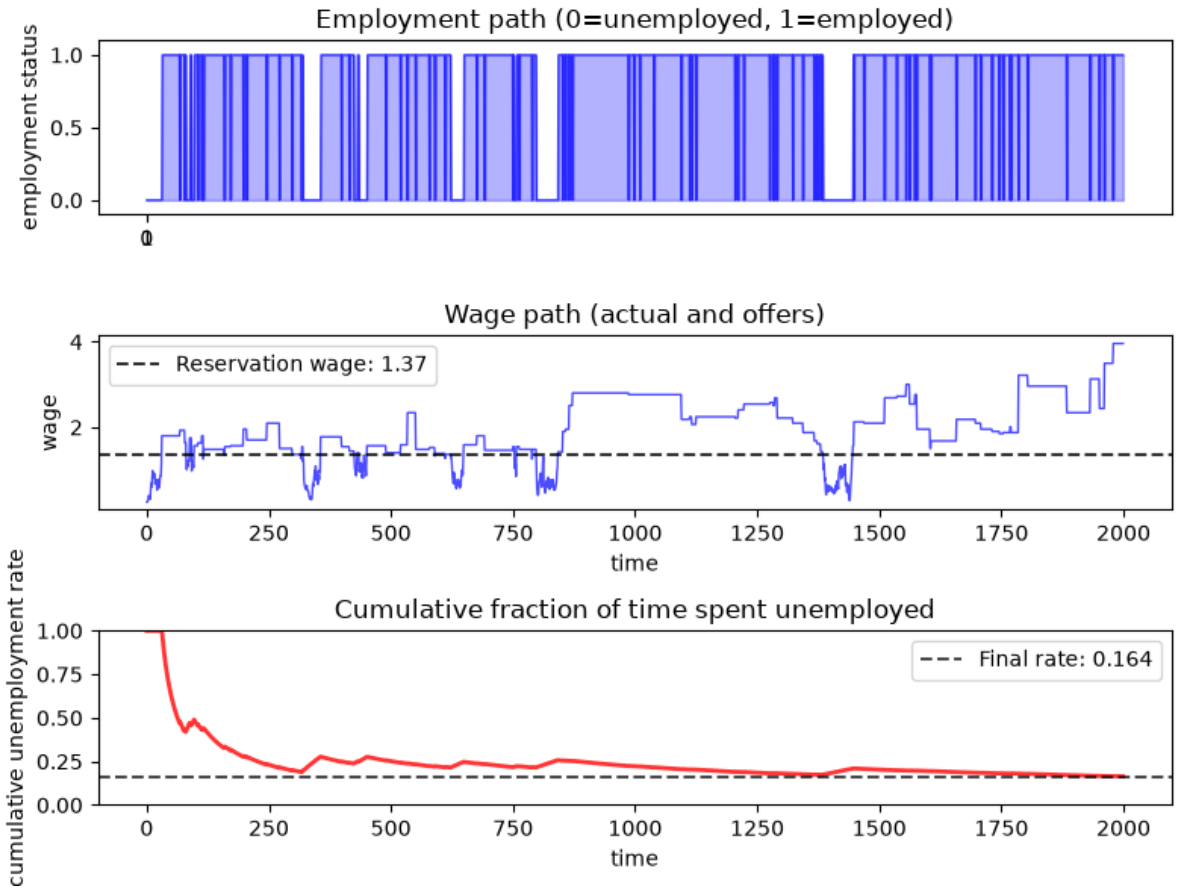
# Plot cumulative fraction of time unemployed
unemployed_indicator = (employment_status == 0).astype(int)
cumulative_unemployment = (
    jnp.cumsum(unemployed_indicator) /
    jnp.arange(1, len(employment_status) + 1)
)

ax3.plot(cumulative_unemployment, 'r-', alpha=0.8, linewidth=2)
ax3.axhline(y=jnp.mean(unemployed_indicator), color='black',
            linestyle='--', alpha=0.7,
            label=f'Final rate: {jnp.mean(unemployed_indicator):.3f}')
ax3.set_xlabel('time')
ax3.set_ylabel('cumulative unemployment rate')
ax3.set_title('Cumulative fraction of time spent unemployed')
ax3.legend()
ax3.set_ylim(0, 1)
```

(continues on next page)

(continued from previous page)

```
plt.tight_layout()
plt.show()
```



The simulation helps to visualize outcomes associated with this model.

The agent follows a reservation wage strategy.

Often the agent loses her job and immediately takes another job at a different wage.

This is because she uses the wage w from her last job to draw a new wage offer via $P(w, \cdot)$, and positive correlation means that a high current w is often leads a high new draw.

56.6 Ergodic property

Below we examine cross-sectional unemployment.

In particular, we will look at the unemployment rate in a cross-sectional simulation and compare it to the time-average unemployment rate, which is the fraction of time an agent spends unemployed over a long time series.

We will see that these two values are approximately equal – in fact they are exactly equal in the limit.

The reason is that the process (S_t, W_t), where

- S_t is the employment status and
- W_t is the wage

is Markovian, since the next pair depends only on the current pair and iid randomness, and ergodic.

Ergodicity holds as a result of irreducibility.

Indeed, from any (status, wage) pair, an agent can eventually reach any other (status, wage) pair.

This holds because:

- Unemployed agents can become employed by accepting offers
- Employed agents can become unemployed through separation (probability α)
- The wage process can transition between all wage states (because P is itself irreducible)

These properties ensure the chain is ergodic with a unique stationary distribution π over states (s, w).

For an ergodic Markov chain, the ergodic theorem guarantees that time averages = cross-sectional averages.

In particular, the fraction of time a single agent spends unemployed (across all wage states) converges to the cross-sectional unemployment rate:

$$\lim_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathbb{1}\{S_t = \text{unemployed}\} = \sum_{w=1}^{n} \pi(\text{unemployed}, w)$$

This holds regardless of initial conditions – provided that we burn in the cross-sectional distribution (run it forward in time from a given initial cross section in order to remove the influence of that initial condition).

As a result, we can study steady-state unemployment either by:

- Following one agent for a long time (time average), or
- Observing many agents at a single point in time (cross-sectional average)

Often the second approach is better for our purposes, since it's easier to parallelize.

56.7 Cross-sectional analysis

Now let's simulate many agents simultaneously to examine the cross-sectional unemployment rate.

To do this efficiently, we need a different approach than `simulate_employment_path` defined above.

The key differences are:

- `simulate_employment_path` records the entire history (all T periods) for a single agent, which is useful for visualization but memory-intensive
- The new function `sim_agent` below only tracks and returns the final state, which is all we need for cross-sectional statistics
- `sim_agent` uses `lax.fori_loop` instead of a Python loop, making it JIT-compilable and suitable for vectorization across many agents

We first define a function that simulates a single agent forward T time steps:

```
@jax.jit
def sim_agent(key, initial_status, initial_wage_idx, model, w_bar, T):
    """
    Simulate a single agent forward T time steps using lax.fori_loop.

    Uses fold_in to generate a new key at each time step.

    Parameters:
```

(continues on next page)

(continued from previous page)

```
    - key: JAX random key for this agent
    - initial_status: Initial employment status (0 or 1)
    - initial_wage_idx: Initial wage index
    - model: Model instance
    - w_bar: Reservation wage
    - T: Number of time periods to simulate

    Returns:
    - final_status: Employment status after T periods
    - final_wage_idx: Wage index after T periods
    """
    def update(t, loop_state):
        status, wage_idx = loop_state
        step_key = jax.random.fold_in(key, t)
        status, wage_idx = update_agent(step_key, status, wage_idx, model, w_bar)
        return status, wage_idx

    initial_loop_state = (initial_status, initial_wage_idx)
    final_loop_state = lax.fori_loop(0, T, update, initial_loop_state)
    final_status, final_wage_idx = final_loop_state
    return final_status, final_wage_idx


# Create vectorized version of sim_agent to process multiple agents in parallel
sim_agents_vmap = jax.vmap(sim_agent, in_axes=(0, 0, 0, None, None, None))


def simulate_cross_section(
        model: Model,               # Model instance with parameters
        n_agents: int = 100_000,    # Number of agents to simulate
        T: int = 200,               # Length of burn-in
        seed: int = 42              # For reproducibility
    ) -> float:
    """
    Simulate cross-section of agents and return unemployment rate.

    This approach:
    1. Generates n_agents random keys
    2. Calls sim_agent for each agent (vectorized via vmap)
    3. Collects the final states to produce the cross-section

    Returns the cross-sectional unemployment rate.
    """
    key = jax.random.key(seed)

    # Solve for optimal reservation wage
    v_u = vfi(model)
    w_bar = get_reservation_wage(v_u, model)

    # Initialize arrays
    initial_wage_indices = jnp.zeros(n_agents, dtype=jnp.int32)
    initial_status_vec = jnp.zeros(n_agents, dtype=jnp.int32)

    # Generate n_agents random keys
    agent_keys = jax.random.split(key, n_agents)

    # Simulate each agent forward T steps (vectorized)
```

(continues on next page)

(continued from previous page)

```
    final_status, final_wage_idx = sim_agents_vmap(
        agent_keys, initial_status_vec, initial_wage_indices, model, w_bar, T
    )

    unemployment_rate = 1 - jnp.mean(final_status)
    return unemployment_rate
```

This function generates a histogram showing the distribution of employment status across many agents:

```
def plot_cross_sectional_unemployment(
        model: Model,
        t_snapshot: int = 200,     # Time of cross-sectional snapshot
        n_agents: int = 20_000     # Number of agents to simulate
    ):
    """
    Generate histogram of cross-sectional unemployment at a specific time.

    """
    # Get final employment state directly
    key = jax.random.key(42)
    v_u = vfi(model)
    w_bar = get_reservation_wage(v_u, model)

    # Initialize arrays
    initial_wage_indices = jnp.zeros(n_agents, dtype=jnp.int32)
    initial_status_vec = jnp.zeros(n_agents, dtype=jnp.int32)

    # Generate n_agents random keys
    agent_keys = jax.random.split(key, n_agents)

    # Simulate each agent forward T steps (vectorized)
    final_status, _ = sim_agents_vmap(
        agent_keys, initial_status_vec, initial_wage_indices, model, w_bar, t_snapshot
    )

    # Calculate unemployment rate
    unemployment_rate = 1 - jnp.mean(final_status)

    fig, ax = plt.subplots(figsize=(8, 5))

    # Plot histogram as density (bars sum to 1)
    weights = jnp.ones_like(final_status) / len(final_status)
    ax.hist(final_status, bins=[-0.5, 0.5, 1.5],
            alpha=0.7, color='blue', edgecolor='black',
            density=True, weights=weights)

    ax.set_xlabel('employment status (0=unemployed, 1=employed)')
    ax.set_ylabel('density')
    ax.set_title(f'Cross-sectional distribution at t={t_snapshot}, ' +
                 f'unemployment rate = {unemployment_rate:.3f}')
    ax.set_xticks([0, 1])

    plt.tight_layout()
    plt.show()
```

Now let's compare the time-average unemployment rate (from a single agent's long simulation) with the cross-sectional unemployment rate (from many agents at a single point in time).

We claimed above that these numbers will be approximately equal in large samples, due to ergodicity.

Let's see if that's true.

```
model = create_js_with_sep_model()
cross_sectional_unemp = simulate_cross_section(
    model, n_agents=20_000, T=200
)

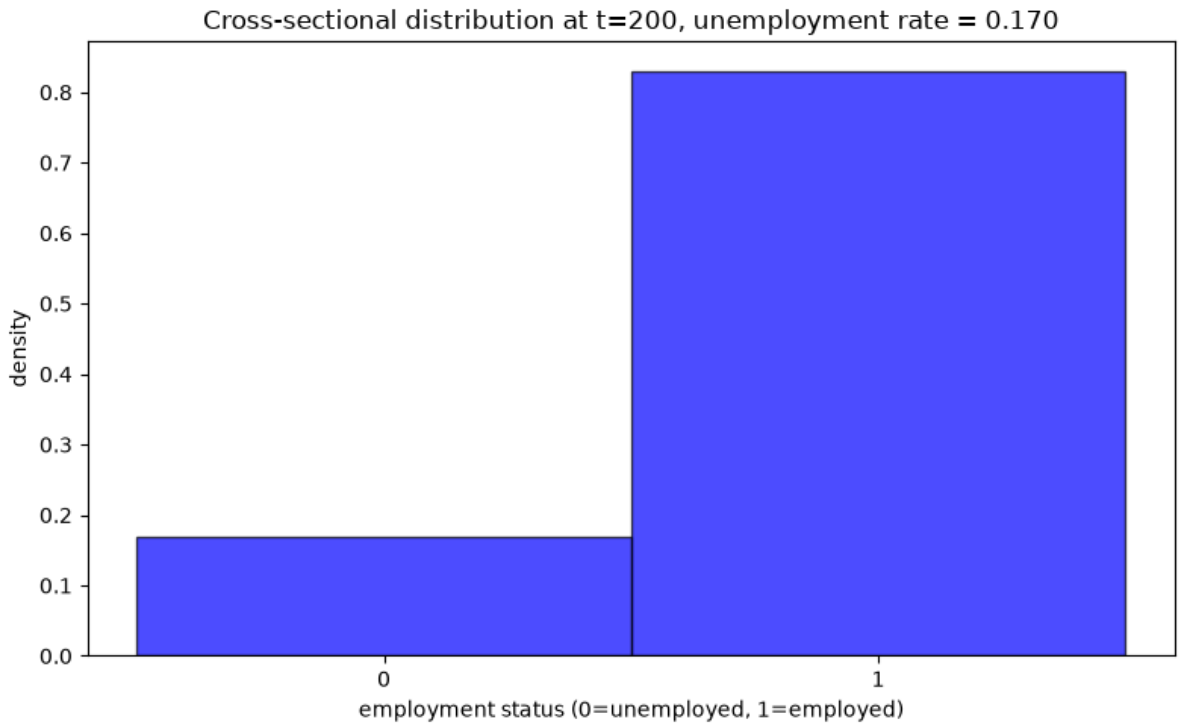
time_avg_unemp = jnp.mean(unemployed_indicator)
print(f"Time-average unemployment rate (single agent): "
      f"{time_avg_unemp:.4f}")
print(f"Cross-sectional unemployment rate (at t=200): "
      f"{cross_sectional_unemp:.4f}")
print(f"Difference: {abs(time_avg_unemp - cross_sectional_unemp):.4f}")
```

```
Time-average unemployment rate (single agent): 0.1640
Cross-sectional unemployment rate (at t=200): 0.1696
Difference: 0.0056
```

Indeed, they are very close.

Now let's visualize the cross-sectional distribution:

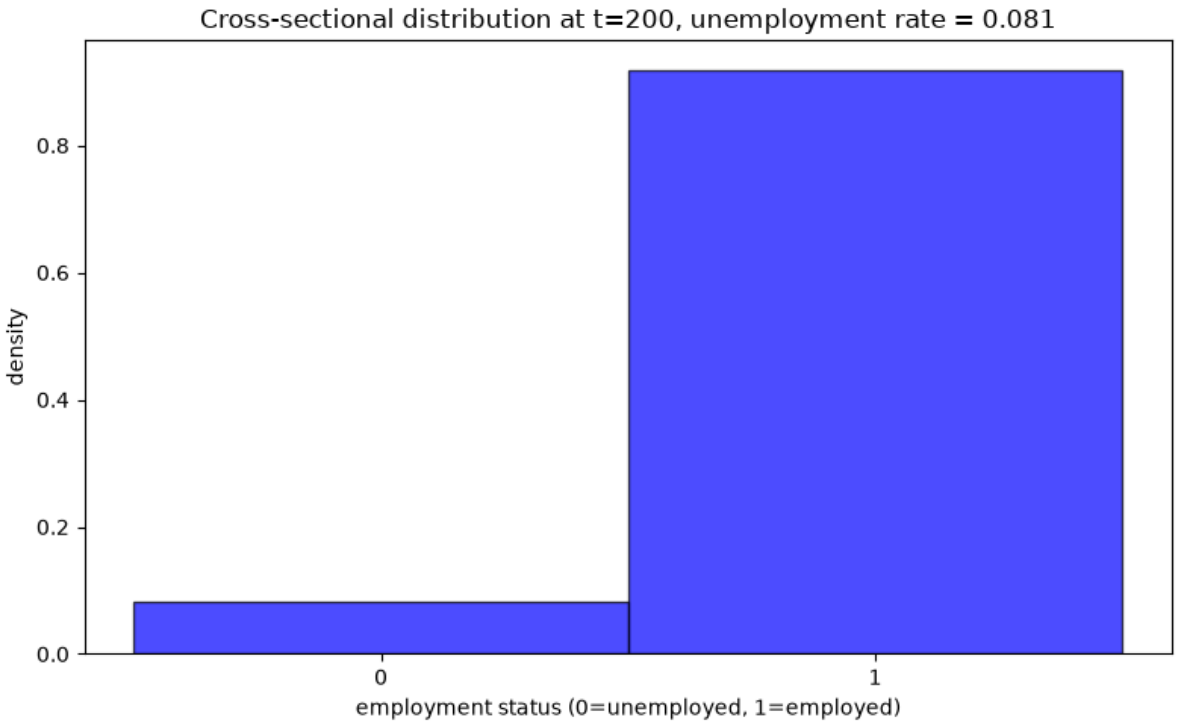
```
plot_cross_sectional_unemployment(model)
```



56.8 Lower unemployment compensation (c=0.5)

What happens to the cross-sectional unemployment rate with lower unemployment compensation?

```
model_low_c = create_js_with_sep_model(c=0.5)
plot_cross_sectional_unemployment(model_low_c)
```



56.9 Exercises

Exercise 56.9.1

Create a plot that investigates more carefully how the steady state cross-sectional unemployment rate changes with unemployment compensation.

Try a range of values for unemployment compensation `c`, such as `c = 0.2, 0.4, 0.6, 0.8, 1.0`. For each value, compute the steady-state cross-sectional unemployment rate and plot it against `c`.

What relationship do you observe between unemployment compensation and the unemployment rate?

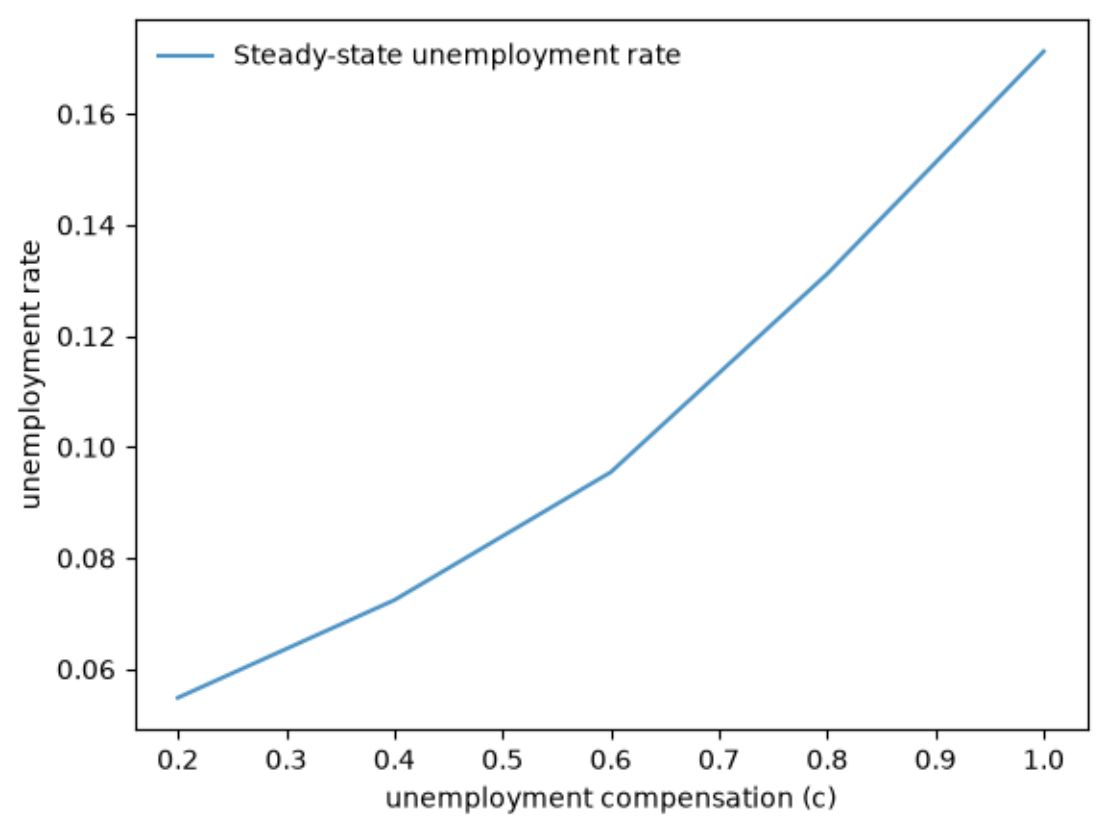
Solution

We compute the steady-state unemployment rate for different values of unemployment compensation:

```
c_values = 1.0, 0.8, 0.6, 0.4, 0.2
rates = []
for c in c_values:
    model = create_js_with_sep_model(c=c)
```

```
    unemployment_rate = simulate_cross_section(model)
    rates.append(unemployment_rate)

fig, ax = plt.subplots()
ax.plot(
    c_values, rates, alpha=0.8,
    linewidth=1.5, label='Steady-state unemployment rate'
)
ax.set_xlabel('unemployment compensation (c)')
ax.set_ylabel('unemployment rate')
ax.legend(frameon=False)
plt.show()
```



CHAPTER

FIFTYSEVEN

JOB SEARCH IV: FITTED VALUE FUNCTION ITERATION

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Job Search IV: Fitted Value Function Iteration*
  - *Overview*
  - *Model*
  - *Solution method*
  - *Implementation*
  - *Simulation*
  - *Exercises*

57.1 Overview

This lecture follows on from the job search model with separation presented in the *previous lecture*.

That lecture combined exogenous job separation events and a Markov wage offer process.

In this lecture we continue with this set and, in addition, allow the wage offer process to be continuous rather than discrete.

In particular,

$$W_t = \exp(X_t) \quad \text{where} \quad X_{t+1} = \rho X_t + \nu Z_{t+1}$$

and $\{Z_t\}$ is IID and standard normal.

While we already considered continuous wage distributions briefly in *Job Search I: The McCall Search Model*, the change was relatively trivial in that case.

The reason is that we were able to reduce the problem to solving for a single scalar value (the continuation value).

Here, in our Markov setting, the change is less trivial, since a continuous wage distribution leads to an uncountably infinite state space.

The infinite state space leads to additional challenges, particularly when it comes to applying value function iteration (VFI).

These challenges will lead us to modify VFI by adding an interpolation step.

The combination of VFI and this interpolation step is called **fitted value function iteration** (fitted VFI).

Fitted VFI is very common in practice, so we will take some time to work through the details.

In addition to what's in Anaconda, this lecture will need the following libraries

```
!pip install quantecon jax
```

We will use the following imports:

```
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
from jax import lax
from typing import NamedTuple
from functools import partial
import quantecon as qe
```

57.2 Model

Assuming that readers are familiar with the content of *Job Search III: Search with Separation and Markov Wages*, the model can be summarized as follows.

- Wage offers follow a continuous Markov process: $W_t = \exp(X_t)$ where $X_{t+1} = \rho X_t + \nu Z_{t+1}$
- $\{Z_t\}$ is IID and standard normal
- Jobs terminate with probability α each period (separation rate)
- Unemployed workers receive compensation c per period
- Workers have CRRA utility $u(x) = \frac{x^{1-\gamma}-1}{1-\gamma}$
- Future payoffs are discounted by factor $\beta \in (0, 1)$

57.3 Solution method

Let's discuss how we can solve this model.

The only real change from *Job Search III: Search with Separation and Markov Wages* is that we replace sums with integrals.

57.3.1 Value function iteration

In the *discrete case*, we ended up iterating on the Bellman operator

$$(Tv_u)(w) = \max\left\{\frac{1}{1-\beta(1-\alpha)} \cdot (u(w) + \alpha\beta(Pv_u)(w)), u(c) + \beta(Pv_u)(w)\right\} \tag{57.1}$$

where

$$(Pv_u)(w) := \sum_{w'} v_u(w')P(w, w')$$

Here we iterate on the same law after changing the definition of the P operator to

$$(Pv_u)(w) := \int v_u(w')p(w, w')dw'$$

where $p(w, \cdot)$ is the conditional density of w' given w.

Here we are thinking of v_u as a function on all of $\mathbb{R}_+$.

After taking ψ to be the standard normal density, we can write the expression above more explicitly as

$$(Pv_u)(w) := \int v_u(w^\rho \exp(\nu z))\psi(z)dz,$$

To understand this expression, recall that $W_t = \exp(X_t)$ where $X_{t+1} = \rho X_t + \nu Z_{t+1}$.

As a result $W_{t+1} = \exp(X_{t+1}) = \exp(\rho \log(W_t) + \nu Z_{t+1}) = W_t^\rho \exp(\nu Z_{t+1})$.

The integral above regards the current wage W_t as fixed at w and takes the expectation of $v_u(w^\rho \exp(\nu Z_{t+1}))$.

57.3.2 Fitting

In theory, we should now proceed as follows:

1. Begin with a guess v
2. Applying T to obtain the update $v' = Tv$
3. Unless some stopping condition is satisfied, set $v = v'$ and go to step 2.

However, there is a problem we must confront before we implement this procedure: The iterates of the value function can neither be calculated exactly nor stored on a computer.

To see the issue, consider (57.1).

Even if v is a known function, the only way to store its update v' is to record its value $v'(w)$ for every $w \in \mathbb{R}_+$.

Clearly, this is impossible.

57.3.3 Fitted value function iteration

What we will do instead is use **fitted value function iteration**.

The procedure is as follows:

Let a current guess v be given.

Now we record the value of the function v' at only finitely many "grid" points $w_1 < w_2 < \cdots < w_I$ and then reconstruct v' from this information when required.

More precisely, the algorithm will be

1. Begin with an array **v** representing the values of an initial guess of the value function on some grid points $\{w_i\}$.
2. Build a function v on the state space $\mathbb{R}_+$ by interpolation or approximation, based on **v** and $\{w_i\}$.
3. Obtain and record the samples of the updated function $v'(w_i)$ on each grid point w_i.
4. Unless some stopping condition is satisfied, take this as the new array and go to step 1.

How should we go about step 2?

This is a problem of function approximation, and there are many ways to approach it.

What's important here is that the function approximation scheme must not only produce a good approximation to each v, but also that it combines well with the broader iteration algorithm described above.

One good choice from both respects is continuous piecewise linear interpolation.

This method

1. combines well with value function iteration (see, e.g., [Gordon, 1995] or [Stachurski, 2008]) and
2. preserves useful shape properties such as monotonicity and concavity/convexity.

Linear interpolation will be implemented using JAX's interpolation function `jnp.interp`.

The next figure illustrates piecewise linear interpolation of an arbitrary function on grid points $0, 0.2, 0.4, 0.6, 0.8, 1$.

```
def f(x):
    y1 = 2 * jnp.cos(6 * x) + jnp.sin(14 * x)
    return y1 + 2.5

c_grid = jnp.linspace(0, 1, 6)
f_grid = jnp.linspace(0, 1, 150)

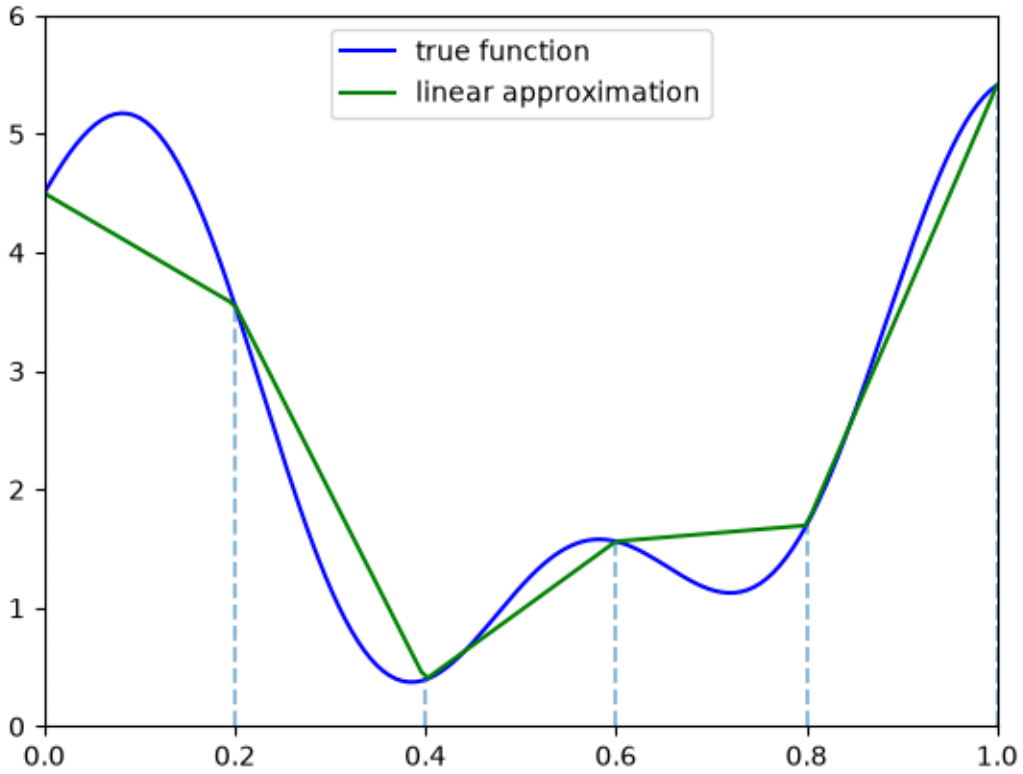
def Af(x):
    return jnp.interp(x, c_grid, f(c_grid))

fig, ax = plt.subplots()

ax.plot(f_grid, f(f_grid), 'b-', label='true function')
ax.plot(f_grid, Af(f_grid), 'g-', label='linear approximation')
ax.vlines(c_grid, c_grid * 0, f(c_grid), linestyle='dashed', alpha=0.5)

ax.legend(loc="upper center")

ax.set(xlim=(0, 1), ylim=(0, 6))
plt.show()
```



57.4 Implementation

Let's code up and solve the model.

57.4.1 Setup

The first step is to build a JAX-compatible structure for the McCall model with separation and a continuous wage offer distribution.

The key computational challenge is evaluating the conditional expectation $(Pv_u)(w) = \int v_u(w')p(w, w')dw'$ at each wage grid point.

Recall that we have:

$$(Pv_u)(w) = \int v_u(w^\rho \exp(\nu z))\psi(z)dz$$

where ψ is the standard normal density.

We will approximate this integral using Monte Carlo integration with draws $\{Z_i\}$ from the standard normal distribution:

$$(Pv_u)(w) \approx \frac{1}{N}\sum_{i=1}^{N} v_u(w^\rho \exp(\nu Z_i))$$

For this reason, our data structure will include a fixed set of IID $N(0, 1)$ draws $\{Z_i\}$.

```
class Model(NamedTuple):
    c: float              # unemployment compensation
    α: float              # job separation rate
    β: float              # discount factor
    ρ: float              # wage persistence
    ν: float              # wage volatility
    γ: float              # utility parameter
    w_grid: jnp.ndarray   # grid of points for fitted VFI
    z_draws: jnp.ndarray  # draws from the standard normal distribution

def create_mccall_model(
        c: float = 1.0,
        α: float = 0.05,
        β: float = 0.96,
        ρ: float = 0.9,
        ν: float = 0.2,
        γ: float = 1.5,
        grid_size: int = 100,
        mc_size: int = 1000,
        seed: int = 1234
    ):
    """Factory function to create a McCall model instance."""

    key = jax.random.key(seed)
    z_draws = jax.random.normal(key, (mc_size,))

    # Discretize just to get a suitable wage grid for interpolation
    mc = qe.markov.tauchen(grid_size, ρ, ν)
    w_grid = jnp.exp(jnp.array(mc.state_values))

    return Model(c, α, β, ρ, ν, γ, w_grid, z_draws)
```

We use the same CRRA utility function as in the discrete case:

```
def u(x, γ):
    return (x**(1 - γ) - 1) / (1 - γ)
```

57.4.2 Iteration

Here is the Bellman operator, where we use Monte Carlo integration to evaluate the expectation.

```
def T(model, v):
    """Update the value function."""

    # Unpack model parameters
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    # Interpolate array represented value function
    vf = lambda x: jnp.interp(x, w_grid, v)

    def compute_expectation(w):
        # Use Monte Carlo to evaluate integral (P v)(w) = E[v(W' | w)]
        # where W' = w^ρ * exp(ν * Z)
        w_next = w**ρ * jnp.exp(ν * z_draws)
        return jnp.mean(vf(w_next))
```

(continues on next page)

(continued from previous page)

```
    compute_exp_on_grid = jax.vmap(compute_expectation)
    Pv = compute_exp_on_grid(w_grid)

    d = 1 / (1 - β * (1 - α))
    v_e = d * (u(w_grid, γ) + α * β * Pv)
    continuation_values = u(c, γ) + β * Pv
    return jnp.maximum(v_e, continuation_values)
```

Here's the solver, which computes an approximate fixed point v_u of T.

```
@jax.jit
def vfi(
        model: Model,
        tolerance: float = 1e-6,    # Error tolerance
        max_iter: int = 100_000,    # Max iteration bound
    ):
    """
    Compute the fixed point v_u of T.

    """

    v_init = jnp.zeros(model.w_grid.shape)

    def cond(loop_state):
        v, error, i = loop_state
        return (error > tolerance) & (i <= max_iter)

    def update(loop_state):
        v, error, i = loop_state
        v_new = T(model, v)
        error = jnp.max(jnp.abs(v_new - v))
        new_loop_state = v_new, error, i + 1
        return new_loop_state

    initial_state = (v_init, tolerance + 1, 1)
    final_loop_state = lax.while_loop(cond, update, initial_state)
    v_final, error, i = final_loop_state

    return v_final
```

Here's a function that uses a solution v_u to compute the remaining functions of interest: v_e, and the continuation value function h.

We use the same expressions as we did in the *discrete case*, after replacing sums with integrals.

```
def compute_solution_functions(model, v_u):

    # Unpack model parameters
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    # Interpolate v_u on the wage grid
    vf = lambda x: jnp.interp(x, w_grid, v_u)

    def compute_expectation(w):
        # Use Monte Carlo to evaluate integral (P v)(w)
        # Compute E[v(w' | w)] where w' = w^ρ * exp(ν * z)
        w_next = w**ρ * jnp.exp(ν * z_draws)
```

(continues on next page)

(continued from previous page)

```
        return jnp.mean(vf(w_next))

    compute_exp_on_grid = jax.vmap(compute_expectation)
    Pv = compute_exp_on_grid(w_grid)

    d = 1 / (1 - β * (1 - α))
    v_e = d * (u(w_grid, γ) + α * β * Pv)
    h = u(c, γ) + β * Pv

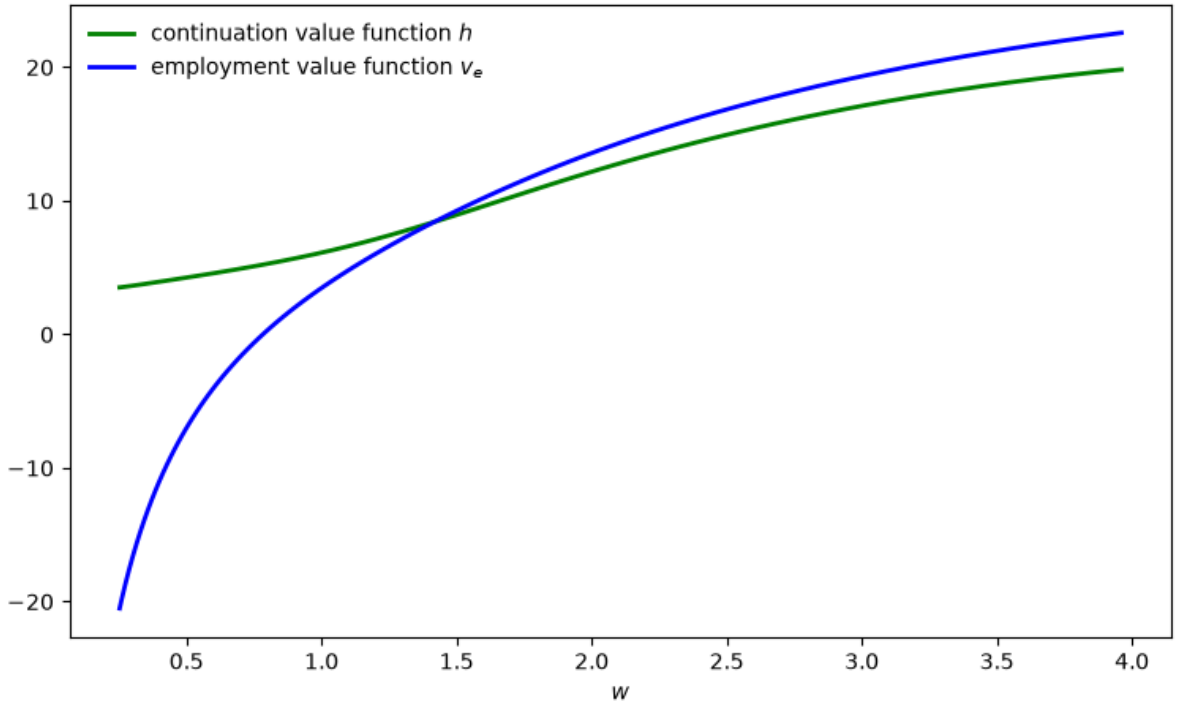
    return v_e, h
```

Let's try solving the model:

```
model = create_mccall_model()
c, α, β, ρ, ν, γ, w_grid, z_draws = model
v_u = vfi(model)
v_e, h = compute_solution_functions(model, v_u)
```

Let's plot our results.

```
fig, ax = plt.subplots(figsize=(9, 5.2))
ax.plot(w_grid, h, 'g-', linewidth=2,
        label="continuation value function $h$")
ax.plot(w_grid, v_e, 'b-', linewidth=2,
        label="employment value function $v_e$")
ax.legend(frameon=False)
ax.set_xlabel(r"$w$")
plt.show()
```



The reservation wage is at the intersection of the employment value function v_e and the continuation value function h.

Here's a function to compute it explicitly.

```
@jax.jit
def get_reservation_wage(model: Model) -> float:
    """
    Calculate the reservation wage for a given model.

    """
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    v_u = vfi(model)
    v_e, h = compute_solution_functions(model, v_u)

    # Compute optimal policy (acceptance indices)
    σ = v_e >= h

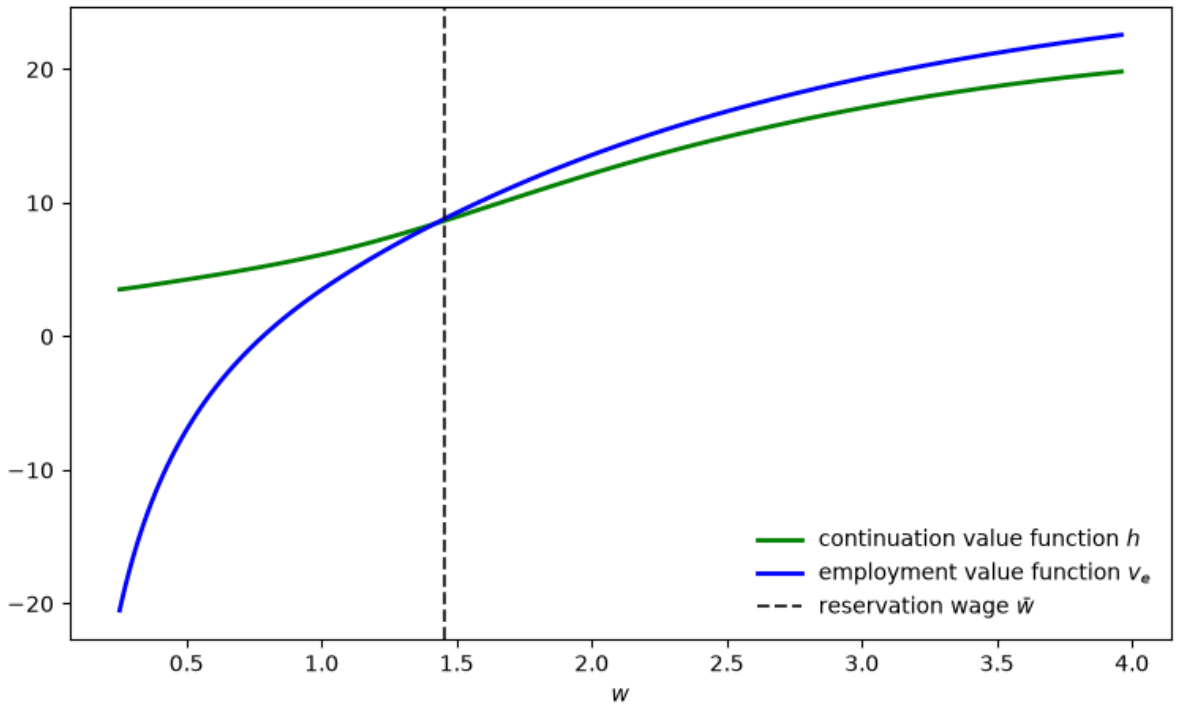
    # Find first index where policy indicates acceptance
    first_accept_idx = jnp.argmax(σ) # returns first True value

    # If no acceptance (all False), return infinity
    # Otherwise return the wage at the first acceptance index
    return jnp.where(jnp.any(σ), w_grid[first_accept_idx], jnp.inf)
```

Let's repeat our plot, but now inserting the reservation wage.

```
w_bar = get_reservation_wage(model)

fig, ax = plt.subplots(figsize=(9, 5.2))
ax.plot(w_grid, h, 'g-', linewidth=2,
        label="continuation value function $h$")
ax.plot(w_grid, v_e, 'b-', linewidth=2,
        label="employment value function $v_e$")
ax.axvline(x=w_bar, color='black', linestyle='--', alpha=0.8,
           label=f'reservation wage $\\bar{{w}}$')
ax.legend(frameon=False)
ax.set_xlabel(r"$w$")
plt.show()
```



57.5 Simulation

Now we run some simulations with a focus on unemployment rate.

57.5.1 Single agent dynamics

Let's simulate the employment path of a single agent under the optimal policy.

We need a function to update the agent's state by one period.

```
def update_agent(key, status, wage, model, w_bar):
    """
    Updates an agent's employment status and current wage by one period.

    Parameters:
    - key: JAX random key
    - status: Current employment status (0 or 1)
    - wage: Current wage if employed, current offer if unemployed
    - model: Model instance
    - w_bar: Reservation wage

    """
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    # Draw new wage offer based on current wage
    key1, key2 = jax.random.split(key)
    z = jax.random.normal(key1)
    new_wage = wage**ρ * jnp.exp(ν * z)
```

(continues on next page)

(continued from previous page)

```
    # Check if separation occurs (for employed workers)
    separation_occurs = jax.random.uniform(key2) < α

    # Accept if current wage meets or exceeds reservation wage
    accepts = wage >= w_bar

    # If employed: status = 1 if no separation, 0 if separation
    # If unemployed: status = 1 if accepts, 0 if rejects
    next_status = jnp.where(
        status,
        1 - separation_occurs.astype(jnp.int32),  # employed path
        accepts.astype(jnp.int32)                 # unemployed path
    )

    # If employed: wage = current if no separation, new if separation
    # If unemployed: wage = current if accepts, new if rejects
    next_wage = jnp.where(
        status,
        jnp.where(separation_occurs, new_wage, wage),  # employed path
        jnp.where(accepts, wage, new_wage)             # unemployed path
    )

    return next_status, next_wage
```

Here's a function to simulate the employment path of a single agent.

```
def simulate_employment_path(
        model: Model,      # Model details
        w_bar: float,      # Reservation wage
        T: int = 2_000,    # Simulation length
        seed: int = 42     # Set seed for simulation
    ):
    """
    Simulate employment path for T periods starting from unemployment.

    """
    key = jax.random.key(seed)
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    # Initial conditions: start unemployed with initial wage draw
    status = 0
    key, subkey = jax.random.split(key)
    wage = jnp.exp(jax.random.normal(subkey) * ν)

    wage_path = []
    status_path = []

    for t in range(T):
        wage_path.append(wage)
        status_path.append(status)

        key, subkey = jax.random.split(key)
        status, wage = update_agent(
            subkey, status, wage, model, w_bar
        )

    return jnp.array(wage_path), jnp.array(status_path)
```

Let's create a comprehensive plot of the employment simulation:

```
model = create_mccall_model()
w_bar = get_reservation_wage(model)

wage_path, employment_status = simulate_employment_path(model, w_bar)

fig, (ax1, ax2, ax3) = plt.subplots(3, 1, figsize=(8, 6))

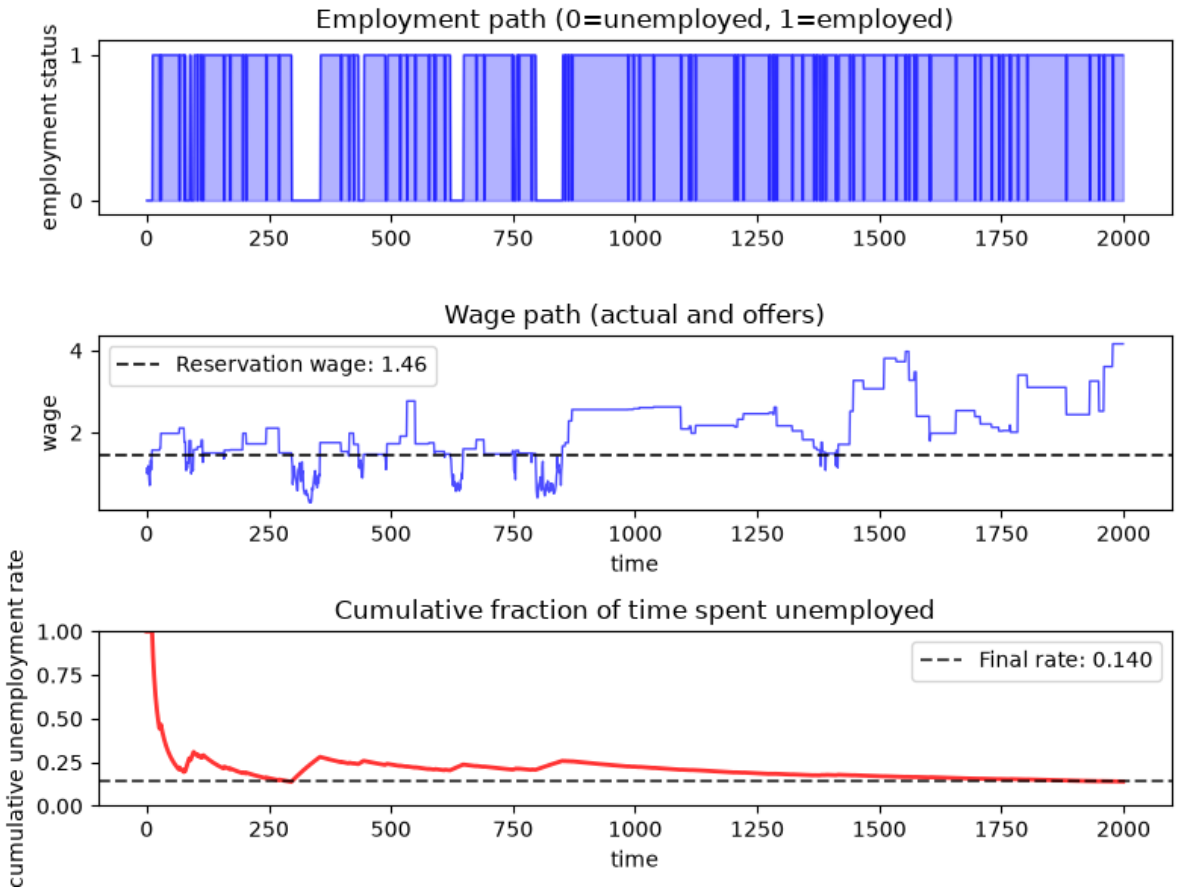
# Plot employment status
ax1.plot(employment_status, 'b-', alpha=0.7, linewidth=1)
ax1.fill_between(
    range(len(employment_status)), employment_status, alpha=0.3, color='blue'
)
ax1.set_ylabel('employment status')
ax1.set_title('Employment path (0=unemployed, 1=employed)')
ax1.set_yticks((0, 1))
ax1.set_ylim(-0.1, 1.1)

# Plot wage path with reservation wage
ax2.plot(wage_path, 'b-', alpha=0.7, linewidth=1)
ax2.axhline(y=w_bar, color='black', linestyle='--', alpha=0.8,
           label=f'Reservation wage: {w_bar:.2f}')
ax2.set_xlabel('time')
ax2.set_ylabel('wage')
ax2.set_title('Wage path (actual and offers)')
ax2.legend()

# Plot cumulative fraction of time unemployed
unemployed_indicator = (employment_status == 0).astype(int)
cumulative_unemployment = (
    jnp.cumsum(unemployed_indicator) /
    jnp.arange(1, len(employment_status) + 1)
)

ax3.plot(cumulative_unemployment, 'r-', alpha=0.8, linewidth=2)
ax3.axhline(y=jnp.mean(unemployed_indicator), color='black',
            linestyle='--', alpha=0.7,
            label=f'Final rate: {jnp.mean(unemployed_indicator):.3f}')
ax3.set_xlabel('time')
ax3.set_ylabel('cumulative unemployment rate')
ax3.set_title('Cumulative fraction of time spent unemployed')
ax3.legend()
ax3.set_ylim(0, 1)

plt.tight_layout()
plt.show()
```



The simulation shows the agent cycling between employment and unemployment.

The agent starts unemployed and receives wage offers according to the Markov process.

When unemployed, the agent accepts offers that exceed the reservation wage.

When employed, the agent faces job separation with probability α each period.

57.5.2 Cross-sectional analysis

Now let's simulate many agents simultaneously to examine the cross-sectional unemployment rate.

To do this efficiently, we need a different approach than `simulate_employment_path` defined above.

The key differences are:

- `simulate_employment_path` records the entire history (all T periods) for a single agent, which is useful for visualization but memory-intensive
- The new function `sim_agent` below only tracks and returns the final state, which is all we need for cross-sectional statistics
- `sim_agent` uses `lax.fori_loop` instead of a Python loop, making it JIT-compilable and suitable for vectorization across many agents

We first define a function that simulates a single agent forward T time steps:

```
@jax.jit
def sim_agent(key, initial_status, initial_wage, model, w_bar, T):
    """
    Simulate a single agent forward T time steps using lax.fori_loop.

    Uses fold_in to generate a new key at each time step.

    Parameters:
    - key: JAX random key for this agent
    - initial_status: Initial employment status (0 or 1)
    - initial_wage: Initial wage
    - model: Model instance
    - w_bar: Reservation wage
    - T: Number of time periods to simulate

    Returns:
    - final_status: Employment status after T periods
    - final_wage: Wage after T periods
    """
    def update(t, loop_state):
        status, wage = loop_state
        step_key = jax.random.fold_in(key, t)
        status, wage = update_agent(step_key, status, wage, model, w_bar)
        return status, wage

    initial_loop_state = (initial_status, initial_wage)
    final_loop_state = lax.fori_loop(0, T, update, initial_loop_state)
    final_status, final_wage = final_loop_state
    return final_status, final_wage


# Create vectorized version of sim_agent to process multiple agents in parallel
sim_agents_vmap = jax.vmap(sim_agent, in_axes=(0, 0, 0, None, None, None))


def simulate_cross_section(
        model: Model,
        n_agents: int = 100_000,
        T: int = 200,
        seed: int = 42
    ) -> float:
    """
    Simulate cross-section of agents and return unemployment rate.

    This approach:
    1. Generates n_agents random keys
    2. Calls sim_agent for each agent (vectorized via vmap)
    3. Collects the final states to produce the cross-section

    Returns the cross-sectional unemployment rate.
    """
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    key = jax.random.key(seed)

    # Solve for optimal reservation wage
    w_bar = get_reservation_wage(model)
```

(continues on next page)

(continued from previous page)

```
    # Initialize arrays
    init_key, subkey = jax.random.split(key)
    initial_wages = jnp.exp(jax.random.normal(subkey, (n_agents,)) * v)
    initial_status_vec = jnp.zeros(n_agents, dtype=jnp.int32)

    # Generate n_agents random keys
    agent_keys = jax.random.split(init_key, n_agents)

    # Simulate each agent forward T steps (vectorized)
    final_status, final_wages = sim_agents_vmap(
        agent_keys, initial_status_vec, initial_wages, model, w_bar, T
    )

    unemployment_rate = 1 - jnp.mean(final_status)
    return unemployment_rate
```

Now let's compare the time-average unemployment rate (from a single agent's long simulation) with the cross-sectional unemployment rate (from many agents at a single point in time).

```
model = create_mccall_model()
cross_sectional_unemp = simulate_cross_section(
    model, n_agents=20_000, T=200
)

time_avg_unemp = jnp.mean(unemployed_indicator)
print(f"Time-average unemployment rate (single agent, T=2000): "
      f"{time_avg_unemp:.4f}")
print(f"Cross-sectional unemployment rate (at t=200): "
      f"{cross_sectional_unemp:.4f}")
print(f"Difference: {abs(time_avg_unemp - cross_sectional_unemp):.4f}")
```

```
Time-average unemployment rate (single agent, T=2000): 0.1395
Cross-sectional unemployment rate (at t=200): 0.2029
Difference: 0.0634
```

The difference above can be further reduced by increasing the simulation length for the single agent.

```
wage_path_long, employment_status_long = simulate_employment_path(model, w_bar, T=10_
↪000)
unemployed_indicator_long = (employment_status_long == 0).astype(int)
time_avg_unemp_long = jnp.mean(unemployed_indicator_long)

print(f"Time-average unemployment rate (single agent, T=10000): "
      f"{time_avg_unemp_long:.4f}")
print(f"Cross-sectional unemployment rate (at t=200): "
      f"{cross_sectional_unemp:.4f}")
print(f"Difference: {abs(time_avg_unemp_long - cross_sectional_unemp):.4f}")
```

```
Time-average unemployment rate (single agent, T=10000): 0.2071
Cross-sectional unemployment rate (at t=200): 0.2029
Difference: 0.0042
```

57.5.3 Visualization

This function generates a histogram showing the distribution of employment status across many agents:

```
def plot_cross_sectional_unemployment(
        model: Model,             # Model instance with parameters
        t_snapshot: int = 200,    # Time for cross-sectional snapshot
        n_agents: int = 20_000    # Number of agents to simulate
    ):
    """
    Generate histogram of cross-sectional unemployment at a specific time.

    """
    c, α, β, ρ, ν, γ, w_grid, z_draws = model

    # Get final employment state directly
    key = jax.random.key(42)
    w_bar = get_reservation_wage(model)

    # Initialize arrays
    init_key, subkey = jax.random.split(key)
    initial_wages = jnp.exp(jax.random.normal(subkey, (n_agents,)) * ν)
    initial_status_vec = jnp.zeros(n_agents, dtype=jnp.int32)

    # Generate n_agents random keys
    agent_keys = jax.random.split(init_key, n_agents)

    # Simulate each agent forward T steps (vectorized)
    final_status, _ = sim_agents_vmap(
        agent_keys, initial_status_vec, initial_wages, model, w_bar, t_snapshot
    )

    # Calculate unemployment rate
    unemployment_rate = 1 - jnp.mean(final_status)

    fig, ax = plt.subplots(figsize=(8, 5))

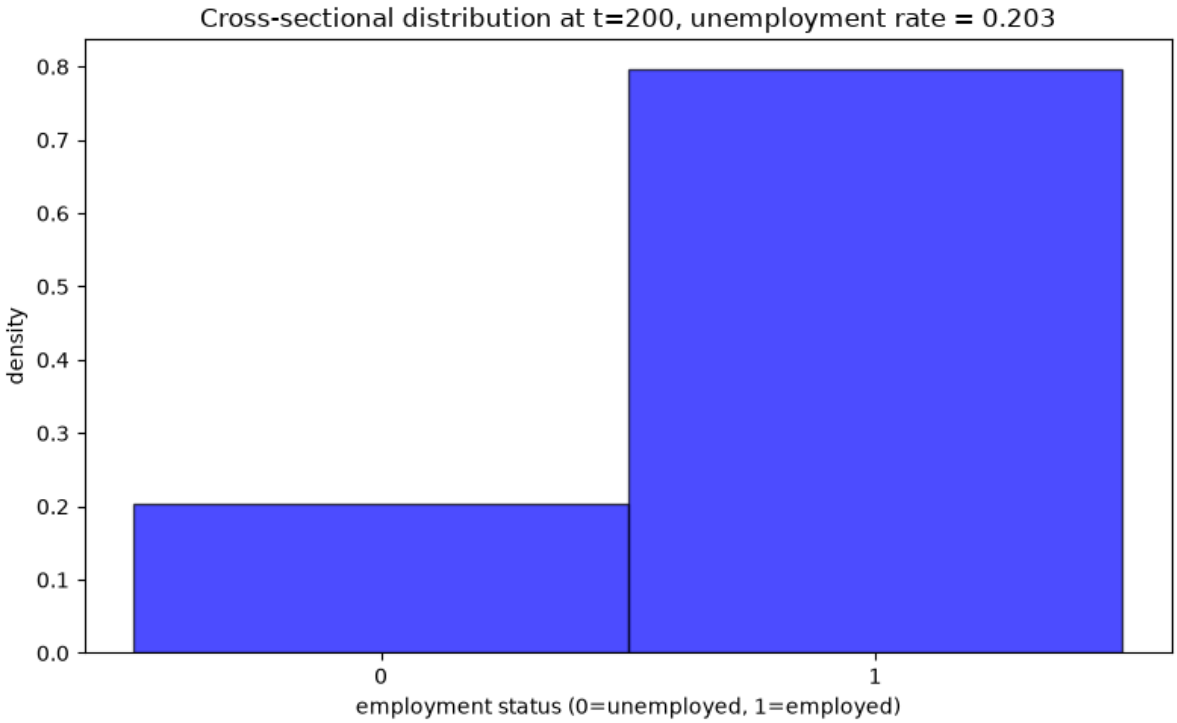
    # Plot histogram as density (bars sum to 1)
    weights = jnp.ones_like(final_status) / len(final_status)
    ax.hist(final_status, bins=[-0.5, 0.5, 1.5],
            alpha=0.7, color='blue', edgecolor='black',
            density=True, weights=weights)

    ax.set_xlabel('employment status (0=unemployed, 1=employed)')
    ax.set_ylabel('density')
    ax.set_title(f'Cross-sectional distribution at t={t_snapshot}, ' +
                 f'unemployment rate = {unemployment_rate:.3f}')
    ax.set_xticks([0, 1])

    plt.tight_layout()
    plt.show()
```

Let's plot the cross-sectional distribution:

```
plot_cross_sectional_unemployment(model)
```



57.6 Exercises

Exercise 57.6.1

Use the code above to explore what happens to the reservation wage when c changes.

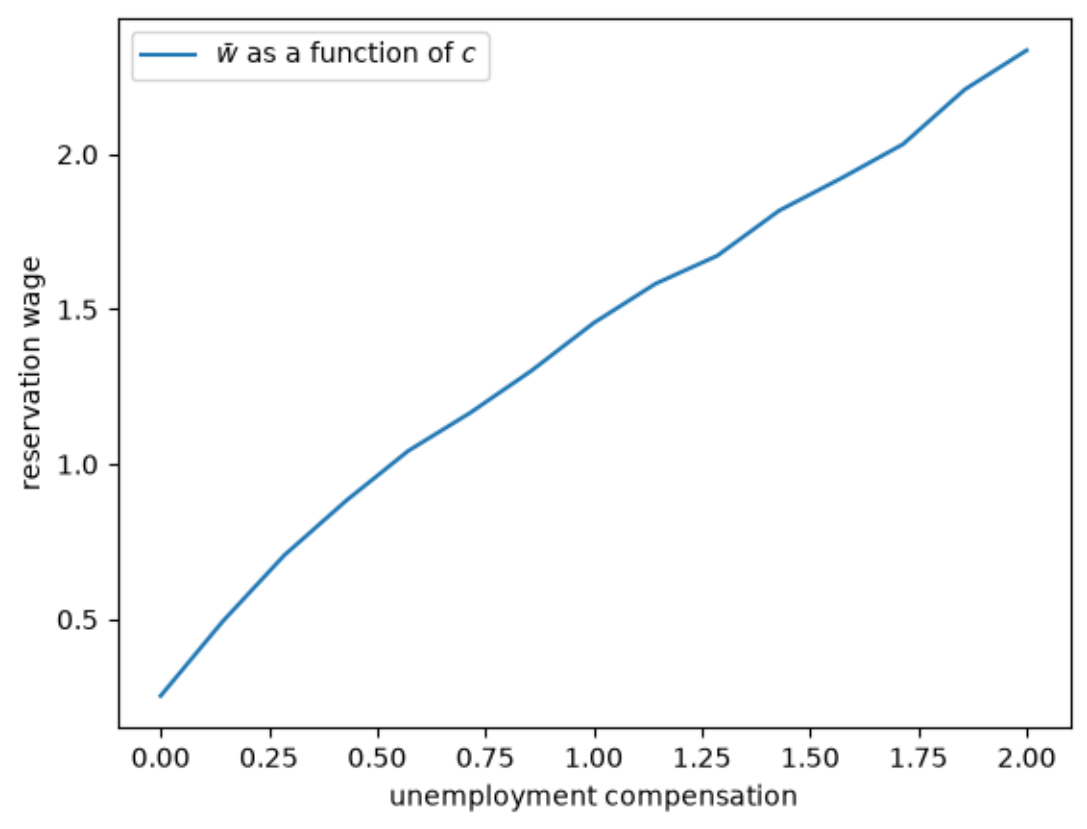
Solution

Here is one solution

```
def compute_res_wage_given_c(c):
    model = create_mccall_model(c=c)
    w_bar = get_reservation_wage(model)
    return w_bar

c_vals = jnp.linspace(0.0, 2.0, 15)
w_bar_vals = jax.vmap(compute_res_wage_given_c)(c_vals)

fig, ax = plt.subplots()
ax.set(xlabel='unemployment compensation', ylabel='reservation wage')
ax.plot(c_vals, w_bar_vals, label=r'$\bar w$ as a function of $c$')
ax.legend()
plt.show()
```



As unemployment compensation increases, the reservation wage also increases.

This makes economic sense: when the value of being unemployed rises (through higher c), workers become more selective about which job offers to accept.

Exercise 57.6.2

Create a plot that shows how the reservation wage changes with the risk aversion parameter γ.

Use `γ_vals = jnp.linspace(1.2, 2.5, 15)` and keep all other parameters at their default values.

How do you expect the reservation wage to vary with γ? Why?

Solution

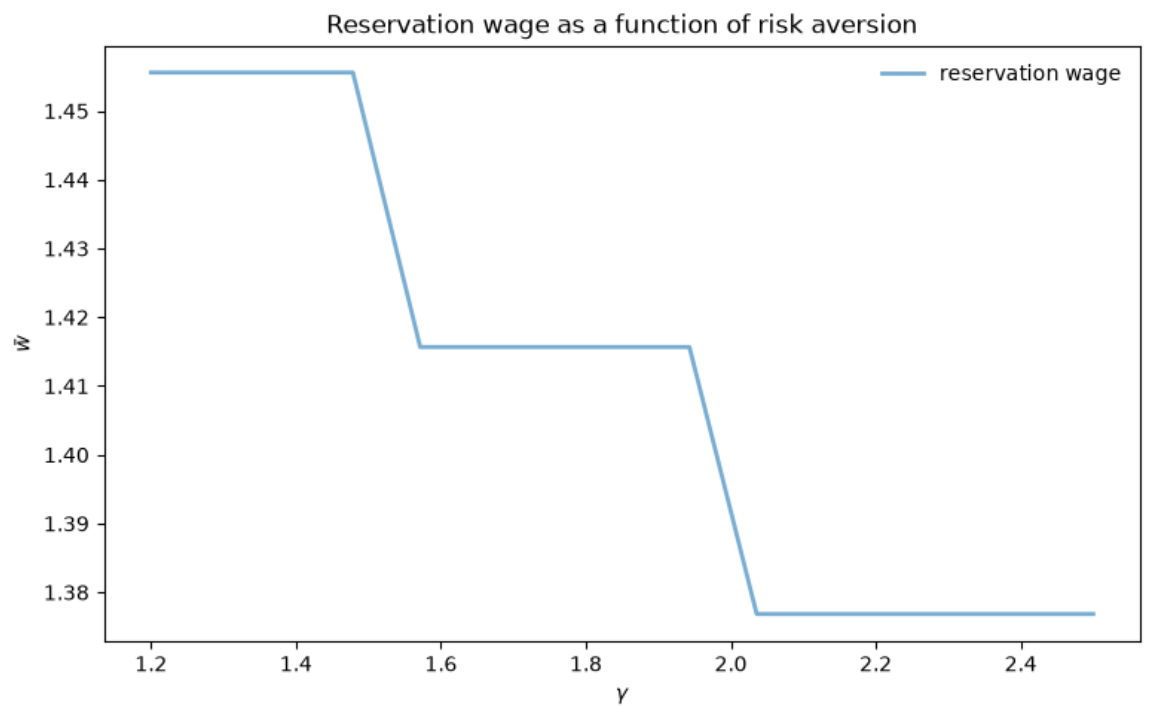
We compute the reservation wage for different values of the risk aversion parameter:

```
γ_vals = jnp.linspace(1.2, 2.5, 15)
w_bar_vec = jnp.empty_like(γ_vals)

for i, γ in enumerate(γ_vals):
    model = create_mccall_model(γ=γ)
    w_bar = get_reservation_wage(model)
    w_bar_vec = w_bar_vec.at[i].set(w_bar)

fig, ax = plt.subplots(figsize=(9, 5.2))
```

```
ax.plot(γ_vals, w_bar_vec, linewidth=2, alpha=0.6,
        label='reservation wage')
ax.legend(frameon=False)
ax.set_xlabel(r'$\gamma$')
ax.set_ylabel(r'$\bar{w}$')
ax.set_title('Reservation wage as a function of risk aversion')
plt.show()
```



As risk aversion (γ) increases, the reservation wage decreases.

This occurs because more risk-averse workers place higher value on the security of employment relative to the uncertainty of continued search.

With higher γ, the utility cost of unemployment (foregone consumption) becomes more severe, making workers more willing to accept lower wages rather than continue searching.

CHAPTER

FIFTYEIGHT

JOB SEARCH V: RISK-SENSITIVE PREFERENCES

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Job Search V: Risk-Sensitive Preferences*
  - *Overview*
  - *Outline*
  - *Introduction to risk-sensitivity*
  - *Back to job search*

58.1 Overview

Risk-sensitive preferences are a common addition to various types of dynamic programming problems.

This lecture gives an introduction to risk-sensitive recursive preferences via job search.

Some motivation is given below.

58.2 Outline

In real-world job-related decisions, individuals and households care about risk.

For example, some individuals might prefer to take a moderate offer already in hand over the risky possibility of a higher offer in the next period, even without discounting future payoffs.

(A bird in the hand is worth two in the bush, etc.)

In previous job search lectures in this series, we inserted some degree of risk aversion by adding a concave flow utility function u.

Unfortunately, this strategy does not isolate preferences over the kind of risk we described above.

This is because adding a concave utility function changes the agent's preferences in other ways, such as in their desire for consumption smoothing.

Hence, if we want to study the pure effects of risk, we need a different solution.

One possibility is to add risk-sensitive preferences.

Here we show how this can be done and study what effects it has on agent choices.

We'll use JAX and the QuantEcon library:

```
!pip install quantecon
```

We use the following imports.

```
import jax
import jax.numpy as jnp
from jax import lax
import quantecon as qe
from typing import NamedTuple
import numpy as np
import matplotlib.pyplot as plt
```

58.3 Introduction to risk-sensitivity

Let's start our discussion in a static environment.

If Y is a random payoff and an agent's evaluation of the payoff is $e := \mathbb{E}Y$, then we say that the agent is **risk neutral**.

Sometimes we want to model agents as risk averse.

One way to do this is to change their evaluation of the payoff Y to

$$e_\theta = \frac{1}{\theta} \ln \left(\mathbb{E}[\exp(\theta Y)] \right)$$

where θ is a number satisfying $\theta < 0$.

The value e_θ is sometimes called the **entropic risk-adjusted expectation** of Y.

58.3.1 A Gaussian example

One way to see the impact is to suppose that Y has the normal distribution $N(\mu, \sigma^2)$, so that its mean is μ and its variance is σ^2.

For this Y we aim to compute the risk-adjusted expectation.

This becomes straightforward if we recognize that $\mathbb{E}[\exp(\theta Y)]$ is the moment generating function (MGF) of the normal distribution.

Using the well-known expression for the MGF of the normal distribution, we get

$$\mathbb{E}[\exp(\theta Y)] = \exp\left(\theta\mu + \frac{\theta^2\sigma^2}{2}\right)$$

Therefore,

$$e_\theta = \frac{1}{\theta}\ln\left(\exp\left(\theta\mu + \frac{\theta^2\sigma^2}{2}\right)\right) = \frac{1}{\theta}\left(\theta\mu + \frac{\theta^2\sigma^2}{2}\right)$$

Simplifying yields

$$e_\theta = \mu + \frac{\theta\sigma^2}{2}$$

We see immediately that the agent prefers a higher average payoff μ.

At the same time, given that $\theta < 0$, the risk-adjusted expectation decreases in σ.

In particular, e_θ decreases as risk increases.

Here is a visualization of e_θ as a function of μ and σ using a contour plot, with $\theta = -1$.

```
θ = -1

μ_vals = np.linspace(-2, 5, 200)
σ_vals = np.linspace(0.1, 3, 200)
μ_grid, σ_grid = np.meshgrid(μ_vals, σ_vals)

e_θ = μ_grid + (θ * σ_grid**2) / 2

# Create contour plot
fig, ax = plt.subplots()
contour = ax.contour(
    μ_grid, σ_grid, e_θ, levels=20, colors='black', linewidths=0.5
)
contourf = ax.contourf(
    μ_grid, σ_grid, e_θ, levels=20, cmap='viridis'
)
ax.clabel(contour, inline=True)
cbar = plt.colorbar(contourf, ax=ax)
cbar.set_label(r'$e_\theta$', rotation=0)

ax.set_xlabel(r'$\mu$ (mean)')
ax.set_ylabel(r'$\sigma$ (standard deviation)')
plt.tight_layout()
plt.show()
```

Again, we see that the agent prefers a higher average payoff but dislikes risk.

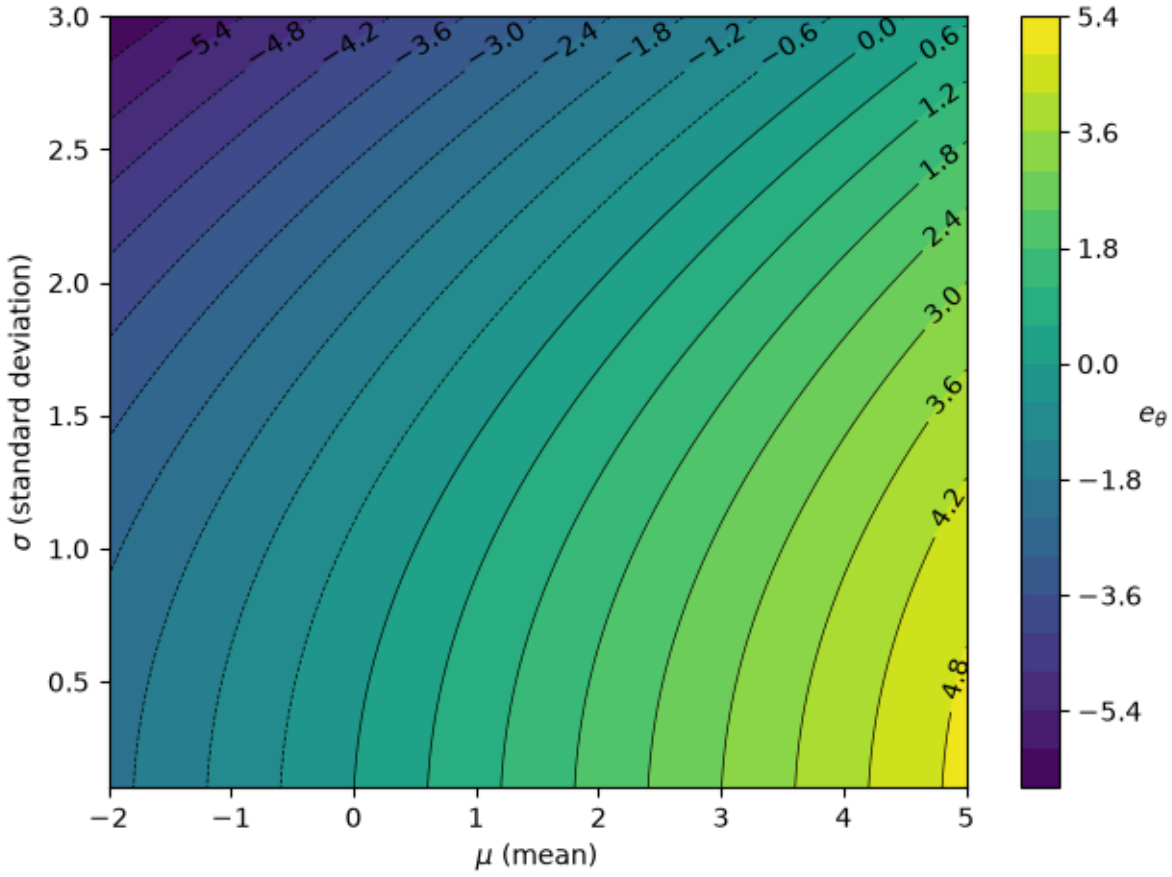


Fig. 58.1: Entropic risk-adjusted expectation, Gaussian case

58.3.2 A more general case

The preceding analysis relies on the Gaussian (normal) assumption to get an analytical solution.

We can investigate all other cases using simulation.

For example, suppose that Y has the Beta(a, b) distribution.

Here we set $a = b = 2.0$ and calculate e_θ using Monte Carlo

The method is:

1. sample $Y_1, \ldots, Y_n$ from Beta$(2, 2)$
2. replace $\mathbb{E}$ with an average over $\exp(\theta Y_i)$

We do this for θ in a grid of 100 points between -2 and -0.1.

Here is a plot of e_θ against θ.

```
# Set parameters
a, b = 2.0, 2.0
mc_size = 1_000_000  # Large number of Monte Carlo samples
θ_grid = jnp.linspace(-2, -0.1, 100)

# Draw samples from Beta(2, 2) distribution using JAX
key = jax.random.key(1234)
Y_samples = jax.random.beta(key, a, b, shape=(mc_size,))

# Define function to compute e_θ for a single θ value
def compute_e_θ(θ):
    """Compute e_θ = (1/θ) * ln(E[exp(θ * Y)])"""
    expectation = jnp.mean(jnp.exp(θ * Y_samples))
    return (1 / θ) * jnp.log(expectation)

# Vectorize over θ_grid using vmap
compute_e_θ_vec = jax.vmap(compute_e_θ)
e_θ_values = compute_e_θ_vec(θ_grid)

# Plot results
fig, ax = plt.subplots()
ax.plot(θ_grid, e_θ_values, lw=2)
ax.set_xlabel(r'$\theta$')
ax.set_ylabel(r'$e_\theta$')
ax.axhline(y=0.5, color='black', linestyle='--',
           linewidth=1, label=r'$\mathbb{E}[Y] = 0.5$')
ax.legend()
plt.tight_layout()
plt.show()
```

The plot shows how the risk-adjusted evaluation e_θ changes with the risk aversion parameter θ.

As $\theta \to 0$, the value e_θ approaches the expected value of Y, which is $\mathbb{E}[Y] = \frac{a}{a+b} = \frac{2}{4} = 0.5$ for Beta(2,2).

This makes sense because when $\theta \to 0$, the agent becomes risk neutral.

As θ becomes more negative, e_θ decreases.

This reflects that a more risk-averse agent values the uncertain payoff Y less than its expected value.

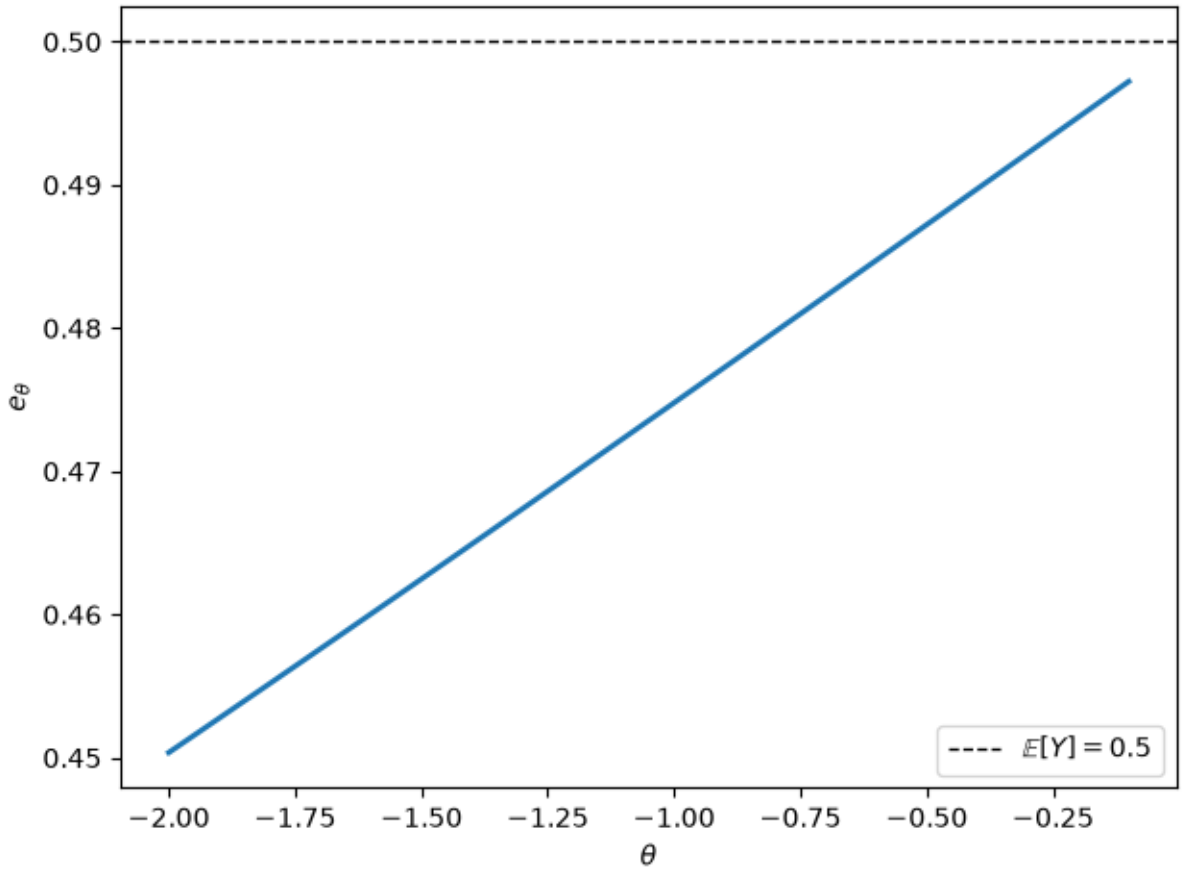


Fig. 58.2: Risk-adjusted evaluation against risk aversion

58.3.3 A mean preserving spread

The next exercise asks you to study the impact of a mean-preserving spread on the risk-adjusted expectation.

Exercise 58.3.1

Keep $Y \sim \text{Beta}(2, 2)$ and fix $\theta = -2$.

Using Monte Carlo again, calculate

$$e_\theta = \frac{1}{\theta} \ln \left(\mathbb{E}[\exp(\theta X)] \right)$$

where $X = Y + \sigma Z$ and Z is standard normal.

How does e_θ change with σ?

Can you provide some intuition for what is happening (given that the agent is risk averse)?

Use a plot to illustrate your results.

Solution

Here's our solution.

```
a, b = 2.0, 2.0
θ = -2.0
mc_size = 1_000_000  # Large number of Monte Carlo samples
σ_grid = jnp.linspace(0.0, 1.0, 50)

# Set random seed for reproducibility
key = jax.random.key(1234)
key_y, key_z = jax.random.split(key)

# Draw samples from Beta(2, 2) distribution
Y_samples = jax.random.beta(key_y, a, b, shape=(mc_size,))

# Draw standard normal samples (reused for all σ values)
Z_samples = jax.random.normal(key_z, shape=(mc_size,))

# Define function to compute e_θ for a single σ value
def compute_e_θ(σ):
    """Compute e_θ for X = Y + σ * Z"""
    # Compute X = Y + σ * Z
    X_samples = Y_samples + σ * Z_samples

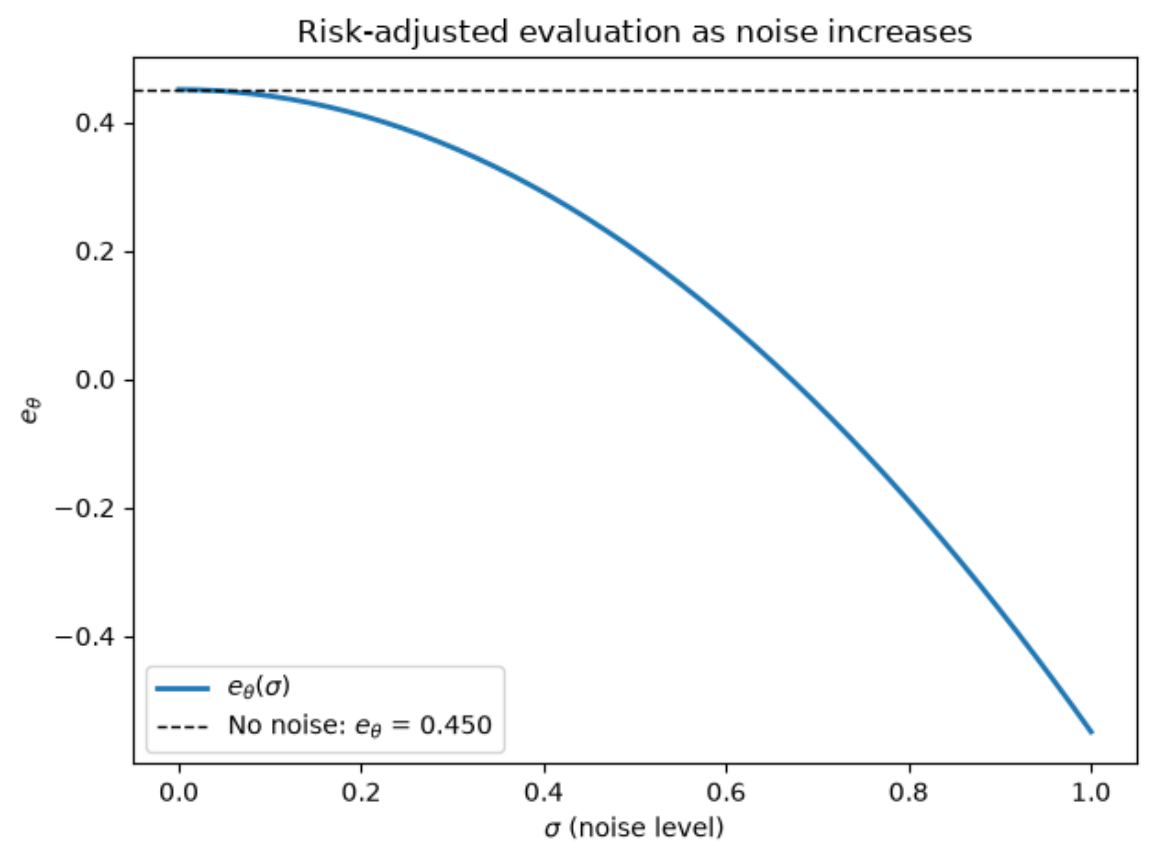
    # Calculate E[exp(θ * X)] using Monte Carlo
    expectation = jnp.mean(jnp.exp(θ * X_samples))

    # Calculate e_θ
    return (1 / θ) * jnp.log(expectation)

# Vectorize over σ_grid using vmap
compute_e_θ_vec = jax.vmap(compute_e_θ)
e_θ_values = compute_e_θ_vec(σ_grid)

# Plot results
fig, ax = plt.subplots()
```

```
ax.plot(σ_grid, e_θ_values, lw=2, label=r'$e_\theta(\sigma)$')
ax.set_xlabel(r'$\sigma$ (noise level)')
ax.set_ylabel(r'$e_\theta$')
ax.set_title(r'Risk-adjusted evaluation as noise increases')
ax.axhline(y=e_θ_values[0], color='black', linestyle='--', linewidth=1,
           label=f'No noise: $e_\\theta$ = {e_θ_values[0]:.3f}')
ax.legend()
plt.tight_layout()
plt.show()
```



The plot clearly shows that e_θ decreases monotonically as σ increases.

Since the agent is risk averse ($\theta = -2 < 0$), she dislikes uncertainty.

As we increase σ, we get more volatility, since

$$\text{Var}(X) = \text{Var}(Y) + \sigma^2 \text{Var}(Z) = \text{Var}(Y) + \sigma^2$$

At the same time, the expected value is unchanged, since

$$\mathbb{E}[X] = \mathbb{E}[Y + \sigma Z] = \mathbb{E}[Y] + \sigma \mathbb{E}[Z] = \mathbb{E}[Y] = 0.5$$

Hence the mean payoff doesn't change with σ.

In other words, the risk averse agent is not compensated for bearing additional risk.

This is why the valuation of the random payoff goes down.

58.4 Back to job search

In the lecture *Job Search IV: Fitted Value Function Iteration* we studied a job search model with separation, Markov wage draws and fitted value function iteration.

The wage offer process is continuous and obeys

$$W_t = \exp(X_t) \quad \text{where} \quad X_{t+1} = \rho X_t + \nu Z_{t+1}$$

and $\{Z_t\}$ is IID and standard normal.

Let's now study the same model but replacing the assumption of risk neutral expectations with risk averse expectations.

In particular, the conditional expectation

$$(Pv_u)(w) = \mathbb{E} v_u(w^\rho \exp(\nu Z))$$

from that lecture is replaced with

$$(P_\theta v_u)(w) = \frac{1}{\theta} \ln \left[\mathbb{E} \exp(\theta v_u(w^\rho \exp(\nu Z))) \right]$$

Otherwise the model is the same.

We now solve the dynamic program and study the impact of θ on the reservation wage.

58.4.1 Setup

Here's a class to store parameters and default parameter values.

```
class Model(NamedTuple):
    c: float              # unemployment compensation
    α: float              # job separation rate
    β: float              # discount factor
    ρ: float              # wage persistence
    ν: float              # wage volatility
    θ: float              # risk aversion parameter
    w_grid: jnp.ndarray   # grid of points for fitted VFI
    z_draws: jnp.ndarray  # draws from the standard normal distribution

def create_mccall_model(
        c: float = 1.0,
        α: float = 0.1,
        β: float = 0.96,
        ρ: float = 0.9,
        ν: float = 0.2,
        θ: float = -1.5,
        grid_size: int = 100,
        mc_size: int = 1000,
        seed: int = 1234
    ):
    """Factory function to create a McCall model instance."""

    key = jax.random.key(seed)
    z_draws = jax.random.normal(key, (mc_size,))

    # Discretize just to get a suitable wage grid for interpolation
```

(continues on next page)

(continued from previous page)

```
    mc = qe.markov.tauchen(grid_size, ρ, ν)
    w_grid = jnp.exp(jnp.array(mc.state_values))

    return Model(c, α, β, ρ, ν, θ, w_grid, z_draws)
```

58.4.2 Bellman equations

Our construction is a direct extension of the Bellman equations in *Job Search IV: Fitted Value Function Iteration*.

First we use the employed worker's Bellman equation to express $v_e(w)$ in terms of $(P_\theta v_u)(w)$:

$$v_e(w) = \frac{1}{1-\beta(1-\alpha)} \cdot (u(w) + \alpha\beta(P_\theta v_u)(w))$$

We substitute into the unemployed agent's Bellman equation to get:

$$v_u(w) = \max\left\{\frac{1}{1-\beta(1-\alpha)} \cdot (u(w) + \alpha\beta(P_\theta v_u)(w)), u(c) + \beta(P_\theta v_u)(w)\right\}$$

We use value function iteration to solve for v_u.

Then we compute the optimal policy: accept if $v_e(w) \geq u(c) + \beta(P_\theta v_u)(w)$

Here's the Bellman operator that updates v_u.

```
def T(model, v):
    # Unpack model parameters
    c, α, β, ρ, ν, θ, w_grid, z_draws = model

    # Interpolate array represented value function
    vf = lambda x: jnp.interp(x, w_grid, v)

    def compute_expectation(w):
        # Use Monte Carlo to evaluate integral (P_θ v)(w)
        inner = jnp.mean(jnp.exp(θ * vf(w**ρ * jnp.exp(ν * z_draws))))
        return (1 / θ) * jnp.log(inner)

    compute_exp_all = jax.vmap(compute_expectation)
    P_θ_v = compute_exp_all(w_grid)

    d = 1 / (1 - β * (1 - α))
    accept = d * (w_grid + α * β * P_θ_v)
    reject = c + β * P_θ_v
    return jnp.maximum(accept, reject)
```

Here's the solver:

```
@jax.jit
def vfi(
        model: Model,
        tolerance: float = 1e-6,    # Error tolerance
        max_iter: int = 100_000,    # Max iteration bound
    ):

    v_init = jnp.zeros(model.w_grid.shape)
```

(continues on next page)

(continued from previous page)

```
    def cond(loop_state):
        v, error, i = loop_state
        return (error > tolerance) & (i <= max_iter)

    def update(loop_state):
        v, error, i = loop_state
        v_new = T(model, v)
        error = jnp.max(jnp.abs(v_new - v))
        new_loop_state = v_new, error, i + 1
        return new_loop_state

    initial_state = (v_init, tolerance + 1, 1)
    final_loop_state = lax.while_loop(cond, update, initial_state)
    v_final, error, i = final_loop_state

    return v_final
```

The next function computes the optimal policy under the assumption that v is the value function:

```
def get_greedy(v: jnp.ndarray, model: Model) -> jnp.ndarray:
    """Get a v-greedy policy."""
    c, α, β, ρ, ν, θ, w_grid, z_draws = model

    # Interpolate array represented value function
    vf = lambda x: jnp.interp(x, w_grid, v)

    def compute_expectation(w):
        # Use Monte Carlo to evaluate integral (P_θ v)(w)
        inner = jnp.mean(jnp.exp(θ * vf(w**ρ * jnp.exp(ν * z_draws))))
        return (1 / θ) * jnp.log(inner)

    compute_exp_all = jax.vmap(compute_expectation)
    P_θ_v = compute_exp_all(w_grid)

    d = 1 / (1 - β * (1 - α))
    accept = d * (w_grid + α * β * P_θ_v)
    reject = c + β * P_θ_v
    σ = accept >= reject
    return σ
```

Here's a function that takes an instance of `Model` and returns the associated reservation wage.

```
@jax.jit
def get_reservation_wage(σ: jnp.ndarray, model: Model) -> float:
    """
    Calculate the reservation wage from a given policy.

    Parameters:
    - σ: Policy array where σ[i] = True means accept wage w_grid[i]
    - model: Model instance containing wage values

    Returns:
    - Reservation wage (lowest wage for which policy indicates acceptance)
    """
    c, α, β, ρ, ν, θ, w_grid, z_draws = model

    # Find the first index where policy indicates acceptance
```

(continues on next page)

(continued from previous page)

```
    # σ is a boolean array, argmax returns the first True value
    first_accept_idx = jnp.argmax(σ)

    # If no acceptance (all False), return infinity
    # Otherwise return the wage at the first acceptance index
    return jnp.where(jnp.any(σ), w_grid[first_accept_idx], jnp.inf)
```

Let's solve the model at the default parameters:

```
# First, let's solve for the default θ = -1.5
model = create_mccall_model()
c, α, β, ρ, ν, θ, w_grid, z_draws = model

print(f"Solving model with θ = {θ}")
v_star = vfi(model)
σ_star = get_greedy(v_star, model)
w_bar = get_reservation_wage(σ_star, model)
print(f"Reservation wage at default parameters: {w_bar:.4f}")
```

```
Solving model with θ = -1.5
Reservation wage at default parameters: 1.0720
```

58.4.3 How does the reservation wage vary with θ?

Now let's examine how the reservation wage changes as we vary the risk aversion parameter.

```
# Create a grid of θ values (all negative for risk aversion)
θ_grid = jnp.linspace(-3.0, -0.1, 25)

# Define function to compute reservation wage for a single θ value
def compute_res_wage_for_theta(θ):
    """Compute reservation wage for a given θ value"""
    model = create_mccall_model(θ=θ)
    v = vfi(model)
    σ = get_greedy(v, model)
    w_bar = get_reservation_wage(σ, model)
    return w_bar

# Vectorize over θ_grid using vmap
compute_res_wages_vec = jax.vmap(compute_res_wage_for_theta)
reservation_wages = compute_res_wages_vec(θ_grid)

# Plot the results
fig, ax = plt.subplots()
ax.plot(θ_grid, reservation_wages,
        lw=2, marker='o', markersize=4)
ax.set_xlabel(r'$\theta$ (risk aversion parameter)')
ax.set_ylabel('Reservation wage')
ax.axvline(x=-1.5, color='black', ls='--',
           linewidth=1, label=r'Default $\theta = -1.5$')
ax.legend()
plt.tight_layout()
plt.show()
```

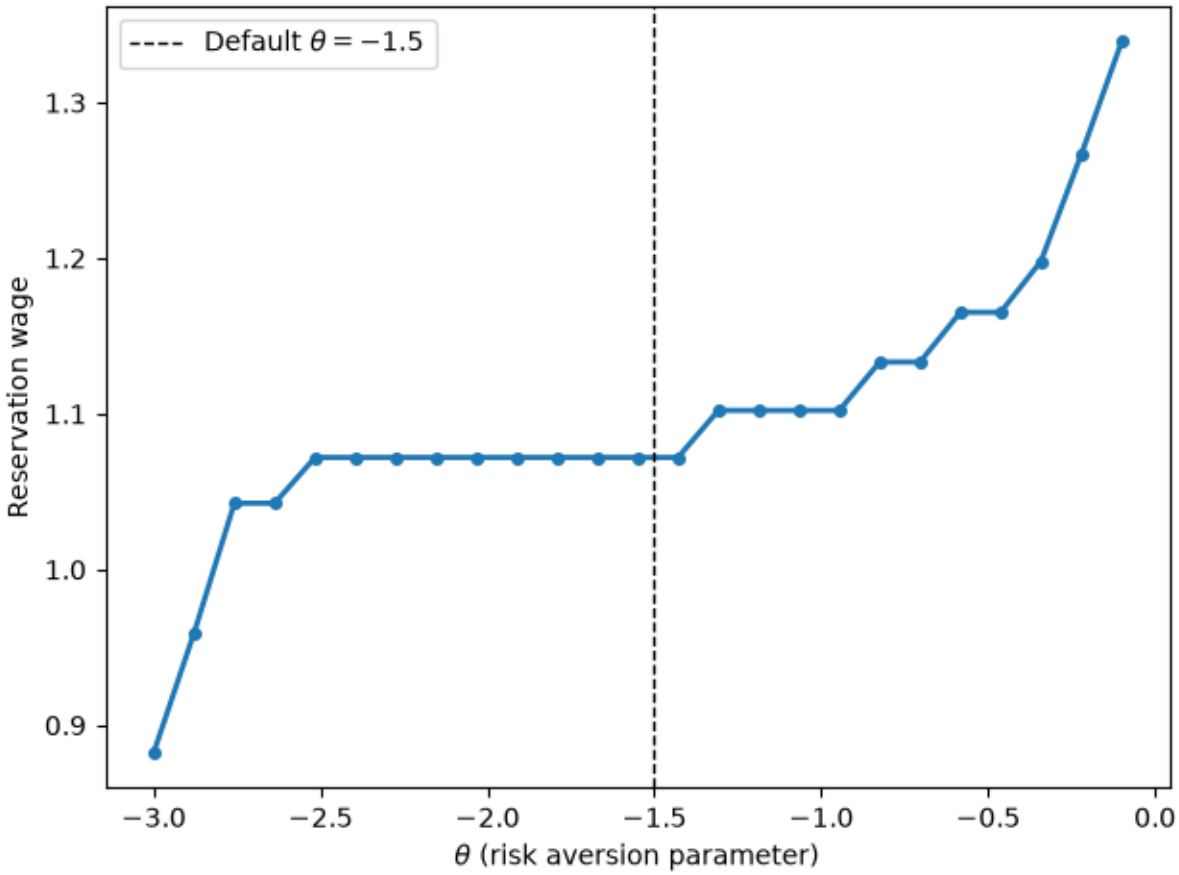


Fig. 58.3: Reservation wage against risk aversion

The reservation wage increases as θ becomes less negative (moves toward zero)

Equivalently, the reservation wage decreases as the agent becomes more risk averse (more negative θ).

The reason is that a more risk-averse agent values the certain income from employment more highly relative to the uncertain future prospects of continued search.

Therefore, they are willing to accept lower wages to escape unemployment.

Exercise 58.4.1

Use simulation to investigate how the long-run unemployment rate varies with θ.

Use the parameters from the previous section, where we studied how the reservation wage varies with θ.

You can use code for simulation from *Job Search IV: Fitted Value Function Iteration*, suitably modified.

Solution

To compute the long-run unemployment rate, we first write a function to update a single agent.

```
@jax.jit
def simulate_single_agent(key, model, w_star, num_periods=200):
    """
    Simulate a single agent for num_periods periods.
    Returns final employment status (1 if employed, 0 if unemployed).
    """
    c, α, β, ρ, ν, θ, w_grid, z_draws = model

    # Start from arbitrary initial conditions
    w = 1.0
    status = 1

    def update(t, loop_state):
        w, status, key = loop_state
        key, k1, k2 = jax.random.split(key, 3)

        # Update wage
        z = jax.random.normal(k2)
        w_new = w**ρ * jnp.exp(ν * z)

        # Employment transitions
        sep_draw = jax.random.uniform(k1)
        becomes_unemployed = sep_draw < α

        # Check if unemployed worker accepts wage
        accepts_job = w >= w_star

        # Update employment status
        new_status = jnp.where(
            status,
            1 - becomes_unemployed,    # employed path
            accepts_job                # unemployed path
        )

        new_wage = jnp.where(
            status,
            jnp.where(becomes_unemployed, w_new, w),   # employed path
```

```
            jnp.where(accepts_job, w, w_new)          # unemployed path
        )

        return (new_wage, new_status, key)

    init_state = (w, status, key)
    final_state = lax.fori_loop(0, num_periods, update, init_state)
    _, final_status, _ = final_state
    return final_status


def compute_unemployment_rate(model, w_star, num_agents=1000, num_periods=200,␣
↪seed=12345):
    """
    Compute unemployment rate via cross-sectional simulation.

    Instead of simulating one agent for a long time series, we simulate
    many agents in parallel for a shorter time period. This is much more
    efficient with JAX parallelization.

    The steady state satisfies:
    - Employed workers lose jobs at rate α
    - Unemployed workers find acceptable jobs at rate (1 - F(w*))

    We simulate num_agents agents for num_periods each, then compute
    the fraction unemployed at the end.
    """
    # Create keys for each agent
    key = jax.random.key(seed)
    keys = jax.random.split(key, num_agents)

    # Vectorize simulation across agents (parallelization!)
    simulate_agents = jax.vmap(
        lambda k: simulate_single_agent(k, model, w_star, num_periods)
    )

    # Run all agents in parallel
    status_cross_section = simulate_agents(keys)

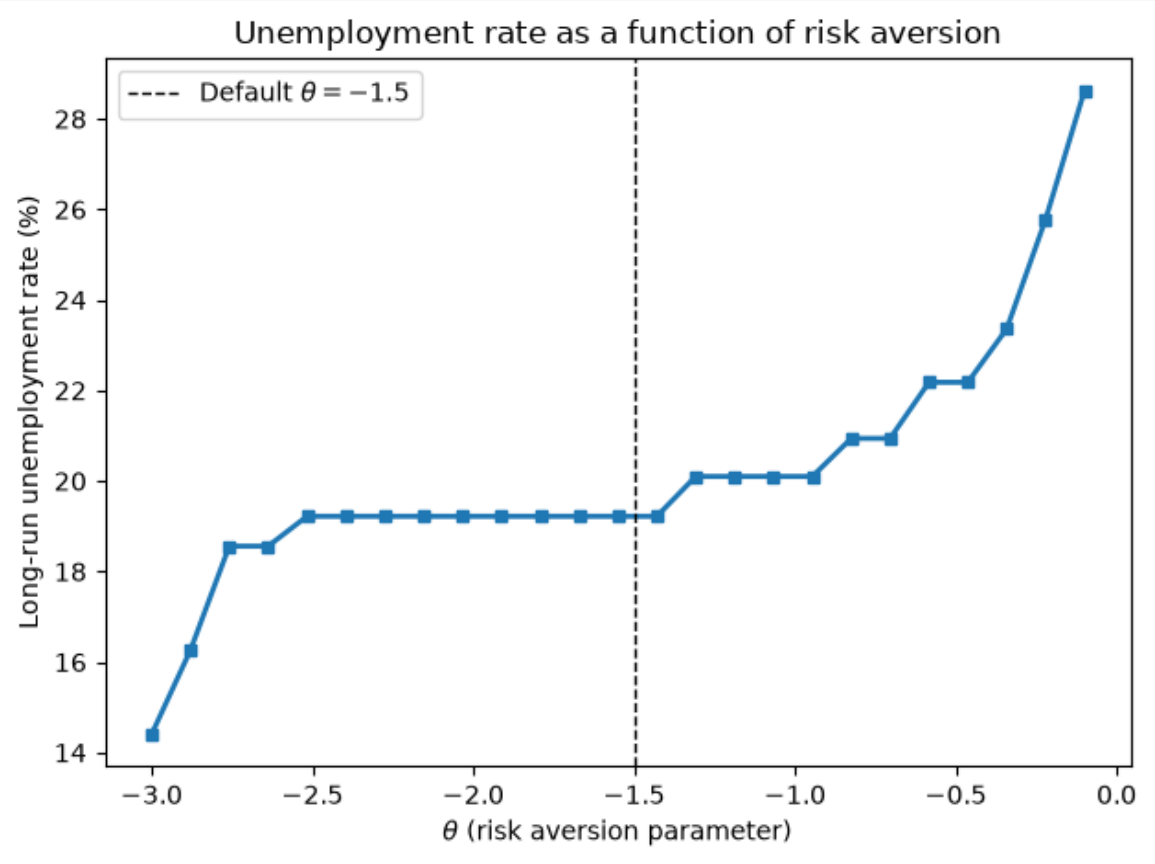
    # Unemployment rate is 1 - mean(status) since status=1 means employed
    unemployment_rate = 1 - jnp.mean(status_cross_section)

    return unemployment_rate


# Define function to compute unemployment rate for a single θ value
def compute_u_rate_for_theta(θ):
    """Compute unemployment rate for a given θ value"""
    model = create_mccall_model(θ=θ)
    v = vfi(model)
    σ = get_greedy(v, model)
    w_star = get_reservation_wage(σ, model)
    u_rate = compute_unemployment_rate(
        model, w_star, num_agents=5000, num_periods=200
    )
    return u_rate
```

```
# Vectorize over θ_grid using vmap
compute_u_rates_vec = jax.vmap(compute_u_rate_for_theta)
unemployment_rates = compute_u_rates_vec(θ_grid)

# Plot the results
fig, ax = plt.subplots()
ax.plot(θ_grid, unemployment_rates * 100,
        lw=2, marker='s', markersize=4)
ax.set_xlabel(r'$\theta$ (risk aversion parameter)')
ax.set_ylabel('Long-run unemployment rate (%)')
ax.set_title('Unemployment rate as a function of risk aversion')
ax.axvline(x=-1.5, color='black', ls='--', linewidth=1, label=r'Default $\theta = -
↪1.5$')
ax.legend()
plt.tight_layout()
plt.show()
```



We see that the unemployment rate decreases as the agent becomes more risk averse (more negative θ).

This is because more risk-averse workers have lower reservation wages, so they accept a wider range of job offers.

As a result, they spend less time unemployed searching for better opportunities

CHAPTER

FIFTYNINE

JOB SEARCH VI: PERSISTENT AND TRANSITORY WAGE SHOCKS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Job Search VI: Persistent and Transitory Wage Shocks*
  - *Overview*
  - *The model*
  - *Implementation*
  - *Unemployment duration*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

59.1 Overview

In this lecture we extend the *McCall job search model* by decomposing wage offers into **persistent** and **transitory** components.

In the *baseline model*, wage offers are IID over time, which is unrealistic.

In *Job Search III*, we introduced correlated wage draws using a Markov chain, but we also added job separation.

Here we take a different approach: we model wage dynamics through an AR(1) process for the persistent component plus a transitory shock, while returning to the assumption that jobs are permanent (as in the *baseline model*).

This persistent-transitory decomposition is:

- More realistic for modeling actual wage processes
- Commonly used in labor economics (see, e.g., [MaCurdy, 1982], [Meghir and Pistaferri, 2004])
- Simple enough to analyze while capturing key features of wage dynamics

By keeping jobs permanent, we can focus on understanding how persistent and transitory wage shocks affect search behavior and reservation wages.

We will solve the model using fitted value function iteration with linear interpolation, as introduced in *Job Search IV*.

We will use the following imports:

```
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
import jax.random
import quantecon as qe
from typing import NamedTuple
```

59.2 The model

Wages at each point in time are given by

$$W_t = \exp(Z_t) + Y_t$$

where

$$Y_t \sim \exp(\mu + s\zeta_t) \quad \text{and} \quad Z_{t+1} = d + \rho Z_t + \sigma \epsilon_{t+1}$$

Here $\{\zeta_t\}$ and $\{\epsilon_t\}$ are both IID and standard normal.

Here $\{Y_t\}$ is a transitory component and $\{Z_t\}$ is persistent.

As before, the worker can either

1. accept an offer and work permanently at that wage, or
2. take unemployment compensation c and wait till next period.

The value function satisfies the Bellman equation

$$v^*(w, z) = \max \left\{ \frac{u(w)}{1-\beta}, u(c) + \beta \, \mathbb{E}_z v^*(w', z') \right\}$$

In this expression, u is a utility function and $\mathbb{E}_z$ is the expectation of next period variables given current z.

The variable z enters as a state in the Bellman equation because its current value helps predict future wages.

59.2.1 A simplification

There is a way that we can reduce dimensionality in this problem, which greatly accelerates computation.

To start, let f^* be the continuation value function, defined by

$$f^*(z) := u(c) + \beta \, \mathbb{E}_z v^*(w', z')$$

The Bellman equation can now be written

$$v^*(w, z) = \max \left\{ \frac{u(w)}{1-\beta}, \, f^*(z) \right\}$$

Combining the last two expressions, we see that the continuation value function satisfies

$$f^*(z) = u(c) + \beta \, \mathbb{E}_z \max \left\{ \frac{u(w')}{1-\beta}, f^*(z') \right\}$$

We'll solve this functional equation for f^* by introducing the operator

$$Qf(z) = u(c) + \beta \, \mathbb{E}_z \max \left\{ \frac{u(w')}{1-\beta}, f(z') \right\}$$

By construction, f^* is a fixed point of Q, in the sense that $Qf^* = f^*$.

Under mild assumptions, it can be shown that Q is a contraction mapping over a suitable space of continuous functions on $\mathbb{R}$.

By Banach's contraction mapping theorem, this means that f^* is the unique fixed point and we can calculate it by iterating with Q from any reasonable initial condition.

Once we have f^*, we can solve the search problem by stopping when the reward for accepting exceeds the continuation value, or

$$\frac{u(w)}{1-\beta} \geq f^*(z)$$

For utility, we take $u(x) = \ln(x)$.

The reservation wage is the wage where equality holds in the last expression.

That is,

$$\bar{w}(z) := \exp(f^*(z)(1-\beta)) \tag{59.1}$$

Our main aim is to solve for the reservation rule and study its properties and implications.

59.3 Implementation

Let f be our initial guess of f^*.

When we iterate, we use the *fitted value function iteration* algorithm.

In particular, f and all subsequent iterates are stored as a vector of values on a grid.

These points are interpolated into a function as required, using piecewise linear interpolation.

The integral in the definition of Qf is calculated by Monte Carlo.

Here's a `NamedTuple` that stores the model parameters and data.

Default parameter values are embedded in the model.

```
class Model(NamedTuple):
    μ: float     # transient shock log mean
    s: float     # transient shock log variance
    d: float     # shift coefficient of persistent state
    ρ: float     # correlation coefficient of persistent state
    σ: float     # state volatility
    β: float     # discount factor
    c: float     # unemployment compensation
    z_grid: jnp.ndarray
    e_draws: jnp.ndarray

def create_job_search_model(μ=0.0, s=1.0, d=0.0, ρ=0.9, σ=0.1, β=0.98, c=5.0,
                            mc_size=1000, grid_size=100, key=jax.random.key(1234)):
    """
    Create a Model with computed grid and draws.
    """
    # Set up grid
    z_mean = d / (1 - ρ)
    z_sd = σ / jnp.sqrt(1 - ρ**2)
    k = 3  # std devs from mean
    a, b = z_mean - k * z_sd, z_mean + k * z_sd
    z_grid = jnp.linspace(a, b, grid_size)

    # Draw and store shocks
    e_draws = jax.random.normal(key, (2, mc_size))

    return Model(μ, s, d, ρ, σ, β, c, z_grid, e_draws)
```

Next, we implement the Q operator.

```
def Q(model, f_in):
    """
    Apply the operator Q.

        * model is an instance of Model
        * f_in is an array that represents f
        * returns Qf

    """
    μ, s, d = model.μ, model.s, model.d
    ρ, σ, β, c = model.ρ, model.σ, model.β, model.c
    z_grid, e_draws = model.z_grid, model.e_draws
    M = e_draws.shape[1]

    def compute_expectation(z):
        def evaluate_shock(e):
            e1, e2 = e[0], e[1]
            z_next = d + ρ * z + σ * e1
            go_val = jnp.interp(z_next, z_grid, f_in)  # f(z')
            y_next = jnp.exp(μ + s * e2)               # y' draw
            w_next = jnp.exp(z_next) + y_next          # w' draw
            stop_val = jnp.log(w_next) / (1 - β)
            return jnp.maximum(stop_val, go_val)

        expectations = jax.vmap(evaluate_shock)(e_draws.T)
        return jnp.mean(expectations)
```

(continues on next page)

(continued from previous page)

```
    expectations = jax.vmap(compute_expectation)(z_grid)
    f_out = jnp.log(c) + β * expectations
    return f_out
```

Here's a function to compute an approximation to the fixed point of Q.

```
@jax.jit
def compute_fixed_point(model, tol=1e-4, max_iter=1000):
    """
    Compute an approximation to the fixed point of Q.
    """

    def cond_fun(loop_state):
        f, i, error = loop_state
        return jnp.logical_and(error > tol, i < max_iter)

    def body_fun(loop_state):
        f, i, error = loop_state
        f_new = Q(model, f)
        error_new = jnp.max(jnp.abs(f_new - f))
        return f_new, i + 1, error_new

    # Initial state
    f_init = jnp.full(len(model.z_grid), jnp.log(model.c))
    init_state = (f_init, 0, tol + 1)

    # Run iteration
    f_final, iterations, final_error = jax.lax.while_loop(
        cond_fun, body_fun, init_state
    )

    return f_final
```

Let's try generating an instance and solving the model.

```
model = create_job_search_model()

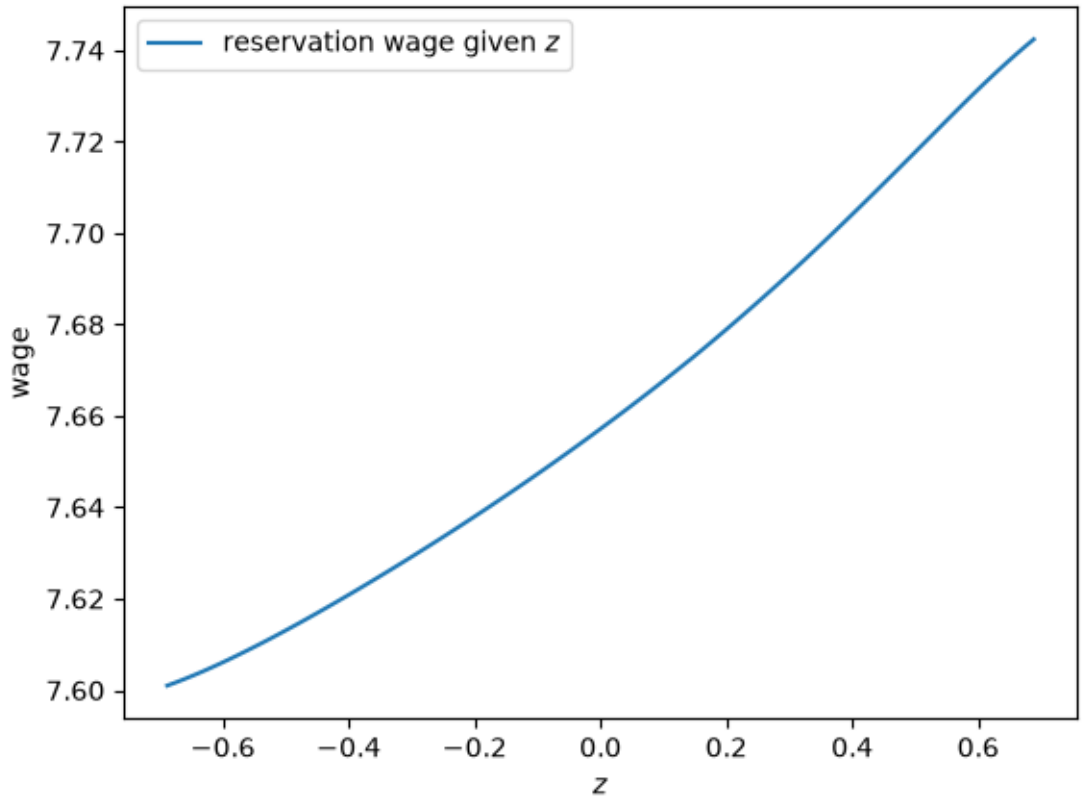
with qe.Timer():
    f_star = compute_fixed_point(model).block_until_ready()
```

```
    0.4899 seconds elapsed
```

Next, we will compute and plot the reservation wage function defined in (59.1).

```
res_wage_function = jnp.exp(f_star * (1 - model.β))

fig, ax = plt.subplots()
ax.plot(
    model.z_grid, res_wage_function, label="reservation wage given $z$"
)
ax.set(xlabel="$z$", ylabel="wage")
ax.legend()
plt.show()
```



Notice that the reservation wage is increasing in the current state z.

This is because a higher state leads the agent to predict higher future wages, increasing the option value of waiting.

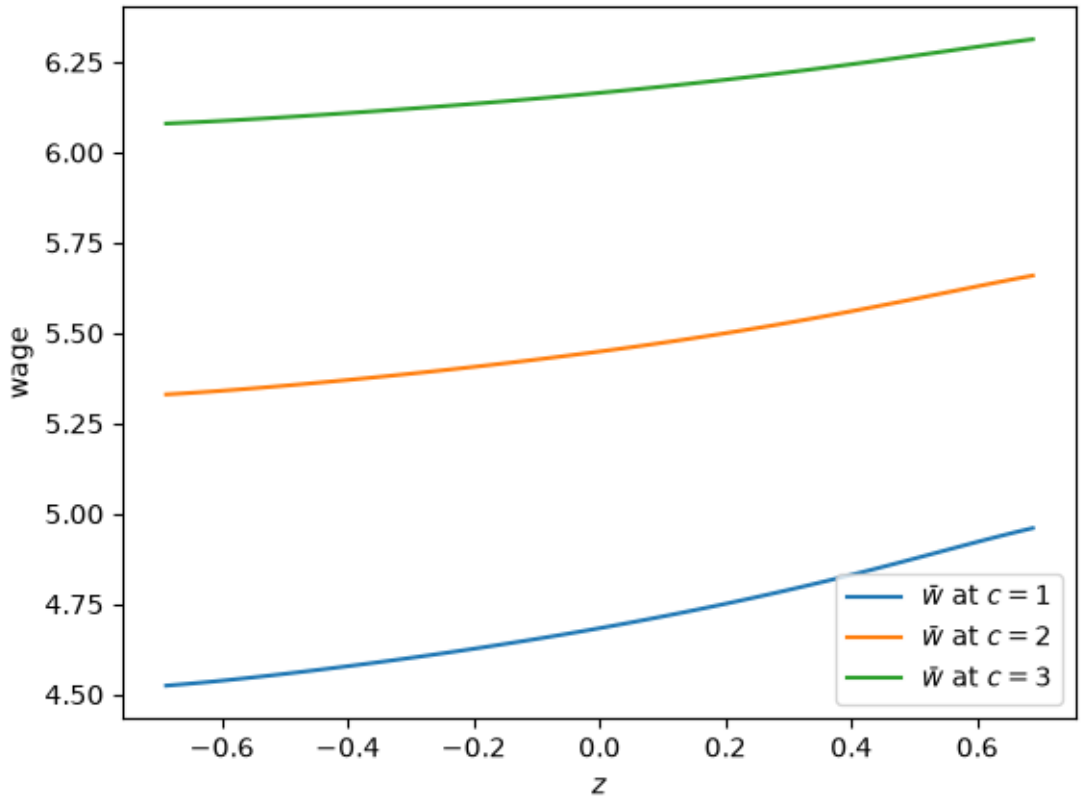
Let's try changing unemployment compensation and looking at its impact on the reservation wage:

```
c_vals = 1, 2, 3

fig, ax = plt.subplots()

for c in c_vals:
    model = create_job_search_model(c=c)
    f_star = compute_fixed_point(model)
    res_wage_function = jnp.exp(f_star * (1 - model.β))
    ax.plot(model.z_grid, res_wage_function,
                label=rf"$\bar w$ at $c = {c}$")

ax.set(xlabel="$z$", ylabel="wage")
ax.legend()
plt.show()
```



As expected, higher unemployment compensation shifts the reservation wage up at all state values.

59.4 Unemployment duration

Next, we study how mean unemployment duration varies with unemployment compensation.

For simplicity, we'll fix the initial state at $Z_0 = 0$.

```
@jax.jit
def draw_duration(key, μ, s, d, ρ, σ, β, z_grid, f_star, t_max=10_000):
    """
    Draw unemployment duration for a single simulation.

    """
    def f_star_function(z):
        return jnp.interp(z, z_grid, f_star)

    def cond_fun(loop_state):
        z, t, unemployed, key = loop_state
        return jnp.logical_and(unemployed, t < t_max)

    def body_fun(loop_state):
        z, t, unemployed, key = loop_state
        key1, key2, key = jax.random.split(key, 3)

        # Draw current wage
        y = jnp.exp(μ + s * jax.random.normal(key1))
```

(continues on next page)

(continued from previous page)

```
        w = jnp.exp(z) + y
        res_wage = jnp.exp(f_star_function(z) * (1 - β))

        # Check if optimal to stop
        accept = w >= res_wage
        τ = jnp.where(accept, t, t_max)

        # Update state if not accepting
        z_new = jnp.where(accept, z,
                          ρ * z + d + σ * jax.random.normal(key2))
        t_new = t + 1
        unemployed_new = jnp.logical_not(accept)

        return z_new, t_new, unemployed_new, key

    # Initial loop_state: (z, t, unemployed, key)
    init_state = (0.0, 0, True, key)
    z_final, t_final, unemployed_final, _ = jax.lax.while_loop(
        cond_fun, body_fun, init_state)

    # Return final time if job found, otherwise t_max
    return jnp.where(unemployed_final, t_max, t_final)


def compute_unemployment_duration(
        model, key=jax.random.key(1234), num_reps=100_000
    ):
    """
    Compute expected unemployment duration.

    """
    f_star = compute_fixed_point(model)
    μ, s, d = model.μ, model.s, model.d
    ρ, σ, β = model.ρ, model.σ, model.β
    z_grid = model.z_grid

    # Generate keys for all simulations
    keys = jax.random.split(key, num_reps)

    # Vectorize over simulations
    τ_vals = jax.vmap(
        lambda k: draw_duration(k, μ, s, d, ρ, σ, β, z_grid, f_star)
    )(keys)

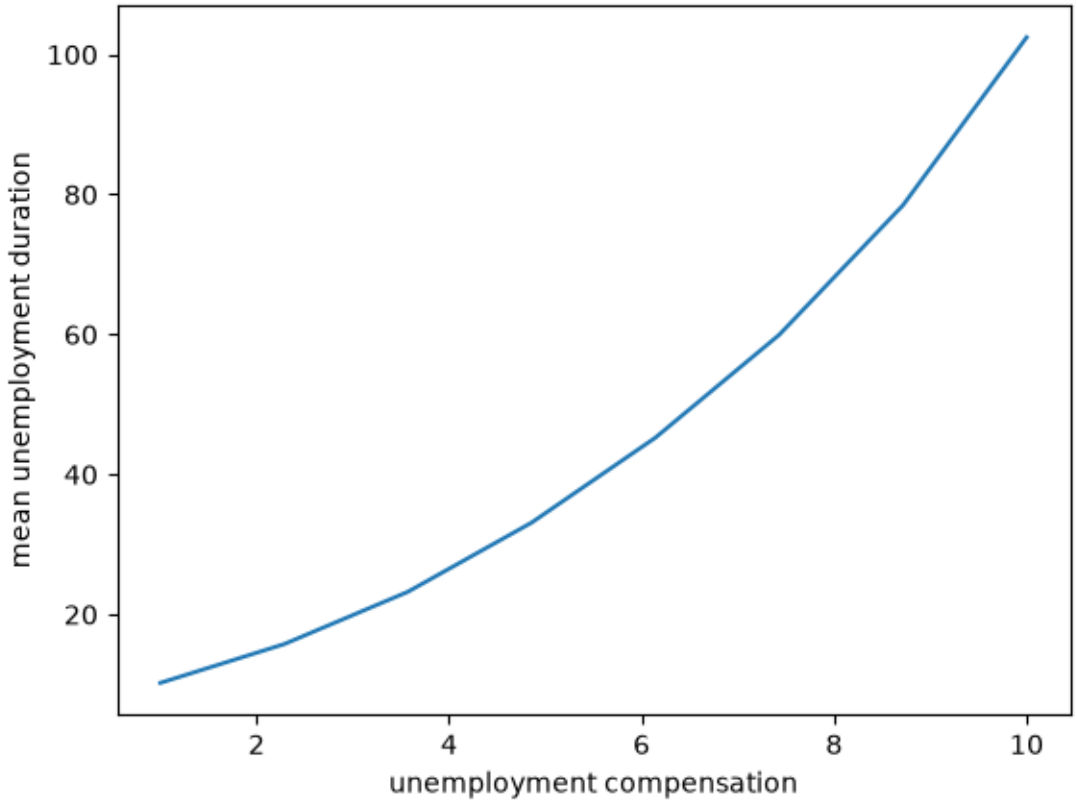
    return jnp.mean(τ_vals)
```

Let's test this out with some possible values for unemployment compensation.

```
c_vals = jnp.linspace(1.0, 10.0, 8)
durations = []
for i, c in enumerate(c_vals):
    model = create_job_search_model(c=c)
    τ = compute_unemployment_duration(model, num_reps=10_000)
    durations.append(τ)
durations = jnp.array(durations)
```

Here is a plot of the results.

```
fig, ax = plt.subplots()
ax.plot(c_vals, durations)
ax.set_xlabel("unemployment compensation")
ax.set_ylabel("mean unemployment duration")
plt.show()
```



Not surprisingly, unemployment duration increases when unemployment compensation is higher.

This is because the value of waiting increases with unemployment compensation.

59.5 Exercises

Exercise 59.5.1

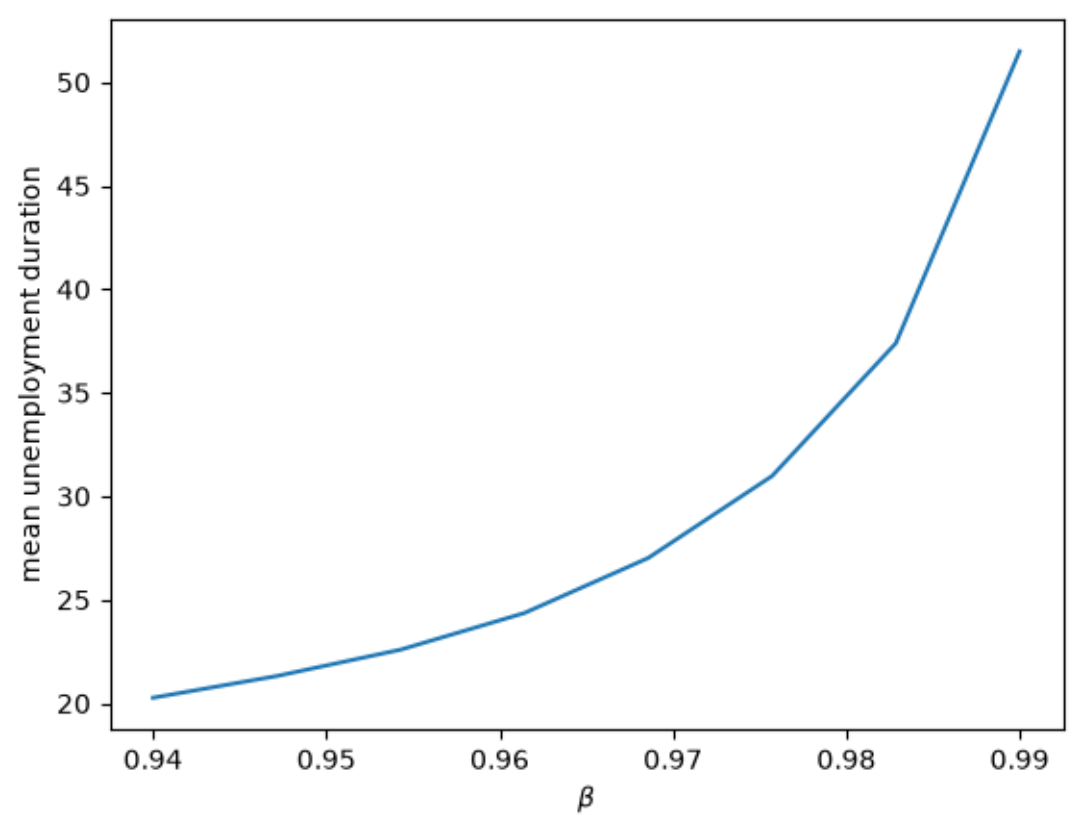
Investigate how mean unemployment duration varies with the discount factor β.

- What is your prior expectation?
- Do your results match up?

Solution

Here is one solution:

```
beta_vals = jnp.linspace(0.94, 0.99, 8)
durations = []
for i, β in enumerate(beta_vals):
    model = create_job_search_model(β=β)
    τ = compute_unemployment_duration(model, num_reps=10_000)
    durations.append(τ)
durations = jnp.array(durations)
fig, ax = plt.subplots()
ax.plot(beta_vals, durations)
ax.set_xlabel(r"$\beta$")
ax.set_ylabel("mean unemployment duration")
plt.show()
```



The figure shows that more patient individuals tend to wait longer before accepting an offer.

CHAPTER

SIXTY

JOB SEARCH VII: MODELING CAREER CHOICE

Contents

- *Job Search VII: Modeling Career Choice*
  - *Overview*
  - *Model*
  - *Implementation*
  - *Exercises*

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

60.1 Overview

Next, we study a computational problem concerning career and job choices.

The model is originally due to Derek Neal [Neal, 1999].

This exposition draws on the presentation in [Ljungqvist and Sargent, 2018], section 6.5.

We begin with some imports:

```
from typing import NamedTuple
```

(continues on next page)

(continued from previous page)

```
import matplotlib.pyplot as plt
from matplotlib import cm
from mpl_toolkits.mplot3d.axes3d import Axes3D
import jax
import jax.numpy as jnp
import jax.random as jr
from quantecon.distributions import BetaBinomial
```

60.1.1 Model Features

- Career and job within career both chosen to maximize expected discounted wage flow.
- Infinite horizon dynamic programming with two state variables.

60.2 Model

In what follows we distinguish between a career and a job, where

- a **career** is understood to be a general field encompassing many possible jobs, and
- a **job** is understood to be a position with a particular firm

For workers, wages can be decomposed into the contribution of job and career

- $w_t = \theta_t + \epsilon_t$, where
  - θ_t is the contribution of career at time t
  - ϵ_t is the contribution of the job at time t

At the start of time t, a worker has the following options

- retain a current (career, job) pair (θ_t, ϵ_t) — referred to hereafter as "stay put"
- retain a current career θ_t but redraw a job ϵ_t — referred to hereafter as "new job"
- redraw both a career θ_t and a job ϵ_t — referred to hereafter as "new life"

Draws of θ and ϵ are independent of each other and past values, with

- $\theta_t \sim F$
- $\epsilon_t \sim G$

Notice that the worker does not have the option to retain a job but redraw a career — starting a new career always requires starting a new job.

A young worker aims to maximize the expected sum of discounted wages

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t w_t \tag{60.1}$$

subject to the choice restrictions specified above.

Let $v(\theta, \epsilon)$ denote the value function, which is the maximum of (60.1) overall feasible (career, job) policies, given the initial state (θ, ϵ).

The value function obeys

$$v(\theta, \epsilon) = \max\{I, II, III\}$$

where

$$I = \theta + \epsilon + \beta v(\theta, \epsilon)$$
$$II = \theta + \int \epsilon' G(d\epsilon') + \beta \int v(\theta, \epsilon') G(d\epsilon')$$
$$III = \int \theta' F(d\theta') + \int \epsilon' G(d\epsilon') + \beta \int \int v(\theta', \epsilon') G(d\epsilon') F(d\theta')$$

Evidently I, II and III correspond to "stay put", "new job" and "new life", respectively.

60.2.1 Parameterization

As in [Ljungqvist and Sargent, 2018], section 6.5, we will focus on a discrete version of the model, parameterized as follows:

- both θ and ϵ take values in the set `jnp.linspace(0, B, grid_size)` — an even grid of points between 0 and B inclusive
- `grid_size = 50`
- `B = 5`
- `β = 0.95`

The distributions F and G are discrete distributions generating draws from the grid points `jnp.linspace(0, B, grid_size)`.

A very useful family of discrete distributions is the Beta-binomial family, with probability mass function

$$p(k \mid n, a, b) = \binom{n}{k} \frac{B(k + a, n - k + b)}{B(a, b)}, \qquad k = 0, \ldots, n$$

Interpretation:

- draw q from a Beta distribution with shape parameters (a, b)
- run n independent binary trials, each with success probability q
- $p(k \mid n, a, b)$ is the probability of k successes in these n trials

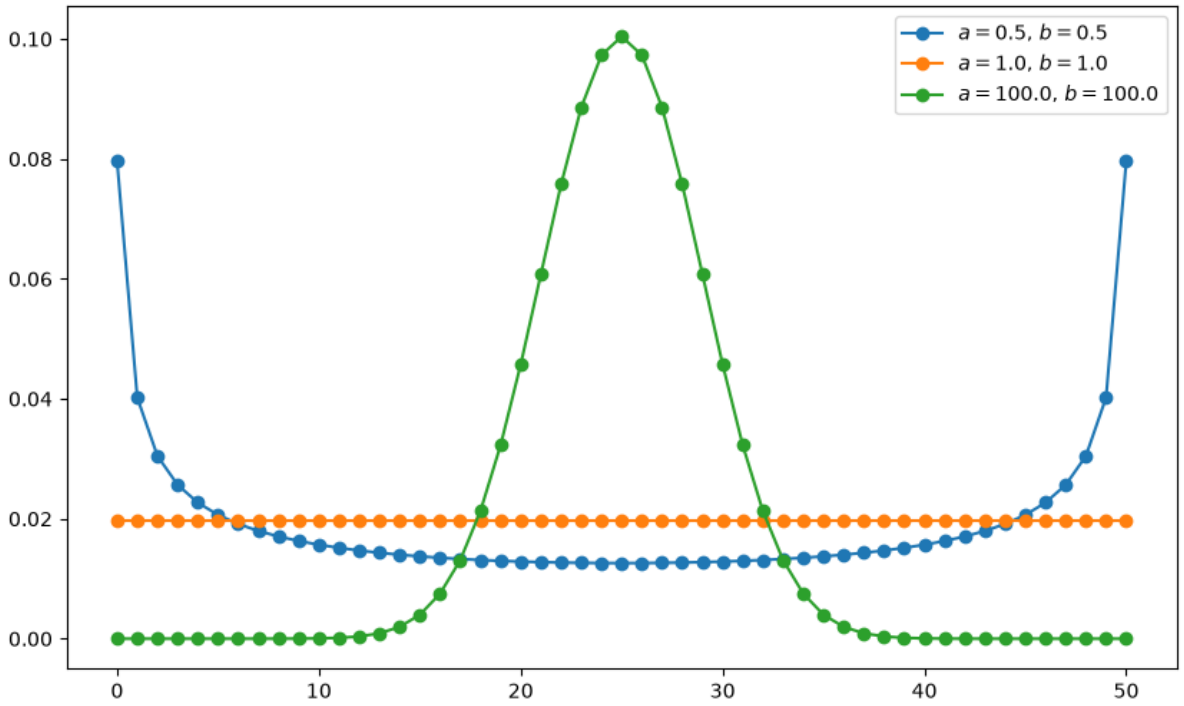
Nice properties:

- very flexible class of distributions, including uniform, symmetric unimodal, etc.
- only three parameters

Here's a figure showing the effect on the pmf of different shape parameters when $n = 50$.

```
n = 50
a_vals = [0.5, 1, 100]
b_vals = [0.5, 1, 100]

fig, ax = plt.subplots(figsize=(10, 6))
for a, b in zip(a_vals, b_vals):
    ab_label = f'$a = {a:.1f}$, $b = {b:.1f}$'
    ax.plot(range(n + 1), BetaBinomial(n, a, b).pdf(), '-o', label=ab_label)
ax.legend()
plt.show()
```



60.3 Implementation

We store the model primitives in a `NamedTuple`, built by a factory function.

```
class CareerWorkerProblem(NamedTuple):
    β: float                    # Discount factor
    θ: jnp.ndarray              # Set of θ values (career)
    ϵ: jnp.ndarray              # Set of ϵ values (job)
    F_probs: jnp.ndarray        # Distribution over new career draws
    G_probs: jnp.ndarray        # Distribution over new job draws
    F_mean: float               # Mean of F
    G_mean: float               # Mean of G


def create_career_worker_problem(B=5.0,          # Upper bound
                                 β=0.95,         # Discount factor
                                 grid_size=50,   # Grid size
                                 F_a=1,
                                 F_b=1,
                                 G_a=1,
                                 G_b=1):
    "Create an instance of the career choice model."
    θ = jnp.linspace(0, B, grid_size)
    ϵ = jnp.linspace(0, B, grid_size)

    F_probs = jnp.array(BetaBinomial(grid_size - 1, F_a, F_b).pdf())
    G_probs = jnp.array(BetaBinomial(grid_size - 1, G_a, G_b).pdf())

    return CareerWorkerProblem(β=β, θ=θ, ϵ=ϵ,
                               F_probs=F_probs, G_probs=G_probs,
```

(continues on next page)

(continued from previous page)

```
                                F_mean=θ @ F_probs, G_mean=ϵ @ G_probs)
```

The Bellman operator is $Tv(\theta, \epsilon) = \max\{I, II, III\}$, where I, II and III are as given in (60.2).

We start by writing those three values for a **single** state (θ_i, ϵ_j), so that the code sits close to the equation.

```
def _B(v, cw, i, j):
    """
    The values of the three options available at state (θ_i, ϵ_j), in the
    order they appear in the Bellman equation.
    """
    stay_put = cw.θ[i] + cw.ϵ[j] + cw.β * v[i, j]                          # I
    new_job = cw.θ[i] + cw.G_mean + cw.β * v[i, :] @ cw.G_probs             # II
    new_life = cw.G_mean + cw.F_mean + cw.β * cw.F_probs @ v @ cw.G_probs # III
    return jnp.array([stay_put, new_job, new_life])
```

Now we evaluate `_B` at every state.

Rather than write two nested loops over i and j, we apply `jax.vmap` twice.

In `in_axes`, a `0` marks the argument being mapped over, while `None` holds an argument fixed.

```
# The argument order of _B is  (v,    cw,   i,    j)
_B_j  = jax.vmap(_B,   in_axes=(None, None, None, 0))   # over j
_B_ij = jax.vmap(_B_j, in_axes=(None, None, 0,    None))  # then over i


@jax.jit
def B(v, cw):
    "Value of each option at each state; shape (grid_size, grid_size, 3)."
    n = len(cw.θ)
    return _B_ij(v, cw, jnp.arange(n), jnp.arange(n))
```

The Bellman operator and the greedy policy are now the maximum and the maximizer of the same array.

```
@jax.jit
def T(v, cw):
    "The Bellman operator."
    return jnp.max(B(v, cw), axis=-1)


@jax.jit
def get_greedy(v, cw):
    "The v-greedy policy, coded as 1 = stay put, 2 = new job, 3 = new life."
    return jnp.argmax(B(v, cw), axis=-1) + 1
```

Lastly, `solve_model` iterates the Bellman operator to find the fixed point.

We use `jax.lax.while_loop` so that the whole iteration compiles into a single operation, and bound the number of steps so the loop always terminates.

```
@jax.jit
def solve_model(cw, tol=1e-4, max_iter=1_000):
    """
    Solve the model by value function iteration.

    Returns the value function, the number of iterations taken and the final
    error, so that the caller can check convergence.
```

(continues on next page)

(continued from previous page)

```
    """
    def condition(loop_state):
        i, v, error = loop_state
        return (error > tol) & (i < max_iter)

    def update(loop_state):
        i, v, error = loop_state
        v_new = T(v, cw)
        return i + 1, v_new, jnp.max(jnp.abs(v_new - v))

    n = len(cw.θ)
    v_init = jnp.full((n, n), 100.0)
    i, v, error = jax.lax.while_loop(condition, update, (0, v_init, tol + 1))
    return v, i, error
```

Note

The grid here is small, and this model would also run perfectly well in NumPy.

We use JAX because the code is almost as readable as the NumPy equivalent while scaling far better — to finer grids, or to richer versions of the model with more state variables, where the same code will make full use of a GPU.

The gain is already visible in Exercise 60.4.2, where we simulate 25,000 independent careers at once.

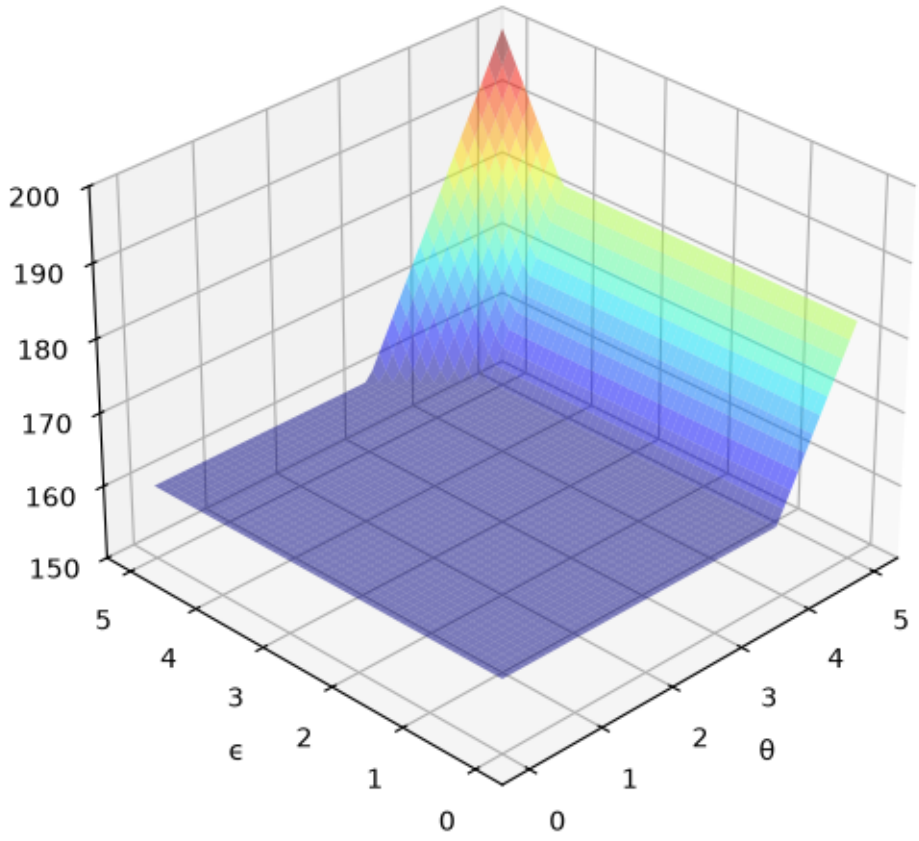
Here's the solution to the model – an approximate value function

```
cw = create_career_worker_problem()
v_star, num_iter, error = solve_model(cw)
greedy_star = get_greedy(v_star, cw)

print(f"Converged in {num_iter} iterations with error {error:.2e}.")

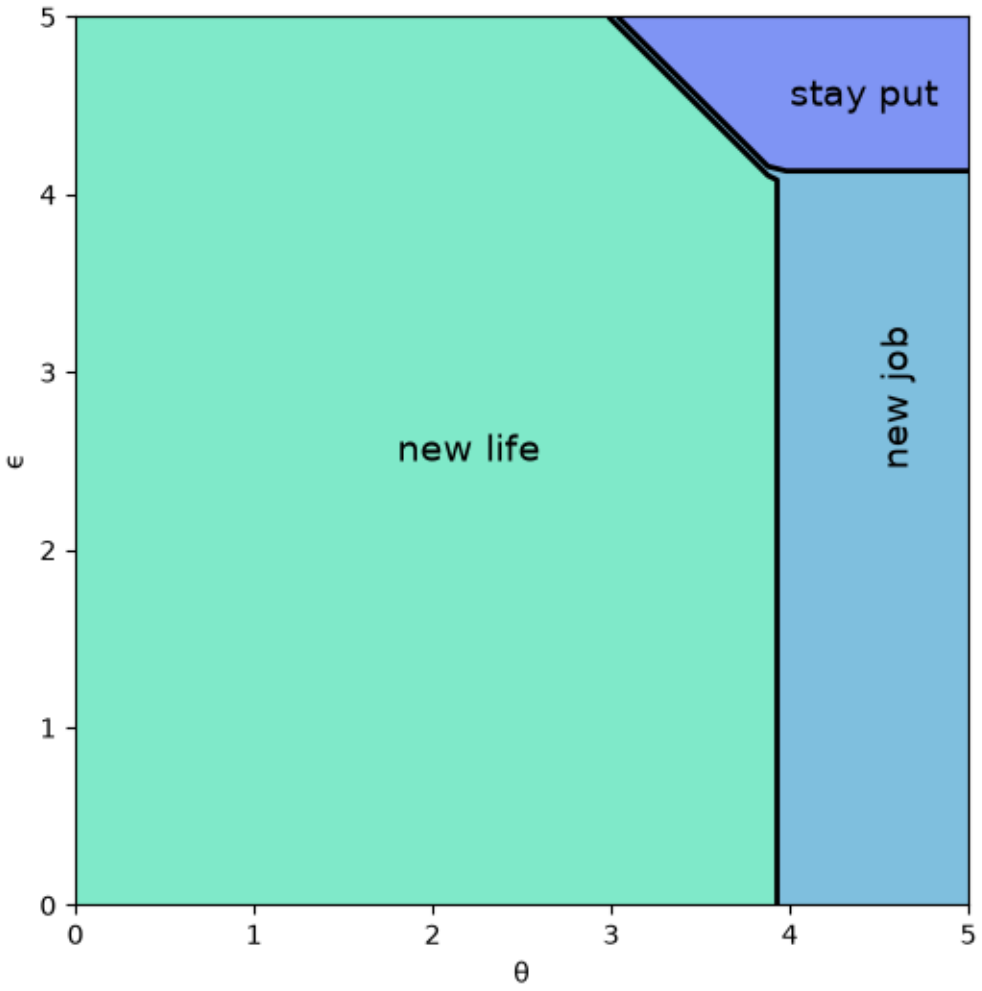
fig = plt.figure(figsize=(8, 6))
ax = fig.add_subplot(111, projection='3d')
tg, eg = jnp.meshgrid(cw.θ, cw.ϵ)
ax.plot_surface(tg,
                eg,
                v_star.T,
                cmap=cm.jet,
                alpha=0.5,
                linewidth=0.25)
ax.set(xlabel='θ', ylabel='ϵ', zlim=(150, 200))
ax.view_init(ax.elev, 225)
plt.show()
```

```
Converged in 213 iterations with error 9.16e-05.
```



And here is the optimal policy

```
fig, ax = plt.subplots(figsize=(6, 6))
tg, eg = jnp.meshgrid(cw.θ, cw.ϵ)
lvls = (0.5, 1.5, 2.5, 3.5)
ax.contourf(tg, eg, greedy_star.T, levels=lvls, cmap=cm.winter, alpha=0.5)
ax.contour(tg, eg, greedy_star.T, colors='k', levels=lvls, linewidths=2)
ax.set(xlabel='θ', ylabel='ϵ')
ax.text(1.8, 2.5, 'new life', fontsize=14)
ax.text(4.5, 2.5, 'new job', fontsize=14, rotation='vertical')
ax.text(4.0, 4.5, 'stay put', fontsize=14)
plt.show()
```



Interpretation:

- If both job and career are poor or mediocre, the worker will experiment with a new job and new career.
- If career is sufficiently good, the worker will hold it and experiment with new jobs until a sufficiently good one is found.
- If both job and career are good, the worker will stay put.

Notice that the worker will always hold on to a sufficiently good career, but not necessarily hold on to even the best paying job.

The reason is that high lifetime wages require both a good job and a good career, but the worker cannot change careers without changing jobs.

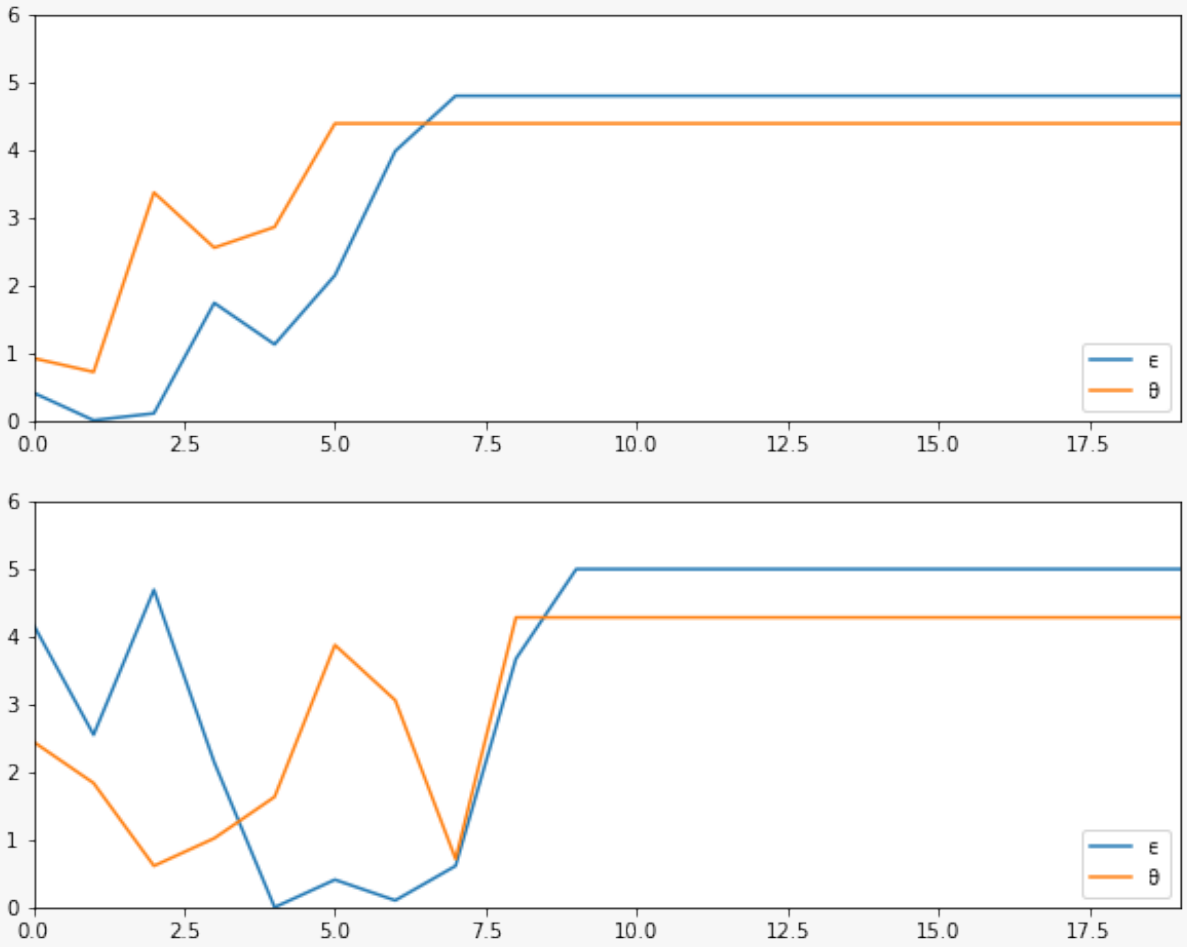
- Sometimes a good job must be sacrificed in order to change to a better career.

60.4 Exercises

Exercise 60.4.1

Using the default parameterization in the function `create_career_worker_problem`, generate and plot typical sample paths for θ and ϵ when the worker follows the optimal policy.

In particular, modulo randomness, reproduce the following figure (where the horizontal axis represents time)



Hint

To draw from F and G, invert their cdfs with `jnp.searchsorted`.

Solution

Simulate job/career paths.

In reading the code, recall that `greedy_star[i, j]` = policy at (θ_i, ϵ_j) = either 1, 2 or 3; meaning 'stay put', 'new job' and 'new life'.

```
def draw(key, cdf):
```

```
    "Draw an index from the distribution with the given cdf."
    return jnp.searchsorted(cdf, jr.uniform(key), side="right")


def simulate_path(cw, greedy_star, key, t=20):
    "Simulate a career/job path of length t under the greedy policy."
    F_cdf, G_cdf = jnp.cumsum(cw.F_probs), jnp.cumsum(cw.G_probs)

    def update(state, key):
        i, j = state
        action = greedy_star[i, j]
        key_F, key_G = jr.split(key)
        # Career changes only under 'new life'; the job changes unless we stay put
        i_new = jnp.where(action == 3, draw(key_F, F_cdf), i)
        j_new = jnp.where(action == 1, j, draw(key_G, G_cdf))
        return (i_new, j_new), (i_new, j_new)

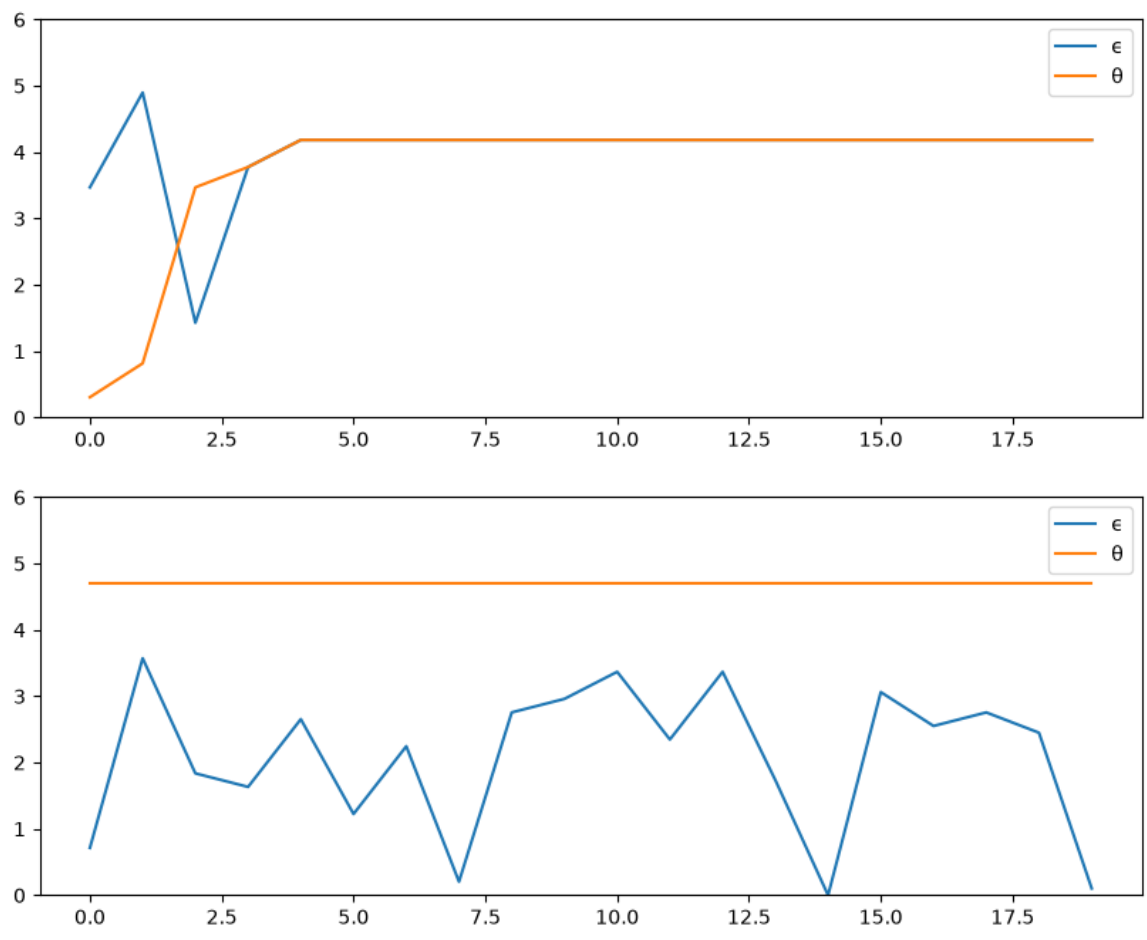
    _, (i_path, j_path) = jax.lax.scan(update, (0, 0), jr.split(key, t))
    return cw.θ[i_path], cw.ϵ[j_path]


cw = create_career_worker_problem()
v_star, _, _ = solve_model(cw)
greedy_star = get_greedy(v_star, cw)

key = jr.key(42)
fig, axes = plt.subplots(2, 1, figsize=(10, 8))

for ax in axes:
    key, subkey = jr.split(key)
    θ_path, ϵ_path = simulate_path(cw, greedy_star, subkey)
    ax.plot(ϵ_path, label='ϵ')
    ax.plot(θ_path, label='θ')
    ax.set_ylim(0, 6)
    ax.legend()

plt.show()
```



Exercise 60.4.2

Let's now consider how long it takes for the worker to settle down to a permanent job, given a starting point of $(\theta, \epsilon) = (0, 0)$.

In other words, we want to study the distribution of the random variable

$$T^* := \text{the first point in time from which the worker's job no longer changes}$$

Evidently, the worker's job becomes permanent if and only if (θ_t, ϵ_t) enters the "stay put" region of (θ, ϵ) space.

Letting S denote this region, T^* can be expressed as the first passage time to S under the optimal policy:

$$T^* := \inf\{t \geq 0 \,|\, (\theta_t, \epsilon_t) \in S\}$$

Collect 25,000 draws of this random variable and compute the median (which should be about 7).

Repeat the exercise with $\beta = 0.99$ and interpret the change.

Solution

The median for the original parameterization can be computed as follows.

Each simulation is an independent sequential search, so we write one with `jax.lax.while_loop` and then run 25,000 of them at once with `jax.vmap`.

```
def passage_time(cw, greedy_star, key, max_t=1_000):
    "Time until the worker first chooses to stay put."
    F_cdf, G_cdf = jnp.cumsum(cw.F_probs), jnp.cumsum(cw.G_probs)

    def condition(state):
        i, j, t, key = state
        return (greedy_star[i, j] != 1) & (t < max_t)

    def update(state):
        i, j, t, key = state
        action = greedy_star[i, j]
        key, key_F, key_G = jr.split(key, 3)
        i_new = jnp.where(action == 3, draw(key_F, F_cdf), i)
        j_new = jnp.where(action == 1, j, draw(key_G, G_cdf))
        return i_new, j_new, t + 1, key

    _, _, t, _ = jax.lax.while_loop(condition, update, (0, 0, 0, key))
    return t


@jax.jit
def median_passage_time(cw, greedy_star, key, M=25_000):
    "Median time to settle down, over M independent simulations."
    keys = jr.split(key, M)
    times = jax.vmap(passage_time, in_axes=(None, None, 0))(cw, greedy_star, keys)
    return jnp.median(times)


median_passage_time(cw, greedy_star, jr.key(42))
```

```
Array(7., dtype=float32)
```

To compute the median with $\beta = 0.99$ instead of the default value $\beta = 0.95$, we create a new instance and solve it again.

```
cw_patient = create_career_worker_problem(β=0.99)
v_patient, _, _ = solve_model(cw_patient)
greedy_patient = get_greedy(v_patient, cw_patient)

median_passage_time(cw_patient, greedy_patient, jr.key(42))
```

```
Array(14., dtype=float32)
```

The medians are subject to randomness but should be about 7 and 14 respectively.

Not surprisingly, more patient workers will wait longer to settle down to their final job.

Exercise 60.4.3

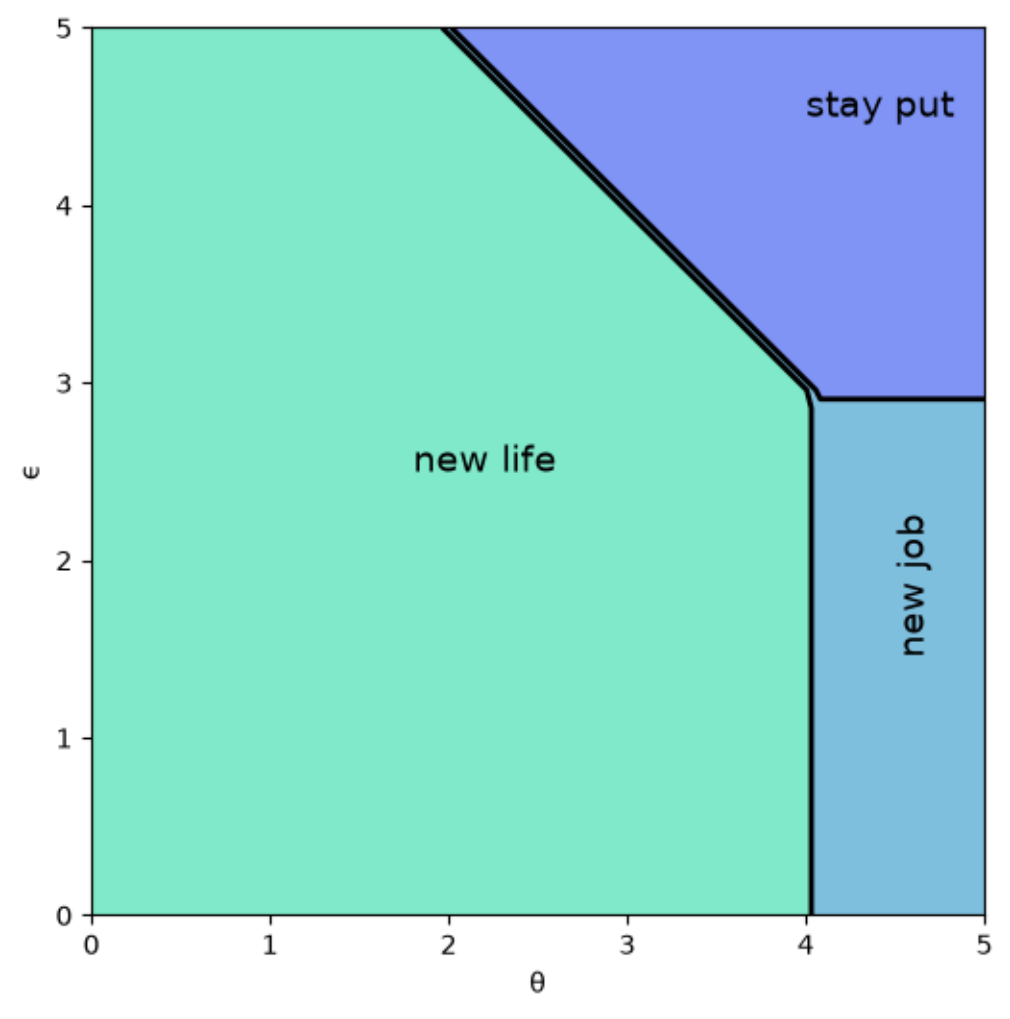
Set the parameterization to `G_a = G_b = 100` and generate a new optimal policy figure – interpret.

Solution

Here is one solution

```
cw = create_career_worker_problem(G_a=100, G_b=100)
v_star, _, _ = solve_model(cw)
greedy_star = get_greedy(v_star, cw)

fig, ax = plt.subplots(figsize=(6, 6))
tg, eg = jnp.meshgrid(cw.θ, cw.ϵ)
lvls = (0.5, 1.5, 2.5, 3.5)
ax.contourf(tg, eg, greedy_star.T, levels=lvls, cmap=cm.winter, alpha=0.5)
ax.contour(tg, eg, greedy_star.T, colors='k', levels=lvls, linewidths=2)
ax.set(xlabel='θ', ylabel='ϵ')
ax.text(1.8, 2.5, 'new life', fontsize=14)
ax.text(4.5, 1.5, 'new job', fontsize=14, rotation='vertical')
ax.text(4.0, 4.5, 'stay put', fontsize=14)
plt.show()
```



In the new figure, you see that the region for which the worker stays put has grown because the distribution for ϵ has become more concentrated around the mean, making high-paying jobs less realistic.

CHAPTER

SIXTYONE

JOB SEARCH VIII: ON-THE-JOB SEARCH

Contents

- *Job Search VIII: On-the-Job Search*
  - *Overview*
  - *Model*
  - *Implementation*
  - *Solving for Policies*
  - *Exercises*

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install jax
```

61.1 Overview

In this section, we solve a simple on-the-job search model

- based on [Ljungqvist and Sargent, 2018], exercise 6.18, and [Jovanovic, 1979]

Let's start with some imports:

```
from typing import NamedTuple
```

(continues on next page)

(continued from previous page)

```
import matplotlib.pyplot as plt
import scipy.stats as stats
import jax
import jax.numpy as jnp
import jax.random as jr
```

61.1.1 Model Features

- job-specific human capital accumulation combined with on-the-job search
- infinite-horizon dynamic programming with one state variable and two controls

61.2 Model

Let x_t denote the time-t job-specific human capital of a worker employed at a given firm and let w_t denote current wages.

Let $w_t = x_t(1 - s_t - \phi_t)$, where

- ϕ_t is investment in job-specific human capital for the current role and
- s_t is search effort, devoted to obtaining new offers from other firms.

For as long as the worker remains in the current job, evolution of $\{x_t\}$ is given by $x_{t+1} = g(x_t, \phi_t)$.

When search effort at t is s_t, the worker receives a new job offer with probability $\pi(s_t) \in [0, 1]$.

The value of the offer, measured in job-specific human capital, is u_{t+1}, where $\{u_t\}$ is IID with common distribution f.

The worker can reject the current offer and continue with existing job.

Hence $x_{t+1} = u_{t+1}$ if he/she accepts and $x_{t+1} = g(x_t, \phi_t)$ otherwise.

Let $b_{t+1} \in \{0, 1\}$ be a binary random variable, where $b_{t+1} = 1$ indicates that the worker receives an offer at the end of time t.

We can write

$$x_{t+1} = (1 - b_{t+1})g(x_t, \phi_t) + b_{t+1} \max\{g(x_t, \phi_t), u_{t+1}\} \tag{61.1}$$

Agent's objective: maximize expected discounted sum of wages via controls $\{s_t\}$ and $\{\phi_t\}$.

Taking the expectation of $v(x_{t+1})$ and using (61.1), the Bellman equation for this problem can be written as

$$v(x) = \max_{s+\phi \leq 1} \left\{ x(1 - s - \phi) + \beta(1 - \pi(s))v[g(x, \phi)] + \beta\pi(s) \int v[g(x, \phi) \vee u]f(du) \right\} \tag{61.2}$$

Here nonnegativity of s and ϕ is understood, while $a \vee b := \max\{a, b\}$.

61.2.1 Parameterization

In the implementation below, we will focus on the parameterization

$$g(x, \phi) = A(x\phi)^\alpha, \quad \pi(s) = \sqrt{s} \quad \text{and} \quad f = \text{Beta}(2, 2)$$

with default parameter values

- $A = 1.4$
- $\alpha = 0.6$
- $\beta = 0.96$

The Beta$(2, 2)$ distribution is supported on $(0, 1)$ - it has a unimodal, symmetric density peaked at 0.5.

61.2.2 Back-of-the-Envelope Calculations

Before we solve the model, let's make some quick calculations that provide intuition on what the solution should look like.

To begin, observe that the worker has two instruments to build capital and hence wages:

1. invest in capital specific to the current job via ϕ
2. search for a new job with better job-specific capital match via s

Since wages are $x(1 - s - \phi)$, marginal cost of investment via either ϕ or s is identical.

Our risk-neutral worker should focus on whatever instrument has the highest expected return.

The relative expected return will depend on x.

For example, suppose first that $x = 0.05$

- If $s = 1$ and $\phi = 0$, then since $g(x, \phi) = 0$, taking expectations of (61.1) gives expected next period capital equal to $\pi(s)\mathbb{E}u = \mathbb{E}u = 0.5$.
- If $s = 0$ and $\phi = 1$, then next period capital is $g(x, \phi) = g(0.05, 1) \approx 0.23$.

Both rates of return are good, but the return from search is better.

Next, suppose that $x = 0.4$

- If $s = 1$ and $\phi = 0$, then expected next period capital is again 0.5
- If $s = 0$ and $\phi = 1$, then $g(x, \phi) = g(0.4, 1) \approx 0.8$

Return from investment via ϕ dominates expected return from search.

Combining these observations gives us two informal predictions:

1. At any given state x, the two controls ϕ and s will function primarily as substitutes — worker will focus on whichever instrument has the higher expected return.
2. For sufficiently small x, search will be preferable to investment in job-specific human capital. For larger x, the reverse will be true.

Now let's turn to implementation, and see if we can match our predictions.

61.3 Implementation

We solve the model with JAX, using a `NamedTuple` to hold the parameters and grids.

```
class JVWorker(NamedTuple):
    A: float                # Scale parameter in g
    α: float                # Curvature parameter in g
    β: float                # Discount factor
    x_grid: jnp.ndarray     # Grid of human capital values
    s_grid: jnp.ndarray     # Grid of search effort values
    ϕ_grid: jnp.ndarray     # Grid of investment values
    f_rvs: jnp.ndarray      # Draws from f, for Monte Carlo integration


def create_jv_worker(A=1.4,                # Scale parameter in g
                     α=0.6,                # Curvature parameter in g
                     β=0.96,               # Discount factor
                     a=2,                  # Parameter of f
                     b=2,                  # Parameter of f
                     grid_size=50,         # Size of the state grid
                     mc_size=100,          # Number of draws from f
                     search_grid_size=15,  # Size of each action grid
                     ε=1e-4,
                     seed=1234):
    """
    Create an instance of the on-the-job search model.
    """
    f_rvs = jr.beta(jr.key(seed), a, b, (mc_size,))

    # Max of grid is the max of a large quantile value for f and the
    # fixed point y = g(y, 1)
    grid_max = max(A**(1 / (1 - α)), stats.beta(a, b).ppf(1 - ε))

    x_grid = jnp.linspace(ε, grid_max, grid_size)
    s_grid = jnp.linspace(ε, 1, search_grid_size)
    ϕ_grid = jnp.linspace(ε, 1, search_grid_size)

    return JVWorker(A=A, α=α, β=β, x_grid=x_grid, s_grid=s_grid,
                    ϕ_grid=ϕ_grid, f_rvs=f_rvs)
```

Here are the transition function g and the offer probability π.

```
@jax.jit
def g(jv, x, ϕ):
    "Transition function for job-specific human capital."
    return jv.A * (x * ϕ)**jv.α


@jax.jit
def π(s):
    "Probability of receiving an offer when search effort is s."
    return jnp.sqrt(s)
```

Next we write the right-hand side of the Bellman equation (61.2), before maximization:

$$B(x, s, \phi) := x(1 - s - \phi) + \beta(1 - \pi(s))v[g(x, \phi)] + \beta\pi(s) \int v[g(x, \phi) \vee u]f(du) \tag{61.3}$$

We represent v by an array giving its values on `x_grid`, and recover a function from it by linear interpolation.

The integral is replaced by a Monte Carlo average over the draws in `f_rvs`.

The function below is written for a **single** state x and a **single** action pair (s, ϕ) — so it reads much like (61.3) itself.

```
def _B(v, jv, x, s, ϕ):
    """
    The right-hand side of the Bellman equation before maximization, for one
    state x and one action pair (s, ϕ).

    Infeasible pairs, where s + ϕ > 1, are given value -∞ so that they are
    never selected by the maximization step.
    """
    v_func = lambda z: jnp.interp(z, jv.x_grid, v)
    gxϕ = g(jv, x, ϕ)

    # Monte Carlo estimate of ∫ v[g(x, ϕ) ∨ u] f(du)
    integral = jnp.mean(v_func(jnp.maximum(gxϕ, jv.f_rvs)))

    q = π(s) * integral + (1 - π(s)) * v_func(gxϕ)
    return jnp.where(s + ϕ <= 1, x * (1 - s - ϕ) + jv.β * q, -jnp.inf)
```

Now we evaluate `_B` at every combination of state and action.

Rather than write three nested loops, we apply `jax.vmap` three times.

Each application vectorizes over one argument, so the stack below plays the role of a triple loop — but the whole thing compiles to code that runs in parallel.

In `in_axes`, a `0` marks the argument being mapped over, while `None` holds an argument fixed.

```
# The argument order of _B is     (v,     jv,    x,     s,     ϕ)
_B_ϕ   = jax.vmap(_B,    in_axes=(None, None, None, None, 0))     # over ϕ
_B_sϕ  = jax.vmap(_B_ϕ,  in_axes=(None, None, None, 0,    None))  # then over s
_B_xsϕ = jax.vmap(_B_sϕ, in_axes=(None, None, 0,    None, None))  # then over x
```

The result is a fully vectorized version of B.

```
@jax.jit
def B(v, jv):
    """
    Evaluate B at every (state, action) combination.

    Returns an array of shape (len(x_grid), len(s_grid), len(ϕ_grid)) where
    entry [i, j, k] holds the value of choosing (s_j, ϕ_k) in state x_i.
    """
    return _B_xsϕ(v, jv, jv.x_grid, jv.s_grid, jv.ϕ_grid)
```

With `B` in hand, the Bellman operator and the greedy policy are both one-liners — we maximize over the two action axes, taking the maximum in one case and the maximizer in the other.

```
@jax.jit
def T(v, jv):
    "The Bellman operator."
    return jnp.max(B(v, jv), axis=(1, 2))


@jax.jit
def get_greedy(v, jv):
    "Compute the v-greedy policy, returned as a pair (s_policy, ϕ_policy)."
```

(continues on next page)

(continued from previous page)

```
    vals = B(v, jv)

    # Flatten the two action axes so that a single argmax picks out the best
    # pair at each state, then convert the flat index back to a (s, ϕ) pair
    n_s, n_ϕ = len(jv.s_grid), len(jv.ϕ_grid)
    best = jnp.argmax(vals.reshape(len(jv.x_grid), n_s * n_ϕ), axis=1)
    j, k = jnp.unravel_index(best, (n_s, n_ϕ))

    return jv.s_grid[j], jv.ϕ_grid[k]
```

To solve the model we iterate T to convergence.

We use `jax.lax.while_loop` so that the entire iteration compiles into a single operation, and bound the number of steps so that the loop always terminates.

```
@jax.jit
def solve_model(jv, tol=1e-4, max_iter=1_000):
    """
    Solve the model by value function iteration.

    Returns the value function, the number of iterations taken, and the final
    error, so that the caller can check convergence.
    """
    def condition(loop_state):
        i, v, error = loop_state
        return (error > tol) & (i < max_iter)

    def update(loop_state):
        i, v, error = loop_state
        v_new = T(v, jv)
        return i + 1, v_new, jnp.max(jnp.abs(v_new - v))

    v_init = jv.x_grid * 0.5
    i, v, error = jax.lax.while_loop(condition, update, (0, v_init, tol + 1))
    return v, i, error
```

Note

The grids here are small, and this model would also run perfectly well in NumPy.

We use JAX because the code is almost as readable as the NumPy equivalent, while scaling far better — to finer grids, or to richer versions of the model with additional state variables, where the same code will make full use of a GPU.

61.4 Solving for Policies

Let's generate the optimal policies and see what they look like.

```
jv = create_jv_worker()
v_star, num_iter, error = solve_model(jv)
s_star, ϕ_star = get_greedy(v_star, jv)

print(f"Converged in {num_iter} iterations with error {error:.2e}.")
```

```
Converged in 205 iterations with error 9.92e-05.
```

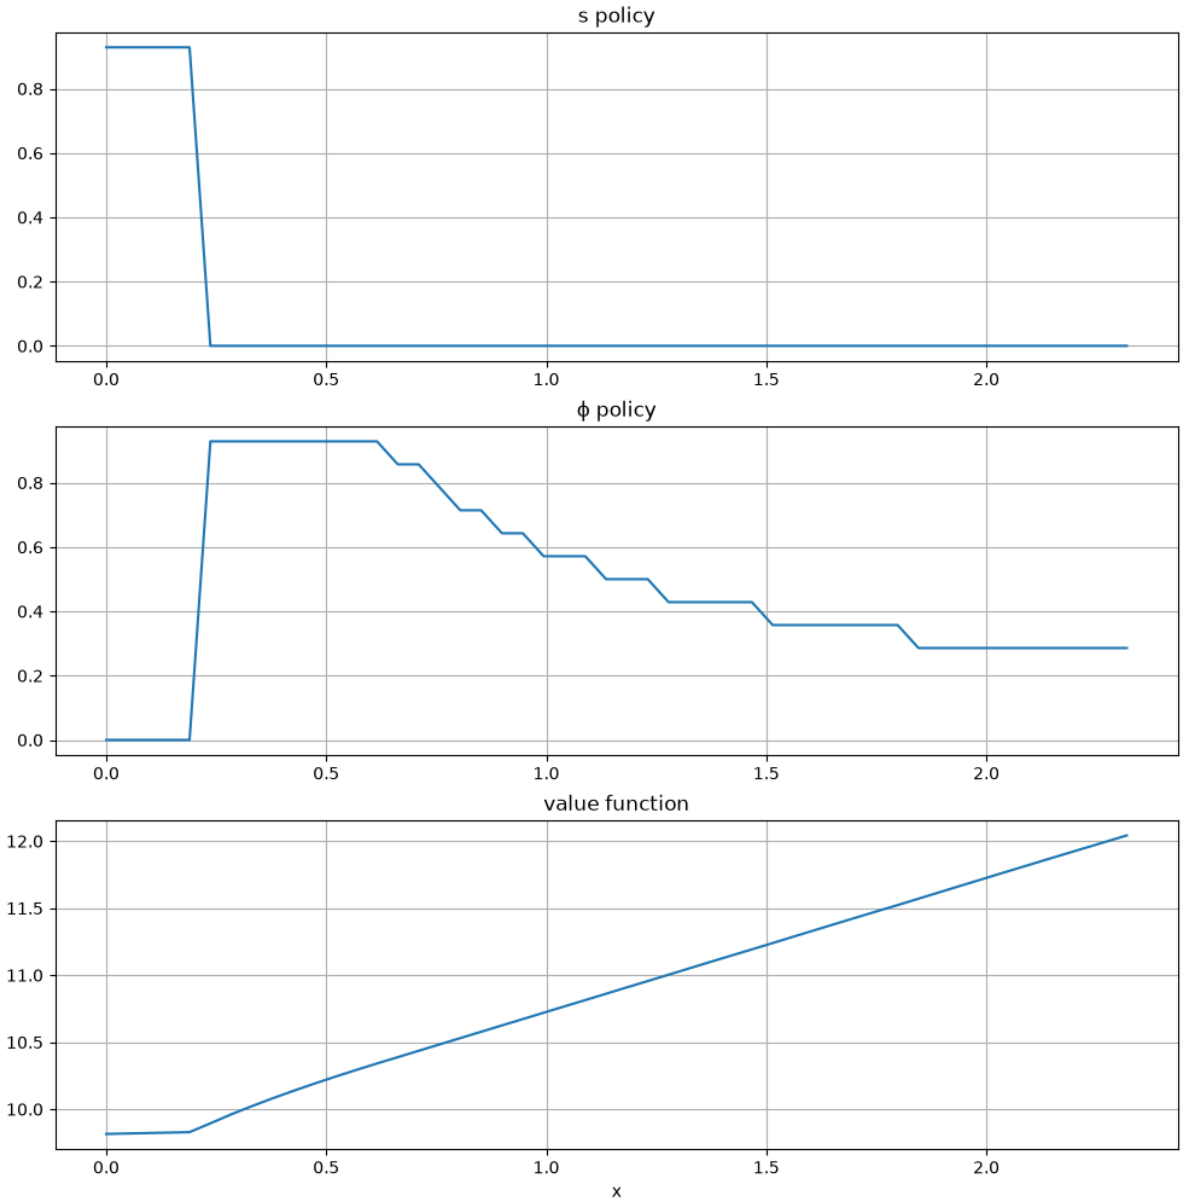
Here are the plots:

```
plots = [s_star, ϕ_star, v_star]
titles = ["s policy", "ϕ policy",  "value function"]

fig, axes = plt.subplots(3, 1, figsize=(12, 12))

for ax, plot, title in zip(axes, plots, titles):
    ax.plot(jv.x_grid, plot)
    ax.set(title=title)
    ax.grid()

axes[-1].set_xlabel("x")
plt.show()
```



The horizontal axis is the state x, while the vertical axis gives $s(x)$ and $\phi(x)$.

Overall, the policies match well with our predictions from *above*

- Worker switches from one investment strategy to the other depending on relative return.
- For low values of x, the best option is to search for a new job.
- Once x is larger, worker does better by investing in human capital specific to the current position.

61.5 Exercises

Exercise 61.5.1

Let's look at the dynamics for the state process $\{x_t\}$ associated with these policies.

The dynamics are given by (61.1) when ϕ_t and s_t are chosen according to the optimal policies, and $\mathbb{P}\{b_{t+1} = 1\} = \pi(s_t)$.

Since the dynamics are random, analysis is a bit subtle.

One way to do it is to plot, for each x in a relatively fine grid called `plot_grid`, a large number K of realizations of x_{t+1} given $x_t = x$.

Plot this with one dot for each realization, in the form of a 45 degree diagram, setting

```
jv = create_jv_worker(grid_size=25, mc_size=50)
plot_grid_max, plot_grid_size = 1.2, 100
plot_grid = jnp.linspace(0, plot_grid_max, plot_grid_size)
fig, ax = plt.subplots()
ax.set_xlim(0, plot_grid_max)
ax.set_ylim(0, plot_grid_max)
```

By examining the plot, argue that under the optimal policies, the state x_t will converge to a constant value $\bar{x}$ close to unity.

Argue that at the steady state, $s_t \approx 0$ and $\phi_t \approx 0.6$.

Solution

Here's code to produce the 45 degree diagram.

Note that we draw all of the realizations at once, rather than looping over states and draws.

```
jv = create_jv_worker(grid_size=25, mc_size=50)
v_star, _, _ = solve_model(jv)
s_policy, ϕ_policy = get_greedy(v_star, jv)

# Turn the policy function arrays into actual functions
s = lambda y: jnp.interp(y, jv.x_grid, s_policy)
ϕ = lambda y: jnp.interp(y, jv.x_grid, ϕ_policy)

plot_grid_max, plot_grid_size = 1.2, 100
plot_grid = jnp.linspace(0, plot_grid_max, plot_grid_size)


@jax.jit
def simulate_next(key, plot_grid):
    """
    Draw realizations of next period capital for every x in plot_grid,
    following the law of motion for x_{t+1} given above.  Returns an array of␣
↪shape (len(plot_grid), mc_size).
    """
    K = len(jv.f_rvs)
    gxϕ = g(jv, plot_grid, ϕ(plot_grid))[:, jnp.newaxis]    # Shape (n, 1)
    u = jv.f_rvs[jnp.newaxis, :]                            # Shape (1, K)

    # An offer arrives with probability π(s(x)), independently across draws
```

```
    b = jr.uniform(key, (len(plot_grid), K)) < π(s(plot_grid))[:, jnp.newaxis]

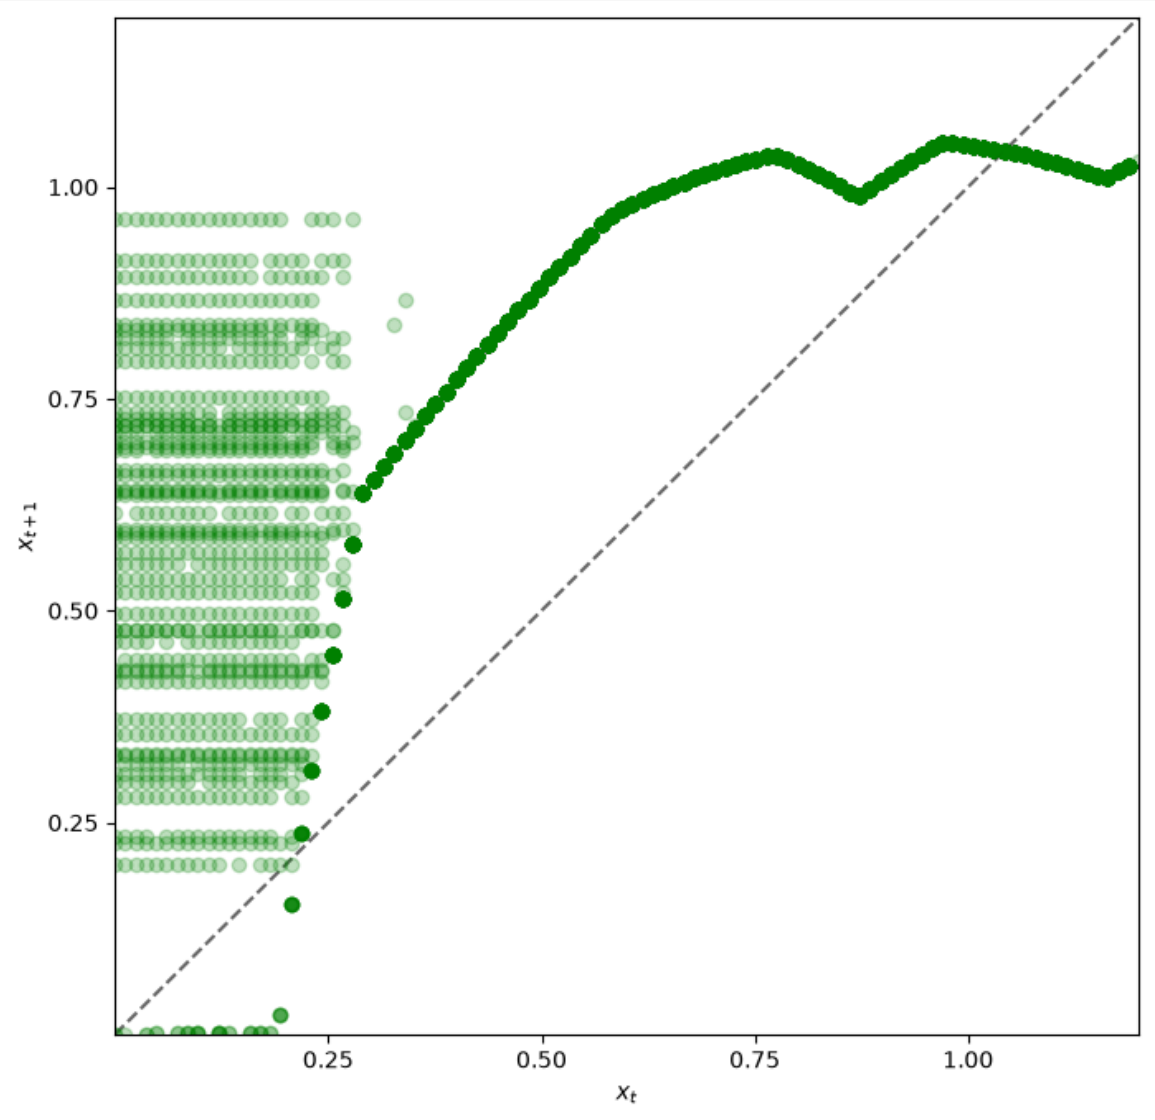
    return jnp.where(b, jnp.maximum(gxϕ, u), gxϕ)


x_next = simulate_next(jr.key(1234), plot_grid)

fig, ax = plt.subplots(figsize=(8, 8))
ticks = (0.25, 0.5, 0.75, 1.0)
ax.set(xticks=ticks, yticks=ticks,
       xlim=(0, plot_grid_max),
       ylim=(0, plot_grid_max),
       xlabel='$x_t$', ylabel='$x_{t+1}$')

ax.plot(plot_grid, plot_grid, 'k--', alpha=0.6)  # 45 degree line
ax.plot(jnp.repeat(plot_grid, x_next.shape[1]), x_next.ravel(),
        'go', alpha=0.25)

plt.show()
```



Looking at the dynamics, we can see that

- If x_t is below about 0.2 the dynamics are random, but $x_{t+1} > x_t$ is very likely.
- As x_t increases the dynamics become deterministic, and x_t converges to a steady state value close to 1.

Referring back to the figure *here* we see that $x_t \approx 1$ means that $s_t = s(x_t) \approx 0$ and $\phi_t = \phi(x_t) \approx 0.6$.

Exercise 61.5.2

In Exercise 61.5.1, we found that s_t converges to zero and ϕ_t converges to about 0.6.

Since these results were calculated at a value of β close to one, let's compare them to the best choice for an *infinitely* patient worker.

Intuitively, an infinitely patient worker would like to maximize steady state wages, which are a function of steady state capital.

You can take it as given—it's certainly true—that the infinitely patient worker does not search in the long run (i.e., $s_t = 0$ for large t).

Thus, given ϕ, steady state capital is the positive fixed point $x^*(\phi)$ of the map $x \mapsto g(x, \phi)$.

Steady state wages can be written as $w^*(\phi) = x^*(\phi)(1 - \phi)$.

Graph $w^*(\phi)$ with respect to ϕ, and examine the best choice of ϕ.

Can you give a rough interpretation for the value that you see?

Solution

The figure can be produced as follows

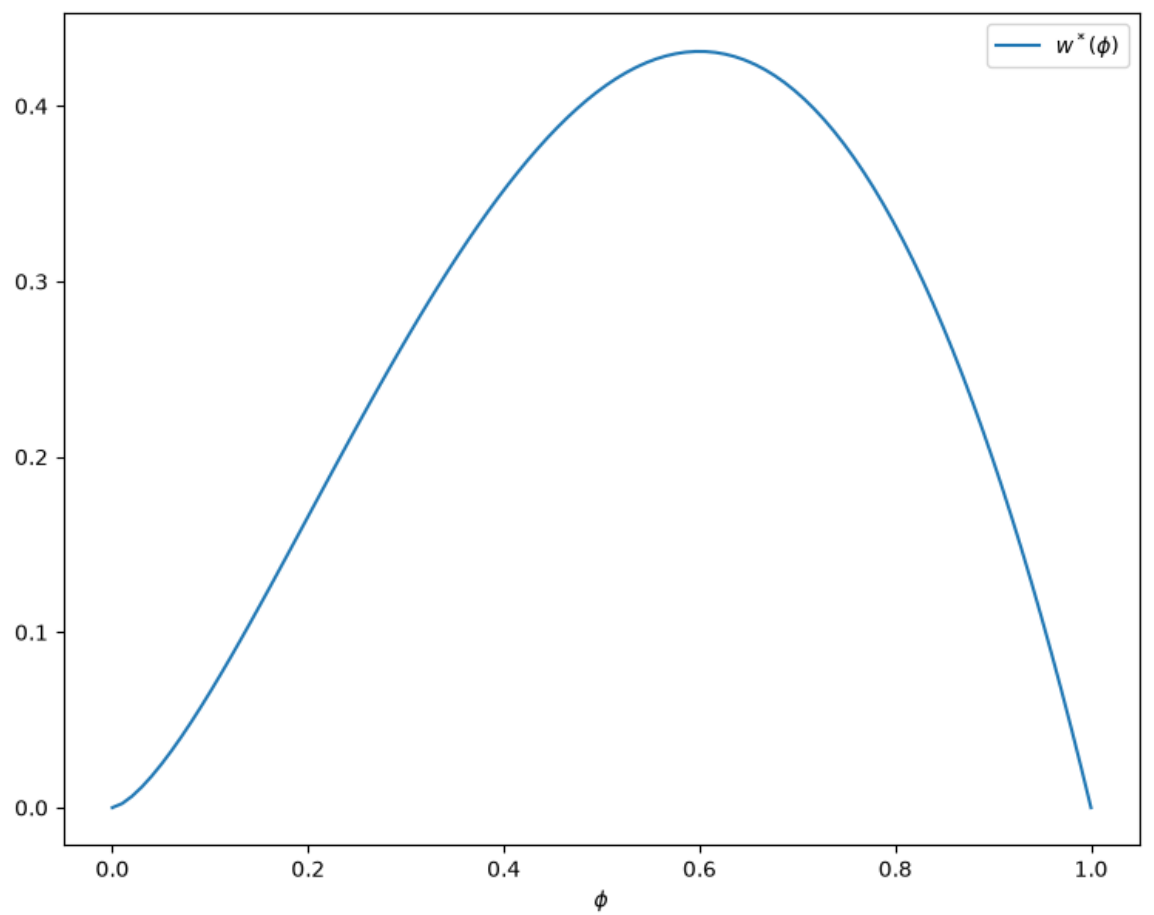
```
jv = create_jv_worker()

def xbar(ϕ):
    return (jv.A * ϕ**jv.α)**(1 / (1 - jv.α))

ϕ_grid = jnp.linspace(0, 1, 100)

fig, ax = plt.subplots(figsize=(9, 7))
ax.set(xlabel=r'$\phi$')
ax.plot(ϕ_grid, xbar(ϕ_grid) * (1 - ϕ_grid), label=r'$w^*(\phi)$')
ax.legend()

plt.show()
```



Observe that the maximizer is around 0.6.

This is similar to the long-run value for ϕ obtained in Exercise 61.5.1.

Hence the behavior of the infinitely patent worker is similar to that of the worker with $\beta = 0.96$.

This seems reasonable and helps us confirm that our dynamic programming solutions are probably correct.

CHAPTER

SIXTYTWO

JOB SEARCH IX: SEARCH WITH LEARNING

Contents

- *Job Search IX: Search with Learning*
  - *Overview*
  - *Model*
  - *Take 1: Solution by VFI*
  - *Take 2: A More Efficient Method*
  - *Another Functional Equation*
  - *Solving the RWFE*
  - *Implementation*
  - *Exercises*
  - *Solutions*
  - *Appendix A*
  - *Appendix B*
  - *Examples*

In addition to what's in Anaconda, this lecture deploys the libraries:

```
!pip install interpolation
```

62.1 Overview

In this lecture, we consider an extension of the previously studied job search model of McCall [McCall, 1970].

We'll build on a model of Bayesian learning discussed in this lecture on the topic of exchangeability and its relationship to the concept of IID (identically and independently distributed) random variables and to Bayesian updating.

In the McCall model, an unemployed worker decides when to accept a permanent job at a specific fixed wage, given

- his or her discount factor
- the level of unemployment compensation

- the distribution from which wage offers are drawn

In the version considered below, the wage distribution is unknown and must be learned.

- The following is based on the presentation in [Ljungqvist and Sargent, 2018], section 6.6.

Let's start with some imports

```
import matplotlib.pyplot as plt
from numba import jit, prange, vectorize
from interpolation import mlinterp
from math import gamma
import numpy as np
from matplotlib import cm
import scipy.optimize as op
from scipy.stats import cumfreq, beta
```

62.1.1 Model Features

- Infinite horizon dynamic programming with two states and one binary control.
- Bayesian updating to learn the unknown distribution.

62.2 Model

Let's first review the basic McCall model [McCall, 1970] and then add the variation we want to consider.

62.2.1 The Basic McCall Model

Recall that, *in the baseline model*, an unemployed worker is presented in each period with a permanent job offer at wage W_t.

At time t, our worker either

1. accepts the offer and works permanently at constant wage W_t
2. rejects the offer, receives unemployment compensation c and reconsiders next period

The wage sequence W_t is IID and generated from known density q.

The worker aims to maximize the expected discounted sum of earnings $\mathbb{E} \sum_{t=0}^{\infty} \beta^t y_t$.

Let $v(w)$ be the optimal value of the problem for a previously unemployed worker who has just received offer w and is yet to decide whether to accept or reject the offer.

The value function v satisfies the recursion

$$v(w) = \max \left\{ \frac{w}{1-\beta}, \, c + \beta \int v(w') q(w') dw' \right\} \tag{62.1}$$

The optimal policy has the form $\mathbf{1}\{w \geq \bar{w}\}$, where $\bar{w}$ is a constant called the **reservation wage**.

62.2.2 Offer Distribution Unknown

Now let's extend the model by considering the variation presented in [Ljungqvist and Sargent, 2018], section 6.6.

The model is as above, apart from the fact that

- the density q is unknown
- the worker learns about q by starting with a prior and updating based on wage offers that he/she observes

The worker knows there are two possible distributions F and G.

These two distributions have densities f and g, repectively.

Just before time starts, "nature" selects q to be either f or g.

This is then the wage distribution from which the entire sequence W_t will be drawn.

The worker does not know which distribution nature has drawn, but the worker does know the two possible distributions f and g.

The worker puts a (subjective) prior probability π_0 on f having been chosen.

The worker's time 0 subjective distribution for the distribution of W_0 is

$$\pi_0 f + (1 - \pi_0) g$$

The worker's time t subjective belief about the the distribution of W_t is

$$\pi_t f + (1 - \pi_t) g,$$

where π_t updates via

$$\pi_{t+1} = \frac{\pi_t f(w_{t+1})}{\pi_t f(w_{t+1}) + (1 - \pi_t) g(w_{t+1})} \tag{62.2}$$

This last expression follows from Bayes' rule, which tells us that

$$\mathbb{P}\{q = f \,|\, W = w\} = \frac{\mathbb{P}\{W = w \,|\, q = f\}\mathbb{P}\{q = f\}}{\mathbb{P}\{W = w\}}$$

and

$$\mathbb{P}\{W = w\} = \sum_{\omega \in \{f, g\}} \mathbb{P}\{W = w \,|\, q = \omega\}\mathbb{P}\{q = \omega\}$$

The fact that (62.2) is recursive allows us to progress to a recursive solution method.

Letting

$$q_\pi(w) := \pi f(w) + (1 - \pi) g(w)$$

and

$$\kappa(w, \pi) := \frac{\pi f(w)}{\pi f(w) + (1 - \pi) g(w)}$$

we can express the value function for the unemployed worker recursively as follows

$$v(w, \pi) = \max \left\{ \frac{w}{1 - \beta}, \, c + \beta \int v(w', \pi') \, q_\pi(w') \, dw' \right\} \quad \text{where} \quad \pi' = \kappa(w', \pi) \tag{62.3}$$

Notice that the current guess π is a state variable, since it affects the worker's perception of probabilities for future rewards.

62.2.3 Parameterization

Following section 6.6 of [Ljungqvist and Sargent, 2018], our baseline parameterization will be

- f is Beta$(1, 1)$
- g is Beta$(3, 1.2)$
- $\beta = 0.95$ and $c = 0.3$

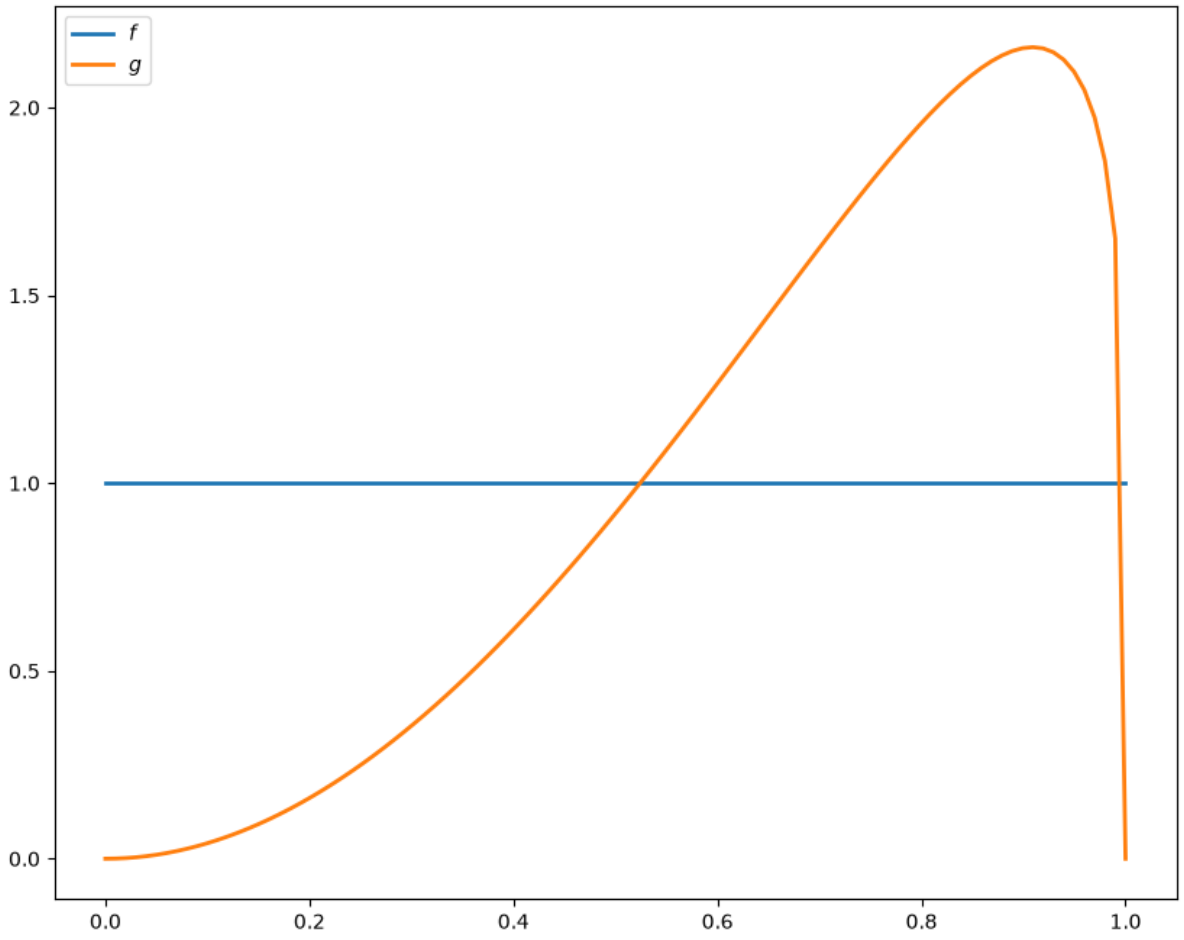
The densities f and g have the following shape

```
@vectorize
def p(x, a, b):
    r = gamma(a + b) / (gamma(a) * gamma(b))
    return r * x**(a-1) * (1 - x)**(b-1)


x_grid = np.linspace(0, 1, 100)
f = lambda x: p(x, 1, 1)
g = lambda x: p(x, 3, 1.2)

fig, ax = plt.subplots(figsize=(10, 8))
ax.plot(x_grid, f(x_grid), label='$f$', lw=2)
ax.plot(x_grid, g(x_grid), label='$g$', lw=2)

ax.legend()
plt.show()
```



62.2.4 Looking Forward

What kind of optimal policy might result from (62.3) and the parameterization specified above?

Intuitively, if we accept at w_a and $w_a \leq w_b$, then – all other things being given – we should also accept at w_b.

This suggests a policy of accepting whenever w exceeds some threshold value $\bar{w}$.

But $\bar{w}$ should depend on π – in fact, it should be decreasing in π because

- f is a less attractive offer distribution than g
- larger π means more weight on f and less on g

Thus, larger π depresses the worker's assessment of her future prospects, so relatively low current offers become more attractive.

Summary: We conjecture that the optimal policy is of the form $\mathbb{1}w \geq \bar{w}(\pi)$ for some decreasing function $\bar{w}$.

62.3 Take 1: Solution by VFI

Let's set about solving the model and see how our results match with our intuition.

We begin by solving via value function iteration (VFI), which is natural but ultimately turns out to be second best.

The class `SearchProblem` is used to store parameters and methods needed to compute optimal actions.

```
class SearchProblem:
    """
    A class to store a given parameterization of the "offer distribution
    unknown" model.

    """

    def __init__(self,
                 β=0.95,            # Discount factor
                 c=0.3,             # Unemployment compensation
                 F_a=1,
                 F_b=1,
                 G_a=3,
                 G_b=1.2,
                 w_max=1,           # Maximum wage possible
                 w_grid_size=100,
                 π_grid_size=100,
                 mc_size=500):

        self.β, self.c, self.w_max = β, c, w_max

        self.f = jit(lambda x: p(x, F_a, F_b))
        self.g = jit(lambda x: p(x, G_a, G_b))

        self.π_min, self.π_max = 1e-3, 1-1e-3    # Avoids instability
        self.w_grid = np.linspace(0, w_max, w_grid_size)
        self.π_grid = np.linspace(self.π_min, self.π_max, π_grid_size)

        self.mc_size = mc_size

        self.w_f = np.random.beta(F_a, F_b, mc_size)
        self.w_g = np.random.beta(G_a, G_b, mc_size)
```

The following function takes an instance of this class and returns jitted versions of the Bellman operator `T`, and a `get_greedy()` function to compute the approximate optimal policy from a guess `v` of the value function

```
def operator_factory(sp, parallel_flag=True):

    f, g = sp.f, sp.g
    w_f, w_g = sp.w_f, sp.w_g
    β, c = sp.β, sp.c
    mc_size = sp.mc_size
    w_grid, π_grid = sp.w_grid, sp.π_grid

    @jit
    def v_func(x, y, v):
        return mlinterp((w_grid, π_grid), v, (x, y))

    @jit
    def κ(w, π):
```

(continues on next page)

(continued from previous page)

```
        """
        Updates π using Bayes' rule and the current wage observation w.
        """
        pf, pg = π * f(w), (1 - π) * g(w)
        π_new = pf / (pf + pg)

        return π_new

    @jit(parallel=parallel_flag)
    def T(v):
        """
        The Bellman operator.

        """
        v_new = np.empty_like(v)

        for i in prange(len(w_grid)):
            for j in prange(len(π_grid)):
                w = w_grid[i]
                π = π_grid[j]

                v_1 = w / (1 - β)

                integral_f, integral_g = 0, 0
                for m in prange(mc_size):
                    integral_f += v_func(w_f[m], κ(w_f[m], π), v)
                    integral_g += v_func(w_g[m], κ(w_g[m], π), v)
                integral = (π * integral_f + (1 - π) * integral_g) / mc_size

                v_2 = c + β * integral
                v_new[i, j] = max(v_1, v_2)

        return v_new

    @jit(parallel=parallel_flag)
    def get_greedy(v):
        """"
        Compute optimal actions taking v as the value function.

        """
        σ = np.empty_like(v)

        for i in prange(len(w_grid)):
            for j in prange(len(π_grid)):
                w = w_grid[i]
                π = π_grid[j]

                v_1 = w / (1 - β)

                integral_f, integral_g = 0, 0
                for m in prange(mc_size):
                    integral_f += v_func(w_f[m], κ(w_f[m], π), v)
                    integral_g += v_func(w_g[m], κ(w_g[m], π), v)
                integral = (π * integral_f + (1 - π) * integral_g) / mc_size

                v_2 = c + β * integral
```

(continues on next page)

(continued from previous page)

```
                σ[i, j] = v_1 > v_2  # Evaluates to 1 or 0

        return σ

    return T, get_greedy
```

We will omit a detailed discussion of the code because there is a more efficient solution method that we will use later.

To solve the model we will use the following function that iterates using T to find a fixed point

```
def solve_model(sp,
                use_parallel=True,
                tol=1e-4,
                max_iter=1000,
                verbose=True,
                print_skip=5):

    """
    Solves for the value function

    * sp is an instance of SearchProblem
    """

    T, _ = operator_factory(sp, use_parallel)

    # Set up loop
    i = 0
    error = tol + 1
    m, n = len(sp.w_grid), len(sp.π_grid)

    # Initialize v
    v = np.zeros((m, n)) + sp.c / (1 - sp.β)

    while i < max_iter and error > tol:
        v_new = T(v)
        error = np.max(np.abs(v - v_new))
        i += 1
        if verbose and i % print_skip == 0:
            print(f"Error at iteration {i} is {error}.")
        v = v_new

    if error > tol:
        print("Failed to converge!")
    elif verbose:
        print(f"\nConverged in {i} iterations.")


    return v_new
```

Let's look at solutions computed from value function iteration

```
sp = SearchProblem()
v_star = solve_model(sp)
fig, ax = plt.subplots(figsize=(6, 6))
ax.contourf(sp.π_grid, sp.w_grid, v_star, 12, alpha=0.6, cmap=cm.jet)
cs = ax.contour(sp.π_grid, sp.w_grid, v_star, 12, colors="black")
ax.clabel(cs, inline=1, fontsize=10)
```

(continues on next page)

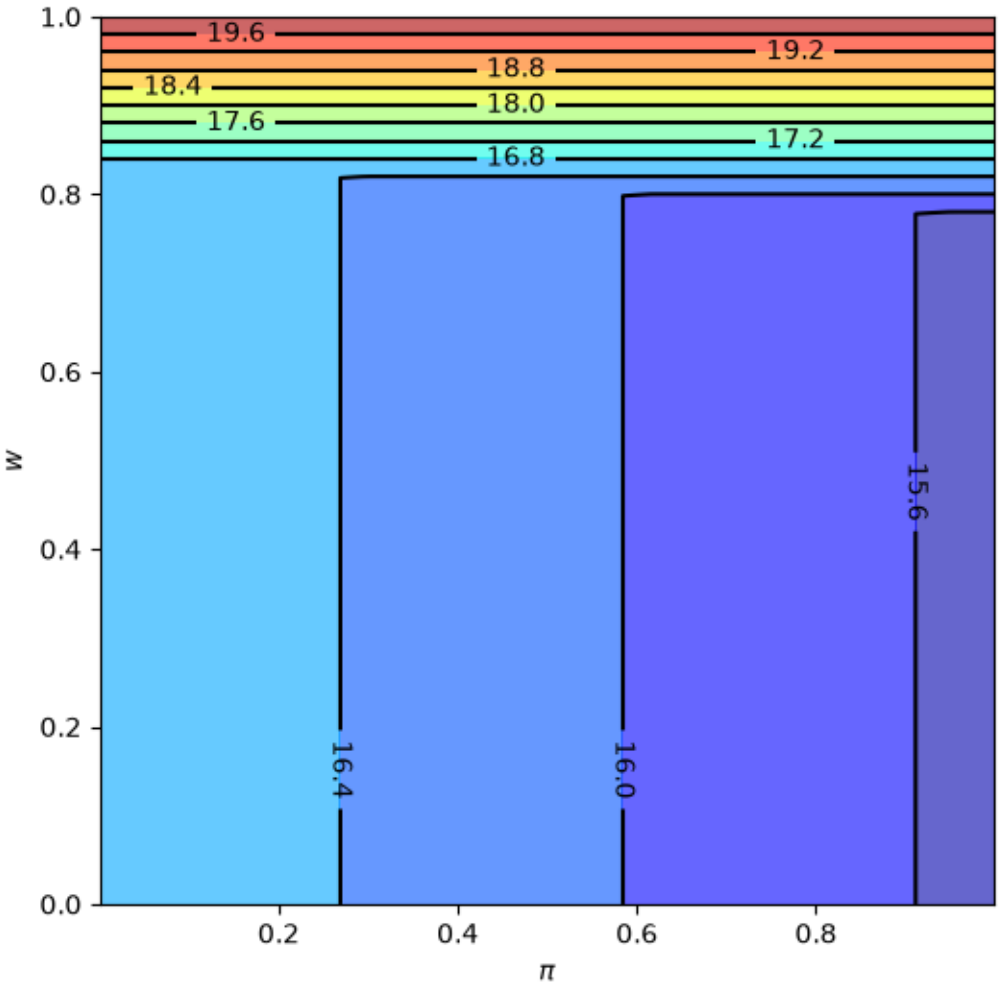
(continued from previous page)

```
ax.set(xlabel=r'$\pi$', ylabel='$w$')

plt.show()
```

```
Error at iteration 5 is 0.620947293686493.
Error at iteration 10 is 0.0975484410521954.
Error at iteration 15 is 0.019667250856013396.
Error at iteration 20 is 0.004029042290257578.
Error at iteration 25 is 0.0008254470732680375.
Error at iteration 30 is 0.00016909894672245684.

Converged in 32 iterations.
```



We will also plot the optimal policy

```
T, get_greedy = operator_factory(sp)
σ_star = get_greedy(v_star)

fig, ax = plt.subplots(figsize=(6, 6))
ax.contourf(sp.π_grid, sp.w_grid, σ_star, 1, alpha=0.6, cmap=cm.jet)
ax.contour(sp.π_grid, sp.w_grid, σ_star, 1, colors="black")
```

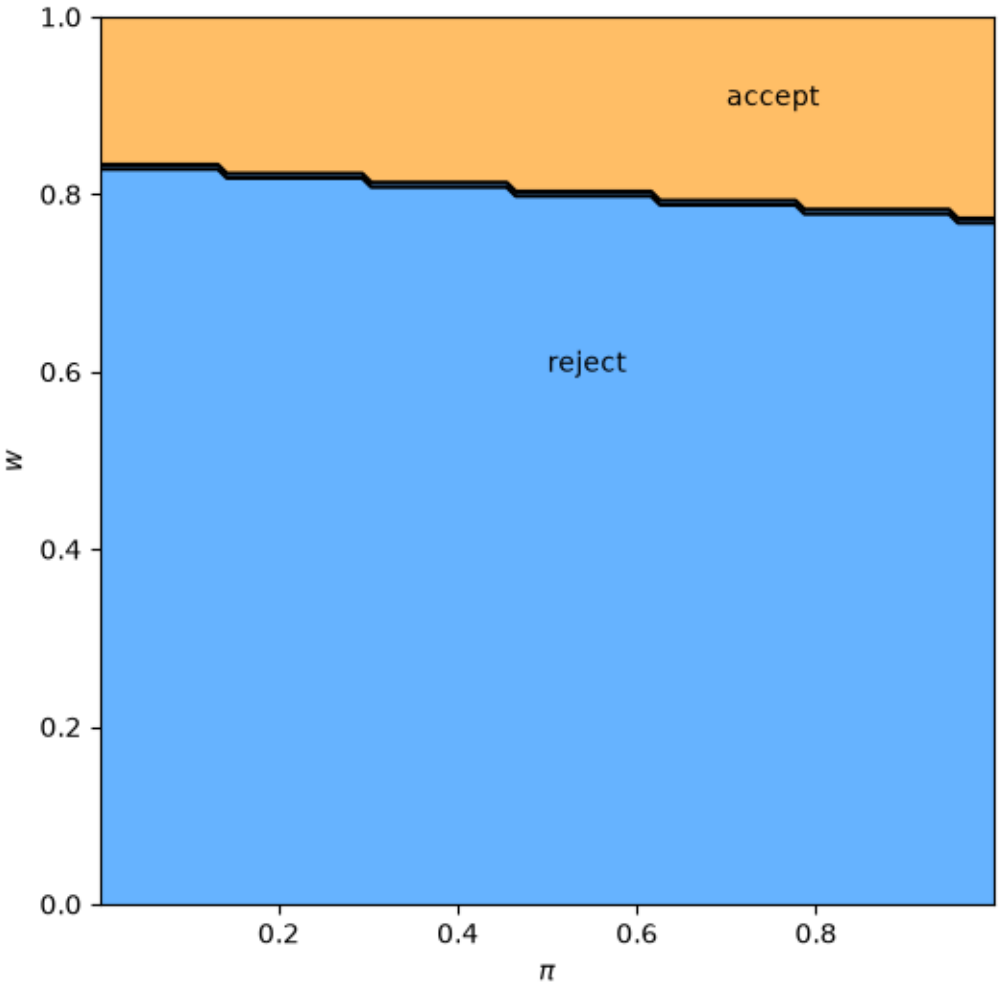
(continues on next page)

(continued from previous page)

```
ax.set(xlabel=r'$\pi$', ylabel='$w$')

ax.text(0.5, 0.6, 'reject')
ax.text(0.7, 0.9, 'accept')

plt.show()
```



The results fit well with our intuition from section *looking forward*.

- The black line in the figure above corresponds to the function $\bar{w}(\pi)$ introduced there.
- It is decreasing as expected.

62.4 Take 2: A More Efficient Method

Let's consider another method to solve for the optimal policy.

We will use iteration with an operator that has the same contraction rate as the Bellman operator, but

- one dimensional rather than two dimensional
- no maximization step

As a consequence, the algorithm is orders of magnitude faster than VFI.

This section illustrates the point that when it comes to programming, a bit of mathematical analysis goes a long way.

62.5 Another Functional Equation

To begin, note that when $w = \bar{w}(\pi)$, the worker is indifferent between accepting and rejecting.

Hence the two choices on the right-hand side of (62.3) have equal value:

$$\frac{\bar{w}(\pi)}{1-\beta} = c + \beta \int v(w', \pi') \, q_\pi(w') \, dw' \tag{62.4}$$

Together, (62.3) and (62.4) give

$$v(w, \pi) = \max \left\{ \frac{w}{1-\beta}, \, \frac{\bar{w}(\pi)}{1-\beta} \right\} \tag{62.5}$$

Combining (62.4) and (62.5), we obtain

$$\frac{\bar{w}(\pi)}{1-\beta} = c + \beta \int \max \left\{ \frac{w'}{1-\beta}, \, \frac{\bar{w}(\pi')}{1-\beta} \right\} q_\pi(w') \, dw'$$

Multiplying by $1-\beta$, substituting in $\pi' = \kappa(w', \pi)$ and using $\circ$ for composition of functions yields

$$\bar{w}(\pi) = (1-\beta)c + \beta \int \max \left\{ w', \bar{w} \circ \kappa(w', \pi) \right\} \, q_\pi(w') \, dw' \tag{62.6}$$

Equation (62.6) can be understood as a functional equation, where $\bar{w}$ is the unknown function.

- Let's call it the **reservation wage functional equation** (RWFE).
- The solution $\bar{w}$ to the RWFE is the object that we wish to compute.

62.6 Solving the RWFE

To solve the RWFE, we will first show that its solution is the fixed point of a contraction mapping.

To this end, let

- $b[0, 1]$ be the bounded real-valued functions on $[0, 1]$
- $\|\omega\| := \sup_{x \in [0,1]} |\omega(x)|$

Consider the operator Q mapping $\omega \in b[0, 1]$ into $Q\omega \in b[0, 1]$ via

$$(Q\omega)(\pi) = (1-\beta)c + \beta \int \max \left\{ w', \omega \circ \kappa(w', \pi) \right\} \, q_\pi(w') \, dw' \tag{62.7}$$

Comparing (62.6) and (62.7), we see that the set of fixed points of Q exactly coincides with the set of solutions to the RWFE.

- If $Q\bar{w} = \bar{w}$ then $\bar{w}$ solves (62.6) and vice versa.

Moreover, for any $\omega, \omega' \in b[0, 1]$, basic algebra and the triangle inequality for integrals tells us that

$$|(Q\omega)(\pi) - (Q\omega')(\pi)| \leq \beta \int |\max\{w', \omega \circ \kappa(w', \pi)\} - \max\{w', \omega' \circ \kappa(w', \pi)\}| \; q_\pi(w') \, dw' \tag{62.8}$$

Working case by case, it is easy to check that for real numbers a, b, c we always have

$$|\max\{a, b\} - \max\{a, c\}| \leq |b - c| \tag{62.9}$$

Combining (62.8) and (62.9) yields

$$|(Q\omega)(\pi) - (Q\omega')(\pi)| \leq \beta \int |\omega \circ \kappa(w', \pi) - \omega' \circ \kappa(w', \pi)| \; q_\pi(w') \, dw' \leq \beta \|\omega - \omega'\| \tag{62.10}$$

Taking the supremum over π now gives us

$$\|Q\omega - Q\omega'\| \leq \beta \|\omega - \omega'\| \tag{62.11}$$

In other words, Q is a contraction of modulus β on the complete metric space $(b[0, 1], \|\cdot\|)$.

Hence

- A unique solution $\bar{w}$ to the RWFE exists in $b[0, 1]$.
- $Q^k \omega \to \bar{w}$ uniformly as $k \to \infty$, for any $\omega \in b[0, 1]$.

62.7 Implementation

The following function takes an instance of `SearchProblem` and returns the operator `Q`

```
def Q_factory(sp, parallel_flag=True):

    f, g = sp.f, sp.g
    w_f, w_g = sp.w_f, sp.w_g
    β, c = sp.β, sp.c
    mc_size = sp.mc_size
    w_grid, π_grid = sp.w_grid, sp.π_grid

    @jit
    def ω_func(p, ω):
        return np.interp(p, π_grid, ω)

    @jit
    def κ(w, π):
        """
        Updates π using Bayes' rule and the current wage observation w.
        """
        pf, pg = π * f(w), (1 - π) * g(w)
        π_new = pf / (pf + pg)

        return π_new

    @jit(parallel=parallel_flag)
    def Q(ω):
        """
```

(continues on next page)

(continued from previous page)

```
        Updates the reservation wage function guess ω via the operator
        Q.

        """
        ω_new = np.empty_like(ω)

        for i in prange(len(π_grid)):
            π = π_grid[i]
            integral_f, integral_g = 0, 0

            for m in prange(mc_size):
                integral_f += max(w_f[m], ω_func(κ(w_f[m], π), ω))
                integral_g += max(w_g[m], ω_func(κ(w_g[m], π), ω))
            integral = (π * integral_f + (1 - π) * integral_g) / mc_size

            ω_new[i] = (1 - β) * c + β * integral

        return ω_new

    return Q
```

In the next exercise, you are asked to compute an approximation to $\bar{w}$.

62.8 Exercises

Exercise 62.8.1

Use the default parameters and `Q_factory` to compute an optimal policy.

Your result should coincide closely with the figure for the optimal policy *shown above*.

Try experimenting with different parameters, and confirm that the change in the optimal policy coincides with your intuition.

62.9 Solutions

Solution

This code solves the "Offer Distribution Unknown" model by iterating on a guess of the reservation wage function.

You should find that the run time is shorter than that of the value function approach.

Similar to above, we set up a function to iterate with `Q` to find the fixed point

```
def solve_wbar(sp,
               use_parallel=True,
               tol=1e-4,
               max_iter=1000,
               verbose=True,
               print_skip=5):

    Q = Q_factory(sp, use_parallel)
```

```
    # Set up loop
    i = 0
    error = tol + 1
    m, n = len(sp.w_grid), len(sp.π_grid)

    # Initialize w
    w = np.ones_like(sp.π_grid)

    while i < max_iter and error > tol:
        w_new = Q(w)
        error = np.max(np.abs(w - w_new))
        i += 1
        if verbose and i % print_skip == 0:
            print(f"Error at iteration {i} is {error}.")
        w = w_new

    if error > tol:
        print("Failed to converge!")
    elif verbose:
        print(f"\nConverged in {i} iterations.")

    return w_new
```

The solution can be plotted as follows

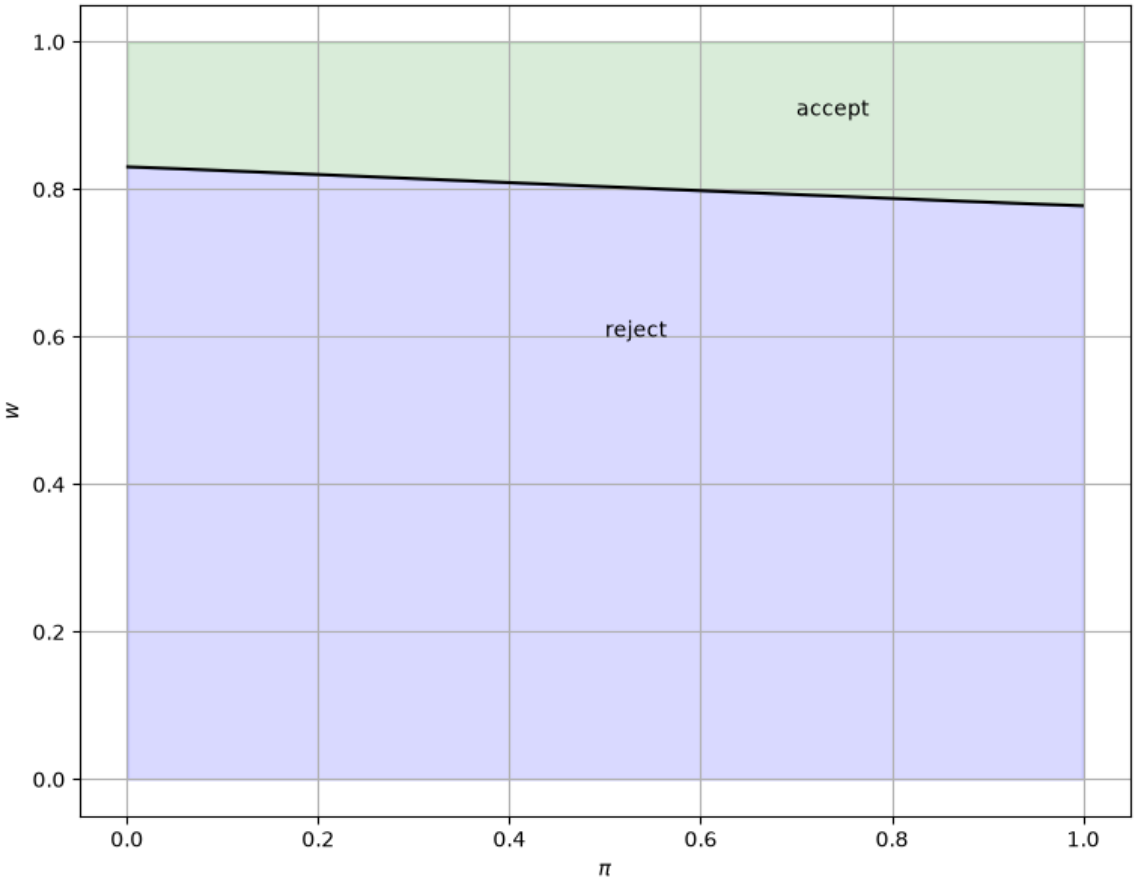
```
sp = SearchProblem()
w_bar = solve_wbar(sp)

fig, ax = plt.subplots(figsize=(9, 7))

ax.plot(sp.π_grid, w_bar, color='k')
ax.fill_between(sp.π_grid, 0, w_bar, color='blue', alpha=0.15)
ax.fill_between(sp.π_grid, w_bar, sp.w_max, color='green', alpha=0.15)
ax.text(0.5, 0.6, 'reject')
ax.text(0.7, 0.9, 'accept')
ax.set(xlabel=r'$\pi$', ylabel='$w$')
ax.grid()
plt.show()
```

```
Error at iteration 5 is 0.021773468257772155.
Error at iteration 10 is 0.00645647791078896.
Error at iteration 15 is 0.0015045191620605136.
Error at iteration 20 is 0.0003281405747559418 3.

Converged in 24 iterations.
```



62.10 Appendix A

The next piece of code generates a fun simulation to see what the effect of a change in the underlying distribution on the unemployment rate is.

At a point in the simulation, the distribution becomes significantly worse.

It takes a while for agents to learn this, and in the meantime, they are too optimistic and turn down too many jobs.

As a result, the unemployment rate spikes

```
F_a, F_b, G_a, G_b = 1, 1, 3, 1.2

sp = SearchProblem(F_a=F_a, F_b=F_b, G_a=G_a, G_b=G_b)
f, g = sp.f, sp.g

# Solve for reservation wage
w_bar = solve_wbar(sp, verbose=False)

# Interpolate reservation wage function
π_grid = sp.π_grid
w_func = jit(lambda x: np.interp(x, π_grid, w_bar))
```

(continues on next page)

(continued from previous page)

```
@jit
def update(a, b, e, π):
    "Update e and π by drawing wage offer from beta distribution with parameters a␣
 ↪and b"

    if e == False:
        w = np.random.beta(a, b)       # Draw random wage
        if w >= w_func(π):
            e = True                   # Take new job
        else:
            π = 1 / (1 + ((1 - π) * g(w)) / (π * f(w)))

    return e, π

@jit
def simulate_path(F_a=F_a,
                  F_b=F_b,
                  G_a=G_a,
                  G_b=G_b,
                  N=5000,        # Number of agents
                  T=600,         # Simulation length
                  d=200,         # Change date
                  s=0.025):      # Separation rate

    """Simulates path of employment for N number of works over T periods"""

    e = np.ones((N, T+1))
    π = np.full((N, T+1), 1e-3)

    a, b = G_a, G_b   # Initial distribution parameters

    for t in range(T+1):

        if t == d:
            a, b = F_a, F_b  # Change distribution parameters

        # Update each agent
        for n in range(N):
            if e[n, t] == 1:                    # If agent is currently employment
                p = np.random.uniform(0, 1)
                if p <= s:                      # Randomly separate with probability s
                    e[n, t] = 0

            new_e, new_π = update(a, b, e[n, t], π[n, t])
            e[n, t+1] = new_e
            π[n, t+1] = new_π

    return e[:, 1:]

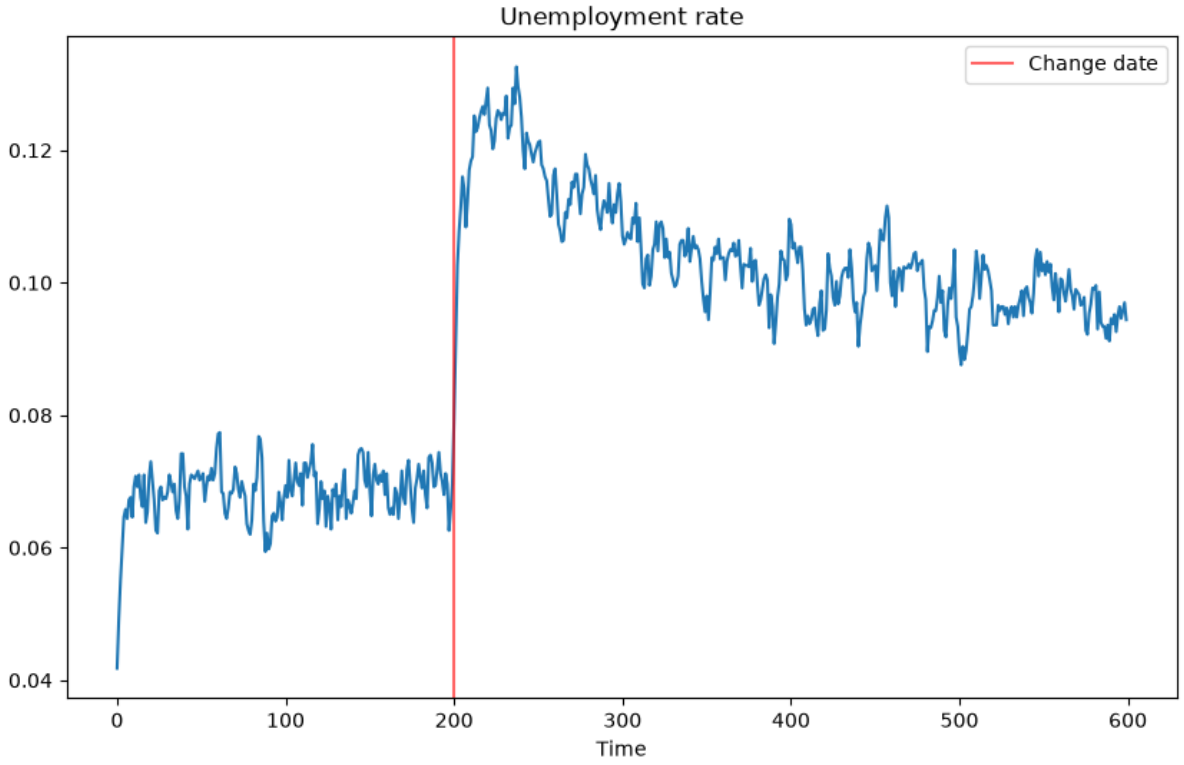
d = 200  # Change distribution at time d
unemployment_rate = 1 - simulate_path(d=d).mean(axis=0)

fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(unemployment_rate)
ax.axvline(d, color='r', alpha=0.6, label='Change date')
ax.set_xlabel('Time')
ax.set_title('Unemployment rate')
```

(continues on next page)

(continued from previous page)

```
ax.legend()
plt.show()
```



62.11 Appendix B

In this appendix we provide more details about how Bayes' Law contributes to the workings of the model.

We present some graphs that bring out additional insights about how learning works.

We build on graphs proposed in this lecture.

In particular, we'll add actions of our searching worker to a key graph presented in that lecture.

To begin, we first define two functions for computing the empirical distributions of unemployment duration and π at the time of employment.

```
@jit
def empirical_dist(F_a, F_b, G_a, G_b, w_bar, π_grid,
                   N=10000, T=600):
    """
    Simulates population for computing empirical cumulative
    distribution of unemployment duration and π at time when
    the worker accepts the wage offer. For each job searching
    problem, we simulate for two cases that either f or g is
    the true offer distribution.

    Parameters
    ----------
```

(continues on next page)

(continued from previous page)

```
    F_a, F_b, G_a, G_b : parameters of beta distributions F and G.
    w_bar : the reservation wage
    π_grid : grid points of π, for interpolation
    N : number of workers for simulation, optional
    T : maximum of time periods for simulation, optional

    Returns
    -------
    accpet_t : 2 by N ndarray. the empirical distribution of
               unemployment duration when f or g generates offers.
    accept_π : 2 by N ndarray. the empirical distribution of
               π at the time of employment when f or g generates offers.
    """

    accept_t = np.empty((2, N))
    accept_π = np.empty((2, N))

    # f or g generates offers
    for i, (a, b) in enumerate([(F_a, F_b), (G_a, G_b)]):
        # update each agent
        for n in range(N):

            # initial priori
            π = 0.5

            for t in range(T+1):

                # Draw random wage
                w = np.random.beta(a, b)
                lw = p(w, F_a, F_b) / p(w, G_a, G_b)
                π = π * lw / (π * lw + 1 - π)

                # move to next agent if accepts
                if w >= np.interp(π, π_grid, w_bar):
                    break

            # record the unemployment duration
            # and π at the time of acceptance
            accept_t[i, n] = t
            accept_π[i, n] = π

    return accept_t, accept_π

def cumfreq_x(res):
    """
    A helper function for calculating the x grids of
    the cumulative frequency histogram.
    """

    cumcount = res.cumcount
    lowerlimit, binsize = res.lowerlimit, res.binsize

    x = lowerlimit + np.linspace(0, binsize*cumcount.size, cumcount.size)

    return x
```

Now we define a wrapper function for analyzing job search models with learning under different parameterizations.

The wrapper takes parameters of beta distributions and unemployment compensation as inputs and then displays various things we want to know to interpret the solution of our search model.

In addition, it computes empirical cumulative distributions of two key objects.

```
def job_search_example(F_a=1, F_b=1, G_a=3, G_b=1.2, c=0.3):
    """
    Given the parameters that specify F and G distributions,
    calculate and display the rejection and acceptance area,
    the evolution of belief π, and the probability of accepting
    an offer at different π level, and simulate and calculate
    the empirical cumulative distribution of the duration of
    unemployment and π at the time the worker accepts the offer.
    """

    # construct a search problem
    sp = SearchProblem(F_a=F_a, F_b=F_b, G_a=G_a, G_b=G_b, c=c)
    f, g = sp.f, sp.g
    π_grid = sp.π_grid

    # Solve for reservation wage
    w_bar = solve_wbar(sp, verbose=False)

    # l(w) = f(w) / g(w)
    l = lambda w: f(w) / g(w)
    # objective function for solving l(w) = 1
    obj = lambda w: l(w) - 1.

    # the mode of beta distribution
    # use this to divide w into two intervals for root finding
    G_mode = (G_a - 1) / (G_a + G_b - 2)
    roots = np.empty(2)
    roots[0] = op.root_scalar(obj, bracket=[1e-10, G_mode]).root
    roots[1] = op.root_scalar(obj, bracket=[G_mode, 1-1e-10]).root

    fig, axs = plt.subplots(2, 2, figsize=(12, 9))

    # part 1: display the details of the model settings and some results
    w_grid = np.linspace(1e-12, 1-1e-12, 100)

    axs[0, 0].plot(l(w_grid), w_grid, label='$l$', lw=2)
    axs[0, 0].vlines(1., 0., 1., linestyle="--")
    axs[0, 0].hlines(roots, 0., 2., linestyle="--")
    axs[0, 0].set_xlim([0., 2.])
    axs[0, 0].legend(loc=4)
    axs[0, 0].set(xlabel='$l(w)=f(w)/g(w)$', ylabel='$w$')

    axs[0, 1].plot(sp.π_grid, w_bar, color='k')
    axs[0, 1].fill_between(sp.π_grid, 0, w_bar, color='blue', alpha=0.15)
    axs[0, 1].fill_between(sp.π_grid, w_bar, sp.w_max, color='green', alpha=0.15)
    axs[0, 1].text(0.5, 0.6, 'reject')
    axs[0, 1].text(0.7, 0.9, 'accept')

    W = np.arange(0.01, 0.99, 0.08)
    Π = np.arange(0.01, 0.99, 0.08)

    ΔW = np.zeros((len(W), len(Π)))
```

(continues on next page)

(continued from previous page)

```
    ΔΠ = np.empty((len(W), len(Π)))
    for i, w in enumerate(W):
        for j, π in enumerate(Π):
            lw = l(w)
            ΔΠ[i, j] = π * (lw / (π * lw + 1 - π) - 1)

    q = axs[0, 1].quiver(Π, W, ΔΠ, ΔW, scale=2, color='r', alpha=0.8)

    axs[0, 1].hlines(roots, 0., 1., linestyle="--")
    axs[0, 1].set(xlabel=r'$\pi$', ylabel='$w$')
    axs[0, 1].grid()

    axs[1, 0].plot(f(x_grid), x_grid, label='$f$', lw=2)
    axs[1, 0].plot(g(x_grid), x_grid, label='$g$', lw=2)
    axs[1, 0].vlines(1., 0., 1., linestyle="--")
    axs[1, 0].hlines(roots, 0., 2., linestyle="--")
    axs[1, 0].legend(loc=4)
    axs[1, 0].set(xlabel='$f(w), g(w)$', ylabel='$w$')

    axs[1, 1].plot(sp.π_grid, 1 - beta.cdf(w_bar, F_a, F_b), label='$f$')
    axs[1, 1].plot(sp.π_grid, 1 - beta.cdf(w_bar, G_a, G_b), label='$g$')
    axs[1, 1].set_ylim([0., 1.])
    axs[1, 1].grid()
    axs[1, 1].legend(loc=4)
    axs[1, 1].set(xlabel=r'$\pi$', ylabel=r'$\mathbb{P}\{w > \overline{w} (\pi)\}$')

    plt.show()

    # part 2: simulate empirical cumulative distribution
    accept_t, accept_π = empirical_dist(F_a, F_b, G_a, G_b, w_bar, π_grid)
    N = accept_t.shape[1]

    cfq_t_F = cumfreq(accept_t[0, :], numbins=100)
    cfq_π_F = cumfreq(accept_π[0, :], numbins=100)

    cfq_t_G = cumfreq(accept_t[1, :], numbins=100)
    cfq_π_G = cumfreq(accept_π[1, :], numbins=100)

    fig, axs = plt.subplots(2, 1, figsize=(12, 9))

    axs[0].plot(cumfreq_x(cfq_t_F), cfq_t_F.cumcount/N, label="f generates")
    axs[0].plot(cumfreq_x(cfq_t_G), cfq_t_G.cumcount/N, label="g generates")
    axs[0].grid(linestyle='--')
    axs[0].legend(loc=4)
    axs[0].title.set_text('CDF of duration of unemployment')
    axs[0].set(xlabel='time', ylabel='Prob(time)')

    axs[1].plot(cumfreq_x(cfq_π_F), cfq_π_F.cumcount/N, label="f generates")
    axs[1].plot(cumfreq_x(cfq_π_G), cfq_π_G.cumcount/N, label="g generates")
    axs[1].grid(linestyle='--')
    axs[1].legend(loc=4)
    axs[1].title.set_text('CDF of π at time worker accepts wage and leaves␣
↪unemployment')
    axs[1].set(xlabel='π', ylabel='Prob(π)')

    plt.show()
```

We now provide some examples that provide insights about how the model works.

62.12 Examples

62.12.1 Example 1 (Baseline)

F ~ Beta(1, 1), G ~ Beta(3, 1.2), c=0.3.

In the graphs below, the red arrows in the upper right figure show how π_t is updated in response to the new information w_t.

Recall the following formula from this lecture

$$\frac{\pi_{t+1}}{\pi_t} = \frac{l\left(w_{t+1}\right)}{\pi_t l\left(w_{t+1}\right) + \left(1 - \pi_t\right)} \begin{cases} > 1 & \text{if } l\left(w_{t+1}\right) > 1 \\ \leq 1 & \text{if } l\left(w_{t+1}\right) \leq 1 \end{cases}$$

The formula implies that the direction of motion of π_t is determined by the relationship between $l(w_t)$ and 1.

The magnitude is small if

- $l(w)$ is close to 1, which means the new w is not very informative for distinguishing two distributions,
- π_{t-1} is close to either 0 or 1, which means the priori is strong.

Will an unemployed worker accept an offer earlier or not, when the actual ruling distribution is g instead of f?

Two countervailing effects are at work.

- if f generates successive wage offers, then w is more likely to be low, but π is moving up toward to 1, which lowers the reservation wage, i.e., the worker becomes less selective the longer he or she remains unemployed.
- if g generates wage offers, then w is more likely to be high, but π is moving downward toward 0, increasing the reservation wage, i.e., the worker becomes more selective the longer he or she remains unemployed.

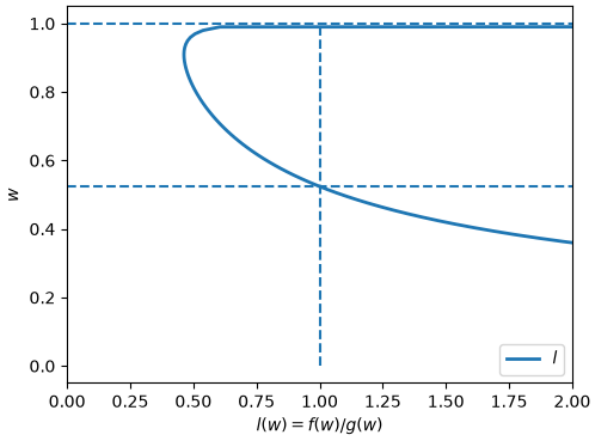
Quantitatively, the lower right figure sheds light on which effect dominates in this example.

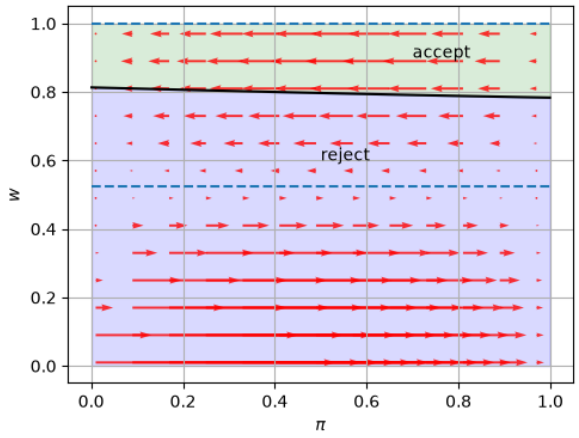
It shows the probability that a previously unemployed worker accepts an offer at different values of π when f or g generates wage offers.

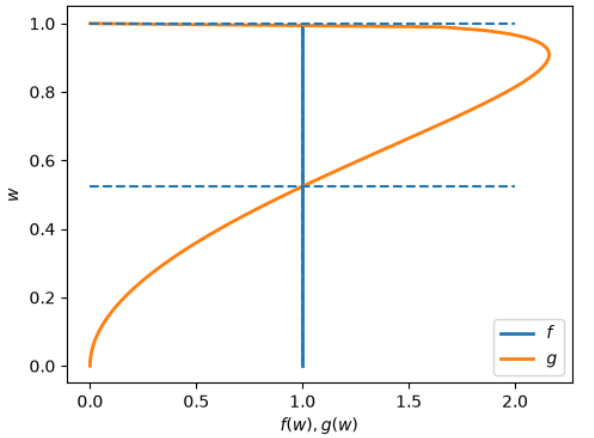
That graph shows that for the particular f and g in this example, the worker is always more likely to accept an offer when f generates the data even when π is close to zero so that the worker believes the true distribution is g and therefore is relatively more selective.

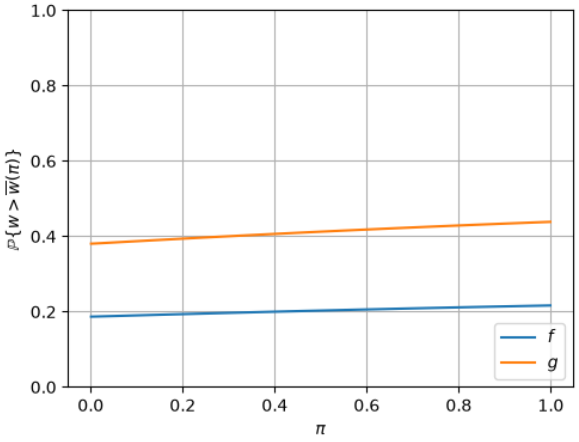
The empirical cumulative distribution of the duration of unemployment verifies our conjecture.

```
job_search_example()
```


w
l(w) = f(w)/g(w)
l


accept
reject
w
π


w
f(w), g(w)
f
g


$\mathbb{P}\{w > \overline{w}(\pi)\}$
π
f
g



62.12.2 Example 2

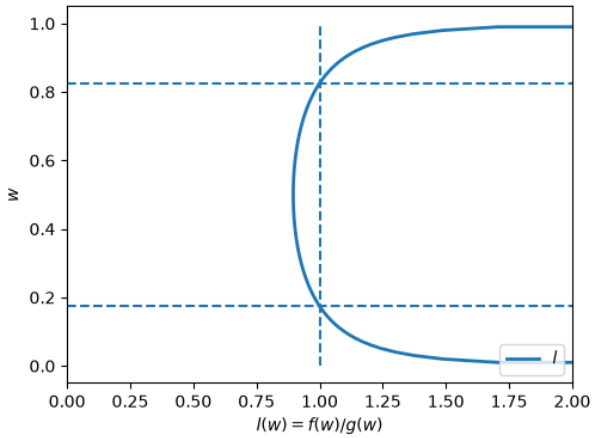
F ~ Beta(1, 1), G ~ Beta(1.2, 1.2), c=0.3.

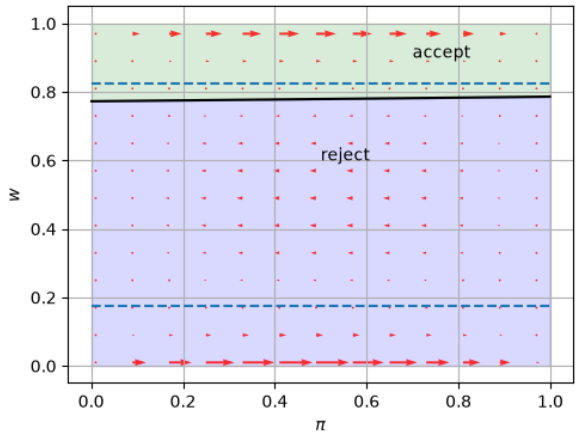
Now G has the same mean as F with a smaller variance.

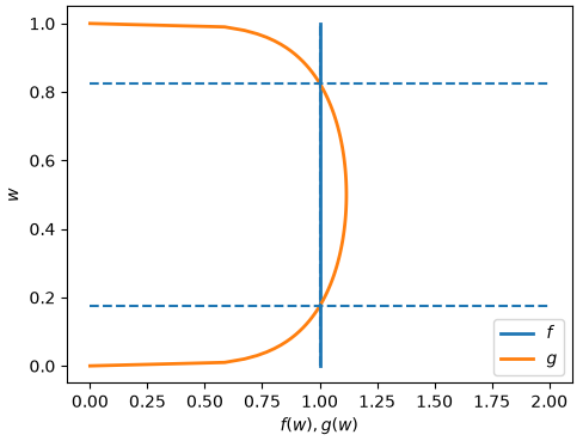
Since the unemployment compensation c serves as a lower bound for bad wage offers, G is now an "inferior" distribution to F.

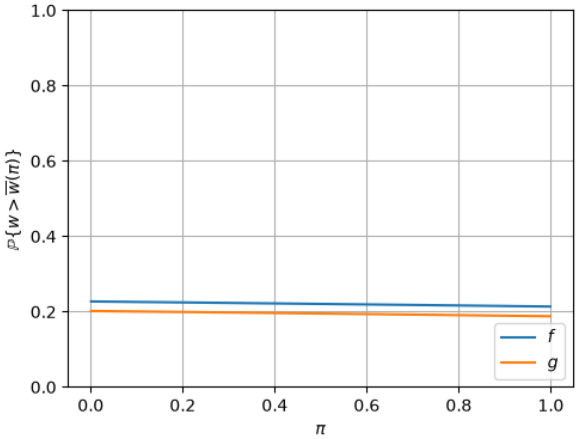
Consequently, we observe that the optimal policy $\overline{w}(\pi)$ is increasing in π.

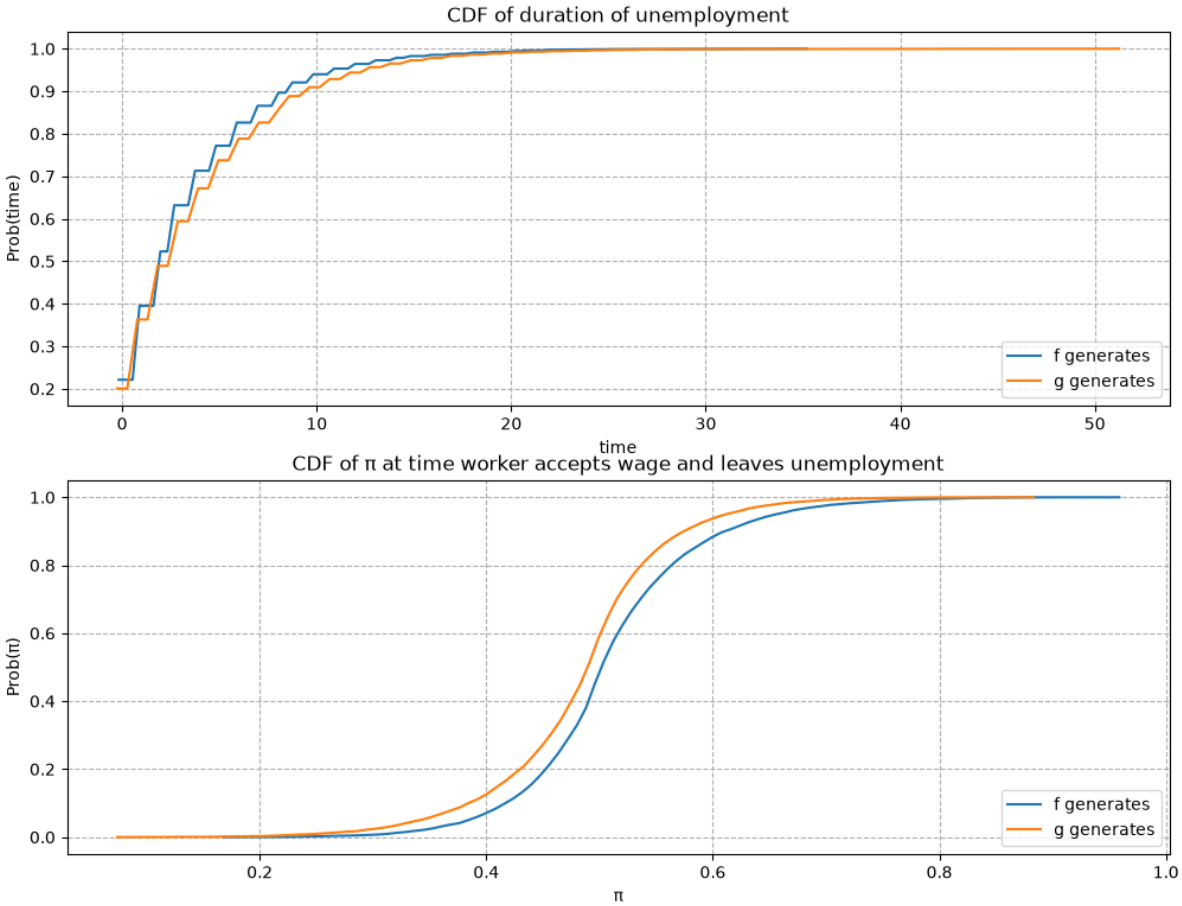
```
job_search_example(1, 1, 1.2, 1.2, 0.3)
```


w
l(w) = f(w)/g(w)
l


accept
reject
w
π


w
f(w), g(w)
f
g


$\mathbb{P}\{w > \bar{w}(\pi)\}$
π
f
g

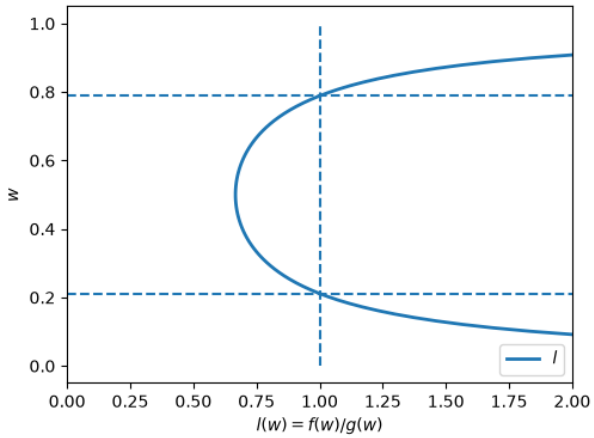


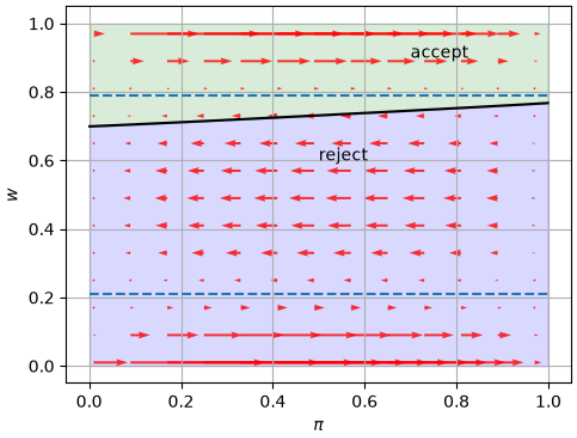
62.12.3 Example 3

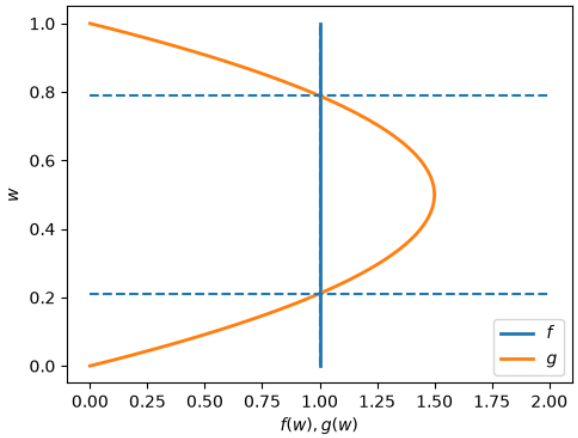
F ~ Beta(1, 1), G ~ Beta(2, 2), c=0.3.

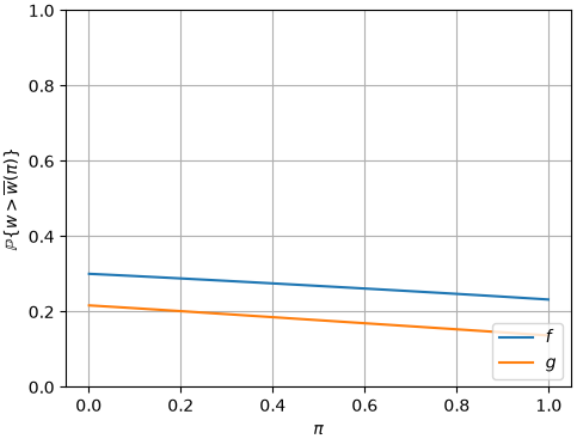
If the variance of G is smaller, we observe in the result that G is even more "inferior" and the slope of $\overline{w}(\pi)$ is larger.

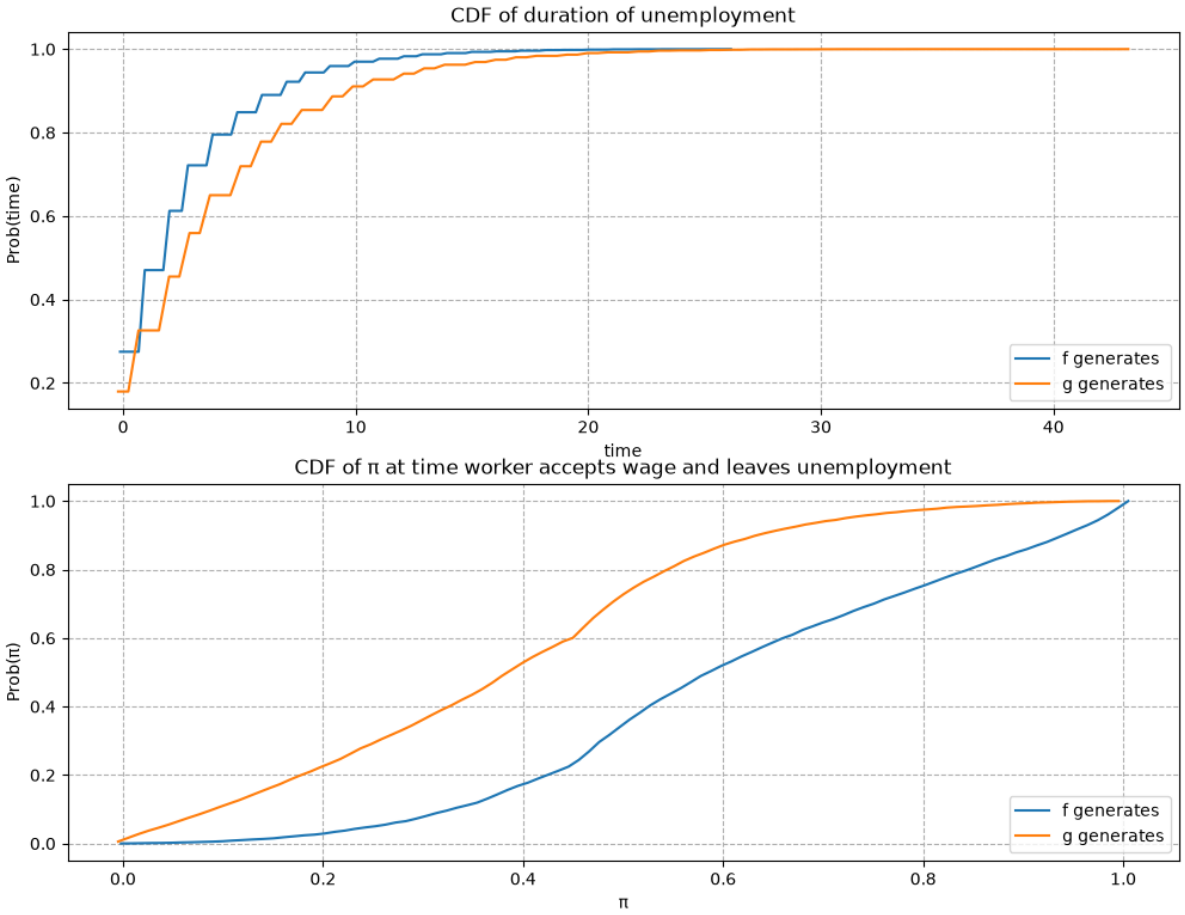
```
job_search_example(1, 1, 2, 2, 0.3)
```


w
l(w) = f(w)/g(w)
l


accept
reject
w
π


w
f(w), g(w)
f
g


$\mathbb{P}\{w > \overline{w}(\pi)\}$
π
f
g



62.12.4 Example 4

F ~ Beta(1, 1), G ~ Beta(3, 1.2), and c=0.8.

In this example, we keep the parameters of beta distributions to be the same with the baseline case but increase the unemployment compensation c.

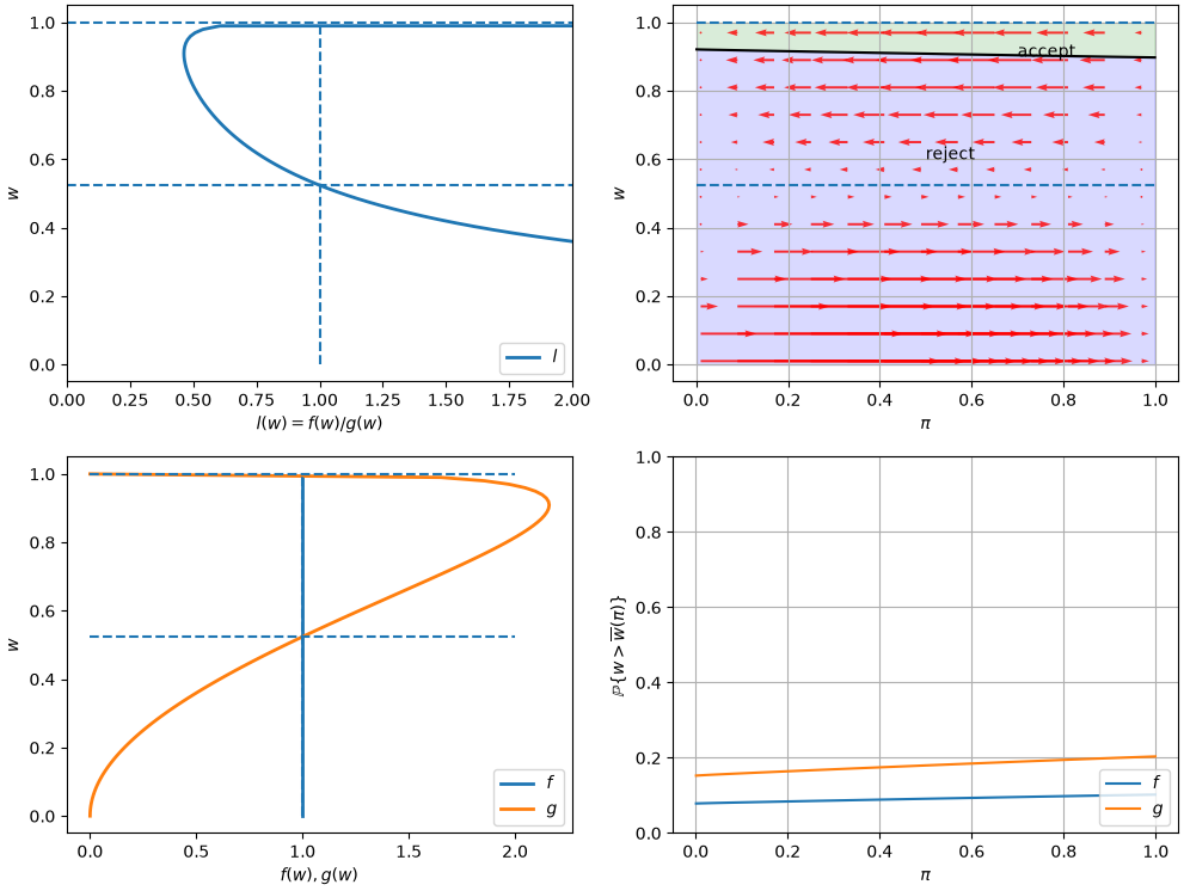
Comparing outcomes to the baseline case (example 1) in which unemployment compensation if low (c=0.3), now the worker can afford a longer learning period.

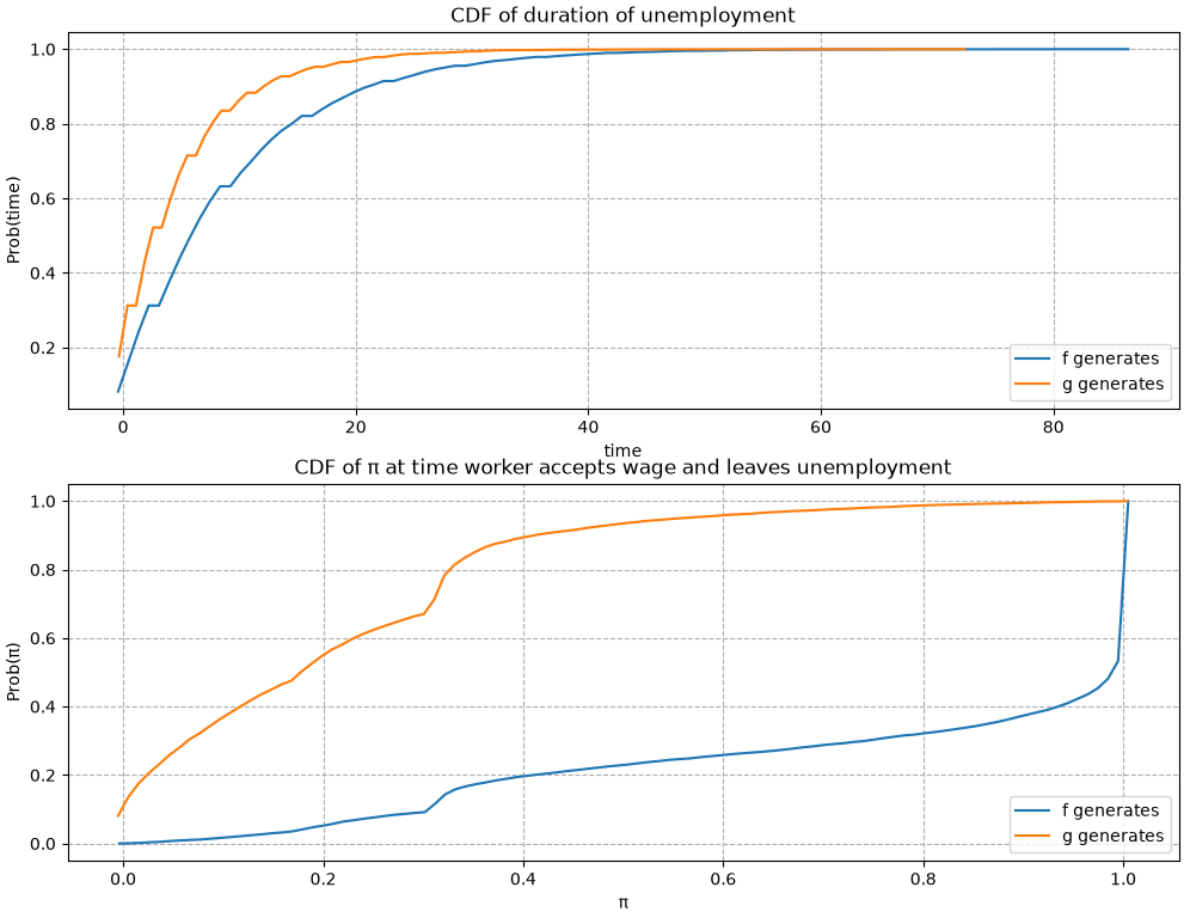
As a result, the worker tends to accept wage offers much later.

Furthermore, at the time of accepting employment, the belief π is closer to either 0 or 1.

That means that the worker has a better idea about what the true distribution is when he eventually chooses to accept a wage offer.

```
job_search_example(1, 1, 3, 1.2, c=0.8)
```



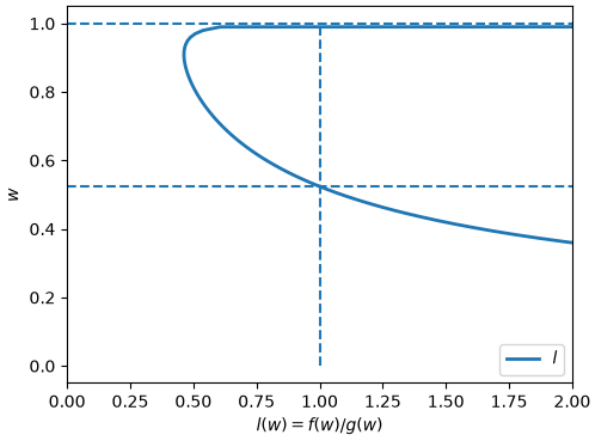


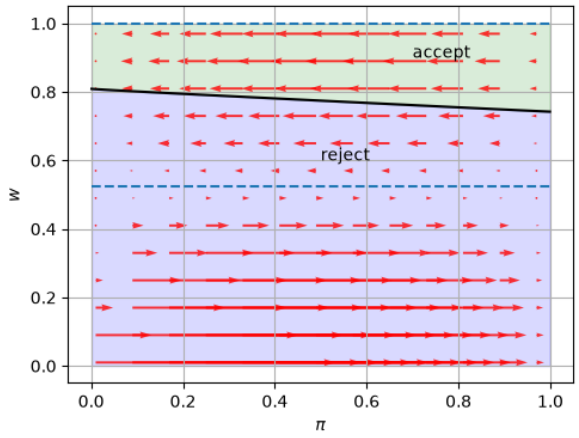
62.12.5 Example 5

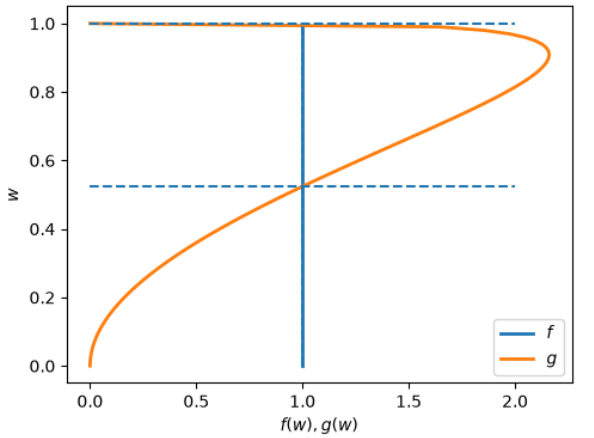
F ~ Beta(1, 1), G ~ Beta(3, 1.2), and c=0.1.

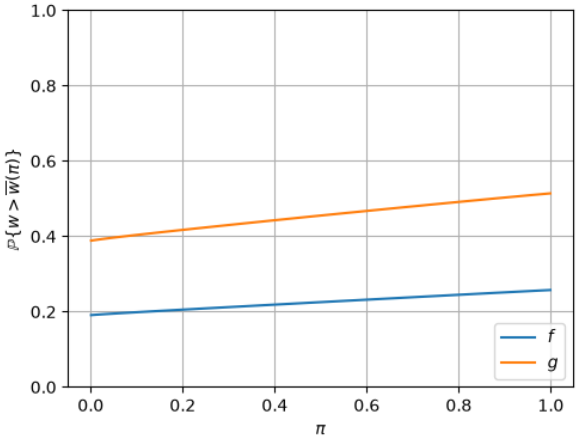
As expected, a smaller c makes an unemployed worker accept wage offers earlier after having acquired less information about the wage distribution.

```
job_search_example(1, 1, 3, 1.2, c=0.1)
```

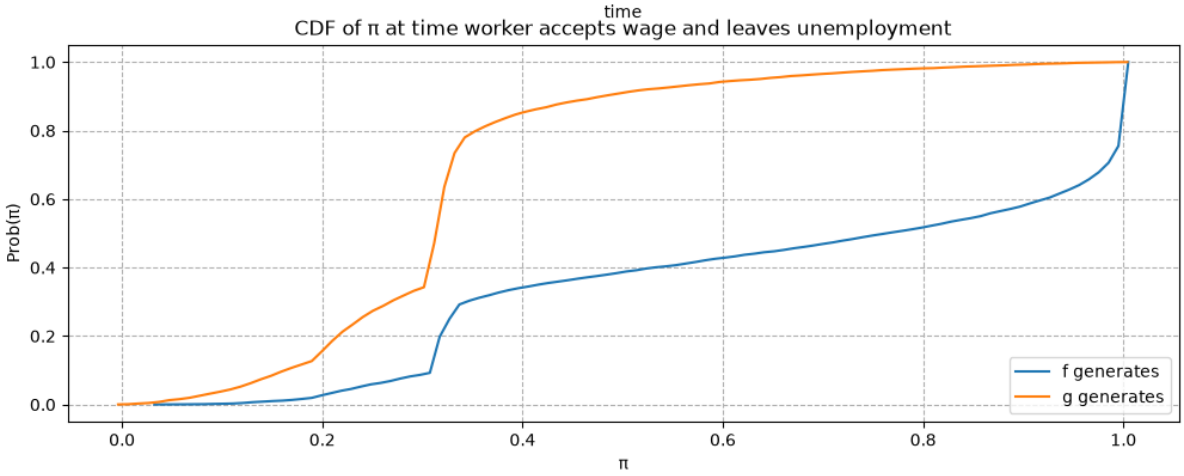

w
l(w) = f(w)/g(w)
l


accept
reject
w
π


w
f(w), g(w)
f
g


$\mathbb{P}\{w > \overline{w}(\pi)\}$
π
f
g


CDF of duration of unemployment
Prob(time)
time
f generates
g generates


CDF of π at time worker accepts wage and leaves unemployment
Prob(π)
π
f generates
g generates

Part VIII

Reinforcement Learning

CHAPTER

SIXTYTHREE

INVENTORY MANAGEMENT VIA Q-LEARNING

Contents

- *Inventory Management via Q-Learning*
  - *Introduction*
  - *The Model*
  - *Solving via Value Function Iteration*
  - *Q-Learning*

63.1 Introduction

In this lecture we study a classic inventory management problem.

A firm must decide how much stock to order each period, facing uncertain demand and a trade-off between lost sales and ordering costs.

We approach the problem in two ways.

First, we solve it exactly using dynamic programming, assuming full knowledge of the model – the demand distribution, cost parameters, and transition dynamics.

Second, we show how a manager can learn the optimal policy from experience alone, using Q-learning.

In this setting, we assume that the manager observes only

- the inventory level,
- the order placed,
- the resulting profit, and
- the next inventory level.

The manager knows the interest rate – and hence the discount factor – but not any of the other underlying parameters.

A key idea is the *Q-factor* representation, which reformulates the Bellman equation so that the optimal policy can be recovered without knowledge of the transition dynamics.

We show that, given enough experience, the manager's learned policy converges to the optimal one.

The lecture proceeds as follows:

1. We set up the inventory model and solve it exactly via value function iteration.
2. We introduce Q-factors and derive the Q-factor Bellman equation.
3. We implement Q-learning and show the learned policy converges to the optimal one.

A risk-sensitive extension of this model is studied in *Risk-Sensitive Inventory Management via Q-Learning*.

We will use the following imports:

```
import numpy as np
import numba
import matplotlib.pyplot as plt
from typing import NamedTuple
```

63.2 The Model

We study a firm where a manager tries to maximize shareholder value by controlling inventories.

To simplify the problem, we assume that the firm only sells one product.

Letting π_t be profits at time t and $r > 0$ be the interest rate, the value of the firm is

$$V_0 = \mathbb{E} \sum_{t \geq 0} \beta^t \pi_t \qquad \text{where} \qquad \beta := \frac{1}{1+r}.$$

Suppose the firm faces exogenous demand process $(D_t)_{t \geq 0}$.

We assume $(D_t)_{t \geq 0}$ is IID with common distribution ϕ on $\{0, 1, ...\}$.

Inventory $(X_t)_{t \geq 0}$ of the product obeys

$$X_{t+1} = h(X_t, A_t, D_{t+1}) \qquad \text{where} \quad h(x, a, d) := \max(x - d, 0) + a.$$

The term A_t is units of stock ordered this period, which arrive at the start of period $t + 1$, after demand D_{t+1} is realized and served:

- observe $X_t \to$ choose $A_t \to$ demand D_{t+1} arrives $\to$ profit realized $\to X_{t+1}$ determined.

(We use a t subscript in A_t to indicate the information set: it is chosen before D_{t+1} is observed.)

We assume that the firm can store at most K items at one time.

Profits are given by

$$\pi(X_t, A_t, D_{t+1}) := \min(X_t, D_{t+1}) - cA_t - \kappa 1\{A_t > 0\}.$$

Here

- the sales price is set to unity (for convenience)
- revenue is the minimum of current stock and demand because orders in excess of inventory are lost (not back-filled)
- c is unit product cost and κ is a fixed cost of ordering inventory

We can map our inventory problem into a dynamic program with state space $\mathsf{X} := \{0, \ldots, K\}$ and action space $\mathsf{A} := \mathsf{X}$.

The feasible correspondence Γ is

$$\Gamma(x) := \{0, \ldots, K - x\}.$$

This represents the set of feasible orders when the current inventory state is x.

The Bellman equation takes the form

$$v(x) = \max_{a \in \Gamma(x)} \mathbb{E}\left[\pi(x, a, D) + \beta v(h(x, a, D))\right]. \tag{63.1}$$

Here D is a random variable with distribution ϕ.

63.3 Solving via Value Function Iteration

Let's start in a setting where the manager knows all parameters, functional forms, and distributions.

She solves the model numerically using value function iteration (VFI).

The idea is to iterate on the Bellman operator T defined by

$$(Tv)(x) = \max_{a \in \Gamma(x)} \sum_d \phi(d) \left[\pi(x, a, d) + \beta\, v(h(x, a, d))\right]$$

starting from an initial guess v_0.

Then, taking the output v from this iteration process, we compute a v**-greedy policy** σ, which obeys

$$\sigma(x) \in \arg\max_{a \in \Gamma(x)} \sum_d \phi(d) \left[\pi(x, a, d) + \beta\, v(h(x, a, d))\right]$$

When $r > 0$ (equivalently, $\beta < 1$), the sequence $v_{k+1} = Tv_k$ converges to the unique fixed point v^*, which is the value function of the optimal policy (see, e.g., [Sargent and Stachurski, 2025]).

63.3.1 Model specification

We store the model primitives in a `NamedTuple`.

Demand follows a geometric distribution with parameter p, so $\phi(d) = (1-p)^d\, p$ for $d = 0, 1, 2, \ldots$.

```
class Model(NamedTuple):
    x_values: np.ndarray        # Inventory values
    d_values: np.ndarray        # Demand values for summation
    ϕ_values: np.ndarray        # Demand probabilities
    p: float                    # Demand parameter
    c: float                    # Unit cost
    κ: float                    # Fixed cost
    β: float                    # Discount factor
```

The function below constructs a `Model` instance with default parameters.

We truncate the demand distribution at `D_MAX` for computational purposes.

```
def create_sdd_inventory_model(
        K: int = 20,          # Max inventory
        D_MAX: int = 21,      # Demand upper bound for summation
        p: float = 0.7,
        c: float = 0.2,
        κ: float = 0.8,
        β: float = 0.98
    ) -> Model:

    def demand_pdf(p, d):
```

(continues on next page)

(continued from previous page)

```
        return (1 - p)**d * p

    d_values = np.arange(D_MAX)
    ϕ_values = demand_pdf(p, d_values)      # ϕ_0, ϕ_1,...
    x_values = np.arange(K + 1)             # 0, 1, ..., K

    return Model(x_values, d_values, ϕ_values, p, c, κ, β)
```

63.3.2 The Bellman operator

The core computation is the Bellman operator T.

For each inventory level x, we loop over all feasible order quantities $a \in \{0, \ldots, K - x\}$ and compute the expected value

$$\sum_d \phi(d) \left[\min(x, d) - c\,a - \kappa\,\mathbf{1}\{a > 0\} + \beta\, v(\max(x - d, 0) + a)\right].$$

We then take the maximum over a.

The inner loops are compiled with Numba for performance.

```
@numba.jit(nopython=True)
def T_kernel(v, d_values, ϕ_values, c, κ, β, K):
    new_v = np.empty(K + 1)
    for x in range(K + 1):
        best = -np.inf
        for a in range(K - x + 1):           # loop over feasible actions
            val = 0.0
            for i in range(len(d_values)):   # compute expectation over demand
                d = d_values[i]
                x_next = max(x - d, 0) + a
                revenue = min(x, d)
                cost = c * a + κ * (a > 0)
                val += ϕ_values[i] * (revenue - cost + β * v[x_next])
            if val > best:
                best = val
        new_v[x] = best
    return new_v
```

The wrapper function `T` unpacks the model and calls the compiled kernel.

```
def T(v, model):
    """The Bellman operator."""
    x_values, d_values, ϕ_values, p, c, κ, β = model
    K = len(x_values) - 1
    return T_kernel(v, d_values, ϕ_values, c, κ, β, K)
```

63.3.3 Computing the greedy policy

Recall that, given a value function v, the **v-greedy policy** is computed via

$$\sigma(x) = \arg\max_{a \in \Gamma(x)} \sum_d \phi(d) \left[\pi(x, a, d) + \beta \, v(h(x, a, d)) \right].$$

The structure is the same as the Bellman operator, except we record the maximizing action rather than the maximized value.

```
@numba.jit(nopython=True)
def get_greedy_kernel(v, d_values, ϕ_values, c, κ, β, K):
    σ = np.empty(K + 1, dtype=np.int32)
    for x in range(K + 1):
        best = -np.inf
        best_a = 0
        for a in range(K - x + 1):
            val = 0.0
            for i in range(len(d_values)):
                d = d_values[i]
                x_next = max(x - d, 0) + a
                revenue = min(x, d)
                cost = c * a + κ * (a > 0)
                val += ϕ_values[i] * (revenue - cost + β * v[x_next])
            if val > best:
                best = val
                best_a = a
        σ[x] = best_a
    return σ
```

```
def get_greedy(v, model):
    """Get a v-greedy policy."""
    x_values, d_values, ϕ_values, p, c, κ, β = model
    K = len(x_values) - 1
    return get_greedy_kernel(v, d_values, ϕ_values, c, κ, β, K)
```

63.3.4 Value function iteration

We iterate $v_{k+1} = Tv_k$ until convergence, starting from $v_0 = 0$.

Once the value function has converged (to within tolerance `tol`), we extract the optimal policy σ^* via `get_greedy`.

```
def solve_inventory_model(v_init, model, max_iter=10_000, tol=1e-6):
    v = v_init.copy()
    i, error = 0, tol + 1

    while i < max_iter and error > tol:
        new_v = T(v, model)
        error = np.max(np.abs(new_v - v))
        i += 1
        v = new_v

    print(f"Converged in {i} iterations with error {error:.2e}")

    σ = get_greedy(v, model)
    return v, σ
```

63.3.5 Creating and solving an instance

```
model = create_sdd_inventory_model()
x_values, d_values, ϕ_values, p, c, κ, β = model
n_x = len(x_values)
v_init = np.zeros(n_x)
```

```
v_star, σ_star = solve_inventory_model(v_init, model)
```

```
Converged in 623 iterations with error 9.95e-07
```

63.3.6 Simulating the optimal policy

To visualize the solution, we simulate the inventory process under the optimal policy σ^*.

At each step, we draw a demand shock from the geometric distribution and update the state via h.

```
@numba.jit(nopython=True)
def sim_inventories(ts_length, σ, p, rng, X_init=0):
    """Simulate inventory dynamics under policy σ."""
    X = np.zeros(ts_length, dtype=np.int32)
    X[0] = X_init
    for t in range(ts_length - 1):
        d = rng.geometric(p) - 1
        X[t+1] = max(X[t] - d, 0) + σ[X[t]]
    return X
```

The plot below shows a typical inventory path under the optimal policy.

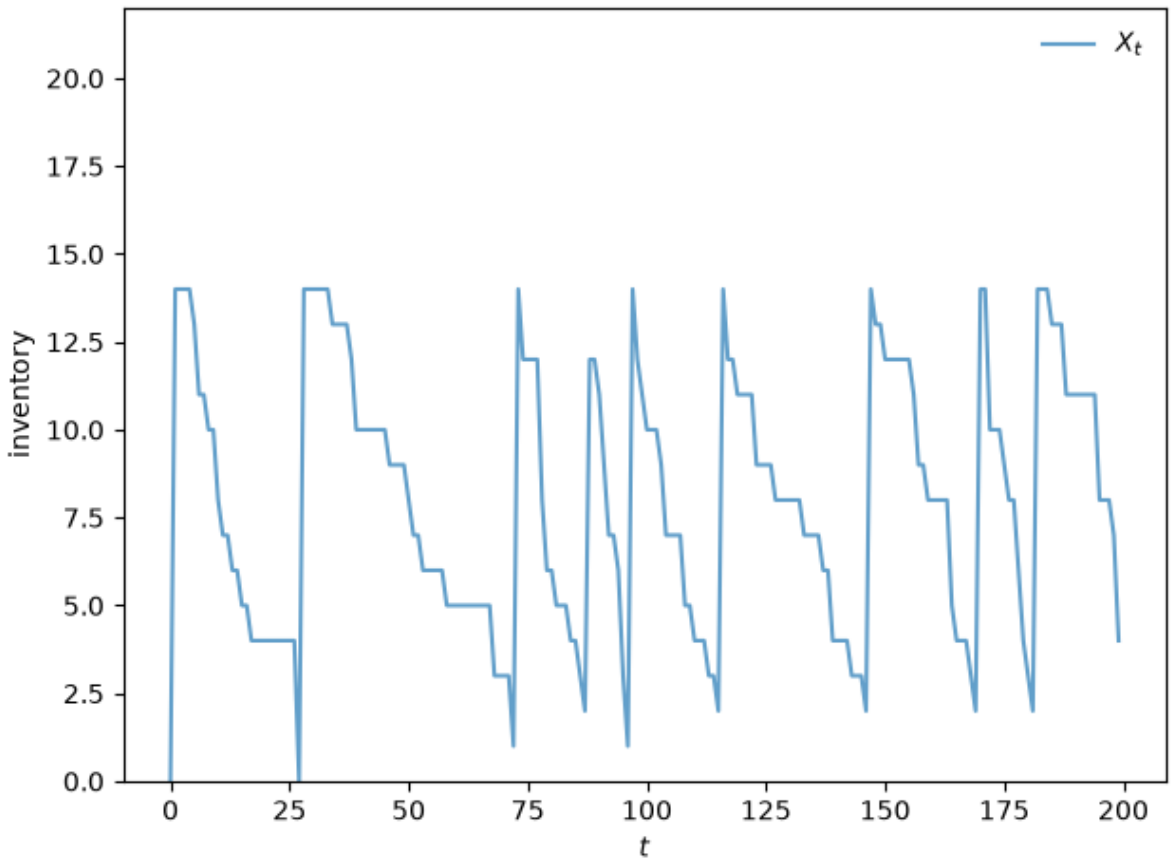
Notice the **S-s pattern**: when inventory falls to a low level, the firm places a large order to replenish stock (the upward jumps), after which inventory gradually declines as demand is served.

```
def plot_ts(ts_length=200, fontsize=10, seed=0):
    X = sim_inventories(ts_length, σ_star, p, np.random.default_rng(seed))
    fig, ax = plt.subplots()

    ax.plot(X, label=r"$X_t$", alpha=0.7)
    ax.set_xlabel(r"$t$", fontsize=fontsize)
    ax.set_ylabel("inventory", fontsize=fontsize)
    ax.legend(fontsize=fontsize, frameon=False)
    ax.set_ylim(0, len(σ_star) + 1)

    plt.tight_layout()
    plt.show()
```

```
plot_ts()
```



63.4 Q-Learning

We now ask: can an agent **learn** the optimal policy without knowing the model?

In particular, suppose the agent does not know the demand distribution ϕ, the cost parameters c and κ, or the transition function h.

Instead, the agent only observes the sequence of states, actions, and profits as it interacts with the environment.

63.4.1 The Q-factor Bellman equation

The first step of Q-learning is to modify the Bellman equation, placing it in a form that allows learning from this limited information.

Rather than working with the value function $v(x)$, we work with the **Q-function** (or Q-factor) $q(x, a)$.

We define q in terms of the value function v^* as

$$q(x, a) := \mathbb{E}\left[\pi(x, a, D) + \beta\, v^*(h(x, a, D))\right].$$

In words, $q(x, a)$ is the expected value of taking action a in state x and then following the optimal policy thereafter.

Note that the Bellman equation (63.1) can be written as

$$v^*(x) = \max_{a \in \Gamma(x)} q(x, a).$$

Substituting this back into the definition of q, we can eliminate v^* and obtain a fixed point equation in q alone:

$$q(x,a) = \mathbb{E}\left[\pi(x,a,D) + \beta \max_{a' \in \Gamma(x')} q(x',a')\right]$$

where $x' = h(x,a,D)$.

One advantage of working with q is that the optimal policy can be read off directly as $\sigma(x) = \arg\max_a q(x,a)$, without needing to know the transition function.

63.4.2 The Q-learning update rule

Q-learning approximates the fixed point of the Q-factor Bellman equation using **stochastic approximation**.

At each step, the agent is in state x, takes action a, observes reward $R_{t+1} = \pi(x,a,D_{t+1})$ and next state $X_{t+1} = h(x,a,D_{t+1})$, and updates

$$q_{t+1}(x,a) = (1-\alpha_t) q_t(x,a) + \alpha_t \left(R_{t+1} + \beta \max_{a' \in \Gamma(X_{t+1})} q_t(X_{t+1}, a')\right)$$

where α_t is the learning rate.

The update blends the current estimate $q_t(x,a)$ with a fresh sample of the Bellman target.

63.4.3 What the manager needs to know

Notice what is **not** required to implement the update.

The manager does not need to know the demand distribution ϕ, the unit cost c, the fixed cost κ, or the transition function h.

All the manager needs to observe at each step is:

1. the current inventory level x,
2. the order quantity a, which they choose,
3. the resulting profit R_{t+1} (which appears on the books),
4. the discount factor β, which is determined by the interest rate, and
5. the next inventory level X_{t+1} (which they can read off the warehouse).

These are all directly observable quantities – no model knowledge is required.

63.4.4 The Q-table and the role of the max

It is important to understand how the update rule relates to the manager's actions.

The manager maintains a **Q-table** – a lookup table storing an estimate $q_t(x,a)$ for every state-action pair (x,a).

At each step, the manager is in some state x and must choose a specific action a to take. Whichever a is chosen, the manager observes profit R_{t+1} and next state X_{t+1}, and updates *that one entry* $q_t(x,a)$ of the table using the rule above.

It is tempting to read the $\max_{a'}$ in the update rule as prescribing the manager's next action – that is, to interpret the update as saying "move to state X_{t+1} and take an action in $\arg\max_{a'} q_t(X_{t+1}, a')$."

But the max plays a different role.

The quantity $\max_{a' \in \Gamma(X_{t+1})} q_t(X_{t+1}, a')$ is just an estimate of the value of being in state X_{t+1} under the best possible continuation.

This scalar enters the update as part of the target value for $q_t(x, a)$.

Which action the manager *actually takes* at time $t+1$ is a separate decision.

In short, the max is doing the work of finding the optimum; it does not dictate the action that the manager actually takes.

63.4.5 The behavior policy

The rule governing how the manager chooses actions is called the **behavior policy**.

Because the max in the update target always points toward q^* regardless of how the manager selects actions, the behavior policy affects only which (x, a) entries get visited – and hence updated – over time.

In the reinforcement learning literature, this property is called **off-policy** learning: the convergence target (q^*) does not depend on the behavior policy.

As long as every (x, a) pair is visited infinitely often (so that every entry of the Q-table receives infinitely many updates) and the learning rates satisfy standard conditions (see below), the Q-table converges to q^*.

The behavior policy affects the *speed* of convergence – visiting important state-action pairs more frequently leads to faster learning – but not the *limit*.

In practice, we want the manager to mostly take good actions (to earn reasonable profits while learning), while still occasionally experimenting to discover better alternatives.

63.4.6 Learning rate

We use $\alpha_t = 1/n_t(x, a)^{0.51}$, where $n_t(x, a)$ is the number of times the pair (x, a) has been visited up to time t.

This decays slowly enough to allow learning from later (better-informed) updates, while still satisfying the Robbins–Monro conditions for convergence.

63.4.7 Exploration: epsilon-greedy

For our behavior policy, we use an ε-greedy strategy:

- With probability ε, choose a random feasible action (explore).
- With probability $1 - \varepsilon$, choose the action with the highest current q-value (exploit).

The exploration ensures that every state-action pair is visited, which is needed for convergence.

The exploitation ensures the manager earns reasonable profits while learning.

We decay ε each step: $\varepsilon_{t+1} = \max(\varepsilon_{\min},\ \varepsilon_t \cdot \lambda)$, so the manager experiments widely early on and increasingly relies on learned q-values as experience accumulates.

The stochastic demand shocks naturally drive the manager across different inventory levels, providing exploration over the state space without any artificial resets.

63.4.8 Optimistic initialization

A simple but powerful technique for accelerating learning is **optimistic initialization**: instead of starting the Q-table at zero, we initialize every entry to a value above the true optimum.

Because every untried action looks optimistically good, the agent is "disappointed" whenever it tries one – the update pulls that entry down toward reality. This drives the agent to try other actions (which still look optimistically high), producing broad exploration of the state-action space early in training.

This idea is sometimes called **optimism in the face of uncertainty** and is widely used in both bandit and reinforcement learning settings.

In our problem, the value function v^* ranges from about 13 to 18. We initialize the Q-table at 20 – modestly above the true maximum – to ensure optimistic exploration without being so extreme as to distort learning.

63.4.9 Implementation

We first define a helper to extract the greedy policy from a Q-table.

```
@numba.jit(nopython=True)
def greedy_policy_from_q(q, K):
    """Extract greedy policy from Q-table."""
    σ = np.empty(K + 1, dtype=np.int32)
    for x in range(K + 1):
        best_val = -np.inf
        best_a = 0
        for a in range(K - x + 1):
            if q[x, a] > best_val:
                best_val = q[x, a]
                best_a = a
        σ[x] = best_a
    return σ
```

The Q-learning loop runs for `n_steps` total steps in a single continuous trajectory – just as a real manager would learn from the ongoing stream of data.

At specified step counts (given by `snapshot_steps`), we record the current greedy policy.

```
@numba.jit(nopython=True)
def q_learning_kernel(K, p, c, κ, β, n_steps, X_init,
                      ε_init, ε_min, ε_decay, q_init, snapshot_steps, rng):
    q = np.full((K + 1, K + 1), q_init)
    n = np.zeros((K + 1, K + 1))       # visit counts for learning rate
    ε = ε_init

    n_snaps = len(snapshot_steps)
    snapshots = np.zeros((n_snaps, K + 1), dtype=np.int32)
    snap_idx = 0

    # Initialize state and action
    x = X_init
    a = rng.integers(0, K - x + 1)

    for t in range(n_steps):
        # Record policy snapshot if needed
        if snap_idx < n_snaps and t == snapshot_steps[snap_idx]:
            snapshots[snap_idx] = greedy_policy_from_q(q, K)
```

(continues on next page)

(continued from previous page)

```
            snap_idx += 1

        # === Draw D_{t+1} and observe outcome ===
        d = rng.geometric(p) - 1
        reward = min(x, d) - c * a - κ * (a > 0)
        x_next = max(x - d, 0) + a

        # === Max over next state (scalar value for update target) ===
        # Also record the argmax action for use by the behavior policy.
        best_next = -np.inf
        a_next = 0
        for aa in range(K - x_next + 1):
            if q[x_next, aa] > best_next:
                best_next = q[x_next, aa]
                a_next = aa

        # === Q-learning update (uses best_next, the max value) ===
        n[x, a] += 1
        α = 1.0 / n[x, a] ** 0.51
        q[x, a] = (1 - α) * q[x, a] + α * (reward + β * best_next)

        # === Behavior policy: ε-greedy (uses a_next, the argmax action) ===
        x = x_next
        if rng.random() < ε:
            a = rng.integers(0, K - x + 1)
        else:
            a = a_next
        ε = max(ε_min, ε * ε_decay)

    return q, snapshots
```

The wrapper function unpacks the model and provides default hyperparameters.

```
def q_learning(model, n_steps=20_000_000, X_init=0,
               ε_init=1.0, ε_min=0.01, ε_decay=0.999999,
               q_init=20.0, snapshot_steps=None, seed=1234):
    x_values, d_values, ϕ_values, p, c, κ, β = model
    K = len(x_values) - 1
    if snapshot_steps is None:
        snapshot_steps = np.array([], dtype=np.int64)
    rng = np.random.default_rng(seed)
    return q_learning_kernel(K, p, c, κ, β, n_steps, X_init,
                             ε_init, ε_min, ε_decay, q_init, snapshot_steps, rng)
```

Next we run n = 5 million steps and take policy snapshots at steps 10,000, 1,000,000, and n.

```
n = 5_000_000
snap_steps = np.array([10_000, 1_000_000, n], dtype=np.int64)
q, snapshots = q_learning(model, n_steps=n+1, snapshot_steps=snap_steps)
```

63.4.10 Comparing with the exact solution

We extract the value function and policy from the final Q-table via

$$v_Q(x) = \max_{a \in \Gamma(x)} q(x,a), \qquad \sigma_Q(x) = \arg\max_{a \in \Gamma(x)} q(x,a)$$

and compare them against v^* and σ^* from VFI.

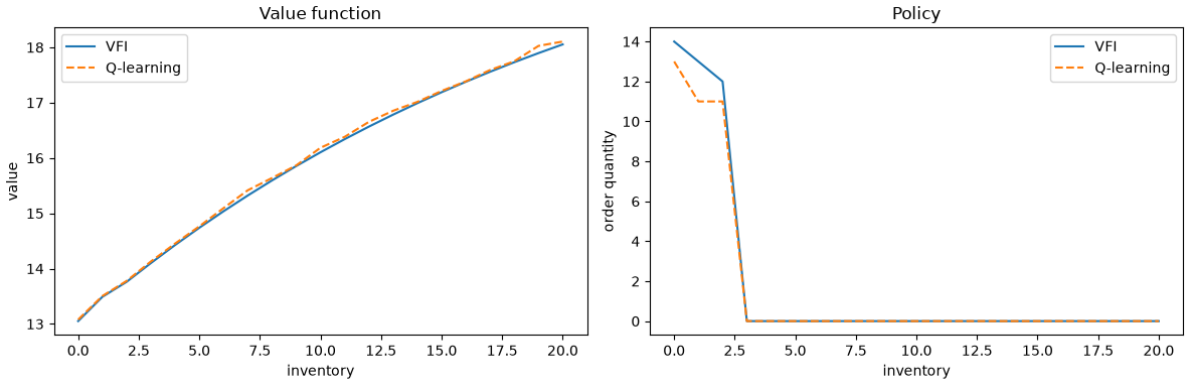
```
K = len(x_values) - 1
# restrict to feasible actions a ∈ {0, ..., K-x}
v_q = np.array([np.max(q[x, :K - x + 1]) for x in range(K + 1)])
σ_q = np.array([np.argmax(q[x, :K - x + 1]) for x in range(K + 1)])
```

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].plot(x_values, v_star, label="VFI")
axes[0].plot(x_values, v_q, '--', label="Q-learning")
axes[0].set_xlabel("inventory")
axes[0].set_ylabel("value")
axes[0].legend()
axes[0].set_title("Value function")

axes[1].plot(x_values, σ_star, label="VFI")
axes[1].plot(x_values, σ_q, '--', label="Q-learning")
axes[1].set_xlabel("inventory")
axes[1].set_ylabel("order quantity")
axes[1].legend()
axes[1].set_title("Policy")

plt.tight_layout()
plt.show()
```



63.4.11 Visualizing learning over time

The panels below show how the agent's behavior evolves during training.

Each panel simulates an inventory path using the greedy policy extracted from the Q-table at a given training step.

All panels use the **same demand sequence** (via a fixed random seed), so differences reflect only changes in the policy.

The top panel shows the optimal policy from VFI for reference.

After 10,000 steps the agent has barely explored and its policy is poor.

By 1,000,000 steps the policy has improved but still differs noticeably from the optimum.

By step 5 million, the learned policy produces inventory dynamics that closely resemble the S-s pattern of the optimal solution.

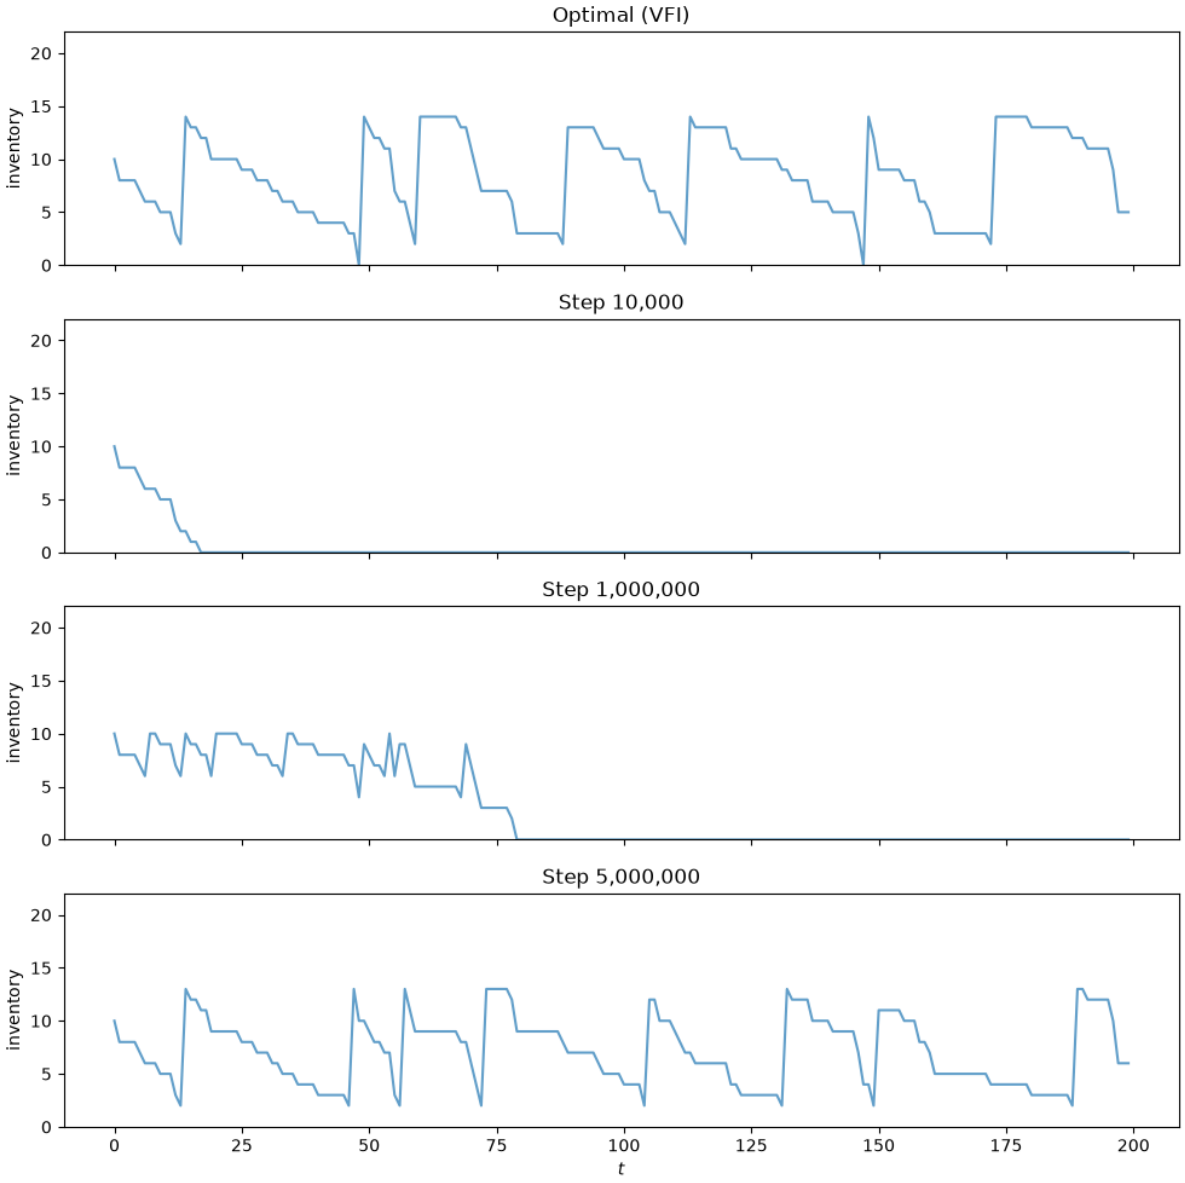
```
ts_length = 200
n_snaps = len(snap_steps)
fig, axes = plt.subplots(n_snaps + 1, 1, figsize=(10, 2.5 * (n_snaps + 1)),
                         sharex=True)

X_init = K // 2
sim_seed = 5678

# Optimal policy
X_opt = sim_inventories(ts_length, σ_star, p,
                        np.random.default_rng(sim_seed), X_init)
axes[0].plot(X_opt, alpha=0.7)
axes[0].set_ylabel("inventory")
axes[0].set_title("Optimal (VFI)")
axes[0].set_ylim(0, K + 2)

# Q-learning snapshots
for i in range(n_snaps):
    σ_snap = snapshots[i]
    X = sim_inventories(ts_length, σ_snap, p,
                        np.random.default_rng(sim_seed), X_init)
    axes[i + 1].plot(X, alpha=0.7)
    axes[i + 1].set_ylabel("inventory")
    axes[i + 1].set_title(f"Step {snap_steps[i]:,}")
    axes[i + 1].set_ylim(0, K + 2)

axes[-1].set_xlabel(r"$t$")
plt.tight_layout()
plt.show()
```


Optimal (VFI)
Step 10,000
Step 1,000,000
Step 5,000,000
inventory
0
5
10
15
20
0
25
50
75
100
125
150
175
200
t

CHAPTER

SIXTYFOUR

RISK-SENSITIVE INVENTORY MANAGEMENT VIA Q-LEARNING

Contents

- *Risk-Sensitive Inventory Management via Q-Learning*
  - *Introduction*
  - *The Model*
  - *Solving via Value Function Iteration*
  - *Interpreting the Outcomes*
  - *Q-Learning*
  - *Conclusion*

64.1 Introduction

In *Inventory Management via Q-Learning*, we looked at an inventory management problem and solved it with both value function iteration and Q-learning.

In this lecture, we consider a risk-sensitive variation.

Injection of risk-sensitivity acknowledges the fact that, in incomplete markets with financial and informational frictions, firms typically take risk into account in their decision making.

In other words, the actions of firms are not, in general, risk neutral.

One natural way to handle this is to use a risk-sensitive version of the Bellman equation.

We show how the model can be solved using value function iteration.

We then investigate how risk sensitivity affects the optimal policy.

```
import numpy as np
import numba
import matplotlib.pyplot as plt
from typing import NamedTuple
```

64.2 The Model

The Bellman equation for the inventory management problem in *Inventory Management via Q-Learning* has the form

$$v(x) = \max_{a \in \Gamma(x)} \mathbb{E}\left[\pi(x, a, D) + \beta v(h(x, a, D))\right].$$

Here D is a random variable with distribution ϕ.

(Primitives and definitions are the same as in *Inventory Management via Q-Learning*.)

The risk-sensitive version of this Bellman equation has the form

$$v(x) = \max_{a \in \Gamma(x)} \psi^{-1}\left\{\mathbb{E}\psi\left[\pi(x, a, D) + \beta v(h(x, a, D))\right]\right\},$$

where $\psi(t) = \exp(-\gamma t)$ for fixed $\gamma > 0$.

Since $\psi^{-1}(y) = -\frac{1}{\gamma}\ln(y)$, the Bellman equation becomes

$$v(x) = \max_{a \in \Gamma(x)} \left(-\frac{1}{\gamma}\right) \ln\left\{\sum_d \phi(d) \exp\left[-\gamma\left(\pi(x, a, d) + \beta\, v(h(x, a, d))\right)\right]\right\}.$$

Here $\phi(d)$ denotes the demand probability mass function, as in *Inventory Management via Q-Learning*.

The parameter γ controls the degree of risk sensitivity.

As $\gamma \to 0$, the certainty equivalent reduces to the ordinary expectation and we recover the risk-neutral case.

Larger γ means more aversion to downside risk.

The Bellman operator, greedy policy, and VFI algorithm all carry over from the risk-neutral case, with the expectation replaced by the certainty equivalent.

64.3 Solving via Value Function Iteration

64.3.1 Model specification

We reuse the same model primitives as in *Inventory Management via Q-Learning*, adding γ as a parameter.

```
class RSModel(NamedTuple):
    x_values: np.ndarray        # Inventory values
    d_values: np.ndarray        # Demand values for summation
    ϕ_values: np.ndarray        # Demand probabilities
    p: float                    # Demand parameter
    c: float                    # Unit cost
    κ: float                    # Fixed cost
    β: float                    # Discount factor
    γ: float                    # Risk-sensitivity parameter
```

```
def create_rs_inventory_model(
        K: int = 20,          # Max inventory
        D_MAX: int = 21,      # Demand upper bound for summation
        p: float = 0.7,
        c: float = 0.2,
        κ: float = 0.8,
        β: float = 0.98,
```

(continues on next page)

(continued from previous page)

```
        γ: float = 1.0
    ) -> RSModel:

    def demand_pdf(p, d):
        return (1 - p)**d * p

    d_values = np.arange(D_MAX)
    ϕ_values = demand_pdf(p, d_values)
    x_values = np.arange(K + 1)

    return RSModel(x_values, d_values, ϕ_values, p, c, κ, β, γ)
```

64.3.2 The Bellman operator

The risk-sensitive Bellman operator replaces the expected value with the certainty equivalent.

For numerical stability, we use the log-sum-exp trick: given values $z_i = \pi(x, a, d_i) + \beta\, v(h(x, a, d_i))$, we compute

$$-\frac{1}{\gamma} \ln \sum_i \phi(d_i) \exp(-\gamma z_i) \;=\; -\frac{1}{\gamma}\left(m + \ln \sum_i \phi(d_i) \exp(-\gamma z_i - m) \right),$$

where $m = \max_i(-\gamma z_i)$.

```
@numba.jit(nopython=True)
def T_rs_kernel(v, d_values, ϕ_values, c, κ, β, γ, K):
    new_v = np.empty(K + 1)
    n_d = len(d_values)
    for x in range(K + 1):
        best = -np.inf
        for a in range(K - x + 1):
            # Compute -γ * z_i for each demand realization
            exponents = np.empty(n_d)
            for i in range(n_d):
                d = d_values[i]
                x_next = max(x - d, 0) + a
                revenue = min(x, d)
                cost = c * a + κ * (a > 0)
                z_i = revenue - cost + β * v[x_next]
                exponents[i] = -γ * z_i
            # Log-sum-exp trick for numerical stability
            m = np.max(exponents)
            weighted_sum = 0.0
            for i in range(n_d):
                weighted_sum += ϕ_values[i] * np.exp(exponents[i] - m)
            val = -(1.0 / γ) * (m + np.log(weighted_sum))
            if val > best:
                best = val
        new_v[x] = best
    return new_v
```

```
def T_rs(v, model):
    """The risk-sensitive Bellman operator."""
    x_values, d_values, ϕ_values, p, c, κ, β, γ = model
    K = len(x_values) - 1
    return T_rs_kernel(v, d_values, ϕ_values, c, κ, β, γ, K)
```

64.3.3 Computing the greedy policy

The greedy policy records the maximizing action instead of the maximized value.

```
@numba.jit(nopython=True)
def get_greedy_rs_kernel(v, d_values, ϕ_values, c, κ, β, γ, K):
    σ = np.empty(K + 1, dtype=np.int32)
    n_d = len(d_values)
    for x in range(K + 1):
        best = -np.inf
        best_a = 0
        for a in range(K - x + 1):
            exponents = np.empty(n_d)
            for i in range(n_d):
                d = d_values[i]
                x_next = max(x - d, 0) + a
                revenue = min(x, d)
                cost = c * a + κ * (a > 0)
                z_i = revenue - cost + β * v[x_next]
                exponents[i] = -γ * z_i
            m = np.max(exponents)
            weighted_sum = 0.0
            for i in range(n_d):
                weighted_sum += ϕ_values[i] * np.exp(exponents[i] - m)
            val = -(1.0 / γ) * (m + np.log(weighted_sum))
            if val > best:
                best = val
                best_a = a
        σ[x] = best_a
    return σ
```

```
def get_greedy_rs(v, model):
    """Get a v-greedy policy for the risk-sensitive model."""
    x_values, d_values, ϕ_values, p, c, κ, β, γ = model
    K = len(x_values) - 1
    return get_greedy_rs_kernel(v, d_values, ϕ_values, c, κ, β, γ, K)
```

64.3.4 Value function iteration

```
def solve_rs_inventory_model(v_init, model, max_iter=10_000, tol=1e-6):
    v = v_init.copy()
    i, error = 0, tol + 1

    while i < max_iter and error > tol:
        new_v = T_rs(v, model)
        error = np.max(np.abs(new_v - v))
        i += 1
        v = new_v

    print(f"Converged in {i} iterations with error {error:.2e}")

    σ = get_greedy_rs(v, model)
    return v, σ
```

64.3.5 Creating and solving an instance

```
model = create_rs_inventory_model()
x_values = model.x_values
n_x = len(x_values)
v_init = np.zeros(n_x)
```

```
v_star, σ_star = solve_rs_inventory_model(v_init, model)
```

```
Converged in 600 iterations with error 9.92e-07
```

64.3.6 Effect of risk sensitivity on the optimal policy

We solve the model for several values of γ and compare the resulting policies.

As we will see, a risk-sensitive firm orders less aggressively than a risk-neutral one.

```
γ_values = [0.01, 1.0, 2.0]
results = {}

for γ in γ_values:
    mod = create_rs_inventory_model(γ=γ)
    v, σ = solve_rs_inventory_model(np.zeros(n_x), mod)
    results[γ] = (v, σ)
```

```
Converged in 623 iterations with error 9.90e-07
Converged in 600 iterations with error 9.92e-07
Converged in 583 iterations with error 9.96e-07
```

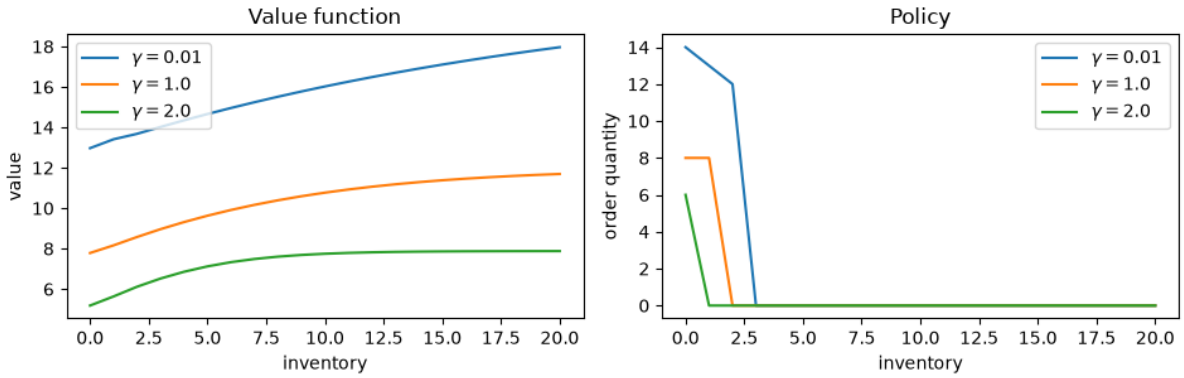
```
fig, axes = plt.subplots(1, 2, figsize=(9.6, 3.2))

for γ in γ_values:
    v, σ = results[γ]
    axes[0].plot(x_values, v, label=f"$\\gamma = {γ}$")
    axes[1].plot(x_values, σ, label=f"$\\gamma = {γ}$")

axes[0].set_xlabel("inventory")
axes[0].set_ylabel("value")
axes[0].legend()
axes[0].set_title("Value function")

axes[1].set_xlabel("inventory")
axes[1].set_ylabel("order quantity")
axes[1].legend()
axes[1].set_title("Policy")

plt.tight_layout()
plt.show()
```



64.3.7 Simulating the optimal policy

We simulate inventory dynamics under the optimal policy for the baseline γ.

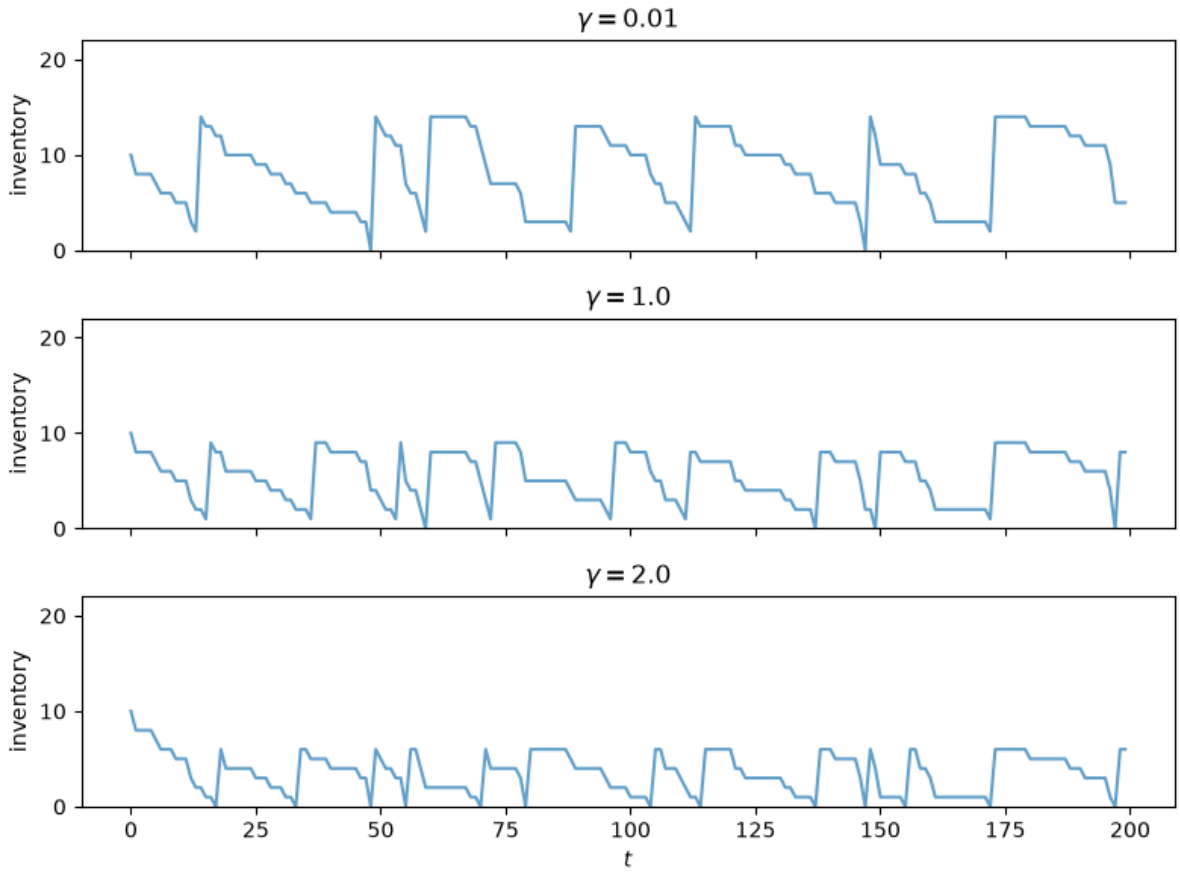
```
@numba.jit(nopython=True)
def sim_inventories(ts_length, σ, p, rng, X_init=0):
    """Simulate inventory dynamics under policy σ."""
    X = np.zeros(ts_length, dtype=np.int32)
    X[0] = X_init
    for t in range(ts_length - 1):
        d = rng.geometric(p) - 1
        X[t+1] = max(X[t] - d, 0) + σ[X[t]]
    return X
```

```
fig, axes = plt.subplots(len(γ_values), 1,
                         figsize=(8, 2.0 * len(γ_values)),
                         sharex=True)

ts_length = 200
sim_seed = 5678
K = len(x_values) - 1

for i, γ in enumerate(γ_values):
    v, σ = results[γ]
    X = sim_inventories(ts_length, σ, model.p,
                        np.random.default_rng(sim_seed), X_init=K // 2)
    axes[i].plot(X, alpha=0.7)
    axes[i].set_ylabel("inventory")
    axes[i].set_title(f"$\\gamma = {γ}$")
    axes[i].set_ylim(0, K + 2)

axes[-1].set_xlabel(r"$t$")
plt.tight_layout()
plt.show()
```



64.4 Interpreting the Outcomes

The plots above show that a more risk-sensitive firm (larger γ) orders less inventory and maintains lower stock levels.

At first glance this may seem surprising: wouldn't holding more inventory reduce variance by ensuring the firm can always meet demand?

The key is to identify where the randomness in profits actually comes from.

Recall that per-period profit is $\pi(x, a, d) = \min(x, d) - ca - \kappa \mathbf{1}\{a > 0\}$.

The ordering cost $ca + \kappa \mathbf{1}\{a > 0\}$ is **deterministic** – it is chosen before the demand shock is realized.

So higher ordering shifts the level of profits down but does not affect their variance.

The variance comes from **revenue**: $\min(x, D)$.

When inventory x is high, $\min(x, D) \approx D$ for most demand realizations – revenue inherits the full variance of demand.

When inventory x is low, $\min(x, D) \approx x$ for most realizations – revenue is nearly deterministic, capped at the inventory level.

A risk-sensitive agent therefore prefers lower inventory because it **caps the randomness of revenue**.

The agent accepts lower expected sales in exchange for more predictable profits.

There is also a continuation value channel: next-period inventory $\max(x - D, 0) + a$ varies with D, and higher x means $x - D$ tracks D more closely, propagating that variance forward through v.

64.5 Q-Learning

We now ask whether the optimal policy can be learned without knowledge of the model, as we did in the risk-neutral case in *Inventory Management via Q-Learning*.

64.5.1 The Q-factor

The first step is to define the Q-factor in a way that is compatible with the risk-sensitive Bellman equation.

We define

$$q(x,a) := \mathbb{E}\left[\exp(-\gamma\,(\pi(x,a,D) + \beta\, v^*(h(x,a,D))))\right].$$

In words, $q(x,a)$ applies the risk-sensitivity transformation $\psi(t) = \exp(-\gamma t)$ inside the expectation, evaluated at the return from taking action a in state x and following the optimal policy thereafter.

64.5.2 Deriving the Q-factor Bellman equation

Our goal is to obtain a fixed point equation in q alone, eliminating v^*.

Step 1. Express v^* in terms of q.

The risk-sensitive Bellman equation says $v^*(x) = \max_{a\in\Gamma(x)} \psi^{-1}(q(x,a))$.

Since $\psi^{-1}(y) = -\frac{1}{\gamma}\ln(y)$ is a **decreasing** function, the maximum over a of $\psi^{-1}(q(x,a))$ corresponds to the **minimum** over a of $q(x,a)$:

$$v^*(x) = \psi^{-1}\left(\min_{a\in\Gamma(x)} q(x,a)\right) = -\frac{1}{\gamma}\ln\left(\min_{a\in\Gamma(x)} q(x,a)\right).$$

Equivalently,

$$\exp(-\gamma\, v^*(x)) = \min_{a\in\Gamma(x)} q(x,a).$$

Step 2. Substitute back into the definition of q to eliminate v^*.

Expanding the exponential in the definition of q,

$$q(x,a) = \mathbb{E}\left[\exp(-\gamma\,\pi(x,a,D)) \;\cdot\; \exp(-\gamma\beta\, v^*(x'))\right]$$

where $x' = h(x,a,D)$.

From Step 1, $\exp(-\gamma\, v^*(x')) = \min_{a'\in\Gamma(x')} q(x',a')$, so $\exp(-\gamma\beta\, v^*(x')) = \left[\min_{a'\in\Gamma(x')} q(x',a')\right]^{\beta}$.

Substituting,

$$q(x,a) = \mathbb{E}\left[\exp(-\gamma\,\pi(x,a,D)) \;\cdot\; \left(\min_{a'\in\Gamma(x')} q(x',a')\right)^{\beta}\right].$$

This is a fixed point equation in q alone – v^* has been eliminated.

64.5.3 The Q-learning update rule

As in the risk-neutral case, we approximate the fixed point using stochastic approximation.

At each step, the agent is in state x, takes action a, observes profit $R_{t+1} = \pi(x, a, D_{t+1})$ and next state $X_{t+1} = h(x, a, D_{t+1})$, and updates

$$q_{t+1}(x,a) = (1-\alpha_t)\, q_t(x,a) + \alpha_t \left[\exp(-\gamma\, R_{t+1}) \cdot \left(\min_{a' \in \Gamma(X_{t+1})} q_t(X_{t+1}, a') \right)^{\beta} \right].$$

The term in brackets is a single-sample estimate of the right-hand side of the Q-factor Bellman equation.

The update blends the current estimate with this fresh sample, just as in standard Q-learning.

Notice several differences from the risk-neutral case:

- The Q-values are **positive** (expectations of exponentials) rather than signed.
- The optimal policy is $\sigma(x) = \arg\min_a q(x,a)$ – we **minimize** rather than maximize, because ψ^{-1} is decreasing.
- The observed profit enters through $\exp(-\gamma R_{t+1})$ rather than additively.
- The continuation value enters as a **power** $(\min_{a'} q_t)^{\beta}$ rather than a scaled sum $\beta \cdot \max_{a'} q_t$.

As before, the agent needs only to observe x, a, R_{t+1}, and X_{t+1} – no model knowledge is required.

64.5.4 Implementation plan

Our implementation follows the same structure as the risk-neutral Q-learning in *Inventory Management via Q-Learning*, with the modifications above:

1. **Initialize** the Q-table q optimistically (see below) and visit counts n to zeros.
2. **At each step:**
   - Draw demand D_{t+1} and compute observed profit R_{t+1} and next state X_{t+1}.
   - Compute $\min_{a'} q_t(X_{t+1}, a')$ over feasible actions (this is a scalar for the update target, and the arg min action is used by the ε-greedy behavior policy).
   - Update $q_t(x,a)$ using the rule above, with learning rate $\alpha_t = 1/n_t(x,a)^{0.51}$.
   - Choose the next action via ε-greedy: with probability ε pick a random feasible action, otherwise pick the arg min action.
   - Decay ε.
3. **Extract the greedy policy** from the final Q-table via $\sigma(x) = \arg\min_{a \in \Gamma(x)} q(x,a)$.
4. **Compare** the learned policy against the VFI solution.

64.5.5 Optimistic initialization

As in *Inventory Management via Q-Learning*, we use optimistic initialization to accelerate learning.

The logic is the same – initialize the Q-table so that every untried action looks attractive, driving the agent to explore broadly – but the direction is reversed.

Since the optimal policy *minimizes* q, "optimistic" means initializing the Q-table *below* the true values. When the agent tries an action, the update pushes q upward toward reality, making that entry look worse and prompting the agent to try other actions that still appear optimistically good.

The true Q-values are on the order of $\exp(-\gamma\, v^*) \approx 10^{-8}$ to 10^{-6}. We initialize the Q-table at 10^{-9}, modestly below this range.

64.5.6 Implementation

We first define a helper to extract the greedy policy from the Q-table.

Since the optimal policy minimizes q, we use arg min rather than arg max.

```
@numba.jit(nopython=True)
def greedy_policy_from_q_rs(q, K):
    """Extract greedy policy from risk-sensitive Q-table (argmin)."""
    σ = np.empty(K + 1, dtype=np.int32)
    for x in range(K + 1):
        best_val = np.inf
        best_a = 0
        for a in range(K - x + 1):
            if q[x, a] < best_val:
                best_val = q[x, a]
                best_a = a
        σ[x] = best_a
    return σ
```

The Q-learning loop mirrors the risk-neutral version, with the key changes: the update target uses $\exp(-\gamma R_{t+1}) \cdot (\min_{a'} q_t)^\beta$, and the behavior policy follows the arg min.

```
@numba.jit(nopython=True)
def q_learning_rs_kernel(K, p, c, κ, β, γ, n_steps, X_init,
                         ε_init, ε_min, ε_decay, q_init, snapshot_steps, rng):
    q = np.full((K + 1, K + 1), q_init)  # optimistic initialization
    n = np.zeros((K + 1, K + 1))         # visit counts for learning rate
    ε = ε_init

    n_snaps = len(snapshot_steps)
    snapshots = np.zeros((n_snaps, K + 1), dtype=np.int32)
    snap_idx = 0

    # Initialize state and action
    x = X_init
    a = rng.integers(0, K - x + 1)

    for t in range(n_steps):
        # Record policy snapshot if needed
        if snap_idx < n_snaps and t == snapshot_steps[snap_idx]:
            snapshots[snap_idx] = greedy_policy_from_q_rs(q, K)
            snap_idx += 1

        # === Draw D_{t+1} and observe outcome ===
        d = rng.geometric(p) - 1
        reward = min(x, d) - c * a - κ * (a > 0)
        x_next = max(x - d, 0) + a

        # === Min over next state (scalar value for update target) ===
        # Also record the argmin action for use by the behavior policy.
        best_next = np.inf
        a_next = 0
        for aa in range(K - x_next + 1):
```

(continues on next page)

(continued from previous page)

```
            if q[x_next, aa] < best_next:
                best_next = q[x_next, aa]
                a_next = aa

        # === Risk-sensitive Q-learning update ===
        target = np.exp(-γ * reward) * best_next ** β
        n[x, a] += 1
        α = 1.0 / n[x, a] ** 0.51
        q[x, a] = (1 - α) * q[x, a] + α * target

        # === Behavior policy: ε-greedy (uses a_next, the argmin action) ===
        x = x_next
        if rng.random() < ε:
            a = rng.integers(0, K - x + 1)
        else:
            a = a_next
        ε = max(ε_min, ε * ε_decay)

    return q, snapshots
```

The wrapper function unpacks the model and provides default hyperparameters.

```
def q_learning_rs(model, n_steps=20_000_000, X_init=0,
                  ε_init=1.0, ε_min=0.01, ε_decay=0.999999,
                  q_init=1e-9, snapshot_steps=None, seed=1234):
    x_values, d_values, ϕ_values, p, c, κ, β, γ = model
    K = len(x_values) - 1
    if snapshot_steps is None:
        snapshot_steps = np.array([], dtype=np.int64)
    rng = np.random.default_rng(seed)
    return q_learning_rs_kernel(K, p, c, κ, β, γ, n_steps, X_init,
                                ε_init, ε_min, ε_decay, q_init, snapshot_steps, rng)
```

64.5.7 Running Q-learning

We run n = 5 million steps and take policy snapshots at steps 10,000, 1,000,000, and n.

```
n = 5_000_000
snap_steps = np.array([10_000, 1_000_000, n], dtype=np.int64)
q_table, snapshots = q_learning_rs(model, n_steps=n+1, snapshot_steps=snap_steps)
```

64.5.8 Comparing with the exact solution

We extract the value function and policy from the final Q-table.

Since Q-values represent $\mathbb{E}[\exp(-\gamma(\cdots))]$, we recover the value function via $v_Q(x) = -\frac{1}{\gamma}\ln(\min_a q(x,a))$ and the policy via $\sigma_Q(x) = \arg\min_a q(x,a)$.

```
K = len(x_values) - 1
γ_base = model.γ
# restrict to feasible actions a ∈ {0, ..., K-x}
v_q = np.array([-(1/γ_base) * np.log(np.min(q_table[x, :K - x + 1]))
                for x in range(K + 1)])
```

(continues on next page)

(continued from previous page)

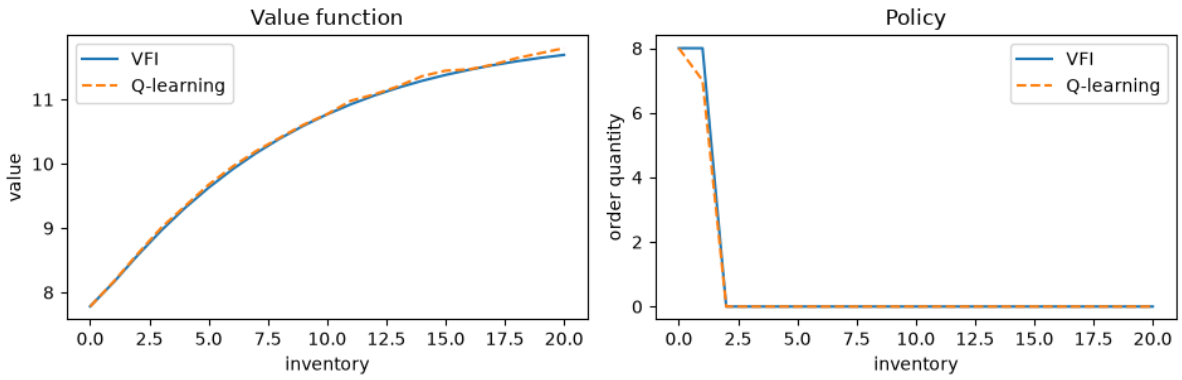
```
σ_q = np.array([np.argmin(q_table[x, :K - x + 1])
                for x in range(K + 1)])
```

```
fig, axes = plt.subplots(1, 2, figsize=(9.6, 3.2))

axes[0].plot(x_values, v_star, label="VFI")
axes[0].plot(x_values, v_q, '--', label="Q-learning")
axes[0].set_xlabel("inventory")
axes[0].set_ylabel("value")
axes[0].legend()
axes[0].set_title("Value function")

axes[1].plot(x_values, σ_star, label="VFI")
axes[1].plot(x_values, σ_q, '--', label="Q-learning")
axes[1].set_xlabel("inventory")
axes[1].set_ylabel("order quantity")
axes[1].legend()
axes[1].set_title("Policy")

plt.tight_layout()
plt.show()
```



64.5.9 Visualizing learning over time

The panels below show how the agent's policy evolves during training.

Each panel simulates an inventory path using the greedy policy extracted from the Q-table at a given training step, with the same demand sequence throughout.

The top panel shows the optimal policy from VFI for reference.

```
ts_length = 200
n_snaps = len(snap_steps)
fig, axes = plt.subplots(n_snaps + 1, 1, figsize=(8, 2.0 * (n_snaps + 1)),
                         sharex=True)

X_init = K // 2
sim_seed = 5678

# Optimal policy
X_opt = sim_inventories(ts_length, σ_star, model.p,
```

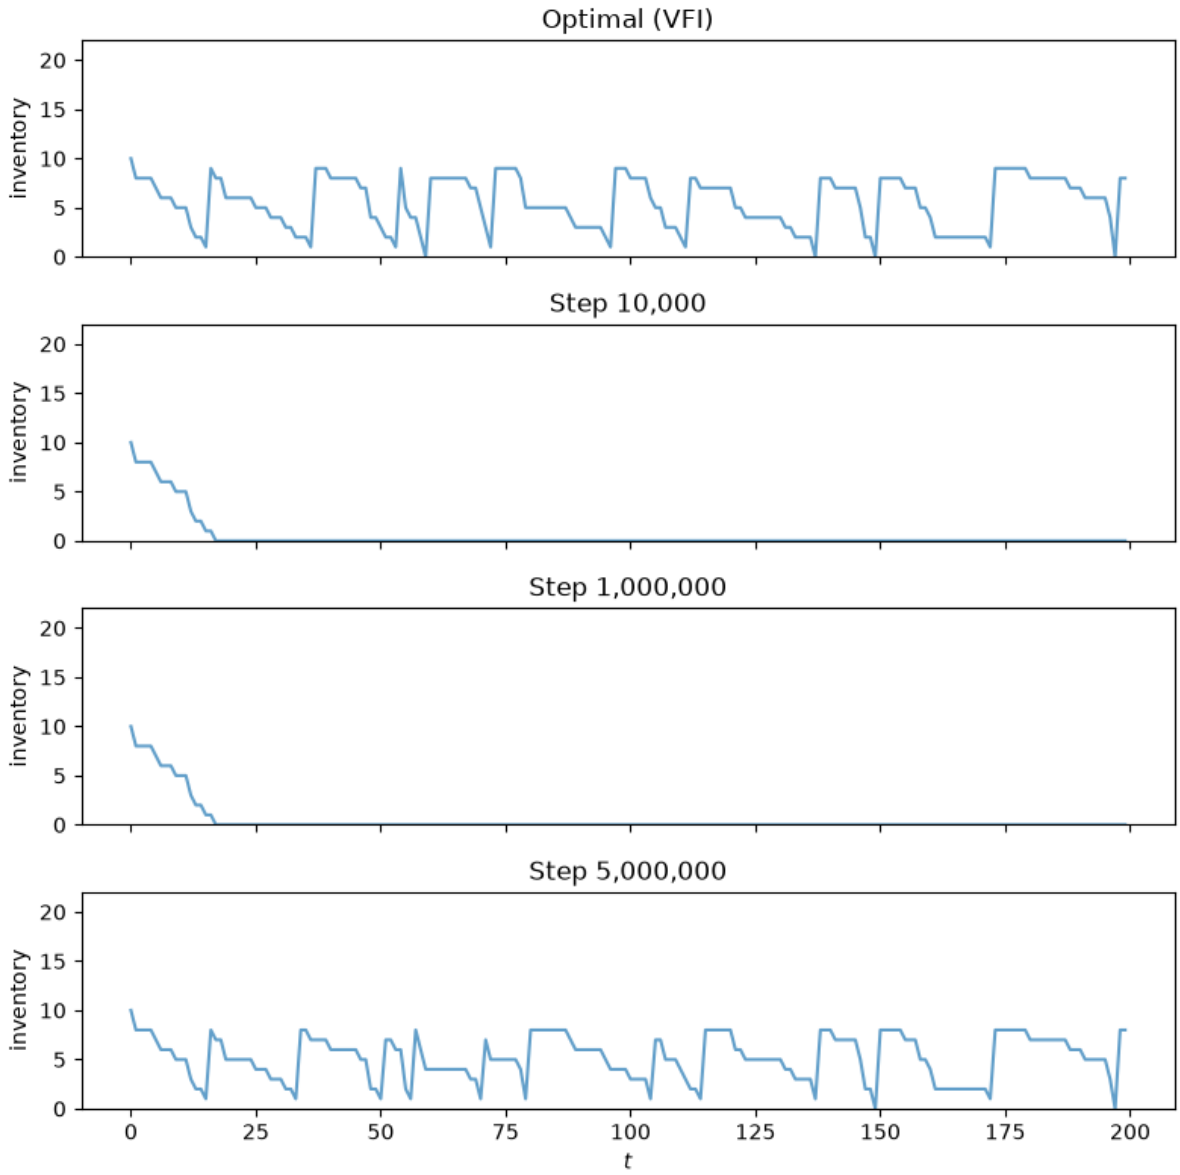
(continues on next page)

(continued from previous page)

```
                        np.random.default_rng(sim_seed), X_init)
axes[0].plot(X_opt, alpha=0.7)
axes[0].set_ylabel("inventory")
axes[0].set_title("Optimal (VFI)")
axes[0].set_ylim(0, K + 2)

# Q-learning snapshots
for i in range(n_snaps):
    σ_snap = snapshots[i]
    X = sim_inventories(ts_length, σ_snap, model.p,
                        np.random.default_rng(sim_seed), X_init)
    axes[i + 1].plot(X, alpha=0.7)
    axes[i + 1].set_ylabel("inventory")
    axes[i + 1].set_title(f"Step {snap_steps[i]:,}")
    axes[i + 1].set_ylim(0, K + 2)

axes[-1].set_xlabel(r"$t$")
plt.tight_layout()
plt.show()
```



After 10,000 steps, the agent has barely explored and its policy is erratic.

By 1,000,000 steps the learned policy has improved but still differs noticeably from the optimum.

By step 5 million the inventory dynamics are nearly indistinguishable from the VFI solution.

Note that the converged policy maintains lower inventory levels than in the risk-neutral case (compare with *Inventory Management via Q-Learning*), consistent with the mechanism discussed above: a risk-sensitive agent caps its exposure to demand variance by holding less stock.

64.6 Conclusion

We extended the inventory management problem from *Inventory Management via Q-Learning* to incorporate risk sensitivity via the certainty equivalent operator $\psi^{-1}(\mathbb{E}[\psi(\cdot)])$ with $\psi(t) = \exp(-\gamma t)$.

Value function iteration confirms that risk-sensitive firms order less aggressively, preferring predictable profits over higher but more volatile returns.

We then showed that Q-learning can be adapted to the risk-sensitive setting by working with the transformed Q-factor $q(x, a) = \mathbb{E}[\exp(-\gamma(\pi + \beta v^*))]$.

The resulting update rule replaces addition with multiplication and max with min, but retains the key property of model-free learning: the agent needs only to observe states, actions, and profits.

CHAPTER

SIXTYFIVE

JOB SEARCH X: SEARCH WITH Q-LEARNING

65.1 Overview

This lecture illustrates a powerful machine learning technique called Q-learning.

[Sutton and Barto, 2018] presents Q-learning and a variety of other statistical learning procedures.

The Q-learning algorithm combines ideas from

- dynamic programming
- a recursive version of least squares known as temporal difference learning.

This lecture applies a Q-learning algorithm to the situation faced by a McCall worker.

This lecture also considers the case where a McCall worker is given an option to quit the current job.

Relative to the dynamic programming formulation of the McCall worker model that we studied in *quantecon lecture*, a Q-learning algorithm gives the worker less knowledge about

- the random process that generates a sequence of wages
- the reward function that tells consequences of accepting or rejecting a job

The Q-learning algorithm invokes a statistical learning model to learn about these things.

Statistical learning often comes down to some version of least squares, and it will be here too.

Any time we say **statistical learning**, we have to say what object is being learned.

For Q-learning, the object that is learned is not the **value function** that is a focus of dynamic programming.

But it is something that is closely affiliated with it.

In the finite-action, finite state context studied in this lecture, the object to be learned statistically is a **Q-table**, an instance of a **Q-function** for finite sets.

Sometimes a Q-function or Q-table is called a quality-function or quality-table.

The rows and columns of a Q-table correspond to possible states that an agent might encounter, and possible actions that he can take in each state.

An equation that resembles a Bellman equation plays an important role in the algorithm.

It differs from the Bellman equation for the McCall model that we have seen in *this quantecon lecture*

In this lecture, we'll learn a little about

- the **Q-function** or **quality function** that is affiliated with any Markov decision problem whose optimal value function satisfies a Bellman equation
- **temporal difference learning**, a key component of a Q-learning algorithm

As usual, let's import some Python modules.

```
!pip install quantecon
```

```
import numpy as np

from numba import jit, float64, int64
from numba.experimental import jitclass
from quantecon.distributions import BetaBinomial

import matplotlib.pyplot as plt

rng = np.random.default_rng(123)
```

65.2 Review of McCall model

We begin by reviewing the McCall model described in *this quantecon lecture*.

We'll compute an optimal value function and a policy that attains it.

We'll eventually compare that optimal policy to what the Q-learning McCall worker learns.

The McCall model is characterized by parameters β, c and a known distribution of wage offers F.

A McCall worker wants to maximize an expected discounted sum of lifetime incomes

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t y_t$$

The worker's income y_t equals his wage w if he is employed, and unemployment compensation c if he is unemployed.

An optimal value $V(w)$ for a McCall worker who has just received a wage offer w and is deciding whether to accept or reject it satisfies the Bellman equation

$$V(w) = \max_{\text{accept, reject}} \left\{ \frac{w}{1-\beta}, c + \beta \int V(w') \, dF(w') \right\} \tag{65.1}$$

To form a benchmark to compare with results from Q-learning, we first approximate the optimal value function.

With possible states residing in a finite discrete state space indexed by $\{1, 2, ..., n\}$, we make an initial guess for the value function of $v \in \mathbb{R}^n$ and then iterate on the Bellman equation:

$$v'(i) = \max \left\{ \frac{w(i)}{1-\beta}, c + \beta \sum_{1 \le j \le n} v(j) q(j) \right\} \quad \text{for } i = 1, \ldots, n$$

Let's use Python code from *this quantecon lecture*.

We use a Python method called `VFI` to compute the optimal value function using value function iterations.

We construct an assumed distribution of wages and plot it with the following Python code

```
n, a, b = 10, 200, 100                       # default parameters
q_default = BetaBinomial(n, a, b).pdf()      # default choice of q

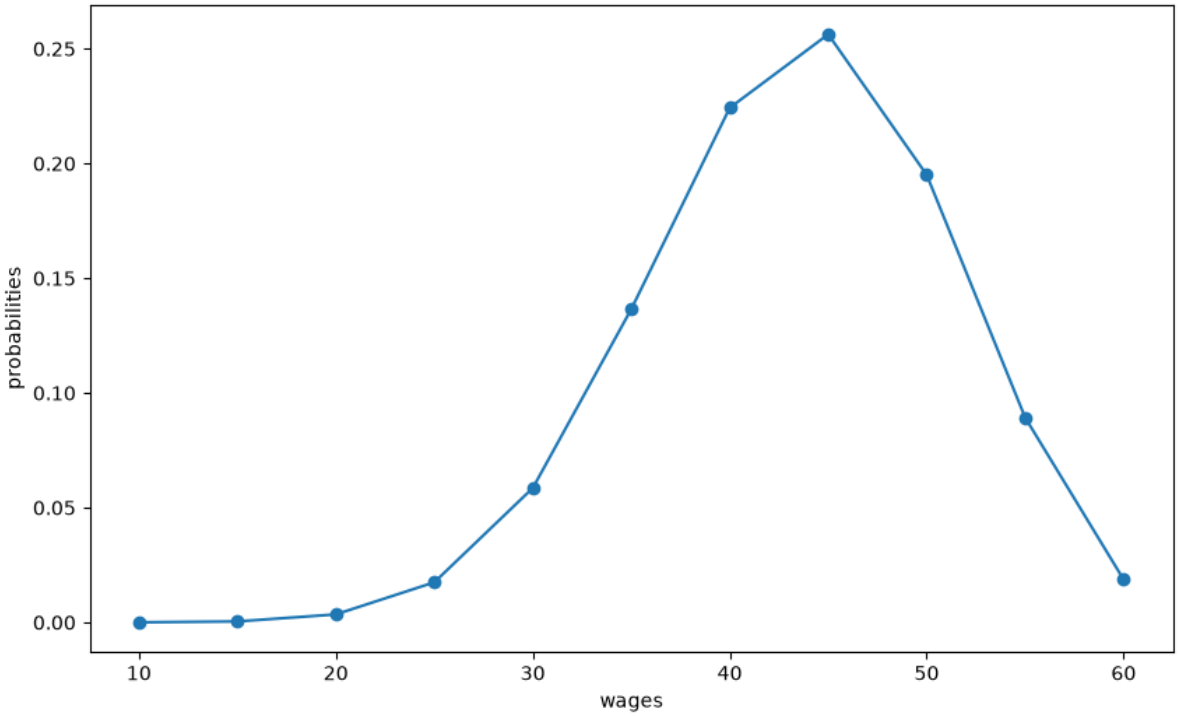
w_min, w_max = 10, 60
w_default = np.linspace(w_min, w_max, n+1)
```

(continues on next page)

(continued from previous page)

```
# plot distribution of wage offer
fig, ax = plt.subplots(figsize=(10,6))
ax.plot(w_default, q_default, '-o', label='$q(w(i))$')
ax.set_xlabel('wages')
ax.set_ylabel('probabilities')

plt.show()
```



Next we'll compute the worker's optimal value function by iterating to convergence on the Bellman equation.

Then we'll plot various iterates on the Bellman operator.

```
mccall_data = [
    ('c', float64),      # unemployment compensation
    ('β', float64),      # discount factor
    ('w', float64[::1]),  # array of wage values, w[i] = wage at state i
    ('q', float64[::1])   # array of probabilities
]


@jitclass(mccall_data)
class McCallModel:

    def __init__(self, c=25, β=0.99, w=w_default, q=q_default):

        self.c, self.β = c, β
        self.w, self.q = w, q

    def state_action_values(self, i, v):
        """
        The values of state-action pairs.
```

(continues on next page)

(continued from previous page)

```
        """
        # Simplify names
        c, β, w, q = self.c, self.β, self.w, self.q
        # Evaluate value for each state-action pair
        # Consider action = accept or reject the current offer
        accept = w[i] / (1 - β)
        reject = c + β * (v @ q)

        return np.array([accept, reject])

    def VFI(self, eps=1e-5, max_iter=500):
        """
        Find the optimal value function.
        """

        n = len(self.w)
        v = self.w / (1 - self.β)
        v_next = np.empty_like(v)
        flag=0

        for i in range(max_iter):
            for j in range(n):
                v_next[j] = np.max(self.state_action_values(j, v))

            if np.max(np.abs(v_next - v))<=eps:
                flag=1
                break
            v[:] = v_next

        return v, flag

def plot_value_function_seq(mcm, ax, num_plots=8):
    """
    Plot a sequence of value functions.

        * mcm is an instance of McCallModel
        * ax is an axes object that implements a plot method.

    """

    n = len(mcm.w)
    v = mcm.w / (1 - mcm.β)
    v_next = np.empty_like(v)
    for i in range(num_plots):
        ax.plot(mcm.w, v, '-', alpha=0.4, label=f"iterate {i}")
        # Update guess
        for i in range(n):
            v_next[i] = np.max(mcm.state_action_values(i, v))
        v[:] = v_next  # copy contents into v

    ax.legend(loc='lower right')
```

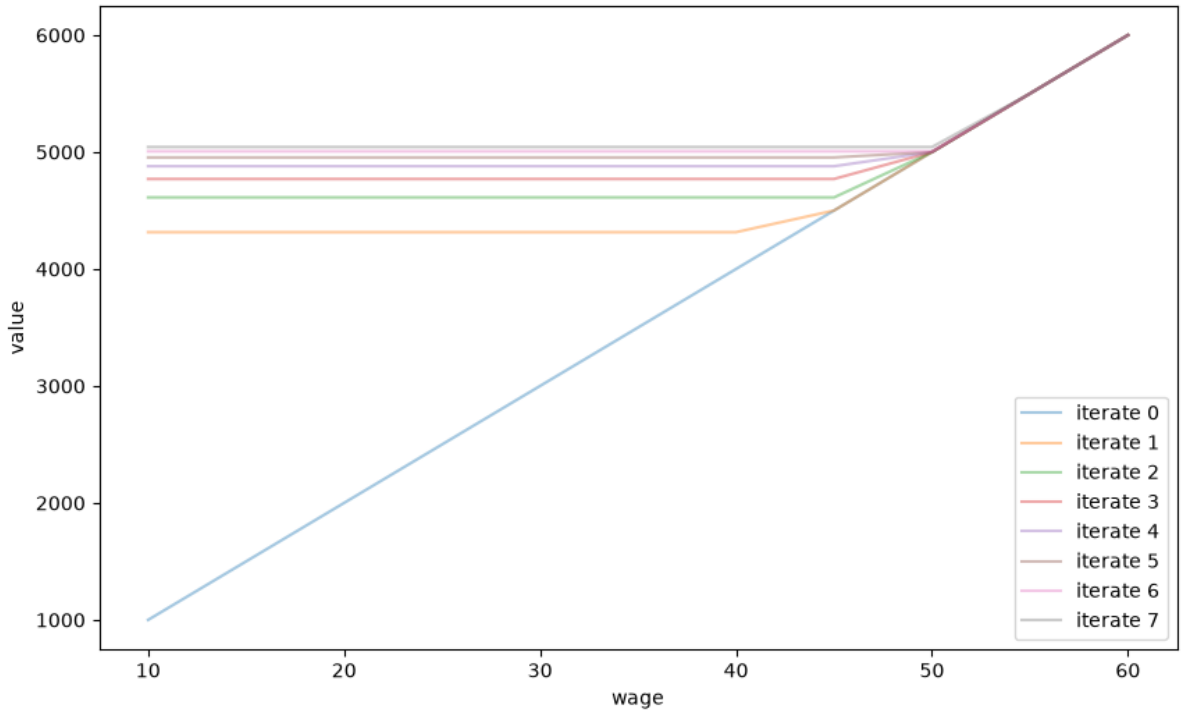
```
mcm = McCallModel()
valfunc_VFI, flag = mcm.VFI()

fig, ax = plt.subplots(figsize=(10,6))
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel('wage')
ax.set_ylabel('value')
plot_value_function_seq(mcm, ax)
plt.show()
```



Next we'll print out the limit of the sequence of iterates.

This is the approximation to the McCall worker's value function that is produced by value function iteration.

We'll use this value function as a benchmark later after we have done some Q-learning.

```
print(valfunc_VFI)
```

```
[5322.27935875 5322.27935875 5322.27935875 5322.27935875 5322.27935875
 5322.27935875 5322.27935875 5322.27935875 5322.27935875 5500.
 6000.        ]
```

65.3 Implied quality function Q

A **quality function** Q maps state-action pairs into optimal values.

They are tightly linked to optimal value functions.

But value functions are functions just of states, and not actions.

For each given state, the quality function gives a list of optimal values that can be attained starting from that state, with each component of the list indicating one of the possible actions that is taken.

For our McCall worker with a finite set of possible wages

- the state space $\mathcal{W} = \{w_1, w_2, ..., w_n\}$ is indexed by integers $1, 2, ..., n$

- the action space is $\mathcal{A} = \{\text{accept}, \text{reject}\}$

Let $a \in \mathcal{A}$ be one of the two possible actions, i.e., accept or reject.

For our McCall worker, an optimal Q-function $Q(w, a)$ equals the maximum value of that a previously unemployed worker who has offer w in hand can attain if he takes action a.

This definition of $Q(w, a)$ presumes that in subsequent periods the worker takes optimal actions.

An optimal Q-function for our McCall worker satisfies

$$\begin{aligned} Q\left(w, \text{accept}\right) &= \frac{w}{1-\beta} \\ Q\left(w, \text{reject}\right) &= c + \beta \int \max_{\text{accept, reject}} \left\{ \frac{w'}{1-\beta}, Q\left(w', \text{reject}\right) \right\} dF\left(w'\right) \end{aligned} \tag{65.2}$$

Note that the first equation of system (65.2) presumes that after the agent has accepted an offer, he will not have the option to reject that same offer in the future.

These equations are aligned with the Bellman equation for the worker's optimal value function that we studied in *this quantecon lecture*.

Evidently, the optimal value function $V(w)$ described in that lecture is related to our Q-function by

$$V(w) = \max_{\text{accept,reject}} \left\{ Q(w, \text{accept}), Q\left(w, \text{reject}\right) \right\}$$

If we stare at the second equation of system (65.2), we notice that since the wage process is identically and independently distributed over time, $Q\left(w, \text{reject}\right)$, the right side of the equation is independent of the current state w.

So we can denote it as a scalar

$$Q_r := Q\left(w, \text{reject}\right) \quad \forall\, w \in \mathcal{W}.$$

This fact provides us with an an alternative, and as it turns out in this case, a faster way to compute an optimal value function and associated optimal policy for the McCall worker.

Instead of using the value function iterations that we deployed above, we can instead iterate to convergence on a version of the second equation in system (65.2) that maps an estimate of Q_r into an improved estimate Q_r':

$$Q_r' = c + \beta \int \max \left\{ \frac{w'}{1-\beta}, Q_r \right\} dF\left(w'\right)$$

After a Q_r sequence has converged, we can recover the optimal value function $V(w)$ for the McCall worker from

$$V\left(w\right) = \max \left\{ \frac{w}{1-\beta}, Q_r \right\}$$

65.4 From probabilities to samples

We noted above that the optimal Q function for our McCall worker satisfies the Bellman equations

$$\begin{aligned} w + \beta \max_{\text{accept, reject}} \left\{ Q(w, \text{accept}), Q(w, \text{reject}) \right\} - Q(w, \text{accept}) &= 0 \\ c + \beta \int \max_{\text{accept, reject}} \left\{ Q(w', \text{accept}), Q\left(w', \text{reject}\right) \right\} dF\left(w'\right) - Q\left(w, \text{reject}\right) &= 0 \end{aligned} \tag{65.3}$$

Notice the integral over $F(w')$ on the second line.

Erasing the integral sign sets the stage for an illegitimate argument that can get us started thinking about Q-learning.

Thus, construct a difference equation system that keeps the first equation of (65.3) but replaces the second by removing integration over $F(w')$:

$$\begin{aligned} w + \beta \max_{\text{accept, reject}} \{Q(w, \text{accept}), Q(w, \text{reject})\} - Q(w, \text{accept}) = 0 \\ c + \beta \max_{\text{accept, reject}} \{Q(w', \text{accept}), Q\left(w', \text{reject}\right)\} - Q\left(w, \text{reject}\right) \approx 0 \end{aligned} \tag{65.4}$$

The second equation can't hold for all w, w' pairs in the appropriate Cartesian product of our state space.

But maybe an appeal to a Law of Large numbers could let us hope that it would hold **on average** for a long time series sequence of draws of w_t, w_{t+1} pairs, where we are thinking of w_t as w and w_{t+1} as w'.

The basic idea of Q-learning is to draw a long sample of wage offers from F (we know F though we assume that the worker doesn't) and iterate on a recursion that maps an estimate $\hat{Q}_t$ of a Q-function at date t into an improved estimate $\hat{Q}_{t+1}$ at date $t+1$.

To set up such an algorithm, we first define some errors or "differences"

$$\begin{aligned} w + \beta \max_{\text{accept, reject}} \left\{\hat{Q}_t(w_t, \text{accept}), \hat{Q}_t(w_t, \text{reject})\right\} - \hat{Q}_t(w_t, \text{accept}) = \text{diff}_{\text{accept},t} \\ c + \beta \max_{\text{accept, reject}} \left\{\hat{Q}_t(w_{t+1}, \text{accept}), \hat{Q}_t\left(w_{t+1}, \text{reject}\right)\right\} - \hat{Q}_t\left(w_t, \text{reject}\right) = \text{diff}_{\text{reject},t} \end{aligned} \tag{65.5}$$

The adaptive learning scheme would then be some version of

$$\hat{Q}_{t+1} = \hat{Q}_t + \alpha \, \text{diff}_t \tag{65.6}$$

where $\alpha \in (0, 1)$ is a small **gain** parameter that governs the rate of learning and $\hat{Q}_t$ and diff_t are 2×1 vectors corresponding to objects in equation system (65.5).

This informal argument takes us to the threshold of Q-learning.

65.5 Q-learning

Let's first describe a Q-learning algorithm precisely.

Then we'll implement it.

The algorithm works by using a Monte Carlo method to update estimates of a Q-function.

We begin with an initial guess for a Q-function.

In the example studied in this lecture, we have a finite action space and also a finite state space.

That means that we can represent a Q-function as a matrix or Q-table, $\widetilde{Q}(w, a)$.

Q-learning proceeds by updating the Q-function as the decision maker acquires experience along a path of wage draws generated by simulation.

During the learning process, our McCall worker takes actions and experiences rewards that are consequences of those actions.

He learns simultaneously about the environment, in this case the distribution of wages, and the reward function, in this case the unemployment compensation c and the present value of wages.

The updating algorithm is based on a slight modification (to be described soon) of a recursion like

$$\widetilde{Q}^{new}\left(w, a\right) = \widetilde{Q}^{old}\left(w, a\right) + \alpha \widetilde{TD}\left(w, a\right) \tag{65.7}$$

where

$$\begin{aligned}\widetilde{TD}\left(w,\text{accept}\right) &= \left[w+\beta\max_{a'\in\mathcal{A}}\widetilde{Q}^{old}\left(w,a'\right)\right]-\widetilde{Q}^{old}\left(w,\text{accept}\right)\\ \widetilde{TD}\left(w,\text{reject}\right) &= \left[c+\beta\max_{a'\in\mathcal{A}}\widetilde{Q}^{old}\left(w',a'\right)\right]-\widetilde{Q}^{old}\left(w,\text{reject}\right),\ w'\sim F\end{aligned} \tag{65.8}$$

The terms $\widetilde{TD}(w,a)$ for $a=\{\text{accept,reject}\}$ are the **temporal difference errors** that drive the updates.

This system is thus a version of the adaptive system that we sketched informally in equation (65.6).

An aspect of the algorithm not yet captured by equation system (65.8) is random **experimentation** that we add by occasionally randomly replacing

$$\text{argmax}_{a'\in\mathcal{A}}\widetilde{Q}^{old}\left(w,a'\right)$$

with

$$\text{argmin}_{a'\in\mathcal{A}}\widetilde{Q}^{old}\left(w,a'\right)$$

and occasionally replacing

$$\text{argmax}_{a'\in\mathcal{A}}\widetilde{Q}^{old}\left(w',a'\right)$$

with

$$\text{argmin}_{a'\in\mathcal{A}}\widetilde{Q}^{old}\left(w',a'\right)$$

We activate such experimentation with probability ϵ in step 3 of the following pseudo-code for our McCall worker to do Q-learning:

1. Set an arbitrary initial Q-table.
2. Draw an initial wage offer w from F.
3. From the appropriate row in the Q-table, choose an action using the following ϵ-greedy algorithm:
    - with probability $1-\epsilon$, choose the action that maximizes the value, and
    - with probability ϵ, choose the alternative action.
4. Update the state associated with the chosen action and compute $\widetilde{TD}$ according to (65.8) and update $\widetilde{Q}$ according to (65.7).
5. Either draw a new state w' if required or else take the existing wage and update the Q-table again according to (65.7).
6. Stop when the old and new Q-tables are close enough, i.e., $\|\tilde{Q}^{new}-\tilde{Q}^{old}\|_\infty\leq\delta$ for given δ or if the worker keeps accepting for T periods for a prescribed T.
7. Return to step 2 with the updated Q-table.

Repeat this procedure for N episodes or until the updated Q-table has converged.

We call one pass through steps 2 to 7 an "episode" or "epoch" of temporal difference learning.

In our context, each episode starts with an agent drawing an initial wage offer, i.e., a new state.

The agent then takes actions based on the preset Q-table, receives rewards, and then enters a new state implied by this period's actions.

The Q-table is updated via temporal difference learning.

We iterate this until convergence of the Q-table or the maximum length of an episode is reached.

Multiple episodes allow the agent to start afresh and visit states that she was less likely to visit from the terminal state of a previous episode.

For example, an agent who has accepted a wage offer based on her Q-table will be less likely to draw a new offer from other parts of the wage distribution.

By using the ϵ-greedy method and also by increasing the number of episodes, the Q-learning algorithm balances gains from exploration and from exploitation.

Remark: Notice that $\widetilde{TD}$ associated with an optimal Q-table defined in (65.7) automatically above satisfies $\widetilde{TD} = 0$ for all state action pairs. Whether a limit of our Q-learning algorithm converges to an optimal Q-table depends on whether the algorithm visits all state-action pairs often enough.

We implement this pseudo code in a Python class.

For simplicity and convenience, we let `s` represent the state index between 0 and $n = 50$ and $w_s = w[s]$.

The first column of the Q-table represents the value associated with rejecting the wage and the second represents accepting the wage.

We use `numba` compilation to accelerate computations.

```
params=[
    ('c', float64),            # unemployment compensation
    ('β', float64),            # discount factor
    ('w', float64[:]),         # array of wage values, w[i] = wage at state i
    ('q', float64[:]),         # array of probabilities
    ('eps', float64),          # for epsilon greedy algorithm
    ('δ', float64),            # Q-table threshold
    ('lr', float64),           # the learning rate α
    ('T', int64),              # maximum periods of accepting
    ('quit_allowed', int64)    # whether quit is allowed after accepting the wage␣
 ↪offer
]

@jitclass(params)
class Qlearning_McCall:
    def __init__(self, c=25, β=0.99, w=w_default, q=q_default, eps=0.1,
                 δ=1e-5, lr=0.5, T=10000, quit_allowed=0):

        self.c, self.β = c, β
        self.w, self.q = w, q
        self.eps, self.δ, self.lr, self.T = eps, δ, lr, T
        self.quit_allowed = quit_allowed


    def draw_offer_index(self, rng):
        """
        Draw a state index from the wage distribution.
        """

        q = self.q
        return np.searchsorted(np.cumsum(q), rng.random(), side="right")

    def temp_diff(self, qtable, state, accept, rng):
        """
        Compute the TD associated with state and action.
        """

        c, β, w = self.c, self.β, self.w
```

(continues on next page)

(continued from previous page)

```
        if accept==0:
            state_next = self.draw_offer_index(rng)
            TD = c + β*np.max(qtable[state_next, :]) - qtable[state, accept]
        else:
            state_next = state
            if self.quit_allowed == 0:
                TD = w[state_next] + β*np.max(qtable[state_next, :]) - qtable[state,␣
↪accept]
            else:
                TD = w[state_next] + β*qtable[state_next, 1] - qtable[state, accept]

        return TD, state_next

    def run_one_epoch(self, qtable, rng, max_times=20000):
        """
        Run an "epoch".
        """

        c, β, w = self.c, self.β, self.w
        eps, δ, lr, T = self.eps, self.δ, self.lr, self.T

        s0 = self.draw_offer_index(rng)
        s = s0
        accept_count = 0

        for t in range(max_times):

            # choose action
            accept = np.argmax(qtable[s, :])
            if rng.random()<=eps:
                accept = 1 - accept

            if accept == 1:
                accept_count += 1
            else:
                accept_count = 0

            TD, s_next = self.temp_diff(qtable, s, accept, rng)

            # update qtable
            qtable_new = qtable.copy()
            qtable_new[s, accept] = qtable[s, accept] + lr*TD

            if np.max(np.abs(qtable_new-qtable))<=δ:
                break

            if accept_count == T:
                break

            s, qtable = s_next, qtable_new

        return qtable_new

@jit
def run_epochs(N, qlmc, qtable, rng):
    """
```

(continues on next page)

(continued from previous page)

```
    Run epochs N times with qtable from the last iteration each time.
    """

    for n in range(N):
        if n % max(1, N // 10) == 0:
            print(f"Progress: EPOCHs = {n}")
        new_qtable = qlmc.run_one_epoch(qtable, rng)
        qtable = new_qtable

    return qtable

def valfunc_from_qtable(qtable):
    return np.max(qtable, axis=1)

def compute_error(valfunc, valfunc_VFI):
    return np.mean(np.abs(valfunc-valfunc_VFI))
```

```
# create an instance of Qlearning_McCall
qlmc = Qlearning_McCall()

# run
qtable0 = np.zeros((len(w_default), 2))
qtable = run_epochs(20000, qlmc, qtable0, rng)
```

```
Progress: EPOCHs = 0
Progress: EPOCHs = 2000
Progress: EPOCHs = 4000
Progress: EPOCHs = 6000
Progress: EPOCHs = 8000
Progress: EPOCHs = 10000
Progress: EPOCHs = 12000
Progress: EPOCHs = 14000
Progress: EPOCHs = 16000
Progress: EPOCHs = 18000
```

```
print(qtable)
```

```
[[2551.23116562    0.        ]
 [5310.29611659 5231.41648921]
 [5461.27287644 5307.2363605 ]
 [5310.97399124 5263.64609929]
 [5463.77676012 5262.04241975]
 [5396.48394892 5300.95680361]
 [5390.7938418  5308.48224118]
 [5293.52306414 5276.79778395]
 [5393.05926154 5276.68135256]
 [5317.30058452 5500.00001218]
 [5379.99537477 6000.        ]]
```

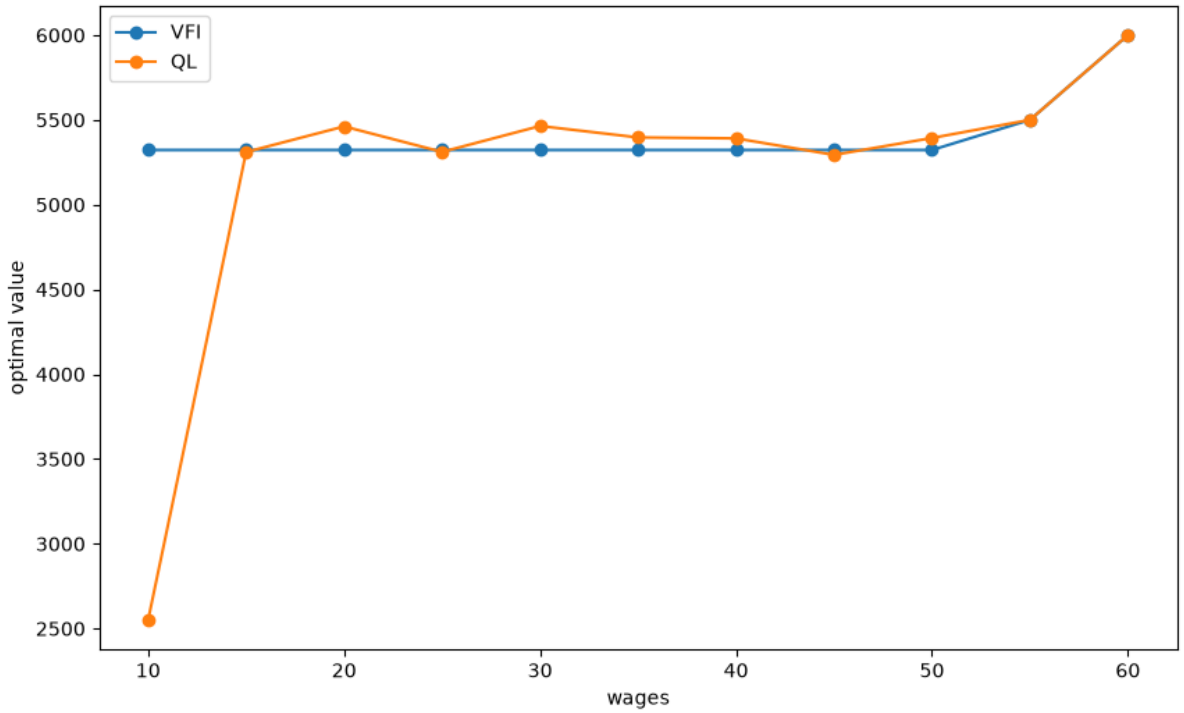
```
# inspect value function
valfunc_qlr = valfunc_from_qtable(qtable)

print(valfunc_qlr)
```

```
[2551.23116562 5310.29611659 5461.27287644 5310.97399124 5463.77676012
 5396.48394892 5390.7938418  5293.52306414 5393.05926154 5500.00001218
 6000.        ]
```

```
# plot
fig, ax = plt.subplots(figsize=(10,6))
ax.plot(w_default, valfunc_VFI, '-o', label='VFI')
ax.plot(w_default, valfunc_qlr, '-o', label='QL')
ax.set_xlabel('wages')
ax.set_ylabel('optimal value')
ax.legend()

plt.show()
```



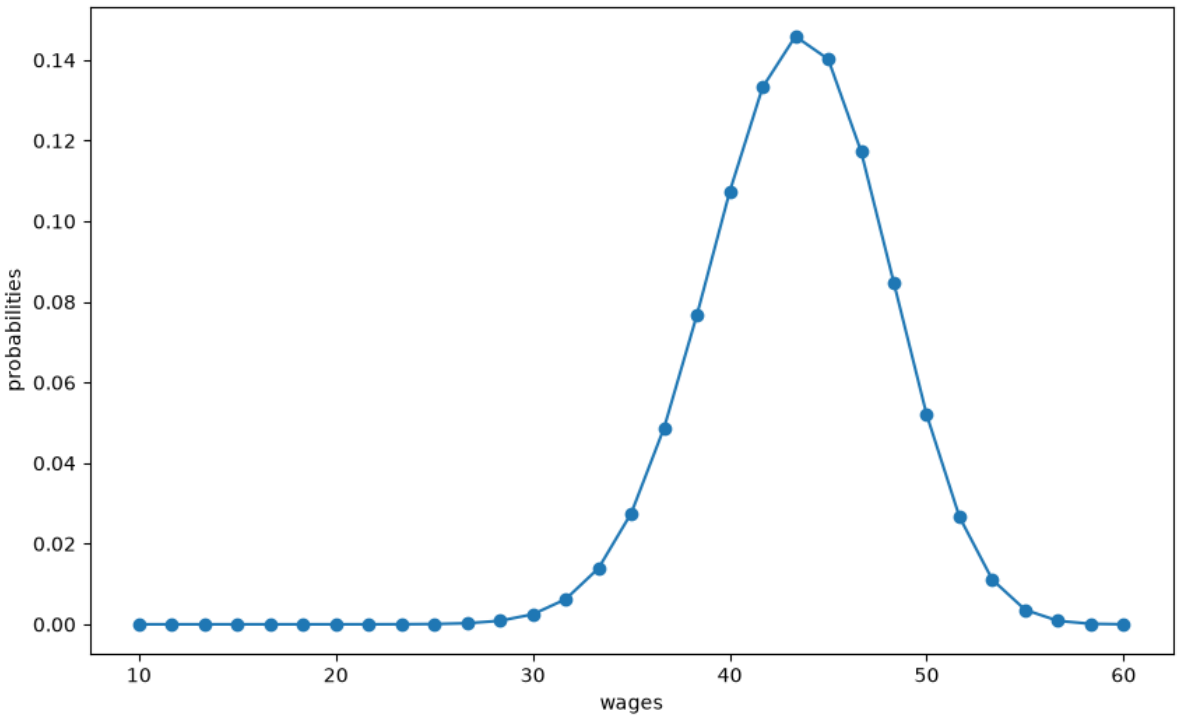
Now, let us compute the case with a larger state space: $n = 30$ instead of $n = 10$.

```
n, a, b = 30, 200, 100                    # default parameters
q_new = BetaBinomial(n, a, b).pdf()       # default choice of q

w_min, w_max = 10, 60
w_new = np.linspace(w_min, w_max, n+1)


# plot distribution of wage offer
fig, ax = plt.subplots(figsize=(10,6))
ax.plot(w_new, q_new, '-o', label='$q(w(i))$')
ax.set_xlabel('wages')
ax.set_ylabel('probabilities')

plt.show()
```




```
mcm = McCallModel(w=w_new, q=q_new)
valfunc_VFI, flag = mcm.VFI()
valfunc_VFI
```

```
array([4859.77015703, 4859.77015703, 4859.77015703, 4859.77015703,
       4859.77015703, 4859.77015703, 4859.77015703, 4859.77015703,
       4859.77015703, 4859.77015703, 4859.77015703, 4859.77015703,
       4859.77015703, 4859.77015703, 4859.77015703, 4859.77015703,
       4859.77015703, 4859.77015703, 4859.77015703, 4859.77015703,
       5000.        , 5166.66666667, 5333.33333333, 5500.        ,
       5666.66666667, 5833.33333333, 6000.        ])
```

```
def plot_epochs(epochs_to_plot, quit_allowed=1):
    "Plot value function implied by outcomes of an increasing number of epochs."
    qlmc_new = Qlearning_McCall(w=w_new, q=q_new, quit_allowed=quit_allowed)
    qtable = np.zeros((len(w_new),2))
    epochs_to_plot = np.asarray(epochs_to_plot)
    # plot
    fig, ax = plt.subplots(figsize=(10,6))
    ax.plot(w_new, valfunc_VFI, '-o', label='VFI')

    max_epochs = np.max(epochs_to_plot)
    # iterate on epoch numbers
    for n in range(max_epochs + 1):
        if n % max(1, max_epochs // 10) == 0:
            print(f"Progress: EPOCHs = {n}")
        if n in epochs_to_plot:
            valfunc_qlr = valfunc_from_qtable(qtable)
            error = compute_error(valfunc_qlr, valfunc_VFI)
```

(continues on next page)

(continued from previous page)

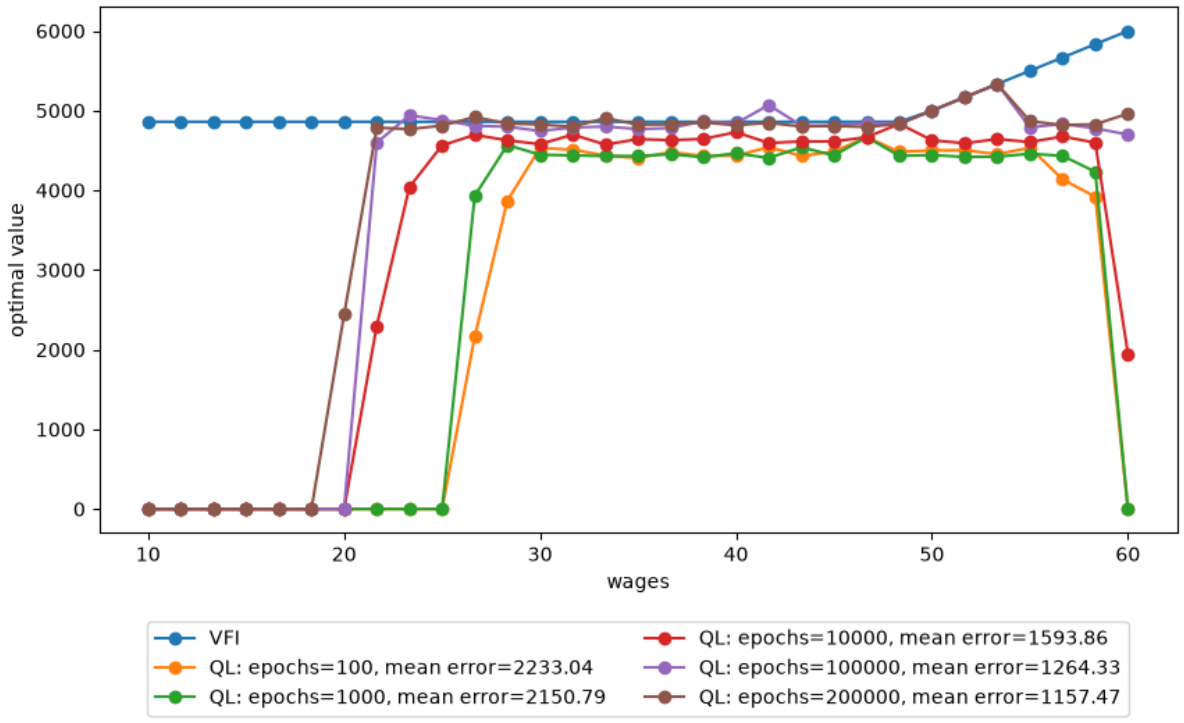
```
            ax.plot(w_new, valfunc_qlr, '-o',
                    label=f'QL: epochs={n}, mean error={error:.2f}')


        new_qtable = qlmc_new.run_one_epoch(qtable, rng)
        qtable = new_qtable

    ax.set_xlabel('wages')
    ax.set_ylabel('optimal value')
    ax.legend(bbox_to_anchor=(0.5, -0.15), loc='upper center', ncol=2)
    plt.subplots_adjust(bottom=0.25)
    plt.show()
```

```
plot_epochs(epochs_to_plot=[100, 1000, 10000, 100000, 200000])
```

```
Progress: EPOCHs = 0
Progress: EPOCHs = 20000
Progress: EPOCHs = 40000
Progress: EPOCHs = 60000
Progress: EPOCHs = 80000
Progress: EPOCHs = 100000
Progress: EPOCHs = 120000
Progress: EPOCHs = 140000
Progress: EPOCHs = 160000
Progress: EPOCHs = 180000
Progress: EPOCHs = 200000
```



The above graphs indicates that

- the Q-learning algorithm has trouble learning the Q-table well for wages that are rarely drawn
- the quality of approximation to the "true" value function computed by value function iteration improves for longer

epochs

65.6 Employed worker can't quit

The preceding version of temporal difference Q-learning described in equation system (65.8) lets an employed worker quit, i.e., reject her wage as an incumbent and instead receive unemployment compensation this period and draw a new offer next period.

This is an option that the McCall worker described in *this quantecon lecture* would not take.

See [Ljungqvist and Sargent, 2018], chapter 6 on search, for a proof.

But in the context of Q-learning, giving the worker the option to quit and get unemployment compensation while unemployed turns out to accelerate the learning process by promoting experimentation versus premature exploitation only.

To illustrate this, we'll amend our formulas for temporal differences to forbid an employed worker from quitting a job she had accepted earlier.

With this understanding about available choices, we obtain the following temporal difference values:

$$\begin{aligned} \widetilde{TD}\left(w, \text{accept}\right) &= \left[w + \beta \widetilde{Q}^{old}\left(w, \text{accept}\right)\right] - \widetilde{Q}^{old}\left(w, \text{accept}\right) \\ \widetilde{TD}\left(w, \text{reject}\right) &= \left[c + \beta \max_{a' \in \mathcal{A}} \widetilde{Q}^{old}\left(w', a'\right)\right] - \widetilde{Q}^{old}\left(w, \text{reject}\right), \; w' \sim F \end{aligned} \tag{65.9}$$

It turns out that formulas (65.9) combined with our Q-learning recursion (65.7) can lead our agent to eventually learn the optimal value function as well as in the case where an option to redraw can be exercised.

But learning is slower because an agent who ends up accepting a wage offer prematurely loses the option to explore new states in the same episode and to adjust the value associated with that state.

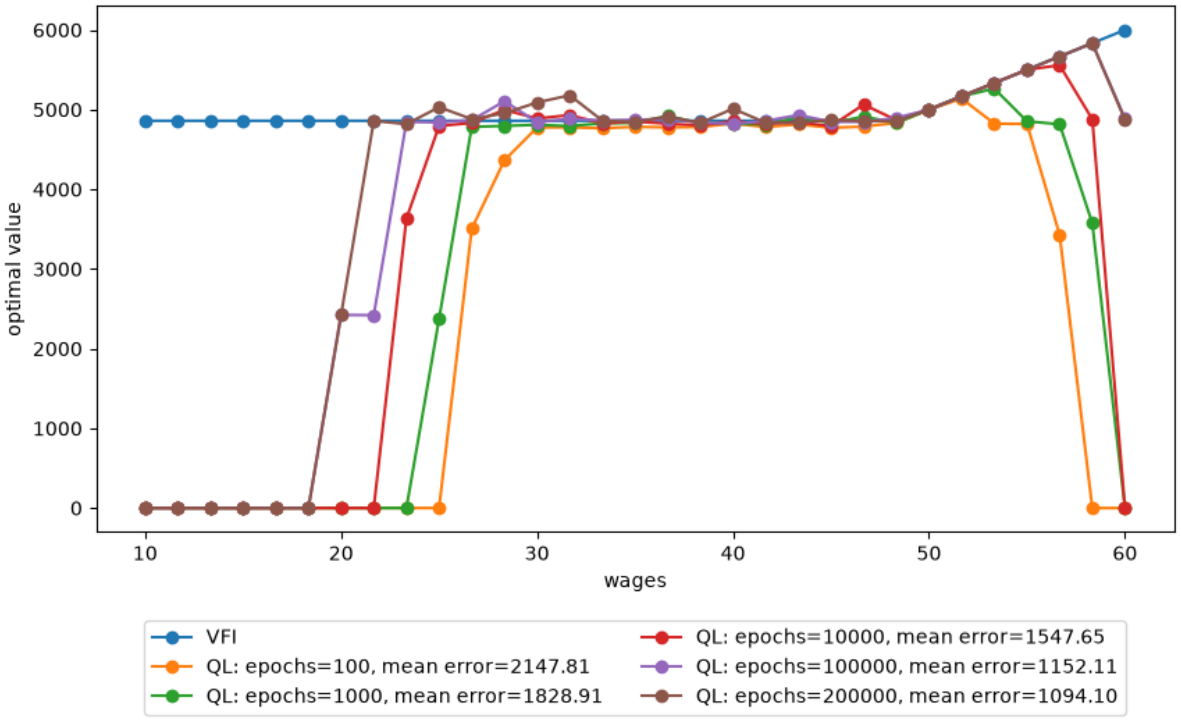
This can lead to inferior outcomes when the number of epochs/episodes is low.

But if we increase the number of epochs/episodes, we can observe that the error decreases and the outcomes get better.

We illustrate these possibilities with the following code and graph.

```
plot_epochs(epochs_to_plot=[100, 1000, 10000, 100000, 200000], quit_allowed=0)
```

```
Progress: EPOCHs = 0
Progress: EPOCHs = 20000
Progress: EPOCHs = 40000
Progress: EPOCHs = 60000
Progress: EPOCHs = 80000
Progress: EPOCHs = 100000
Progress: EPOCHs = 120000
Progress: EPOCHs = 140000
Progress: EPOCHs = 160000
Progress: EPOCHs = 180000
Progress: EPOCHs = 200000
```



65.7 Possible extensions

To extend the algorithm to handle problems with continuous state spaces, a typical approach is to restrict Q-functions and policy functions to take particular functional forms.

This is the approach in **deep Q-learning** where the idea is to use a multilayer neural network as a good function approximator.

We will take up this topic in a subsequent quantecon lecture.

Part IX

Introduction to Optimal Savings

CHAPTER

SIXTYSIX

OPTIMAL SAVINGS I: CAKE EATING

Contents

- *Optimal Savings I: Cake Eating*
  - *Overview*
  - *The model*
  - *The value function*
  - *The optimal policy*
  - *The Euler equation*
  - *Exercises*

66.1 Overview

In this lecture we introduce a simple "cake eating" problem.

The intertemporal problem is: how much to enjoy today and how much to leave for the future?

This trade-off between current and future rewards is at the heart of many savings and consumption problems.

Once we master the ideas in a simple environment, we will apply them to progressively more challenging problems.

The main tool we will use to solve the cake eating problem is dynamic programming.

The following lectures contain background on dynamic programming and might be worth reviewing:

- The shortest paths lecture
- The *basic McCall model*
- The *McCall model with separation*
- The *McCall model with separation and a continuous wage distribution*

We require the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

66.2 The model

We consider a discrete time model with an infinite time horizon.

At $t = 0$ the agent is given a complete cake with size $\bar{x}$.

Let x_t denote the size of the cake at the beginning of each period, so that, in particular, $x_0 = \bar{x}$.

We choose how much of the cake to eat in any given period t.

After choosing to consume c_t of the cake in period t, the amount left in period $t + 1$ is

$$x_{t+1} = x_t - c_t$$

Consuming quantity c of the cake gives current utility $u(c)$.

We adopt the CRRA utility function

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma} \qquad (\gamma > 0,\ \gamma \neq 1) \tag{66.1}$$

In Python this is

```
def u(c, γ):

    return c**(1 - γ) / (1 - γ)
```

Future cake consumption utility is discounted according to $\beta \in (0, 1)$.

In particular, consumption of c units t periods hence has present value $\beta^t u(c)$

The agent's problem can be written as

$$\max_{\{c_t\}} \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{66.2}$$

subject to

$$x_{t+1} = x_t - c_t \quad \text{and} \quad 0 \leq c_t \leq x_t \tag{66.3}$$

for all t.

A consumption path $\{c_t\}$ satisfying (66.3) and $x_0 = \bar{x}$ is called **feasible**.

In this problem, the following terminology is standard:

- x_t is called the **state variable**
- c_t is called the **control variable** or the **action**
- β and γ are **parameters**

66.2.1 Trade-off

The key trade-off in the cake-eating problem is this:

- Delaying consumption is costly because of the discount factor.
- But delaying some consumption is also attractive because u is concave.

The concavity of u implies that the consumer gains value from **consumption smoothing**, which means spreading consumption out over time.

This is because concavity implies diminishing marginal utility—a progressively smaller gain in utility for each additional spoonful of cake consumed within one period.

66.2.2 Intuition

The reasoning given above suggests that the discount factor β and the curvature parameter γ will play a key role in determining the rate of consumption.

Here's an educated guess as to what impact these parameters will have.

1. Higher β implies less discounting, and hence the agent is more patient, which should reduce the rate of consumption.
2. Higher γ implies more curvature in u, more desire for consumption smoothing, and hence a lower rate of consumption.

Note

More formally, higher γ implies a lower intertemporal elasticity of substitution, since IES = $1/\gamma$.

This means the consumer is less willing to substitute consumption between periods.

This stronger preference for consumption smoothing results in a lower consumption rate.

In summary, we expect the rate of consumption to be decreasing in both parameters.

Let's see if this is true.

66.3 The value function

The first step of our dynamic programming treatment is to obtain the Bellman equation.

The next step is to use it to calculate the solution.

66.3.1 The Bellman equation

To this end, we let $v(x)$ be maximum lifetime utility attainable from the current time when x units of cake are left.

That is,

$$v(x) = \max \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{66.4}$$

where the maximization is over all paths $\{c_t\}$ that are feasible from $x_0 = x$.

At this point, we do not have an expression for v, but we can still make inferences about it.

For example, as was the case with the *McCall model*, the value function will satisfy a version of the *Bellman equation*.

In the present case, this equation states that v satisfies

$$v(x) = \max_{0 \leq c \leq x} \{u(c) + \beta v(x - c)\} \quad \text{for any given } x \geq 0. \tag{66.5}$$

The intuition here is essentially the same as it was for the McCall model.

Choosing c optimally means trading off current vs future rewards.

Current rewards from choice c are just $u(c)$.

Future rewards given current cake size x, measured from next period and assuming optimal behavior, are $v(x - c)$.

These are the two terms on the right hand side of (66.5), after suitable discounting.

If c is chosen optimally using this trade off strategy, then we obtain maximal lifetime rewards from our current state x.

Hence, $v(x)$ equals the right hand side of (66.5), as claimed.

66.3.2 An analytical solution

It has been shown that, with u as the CRRA utility function in (66.1), the function v^* given by

$$v^*(x) = \left(1 - \beta^{1/\gamma}\right)^{-\gamma} u(x) \tag{66.6}$$

solves the Bellman equation and hence is equal to the value function.

You are asked to confirm that this is true in the exercises below.

Note

The solution (66.6) depends heavily on the CRRA utility function.

In fact, if we move away from CRRA utility, usually there is no analytical solution at all.

In other words, beyond CRRA utility, we know that the value function still satisfies the Bellman equation, but we do not have a way of writing it explicitly, as a function of the state variable and the parameters.

We will deal with that situation numerically in the following lectures.

Here is a Python representation of the value function:

```
def v_star(x, β, γ):

    return (1 - β**(1 / γ))**(-γ) * u(x, γ)
```

And here's a figure showing the function for fixed parameters:

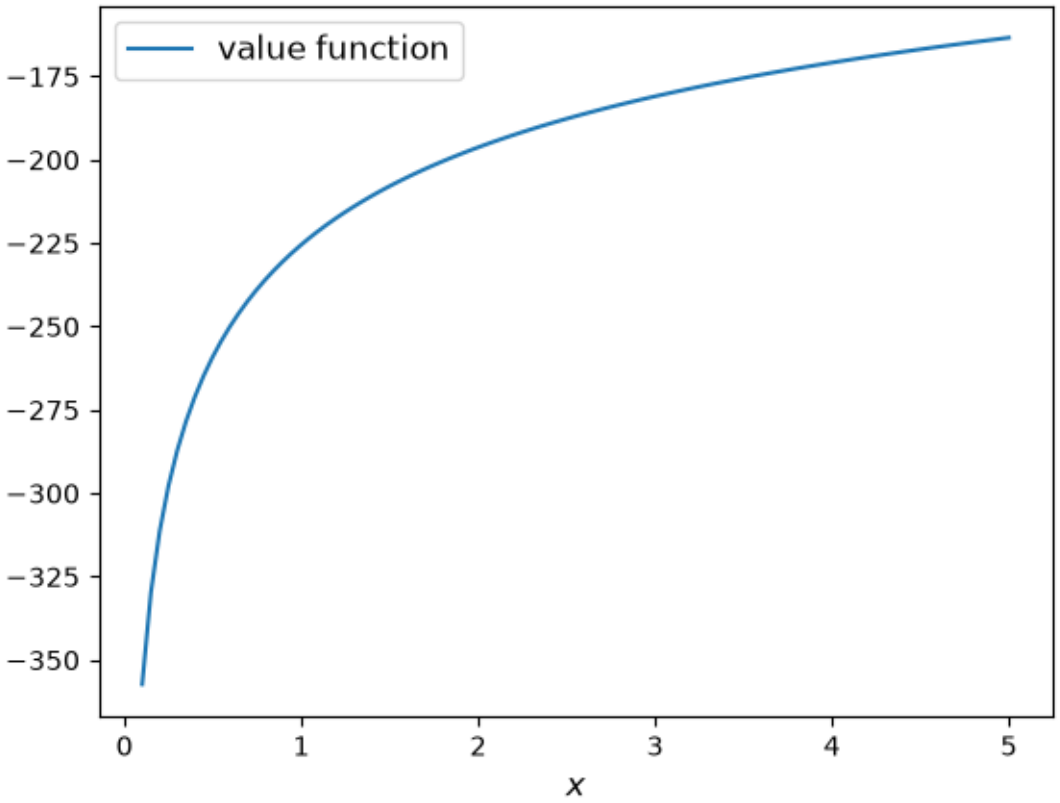
```
β, γ = 0.95, 1.2
x_grid = np.linspace(0.1, 5, 100)

fig, ax = plt.subplots()

ax.plot(x_grid, v_star(x_grid, β, γ), label='value function')

ax.set_xlabel('$x$', fontsize=12)
ax.legend(fontsize=12)

plt.show()
```



66.4 The optimal policy

Now that we have the value function v^*, it is straightforward to calculate the optimal action at each state.

We should choose consumption to maximize the right hand side of the Bellman equation (66.5).

$$c^* = \arg \max_{0 \leq c \leq x} \{u(c) + \beta v^*(x - c)\}$$

We can think of this optimal choice as a *function* of the state x, in which case we call it the **optimal policy**.

We denote the optimal policy by σ^*, so that

$$\sigma^*(x) := \arg \max_c \{u(c) + \beta v^*(x - c)\} \quad \text{for all} \quad x \geq 0$$

If we plug the analytical expression (66.6) for the value function into the right hand side and compute the optimum, we find that

$$\sigma^*(x) = \left(1 - \beta^{1/\gamma}\right) x \tag{66.7}$$

Now let's recall our intuition on the impact of parameters.

We guessed that consumption would be decreasing in both parameters.

This is in fact the case, as can be seen from (66.7).

Here are some plots that illustrate.

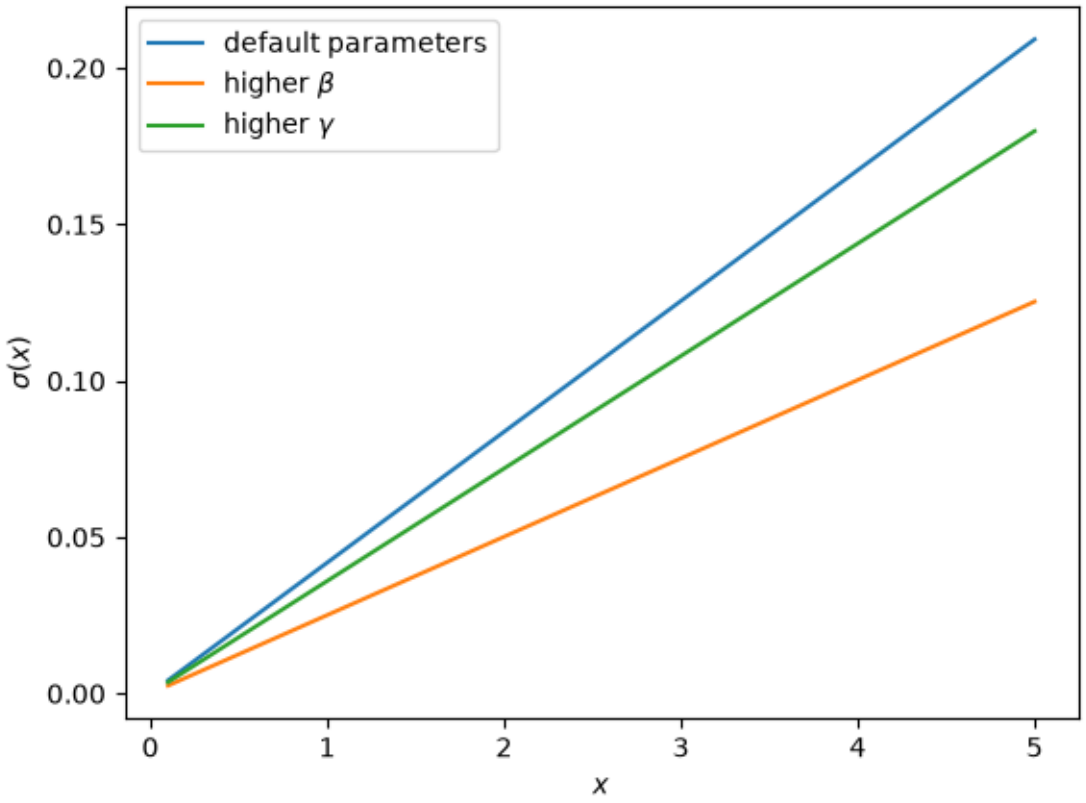
```
def c_star(x, β, γ):

    return (1 - β ** (1/γ)) * x
```

Continuing with the values for β and γ used above, the plot is

```
fig, ax = plt.subplots()
ax.plot(x_grid, c_star(x_grid, β, γ), label='default parameters')
ax.plot(x_grid, c_star(x_grid, β + 0.02, γ), label=r'higher $\beta$')
ax.plot(x_grid, c_star(x_grid, β, γ + 0.2), label=r'higher $\gamma$')
ax.set_ylabel(r'$\sigma(x)$')
ax.set_xlabel('$x$')
ax.legend()

plt.show()
```



66.5 The Euler equation

In the discussion above we have provided a complete solution to the cake eating problem in the case of CRRA utility.

There is in fact another way to solve for the optimal policy, based on the so-called **Euler equation**.

Although we already have a complete solution, now is a good time to study the Euler equation.

This is because, for more difficult problems, this equation provides key insights that are hard to obtain by other methods.

66.5.1 Statement and implications

The Euler equation for the present problem can be stated as

$$u'(c_t^*) = \beta u'(c_{t+1}^*) \tag{66.8}$$

This is a necessary condition for an optimal consumption path $\{c_t^*\}_{t\geq 0}$.

It says that, along the optimal path, marginal rewards are equalized across time, after appropriate discounting.

This makes sense: optimality is obtained by smoothing consumption up to the point where no marginal gains remain.

We can also state the Euler equation in terms of the policy function.

A **feasible consumption policy** is a map $x \mapsto \sigma(x)$ satisfying $0 \leq \sigma(x) \leq x$.

(The last restriction says that we cannot consume more than the remaining quantity of cake.)

A feasible consumption policy σ is said to **satisfy the Euler equation** if, for all $x > 0$,

$$u'(\sigma(x)) = \beta u'(\sigma(x - \sigma(x))) \tag{66.9}$$

Evidently (66.9) is just the policy equivalent of (66.8).

It turns out that a feasible policy is optimal if and only if it satisfies the Euler equation.

In the exercises, you are asked to verify that the optimal policy (66.7) does indeed satisfy this functional equation.

Note

A **functional equation** is an equation where the unknown object is a function.

For a proof of sufficiency of the Euler equation in a very general setting, see proposition 2.2 of [Ma *et al.*, 2020].

The following arguments focus on necessity, explaining why an optimal path or policy should satisfy the Euler equation.

66.5.2 Derivation I: a perturbation approach

Let's write c as a shorthand for consumption path $\{c_t\}_{t=0}^{\infty}$.

The overall cake-eating maximization problem can be written as

$$\max_{c \in F} U(c) \quad \text{where } U(c) := \sum_{t=0}^{\infty} \beta^t u(c_t)$$

and F is the set of feasible consumption paths.

We know that differentiable functions have a zero gradient at a maximizer.

So the optimal path $c^* := \{c_t^*\}_{t=0}^{\infty}$ must satisfy $U'(c^*) = 0$.

Note

If you want to know exactly how the derivative $U'(c^*)$ is defined, given that the argument c^* is a vector of infinite length, you can start by learning about Gateaux derivatives. However, such knowledge is not assumed in what follows.

In other words, the rate of change in U must be zero for any infinitesimally small (and feasible) perturbation away from the optimal path.

So consider a feasible perturbation that reduces consumption at time t to $c_t^* - h$ and increases it in the next period to $c_{t+1}^* + h$.

Consumption does not change in any other period.

We call this perturbed path c^h.

By the preceding argument about zero gradients, we have

$$\lim_{h \to 0} \frac{U(c^h) - U(c^*)}{h} = U'(c^*) = 0$$

Recalling that consumption only changes at t and $t+1$, this becomes

$$\lim_{h \to 0} \frac{\beta^t u(c_t^* - h) + \beta^{t+1} u(c_{t+1}^* + h) - \beta^t u(c_t^*) - \beta^{t+1} u(c_{t+1}^*)}{h} = 0$$

After rearranging, the same expression can be written as

$$\lim_{h \to 0} \frac{u(c_t^* - h) - u(c_t^*)}{h} + \beta \lim_{h \to 0} \frac{u(c_{t+1}^* + h) - u(c_{t+1}^*)}{h} = 0$$

or, taking the limit,

$$-u'(c_t^*) + \beta u'(c_{t+1}^*) = 0$$

This is just the Euler equation.

66.5.3 Derivation II: using the Bellman equation

Another way to derive the Euler equation is to use the Bellman equation (66.5).

Note

The argument that follows assumes that the value function is differentiable. A proof of differentiability of the value function can be found in EDTC, theorem 10.1.13.

Taking the derivative on the right hand side of the Bellman equation with respect to c and setting it to zero, we get

$$u'(c) = \beta v'(x - c) \tag{66.10}$$

To obtain $v'(x-c)$, we set $g(c, x) = u(c) + \beta v(x - c)$, so that, at the optimal choice of consumption,

$$v(x) = g(c, x) \tag{66.11}$$

Differentiating both sides while acknowledging that the maximizing consumption will depend on x, we get

$$v'(x) = \frac{\partial}{\partial c} g(c, x) \frac{\partial c}{\partial x} + \frac{\partial}{\partial x} g(c, x)$$

When $g(c, x)$ is maximized at c, we have $\frac{\partial}{\partial c} g(c, x) = 0$.

Hence the derivative simplifies to

$$v'(x) = \frac{\partial g(c, x)}{\partial x} = \frac{\partial}{\partial x} \beta v(x - c) = \beta v'(x - c) \tag{66.12}$$

(This argument is an example of the Envelope Theorem.)

But now an application of (66.10) gives

$$u'(c) = v'(x) \tag{66.13}$$

Thus, the derivative of the value function is equal to marginal utility.

Combining this fact with (66.12) recovers the Euler equation.

66.6 Exercises

Exercise 66.6.1

How does one obtain the expressions for the value function and optimal policy given in (66.6) and (66.7) respectively?

The first step is to make a guess of the functional form for the consumption policy.

So suppose that we do not know the solutions and start with a guess that the optimal policy is linear.

In other words, we conjecture that there exists a positive θ such that setting $c_t^* = \theta x_t$ for all t produces an optimal path.

Starting from this conjecture, try to obtain the solutions (66.6) and (66.7).

In doing so, you will need to use the definition of the value function and the Bellman equation.

Solution

We start with the conjecture $c_t^* = \theta x_t$, which leads to a path for the state variable (cake size) given by

$$x_{t+1} = x_t(1-\theta)$$

Then $x_t = x_0(1-\theta)^t$ and hence

$$\begin{aligned} v(x_0) &= \sum_{t=0}^{\infty} \beta^t u(\theta x_t) \\ &= \sum_{t=0}^{\infty} \beta^t u(\theta x_0 (1-\theta)^t) \\ &= \sum_{t=0}^{\infty} \theta^{1-\gamma} \beta^t (1-\theta)^{t(1-\gamma)} u(x_0) \\ &= \frac{\theta^{1-\gamma}}{1-\beta(1-\theta)^{1-\gamma}} u(x_0) \end{aligned}$$

From the Bellman equation, then,

$$\begin{aligned} v(x) &= \max_{0 \le c \le x} \left\{ u(c) + \beta \frac{\theta^{1-\gamma}}{1-\beta(1-\theta)^{1-\gamma}} \cdot u(x-c) \right\} \\ &= \max_{0 \le c \le x} \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \beta \frac{\theta^{1-\gamma}}{1-\beta(1-\theta)^{1-\gamma}} \cdot \frac{(x-c)^{1-\gamma}}{1-\gamma} \right\} \end{aligned}$$

From the first order condition, we obtain

$$c^{-\gamma} + \beta \frac{\theta^{1-\gamma}}{1-\beta(1-\theta)^{1-\gamma}} \cdot (x-c)^{-\gamma}(-1) = 0$$

or

$$c^{-\gamma} = \beta \frac{\theta^{1-\gamma}}{1-\beta(1-\theta)^{1-\gamma}} \cdot (x-c)^{-\gamma}$$

With $c = \theta x$ we get

$$(\theta x)^{-\gamma} = \beta \frac{\theta^{1-\gamma}}{1-\beta(1-\theta)^{1-\gamma}} \cdot (x(1-\theta))^{-\gamma}$$

Some rearrangement produces

$$\theta = 1 - \beta^{\frac{1}{\gamma}}$$

This confirms our earlier expression for the optimal policy:

$$c_t^* = \left(1 - \beta^{\frac{1}{\gamma}}\right) x_t$$

Substituting θ into the value function above gives

$$v^*(x_t) = \frac{\left(1-\beta^{\frac{1}{\gamma}}\right)^{1-\gamma}}{1-\beta\left(\beta^{\frac{1-\gamma}{\gamma}}\right)} u(x_t)$$

Rearranging gives

$$v^*(x_t) = \left(1-\beta^{\frac{1}{\gamma}}\right)^{-\gamma} u(x_t)$$

Our claims are now verified.

Exercise 66.6.2

Verify that the optimal policy (66.7) satisfies the Euler equation (66.9).

Solution

Recall that the optimal policy is

$$\sigma^*(x) = (1 - \beta^{1/\gamma})x$$

and the Euler equation in policy form is

$$u'(\sigma(x)) = \beta u'(\sigma(x - \sigma(x)))$$

With CRRA utility $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$, the marginal utility is

$$u'(c) = c^{-\gamma}$$

Now let's verify the Euler equation. The left-hand side is

$$u'(\sigma^*(x)) = [\sigma^*(x)]^{-\gamma} = [(1 - \beta^{1/\gamma})x]^{-\gamma} = (1 - \beta^{1/\gamma})^{-\gamma} x^{-\gamma}$$

For the right-hand side, we first compute the state in the next period:

$$x - \sigma^*(x) = x - (1 - \beta^{1/\gamma})x = x\beta^{1/\gamma}$$

Next period's consumption under the optimal policy is

$$\sigma^*(x - \sigma^*(x)) = \sigma^*(x\beta^{1/\gamma}) = (1 - \beta^{1/\gamma}) \cdot x\beta^{1/\gamma}$$

Therefore, the right-hand side of the Euler equation is

$$\begin{aligned}
\beta u'(\sigma^*(x - \sigma^*(x))) &= \beta[(1 - \beta^{1/\gamma}) \cdot x\beta^{1/\gamma}]^{-\gamma} \\
&= \beta(1 - \beta^{1/\gamma})^{-\gamma} x^{-\gamma} (\beta^{1/\gamma})^{-\gamma} \\
&= \beta(1 - \beta^{1/\gamma})^{-\gamma} x^{-\gamma} \beta^{-1} \\
&= (1 - \beta^{1/\gamma})^{-\gamma} x^{-\gamma}
\end{aligned}$$

Since the left-hand side equals the right-hand side, the optimal policy σ^* satisfies the Euler equation.

CHAPTER

SIXTYSEVEN

OPTIMAL SAVINGS II: NUMERICAL CAKE EATING

Contents

- *Optimal Savings II: Numerical Cake Eating*
  - *Overview*
  - *Reviewing the Model*
  - *Value Function Iteration*
  - *Exercises*

67.1 Overview

In this lecture we continue the study of the problem described in *Optimal Savings I: Cake Eating*.

The aim of this lecture is to solve the problem using numerical methods.

At first this might appear unnecessary, since we already obtained the optimal policy analytically.

However, the cake eating problem is too simple to be useful without modifications, and once we start modifying the problem, numerical methods become essential.

Hence it makes sense to introduce numerical methods now, and test them on this simple problem.

Since we know the analytical solution, this will allow us to assess the accuracy of alternative numerical methods.

Note

The code below aims for clarity rather than maximum efficiency.

In the lectures below we will explore best practice for speed and efficiency.

Let's put these algorithm and code optimizations to one side for now.

We will use the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.optimize import minimize_scalar, bisect
from typing import NamedTuple
```

67.2 Reviewing the Model

You might like to review the details in *Optimal Savings I: Cake Eating* before we start.

Recall in particular that the Bellman equation is

$$v(x) = \max_{0 \leq c \leq x} \{u(c) + \beta v(x - c)\} \quad \text{for all } x \geq 0. \tag{67.1}$$

where u is the CRRA utility function.

The analytical solutions for the value function and optimal policy were found to be as follows.

```
def c_star(x, β, γ):

    return (1 - β ** (1/γ)) * x


def v_star(x, β, γ):

    return (1 - β**(1 / γ))**(-γ) * (x**(1-γ) / (1-γ))
```

Our first aim is to obtain these analytical solutions numerically.

67.3 Value Function Iteration

The first approach we will take is **value function iteration**.

This is a form of **successive approximation**, and was discussed in our *lecture on job search*.

The basic idea is:

1. Take an arbitrary initial guess of v.
2. Obtain an update $\hat{v}$ defined by

$$\hat{v}(x) = \max_{0 \leq c \leq x} \{u(c) + \beta v(x - c)\}$$

3. Stop if $\hat{v}$ is approximately equal to v, otherwise set $v = \hat{v}$ and go back to step 2.

Let's write this a bit more mathematically.

67.3.1 The Bellman Operator

We introduce the **Bellman operator** T that takes a function v as an argument and returns a new function Tv defined by

$$Tv(x) = \max_{0 \leq c \leq x} \{u(c) + \beta v(x - c)\}$$

From v we get Tv, and applying T to this yields $T^2 v := T(Tv)$ and so on.

This is called **iterating with the Bellman operator** from initial guess v.

As we discuss in more detail in later lectures, one can use Banach's contraction mapping theorem to prove that the sequence of functions $T^n v$ converges to the solution to the Bellman equation.

67.3.2 Fitted Value Function Iteration

Both consumption c and the state variable x are continuous.

This causes complications when it comes to numerical work.

For example, we need to store each function $T^n v$ in order to compute the next iterate $T^{n+1} v$.

But this means we have to store $T^n v(x)$ at infinitely many x, which is, in general, impossible.

To circumvent this issue we will use fitted value function iteration, as discussed previously in *one of the lectures* on job search.

The process looks like this:

1. Begin with an array of values $\{v_0, \ldots, v_I\}$ representing the values of some initial function v on the grid points $\{x_0, \ldots, x_I\}$.
2. Build a function $\hat{v}$ on the state space $\mathbb{R}_+$ by interpolation, based on the interpolation points $\{(x_i, v_i)\}$.
3. Insert $\hat{v}$ into the right hand side of the Bellman equation and obtain and record the value $T\hat{v}(x_i)$ on each grid point x_i.
4. Unless some stopping condition is satisfied, set $\{v_0, \ldots, v_I\} = \{T\hat{v}(x_0), \ldots, T\hat{v}(x_I)\}$ and go to step 2.

In step 2 we'll use piecewise linear interpolation.

67.3.3 Implementation

The `maximize` function below is a small helper function that converts a SciPy minimization routine into a maximization routine.

```
def maximize(g, upper_bound):
    """
    Maximize the function g over the interval [0, upper_bound].

    We use the fact that the maximizer of g on any interval is
    also the minimizer of -g.

    """

    objective = lambda x: -g(x)
    bounds = (0, upper_bound)
    result = minimize_scalar(objective, bounds=bounds, method='bounded')
    maximizer, maximum = result.x, -result.fun
    return maximizer, maximum
```

We'll store the parameters β and γ and the grid in a `NamedTuple` called `Model`.

We'll also create a helper function called `create_cake_eating_model` to store default parameters and build an instance of `Model`.

```
class Model(NamedTuple):
    β: float
    γ: float
    x_grid: np.ndarray

def create_cake_eating_model(
        β: float = 0.96,           # discount factor
        γ: float = 1.5,            # degree of relative risk aversion
```

(continues on next page)

(continued from previous page)

```
        x_grid_min: float = 1e-3,  # exclude zero for numerical stability
        x_grid_max: float = 2.5,   # size of cake
        x_grid_size: int = 120
    ):
    """
    Creates an instance of the cake eating model.

    """
    x_grid = np.linspace(x_grid_min, x_grid_max, x_grid_size)
    return Model(β, γ, x_grid)
```

Here's the CRRA utility function.

```
def u(c, γ):
    return (c ** (1 - γ)) / (1 - γ)
```

To work with the Bellman equation, let's write it as

$$v(x) = \max_{0 \leq c \leq x} B(x, c, v)$$

where

$$B(x, c, v) := u(c) + \beta v(x - c)$$

Now we implement the function B.

```
def B(
        x: float,                # the current state (remaining cake)
        c: float,                # current consumption
        v: np.ndarray,           # current guess of the value function
        model: Model             # instance of cake eating model
    ):
    """
    Right hand side of the Bellman equation given x and c.

    """
    # Unpack (simplify names)
    β, γ, x_grid = model

    # Convert array v into a function by linear interpolation
    vf = lambda x: np.interp(x, x_grid, v)

    # Return B(x, c, v)
    return u(c, γ) + β * vf(x - c)
```

We now define the Bellman operator acting on grid points:

$$Tv(x_i) = \max_{0 \leq c \leq x_i} B(x_i, c, v) \qquad \text{for all } i$$

```
def T(
        v: np.ndarray,           # Current guess of the value function
        model: Model             # Instance of the cake eating model
    ):
    " The Bellman operator.  Updates the guess of the value function. "

    # Allocate memory for the new array v_new = Tv
```

(continues on next page)

(continued from previous page)

```
    v_new = np.empty_like(v)

    # Calculate Tv(x) for all x
    for i, x in enumerate(model.x_grid):
        # Maximize RHS of Bellman equation with respect to c over [0, x]
        _, v_new[i] = maximize(lambda c: B(x, c, v, model), x)

    return v_new
```

After defining the Bellman operator, we are ready to solve the model.

Let's start by creating a model using the default parameterization.

```
model = create_cake_eating_model()
β, γ, x_grid = model
```

Now let's see iteration of the value function in action.

We start from guess v given by $v(x) = u(x)$ for every x grid point.

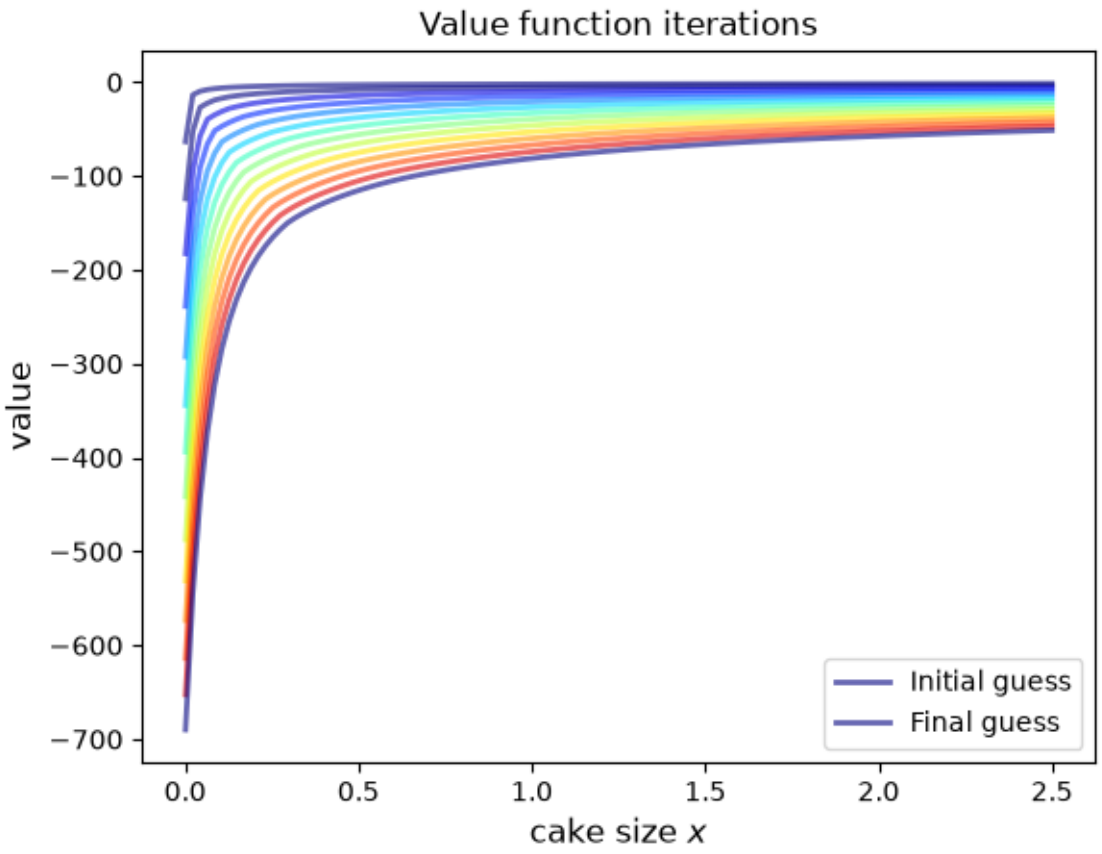
```
v = u(x_grid, γ)  # Initial guess
n = 12                  # Number of iterations
fig, ax = plt.subplots()

# Initial plot
ax.plot(x_grid, v, color=plt.cm.jet(0),
        lw=2, alpha=0.6, label='Initial guess')

# Iterate
for i in range(n):
    v = T(v, model)  # Apply the Bellman operator
    ax.plot(x_grid, v, color=plt.cm.jet(i / n), lw=2, alpha=0.6)

# One last update and plot
v = T(v, model)
ax.plot(x_grid, v, color=plt.cm.jet(0),
        lw=2, alpha=0.6, label='Final guess')

ax.legend()
ax.set_ylabel('value', fontsize=12)
ax.set_xlabel('cake size $x$', fontsize=12)
ax.set_title('Value function iterations')
plt.show()
```



To iterate more systematically, we introduce a wrapper function called `compute_value_function`.

It's task is to iterate using T until some convergence conditions are satisfied.

```
def compute_value_function(
        model: Model,
        tol: float = 1e-4,
        max_iter: int = 1_000,
        verbose: bool = True,
        print_skip: int = 25
    ):

    # Set up loop
    v = np.zeros(len(model.x_grid)) # Initial guess
    i = 0
    error = tol + 1

    while i < max_iter and error > tol:
        v_new = T(v, model)

        error = np.max(np.abs(v - v_new))
        i += 1

        if verbose and i % print_skip == 0:
            print(f"Error at iteration {i} is {error}.")

        v = v_new
```

(continues on next page)

(continued from previous page)

```
    if error > tol:
        print("Failed to converge!")
    elif verbose:
        print(f"\nConverged in {i} iterations.")

    return v_new
```

Now let's call it, noting that it takes a little while to run.

```
v = compute_value_function(model)
```

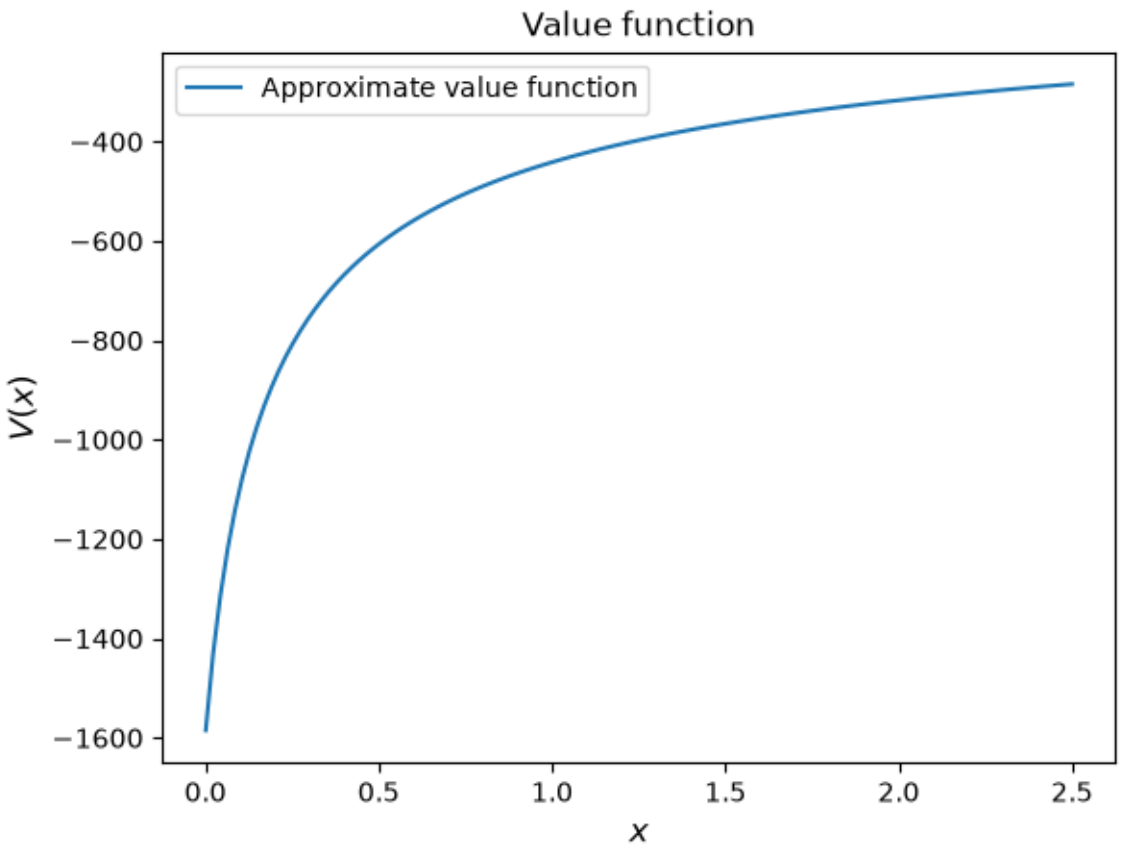
```
Error at iteration 25 is 23.80037552320357.
Error at iteration 50 is 8.577577198550443.
Error at iteration 75 is 3.091330660953872.
Error at iteration 100 is 1.1141054209301728.
Error at iteration 125 is 0.4015199359369035.
Error at iteration 150 is 0.1447064666647293.
Error at iteration 175 is 0.05215173549413521.
Error at iteration 200 is 0.018795314250610318.
Error at iteration 225 is 0.006773769548317432.
Error at iteration 250 is 0.0024412443060555233.
Error at iteration 275 is 0.0008798164330414693.
Error at iteration 300 is 0.0003170829540977138.
Error at iteration 325 is 0.00011427565573285392.

Converged in 329 iterations.
```

Now we can plot and see what the converged value function looks like.

```
fig, ax = plt.subplots()

ax.plot(x_grid, v, label='Approximate value function')
ax.set_ylabel('$V(x)$', fontsize=12)
ax.set_xlabel('$x$', fontsize=12)
ax.set_title('Value function')
ax.legend()
plt.show()
```

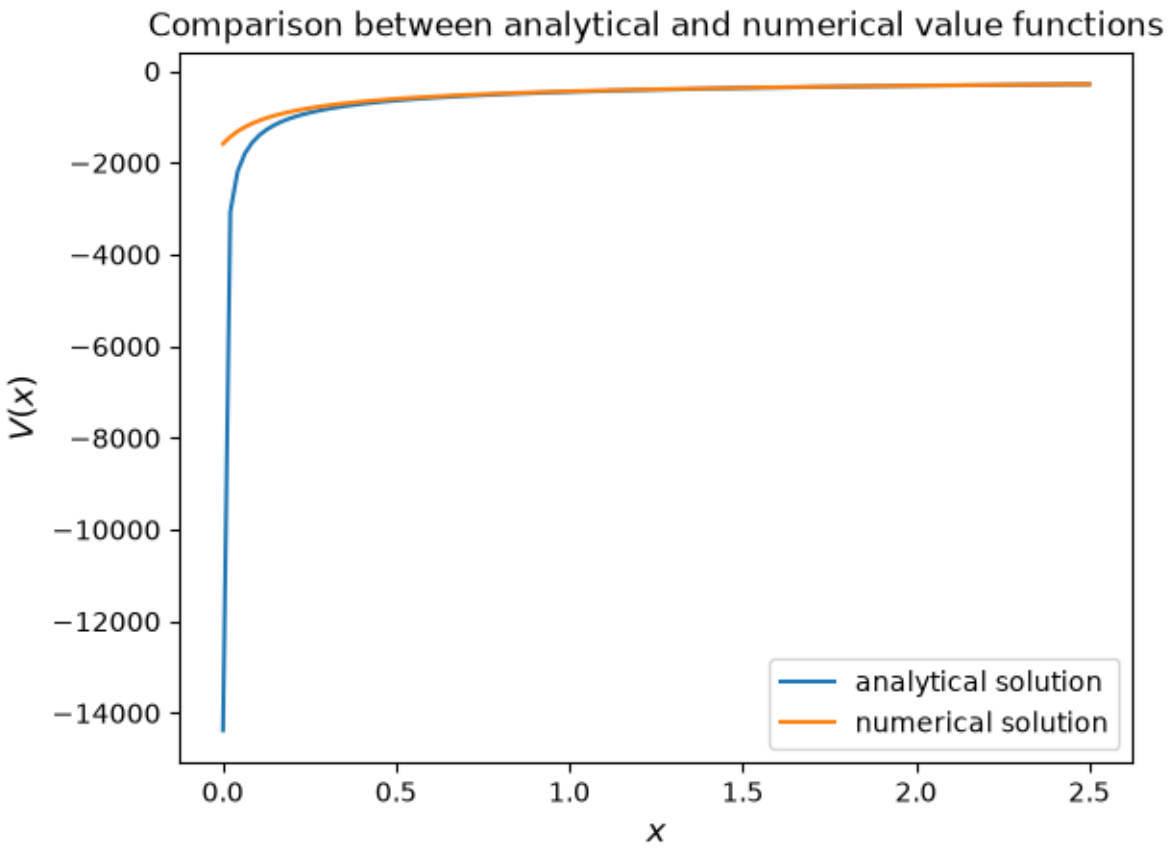


Next let's compare it to the analytical solution.

```
v_analytical = v_star(x_grid, β, γ)
```

```
fig, ax = plt.subplots()

ax.plot(x_grid, v_analytical, label='analytical solution')
ax.plot(x_grid, v, label='numerical solution')
ax.set_ylabel('$V(x)$', fontsize=12)
ax.set_xlabel('$x$', fontsize=12)
ax.legend()
ax.set_title('Comparison between analytical and numerical value functions')
plt.show()
```



The quality of approximation is reasonably good for large x, but less so near the lower boundary.

The reason is that the utility function and hence value function is very steep near the lower boundary, and hence hard to approximate.

Note

One way to fix this issue is to use a nonlinear grid, with more points in the neighborhood of zero.

Instead of pursuing this idea, however, we will turn our attention to working with policy functions.

We will see that value function iteration can be avoided by iterating on a guess of the policy function instead.

The policy function has less curvature and hence is easier to interpolate than the value function.

These ideas will be explored over the next few lectures.

67.3.4 Policy Function

Let's try computing the optimal policy.

In *Optimal Savings I: Cake Eating*, the optimal consumption policy was shown to be

$$\sigma^*(x) = \left(1 - \beta^{1/\gamma}\right) x$$

Let's see if our numerical results lead to something similar.

Our numerical strategy will be to compute, for any given v, the policy

$$\sigma(x) = \arg \max_{0 \leq c \leq x} \{u(c) + \beta v(x - c)\}$$

This policy is called the v-**greedy policy**.

In practice we will compute σ on a grid of x points and then interpolate.

For v we will use the approximation of the value function we obtained above.

Here's the function:

```
def get_greedy(
        v: np.ndarray,          # current guess of the value function
        model: Model            # instance of cake eating model
    ):
    " Compute the v-greedy policy on x_grid."

    σ = np.empty_like(v)

    for i, x in enumerate(model.x_grid):
        # Maximize RHS of Bellman equation at state x
        σ[i], _ = maximize(lambda c: B(x, c, v, model), x)

    return σ
```

Now let's pass the approximate value function and compute optimal consumption:

```
σ = get_greedy(v, model)
```

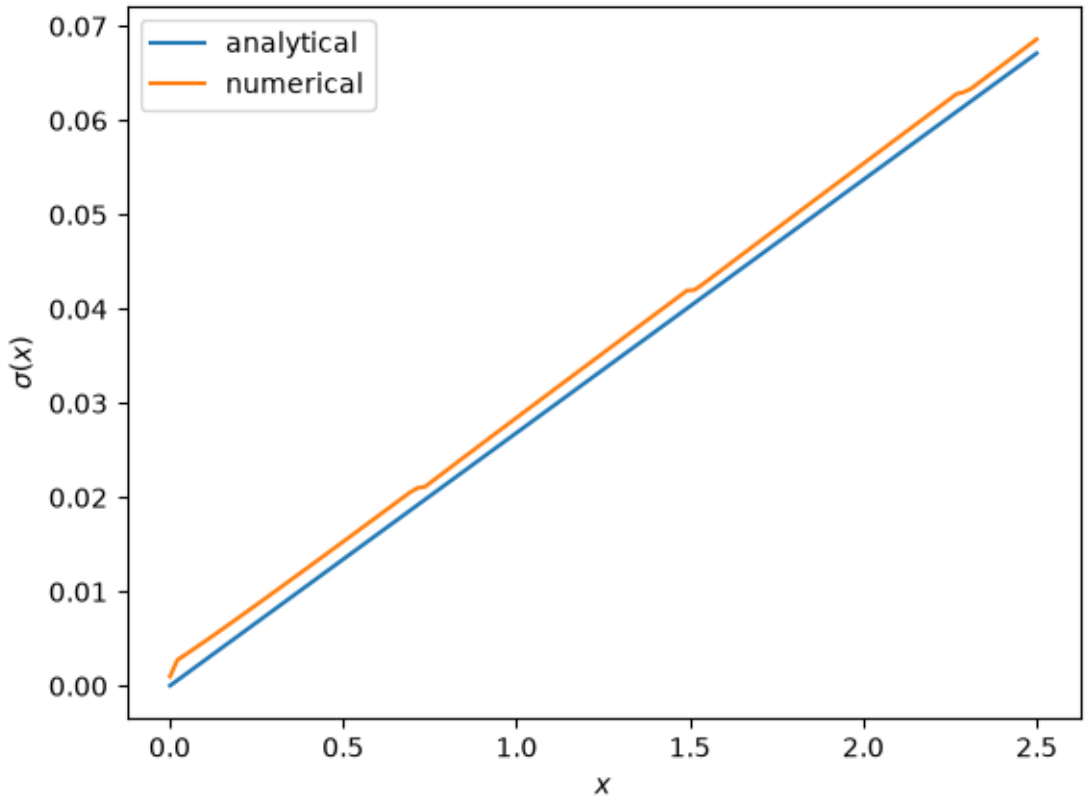
Let's plot this next to the true analytical solution

```
c_analytical = c_star(model.x_grid, model.β, model.γ)

fig, ax = plt.subplots()

ax.plot(model.x_grid, c_analytical, label='analytical')
ax.plot(model.x_grid, σ, label='numerical')
ax.set_ylabel(r'$\sigma(x)$')
ax.set_xlabel('$x$')
ax.legend()

plt.show()
```



The fit is reasonable but not perfect.

We can improve it by increasing the grid size or reducing the error tolerance in the value function iteration routine.

However, both changes will lead to a longer compute time.

Another possibility is to use an alternative algorithm, which offers the possibility of faster compute time and, at the same time, more accuracy.

We explore this in *Optimal Savings IV: Time Iteration*.

67.4 Exercises

Exercise 67.4.1

Try the following modification of the problem.

Instead of the cake size changing according to $x_{t+1} = x_t - c_t$, let it change according to

$$x_{t+1} = (x_t - c_t)^\alpha$$

where α is a parameter satisfying $0 < \alpha < 1$.

(We will see this kind of update rule when we study optimal growth models.)

Make the required changes to value function iteration code and plot the value and policy functions.

Try to reuse as much code as possible.

Solution

We need to create an extended version of our model and state-action value function.

We'll create a new `NamedTuple` for the extended cake model and a helper function.

```
# Create extended cake model data structure
class ExtendedModel(NamedTuple):
    β: float
    γ: float
    α: float
    x_grid: np.ndarray

def create_extended_model(β=0.96,           # discount factor
                          γ=1.5,            # degree of relative risk aversion
                          α=0.4,            # productivity parameter
                          x_grid_min=1e-3,  # exclude zero for numerical stability
                          x_grid_max=2.5,   # size of cake
                          x_grid_size=120):
    """
    Creates an instance of the extended cake eating model.
    """
    x_grid = np.linspace(x_grid_min, x_grid_max, x_grid_size)
    return ExtendedModel(β, γ, α, x_grid)

def extended_B(c, x, v, model):
    """
    Right hand side of the Bellman equation for the extended cake model given x␣
 ↪and c.

    """
    β, γ, α, x_grid = model
    vf = lambda x: np.interp(x, x_grid, v)
    return u(c, γ) + β * vf((x - c)**α)
```

We also need a modified Bellman operator:

```
def extended_T(v, model):
    " The Bellman operator for the extended cake model. "

    v_new = np.empty_like(v)
    for i, x in enumerate(model.x_grid):
        _, v_new[i] = maximize(lambda c: extended_B(c, x, v, model), x)
    return v_new
```

Now create the model:

```
model = create_extended_model()
```

Here's a function to compute the value function.

```
def compute_value_function_extended(model,
                                    tol=1e-4,
                                    max_iter=1000,
                                    verbose=True,
                                    print_skip=25):
    """
    Compute value function for extended cake model.
    """
    v = np.zeros(len(model.x_grid))
    i = 0
    error = tol + 1

    while i < max_iter and error > tol:
        v_new = extended_T(v, model)
        error = np.max(np.abs(v - v_new))
        i += 1
        if verbose and i % print_skip == 0:
            print(f"Error at iteration {i} is {error}.")
        v = v_new

    if error > tol:
        print("Failed to converge!")
    elif verbose:
        print(f"\nConverged in {i} iterations.")

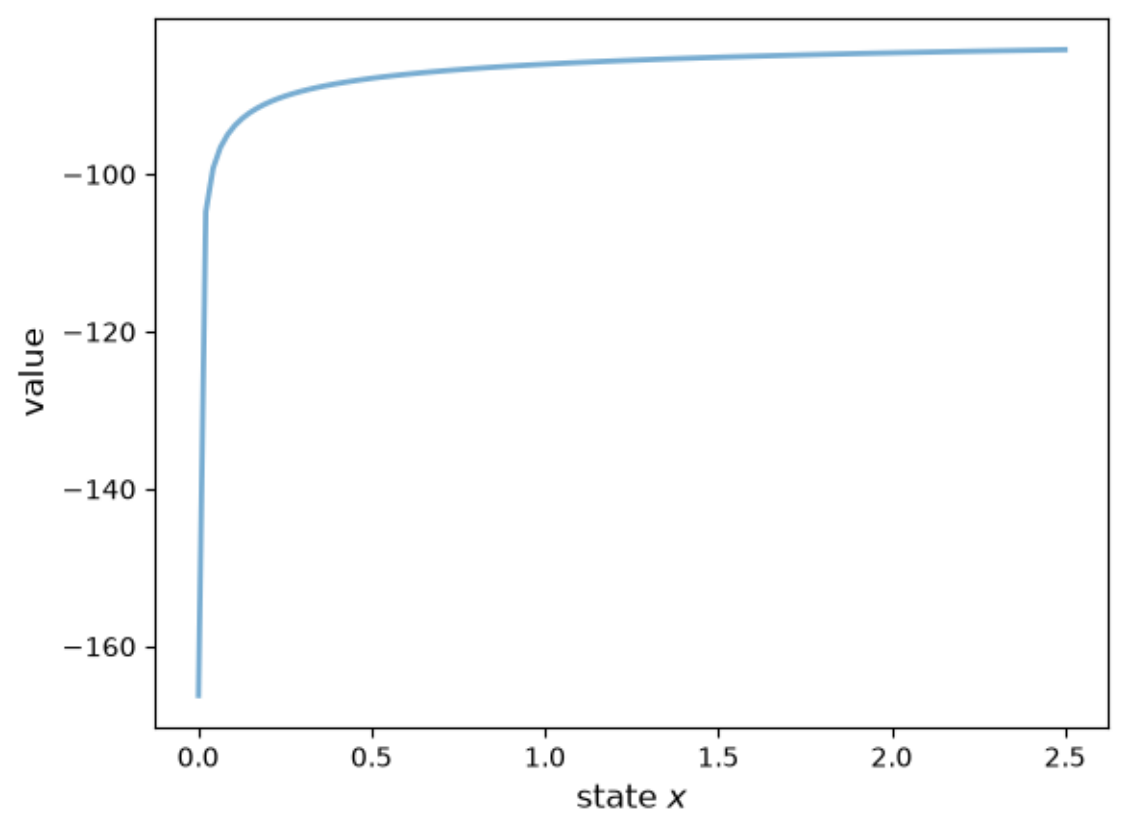
    return v_new

v = compute_value_function_extended(model, verbose=False)

fig, ax = plt.subplots()

ax.plot(model.x_grid, v, lw=2, alpha=0.6)
ax.set_ylabel('value', fontsize=12)
ax.set_xlabel('state $x$', fontsize=12)

plt.show()
```



Here's the computed policy, combined with the solution we derived above for the standard cake eating case $\alpha = 1$.

```
def extended_get_greedy(model, v):
    """
    The optimal policy function for the extended cake model.
    """
    σ = np.empty_like(v)

    for i, x in enumerate(model.x_grid):
        # Maximize extended_B with respect to c over [0, x]
        σ[i], _ = maximize(lambda c: extended_B(c, x, v, model), x)

    return σ

σ = extended_get_greedy(model, v)

# Get the baseline model for comparison
baseline_model = create_cake_eating_model()
c_analytical = c_star(baseline_model.x_grid, baseline_model.β, baseline_model.γ)

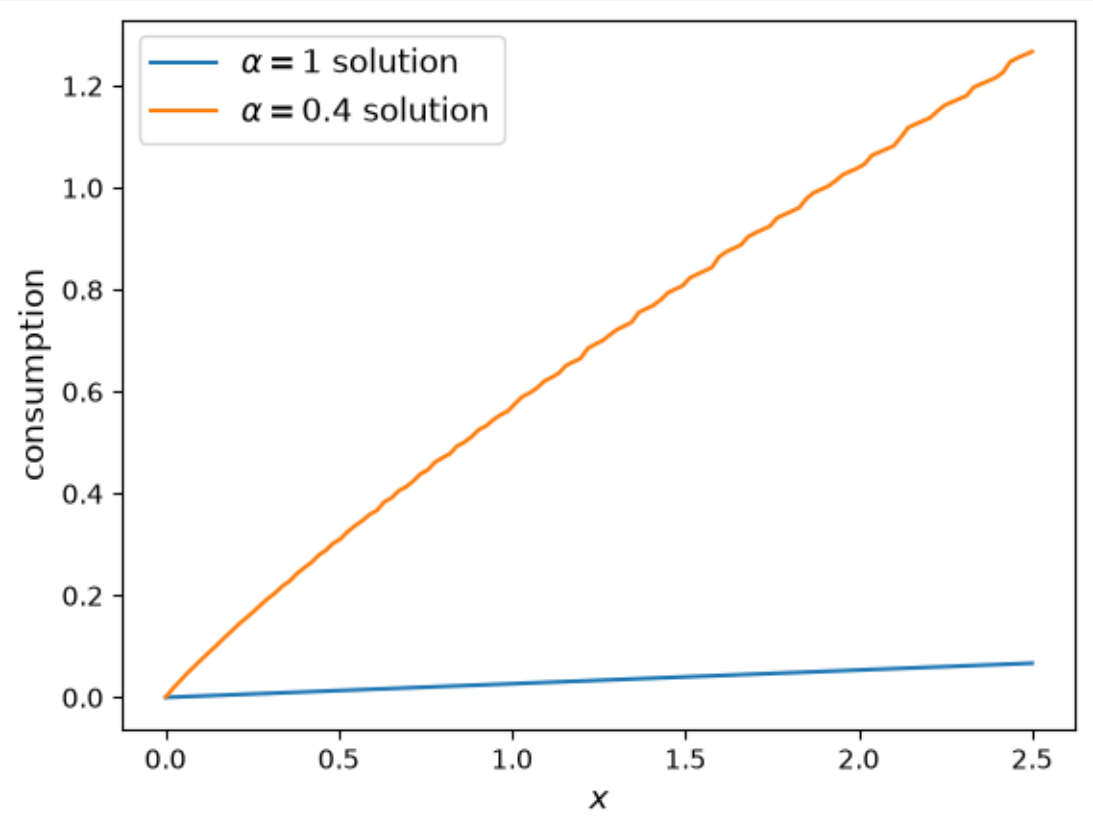
fig, ax = plt.subplots()

ax.plot(baseline_model.x_grid, c_analytical, label=r'$\alpha=1$ solution')
ax.plot(model.x_grid, σ, label=fr'$\alpha={model.α}$ solution')

ax.set_ylabel('consumption', fontsize=12)
ax.set_xlabel('$x$', fontsize=12)

ax.legend(fontsize=12)

plt.show()
```



Consumption is higher when $\alpha < 1$ because, at least for large x, the return to savings is lower.

CHAPTER

SIXTYEIGHT

OPTIMAL SAVINGS III: STOCHASTIC RETURNS

Contents

- *Optimal Savings III: Stochastic Returns*
  - *Overview*
  - *The Model*
  - *Computation*
  - *Exercises*

68.1 Overview

In this lecture, we continue our study of optimal savings problems, building on *Optimal Savings I: Cake Eating* and *Optimal Savings II: Numerical Cake Eating*.

The key difference from the previous lectures is that wealth now evolves stochastically.

We can think of wealth as a harvest that regrows if we save some seeds.

Specifically, if we save and invest part of today's harvest x_t, it grows into next period's harvest x_{t+1} according to a stochastic production process.

The extensions in this lecture introduce several new elements:

- nonlinear returns to saving, through a production function, and
- stochastic returns, due to shocks to production.

Despite these additions, the model remains relatively tractable.

As a first pass, we will solve the model using dynamic programming and value function iteration (VFI).

Note

In later lectures we'll explore more efficient methods for this class of problems.

At the same time, VFI is foundational and globally convergent.

Hence we want to be sure we can use this method too.

More information on this savings problem can be found in

- [Ljungqvist and Sargent, 2018], Section 3.1
- EDTC, Chapter 1
- [Sundaram, 1996], Chapter 12

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.interpolate import interp1d
from scipy.optimize import minimize_scalar
from typing import NamedTuple, Callable
```

68.2 The Model

Here we described the new model and the optimization problem.

68.2.1 Setup

Consider an agent who owns an amount $x_t \in \mathbb{R}_+ := [0, \infty)$ of a consumption good at time t.

This output can either be consumed or saved and used for production.

Production is stochastic, in that it also depends on a shock ξ_{t+1} realized at the end of the current period.

Next period output is

$$x_{t+1} := f(s_t)\xi_{t+1}$$

where $f \colon \mathbb{R}_+ \to \mathbb{R}_+$ is the **production function** and

$$s_t = x_t - c_t \tag{68.1}$$

is **current savings**.

and all variables are required to be nonnegative.

In what follows,

- The sequence $\{\xi_t\}$ is assumed to be IID.
- The common distribution of each ξ_t will be denoted by ϕ.
- The production function f is assumed to be increasing and continuous.

68.2.2 Optimization

Taking x_0 as given, the agent wishes to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{68.2}$$

subject to

$$x_{t+1} = f(x_t - c_t)\xi_{t+1} \quad \text{and} \quad 0 \leq c_t \leq x_t \quad \text{for all } t \tag{68.3}$$

where

- u is a bounded, continuous and strictly increasing utility function and
- $\beta \in (0, 1)$ is a discount factor.

In summary, the agent's aim is to select a path $c_0, c_1, c_2, \ldots$ for consumption that is

1. nonnegative,
2. feasible,
3. optimal, in the sense that it maximizes (68.2) relative to all other feasible consumption sequences, and
4. **adapted**, in the sense that the current action c_t depends only on current and historical outcomes, not on future outcomes such as ξ_{t+1}.

In the present context

- x_t is called the **state** variable — it summarizes the "state of the world" at the start of each period.
- c_t is called the **control** variable — a value chosen by the agent each period after observing the state.

68.2.3 Optimal Policies

Let us look at **policy functions**, each one of which is a map σ from the current state x_t into a current action c_t.

Note

These kinds of policies are called Markov policies (or stationary Markov policies).

For this dynamic program, the optimal policy is always a Markov policy (see, e.g., DP1).

In essence, the current state x_t provides a sufficient statistic for the history in terms of making an optimal decision today.

In what follows, we will call σ a **feasible consumption policy** if it satisfies

$$0 \leq \sigma(x) \leq x \quad \text{for all} \quad x \in \mathbb{R}_+ \tag{68.4}$$

In other words, a feasible policy is a policy function that respects the resource constraint.

The set of all feasible consumption policies will be denoted by Σ.

Each $\sigma \in \Sigma$ determines a Markov dynamics for output $\{x_t\}$ via

$$x_{t+1} = f(x_t - \sigma(x_t))\xi_{t+1}, \quad x_0 \text{ given} \tag{68.5}$$

This is the time path for output when we choose and stick with the policy σ.

We insert this process into the objective function to get

$$\mathbb{E}\sum_{t=0}^{\infty} \beta^t u(c_t) = \mathbb{E}\sum_{t=0}^{\infty} \beta^t u(\sigma(x_t)) \tag{68.6}$$

This is the total expected present value of following policy σ forever, given initial income x_0.

The aim is to select a policy that makes this number as large as possible.

The next section covers these ideas more formally.

68.2.4 Optimality

The lifetime value v_σ associated with a given policy σ is the mapping defined by

$$v_\sigma(x) = \mathbb{E} \sum_{t=0}^{\infty} \beta^t u(\sigma(x_t)) \tag{68.7}$$

when $\{x_t\}$ is given by (68.5) with $x_0 = x$.

In other words, it is the lifetime value of following policy σ forever, starting at initial condition x.

The **value function** is then defined as

$$v^*(x) := \sup_{\sigma \in \Sigma} \; v_\sigma(x) \tag{68.8}$$

The value function gives the maximal value that can be obtained from state x, after considering all feasible policies.

A policy $\sigma \in \Sigma$ is called **optimal** if $v_\sigma(x) = v^*(x)$ for all $x \in \mathbb{R}_+$.

68.2.5 The Bellman Equation

The following equation is called the **Bellman equation** associated with this dynamic programming problem.

$$v(x) = \max_{0 \leq c \leq x} \left\{ u(c) + \beta \int v(f(x-c)z)\phi(dz) \right\} \qquad (x \in \mathbb{R}_+) \tag{68.9}$$

This is a *functional equation in* v, in the sense that a given v can either satisfy it or not satisfy it.

The term $\int v(f(x-c)z)\phi(dz)$ can be understood as the expected next period value when

- v is used to measure value
- the state is x
- consumption is set to c

As shown in DP1 and a range of other texts, the value function v^* satisfies the Bellman equation.

In other words, (68.9) holds when $v = v^*$.

The intuition is that maximal value from a given state can be obtained by optimally trading off

- current reward from a given action, vs
- expected discounted future value of the state resulting from that action

The Bellman equation is important because it

1. gives us more information about the value function and
2. suggests a way of computing the value function, which we discuss below.

68.2.6 Greedy Policies

The value function can be used to compute optimal policies.

Given a continuous function v on $\mathbb{R}_+$, we say that $\sigma \in \Sigma$ is v-**greedy** if

$$\sigma(x) \in \arg\max_{0 \leq c \leq x} \left\{ u(c) + \beta \int v(f(x-c)z)\phi(dz) \right\} \tag{68.10}$$

for every $x \in \mathbb{R}_+$.

In other words, $\sigma \in \Sigma$ is v-greedy if it optimally trades off current and future rewards when v is taken to be the value function.

In our setting, we have the following key result

Theorem 68.2.1

A feasible consumption policy is optimal if and only if it is v^*-greedy.

See, for example, Theorem 10.1.11 of EDTC.

Hence, once we have a good approximation to v^*, we can compute the (approximately) optimal policy by computing the corresponding greedy policy.

The advantage is that we are now solving a much lower dimensional optimization problem.

68.2.7 The Bellman Operator

How, then, should we compute the value function?

One way is to use the so-called **Bellman operator**.

(The term **operator** is usually reserved for functions that send functions into functions!)

The Bellman operator is denoted by T and defined by

$$Tv(x) := \max_{0 \leq c \leq x} \left\{ u(c) + \beta \int v(f(x-c)z)\phi(dz) \right\} \qquad (x \in \mathbb{R}_+) \tag{68.11}$$

In other words, T sends the function v into the new function Tv defined by (68.11).

By construction, the set of solutions to the Bellman equation (68.9) *exactly coincides with* the set of fixed points of T.

For example, if $Tv = v$, then, for any $x \geq 0$,

$$v(x) = Tv(x) = \max_{0 \leq c \leq x} \left\{ u(c) + \beta \int v(f(x-c)z)\phi(dz) \right\}$$

which says precisely that v is a solution to the Bellman equation.

It follows that v^* is a fixed point of T.

68.2.8 Review of Theoretical Results

One can also show that T is a contraction mapping on the set of continuous bounded functions on $\mathbb{R}_+$ under the supremum distance

$$\rho(g, h) = \sup_{x \geq 0} |g(x) - h(x)|$$

See EDTC, Lemma 10.1.18.

Hence, it has exactly one fixed point in this set, which we know is equal to the value function.

It follows that

- The value function v^* is bounded and continuous.
- Starting from any bounded and continuous v, the sequence $v, Tv, T^2v, \ldots$ generated by iteratively applying T converges uniformly to v^*.

This iterative method is called **value function iteration**.

We also know that a feasible policy is optimal if and only if it is v^*-greedy.

It's not too hard to show that a v^*-greedy policy exists.

Hence, at least one optimal policy exists.

Our problem now is how to compute it.

68.2.9 Unbounded Utility

The results stated above assume that u is bounded.

In practice economists often work with unbounded utility functions — and so will we.

In the unbounded setting, various optimality theories exist.

Nevertheless, their main conclusions are usually in line with those stated for the bounded case just above (as long as we drop the word "bounded").

Note

Consult the following references for more on the unbounded case:

- The lecture *The Income Fluctuation Problem V: Stochastic Returns on Assets*.
- Section 12.2 of EDTC.

68.3 Computation

Let's now look at computing the value function and the optimal policy.

Our implementation in this lecture will focus on clarity and flexibility.

(In subsequent lectures we will focus on efficiency and speed.)

We will use fitted value function iteration, which was already described in *Optimal Savings II: Numerical Cake Eating*.

68.3.1 Scalar Maximization

To maximize the right hand side of the Bellman equation (68.9), we are going to use the `minimize_scalar` routine from SciPy.

To keep the interface tidy, we will wrap `minimize_scalar` in an outer function as follows:

```
def maximize(g, upper_bound):
    """
    Maximize the function g over the interval [0, upper_bound].

    We use the fact that the maximizer of g on any interval is
    also the minimizer of -g.

    """

    objective = lambda x: -g(x)
```

(continues on next page)

(continued from previous page)

```
    bounds = (0, upper_bound)
    result = minimize_scalar(objective, bounds=bounds, method='bounded')
    maximizer, maximum = result.x, -result.fun
    return maximizer, maximum
```

68.3.2 Model

We will assume for now that ϕ is the distribution of $\xi := \exp(\mu + \nu\zeta)$ where

- ζ is standard normal,
- μ is a shock location parameter and
- ν is a shock scale parameter.

We will store the primitives of the model in a `NamedTuple`.

```
class Model(NamedTuple):
    u: Callable        # utility function
    f: Callable        # production function
    β: float           # discount factor
    μ: float           # shock location parameter
    ν: float           # shock scale parameter
    x_grid: np.ndarray # state grid
    shocks: np.ndarray # shock draws


def create_model(
        u: Callable,
        f: Callable,
        β: float = 0.96,
        μ: float = 0.0,
        ν: float = 0.1,
        grid_max: float = 4.0,
        grid_size: int = 120,
        shock_size: int = 250,
        seed: int = 1234
    ) -> Model:
    """
    Creates an instance of the optimal savings model.
    """
    # Set up grid
    x_grid = np.linspace(1e-4, grid_max, grid_size)

    # Store shocks (with a seed, so results are reproducible)
    rng = np.random.default_rng(seed)
    shocks = np.exp(μ + ν * rng.standard_normal(shock_size))

    return Model(u, f, β, μ, ν, x_grid, shocks)
```

We set up the right-hand side of the Bellman equation

$$B(x, c, v) := u(c) + \beta \int v(f(x - c)z)\phi(dz)$$

```
def B(
        x: float,               # State
        c: float,               # Action
        v_array: np.ndarray,    # Array representing a guess of the value fn
        model: Model            # An instance of Model containing parameters
    ):

    u, f, β, μ, v, x_grid, shocks = model
    v = interp1d(x_grid, v_array)

    return u(c) + β * np.mean(v(f(x - c) * shocks))
```

In the second last line we are using linear interpolation.

In the last line, the expectation in (68.11) is computed via Monte Carlo, using the approximation

$$\int v(f(x-c)z)\phi(dz) \approx \frac{1}{n}\sum_{i=1}^{n} v(f(x-c)\xi_i)$$

where $\{\xi_i\}_{i=1}^n$ are IID draws from ϕ.

Monte Carlo is not always the most efficient way to compute integrals numerically but it does have some theoretical advantages in the present setting.

(For example, it preserves the contraction mapping property of the Bellman operator — see, e.g., [Pál and Stachurski, 2013].)

68.3.3 The Bellman Operator

The next function implements the Bellman operator.

```
def T(v: np.ndarray, model: Model) -> tuple[np.ndarray, np.ndarray]:
    """
    The Bellman operator.  Updates the guess of the value function.

      * model is an instance of Model
      * v is an array representing a guess of the value function

    """
    x_grid = model.x_grid
    v_new = np.empty_like(v)

    for i in range(len(x_grid)):
        x = x_grid[i]
        _, v_max = maximize(lambda c: B(x, c, v, model), x)
        v_new[i] = v_max

    return v_new
```

Here's the function:

```
def get_greedy(
        v: np.ndarray,          # current guess of the value function
        model: Model            # instance of optimal savings model
    ):
    " Compute the v-greedy policy on x_grid."
```

(continues on next page)

(continued from previous page)

```
    σ = np.empty_like(v)

    for i, x in enumerate(model.x_grid):
        # Maximize RHS of Bellman equation at state x
        σ[i], _ = maximize(lambda c: B(x, c, v, model), x)

    return σ
```

68.3.4 An Example

Let's suppose now that

$$f(x-c) = (x-c)^\alpha \quad \text{and} \quad u(c) = \ln c$$

For this particular problem, an exact analytical solution is available (see [Ljungqvist and Sargent, 2018], section 3.1.2), with

$$v^*(x) = \frac{\ln(1-\alpha\beta)}{1-\beta} + \frac{(\mu + \alpha \ln(\alpha\beta))}{1-\alpha}\left[\frac{1}{1-\beta} - \frac{1}{1-\alpha\beta}\right] + \frac{1}{1-\alpha\beta}\ln x \tag{68.12}$$

and optimal consumption policy

$$\sigma^*(x) = (1-\alpha\beta)x$$

It is valuable to have these closed-form solutions because it lets us check whether our code works for this particular case.

In Python, the functions above can be expressed as:

```
def v_star(x, α, β, μ):
    """
    True value function
    """
    c1 = np.log(1 - α * β) / (1 - β)
    c2 = (μ + α * np.log(α * β)) / (1 - α)
    c3 = 1 / (1 - β)
    c4 = 1 / (1 - α * β)
    return c1 + c2 * (c3 - c4) + c4 * np.log(x)

def σ_star(x, α, β):
    """
    True optimal policy
    """
    return (1 - α * β) * x
```

Next let's create an instance of the model with the above primitives and assign it to the variable `model`.

```
α = 0.4
def fcd(s):
    return s**α

model = create_model(u=np.log, f=fcd)
```

Now let's see what happens when we apply our Bellman operator to the exact solution v^* in this case.

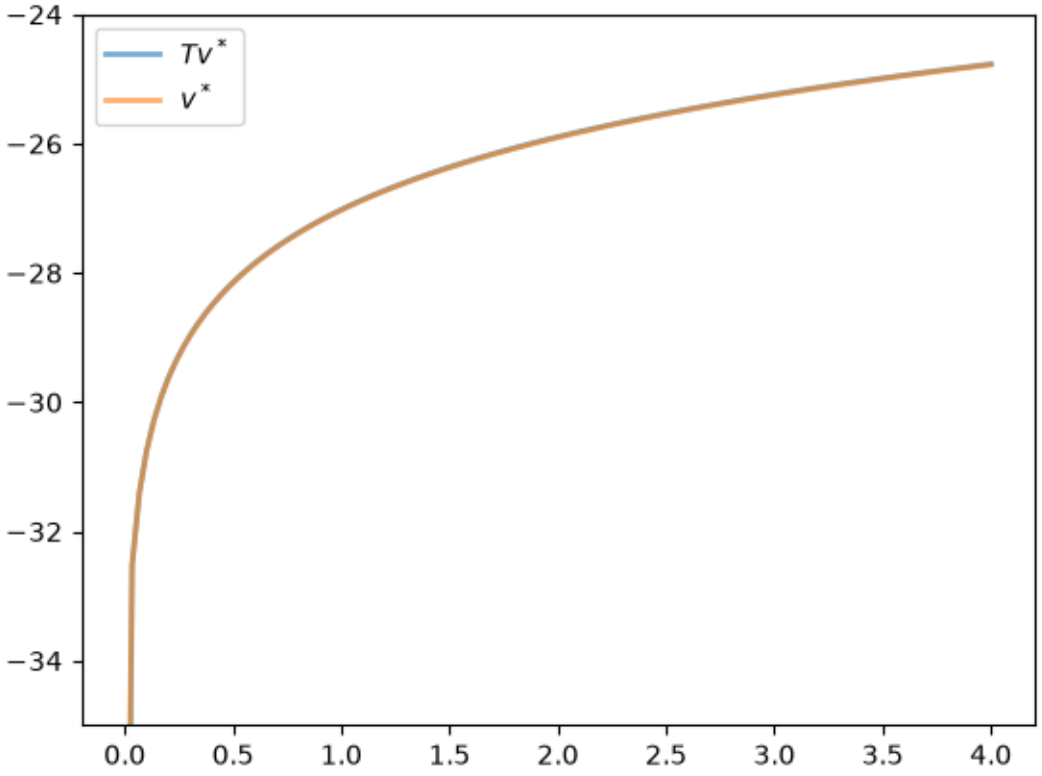
In theory, since v^* is a fixed point, the resulting function should again be v^*.

In practice, we expect some small numerical error.

```
x_grid = model.x_grid

v_init = v_star(x_grid, α, model.β, model.μ)     # Start at the solution
v = T(v_init, model)              # Apply T once

fig, ax = plt.subplots()
ax.set_ylim(-35, -24)
ax.plot(x_grid, v, lw=2, alpha=0.6, label='$Tv^*$')
ax.plot(x_grid, v_init, lw=2, alpha=0.6, label='$v^*$')
ax.legend()
plt.show()
```



The two functions are essentially indistinguishable, so we are off to a good start.

Now let's have a look at iterating with the Bellman operator, starting from an arbitrary initial condition.

The initial condition we'll start with is, somewhat arbitrarily, $v(x) = 5 \ln(x)$.

```
v = 5 * np.log(x_grid)  # An initial condition
n = 35

fig, ax = plt.subplots()

ax.plot(x_grid, v, color=plt.cm.jet(0),
        lw=2, alpha=0.6, label='Initial condition')

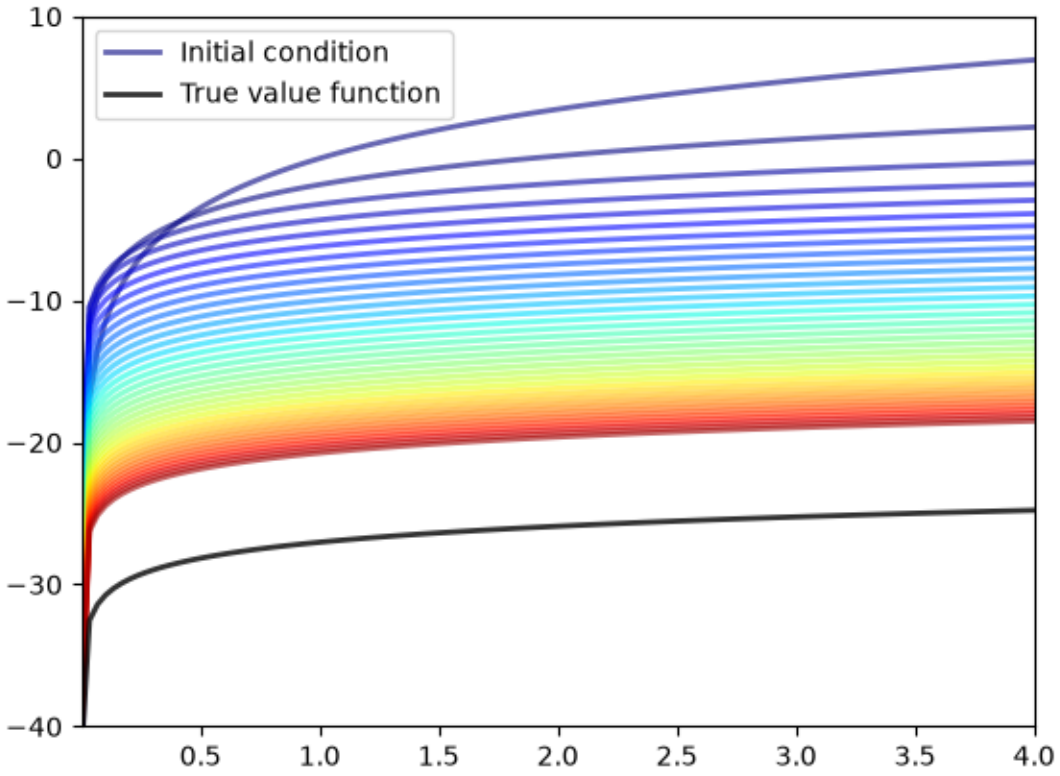
for i in range(n):
    v = T(v, model)  # Apply the Bellman operator
    ax.plot(x_grid, v, color=plt.cm.jet(i / n), lw=2, alpha=0.6)
```

(continues on next page)

(continued from previous page)

```
ax.plot(x_grid, v_star(x_grid, α, model.β, model.μ), 'k-', lw=2,
        alpha=0.8, label='True value function')

ax.legend()
ax.set(ylim=(-40, 10), xlim=(np.min(x_grid), np.max(x_grid)))
plt.show()
```



The figure shows

1. the first 36 functions generated by the fitted value function iteration algorithm, with hotter colors given to higher iterates
2. the true value function v^* drawn in black

The sequence of iterates converges towards v^*.

We are clearly getting closer.

68.3.5 Iterating to Convergence

We can write a function that iterates until the difference is below a particular tolerance level.

```
def solve_model(
        model: Model,           # instance of optimal savings model
        tol: float = 1e-4,      # convergence tolerance
        max_iter: int = 1000,   # maximum iterations
        verbose: bool = True,   # print iteration info
        print_skip: int = 25    # iterations between prints
```

(continues on next page)

(continued from previous page)

```
    ):
    " Solve by value function iteration. "

    v = model.u(model.x_grid)  # Initial condition
    i = 0
    error = tol + 1

    while i < max_iter and error > tol:
        v_new = T(v, model)
        error = np.max(np.abs(v - v_new))
        i += 1
        if verbose and i % print_skip == 0:
            print(f"Error at iteration {i} is {error}.")
        v = v_new

    if error > tol:
        print("Failed to converge!")
    elif verbose:
        print(f"\nConverged in {i} iterations.")

    v_greedy = get_greedy(v_new, model)
    return v_greedy, v_new
```

Let's use this function to compute an approximate solution at the defaults.

```
v_greedy, v_solution = solve_model(model)
```

```
Error at iteration 25 is 0.4089748660600776.
Error at iteration 50 is 0.14739319864069955.
Error at iteration 75 is 0.05312002487731604.
Error at iteration 100 is 0.019144282565221715.
Error at iteration 125 is 0.006899536583134136.
Error at iteration 150 is 0.0024865703324401522.
Error at iteration 175 is 0.0008961517840404554.
Error at iteration 200 is 0.0003229701607807556.
Error at iteration 225 is 0.0001163973856428413.

Converged in 229 iterations.
```

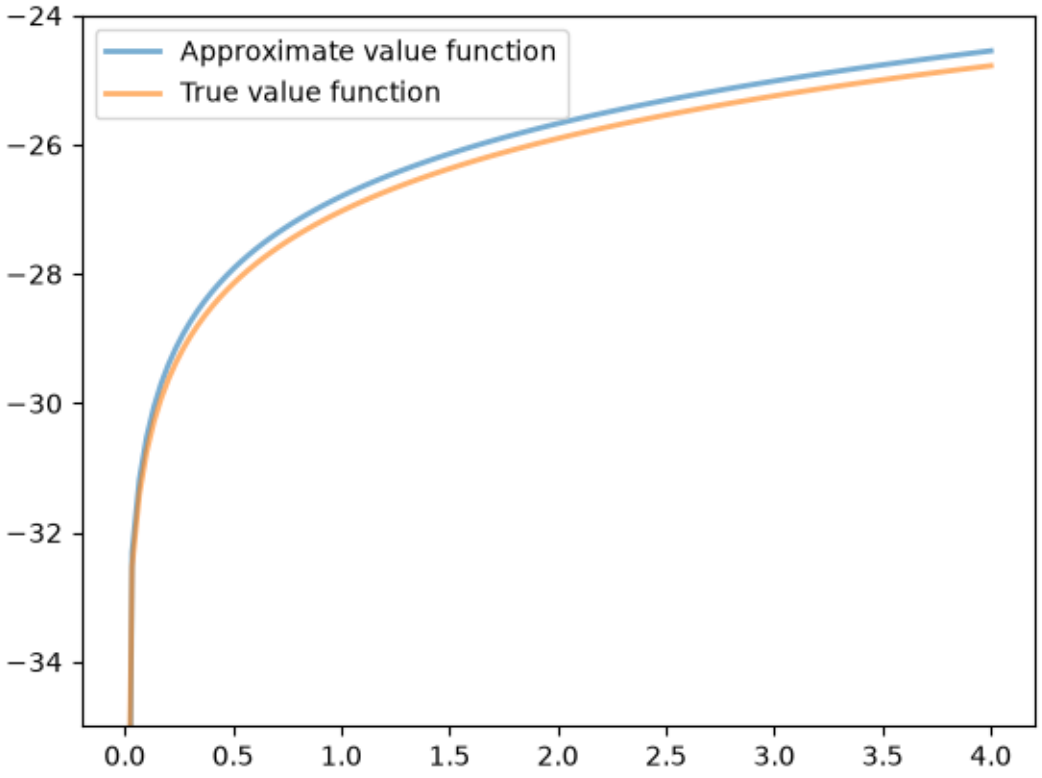
Now we check our result by plotting it against the true value:

```
fig, ax = plt.subplots()

ax.plot(x_grid, v_solution, lw=2, alpha=0.6,
        label='Approximate value function')

ax.plot(x_grid, v_star(x_grid, α, model.β, model.μ), lw=2,
        alpha=0.6, label='True value function')

ax.legend()
ax.set_ylim(-35, -24)
plt.show()
```



The figure shows that we are pretty much on the money.

68.3.6 The Policy Function

The policy `v_greedy` computed above corresponds to an approximate optimal policy.

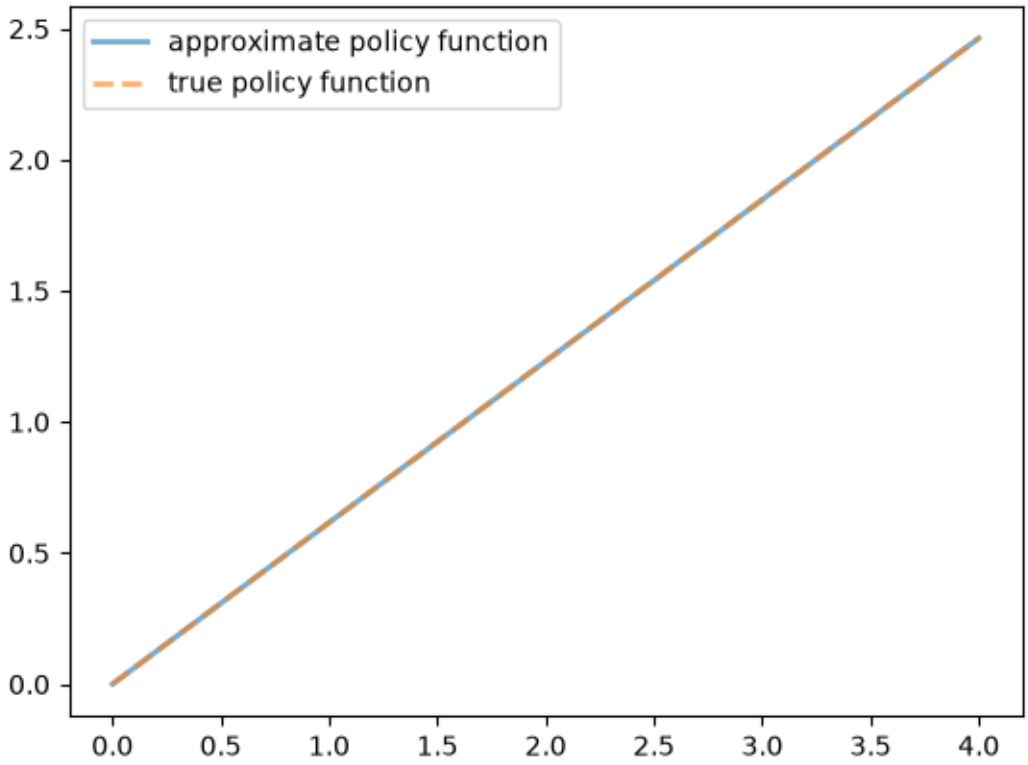
The next figure compares it to the exact solution, which, as mentioned above, is $\sigma(x) = (1 - \alpha\beta)x$

```
fig, ax = plt.subplots()

ax.plot(x_grid, v_greedy, lw=2,
        alpha=0.6, label='approximate policy function')

ax.plot(x_grid, σ_star(x_grid, α, model.β), '--',
        lw=2, alpha=0.6, label='true policy function')

ax.legend()
plt.show()
```



The figure shows that we've done a good job in this instance of approximating the true policy.

68.4 Exercises

Exercise 68.4.1

A common choice for utility function in this kind of work is the CRRA specification

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

Maintaining the other defaults, including the Cobb-Douglas production function, solve the optimal savings model with this utility specification.

Setting $\gamma = 1.5$, compute and plot an estimate of the optimal policy.

Solution

Here we set up the model.

```
γ = 1.5    # Preference parameter

def u_crra(c):
    return (c**(1 - γ) - 1) / (1 - γ)

model = create_model(u=u_crra, f=fcd)
```

Now let's run it, with a timer.

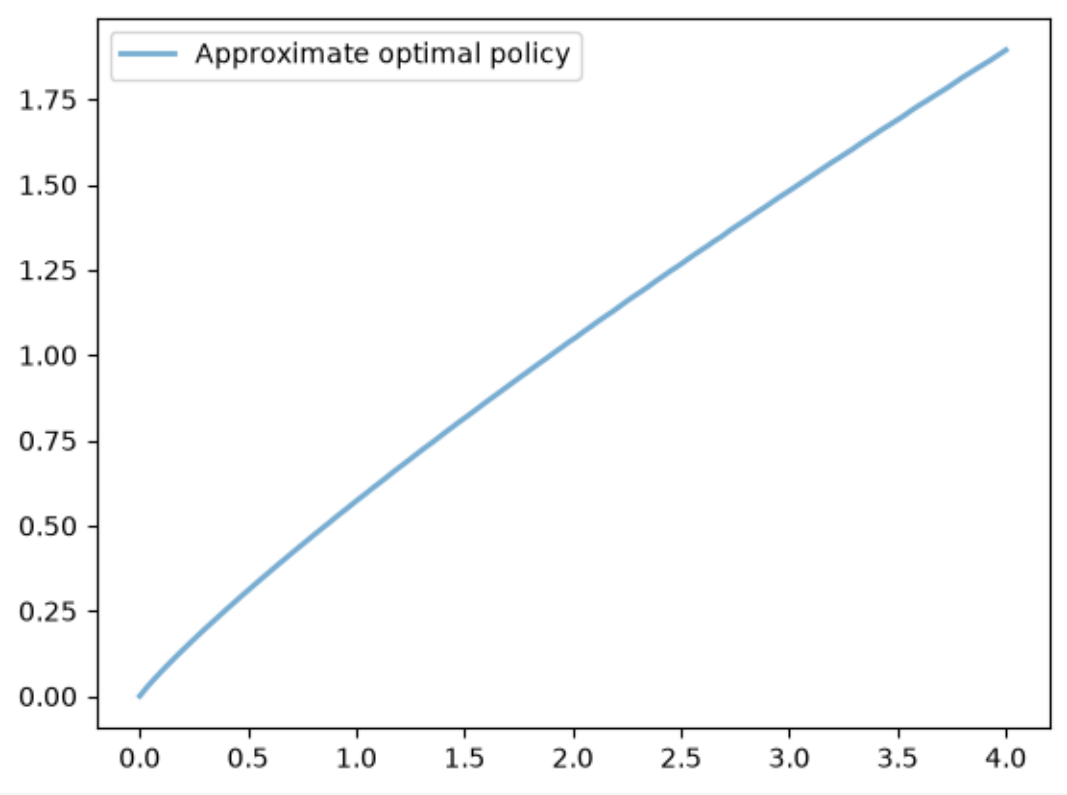
```
%%time
v_greedy, v_solution = solve_model(model)
```

```
Error at iteration 25 is 0.5516510200263269.
Error at iteration 50 is 0.19881274768354373.
Error at iteration 75 is 0.07165146153470658.
Error at iteration 100 is 0.025822951495229063.
Error at iteration 125 is 0.009306506938493442.
Error at iteration 150 is 0.003354034546092066.
Error at iteration 175 is 0.001208783038691763.
Error at iteration 200 is 0.0004356414385711105.
Error at iteration 225 is 0.00015700374421356855.

Converged in 237 iterations.
CPU times: user 26.6 s, sys: 7.96 ms, total: 26.6 s
Wall time: 26.6 s
```

Let's plot the policy function just to see what it looks like:

```
fig, ax = plt.subplots()
ax.plot(x_grid, v_greedy, lw=2,
        alpha=0.6, label='Approximate optimal policy')
ax.legend()
plt.show()
```



CHAPTER

SIXTYNINE

OPTIMAL SAVINGS IV: TIME ITERATION

Contents

- *Optimal Savings IV: Time Iteration*
  - *Overview*
  - *The Euler Equation*
  - *Implementation*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

69.1 Overview

In this lecture, we introduce the core idea of **time iteration**: iterating on a guess of the optimal policy using the Euler equation.

This approach differs from the value function iteration we used in *Optimal Savings III: Stochastic Returns*, where we iterated on the value function itself.

Time iteration exploits the structure of the Euler equation to find the optimal policy directly, rather than computing the value function as an intermediate step.

The key advantage is computational efficiency: by working directly with the policy function, we can often solve problems faster than with value function iteration.

However, time iteration is not the most efficient Euler equation-based method available.

In *Optimal Savings V: The Endogenous Grid Method*, we'll introduce the **endogenous grid method** (EGM), which provides an even more efficient way to solve the problem.

For now, our goal is to understand the basic mechanics of time iteration and how it leverages the Euler equation.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

(continues on next page)

(continued from previous page)

```
from scipy.optimize import brentq
from typing import NamedTuple, Callable
```

69.2 The Euler Equation

Our first step is to derive the Euler equation, which is a generalization of the Euler equation we obtained in *Optimal Savings I: Cake Eating*.

We take the model set out in *Optimal Savings III: Stochastic Returns* and add the following assumptions:

1. u and f are continuously differentiable and strictly concave
2. $f(0) = 0$
3. $\lim_{c \to 0} u'(c) = \infty$ and $\lim_{c \to \infty} u'(c) = 0$
4. $\lim_{k \to 0} f'(k) = \infty$ and $\lim_{k \to \infty} f'(k) = 0$

The last two conditions are usually called **Inada conditions**.

Recall the Bellman equation

$$v(x) = \max_{0 \leq c \leq x} \left\{ u(c) + \beta \int v(f(x-c)z)\phi(dz) \right\} \quad \text{for all} \quad x \in \mathbb{R}_+ \tag{69.1}$$

Let v^* be the value function and let σ^* be the optimal consumption policy.

We know that σ^* is a v^*-greedy policy.

The conditions above imply that

- σ^* is the unique optimal policy for the optimal savings problem
- the optimal policy is continuous, strictly increasing and also **interior**, in the sense that $0 < \sigma^*(x) < x$ for all strictly positive x, and
- the value function is strictly concave and continuously differentiable, with

$$(v^*)'(x) = u'(\sigma^*(x)) := (u' \circ \sigma^*)(x) \tag{69.2}$$

The last result is called the **envelope condition** due to its relationship with the envelope theorem.

To see why (69.2) holds, write the Bellman equation in the equivalent form

$$v^*(x) = \max_{0 \leq k \leq x} \left\{ u(x-k) + \beta \int v^*(f(k)z)\phi(dz) \right\},$$

Differentiating with respect to x, and then evaluating at the optimum yields (69.2).

(Section 12.1 of EDTC contains full proofs of these results, and closely related discussions can be found in many other texts.)

Differentiability of the value function and interiority of the optimal policy imply that optimal consumption satisfies the first-order condition associated with (69.1), which is

$$u'(\sigma^*(x)) = \beta \int (v^*)'(f(x - \sigma^*(x))z) f'(x - \sigma^*(x)) z \phi(dz) \tag{69.3}$$

Combining (69.2) and the first-order condition (69.3) gives the **Euler equation**

$$(u' \circ \sigma^*)(x) = \beta \int (u' \circ \sigma^*)(f(x - \sigma^*(x))z) f'(x - \sigma^*(x)) z \phi(dz) \tag{69.4}$$

We can think of the Euler equation as a functional equation

$$(u' \circ \sigma)(x) = \beta \int (u' \circ \sigma)(f(x - \sigma(x))z) f'(x - \sigma(x)) z \phi(dz) \tag{69.5}$$

over interior consumption policies σ, one solution of which is the optimal policy σ^*.

Our aim is to solve the functional equation (69.5) and hence obtain σ^*.

69.2.1 The Coleman-Reffett Operator

Recall the Bellman operator

$$Tv(x) := \max_{0 \leq c \leq x} \left\{ u(c) + \beta \int v(f(x - c)z) \phi(dz) \right\} \tag{69.6}$$

Just as we introduced the Bellman operator to solve the Bellman equation, we will now introduce an operator over policies to help us solve the Euler equation.

This operator K will act on the set of all $\sigma \in \Sigma$ that are continuous, strictly increasing and interior.

Henceforth we denote this set of policies by $\mathcal{P}$

1. The operator K takes as its argument a $\sigma \in \mathcal{P}$ and
2. returns a new function $K\sigma$, where $K\sigma(x)$ is the $c \in (0, x)$ that solves

$$u'(c) = \beta \int (u' \circ \sigma)(f(x - c)z) f'(x - c) z \phi(dz) \tag{69.7}$$

We call this operator the **Coleman-Reffett operator** to acknowledge the work of [Coleman, 1990] and [Reffett, 1996].

In essence, $K\sigma$ is the consumption policy that the Euler equation tells you to choose today when your future consumption policy is σ.

The important thing to note about K is that, by construction, its fixed points coincide with solutions to the functional equation (69.5).

In particular, the optimal policy σ^* is a fixed point.

Indeed, for fixed x, the value $K\sigma^*(x)$ is the c that solves

$$u'(c) = \beta \int (u' \circ \sigma^*)(f(x - c)z) f'(x - c) z \phi(dz)$$

In view of the Euler equation, this is exactly $\sigma^*(x)$.

69.2.2 Is the Coleman-Reffett Operator Well Defined?

In particular, is there always a unique $c \in (0, x)$ that solves (69.7)?

The answer is yes, under our assumptions.

For any $\sigma \in \mathcal{P}$, the right side of (69.7)

- is continuous and strictly increasing in c on $(0, x)$
- diverges to $+\infty$ as $c \uparrow x$

The left side of (69.7)

- is continuous and strictly decreasing in c on $(0, x)$

- diverges to $+\infty$ as $c \downarrow 0$

Sketching these curves and using the information above will convince you that they cross exactly once as c ranges over $(0, x)$.

With a bit more analysis, one can show in addition that $K\sigma \in \mathcal{P}$ whenever $\sigma \in \mathcal{P}$.

69.2.3 Comparison with VFI (Theory)

It is possible to prove that there is a tight relationship between iterates of K and iterates of the Bellman operator.

Mathematically, T and K are **topologically conjugate** under a translation that involves differentiation in one direction and integration in the other.

This conjugacy implies that if iterates of one operator converge then so do iterates of the other, and vice versa.

Moreover, there is a sense in which they converge *at the same rate*, at least in theory.

However, it turns out that the operator K is more stable *numerically* and hence more efficient in the applications we consider.

This is because

- K exploits additional structure because it uses first-order conditions, and
- policies near the optimal policy have less curvature and hence are easier to approximate than value functions near the optimal value function.

Examples are given below.

69.3 Implementation

Let's turn to implementation.

Note

In this lecture we mainly focus on the algorithm, favoring clarity over efficiency in the code.

In later lectures we will optimize both the algorithm and the code.

As in *Optimal Savings III: Stochastic Returns*, we assume that

- $u(c) = \ln c$
- $f(x - c) = (x - c)^\alpha$
- ϕ is the distribution of $\xi := \exp(\mu + \nu\zeta)$ when ζ is standard normal

This allows us to compare our results to the analytical solutions we obtained in that lecture:

```
def v_star(x, α, β, μ):
    """
    True value function
    """
    c1 = np.log(1 - α * β) / (1 - β)
    c2 = (μ + α * np.log(α * β)) / (1 - α)
    c3 = 1 / (1 - β)
    c4 = 1 / (1 - α * β)
```

(continues on next page)

(continued from previous page)

```
    return c1 + c2 * (c3 - c4) + c4 * np.log(x)

def σ_star(x, α, β):
    """
    True optimal policy
    """
    return (1 - α * β) * x
```

As discussed above, our plan is to solve the model using time iteration, which means iterating with the operator K.

For this we need access to the functions u' and f, f'.

We use the same `Model` structure from *Optimal Savings III: Stochastic Returns*.

```
class Model(NamedTuple):
    u: Callable        # utility function
    f: Callable        # production function
    β: float           # discount factor
    μ: float           # shock location parameter
    ν: float           # shock scale parameter
    grid: np.ndarray   # state grid
    shocks: np.ndarray # shock draws
    α: float = 0.4     # production function parameter
    u_prime: Callable = None        # derivative of utility
    f_prime: Callable = None        # derivative of production


def create_model(
        u: Callable,
        f: Callable,
        β: float = 0.96,
        μ: float = 0.0,
        ν: float = 0.1,
        grid_max: float = 4.0,
        grid_size: int = 120,
        shock_size: int = 250,
        seed: int = 1234,
        α: float = 0.4,
        u_prime: Callable = None,
        f_prime: Callable = None
    ) -> Model:
    """
    Creates an instance of the optimal savings model.
    """
    # Set up grid
    grid = np.linspace(1e-4, grid_max, grid_size)

    # Store shocks (with a seed, so results are reproducible)
    rng = np.random.default_rng(seed)
    shocks = np.exp(μ + ν * rng.standard_normal(shock_size))

    return Model(u, f, β, μ, ν, grid, shocks, α, u_prime, f_prime)
```

Now we implement a method called `euler_diff`, which returns

$$u'(c) - \beta \int (u' \circ \sigma)(f(x-c)z) f'(x-c) z \phi(dz) \tag{69.8}$$

```
def euler_diff(c: float, σ: np.ndarray, x: float, model: Model) -> float:
    """
    Set up a function such that the root with respect to c,
    given x and σ, is equal to Kσ(x).

    """

    # Unpack
    u, f, β, μ, ν, grid, shocks, α, u_prime, f_prime = model

    # Turn σ into a function via interpolation
    σ_func = lambda x: np.interp(x, grid, σ)

    # Now set up the function we need to find the root of.
    vals = u_prime(σ_func(f(x - c, α) * shocks)) * f_prime(x - c, α) * shocks
    return u_prime(c) - β * np.mean(vals)
```

The function `euler_diff` evaluates integrals by Monte Carlo and approximates functions using linear interpolation.

We will use a root-finding algorithm to solve (69.8) for c given state x and σ, the current guess of the policy.

Here's the operator K, that implements the root-finding step.

```
def K(σ: np.ndarray, model: Model) -> np.ndarray:
    """
    The Coleman-Reffett operator

    """

    # Unpack
    u, f, β, μ, ν, grid, shocks, α, u_prime, f_prime = model

    σ_new = np.empty_like(σ)
    for i, x in enumerate(grid):
        # Solve for optimal c at x
        c_star = brentq(euler_diff, 1e-10, x-1e-10, args=(σ, x, model))
        σ_new[i] = c_star

    return σ_new
```

69.3.1 Testing

Let's generate an instance and plot some iterates of K, starting from $\sigma(x) = x$.

```
# Define utility and production functions with derivatives
α = 0.4
u = lambda c: np.log(c)
u_prime = lambda c: 1 / c
f = lambda k, α: k**α
f_prime = lambda k, α: α * k**(α - 1)

model = create_model(u=u, f=f, α=α, u_prime=u_prime, f_prime=f_prime)
grid = model.grid

n = 15
σ = grid.copy()  # Set initial condition
```

(continues on next page)

(continued from previous page)

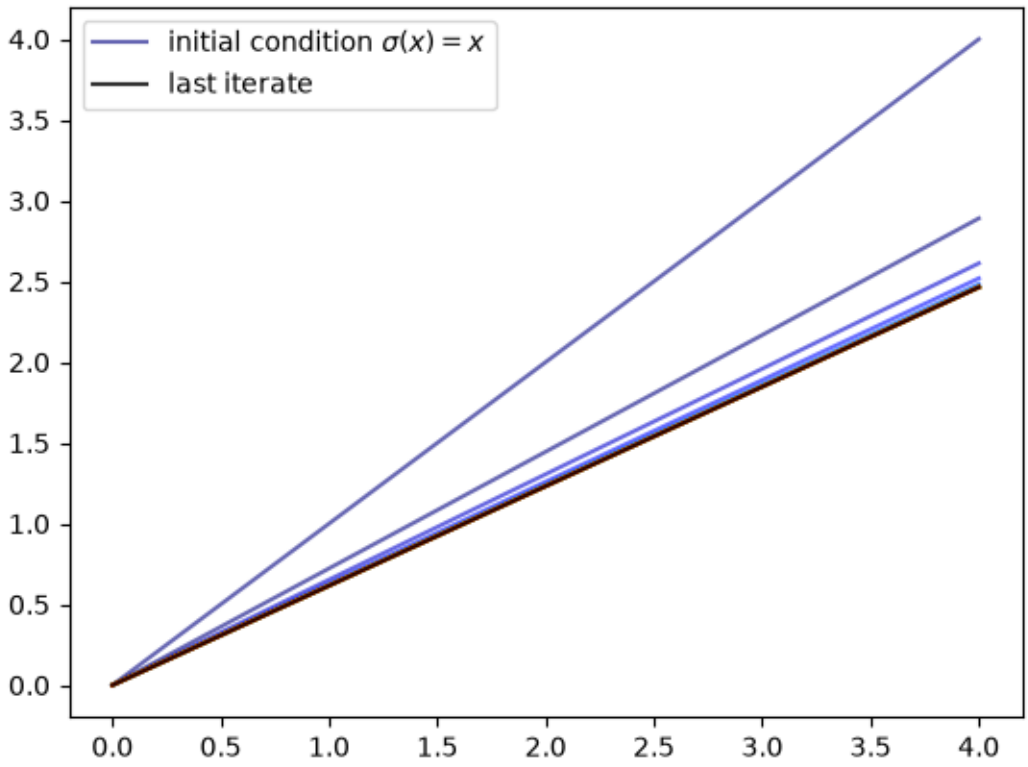
```
fig, ax = plt.subplots()
lb = r'initial condition $\sigma(x) = x$'
ax.plot(grid, σ, color=plt.cm.jet(0), alpha=0.6, label=lb)

for i in range(n):
    σ = K(σ, model)
    ax.plot(grid, σ, color=plt.cm.jet(i / n), alpha=0.6)

# Update one more time and plot the last iterate in black
σ = K(σ, model)
ax.plot(grid, σ, color='k', alpha=0.8, label='last iterate')

ax.legend()

plt.show()
```



We see that the iteration process converges quickly to a limit that resembles the solution we obtained in *Optimal Savings III: Stochastic Returns*.

Here is a function called `solve_model_time_iter` that takes an instance of `Model` and returns an approximation to the optimal policy, using time iteration.

```
def solve_model_time_iter(
        model: Model,
        σ_init: np.ndarray,
        tol: float = 1e-5,
        max_iter: int = 1000,
        verbose: bool = True
```

(continues on next page)

(continued from previous page)

```
    ) -> np.ndarray:
    """
    Solve the model using time iteration.

    """
    σ = σ_init
    error = tol + 1
    i = 0

    while error > tol and i < max_iter:
        σ_new = K(σ, model)
        error = np.max(np.abs(σ_new - σ))
        σ = σ_new
        i += 1
        if verbose:
            print(f"Iteration {i}, error = {error}")

    if i == max_iter:
        print("Warning: maximum iterations reached")

    return σ
```

Let's call it:

```
# Unpack
grid = model.grid

σ_init = np.copy(grid)
σ = solve_model_time_iter(model, σ_init)
```

```
Iteration 1, error = 1.1098265895953756
Iteration 2, error = 0.2782798920795728
Iteration 3, error = 0.09312729948558696
Iteration 4, error = 0.034020038271348696
Iteration 5, error = 0.012820752818721193
Iteration 6, error = 0.00488808156053766
Iteration 7, error = 0.0018718902256100733
Iteration 8, error = 0.0007180512309559184
Iteration 9, error = 0.0002756205293255043
Iteration 10, error = 0.0001058219018155171
Iteration 11, error = 4.063319516767194e-05
Iteration 12, error = 1.560279084245053e-05
Iteration 13, error = 5.991419175899182e-06
```

Here is a plot of the resulting policy, compared with the true policy:

```
# Unpack
grid, α, β = model.grid, model.α, model.β

fig, ax = plt.subplots()

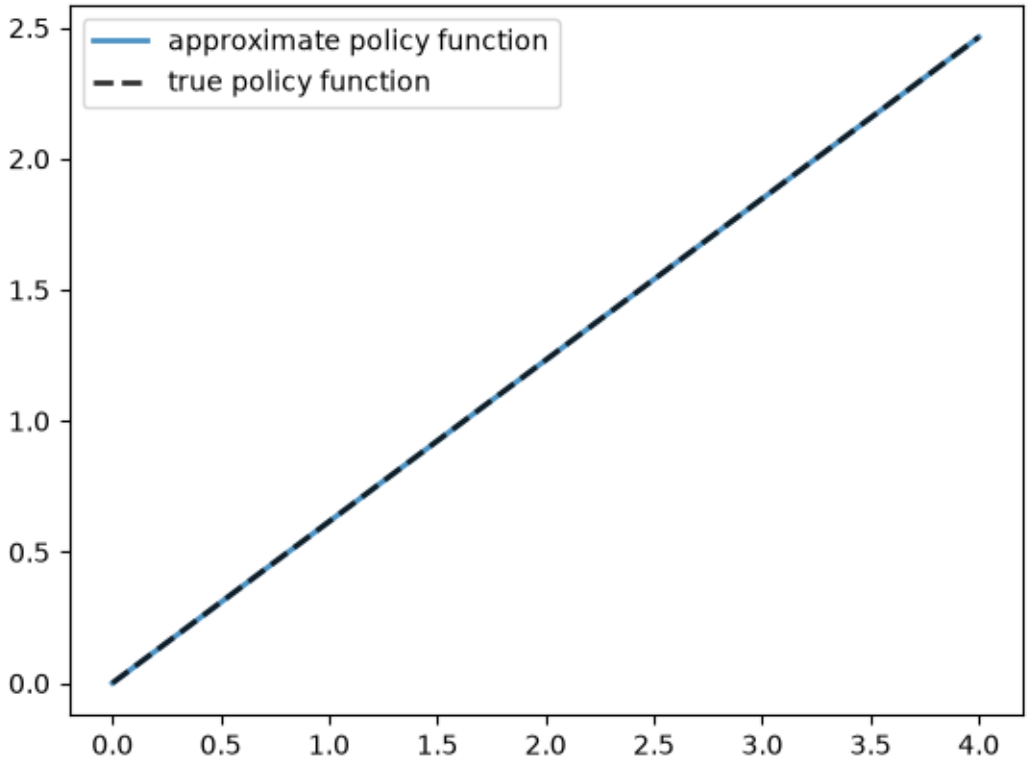
ax.plot(grid, σ, lw=2,
        alpha=0.8, label='approximate policy function')

ax.plot(grid, σ_star(grid, α, β), 'k--',
        lw=2, alpha=0.8, label='true policy function')
```

(continues on next page)

(continued from previous page)

```
ax.legend()
plt.show()
```



Again, the fit is excellent.

The maximal absolute deviation between the two policies is

```
# Unpack
grid, α, β = model.grid, model.α, model.β

np.max(np.abs(σ - σ_star(grid, α, β)))
```

```
np.float64(3.7348959640581825e-06)
```

Time iteration runs faster than value function iteration, as discussed in *Optimal Savings III: Stochastic Returns*.

This is because time iteration exploits differentiability and the first-order conditions, while value function iteration does not use this available structure.

At the same time, there is a variation of time iteration that runs even faster.

This is the endogenous grid method, which we will introduce in *Optimal Savings V: The Endogenous Grid Method*.

69.4 Exercises

Exercise 69.4.1

Solve the optimal savings problem with CRRA utility

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

Set `γ = 1.5`.

Compute and plot the optimal policy.

Solution

We define the CRRA utility function and its derivative.

```
γ = 1.5

def u_crra(c):
    return c**(1 - γ) / (1 - γ)

def u_prime_crra(c):
    return c**(-γ)

# Use the same production function as before
model_crra = create_model(u=u_crra, f=f, α=α,
                          u_prime=u_prime_crra, f_prime=f_prime)
```

Now we solve and plot the policy:

```
%%time
# Unpack
grid = model_crra.grid

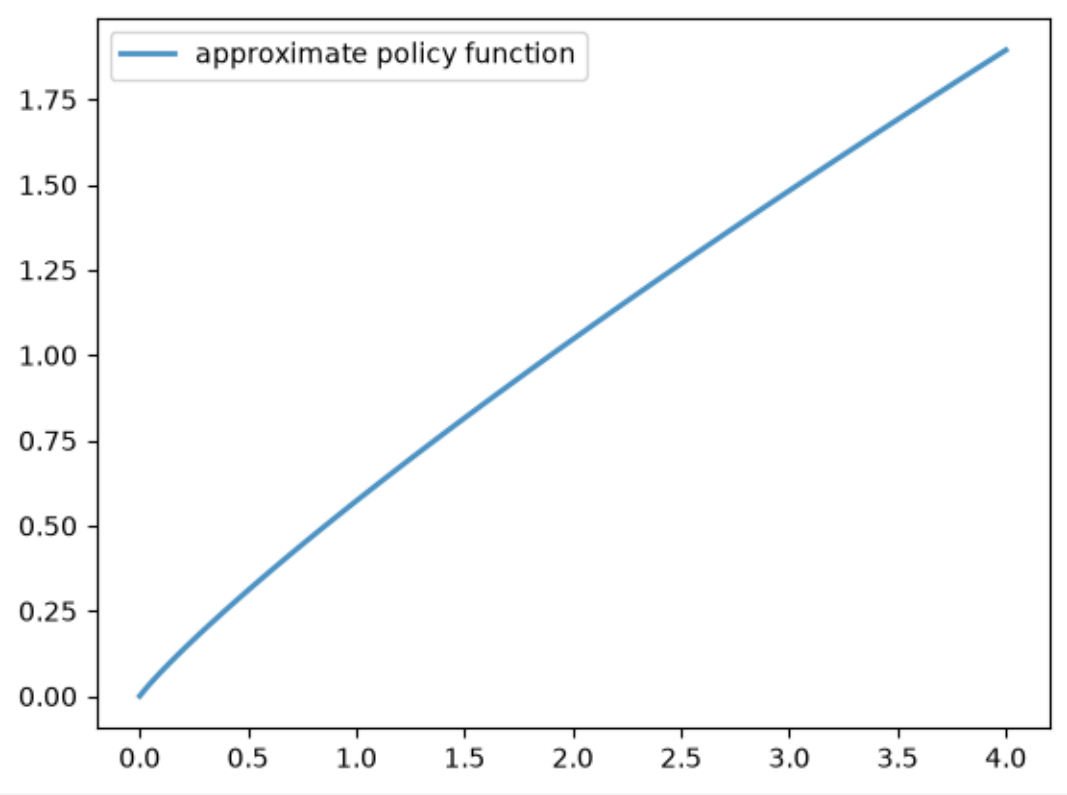
σ_init = np.copy(grid)
σ = solve_model_time_iter(model_crra, σ_init)


fig, ax = plt.subplots()

ax.plot(grid, σ, lw=2,
        alpha=0.8, label='approximate policy function')

ax.legend()
plt.show()
```

```
Iteration 1, error = 1.449604914521836
Iteration 2, error = 0.3966806986253908
Iteration 3, error = 0.14842065541477645
Iteration 4, error = 0.06191571110718308
Iteration 5, error = 0.027010882945130543
Iteration 6, error = 0.012016369353877288
Iteration 7, error = 0.005392555225094853
Iteration 8, error = 0.0024294216912608135
Iteration 9, error = 0.001096446547679708
Iteration 10, error = 0.0004952512989286895
Iteration 11, error = 0.00022378170290093813
Iteration 12, error = 0.00010113431534053419
Iteration 13, error = 4.570909652401589e-05
Iteration 14, error = 2.06595703444723e-05
Iteration 15, error = 9.337873995374224e-06
CPU times: user 654 ms, sys: 5.98 ms, total: 660 ms
```


approximate policy function
1.75
1.50
1.25
1.00
0.75
0.50
0.25
0.00
0.0
0.5
1.0
1.5
2.0
2.5
3.0
3.5
4.0

CHAPTER

SEVENTY

OPTIMAL SAVINGS V: THE ENDOGENOUS GRID METHOD

Contents

- *Optimal Savings V: The Endogenous Grid Method*
  - *Overview*
  - *Key Idea*
  - *Implementation*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

70.1 Overview

Previously, we solved the optimal savings problem using

1. *value function iteration*
2. *Euler equation based time iteration*

We found time iteration to be significantly more accurate and efficient.

In this lecture, we'll look at a clever twist on time iteration called the **endogenous grid method** (EGM).

EGM is a numerical method for implementing policy iteration invented by Chris Carroll.

The original reference is [Carroll, 2006].

For now we will focus on a clean and simple implementation of EGM that stays close to the underlying mathematics.

Then, in *Optimal Savings VI: EGM with JAX*, we will construct a fully vectorized and parallelized version of EGM based on JAX.

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
```

70.2 Key Idea

First we remind ourselves of the theory and then we turn to numerical methods.

70.2.1 Theory

We work with the model set out in *Optimal Savings IV: Time Iteration*, following the same terminology and notation.

As we saw, the Coleman-Reffett operator is a nonlinear operator K engineered so that the optimal policy σ^* is a fixed point of K.

It takes as its argument a continuous strictly increasing consumption policy $\sigma \in \Sigma$.

It returns a new function $K\sigma$, where $(K\sigma)(x)$ is the $c \in (0, \infty)$ that solves

$$u'(c) = \beta \int (u' \circ \sigma)(f(x-c)z) f'(x-c) z \phi(dz) \tag{70.1}$$

70.2.2 Exogenous Grid

As discussed in *Optimal Savings IV: Time Iteration*, to implement the method on a computer, we represent a policy function by a set of values on a finite grid.

The function itself is reconstructed from this representation when necessary, using interpolation or some other method.

Our previous strategy in *Optimal Savings IV: Time Iteration* for obtaining a finite representation of an updated consumption policy was to

- fix a grid of income points $\{x_i\}$
- calculate the consumption value c_i corresponding to each x_i using (70.1) and a root-finding routine

Each c_i is then interpreted as the value of the function $K\sigma$ at x_i.

Thus, with the pairs $\{(x_i, c_i)\}$ in hand, we can reconstruct $K\sigma$ via approximation.

Iteration then continues…

70.2.3 Endogenous Grid

The method discussed above requires a root-finding routine to find the c_i corresponding to a given income value x_i.

Root-finding is costly because it typically involves a significant number of function evaluations.

As pointed out by Carroll [Carroll, 2006], we can avoid this step if x_i is chosen endogenously.

The only assumption required is that u' is invertible on $(0, \infty)$.

Let $(u')^{-1}$ be the inverse function of u'.

The idea is this:

- First, we fix an *exogenous* grid $\{s_i\}$ for savings ($s = x - c$).
- Then we obtain c_i via

$$c_i = (u')^{-1} \left\{ \beta \int (u' \circ \sigma)(f(s_i)z) \, f'(s_i) \, z \, \phi(dz) \right\} \tag{70.2}$$

- Finally, for each c_i we set $x_i = c_i + s_i$.

Importantly, each (x_i, c_i) pair constructed in this manner satisfies (70.1).

With the points $\{x_i, c_i\}$ in hand, we can reconstruct $K\sigma$ via approximation as before.

The name EGM comes from the fact that the grid $\{x_i\}$ is determined **endogenously**.

70.3 Implementation

As in *Optimal Savings IV: Time Iteration*, we will start with a simple setting where

- $u(c) = \ln c$,
- the function f has a Cobb-Douglas specification, and
- the shocks are lognormal.

This will allow us to make comparisons with the analytical solutions.

```
def v_star(x, α, β, μ):
    """
    True value function
    """
    c1 = np.log(1 - α * β) / (1 - β)
    c2 = (μ + α * np.log(α * β)) / (1 - α)
    c3 = 1 / (1 - β)
    c4 = 1 / (1 - α * β)
    return c1 + c2 * (c3 - c4) + c4 * np.log(x)

def σ_star(x, α, β):
    """
    True optimal policy
    """
    return (1 - α * β) * x
```

We reuse the `Model` structure from *Optimal Savings IV: Time Iteration*.

```
from typing import NamedTuple, Callable

class Model(NamedTuple):
    u: Callable           # utility function
    f: Callable           # production function
    β: float              # discount factor
    μ: float              # shock location parameter
    ν: float              # shock scale parameter
    s_grid: np.ndarray    # exogenous savings grid
    shocks: np.ndarray    # shock draws
    α: float              # production function parameter
    u_prime: Callable     # derivative of utility
    f_prime: Callable     # derivative of production
    u_prime_inv: Callable # inverse of u_prime


def create_model(
        u: Callable,
        f: Callable,
        β: float = 0.96,
        μ: float = 0.0,
        ν: float = 0.1,
```

(continues on next page)

(continued from previous page)

```
        grid_max: float = 4.0,
        grid_size: int = 120,
        shock_size: int = 250,
        seed: int = 1234,
        α: float = 0.4,
        u_prime: Callable = None,
        f_prime: Callable = None,
        u_prime_inv: Callable = None
    ) -> Model:
    """
    Creates an instance of the optimal savings model.
    """
    # Set up exogenous savings grid
    s_grid = np.linspace(1e-4, grid_max, grid_size)

    # Store shocks (with a seed, so results are reproducible)
    rng = np.random.default_rng(seed)
    shocks = np.exp(μ + ν * rng.standard_normal(shock_size))

    return Model(
        u, f, β, μ, ν, s_grid, shocks, α, u_prime, f_prime, u_prime_inv
    )
```

70.3.1 The Operator

Here's an implementation of K using EGM as described above.

```
def K(
        c_in: np.ndarray,    # Consumption values on the endogenous grid
        x_in: np.ndarray,    # Current endogenous grid
        model: Model         # Model specification
    ):
    """
    An implementation of the Coleman-Reffett operator using EGM.

    """

    # Simplify names
    u, f, β, μ, ν, s_grid, shocks, α, u_prime, f_prime, u_prime_inv = model

    # Linear interpolation of policy on the endogenous grid
    σ = lambda x: np.interp(x, x_in, c_in)

    # Allocate memory for new consumption array
    c_out = np.empty_like(s_grid)

    for i, s in enumerate(s_grid):
        # Approximate marginal utility ∫ u'(σ(f(s, α)z)) f'(s, α) z ϕ(z)dz
        vals = u_prime(σ(f(s, α) * shocks)) * f_prime(s, α) * shocks
        mu = np.mean(vals)
        # Compute consumption
        c_out[i] = u_prime_inv(β * mu)

    # Determine corresponding endogenous grid
    x_out = s_grid + c_out        # x_i = s_i + c_i
```

(continues on next page)

(continued from previous page)

```
    return c_out, x_out
```

Note the lack of any root-finding algorithm.

Note

The routine is still not particularly fast because we are using pure Python loops.

But in the next lecture (*Optimal Savings VI: EGM with JAX*) we will use a fully vectorized and efficient solution.

70.3.2 Testing

First we create an instance.

```
# Define utility and production functions with derivatives
u = lambda c: np.log(c)
u_prime = lambda c: 1 / c
u_prime_inv = lambda x: 1 / x
f = lambda k, α: k**α
f_prime = lambda k, α: α * k**(α - 1)

model = create_model(u=u, f=f, u_prime=u_prime,
                     f_prime=f_prime, u_prime_inv=u_prime_inv)
s_grid = model.s_grid
```

Here's our solver routine:

```
def solve_model_time_iter(
        model: Model,              # Model details
        c_init: np.ndarray,        # initial guess of consumption on EG
        x_init: np.ndarray,        # initial guess of endogenous grid
        tol: float = 1e-5,         # Error tolerance
        max_iter: int = 1000,      # Max number of iterations of K
        verbose: bool = True       # If true print output
    ):
    """
    Solve the model using time iteration with EGM.
    """
    c, x = c_init, x_init
    error = tol + 1
    i = 0

    while error > tol and i < max_iter:
        c_new, x_new = K(c, x, model)
        error = np.max(np.abs(c_new - c))
        c, x = c_new, x_new
        i += 1
        if verbose:
            print(f"Iteration {i}, error = {error}")

    if i == max_iter:
        print("Warning: maximum iterations reached")
```

(continues on next page)

(continued from previous page)

```
    return c, x
```

Let's call it:

```
c_init = np.copy(s_grid)
x_init = s_grid + c_init
c, x = solve_model_time_iter(model, c_init, x_init)
```

```
Iteration 1, error = 1.208333333333333
Iteration 2, error = 0.6834464555052788
Iteration 3, error = 0.3126351338414741
Iteration 4, error = 0.12905177785629984
Iteration 5, error = 0.05099394329538409
Iteration 6, error = 0.01980055511172729
Iteration 7, error = 0.007636088144566955
Iteration 8, error = 0.002937098891578671
Iteration 9, error = 0.0011285611196765188
Iteration 10, error = 0.00043347299641638415
Iteration 11, error = 0.00016646919532892213
Iteration 12, error = 6.392646635156041e-05
Iteration 13, error = 2.454810155327891e-05
Iteration 14, error = 9.426520909627811e-06
```

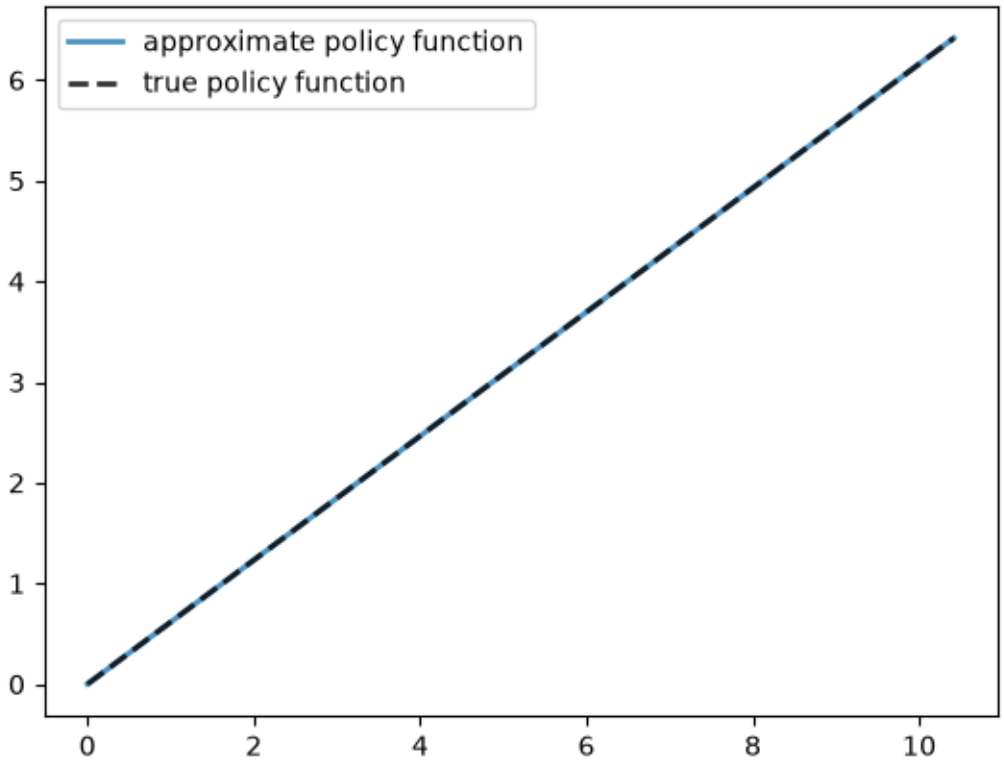
Here is a plot of the resulting policy, compared with the true policy:

```
fig, ax = plt.subplots()

ax.plot(x, c, lw=2,
        alpha=0.8, label='approximate policy function')

ax.plot(x, σ_star(x, model.α, model.β), 'k--',
        lw=2, alpha=0.8, label='true policy function')

ax.legend()
plt.show()
```



The maximal absolute deviation between the two policies is

```
np.max(np.abs(c - σ_star(x, model.α, model.β)))
```

```
np.float64(2.2564941266622895e-06)
```

Here's the execution time:

```
with qe.Timer():
    c, x = solve_model_time_iter(model, c_init, x_init, verbose=False)
```

```
0.0321 seconds elapsed
```

EGM is faster than time iteration because it avoids numerical root-finding.

Instead, we invert the marginal utility function directly, which is much more efficient.

In *Optimal Savings VI: EGM with JAX*, we will use a fully vectorized and efficient version of EGM that is also parallelized using JAX.

This provides an extremely fast way to solve the optimal consumption problem we have been studying for the last few lectures.

CHAPTER

SEVENTYONE

OPTIMAL SAVINGS VI: EGM WITH JAX

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Optimal Savings VI: EGM with JAX*
  - *Overview*
  - *Implementation*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

71.1 Overview

In this lecture, we'll implement the endogenous grid method (EGM) using JAX.

This lecture builds on *Optimal Savings V: The Endogenous Grid Method*, which introduced EGM using NumPy.

By converting to JAX, we can leverage fast linear algebra, hardware accelerators, and JIT compilation for improved performance.

We'll also use JAX's `vmap` function to fully vectorize the Coleman-Reffett operator.

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
import quantecon as qe
from typing import NamedTuple
```

71.2 Implementation

For details on the savings problem and the endogenous grid method (EGM), please see *Optimal Savings V: The Endogenous Grid Method*.

Here we focus on the JAX implementation of EGM.

We use the same setting as in *Optimal Savings V: The Endogenous Grid Method*:

- $u(c) = \ln c$,
- production is Cobb-Douglas, and
- the shocks are lognormal.

Here are the analytical solutions for comparison.

```
def v_star(x, α, β, μ):
    """
    True value function
    """
    c1 = jnp.log(1 - α * β) / (1 - β)
    c2 = (μ + α * jnp.log(α * β)) / (1 - α)
    c3 = 1 / (1 - β)
    c4 = 1 / (1 - α * β)
    return c1 + c2 * (c3 - c4) + c4 * jnp.log(x)

def σ_star(x, α, β):
    """
    True optimal policy
    """
    return (1 - α * β) * x
```

The `Model` class stores only the data (grids, shocks, and parameters).

Utility and production functions will be defined globally to work with JAX's JIT compiler.

```
class Model(NamedTuple):
    β: float              # discount factor
    μ: float              # shock location parameter
    s: float              # shock scale parameter
    s_grid: jnp.ndarray   # exogenous savings grid
    shocks: jnp.ndarray   # shock draws
    α: float              # production function parameter


def create_model(
        β: float = 0.96,
        μ: float = 0.0,
        s: float = 0.1,
        grid_max: float = 4.0,
```

(continues on next page)

(continued from previous page)

```
        grid_size: int = 120,
        shock_size: int = 250,
        seed: int = 1234,
        α: float = 0.4
    ) -> Model:
    """
    Creates an instance of the optimal savings model.
    """
    # Set up exogenous savings grid
    s_grid = jnp.linspace(1e-4, grid_max, grid_size)

    # Store shocks (with a seed, so results are reproducible)
    key = jax.random.key(seed)
    shocks = jnp.exp(μ + s * jax.random.normal(key, shape=(shock_size,)))

    return Model(β, μ, s, s_grid, shocks, α)
```

We define utility and production functions globally.

```
# Define utility and production functions with derivatives
u = lambda c: jnp.log(c)
u_prime = lambda c: 1 / c
u_prime_inv = lambda x: 1 / x
f = lambda k, α: k**α
f_prime = lambda k, α: α * k**(α - 1)
```

Here's the Coleman-Reffett operator using EGM.

The key JAX feature here is `vmap`, which vectorizes the computation over the grid points.

```
def K(
        c_in: jnp.ndarray,  # Consumption values on the endogenous grid
        x_in: jnp.ndarray,  # Current endogenous grid
        model: Model        # Model specification
    ):
    """
    The Coleman-Reffett operator using EGM

    """
    β, μ, s, s_grid, shocks, α = model
    σ = lambda x_val: jnp.interp(x_val, x_in, c_in)

    # Define function to compute consumption at a single grid point
    def compute_c(s):
        # Approximate marginal utility ∫ u'(σ(f(s, α)z)) f'(s, α) z ϕ(z)dz
        vals = u_prime(σ(f(s, α) * shocks)) * f_prime(s, α) * shocks
        mu = jnp.mean(vals)
        # Calculate consumption
        return u_prime_inv(β * mu)

    # Vectorize and calculate on all exogenous grid points
    compute_c_vectorized = jax.vmap(compute_c)
    c_out = compute_c_vectorized(s_grid)

    # Determine corresponding endogenous grid
    x_out = s_grid + c_out  # x_i = s_i + c_i
```

(continues on next page)

(continued from previous page)

```
    return c_out, x_out
```

Now we create a model instance.

```
model = create_model()
s_grid = model.s_grid
```

The solver uses JAX's `jax.lax.while_loop` for the iteration and is JIT-compiled for speed.

```
@jax.jit
def solve_model_time_iter(
        model: Model,
        c_init: jnp.ndarray,
        x_init: jnp.ndarray,
        tol: float = 1e-5,
        max_iter: int = 1000
    ):
    """
    Solve the model using time iteration with EGM.
    """

    def condition(loop_state):
        i, c, x, error = loop_state
        return (error > tol) & (i < max_iter)

    def body(loop_state):
        i, c, x, error = loop_state
        c_new, x_new = K(c, x, model)
        error = jnp.max(jnp.abs(c_new - c))
        return i + 1, c_new, x_new, error

    # Initialize loop state
    initial_state = (0, c_init, x_init, tol + 1)

    # Run the loop
    i, c, x, error = jax.lax.while_loop(condition, body, initial_state)

    return c, x
```

We solve the model starting from an initial guess.

```
c_init = jnp.copy(s_grid)
x_init = s_grid + c_init
c, x = solve_model_time_iter(model, c_init, x_init)
```

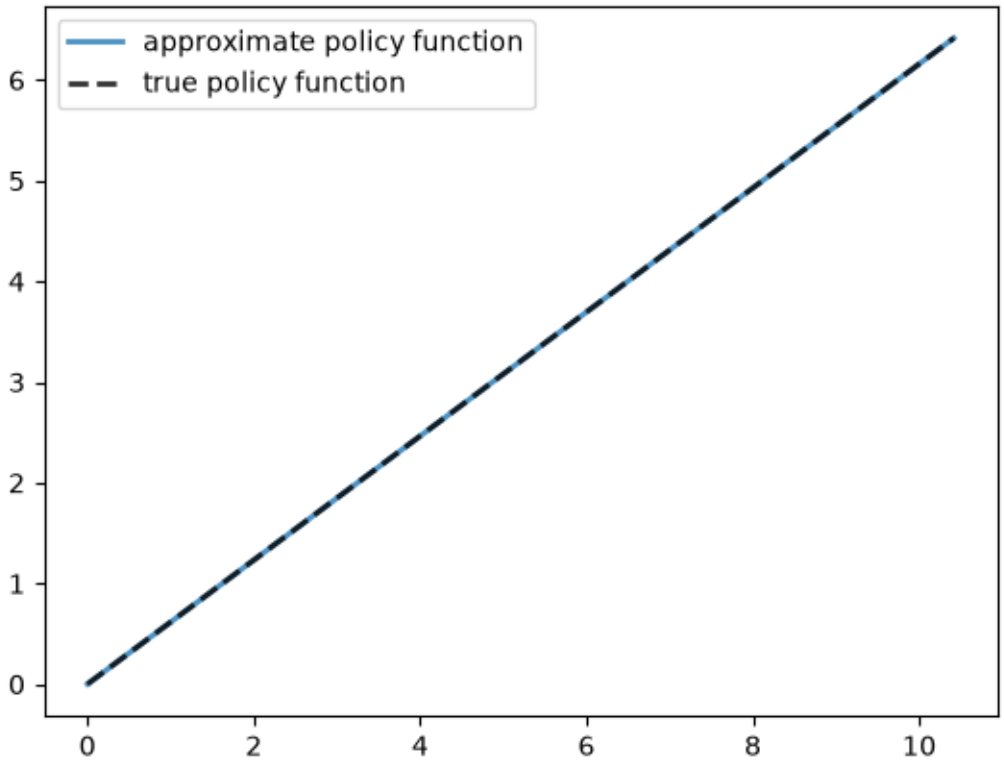
Let's plot the resulting policy against the analytical solution.

```
fig, ax = plt.subplots()

ax.plot(x, c, lw=2,
        alpha=0.8, label='approximate policy function')

ax.plot(x, σ_star(x, model.α, model.β), 'k--',
        lw=2, alpha=0.8, label='true policy function')

ax.legend()
plt.show()
```



The fit is very good.

```
max_dev = jnp.max(jnp.abs(c - σ_star(x, model.α, model.β)))
print(f"Maximum absolute deviation: {max_dev:.7}")
```

```
Maximum absolute deviation: 2.861023e-06
```

The JAX implementation is very fast thanks to JIT compilation and vectorization.

```
with qe.Timer(precision=8):
    c, x = solve_model_time_iter(model, c_init, x_init)
    jax.block_until_ready(c)
```

```
0.00276637 seconds elapsed
```

This speed comes from:

- JIT compilation of the entire solver
- Vectorization via `vmap` in the Coleman-Reffett operator
- Use of `jax.lax.while_loop` instead of a Python loop
- Efficient JAX array operations throughout

71.3 Exercises

Exercise 71.3.1

Solve the optimal savings problem with CRRA utility

$$u(c) = \frac{c^{1-\gamma} - 1}{1 - \gamma}$$

Compare the optimal policies for values of γ approaching 1 from above (e.g., 1.05, 1.1, 1.2).

Show that as $\gamma \to 1$, the optimal policy converges to the policy obtained with log utility ($\gamma = 1$).

Hint: Use values of γ close to 1 to ensure the endogenous grids have similar coverage and make visual comparison easier.

Solution

We need to create a version of the Coleman-Reffett operator and solver that work with CRRA utility.

The key is to parameterize the utility functions by γ.

```
def u_crra(c, γ):
    return (c**(1 - γ) - 1) / (1 - γ)

def u_prime_crra(c, γ):
    return c**(-γ)

def u_prime_inv_crra(x, γ):
    return x**(-1/γ)
```

Now we create a version of the Coleman-Reffett operator that takes γ as a parameter.

```
def K_crra(
        c_in: jnp.ndarray,  # Consumption values on the endogenous grid
        x_in: jnp.ndarray,  # Current endogenous grid
        model: Model,       # Model specification
        γ: float            # CRRA parameter
    ):
    """
    The Coleman-Reffett operator using EGM with CRRA utility
    """
    # Simplify names
    β, α = model.β, model.α
    s_grid, shocks = model.s_grid, model.shocks

    # Linear interpolation of policy using endogenous grid
    σ = lambda x_val: jnp.interp(x_val, x_in, c_in)

    # Define function to compute consumption at a single grid point
    def compute_c(s):
        vals = u_prime_crra(σ(f(s, α) * shocks), γ) * f_prime(s, α) * shocks
        return u_prime_inv_crra(β * jnp.mean(vals), γ)

    # Vectorize over grid using vmap
    compute_c_vectorized = jax.vmap(compute_c)
    c_out = compute_c_vectorized(s_grid)

    # Determine corresponding endogenous grid
    x_out = s_grid + c_out  # x_i = s_i + c_i

    return c_out, x_out
```

We also need a solver that uses this operator.

```
@jax.jit
def solve_model_crra(model: Model,
                     c_init: jnp.ndarray,
                     x_init: jnp.ndarray,
                     γ: float,
                     tol: float = 1e-5,
                     max_iter: int = 1000):
    """
    Solve the model using time iteration with EGM and CRRA utility.
    """

    def condition(loop_state):
        i, c, x, error = loop_state
        return (error > tol) & (i < max_iter)

    def body(loop_state):
        i, c, x, error = loop_state
        c_new, x_new = K_crra(c, x, model, γ)
        error = jnp.max(jnp.abs(c_new - c))
        return i + 1, c_new, x_new, error

    # Initialize loop state
    initial_state = (0, c_init, x_init, tol + 1)

    # Run the loop
    i, c, x, error = jax.lax.while_loop(condition, body, initial_state)

    return c, x
```

Now we solve for $\gamma = 1$ (log utility) and values approaching 1 from above.

```
γ_values = [1.0, 1.05, 1.1, 1.2]
policies = {}
endogenous_grids = {}

model_crra = create_model()

for γ in γ_values:
    c_init = jnp.copy(model_crra.s_grid)
    x_init = model_crra.s_grid + c_init
    c_gamma, x_gamma = solve_model_crra(model_crra, c_init, x_init, γ)
    jax.block_until_ready(c_gamma)
    policies[γ] = c_gamma
    endogenous_grids[γ] = x_gamma
    print(f"Solved for γ = {γ}")
```

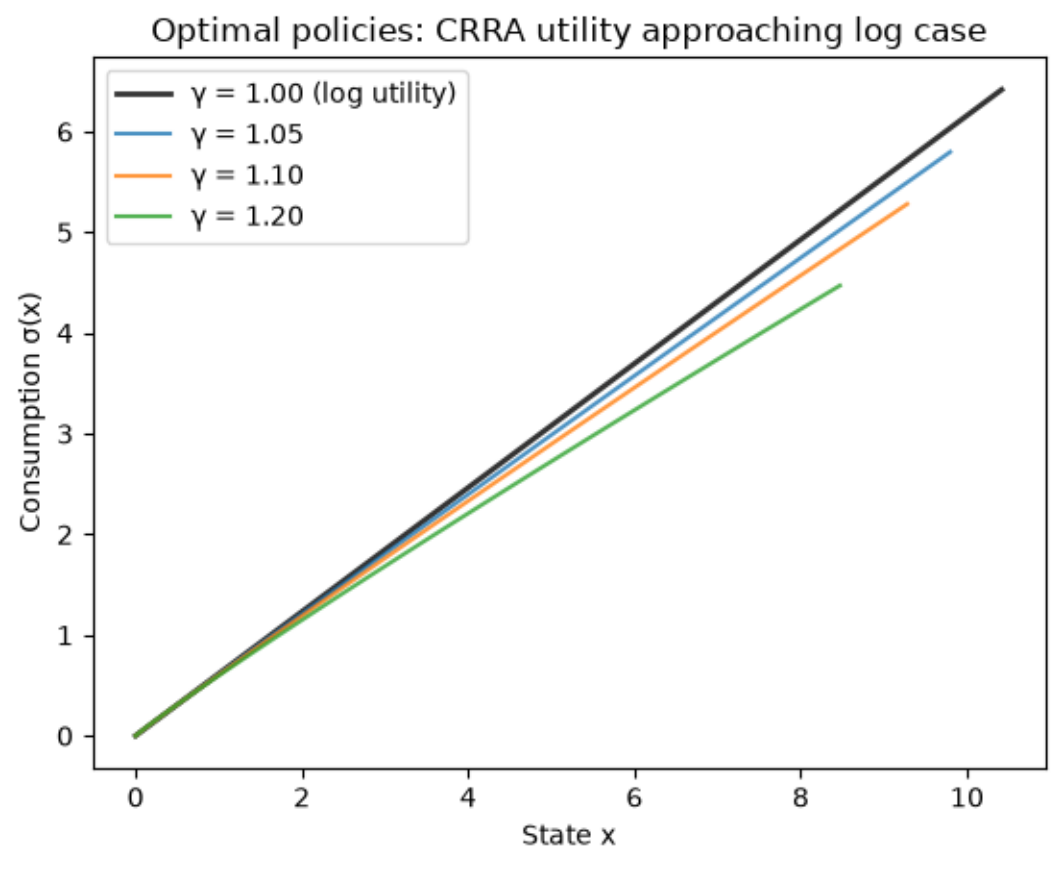
```
Solved for γ = 1.0
Solved for γ = 1.05
Solved for γ = 1.1
Solved for γ = 1.2
```

Plot the policies on their endogenous grids.

```
fig, ax = plt.subplots()

for γ in γ_values:
    x = endogenous_grids[γ]
    if γ == 1.0:
        ax.plot(x, policies[γ], 'k-', linewidth=2,
                label=f'γ = {γ:.2f} (log utility)', alpha=0.8)
    else:
        ax.plot(x, policies[γ], label=f'γ = {γ:.2f}', alpha=0.8)

ax.set_xlabel('State x')
ax.set_ylabel('Consumption σ(x)')
ax.legend()
ax.set_title('Optimal policies: CRRA utility approaching log case')
plt.show()
```



Note that the plots for $\gamma > 1$ do not cover the entire x-axis range shown.

This is because the endogenous grid $x = s + \sigma(s)$ depends on the consumption policy, which varies with γ.

Let's check the maximum deviation between the log utility case ($\gamma = 1.0$) and values approaching from above.

```
for γ in [1.05, 1.1, 1.2]:
    max_diff = jnp.max(jnp.abs(policies[1.0] - policies[γ]))
    print(f"Max difference between γ=1.0 and γ={γ}: {max_diff:.6}")
```

```
Max difference between γ=1.0 and γ=1.05: 0.6192
Max difference between γ=1.0 and γ=1.1: 1.1362
Max difference between γ=1.0 and γ=1.2: 1.94592
```

As expected, the differences decrease as γ approaches 1 from above, confirming convergence.

Part X

Household Problems

CHAPTER

SEVENTYTWO

THE INCOME FLUCTUATION PROBLEM I: DISCRETIZATION AND VFI

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

72.1 Overview

In this lecture, we study an optimal savings problem for an infinitely lived consumer—the "common ancestor" described in [Ljungqvist and Sargent, 2018], section 1.3.

This savings problem is often called an **income fluctuation problem** or a **household problem**.

It is an essential sub-problem for many representative macroeconomic models

- [Aiyagari, 1994]
- [Huggett, 1993]
- etc.

It is related to the decision problem in *Optimal Savings III: Stochastic Returns* but differs in significant ways.

For example,

1. The choice problem for the agent includes an additive income term that leads to an occasionally binding constraint.
2. Shocks affecting the budget constraint are correlated, forcing us to track an extra state variable.

We will begin by working with a relatively basic version of the model and solving it via old-fashioned discretization + value function iteration.

Although this approach is not the fastest or the most efficient, it is very robust and flexible.

For example, if we suddenly decided to add Epstein–Zin preferences, or modify ordinary conditional expectations to quantiles, the technique would continue to work well.

> **Note**
>
> The same is not true of some other methods we will deploy, such as the endogenous grid method.
>
> This is a general rule of computation and analysis — while we can often come up with faster algorithms by exploiting structure, these new algorithms are typically less robust.
>
> They are less robust precisely because they exploit more structure — which implies that they are, inevitably, more vulnerable to change.

In addition to Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

We will use the following imports:

```
import quantecon as qe
import jax
import jax.numpy as jnp
import matplotlib.pyplot as plt
from typing import NamedTuple
from time import time
```

We'll use 64 bit floats to gain extra precision.

```
jax.config.update("jax_enable_x64", True)
```

72.2 Set Up

We study a household that chooses a state-contingent consumption plan $\{c_t\}_{t \geq 0}$ to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$a_{t+1} + c_t \leq R a_t + y_t$$

Here

- c_t is consumption and $c_t \geq 0$,
- a_t is assets and $a_t \geq 0$,
- $R = 1 + r$ is a gross rate of return, and
- $(y_t)_{t \geq 0}$ is labor income, taking values in some finite set Y.

We assume below that labor income dynamics follow a discretized AR(1) process.

We set $\mathsf{S} := \mathbb{R}_+ \times \mathsf{Y}$, which represents the state space.

The **value function** $V \colon \mathsf{S} \to \mathbb{R}$ is defined by

$$V(a, y) := \max \, \mathbb{E} \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t) \right\} \tag{72.1}$$

where the maximization is over all feasible consumption sequences given $(a_0, y_0) = (a, y)$.

The Bellman equation is

$$v(a, y) = \max_{0 \leq a' \leq Ra+y} \left\{ u(Ra + y - a') + \beta \sum_{y'} v(a', y') Q(y, y') \right\}$$

where

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

In the code we use the function

$$B((a, y), a', v) = u(Ra + y - a') + \beta \sum_{y'} v(a', y') Q(y, y').$$

the encapsulate the right hand side of the Bellman equation.

72.3 Code

The following code defines a `NamedTuple` to store the model parameters and grids.

```
class Model(NamedTuple):
    β: float              # Discount factor
    R: float              # Gross interest rate
    γ: float              # CRRA parameter
    a_grid: jnp.ndarray   # Asset grid
    y_grid: jnp.ndarray   # Income grid
    Q: jnp.ndarray        # Markov matrix for income


def create_consumption_model(
        R=1.01,                    # Gross interest rate
        β=0.98,                    # Discount factor
        γ=2,                       # CRRA parameter
        a_min=0.01,                # Min assets
        a_max=10.0,                # Max assets
        a_size=150,                # Grid size
        ρ=0.9, ν=0.1, y_size=100   # Income parameters
    ):
    """
    Creates an instance of the consumption-savings model.

    """
    a_grid = jnp.linspace(a_min, a_max, a_size)
    mc = qe.tauchen(n=y_size, rho=ρ, sigma=ν)
    y_grid, Q = jnp.exp(mc.state_values), jax.device_put(mc.P)
    return Model(β, R, γ, a_grid, y_grid, Q)
```

Now we define the right hand side of the Bellman equation.

We'll use a vectorized coding style reminiscent of Matlab and NumPy (avoiding all loops).

Your are invited to explore an alternative style based around `jax.vmap` in the Exercises.

```
@jax.jit
def B(v, model):
    """
    A vectorized version of the right-hand side of the Bellman equation
    (before maximization), which is a 3D array representing

        B(a, y, a′) = u(Ra + y - a′) + β Σ_y′ v(a′, y′) Q(y, y′)

    for all (a, y, a′).
    """

    # Unpack
    β, R, γ, a_grid, y_grid, Q = model
    a_size, y_size = len(a_grid), len(y_grid)

    # Compute current rewards r(a, y, ap) as array r[i, j, ip]
    a  = jnp.reshape(a_grid, (a_size, 1, 1))    # a[i]   ->  a[i, j, ip]
    y  = jnp.reshape(y_grid, (1, y_size, 1))    # z[j]   ->  z[i, j, ip]
    ap = jnp.reshape(a_grid, (1, 1, a_size))    # ap[ip] -> ap[i, j, ip]
    c = R * a + y - ap

    # Calculate continuation rewards at all combinations of (a, y, ap)
    v = jnp.reshape(v, (1, 1, a_size, y_size))  # v[ip, jp] -> v[i, j, ip, jp]
    Q = jnp.reshape(Q, (1, y_size, 1, y_size))  # Q[j, jp]  -> Q[i, j, ip, jp]
    EV = jnp.sum(v * Q, axis=3)                 # sum over last index jp

    # Compute the right-hand side of the Bellman equation
    return jnp.where(c > 0, c**(1-γ)/(1-γ) + β * EV, -jnp.inf)
```

Some readers might be concerned that we are creating high dimensional arrays, leading to inefficiency.

Could they be avoided by more careful vectorization?

In fact this is not necessary: this function will be JIT-compiled by JAX, and the JIT compiler will optimize compiled code to minimize memory use.

The Bellman operator T can be implemented by

```
@jax.jit
def T(v, model):
    "The Bellman operator."
    return jnp.max(B(v, model), axis=2)
```

The next function computes a v-greedy policy given v (i.e., the policy that maximizes the right-hand side of the Bellman equation.)

```
@jax.jit
def get_greedy(v, model):
    "Computes a v-greedy policy, returned as a set of indices."
    return jnp.argmax(B(v, model), axis=2)
```

72.3.1 Value function iteration

Now we define a solver that implements VFI.

First we write a simple version using a standard Python loop.

```
def value_function_iteration_python(model, tol=1e-5, max_iter=10_000):
    """
    Implements VFI using successive approximation with a Python loop.
    """
    v = jnp.zeros((len(model.a_grid), len(model.y_grid)))
    error = tol + 1
    k = 0

    while error > tol and k < max_iter:
        v_new = T(v, model)
        error = jnp.max(jnp.abs(v_new - v))
        v = v_new
        k += 1

    return v, get_greedy(v, model)
```

Next we write a version that uses `jax.lax.while_loop`.

```
@jax.jit
def value_function_iteration(model, tol=1e-5, max_iter=10_000):
    """
    Implements VFI using successive approximation.
    """
    def body_fun(k_v_err):
        k, v, error = k_v_err
        v_new = T(v, model)
        error = jnp.max(jnp.abs(v_new - v))
        return k + 1, v_new, error

    def cond_fun(k_v_err):
        k, v, error = k_v_err
        return jnp.logical_and(error > tol, k < max_iter)

    v_init = jnp.zeros((len(model.a_grid), len(model.y_grid)))
    k, v_star, error = jax.lax.while_loop(cond_fun, body_fun,
                                          (1, v_init, tol + 1))
    return v_star, get_greedy(v_star, model)
```

72.3.2 Timing

Let's create an instance and compare the two implementations.

```
model = create_consumption_model()
```

First let's time the Python version.

```
print("Starting VFI using Python loop.")
start = time()
v_star_python, σ_star_python = value_function_iteration_python(model)
python_time = time() - start
print(f"VFI completed in {python_time} seconds.")
```

```
Starting VFI using Python loop.
VFI completed in 2.4602842330932617 seconds.
```

Now let's time the `jax.lax.while_loop` version.

```
print("Starting VFI using jax.lax.while_loop.")
start = time()
v_star_jax, σ_star_jax = value_function_iteration(model)
v_star_jax.block_until_ready()
jax_with_compile = time() - start
print(f"VFI completed in {jax_with_compile} seconds.")
```

```
Starting VFI using jax.lax.while_loop.
VFI completed in 1.6095316410064697 seconds.
```

Let's run it again to eliminate compile time.

```
start = time()
v_star_jax, σ_star_jax = value_function_iteration(model)
v_star_jax.block_until_ready()
jax_without_compile = time() - start
print(f"VFI completed in {jax_without_compile} seconds.")
```

```
VFI completed in 1.385802984237671 seconds.
```

Let's check that the two implementations produce the same result.

```
print(f"Values match: {jnp.allclose(v_star_python, v_star_jax)}")
print(f"Policies match: {jnp.allclose(σ_star_python, σ_star_jax)}")
```

```
Values match: True
Policies match: True
```

Here's the speedup from using `jax.lax.while_loop`.

```
print(f"Relative speed = {python_time / jax_without_compile:.2f}")
```

```
Relative speed = 1.78
```

72.3.3 Asset Dynamics

To understand long-run behavior, let's examine the asset accumulation dynamics under the optimal policy.

The following 45-degree diagram shows how assets evolve over time:

```
fig, ax = plt.subplots()

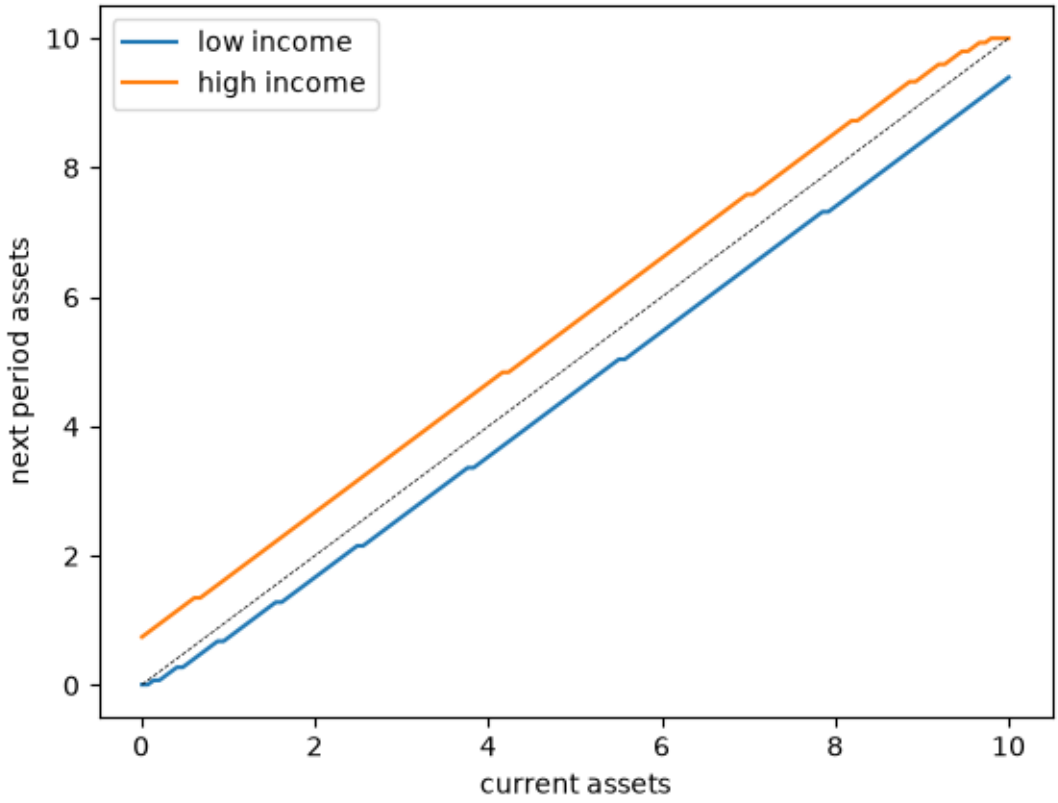
# Plot asset accumulation for first and last income states
for j, label in zip([0, -1], ['low income', 'high income']):
    # Get next-period assets for each current asset level
    a_next = model.a_grid[σ_star_jax[:, j]]
    ax.plot(model.a_grid, a_next, label=label)

# Add 45-degree line
```

(continues on next page)

(continued from previous page)

```
ax.plot(model.a_grid, model.a_grid, 'k--', linewidth=0.5)
ax.set(xlabel='current assets', ylabel='next period assets')
ax.legend()
plt.show()
```



The plot shows the asset accumulation rule for each income state.

The dotted line is the 45-degree line, representing points where $a_{t+1} = a_t$.

We see that:

- For low income levels, assets tend to decrease (points below the 45-degree line)
- For high income levels, assets tend to increase at low asset levels
- The dynamics suggest convergence to a stationary distribution

72.4 Exercises

Exercise 72.4.1

In this exercise, we explore an alternative approach to implementing value function iteration using `jax.vmap`.

For this simple optimal savings problem, direct vectorization is relatively easy.

In particular, it's straightforward to express the right hand side of the Bellman equation as an array that stores evaluations of the function at every state and control.

However, for more complex models, direct vectorization can be much harder.

For this reason, it helps to have another approach to fast JAX implementations up our sleeves.

Your task is to implement a version that:

1. writes the right hand side of the Bellman operator as a function of individual states and controls, and
2. applies `jax.vmap` on the outside to achieve a parallelized solution.

Specifically:

1. Rewrite `B` to take indices `(i, j, ip)` corresponding to `(a, y, a')` and compute the Bellman equation for those specific indices.
2. Use `jax.vmap` successively to vectorize over all indices (use staged vmap as shown in earlier examples).
3. Implement `T_vmap` and `get_greedy_vmap` functions using the vectorized `B`.
4. Implement `value_iteration_vmap` using `jax.lax.while_loop`.
5. Test that your implementation produces the same results as the direct vectorization approach.
6. Compare the execution times of both approaches.

Solution

Here's one solution.

First let's rewrite `B` to work with individual indices:

```
def B(v, model, i, j, ip):
    """
    The right-hand side of the Bellman equation before maximization, which takes
    the form

        B(a, y, a') = u(Ra + y - a') + β Σ_y' v(a', y') Q(y, y')

    The indices are (i, j, ip) -> (a, y, a').
    """
    β, R, γ, a_grid, y_grid, Q = model
    a, y, ap  = a_grid[i], y_grid[j], a_grid[ip]
    c = R * a + y - ap
    EV = jnp.sum(v[ip, :] * Q[j, :])
    return jnp.where(c > 0, c**(1-γ)/(1-γ) + β * EV, -jnp.inf)
```

Now we successively apply `vmap` to simulate nested loops.

```
B_1    = jax.vmap(B,   in_axes=(None, None, None, None, 0))
B_2    = jax.vmap(B_1, in_axes=(None, None, None, 0,    None))
B_vmap = jax.vmap(B_2, in_axes=(None, None, 0,    None, None))
```

Here's the Bellman operator and the `get_greedy` functions for the `vmap` case.

```
@jax.jit
def T_vmap(v, model):
    "The Bellman operator."
    a_indices = jnp.arange(len(model.a_grid))
    y_indices = jnp.arange(len(model.y_grid))
    B_values = B_vmap(v, model, a_indices, y_indices, a_indices)
    return jnp.max(B_values, axis=-1)

@jax.jit
def get_greedy_vmap(v, model):
    "Computes a v-greedy policy, returned as a set of indices."
    a_indices = jnp.arange(len(model.a_grid))
    y_indices = jnp.arange(len(model.y_grid))
    B_values = B_vmap(v, model, a_indices, y_indices, a_indices)
    return jnp.argmax(B_values, axis=-1)
```

Here's the iteration routine.

```
def value_iteration_vmap(model, tol=1e-5, max_iter=10_000):
    """
    Implements VFI using vmap and successive approximation.
    """
    def body_fun(k_v_err):
        k, v, error = k_v_err
        v_new = T_vmap(v, model)
        error = jnp.max(jnp.abs(v_new - v))
        return k + 1, v_new, error

    def cond_fun(k_v_err):
        k, v, error = k_v_err
        return jnp.logical_and(error > tol, k < max_iter)

    v_init = jnp.zeros((len(model.a_grid), len(model.y_grid)))
    k, v_star, error = jax.lax.while_loop(cond_fun, body_fun,
                                          (1, v_init, tol + 1))
    return v_star, get_greedy_vmap(v_star, model)
```

Let's see how long it takes to solve the model using the `vmap` method.

```
print("Starting VFI using vmap.")
start = time()
v_star_vmap, σ_star_vmap = value_iteration_vmap(model)
v_star_vmap.block_until_ready()
jax_vmap_with_compile = time() - start
print(f"VFI completed in {jax_vmap_with_compile} seconds.")
```

```
Starting VFI using vmap.
VFI completed in 0.6106016635894775 seconds.
```

Let's run it again to get rid of compile time.

```
start = time()
v_star_vmap, σ_star_vmap = value_iteration_vmap(model)
v_star_vmap.block_until_ready()
jax_vmap_without_compile = time() - start
print(f"VFI completed in {jax_vmap_without_compile} seconds.")
```

```
VFI completed in 0.3820841312408447 seconds.
```

We need to make sure that we got the same result.

```
print(jnp.allclose(v_star_vmap, v_star_jax))
print(jnp.allclose(σ_star_vmap, σ_star_jax))
```

```
True
True
```

Here's the comparison with the first JAX implementation (which used direct vectorization).

```
print(f"Relative speed = {jax_without_compile / jax_vmap_without_compile}")
```

```
Relative speed = 3.6269577062443803
```

The execution times for the two JAX versions are relatively similar.

However, as emphasized above, having a second method up our sleeves (i.e, the `vmap` approach) will be helpful when confronting dynamic programs with more sophisticated Bellman equations.

CHAPTER

SEVENTYTHREE

THE INCOME FLUCTUATION PROBLEM II: OPTIMISTIC POLICY ITERATION

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

73.1 Overview

In *The Income Fluctuation Problem I: Discretization and VFI* we studied the income fluctuation problem and solved it using value function iteration (VFI).

In this lecture we'll solve the same problem using **optimistic policy iteration** (OPI), which is very general, typically faster than VFI and only slightly more complex.

OPI combines elements of both value function iteration and policy iteration.

A detailed discussion of the algorithm can be found in DP1.

Here our aim is to implement OPI and test whether or not it yields significant speed improvements over standard VFI for the income fluctuation problem.

In addition to Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

We will use the following imports:

```
import quantecon as qe
import jax
import jax.numpy as jnp
import matplotlib.pyplot as plt
from typing import NamedTuple
from time import time
```

73.2 Model and Primitives

The model and parameters are the same as in *The Income Fluctuation Problem I: Discretization and VFI*.

We repeat the key elements here for convenience.

The household's problem is to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$a_{t+1} + c_t \leq R a_t + y_t$$

where $u(c) = c^{1-\gamma}/(1-\gamma)$.

Here's the model structure:

```
class Model(NamedTuple):
    β: float              # Discount factor
    R: float              # Gross interest rate
    γ: float              # CRRA parameter
    a_grid: jnp.ndarray   # Asset grid
    y_grid: jnp.ndarray   # Income grid
    Q: jnp.ndarray        # Markov matrix for income


def create_consumption_model(
        R=1.01,                    # Gross interest rate
        β=0.98,                    # Discount factor
        γ=2,                       # CRRA parameter
        a_min=0.01,                # Min assets
        a_max=10.0,                # Max assets
        a_size=150,                # Grid size
        ρ=0.9, ν=0.1, y_size=100   # Income parameters
    ):
    """
    Creates an instance of the consumption-savings model.

    """
    a_grid = jnp.linspace(a_min, a_max, a_size)
    mc = qe.tauchen(n=y_size, rho=ρ, sigma=ν)
    y_grid, Q = jnp.exp(mc.state_values), jax.device_put(mc.P)
    return Model(β, R, γ, a_grid, y_grid, Q)
```

73.3 Operators and Policies

We repeat some functions from *The Income Fluctuation Problem I: Discretization and VFI*.

Here is the right hand side of the Bellman equation:

```
def B(v, model, i, j, ip):
    """
    The right-hand side of the Bellman equation before maximization, which takes
    the form
```

(continues on next page)

(continued from previous page)

```
        B(a, y, a') = u(Ra + y - a') + β Σ_y' v(a', y') Q(y, y')

    The indices are (i, j, ip) -> (a, y, a').
    """
    β, R, γ, a_grid, y_grid, Q = model
    a, y, ap  = a_grid[i], y_grid[j], a_grid[ip]
    c = R * a + y - ap
    EV = jnp.sum(v[ip, :] * Q[j, :])
    return jnp.where(c > 0, c**(1-γ)/(1-γ) + β * EV, -jnp.inf)
```

Now we successively apply `vmap` to vectorize over all indices:

```
B_1    = jax.vmap(B,   in_axes=(None, None, None, None, 0))
B_2    = jax.vmap(B_1, in_axes=(None, None, None, 0,    None))
B_vmap = jax.vmap(B_2, in_axes=(None, None, 0,    None, None))
```

Here's the Bellman operator:

```
def T(v, model):
    "The Bellman operator."
    a_indices = jnp.arange(len(model.a_grid))
    y_indices = jnp.arange(len(model.y_grid))
    B_values = B_vmap(v, model, a_indices, y_indices, a_indices)
    return jnp.max(B_values, axis=-1)
```

Here's the function that computes a v-greedy policy:

```
def get_greedy(v, model):
    "Computes a v-greedy policy, returned as a set of indices."
    a_indices = jnp.arange(len(model.a_grid))
    y_indices = jnp.arange(len(model.y_grid))
    B_values = B_vmap(v, model, a_indices, y_indices, a_indices)
    return jnp.argmax(B_values, axis=-1)
```

Now we define the policy operator T_σ, which is the Bellman operator with policy σ fixed.

For a given policy σ, the policy operator is defined by

$$(T_\sigma v)(a, y) = u(Ra + y - \sigma(a, y)) + \beta \sum_{y'} v(\sigma(a, y), y') Q(y, y')$$

```
def T_σ(v, σ, model, i, j):
    """
    The σ-policy operator for indices (i, j) -> (a, y).
    """
    β, R, γ, a_grid, y_grid, Q = model

    # Get values at current state
    a, y = a_grid[i], y_grid[j]
    # Get policy choice
    ap = a_grid[σ[i, j]]

    # Compute current reward
    c = R * a + y - ap
    r = jnp.where(c > 0, c**(1-γ)/(1-γ), -jnp.inf)
```

(continues on next page)

(continued from previous page)

```
    # Compute expected value
    EV = jnp.sum(v[σ[i, j], :] * Q[j, :])

    return r + β * EV
```

Apply vmap to vectorize:

```
T_σ_1    = jax.vmap(T_σ,   in_axes=(None, None, None, None, 0))
T_σ_vmap = jax.vmap(T_σ_1, in_axes=(None, None, None, 0,    None))

def T_σ_vec(v, σ, model):
    """Vectorized version of T_σ."""
    a_size, y_size = len(model.a_grid), len(model.y_grid)
    a_indices = jnp.arange(a_size)
    y_indices = jnp.arange(y_size)
    return T_σ_vmap(v, σ, model, a_indices, y_indices)
```

Now we need a function to apply the policy operator m times:

```
def iterate_policy_operator(σ, v, m, model):
    """
    Apply the policy operator T_σ exactly m times to v.
    """
    def update(i, v):
        return T_σ_vec(v, σ, model)

    v = jax.lax.fori_loop(0, m, update, v)
    return v
```

73.4 Value Function Iteration

For comparison, here's VFI from *The Income Fluctuation Problem I: Discretization and VFI*:

```
@jax.jit
def value_function_iteration(model, tol=1e-5, max_iter=10_000):
    """
    Implements VFI using successive approximation.
    """
    def body_fun(k_v_err):
        k, v, error = k_v_err
        v_new = T(v, model)
        error = jnp.max(jnp.abs(v_new - v))
        return k + 1, v_new, error

    def cond_fun(k_v_err):
        k, v, error = k_v_err
        return jnp.logical_and(error > tol, k < max_iter)

    v_init = jnp.zeros((len(model.a_grid), len(model.y_grid)))
    k, v_star, error = jax.lax.while_loop(cond_fun, body_fun,
                                          (1, v_init, tol + 1))
    return v_star, get_greedy(v_star, model)
```

73.5 Optimistic Policy Iteration

Now we implement OPI.

The algorithm alternates between

1. Performing m policy operator iterations to update the value function
2. Computing a new greedy policy based on the updated value function

```
@jax.jit
def optimistic_policy_iteration(model, m=10, tol=1e-5, max_iter=10_000):
    """
    Implements optimistic policy iteration with step size m.

    Parameters:
    -----------
    model : Model
        The consumption-savings model
    m : int
        Number of policy operator iterations per step
    tol : float
        Tolerance for convergence
    max_iter : int
        Maximum number of iterations
    """
    v_init = jnp.zeros((len(model.a_grid), len(model.y_grid)))

    def condition_function(inputs):
        i, v, error = inputs
        return jnp.logical_and(error > tol, i < max_iter)

    def update(inputs):
        i, v, error = inputs
        last_v = v
        σ = get_greedy(v, model)
        v = iterate_policy_operator(σ, v, m, model)
        error = jnp.max(jnp.abs(v - last_v))
        i += 1
        return i, v, error

    num_iter, v, error = jax.lax.while_loop(condition_function,
                                            update,
                                            (0, v_init, tol + 1))

    return v, get_greedy(v, model)
```

73.6 Timing Comparison

Let's create a model and compare the performance of VFI and OPI.

```
model = create_consumption_model()
```

First, let's time VFI:

```
print("Starting VFI.")
start = time()
v_star_vfi, σ_star_vfi = value_function_iteration(model)
v_star_vfi.block_until_ready()
vfi_time_with_compile = time() - start
print(f"VFI completed in {vfi_time_with_compile:.2f} seconds.")
```

```
Starting VFI.
VFI completed in 0.37 seconds.
```

Run it again to eliminate compile time:

```
start = time()
v_star_vfi, σ_star_vfi = value_function_iteration(model)
v_star_vfi.block_until_ready()
vfi_time = time() - start
print(f"VFI completed in {vfi_time:.2f} seconds.")
```

```
VFI completed in 0.10 seconds.
```

Now let's time OPI with different values of m:

```
print("Starting OPI with m=50.")
start = time()
v_star_opi, σ_star_opi = optimistic_policy_iteration(model, m=50)
v_star_opi.block_until_ready()
opi_time_with_compile = time() - start
print(f"OPI completed in {opi_time_with_compile:.2f} seconds.")
```

```
Starting OPI with m=50.
OPI completed in 0.33 seconds.
```

Run it again:

```
start = time()
v_star_opi, σ_star_opi = optimistic_policy_iteration(model, m=50)
v_star_opi.block_until_ready()
opi_time = time() - start
print(f"OPI completed in {opi_time:.2f} seconds.")
```

```
OPI completed in 0.03 seconds.
```

Check that we get the same result:

```
print(f"Values match: {jnp.allclose(v_star_vfi, v_star_opi)}")
```

```
Values match: True
```

The value functions match, confirming both algorithms converge to the same solution.

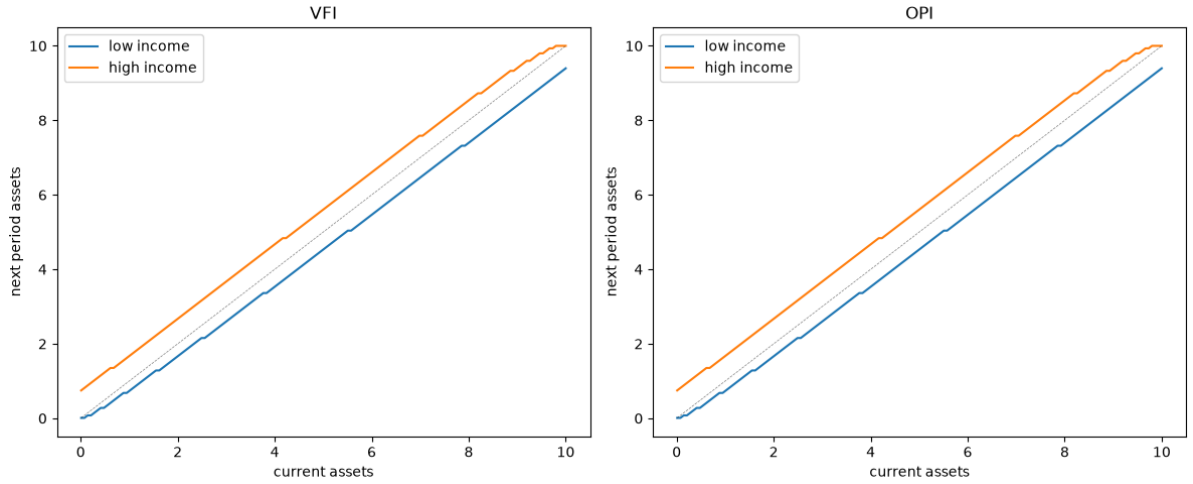
Let's visually compare the asset dynamics under both policies:

```
fig, axes = plt.subplots(1, 2, figsize=(12, 5))

# VFI policy
for j, label in zip([0, -1], ['low income', 'high income']):
    a_next_vfi = model.a_grid[σ_star_vfi[:, j]]
    axes[0].plot(model.a_grid, a_next_vfi, label=label)
axes[0].plot(model.a_grid, model.a_grid, 'k--', linewidth=0.5, alpha=0.5)
axes[0].set(xlabel='current assets', ylabel='next period assets', title='VFI')
axes[0].legend()

# OPI policy
for j, label in zip([0, -1], ['low income', 'high income']):
    a_next_opi = model.a_grid[σ_star_opi[:, j]]
    axes[1].plot(model.a_grid, a_next_opi, label=label)
axes[1].plot(model.a_grid, model.a_grid, 'k--', linewidth=0.5, alpha=0.5)
axes[1].set(xlabel='current assets', ylabel='next period assets', title='OPI')
axes[1].legend()

plt.tight_layout()
plt.show()
```



The policies are visually indistinguishable, confirming both methods produce the same solution.

Here's the speedup:

```
print(f"Speedup factor: {vfi_time / opi_time:.2f}")
```

```
Speedup factor: 3.21
```

Let's try different values of m to see how it affects performance:

```
m_vals = [1, 5, 10, 25, 50, 100, 200, 400]
opi_times = []
```

(continues on next page)

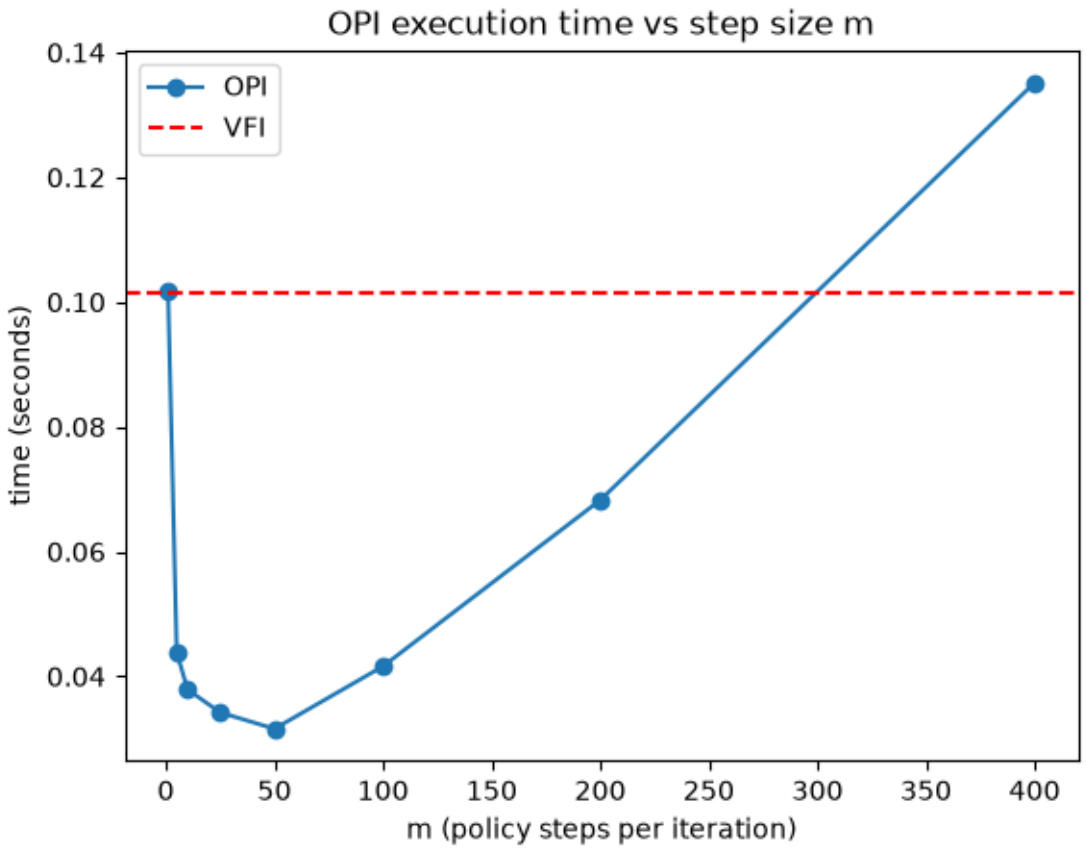
(continued from previous page)

```
for m in m_vals:
    start = time()
    v_star, σ_star = optimistic_policy_iteration(model, m=m)
    v_star.block_until_ready()
    elapsed = time() - start
    opi_times.append(elapsed)
    print(f"OPI with m={m:3d} completed in {elapsed:.2f} seconds.")
```

```
OPI with m=  1 completed in 0.10 seconds.
OPI with m=  5 completed in 0.04 seconds.
OPI with m= 10 completed in 0.04 seconds.
OPI with m= 25 completed in 0.03 seconds.
OPI with m= 50 completed in 0.03 seconds.
OPI with m=100 completed in 0.04 seconds.
OPI with m=200 completed in 0.07 seconds.
OPI with m=400 completed in 0.14 seconds.
```

Plot the results:

```
fig, ax = plt.subplots()
ax.plot(m_vals, opi_times, 'o-', label='OPI')
ax.axhline(vfi_time, linestyle='--', color='red', label='VFI')
ax.set_xlabel('m (policy steps per iteration)')
ax.set_ylabel('time (seconds)')
ax.legend()
ax.set_title('OPI execution time vs step size m')
plt.show()
```



Here's a summary of the results

- OPI outperforms VFI for a large range of m values.
- For very large m, OPI performance begins to degrade as we spend too much time iterating the policy operator.

73.7 Exercises

Exercise 73.7.1

The speed gains achieved by OPI are quite robust to parameter changes.

Confirm this by experimenting with different parameter values for the income process (ρ and ν).

Measure how they affect the relative performance of VFI vs OPI.

Try:

- $\rho \in \{0.8, 0.9, 0.95\}$
- $\nu \in \{0.05, 0.1, 0.2\}$

For each combination, compute the speedup factor (VFI time / OPI time) and report your findings.

Solution

Here's one solution:

```
ρ_vals = [0.8, 0.9, 0.95]
ν_vals = [0.05, 0.1, 0.2]

results = []

for ρ in ρ_vals:
    for ν in ν_vals:
        print(f"\nTesting ρ={ρ}, ν={ν}")

        # Create model
        model = create_consumption_model(ρ=ρ, ν=ν)

        # Time VFI
        start = time()
        v_vfi, σ_vfi = value_function_iteration(model)
        v_vfi.block_until_ready()
        vfi_t = time() - start

        # Time OPI
        start = time()
        v_opi, σ_opi = optimistic_policy_iteration(model, m=10)
        v_opi.block_until_ready()
        opi_t = time() - start

        speedup = vfi_t / opi_t
        results.append((ρ, ν, speedup))
        print(f"  VFI: {vfi_t:.2f}s, OPI: {opi_t:.2f}s, Speedup: {speedup:.2f}x")

# Print summary
print("\nSummary of speedup factors:")
for ρ, ν, speedup in results:
    print(f"ρ={ρ}, ν={ν}: {speedup:.2f}x")
```

```
Testing ρ=0.8, ν=0.05
  VFI: 0.06s, OPI: 0.03s, Speedup: 1.68x

Testing ρ=0.8, ν=0.1
  VFI: 0.06s, OPI: 0.03s, Speedup: 1.80x

Testing ρ=0.8, ν=0.2
  VFI: 0.06s, OPI: 0.03s, Speedup: 1.73x

Testing ρ=0.9, ν=0.05
  VFI: 0.06s, OPI: 0.03s, Speedup: 1.80x

Testing ρ=0.9, ν=0.1
  VFI: 0.06s, OPI: 0.03s, Speedup: 1.70x

Testing ρ=0.9, ν=0.2
  VFI: 0.06s, OPI: 0.03s, Speedup: 1.76x

Testing ρ=0.95, ν=0.05
  VFI: 0.06s, OPI: 0.04s, Speedup: 1.69x

Testing ρ=0.95, ν=0.1
  VFI: 0.06s, OPI: 0.04s, Speedup: 1.74x

Testing ρ=0.95, ν=0.2
  VFI: 0.07s, OPI: 0.04s, Speedup: 1.87x

Summary of speedup factors:
ρ=0.8, ν=0.05: 1.68x
```

CHAPTER

SEVENTYFOUR

THE INCOME FLUCTUATION PROBLEM III: THE ENDOGENOUS GRID METHOD

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *The Income Fluctuation Problem III: The Endogenous Grid Method*
  - *Overview*
  - *The Household Problem*
  - *Computation*
  - *NumPy Implementation*
  - *JAX Implementation*
  - *Simulation*

74.1 Overview

In this lecture we continue examining a version of the IFP from

- *The Income Fluctuation Problem I: Discretization and VFI* and
- *The Income Fluctuation Problem II: Optimistic Policy Iteration*.

We will make two changes.

1. Change the timing to one that is more efficient for our set up.
2. Use the endogenous grid method (EGM) to solve the model.

We use EGM because we know it to be fast and accurate from *Optimal Savings VI: EGM with JAX*.

The primary source for the technical details discussed below is [Ma *et al.*, 2020].

Other references include [Deaton, 1991], [Den Haan, 2010], [Kuhn, 2013], [Rabault, 2002], [Reiter, 2009] and [Schechtman and Escudero, 1977].

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

We'll also need the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
import numba
from quantecon import MarkovChain
import jax
import jax.numpy as jnp
from typing import NamedTuple
```

74.2 The Household Problem

Let's write down the model and then discuss how to solve it.

74.2.1 Set-Up

A household chooses a state-contingent consumption plan $\{c_t\}_{t \geq 0}$ to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$a_{t+1} = R(a_t - c_t) + Y_{t+1} \quad c_t \geq 0, \quad a_t \geq 0 \quad t = 0, 1, \ldots \tag{74.1}$$

Here

- $\beta \in (0, 1)$ is the discount factor
- a_t is asset holdings at time t, with borrowing constraint $a_t \geq 0$
- c_t is consumption
- Y_t is non-capital income (wages, unemployment compensation, etc.)
- $R := 1 + r$, where $r > 0$ is the interest rate on savings

The timing here is as follows:

1. At the start of period t, the household observes current asset holdings a_t.
2. The household chooses current consumption c_t.
3. Savings $s_t := a_t - c_t$ earns interest at rate r.
4. Labor income Y_{t+1} is realized and time shifts to $t + 1$.

Non-capital income Y_t is given by $Y_t = y(Z_t)$, where

- $\{Z_t\}$ is an exogenous state process, and
- y is a function taking values in $\mathbb{R}_+$.

We take $\{Z_t\}$ to be a finite state Markov chain taking values in Z with Markov matrix Π.

Note

In previous lectures we used the more standard household budget constraint $a_{t+1} + c_t \leq Ra_t + Y_t$.

This setup, which is pervasive in quantitative economics, was developed for discretization.

It means that the control variable is also the next period state a_{t+1}, which makes it straightforward to restrict assets to a finite grid.

But fixing the control to be the next period state forces us to include more information in the current state, which expands the size of the state space.

Moreover, aiming for discretization is not always a good idea, since it suffers heavily from the curse of dimensionality.

These ideas will become clearer in the *next lecture*.

We further assume that

1. $\beta R < 1$
2. u is smooth, strictly increasing and strictly concave with $\lim_{c \to 0} u'(c) = \infty$ and $\lim_{c \to \infty} u'(c) = 0$
3. $y(z) = \exp(z)$

The asset space is $\mathbb{R}_+$ and the state is the pair $(a, z) \in \mathsf{S} := \mathbb{R}_+ \times \mathsf{Z}$.

A **feasible consumption path** from $(a, z) \in \mathsf{S}$ is a consumption sequence $\{c_t\}$ such that $\{c_t\}$ and its induced asset path $\{a_t\}$ satisfy

1. $(a_0, z_0) = (a, z)$
2. the feasibility constraints in (74.1), and
3. adaptedness, which means that c_t is a function of random outcomes up to date t but not after.

The meaning of the third point is just that consumption at time t cannot be a function of outcomes are yet to be observed.

In fact, for this problem, consumption can be chosen optimally by taking it to be contingent only on the current state.

Optimality is defined below.

74.2.2 Value Function and Euler Equation

The **value function** $V \colon \mathsf{S} \to \mathbb{R}$ is defined by

$$V(a, z) := \max \mathbb{E} \left\{ \sum_{t=0}^{\infty} \beta^t u(c_t) \right\} \tag{74.2}$$

where the maximization is overall feasible consumption paths from (a, z).

An **optimal consumption path** from (a, z) is a feasible consumption path from (a, z) that maximizes (72.1).

To pin down such paths we can use a version of the Euler equation, which in the present setting is

$$u'(c_t) \geq \beta R \, \mathbb{E}_t u'(c_{t+1}) \tag{74.3}$$

with

$$c_t < a_t \implies u'(c_t) = \beta R \, \mathbb{E}_t u'(c_{t+1}) \tag{74.4}$$

When c_t hits the upper bound a_t, the strict inequality $u'(c_t) > \beta R \, \mathbb{E}_t u'(c_{t+1})$ can occur because c_t cannot increase sufficiently to attain equality.

The case $c_t = 0$ never arises along the optimal path because $u'(0) = \infty$.

74.2.3 Optimality Results

As shown in [Ma *et al.*, 2020],

1. For each $(a, z) \in \mathsf{S}$, a unique optimal consumption path from (a, z) exists
2. This path is the unique feasible path from (a, z) satisfying the Euler equations (74.3)-(74.4) and the transversality condition

$$\lim_{t \to \infty} \beta^t \, \mathbb{E} \left[u'(c_t) a_{t+1} \right] = 0 \tag{74.5}$$

Moreover, there exists an **optimal consumption policy** $\sigma^* \colon \mathsf{S} \to \mathbb{R}_+$ such that the path from (a, z) generated by

$$(a_0, z_0) = (a, z), \quad c_t = \sigma^*(a_t, Z_t) \quad \text{and} \quad a_{t+1} = R(a_t - c_t) + Y_{t+1}$$

satisfies both the Euler equations (74.3)-(74.4) and (74.5), and hence is the unique optimal path from (a, z).

Thus, to solve the optimization problem, we need to compute the policy σ^*.

74.3 Computation

We solve for the optimal consumption policy using time iteration and the endogenous grid method, which were previously discussed in

- *Optimal Savings IV: Time Iteration*
- *Optimal Savings V: The Endogenous Grid Method*

74.3.1 Solution Method

We rewrite (74.4) to make it a statement about functions rather than random variables:

$$(u' \circ \sigma)(a, z) = \beta R \sum_{z'} (u' \circ \sigma)[R(a - \sigma(a, z)) + y(z'), \, z'] \, \Pi(z, z') \tag{74.6}$$

Here

- $(u' \circ \sigma)(s) := u'(\sigma(s))$,
- primes indicate next period states (as well as derivatives), and
- σ is the unknown function.

The equality (74.6) holds at all interior choices, meaning $\sigma(a, z) < a$.

We aim to find a fixed point σ of (74.6).

To do so we use the EGM.

Below we use the relationships $a_t = c_t + s_t$ and $a_{t+1} = R s_t + Y_{t+1}$.

We begin with an exogenous savings grid $s_0 < s_1 < \cdots < s_m$ with $s_0 = 0$.

We fix a current guess of the policy function σ.

For each exogenous savings level s_i with $i \geq 1$ and current state z_j, we set

$$c_{ij} := (u')^{-1} \left[\beta R \sum_{z'} u'[\sigma(Rs_i + y(z'), z')] \, \Pi(z_j, z') \right] \tag{74.7}$$

The Euler equation holds here because $i \geq 1$ implies $s_i > 0$ and hence consumption is interior.

For the boundary case $s_0 = 0$ we set

$$c_{0j} := 0 \quad \text{for all j}$$

We then obtain a corresponding endogenous grid of current assets via

$$a_{ij} := c_{ij} + s_i.$$

Notice that, for each j, we have $a_{0j} = c_{0j} = 0$.

This anchors the interpolation at the correct value at the origin, since, without borrowing, consumption is zero when assets are zero.

Our next guess of the policy function, which we write as $K\sigma$, is the linear interpolation of the interpolation points

$$\{(a_{0j}, c_{0j}), \ldots, (a_{mj}, c_{mj})\}$$

for each j.

(The number of one-dimensional linear interpolations is equal to the size of Z.)

74.4 NumPy Implementation

In this section we'll code up a NumPy version of the code that aims only for clarity, rather than efficiency.

Once we have it working, we'll produce a JAX version that's far more efficient and check that we obtain the same results.

We use the CRRA utility specification

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

74.4.1 Set Up

Here we build a class called `IFPNumPy` that stores the model primitives.

The exogenous state process $\{Z_t\}$ defaults to a two-state Markov chain with transition matrix Π.

```
class IFPNumPy(NamedTuple):
    R: float                  # Gross interest rate R = 1 + r
    β: float                  # Discount factor
    γ: float                  # Preference parameter
    Π: np.ndarray             # Markov matrix for exogenous shock
    z_grid: np.ndarray        # Markov state values for Z_t
    s: np.ndarray             # Exogenous savings grid
```

(continues on next page)

(continued from previous page)

```
def create_ifp(r=0.01,
               β=0.96,
               γ=1.5,
               Π=((0.6, 0.4),
                  (0.05, 0.95)),
               z_grid=(-10.0, np.log(2.0)),
               savings_grid_max=16,
               savings_grid_size=200):

    s = np.linspace(0, savings_grid_max, savings_grid_size)
    Π, z_grid = np.array(Π), np.array(z_grid)
    R = 1 + r
    assert R * β < 1, "Stability condition violated."
    return IFPNumPy(R, β, γ, Π, z_grid, s)
```

74.4.2 Solver

Here is the operator K that transforms current guess σ into next period guess $K\sigma$.

In practice, it takes in

- a guess of optimal consumption values c_{ij}, stored as `c_vec`
- and a corresponding set of endogenous grid points a^e_{ij}, stored as `a_vec`

These are converted into a consumption policy $a \mapsto \sigma(a, z_j)$ by linear interpolation of (a^e_{ij}, c_{ij}) over i for each j.

Since there are no shocks to integrate out in this version of the model, we can compute (74.7) directly by summing over the finite state space Z.

```
@numba.jit
def K_numpy(
        c_in: np.ndarray,    # Initial guess of σ on grid endogenous grid
        a_in: np.ndarray,    # Initial endogenous grid
        ifp_numpy: IFPNumPy
    ) -> np.ndarray:
    """
    The Euler equation operator for the IFP model using the
    Endogenous Grid Method.

    This operator implements one iteration of the EGM algorithm to
    update the consumption policy function.

    """
    R, β, γ, Π, z_grid, s = ifp_numpy
    n_a, n_z = len(s), len(z_grid)
    c_out = np.zeros_like(c_in)
    u_prime = lambda c: c**(-γ)
    u_prime_inv = lambda c: c**(-1/γ)
    y = lambda z: np.exp(z)

    for i in range(1, n_a):  # Start from 1 for positive savings levels
        for j in range(n_z):

            # Compute Σ_z' u'(σ(R s_i + y(z'), z')) Π[z_j, z']
```

(continues on next page)

(continued from previous page)

```
            expectation = 0.0
            for k in range(n_z):
                z_prime = z_grid[k]
                # Calculate next period assets
                next_a = R * s[i] + y(z_prime)
                # Interpolate to get σ(R s_i + y(z'), z')
                next_c = np.interp(next_a, a_in[:, k], c_in[:, k])
                # Weight by transition probability and add to the expectation
                expectation += u_prime(next_c) * Π[j, k]

            # Calculate updated c_{ij} values
            c_out[i, j] = u_prime_inv(β * R * expectation)

    a_out = c_out + s[:, None]
    return c_out, a_out
```

To solve the model we use a simple while loop.

```
def solve_model_numpy(
        ifp_numpy: IFPNumPy,
        c_init: np.ndarray,
        a_init: np.ndarray,
        tol: float = 1e-5,
        max_iter: int = 1_000
    ) -> np.ndarray:
    """
    Solve the model using time iteration with EGM.

    """
    c_in, a_in = c_init, a_init
    i = 0
    error = tol + 1

    while error > tol and i < max_iter:
        c_out, a_out = K_numpy(c_in, a_in, ifp_numpy)
        error = np.max(np.abs(c_out - c_in))
        i = i + 1
        c_in, a_in = c_out, a_out

    return c_out, a_out
```

Let's road test the EGM code.

```
ifp_numpy = create_ifp()
R, β, γ, Π, z_grid, s = ifp_numpy
# Initial conditions -- agent consumes everything
a_init = s[:, None] * np.ones(len(z_grid))
c_init = a_init
# Solve from these initial conditions
c_vec, a_vec = solve_model_numpy(
    ifp_numpy, c_init, a_init
)
```

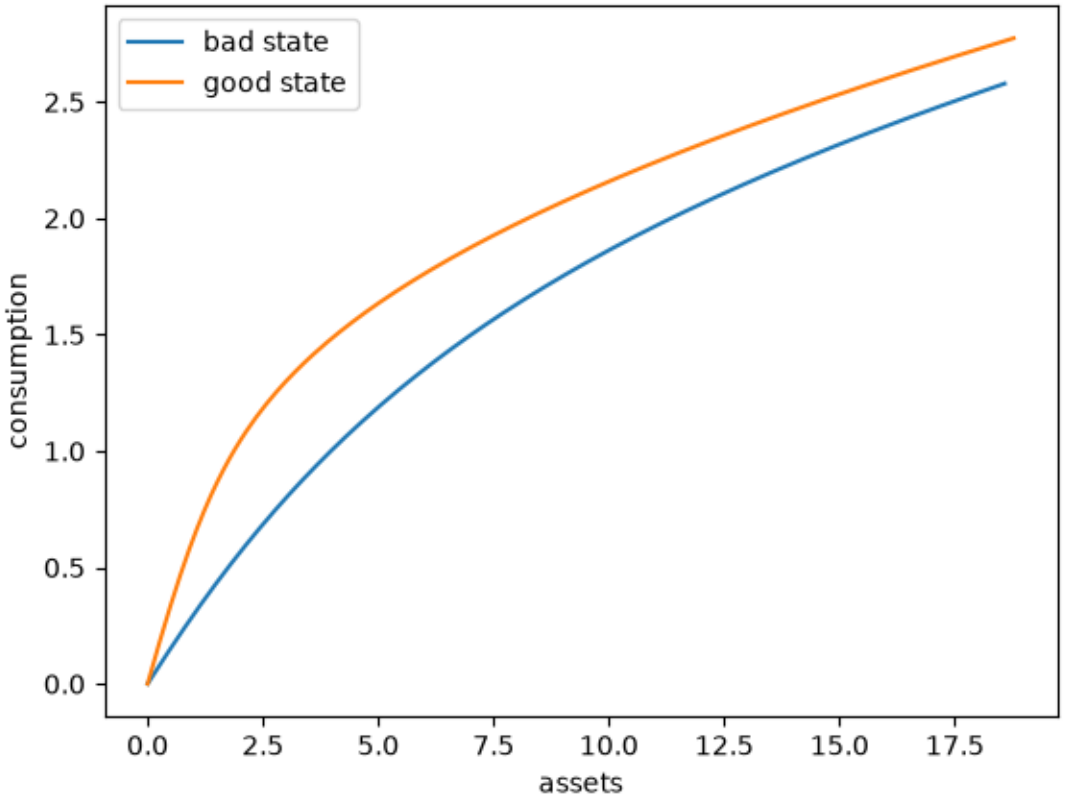
Here's a plot of the optimal consumption policy for each z state

```
fig, ax = plt.subplots()
```

(continues on next page)

(continued from previous page)

```
ax.plot(a_vec[:, 0], c_vec[:, 0], label='bad state')
ax.plot(a_vec[:, 1], c_vec[:, 1], label='good state')
ax.set(xlabel='assets', ylabel='consumption')
ax.legend()
plt.show()
```



74.5 JAX Implementation

Now we write a more efficient JAX version, which can run on a GPU.

74.5.1 Set Up

We start with a class called `IFP` that stores the model primitives.

```
class IFP(NamedTuple):
    R: float                  # Gross interest rate R = 1 + r
    β: float                  # Discount factor
    γ: float                  # Preference parameter
    Π: jnp.ndarray            # Markov matrix for exogenous shock
    z_grid: jnp.ndarray       # Markov state values for Z_t
    s: jnp.ndarray            # Exogenous savings grid


def create_ifp(r=0.01,
```

(continues on next page)

(continued from previous page)

```
                β=0.94,
                γ=1.5,
                Π=((0.6, 0.4),
                   (0.05, 0.95)),
                z_grid=(-10.0, jnp.log(2.0)),
                savings_grid_max=16,
                savings_grid_size=200):

    s = jnp.linspace(0, savings_grid_max, savings_grid_size)
    Π, z_grid = jnp.array(Π), jnp.array(z_grid)
    R = 1 + r
    assert R * β < 1, "Stability condition violated."
    return IFP(R, β, γ, Π, z_grid, s)
```

74.5.2 Solver

Here is the operator K that transforms current guess σ into next period guess $K\sigma$.

```
def K(
        c_in: jnp.ndarray,
        a_in: jnp.ndarray,
        ifp: IFP
    ) -> jnp.ndarray:
    """
    The Euler equation operator for the IFP model using the
    Endogenous Grid Method.

    This operator implements one iteration of the EGM algorithm to
    update the consumption policy function.

    """
    R, β, γ, Π, z_grid, s = ifp
    n_z = len(z_grid)
    z_indices = jnp.arange(n_z)
    u_prime = lambda c: c**(-γ)
    u_prime_inv = lambda c: c**(-1/γ)
    y = lambda z: jnp.exp(z)

    def compute_c(i, j):
        " Computes consumption for one (i, j) pair where i >= 1. "

        def compute_mu_k(k):
            " Given i, compute marginal utility u'(σ(R s_i + y(z_k), z_k)) "
            next_a = R * s[i] + y(z_grid[k])
            # Interpolate to get σ(R * s_i + y(z_k), z_k)
            next_c = jnp.interp(next_a, a_in[:, k], c_in[:, k])
            # Return u'(σ(R * s_i + y(z_k), z_k))
            return u_prime(next_c)

        # Compute marginal utility u'(σ(R * s_i + y(z_k), z_k)) for all k
        mu_values = jax.vmap(compute_mu_k)(z_indices)
        # Compute expectation Σ_k u'(σ(...)) * Π[j, k]
        expectation = jnp.sum(mu_values * Π[j, :])
        # Invert to get consumption c_{ij} at (s_i, z_j)
        return u_prime_inv(β * R * expectation)
```

(continues on next page)

(continued from previous page)

```
    # vmap over j for each i
    compute_c_i = jax.vmap(compute_c, in_axes=(None, 0))
    # vmap over i
    compute_c = jax.vmap(lambda i: compute_c_i(i, z_indices))
    # Compute consumption for i >= 1
    c_out_interior = compute_c(jnp.arange(1, len(s)))
    # For i = 0, set consumption to 0
    c_out_boundary = jnp.zeros((1, n_z))
    # Concatenate boundary and interior
    c_out = jnp.concatenate([c_out_boundary, c_out_interior], axis=0)
    # Compute endogenous asset grid a_{ij} = c_{ij} + s_i
    a_out = c_out + s[:, None]
    return c_out, a_out
```

Here's a jit-accelerated iterative routine to solve the model using this operator.

```
@jax.jit
def solve_model(
        ifp: IFP,
        c_init: jnp.ndarray,  # Initial guess of σ on grid endogenous grid
        a_init: jnp.ndarray,  # Initial endogenous grid
        tol: float = 1e-5,
        max_iter: int = 1000
    ) -> jnp.ndarray:
    """
    Solve the model using time iteration with EGM.

    """

    def condition(loop_state):
        c_in, a_in, i, error = loop_state
        return (error > tol) & (i < max_iter)

    def body(loop_state):
        c_in, a_in, i, error = loop_state
        c_out, a_out = K(c_in, a_in, ifp)
        error = jnp.max(jnp.abs(c_out - c_in))
        i += 1
        return c_out, a_out, i, error

    i, error = 0, tol + 1
    initial_state = (c_init, a_init, i, error)
    final_loop_state = jax.lax.while_loop(condition, body, initial_state)
    c_out, a_out, i, error = final_loop_state

    return c_out, a_out
```

74.5.3 Test run

Let's road test the EGM code.

```
ifp = create_ifp()
R, β, γ, Π, z_grid, s = ifp
# Set initial conditions where the agent consumes everything
a_init = s[:, None] * jnp.ones(len(z_grid))
c_init = a_init
# Solve starting from these initial conditions
c_vec_jax, a_vec_jax = solve_model(ifp, c_init, a_init)
```

To verify the correctness of our JAX implementation, let's compare it with the NumPy version we developed earlier.

```
# Compare the results
max_c_diff = np.max(np.abs(np.array(c_vec) - c_vec_jax))
max_ae_diff = np.max(np.abs(np.array(a_vec) - a_vec_jax))

print(f"Maximum difference in consumption policy: {max_c_diff:.2e}")
print(f"Maximum difference in asset grid:        {max_ae_diff:.2e}")
```

```
Maximum difference in consumption policy: 4.45e-01
Maximum difference in asset grid:        4.45e-01
```

These numbers confirm that we are computing essentially the same policy using the two approaches.

74.5.4 Timing

Now let's compare the execution time between NumPy and JAX implementations.

```
import time

# Set up initial conditions for NumPy version
s_np = np.array(s)
z_grid_np = np.array(z_grid)
a_init_np = s_np[:, None] * np.ones(len(z_grid_np))
c_init_np = a_init_np.copy()

# Set up initial conditions for JAX version
a_init_jx = s[:, None] * jnp.ones(len(z_grid))
c_init_jx = a_init_jx

# Time NumPy version
start = time.time()
c_vec_np, a_vec_np = solve_model_numpy(ifp_numpy, c_init_np, a_init_np)
numpy_time = time.time() - start

# Time JAX version (with compilation)
start = time.time()
c_vec_jx, a_vec_jx = solve_model(ifp, c_init_jx, a_init_jx)
c_vec_jx.block_until_ready()
jax_time_with_compile = time.time() - start

# Time JAX version (without compilation - second run)
start = time.time()
c_vec_jx, a_vec_jx = solve_model(ifp, c_init_jx, a_init_jx)
```

(continues on next page)

(continued from previous page)

```
c_vec_jx.block_until_ready()
jax_time = time.time() - start

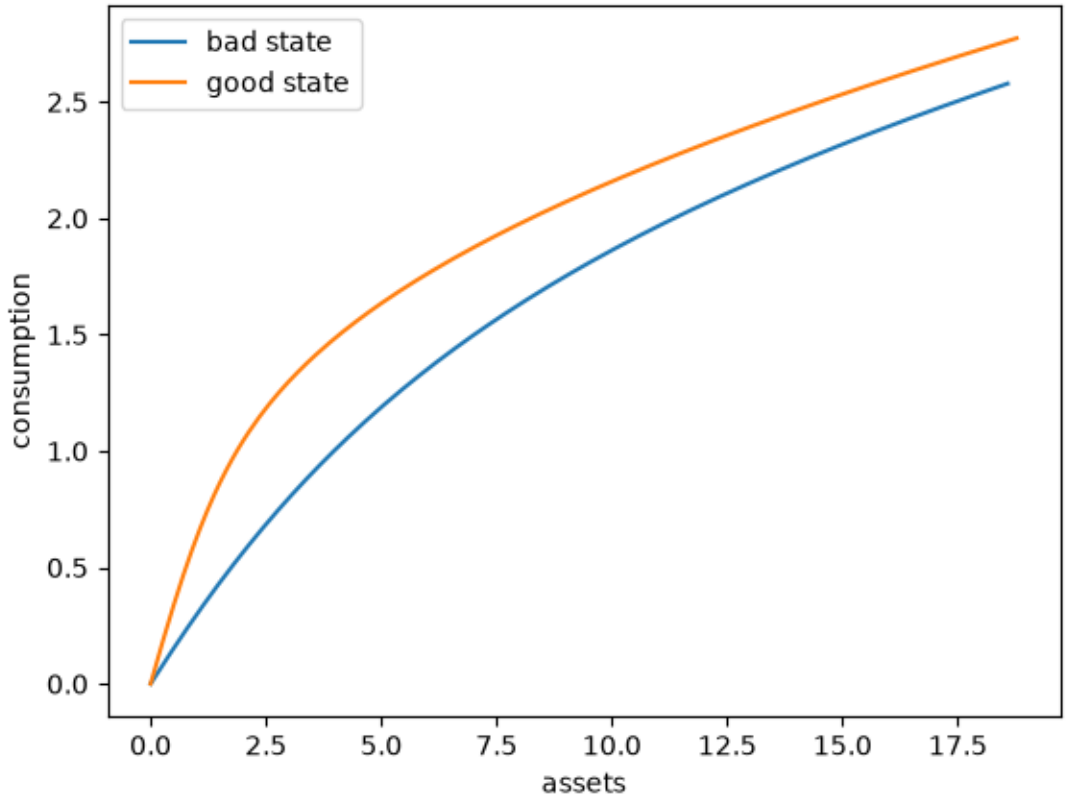
print(f"NumPy time:                 {numpy_time:.4f} seconds")
print(f"JAX time (with compile):    {jax_time_with_compile:.4f} seconds")
print(f"JAX time (without compile): {jax_time:.4f} seconds")
print(f"Speedup (NumPy/JAX):        {numpy_time/jax_time:.2f}x")
```

```
NumPy time:                 0.0207 seconds
JAX time (with compile):    0.0103 seconds
JAX time (without compile): 0.0101 seconds
Speedup (NumPy/JAX):        2.06x
```

The JAX implementation is faster due to JIT compilation and GPU/TPU acceleration (if available).

Here's a plot of the optimal policy for each z state

```
fig, ax = plt.subplots()
ax.plot(a_vec[:, 0], c_vec[:, 0], label='bad state')
ax.plot(a_vec[:, 1], c_vec[:, 1], label='good state')
ax.set(xlabel='assets', ylabel='consumption')
ax.legend()
plt.show()
```



74.5.5 Dynamics

To begin to understand the long run asset levels held by households under the default parameters, let's look at the 45 degree diagram showing the law of motion for assets under the optimal consumption policy.

```
fig, ax = plt.subplots()

y = lambda z: jnp.exp(z)

def y_bar(k):
    """
    Taking z = z_grid[k], compute

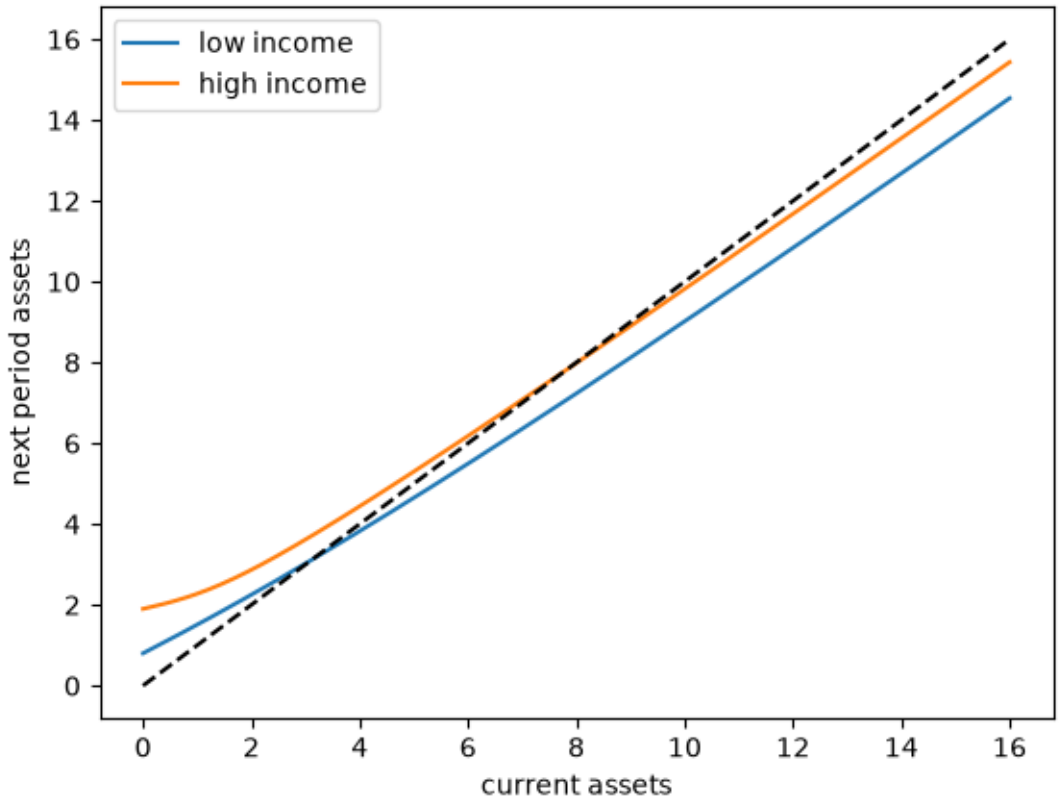
            E_z Y' = Σ_{z'} y(z') Π[z, z']

    This is the expectation of Y_{t+1} given Z_t = z.
    """
    # Compute y(z') for all z'
    y_values = jax.vmap(y)(z_grid)
    # Weight by transition probabilities and sum
    return jnp.sum(y_values * Π[k, :])

for k, label in zip((0, 1), ('low income', 'high income')):
    # Interpolate consumption policy on the savings grid
    c_on_grid = jnp.interp(s, a_vec[:, k], c_vec[:, k])
    ax.plot(s, R * (s - c_on_grid) + y_bar(k) , label=label)

ax.plot(s, s, 'k--')
ax.set(xlabel='current assets', ylabel='next period assets')

ax.legend()
plt.show()
```



The unbroken lines show the update function for assets at each z, which is

$$a \mapsto R(a - \sigma^*(a, z)) + \bar{y}(z)$$

where

$$\bar{y}(z) := \sum_{z'} y(z')\Pi(z, z')$$

is the expected labor income conditional on current state z.

The dashed line is the 45 degree line.

The figure suggests that, on average, the dynamics will be stable — assets do not diverge even in the highest state.

This turns out to be true: there is a unique stationary distribution of assets.

- For details see [Ma *et al.*, 2020]

This stationary distribution represents the long run dispersion of assets across households when households have idiosyncratic shocks.

74.5.6 A Sanity Check

One way to check our results is to

- set labor income to zero in each state and
- set the gross interest rate R to unity.

In this case, our income fluctuation problem is just a CRRA cake eating problem.

Then the value function and optimal consumption policy are given by

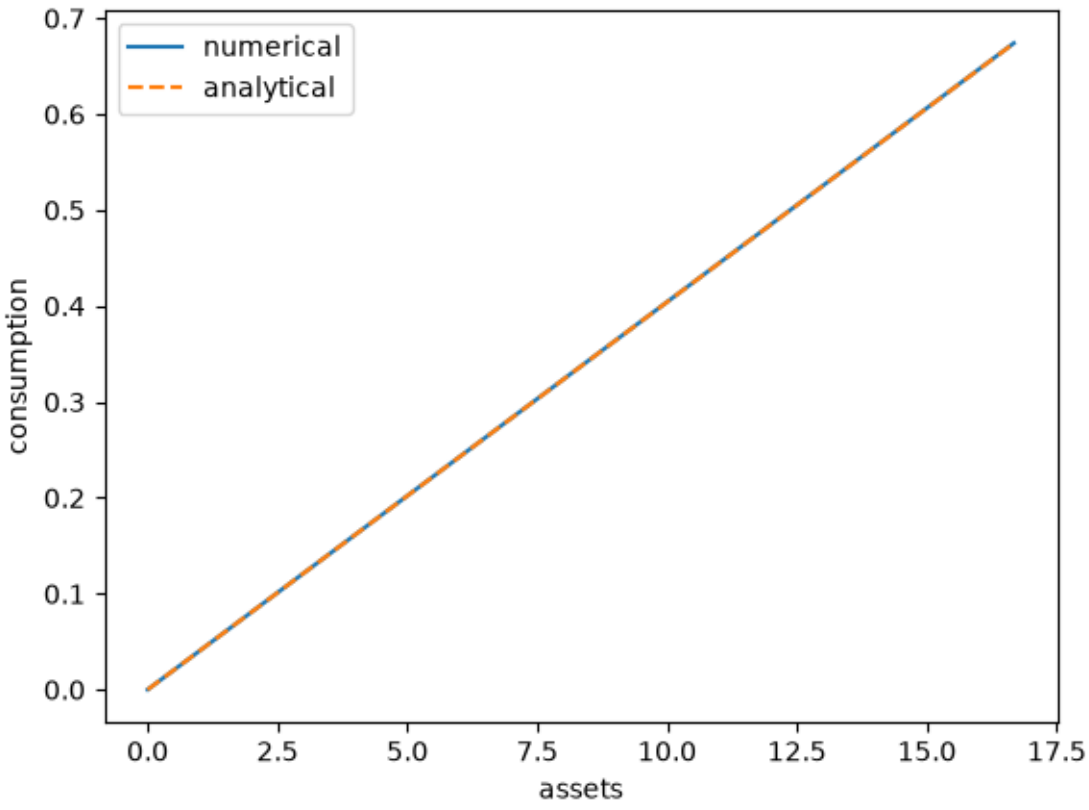
```
def c_star(x, β, γ):
    return (1 - β ** (1/γ)) * x


def v_star(x, β, γ):
    return (1 - β**(1 / γ))**(-γ) * (x**(1-γ) / (1-γ))
```

Let's see if we match up:

```
ifp_cake_eating = create_ifp(r=0.0, z_grid=(-jnp.inf, -jnp.inf))
R, β, γ, Π, z_grid, s = ifp_cake_eating
a_init = s[:, None] * jnp.ones(len(z_grid))
c_init = a_init
c_vec, a_vec = solve_model(ifp_cake_eating, c_init, a_init)

fig, ax = plt.subplots()
ax.plot(a_vec[:, 0], c_vec[:, 0], label='numerical')
ax.plot(a_vec[:, 0],
        c_star(a_vec[:, 0], ifp_cake_eating.β, ifp_cake_eating.γ),
        '--', label='analytical')
ax.set(xlabel='assets', ylabel='consumption')
ax.legend()
plt.show()
```



This looks pretty good.

74.6 Simulation

Let's return to the default model and study the stationary distribution of assets.

Our plan is to run a large number of households forward for T periods and then histogram the cross-sectional distribution of assets.

Set `num_households=50_000, T=500`.

First we write a function to run a single household forward in time and record the final value of assets.

The function takes a solution pair `c_vec` and `a_vec`, understanding them as representing an optimal policy associated with a given model `ifp`

```
@jax.jit
def simulate_household(
        key, a_0, z_idx_0, c_vec, a_vec, ifp, T
    ):
    """
    Simulates a single household for T periods to approximate the stationary
    distribution of assets.

    - key is the state of the random number generator
    - ifp is an instance of IFP
    - c_vec, a_vec are the optimal consumption policy, endogenous grid for ifp
```

(continues on next page)

(continued from previous page)

```
    """
    R, β, γ, Π, z_grid, s = ifp
    n_z = len(z_grid)

    y = lambda z: jnp.exp(z)
    σ = lambda a, z_idx: jnp.interp(a, a_vec[:, z_idx], c_vec[:, z_idx])

    # Simulate forward T periods
    def update(t, state):
        a, z_idx = state
        # Draw next shock z' from Π[z, z']
        current_key = jax.random.fold_in(key, t)
        z_next_idx = jax.random.choice(current_key, n_z, p=Π[z_idx]).astype(jnp.int32)
        z_next = z_grid[z_next_idx]
        # Update assets: a' = R * (a - c) + Y'
        a_next = R * (a - σ(a, z_idx)) + y(z_next)
        # Return updated state
        return a_next, z_next_idx

    initial_state = a_0, z_idx_0
    final_state = jax.lax.fori_loop(0, T, update, initial_state)
    a_final, _ = final_state
    return a_final
```

Now we write a function to simulate many households in parallel.

```
def compute_asset_stationary(
        c_vec, a_vec, ifp, num_households=50_000, T=500, seed=1234
    ):
    """
    Simulates num_households households for T periods to approximate
    the stationary distribution of assets.

    Returns the final cross-section of asset holdings.

    - ifp is an instance of IFP
    - c_vec, a_vec are the optimal consumption policy and endogenous grid.

    """
    R, β, γ, Π, z_grid, s = ifp
    n_z = len(z_grid)

    # Create interpolation function for consumption policy
    # Interpolate on the endogenous grid
    σ = lambda a, z_idx: jnp.interp(a, a_vec[:, z_idx], c_vec[:, z_idx])

    # Start with assets = savings_grid_max / 2
    a_0_vector = jnp.full(num_households, s[-1] / 2)
    # Initialize the exogenous state of each household
    z_idx_0_vector = jnp.zeros(num_households).astype(jnp.int32)

    # Vectorize over many households
    key = jax.random.key(seed)
    keys = jax.random.split(key, num_households)
    # Vectorize simulate_household in (key, a_0, z_idx_0)
    sim_all_households = jax.vmap(
        simulate_household, in_axes=(0, 0, 0, None, None, None, None)
```

(continues on next page)

(continued from previous page)

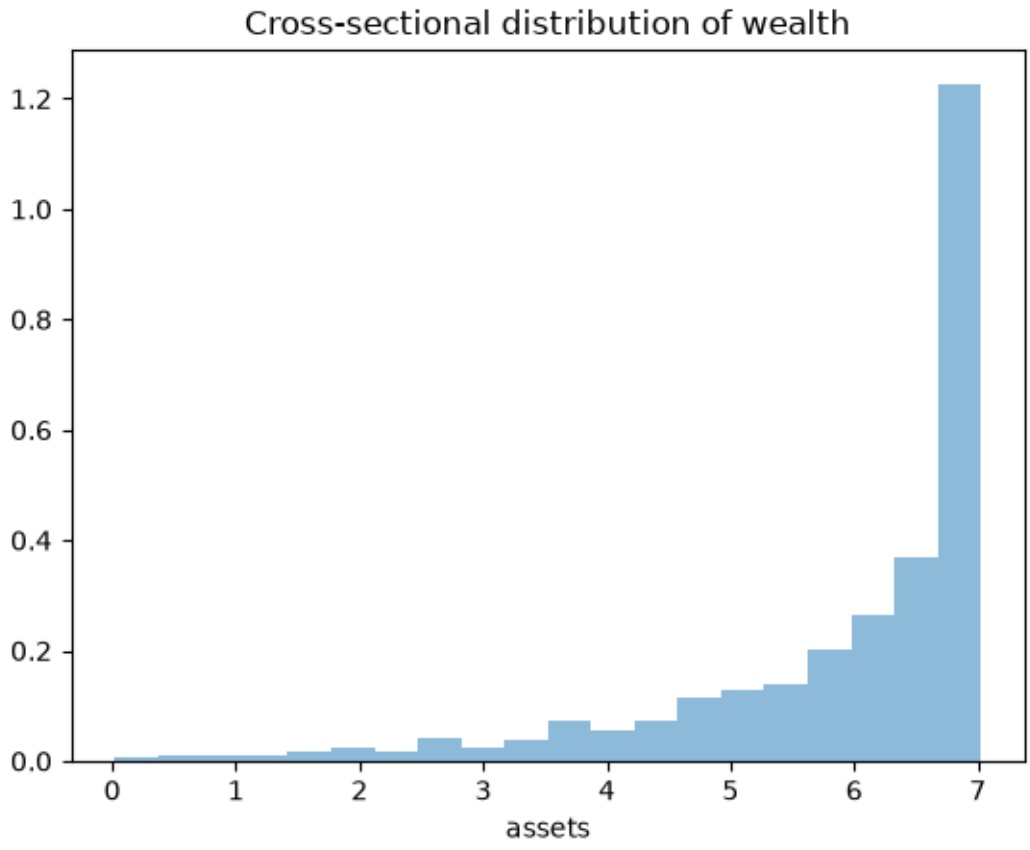
```
    )
    assets = sim_all_households(keys, a_0_vector, z_idx_0_vector, c_vec, a_vec, ifp,␣
↪T)

    return np.array(assets)
```

Now we call the function, generate the asset distribution and histogram it:

```
ifp = create_ifp()
R, β, γ, Π, z_grid, s = ifp
a_init = s[:, None] * jnp.ones(len(z_grid))
c_init = a_init
c_vec, a_vec = solve_model(ifp, c_init, a_init)
assets = compute_asset_stationary(c_vec, a_vec, ifp)

fig, ax = plt.subplots()
ax.hist(assets, bins=20, alpha=0.5, density=True)
ax.set(xlabel='assets', title="Cross-sectional distribution of wealth")
plt.show()
```



The wealth distribution looks very different to a typical wealth distribution in the data.

For one thing it is left-skewed rather than right-skewed.

In fact there is essentially no right-hand tail, even though real-world wealth distributions have long right-hand tails.

We'll do our best to fix these issues in the next few lectures.

CHAPTER

SEVENTYFIVE

THE INCOME FLUCTUATION PROBLEM IV: TRANSIENT INCOME SHOCKS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *The Income Fluctuation Problem IV: Transient Income Shocks*
  - *Overview*
  - *The Household Problem*
  - *NumPy Implementation*
  - *JAX Implementation*
  - *Simulation*
  - *Wealth Inequality*
  - *Exercises*

75.1 Overview

In this lecture we continue extend the IFP from *The Income Fluctuation Problem III: The Endogenous Grid Method* by adding transient shocks to the income process.

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

We'll also need the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
import numba
from quantecon import MarkovChain
import jax
import jax.numpy as jnp
from typing import NamedTuple
```

75.2 The Household Problem

We briefly outline the model and then discuss how to solve it.

Readers seeking a more extensive discussion of the model and the EGM solution method can review *The Income Fluctuation Problem III: The Endogenous Grid Method*.

75.2.1 Set-Up

A household chooses a state-contingent consumption plan $\{c_t\}_{t \geq 0}$ to maximize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$a_{t+1} = R(a_t - c_t) + Y_{t+1} \quad c_t \geq 0, \quad a_t \geq 0 \quad t = 0, 1, \ldots \tag{75.1}$$

The definitions of symbols and the timing are the same as in *The Income Fluctuation Problem III: The Endogenous Grid Method*.

Now, non-capital income Y_t is given by $Y_t = y(Z_t, \eta_t)$, where

- $\{Z_t\}$ is an exogenous state process (persistent component),
- $\{\eta_t\}$ is an IID shock process, and
- y is a function taking values in $\mathbb{R}_+$.

Throughout this lecture, we assume that $\eta_t \sim N(0, 1)$.

We again take $\{Z_t\}$ to be a finite state Markov chain taking values in Z with Markov matrix Π.

The shock process $\{\eta_t\}$ is independent of $\{Z_t\}$ and represents transient income fluctuations.

In addition to previous assumptions, we suppose that $y(z, \eta) = \exp(a_y \eta + z b_y)$ where a_y, b_y are positive constants

The asset space and state space are unchanged, as is the definition of an optimal path.

The functional Euler equation has the form

$$(u' \circ \sigma)(a, z) = \beta R \sum_{z'} \int (u' \circ \sigma)[R(a - \sigma(a, z)) + y(z', \eta'),\, z'] \phi(\eta') d\eta' \, \Pi(z, z') \tag{75.2}$$

Here

- $(u' \circ \sigma)(s) := u'(\sigma(s))$,
- primes indicate next period states (as well as derivatives),
- ϕ is the density of the shock η_t (standard normal), and

- σ is the unknown function.

The equality (75.2) holds at all interior choices, meaning $\sigma(a, z) < a$.

We aim to find a fixed point σ of (75.2).

To do so we use the EGM.

Below we use the relationships $a_t = c_t + s_t$ and $a_{t+1} = Rs_t + Y_{t+1}$.

We begin with an exogenous savings grid $s_0 < s_1 < \cdots < s_m$ with $s_0 = 0$.

We fix a current guess of the policy function σ.

For each exogenous savings level s_i with $i \geq 1$ and current state z_j, we set

$$c_{ij} := (u')^{-1} \left[\beta R \sum_{z'} \int u'[\sigma(Rs_i + y(z', \eta'), z')] \phi(\eta') d\eta' \, \Pi(z_j, z') \right] \tag{75.3}$$

The Euler equation holds here because $i \geq 1$ implies $s_i > 0$ and hence consumption is interior.

For the boundary case $s_0 = 0$ we set

$$c_{0j} := 0 \quad \text{for all j}$$

We then obtain a corresponding endogenous grid of current assets via

$$a_{ij} := c_{ij} + s_i.$$

Our next guess of the policy function, which we write as $K\sigma$, is the linear interpolation of the interpolation points

$$\{(a_{0j}, c_{0j}), \ldots, (a_{mj}, c_{mj})\}$$

for each j.

75.3 NumPy Implementation

In this section we'll code up a NumPy version of the code that aims only for clarity, rather than efficiency.

Once we have it working, we'll produce a JAX version that's far more efficient and check that we obtain the same results.

We use the CRRA utility specification

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

75.3.1 Set Up

Here we build a class called `IFPNumPy` that stores the model primitives.

The exogenous state process $\{Z_t\}$ defaults to a two-state Markov chain with transition matrix Π.

```
class IFPNumPy(NamedTuple):
    R: float                  # Gross interest rate R = 1 + r
    β: float                  # Discount factor
    γ: float                  # Preference parameter
    Π: np.ndarray             # Markov matrix for exogenous shock
```

(continues on next page)

(continued from previous page)

```
    z_grid: np.ndarray        # Markov state values for Z_t
    s: np.ndarray             # Exogenous savings grid
    a_y: float                # Scale parameter for Y_t
    b_y: float                # Additive parameter for Y_t
    η_draws: np.ndarray       # Draws of innovation η for MC


def create_ifp(r=0.01,
               β=0.96,
               γ=1.5,
               Π=((0.6, 0.4),
                  (0.05, 0.95)),
               z_grid=(-10.0, np.log(2.0)),
               savings_grid_max=16,
               savings_grid_size=50,
               a_y=0.2,
               b_y=0.5,
               shock_draw_size=100,
               seed=1234):

    np.random.seed(seed)
    s = np.linspace(0, savings_grid_max, savings_grid_size)
    Π, z_grid = np.array(Π), np.array(z_grid)
    R = 1 + r
    η_draws = np.random.randn(shock_draw_size)
    assert R * β < 1, "Stability condition violated."
    return IFPNumPy(R, β, γ, Π, z_grid, s, a_y, b_y, η_draws)
```

75.3.2 Solver

Here is the operator K that transforms current guess σ into next period guess $K\sigma$.

In practice, it takes in

- a guess of optimal consumption values c_{ij}, stored as `c_vec`
- and a corresponding set of endogenous grid points a^e_{ij}, stored as `a_vec`

These are converted into a consumption policy $a \mapsto \sigma(a, z_j)$ by linear interpolation of (a^e_{ij}, c_{ij}) over i for each j.

When we compute consumption in (75.3), we will use Monte Carlo over η', so that the expression becomes

$$c_{ij} := (u')^{-1} \left[\beta R \sum_{z'} \frac{1}{m} \sum_{\ell=1}^{m} u'[\sigma(Rs_i + y(z', \eta_\ell), z')] \, \Pi(z_j, z') \right] \tag{75.4}$$

with each η_ℓ being a standard normal draw.

```
@numba.jit
def K_numpy(
        c_in: np.ndarray,   # Initial guess of σ on grid endogenous grid
        a_in: np.ndarray,   # Initial endogenous grid
        ifp_numpy: IFPNumPy
    ) -> np.ndarray:
    """
    The Euler equation operator for the IFP model using the
    Endogenous Grid Method.
```

(continues on next page)

(continued from previous page)

```
    This operator implements one iteration of the EGM algorithm to
    update the consumption policy function.

    """
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp_numpy
    n_a = len(s)
    n_z = len(z_grid)

    # Utility functions
    def u_prime(c):
        return c**(-γ)

    def u_prime_inv(c):
        return c**(-1/γ)

    def y(z, η):
        return np.exp(a_y * η + z * b_y)

    c_out = np.zeros_like(c_in)

    for i in range(1, n_a):  # Start from 1 for positive savings levels
        for j in range(n_z):

            # Compute Σ_z' ∫ u'(σ(R s_i + y(z', η'), z')) ψ(η') dη' Π[z_j, z']
            expectation = 0.0
            for k in range(n_z):
                z_prime = z_grid[k]
                # Integrate over η draws (Monte Carlo)
                inner_sum = 0.0
                for η in η_draws:
                    # Calculate next period assets
                    next_a = R * s[i] + y(z_prime, η)
                    # Interpolate to get σ(R s_i + y(z', η), z')
                    next_c = np.interp(next_a, a_in[:, k], c_in[:, k])
                    # Add to the inner sum
                    inner_sum += u_prime(next_c)
                # Average over η draws to approximate the integral
                # ∫ u'(σ(R s_i + y(z', η'), z')) ψ(η') dη' when z' = z_grid[k]
                inner_mean_k = (inner_sum / len(η_draws))
                # Weight by transition probability and add to the expectation
                expectation += inner_mean_k * Π[j, k]

            # Calculate updated c_{ij} values
            c_out[i, j] = u_prime_inv(β * R * expectation)

    a_out = c_out + s[:, None]

    return c_out, a_out
```

To solve the model we use a simple while loop.

```
def solve_model_numpy(
        ifp_numpy: IFPNumPy,
        c_init: np.ndarray,
        a_init: np.ndarray,
        tol: float = 1e-5,
```

(continues on next page)

(continued from previous page)

```
        max_iter: int = 1_000
    ) -> np.ndarray:
    """
    Solve the model using time iteration with EGM.

    """
    c_in, a_in = c_init, a_init
    i = 0
    error = tol + 1

    while error > tol and i < max_iter:
        c_out, a_out = K_numpy(c_in, a_in, ifp_numpy)
        error = np.max(np.abs(c_out - c_in))
        i = i + 1
        c_in, a_in = c_out, a_out

    return c_out, a_out
```

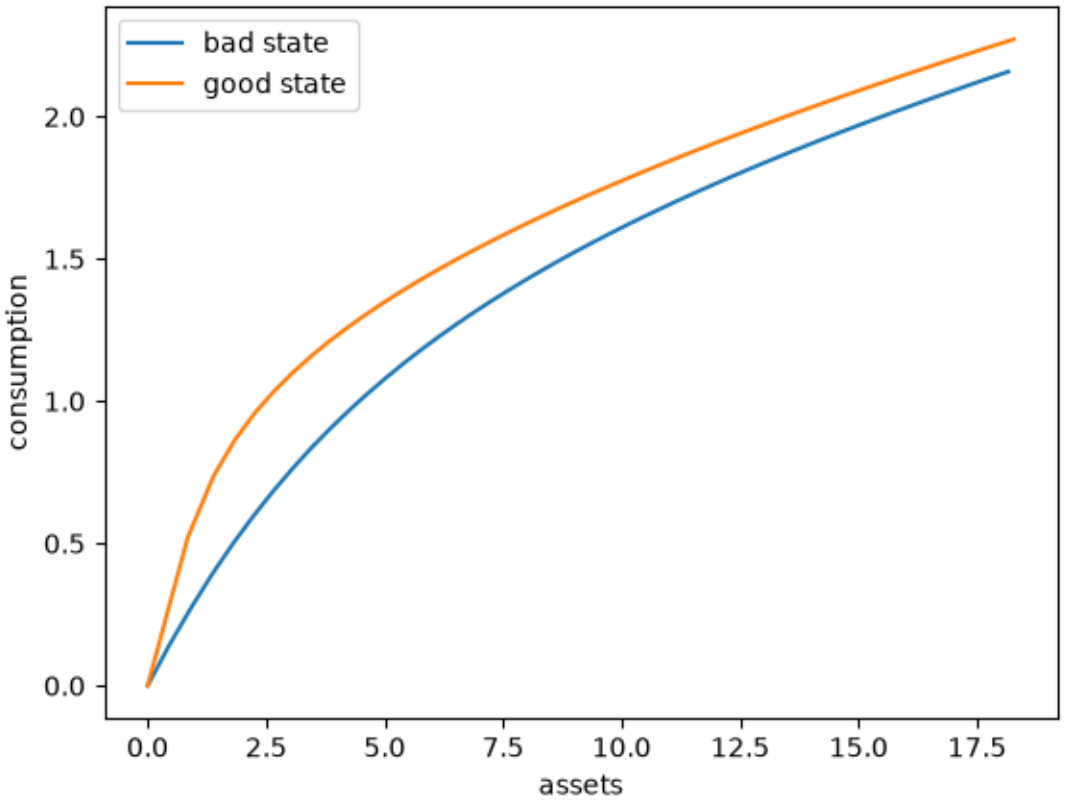
Let's road test the EGM code.

```
ifp_numpy = create_ifp()
R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp_numpy
# Initial conditions -- agent consumes everything
a_init = s[:, None] * np.ones(len(z_grid))
c_init = a_init
# Solve from these initial conditions
c_vec, a_vec = solve_model_numpy(
    ifp_numpy, c_init, a_init
)
```

Here's a plot of the optimal consumption policy for each z state

```
fig, ax = plt.subplots()

ax.plot(a_vec[:, 0], c_vec[:, 0], label='bad state')
ax.plot(a_vec[:, 1], c_vec[:, 1], label='good state')
ax.set(xlabel='assets', ylabel='consumption')
ax.legend()
plt.show()
```



75.4 JAX Implementation

Now we write a more efficient JAX version, which can run on a GPU.

75.4.1 Set Up

We start with a class called `IFP` that stores the model primitives.

```
class IFP(NamedTuple):
    R: float                  # Gross interest rate R = 1 + r
    β: float                  # Discount factor
    γ: float                  # Preference parameter
    Π: jnp.ndarray            # Markov matrix for exogenous shock
    z_grid: jnp.ndarray       # Markov state values for Z_t
    s: jnp.ndarray            # Exogenous savings grid
    a_y: float                # Scale parameter for Y_t
    b_y: float                # Additive parameter for Y_t
    η_draws: jnp.ndarray      # Draws of innovation η for MC


def create_ifp(r=0.01,
               β=0.94,
               γ=1.5,
               Π=((0.6, 0.4),
                  (0.05, 0.95)),
```

(continues on next page)

(continued from previous page)

```
                z_grid=(-10.0, jnp.log(2.0)),
                savings_grid_max=16,
                savings_grid_size=50,
                a_y=0.2,
                b_y=0.5,
                shock_draw_size=100,
                seed=1234):

    key = jax.random.key(seed)
    s = jnp.linspace(0, savings_grid_max, savings_grid_size)
    Π, z_grid = jnp.array(Π), jnp.array(z_grid)
    R = 1 + r
    η_draws = jax.random.normal(key, (shock_draw_size,))
    assert R * β < 1, "Stability condition violated."
    return IFP(R, β, γ, Π, z_grid, s, a_y, b_y, η_draws)
```

75.4.2 Solver

Here is the operator K that transforms current guess σ into next period guess $K\sigma$.

```
def K(
        c_in: jnp.ndarray,
        a_in: jnp.ndarray,
        ifp: IFP
    ) -> jnp.ndarray:
    """
    The Euler equation operator for the IFP model using the
    Endogenous Grid Method.

    This operator implements one iteration of the EGM algorithm to
    update the consumption policy function.

    """
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
    n_a = len(s)
    n_z = len(z_grid)

    # Utility functions
    def u_prime(c):
        return c**(-γ)

    def u_prime_inv(c):
        return c**(-1/γ)

    def y(z, η):
        return jnp.exp(a_y * η + z * b_y)

    def compute_c(i, j):
        " Compute c_ij when i >= 1 (interior choice). "

        def expected_mu(k):
            " Approximate ∫ u'(σ(R s_i + y(z_k, η'), z_k)) ψ(η') dη' "

            def compute_mu_at_eta(η):
                " Compute u'(σ(R * s_i + y(z_k, η), z_k)) "
```

(continues on next page)

(continued from previous page)

```
                next_a = R * s[i] + y(z_grid[k], η)
                # Interpolate to get σ(R * s_i + y(z_k, η), z_k)
                next_c = jnp.interp(next_a, a_in[:, k], c_in[:, k])
                # Return u'(σ(R * s_i + y(z_k, η), z_k))
                return u_prime(next_c)

            # Average over η draws to approximate the inner integral
            # ∫ u'(σ(R s_i + y(z_k, η'), z_k)) φ(η') dη'
            all_draws = jax.vmap(compute_mu_at_eta)(η_draws)
            return jnp.mean(all_draws)

        # Compute expectation: Σ_k [∫ u'(σ(...)) φ(η) dη] * Π[j, k]
        expectations = jax.vmap(expected_mu)(jnp.arange(n_z))
        expectation = jnp.sum(expectations * Π[j, :])
        # Invert to get consumption c_ij at (s_i, z_j)
        return u_prime_inv(β * R * expectation)

    # Set up index grids for vmap computation of all c_{ij}
    i_grid = jnp.arange(1, n_a)
    j_grid = jnp.arange(n_z)

    # vmap over j for each i
    compute_c_i = jax.vmap(compute_c, in_axes=(None, 0))
    # vmap over i
    compute_c = jax.vmap(lambda i: compute_c_i(i, j_grid))
    # Compute consumption for i >= 1
    c_out_interior = compute_c(i_grid)  # Shape: (n_a-1, n_z)
    # For i = 0, set consumption to 0
    c_out_boundary = jnp.zeros((1, n_z))

    # Concatenate boundary and interior
    c_out = jnp.concatenate([c_out_boundary, c_out_interior], axis=0)

    # Compute endogenous asset grid: a^e_{ij} = c_{ij} + s_i
    a_out = c_out + s[:, None]

    return c_out, a_out
```

Here's a jit-accelerated iterative routine to solve the model using this operator.

```
@jax.jit
def solve_model(
        ifp: IFP,
        c_init: jnp.ndarray,  # Initial guess of σ on grid endogenous grid
        a_init: jnp.ndarray,  # Initial endogenous grid
        tol: float = 1e-5,
        max_iter: int = 1000
    ) -> jnp.ndarray:
    """
    Solve the model using time iteration with EGM.

    """

    def condition(loop_state):
        c_in, a_in, i, error = loop_state
        return (error > tol) & (i < max_iter)
```

(continues on next page)

(continued from previous page)

```
    def body(loop_state):
        c_in, a_in, i, error = loop_state
        c_out, a_out = K(c_in, a_in, ifp)
        error = jnp.max(jnp.abs(c_out - c_in))
        i += 1
        return c_out, a_out, i, error

    i, error = 0, tol + 1
    initial_state = (c_init, a_init, i, error)
    final_loop_state = jax.lax.while_loop(condition, body, initial_state)
    c_out, a_out, i, error = final_loop_state

    return c_out, a_out
```

75.4.3 Test run

Let's road test the EGM code.

```
ifp = create_ifp()
R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
# Set initial conditions where the agent consumes everything
a_init = s[:, None] * jnp.ones(len(z_grid))
c_init = a_init
# Solve starting from these initial conditions
c_vec_jax, a_vec_jax = solve_model(ifp, c_init, a_init)
```

To verify the correctness of our JAX implementation, let's compare it with the NumPy version we developed earlier.

```
# Compare the results
max_c_diff = np.max(np.abs(np.array(c_vec) - c_vec_jax))
max_ae_diff = np.max(np.abs(np.array(a_vec) - a_vec_jax))

print(f"Maximum difference in consumption policy: {max_c_diff:.2e}")
print(f"Maximum difference in asset grid:        {max_ae_diff:.2e}")
```

```
Maximum difference in consumption policy: 3.94e-01
Maximum difference in asset grid:        3.94e-01
```

These numbers confirm that we are computing essentially the same policy using the two approaches.

(Remaining differences are mainly due to different Monte Carlo integration outcomes over relatively small samples.)

75.4.4 Timing

Now let's compare the execution time between NumPy and JAX implementations.

```
import time

# Set up initial conditions for NumPy version
s_np = np.array(s)
z_grid_np = np.array(z_grid)
a_init_np = s_np[:, None] * np.ones(len(z_grid_np))
c_init_np = a_init_np.copy()
```

(continues on next page)

(continued from previous page)

```
# Set up initial conditions for JAX version
a_init_jx = s[:, None] * jnp.ones(len(z_grid))
c_init_jx = a_init_jx

# Time NumPy version
start = time.time()
c_vec_np, a_vec_np = solve_model_numpy(ifp_numpy, c_init_np, a_init_np)
numpy_time = time.time() - start

# Time JAX version (with compilation)
start = time.time()
c_vec_jx, a_vec_jx = solve_model(ifp, c_init_jx, a_init_jx)
c_vec_jx.block_until_ready()
jax_time_with_compile = time.time() - start

# Time JAX version (without compilation - second run)
start = time.time()
c_vec_jx, a_vec_jx = solve_model(ifp, c_init_jx, a_init_jx)
c_vec_jx.block_until_ready()
jax_time = time.time() - start

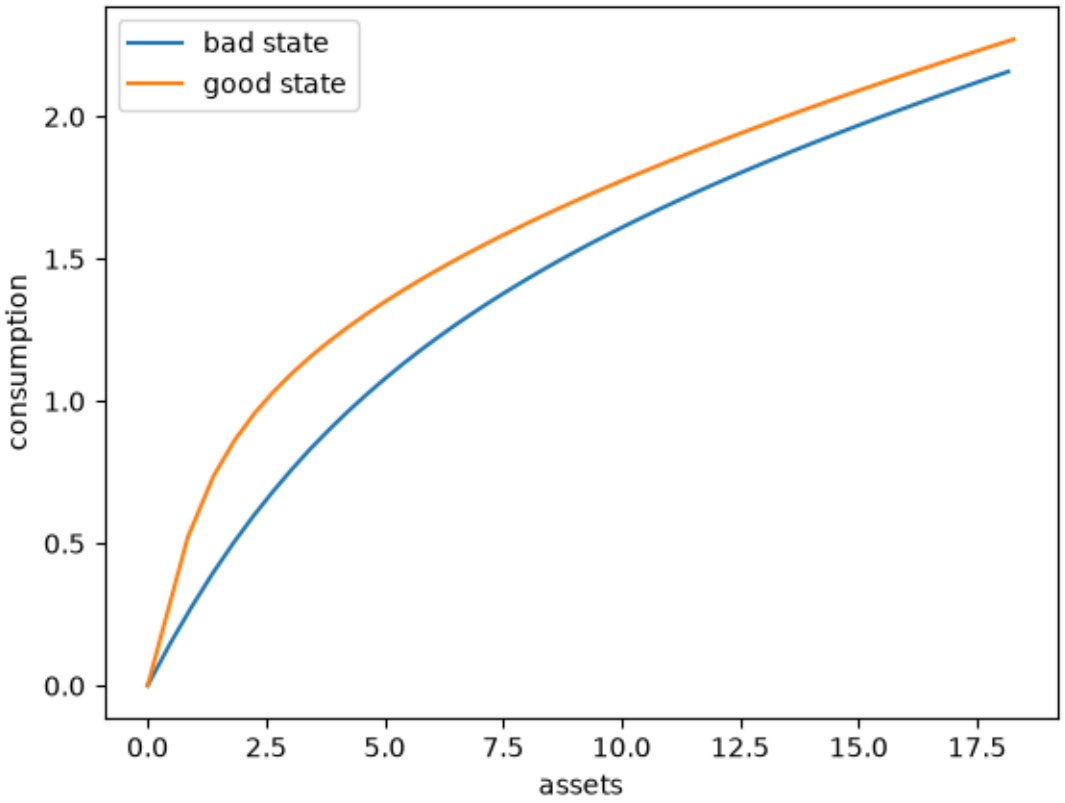
print(f"NumPy time:                 {numpy_time:.4f} seconds")
print(f"JAX time (with compile):    {jax_time_with_compile:.4f} seconds")
print(f"JAX time (without compile): {jax_time:.4f} seconds")
print(f"Speedup (NumPy/JAX):        {numpy_time/jax_time:.2f}x")
```

```
NumPy time:                 0.3855 seconds
JAX time (with compile):    0.0097 seconds
JAX time (without compile): 0.0092 seconds
Speedup (NumPy/JAX):        41.84x
```

The JAX implementation is significantly faster due to JIT compilation and GPU/TPU acceleration (if available).

Here's a plot of the optimal policy for each z state

```
fig, ax = plt.subplots()
ax.plot(a_vec[:, 0], c_vec[:, 0], label='bad state')
ax.plot(a_vec[:, 1], c_vec[:, 1], label='good state')
ax.set(xlabel='assets', ylabel='consumption')
ax.legend()
plt.show()
```



75.4.5 Dynamics

To begin to understand the long run asset levels held by households under the default parameters, let's look at the 45 degree diagram showing the law of motion for assets under the optimal consumption policy.

```
fig, ax = plt.subplots()

def y(z, η):
    return jnp.exp(a_y * η + z * b_y)

def y_bar(k):
    """
    Taking z = z_grid[k], compute an approximation to

            E_z Y' = Σ_{z'} ∫ y(z', η') ψ(η') dη' Π[z, z']

    This is the expectation of Y_{t+1} given Z_t = z.
    """
    # Approximate ∫ y(z', η') ψ(η') dη' at given z'
    def mean_y_at_z(z_prime):
        return jnp.mean(y(z_prime, η_draws))
    # Evaluate this integral across all z'
    y_means = jax.vmap(mean_y_at_z)(z_grid)
    # Weight by transition probabilities and sum
    return jnp.sum(y_means * Π[k, :])

for k, label in zip((0, 1), ('low income', 'high income')):
```

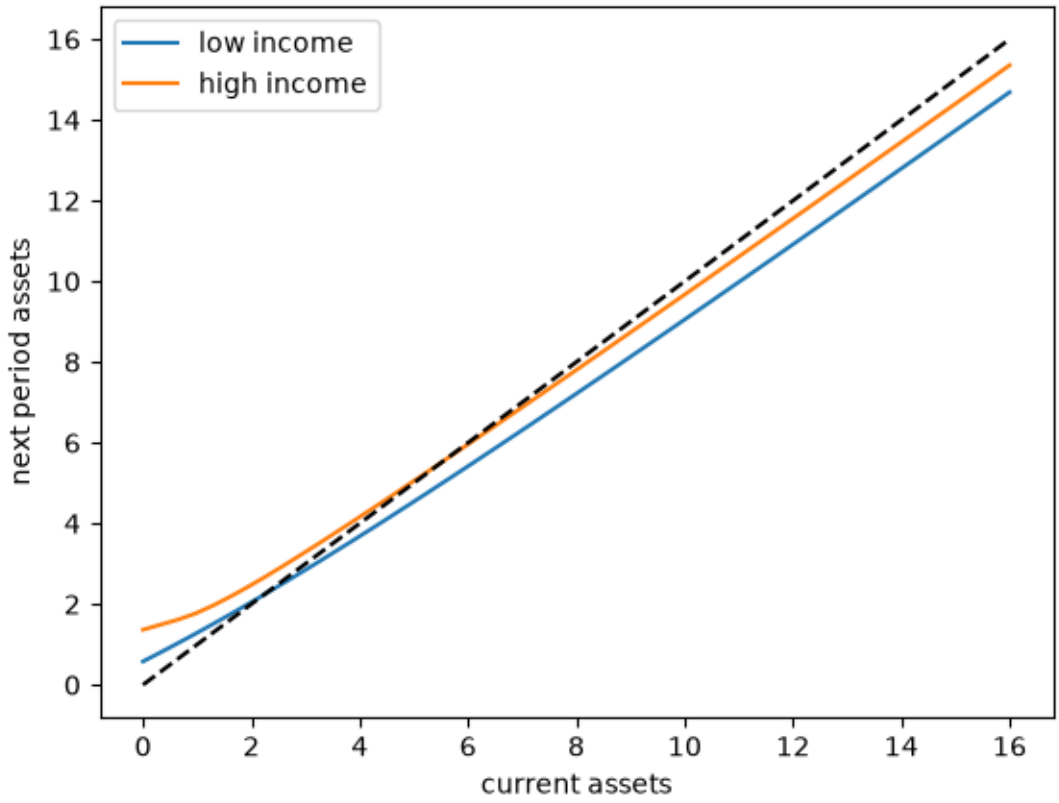
(continues on next page)

(continued from previous page)

```
    # Interpolate consumption policy on the savings grid
    c_on_grid = jnp.interp(s, a_vec[:, k], c_vec[:, k])
    ax.plot(s, R * (s - c_on_grid) + y_bar(k) , label=label)

ax.plot(s, s, 'k--')
ax.set(xlabel='current assets', ylabel='next period assets')

ax.legend()
plt.show()
```



The unbroken lines show the update function for assets at each z, which is

$$a \mapsto R(a - \sigma^*(a, z)) + \bar{y}(z)$$

where

$$\bar{y}(z) := \sum_{z'} \frac{1}{m} \sum_{\ell=1}^{m} y(z', \eta_\ell) \Pi(z, z')$$

is a Monte Carlo approximation to expected labor income conditional on current state z.

The dashed line is the 45 degree line.

The figure suggests that, on average, the dynamics will be stable — assets do not diverge even in the highest state.

This turns out to be true: there is a unique stationary distribution of assets.

- For details see [Ma *et al.*, 2020]

This stationary distribution represents the long run dispersion of assets across households when households have idiosyncratic shocks.

75.5 Simulation

Let's return to the default model and study the stationary distribution of assets.

Our plan is to run a large number of households forward for T periods and then histogram the cross-sectional distribution of assets.

Set `num_households=50_000, T=500`.

First we write a function to run a single household forward in time and record the final value of assets.

The function takes a solution pair `c_vec` and `a_vec`, understanding them as representing an optimal policy associated with a given model `ifp`

```
@jax.jit
def simulate_household(
        key, a_0, z_idx_0, c_vec, a_vec, ifp, T
    ):
    """
    Simulates a single household for T periods to approximate the stationary
    distribution of assets.

    - key is the state of the random number generator
    - ifp is an instance of IFP
    - c_vec, a_vec are the optimal consumption policy, endogenous grid for ifp

    """
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
    n_z = len(z_grid)

    def y(z, η):
        return jnp.exp(a_y * η + z * b_y)

    # Create interpolation function for consumption policy
    σ = lambda a, z_idx: jnp.interp(a, a_vec[:, z_idx], c_vec[:, z_idx])

    # Simulate forward T periods
    def update(t, state):
        a, z_idx = state
        # Draw next shock z' from Π[z, z']
        current_key = jax.random.fold_in(key, 2*t)
        z_next_idx = jax.random.choice(current_key, n_z, p=Π[z_idx]).astype(jnp.int32)
        z_next = z_grid[z_next_idx]
        # Draw η shock
        η_key = jax.random.fold_in(key, 2*t + 1)
        η = jax.random.normal(η_key)
        # Update assets: a' = R * (a - c) + Y'
        a_next = R * (a - σ(a, z_idx)) + y(z_next, η)
        # Return updated state
        return a_next, z_next_idx

    initial_state = a_0, z_idx_0
    final_state = jax.lax.fori_loop(0, T, update, initial_state)
    a_final, _ = final_state
    return a_final
```

Now we write a function to simulate many households in parallel.

```
def compute_asset_stationary(
        c_vec, a_vec, ifp, num_households=50_000, T=500, seed=1234
    ):
    """
    Simulates num_households households for T periods to approximate
    the stationary distribution of assets.

    Returns the final cross-section of asset holdings.

    - ifp is an instance of IFP
    - c_vec, a_vec are the optimal consumption policy and endogenous grid.

    """
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
    n_z = len(z_grid)

    # Create interpolation function for consumption policy
    # Interpolate on the endogenous grid
    σ = lambda a, z_idx: jnp.interp(a, a_vec[:, z_idx], c_vec[:, z_idx])

    # Start with assets = savings_grid_max / 2
    a_0_vector = jnp.full(num_households, s[-1] / 2)
    # Initialize the exogenous state of each household
    z_idx_0_vector = jnp.zeros(num_households).astype(jnp.int32)

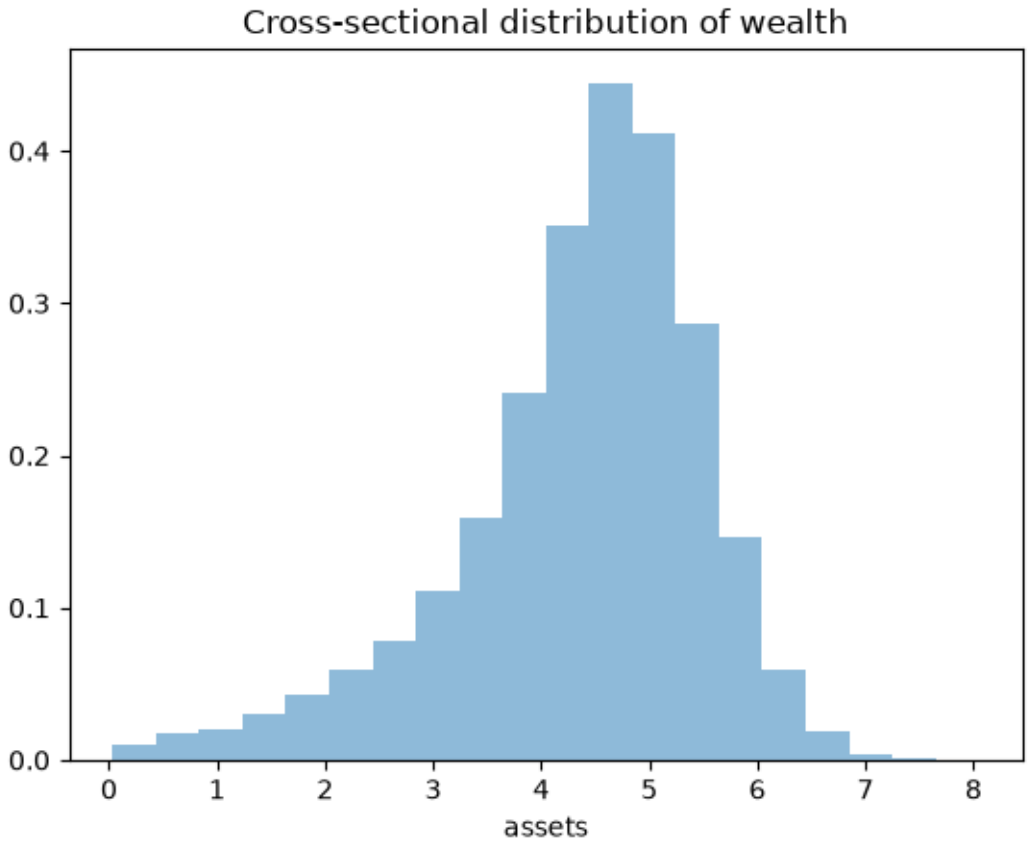
    # Vectorize over many households
    key = jax.random.key(seed)
    keys = jax.random.split(key, num_households)
    # Vectorize simulate_household in (key, a_0, z_idx_0)
    sim_all_households = jax.vmap(
        simulate_household, in_axes=(0, 0, 0, None, None, None, None)
    )
    assets = sim_all_households(keys, a_0_vector, z_idx_0_vector, c_vec, a_vec, ifp,␣
↪T)

    return np.array(assets)
```

Now we call the function, generate the asset distribution and histogram it:

```
ifp = create_ifp()
R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
a_init = s[:, None] * jnp.ones(len(z_grid))
c_init = a_init
c_vec, a_vec = solve_model(ifp, c_init, a_init)
assets = compute_asset_stationary(c_vec, a_vec, ifp)

fig, ax = plt.subplots()
ax.hist(assets, bins=20, alpha=0.5, density=True)
ax.set(xlabel='assets', title="Cross-sectional distribution of wealth")
plt.show()
```



As was the case in *The Income Fluctuation Problem III: The Endogenous Grid Method*, the wealth distribution looks implausible.

While we have at least gained a nontrivial right tail, we still have a left skew.

75.6 Wealth Inequality

Lets' look at wealth inequality by computing some standard measures of this phenomenon.

We will also examine how inequality varies with the interest rate.

75.6.1 Measuring Inequality

We'll compute two common measures of wealth inequality:

1. **Gini coefficient**: A measure of inequality ranging from 0 (perfect equality) to 1 (perfect inequality)
2. **Top 1% wealth share**: The fraction of total wealth held by the richest 1% of households

Here are functions to compute these measures:

```
def gini_coefficient(x):
    """
    Compute the Gini coefficient for array x.

    """
```

(continues on next page)

(continued from previous page)

```
    x = jnp.asarray(x)
    n = len(x)
    x_sorted = jnp.sort(x)
    # Compute Gini coefficient
    cumsum = jnp.cumsum(x_sorted)
    a = (2 * jnp.sum((jnp.arange(1, n+1)) * x_sorted)) / (n * cumsum[-1])
    return a - (n + 1) / n


def top_share(
        x: jnp.array,    # array of wealth values
        p: float=0.01    # fraction of top households (default 0.01 for top 1%)
    ):
    """
    Compute the share of total wealth held by the top p fraction of households.

    """
    x = jnp.asarray(x)
    x_sorted = jnp.sort(x)
    # Number of households in top p%
    n_top = int(jnp.ceil(len(x) * p))
    # Wealth held by top p%
    wealth_top = jnp.sum(x_sorted[-n_top:])
    # Total wealth
    wealth_total = jnp.sum(x_sorted)
    return wealth_top / wealth_total
```

Let's compute these measures for our baseline simulation:

```
gini = gini_coefficient(assets)
top1 = top_share(assets, p=0.01)

print(f"Gini coefficient: {gini:.4f}")
print(f"Top 1% wealth share: {top1:.4f}")
```

```
Gini coefficient: 0.1428
Top 1% wealth share: 0.0154
```

These numbers are a long way out, at least for a country such as the US!

Recent numbers suggest that

- the Gini coefficient for wealth in the US is around 0.8
- the top 1% wealth share is over 0.3

Of course we have not made much effort to accurately estimate or calibrate our parameters.

But actually the cause is deeper — a model with this structure will always struggle to replicate the observed wealth distribution.

In a *later lecture* we'll see if we can improve on these numbers.

75.6.2 Interest Rate and Inequality

Let's examine how wealth inequality varies with the interest rate r.

We conjecture that higher interest rates will increase wealth inequality, as wealthier households benefit more from returns on their assets.

Let's investigate empirically:

```
# Test over 8 interest rate values
M = 8
r_vals = np.linspace(0, 0.05, M)

gini_vals = []
top1_vals = []

# Solve and simulate for each r
for r in r_vals:
    print(f'Analyzing inequality at r = {r:.4f}')
    ifp = create_ifp(r=r)
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
    a_init = s[:, None] * jnp.ones(len(z_grid))
    c_init = a_init
    c_vec, a_vec = solve_model(ifp, c_init, a_init)
    assets = compute_asset_stationary(
        c_vec, a_vec, ifp, num_households=50_000, T=500
    )
    gini = gini_coefficient(assets)
    top1 = top_share(assets, p=0.01)
    gini_vals.append(gini)
    top1_vals.append(top1)
    # Use last solution as initial conditions for the policy solver
    c_init = c_vec
    a_init = a_vec
```

```
Analyzing inequality at r = 0.0000
Analyzing inequality at r = 0.0071
Analyzing inequality at r = 0.0143
Analyzing inequality at r = 0.0214
Analyzing inequality at r = 0.0286
Analyzing inequality at r = 0.0357
Analyzing inequality at r = 0.0429
Analyzing inequality at r = 0.0500
```

Now let's visualize the results:

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4))

# Plot Gini coefficient vs interest rate
axes[0].plot(r_vals, gini_vals, 'o-')
axes[0].set_xlabel('interest rate $r$')
axes[0].set_ylabel('Gini coefficient')
axes[0].set_title('Wealth Inequality vs Interest Rate')
axes[0].grid(alpha=0.3)

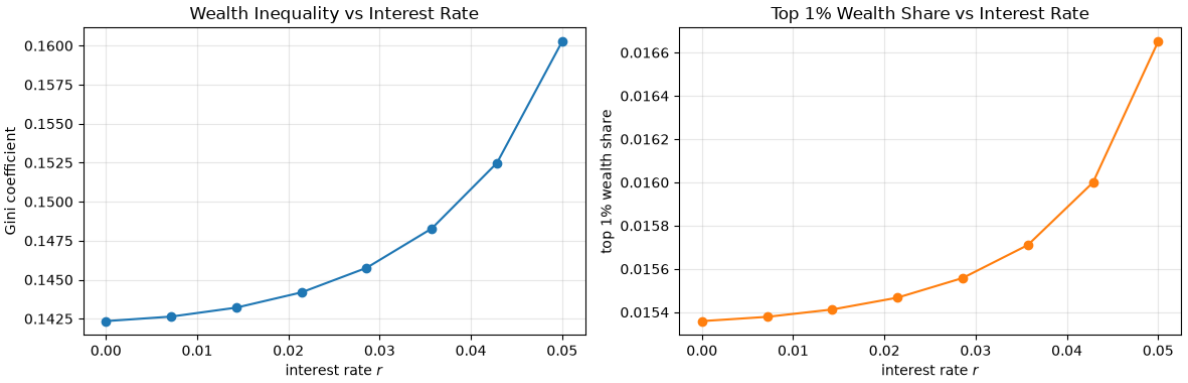
# Plot top 1% share vs interest rate
axes[1].plot(r_vals, top1_vals, 'o-', color='C1')
axes[1].set_xlabel('interest rate $r$')
```

(continues on next page)

(continued from previous page)

```
axes[1].set_ylabel('top 1% wealth share')
axes[1].set_title('Top 1% Wealth Share vs Interest Rate')
axes[1].grid(alpha=0.3)

plt.tight_layout()
plt.show()
```



The results show that these two inequality measures increase with the interest rate.

However the differences are minor and we cannot increase r much more without violating the stability constraint.

Certainly changing the interest rate cannot produce the kinds of numbers that we see in the data.

75.7 Exercises

Exercise 75.7.1

Let's consider how the interest rate affects consumption.

- Step `r` through `np.linspace(0, 0.016, 4)`.
- Other than `r`, hold all parameters at their default values.
- Plot consumption against assets for income shock fixed at the smallest value.

Your figure should show that, for this model, higher interest rates suppress consumption (because they encourage more savings).

Solution

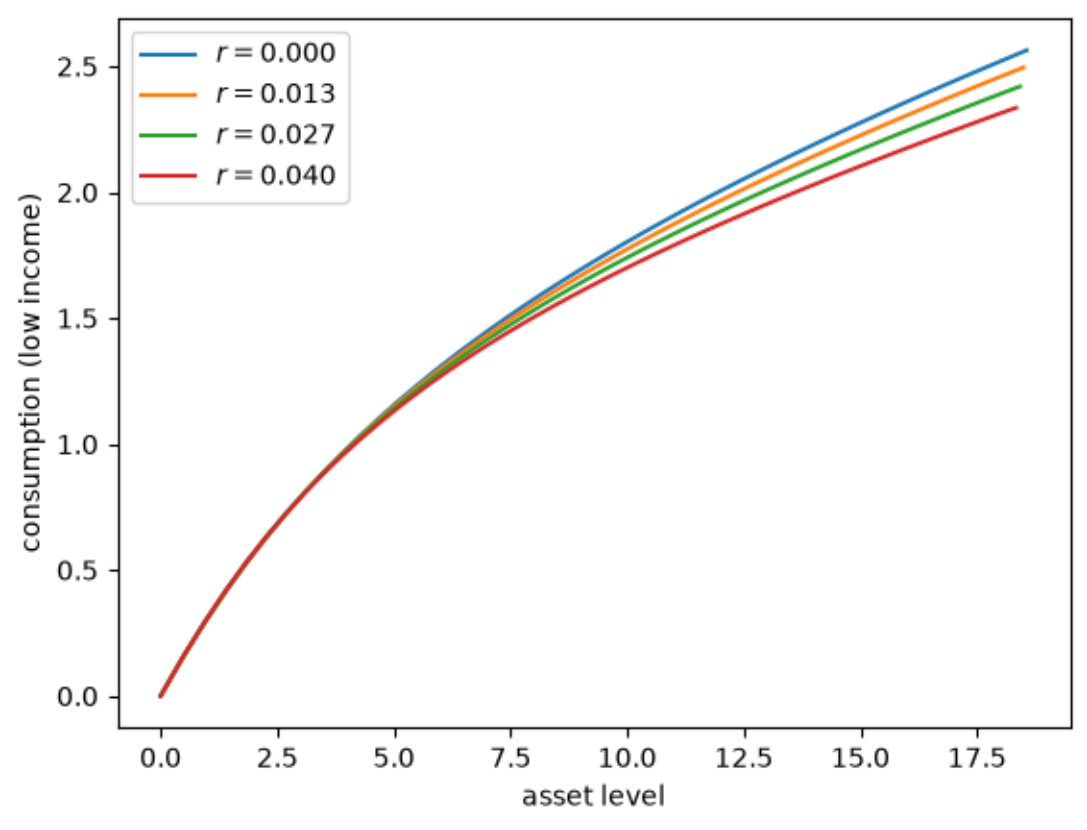
Here's one solution:

```
# With β=0.96, we need R*β < 1, so r < 0.0416
r_vals = np.linspace(0, 0.04, 4)

fig, ax = plt.subplots()
for r_val in r_vals:
    ifp = create_ifp(r=r_val)
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
    a_init = s[:, None] * jnp.ones(len(z_grid))
    c_init = a_init
```

```
    c_vec, a_vec = solve_model(ifp, c_init, a_init)
    # Plot policy
    ax.plot(a_vec[:, 0], c_vec[:, 0], label=f'$r = {r_val:.3f}$')
    # Start next round with last solution
    c_init = c_vec
    a_init = a_vec

ax.set(xlabel='asset level', ylabel='consumption (low income)')
ax.legend()
plt.show()
```



Exercise 75.7.2

Following on from Exercises 1, let's look at how savings and aggregate asset holdings vary with the interest rate

Note

[Ljungqvist and Sargent, 2018] section 18.6 can be consulted for more background on the topic treated in this exercise.

For a given parameterization of the model, the mean of the stationary distribution of assets can be interpreted as aggregate capital in an economy with a unit mass of *ex-ante* identical households facing idiosyncratic shocks.

Your task is to investigate how this measure of aggregate capital varies with the interest rate.

Intuition suggests that a higher interest rate should encourage capital formation — test this.

For the interest rate grid, use

```
M = 8
r_vals = np.linspace(0, 0.05, M)
```

Solution

Here's one solution

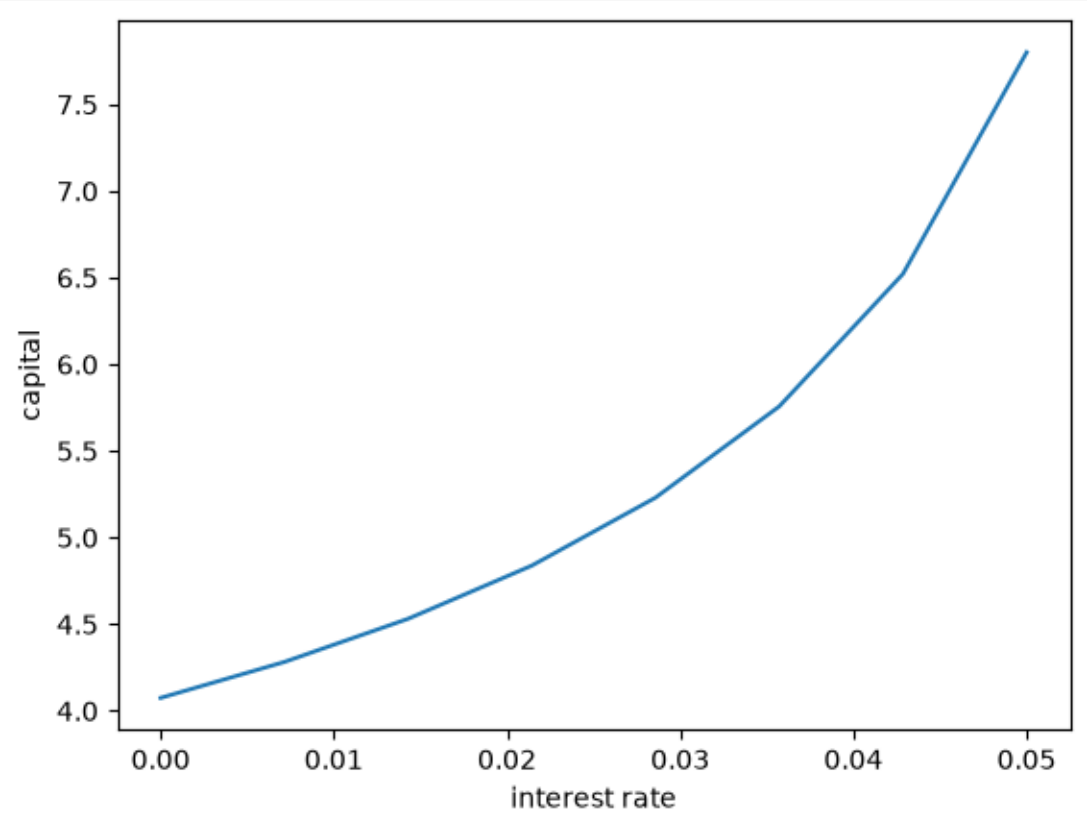
```
fig, ax = plt.subplots()

asset_mean = []
for r in r_vals:
    print(f'Solving model at r = {r}')
    ifp = create_ifp(r=r)
    R, β, γ, Π, z_grid, s, a_y, b_y, η_draws = ifp
    a_init = s[:, None] * jnp.ones(len(z_grid))
    c_init = a_init
    c_vec, a_vec = solve_model(ifp, c_init, a_init)
    assets = compute_asset_stationary(
        c_vec, a_vec, ifp, num_households=10_000, T=500
    )
    mean = np.mean(assets)
    asset_mean.append(mean)
    print(f'  Mean assets: {mean:.4f}')
    # Start next round with last solution
    c_init = c_vec
    a_init = a_vec
ax.plot(r_vals, asset_mean)

ax.set(xlabel='interest rate', ylabel='capital')

plt.show()
```

```
Solving model at r = 0.0
  Mean assets: 4.0748
Solving model at r = 0.0071428571428571435
  Mean assets: 4.2822
Solving model at r = 0.014285714285714287
  Mean assets: 4.5313
Solving model at r = 0.02142857142857143
  Mean assets: 4.8383
Solving model at r = 0.028571428571428574
  Mean assets: 5.2304
Solving model at r = 0.03571428571428572
  Mean assets: 5.7570
Solving model at r = 0.04285714285714286
  Mean assets: 6.5214
Solving model at r = 0.05
  Mean assets: 7.8005
```



As expected, aggregate savings increases with the interest rate.

CHAPTER

SEVENTYSIX

THE INCOME FLUCTUATION PROBLEM V: STOCHASTIC RETURNS ON ASSETS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *The Income Fluctuation Problem V: Stochastic Returns on Assets*
  - *Overview*
  - *The Model*
  - *Solution Algorithm*
  - *Implementation*
  - *Simulation*
  - *Wealth Inequality*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

76.1 Overview

In this lecture, we continue our study of the income fluctuation problem described in *The Income Fluctuation Problem III: The Endogenous Grid Method*.

While the interest rate was previously taken to be fixed, we now allow returns on assets to be state-dependent.

This matches the fact that most households with a positive level of assets face some capital income risk.

It has been argued that modeling capital income risk is essential for understanding the joint distribution of income and wealth (see, e.g., [Benhabib *et al.*, 2015] or [Stachurski and Toda, 2019]).

Theoretical properties of the household savings model presented here are analyzed in detail in [Ma *et al.*, 2020].

In terms of computation, we use a combination of time iteration and the endogenous grid method to solve the model quickly and accurately.

We require the following imports:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
import jax
import jax.numpy as jnp
from jax import vmap
from typing import NamedTuple
from functools import partial
```

76.2 The Model

In this section we review the household problem and optimality results.

76.2.1 Set Up

A household chooses a consumption-asset path $\{(c_t, a_t)\}$ to maximize

$$\mathbb{E}\left\{\sum_{t=0}^{\infty} \beta^t u(c_t)\right\} \tag{76.1}$$

subject to

$$a_{t+1} = R_{t+1}(a_t - c_t) + Y_{t+1} \text{ and } 0 \leq c_t \leq a_t, \tag{76.2}$$

with initial condition $(a_0, Z_0) = (a, z)$ treated as given.

The only difference from *The Income Fluctuation Problem IV: Transient Income Shocks* is that $\{R_t\}_{t\geq 1}$, the gross rate of return on wealth, is allowed to be stochastic.

In particular, we assume that

$$R_t = R(Z_t, \zeta_t) \quad \text{and} \quad Y_t = Y(Z_t, \eta_t), \tag{76.3}$$

where

- R and Y are time-invariant nonnegative functions,
- the innovation processes $\{\zeta_t\}$ and $\{\eta_t\}$ are IID and independent of each other, and

- $\{Z_t\}_{t \geq 0}$ is a Markov chain on a finite set Z

Let P represent the Markov matrix for the chain $\{Z_t\}_{t \geq 0}$.

In what follows, $\mathbb{E}_z \hat{X}$ means expectation of next period value $\hat{X}$ given current value $Z = z$.

76.2.2 Assumptions

We need restrictions to ensure that the objective (76.1) is finite and the solution methods described below converge.

We also need to ensure that the present discounted value of wealth does not grow too quickly.

When $\{R_t\}$ was constant we required that $\beta R < 1$.

Since it is now stochastic, we require (see [Ma *et al.*, 2020]) that

$$\beta G_R < 1, \quad \text{where} \quad G_R := \lim_{n \to \infty} \left(\mathbb{E} \prod_{t=1}^{n} R_t \right)^{1/n} \tag{76.4}$$

The value G_R can be thought of as the long run (geometric) average gross rate of return.

To simplify this lecture, we will *assume that the interest rate process is IID*.

In that case, it is clear from the definition of G_R that G_R is just $\mathbb{E} R_t$.

We test the condition $\beta \mathbb{E} R_t < 1$ in the code below.

Finally, we impose some routine technical restrictions on non-financial income.

$$\mathbb{E}\, Y_t < \infty \text{ and } \mathbb{E}\, u'(Y_t) < \infty$$

One relatively simple setting where all these restrictions are satisfied is the IID and CRRA environment of [Benhabib *et al.*, 2015].

76.2.3 Optimality

Let the class of candidate consumption policies $\mathcal{C}$ be defined as in *The Income Fluctuation Problem III: The Endogenous Grid Method*.

In [Ma *et al.*, 2020] it is shown that, under the stated assumptions,

- any $\sigma \in \mathcal{C}$ satisfying the Euler equation is an optimal policy and
- exactly one such policy exists in $\mathcal{C}$.

In the present setting, the Euler equation takes the form

$$(u' \circ \sigma)(a, z) = \max \left\{ \beta\, \mathbb{E}_z\, \hat{R}\, (u' \circ \sigma)[\hat{R}(a - \sigma(a, z)) + \hat{Y},\ \hat{Z}],\ u'(a) \right\} \tag{76.5}$$

(Intuition and derivation are similar to *The Income Fluctuation Problem III: The Endogenous Grid Method*.)

We again solve the Euler equation using time iteration, iterating with a Coleman–Reffett operator K defined to match the Euler equation (76.5).

76.3 Solution Algorithm

76.3.1 A Time Iteration Operator

Our definition of the candidate class $\sigma \in \mathcal{C}$ of consumption policies is the same as in *The Income Fluctuation Problem III: The Endogenous Grid Method*.

For fixed $\sigma \in \mathcal{C}$ and $(a, z) \in \mathbf{S}$, the value $K\sigma(a, z)$ of the function $K\sigma$ at (a, z) is defined as the $\xi \in (0, a]$ that solves

$$u'(\xi) = \max \left\{ \beta \, \mathbb{E}_z \, \hat{R} \, (u' \circ \sigma)[\hat{R}(a - \xi) + \hat{Y}, \, \hat{Z}], \, u'(a) \right\} \tag{76.6}$$

The idea behind K is that, as can be seen from the definitions, $\sigma \in \mathcal{C}$ satisfies the Euler equation if and only if $K\sigma(a, z) = \sigma(a, z)$ for all $(a, z) \in \mathbf{S}$.

This means that fixed points of K in $\mathcal{C}$ and optimal consumption policies exactly coincide (see [Ma *et al.*, 2020] for more details).

76.3.2 Convergence Properties

As before, we pair $\mathcal{C}$ with the distance

$$\rho(c, d) := \sup_{(a,z) \in \mathbf{S}} \left| (u' \circ c) \, (a, z) - (u' \circ d) \, (a, z) \right|,$$

It can be shown that

1. $(\mathcal{C}, \rho)$ is a complete metric space,
2. there exists an integer n such that K^n is a contraction mapping on $(\mathcal{C}, \rho)$, and
3. The unique fixed point of K in $\mathcal{C}$ is the unique optimal policy in $\mathcal{C}$.

We now have a clear path to successfully approximating the optimal policy: choose some $\sigma \in \mathcal{C}$ and then iterate with K until convergence (as measured by the distance ρ).

76.3.3 Using an Endogenous Grid

In the study of that model we found that it was possible to further accelerate time iteration via the *endogenous grid method*.

We will use the same method here.

The methodology is the same as it was for the optimal growth model, with the minor exception that we need to remember that consumption is not always interior.

In particular, optimal consumption can be equal to assets when the level of assets is low.

Finding Optimal Consumption

The endogenous grid method (EGM) calls for us to take a grid of *savings* values s_i, where each such s is interpreted as $s = a - c$.

For the lowest grid point we take $s_0 = 0$.

For the corresponding a_0, c_0 pair we have $a_0 = c_0$.

This happens close to the origin, where assets are low and the household consumes all that it can.

Although there are many solutions, the one we take is $a_0 = c_0 = 0$, which pins down the policy at the origin, aiding interpolation.

For $s > 0$, we have, by definition, $c < a$, and hence consumption is interior.

Hence the max component of (76.5) drops out, and we solve for

$$c_i = (u')^{-1} \left\{ \beta \, \mathbb{E}_z \hat{R} (u' \circ \sigma) \left[\hat{R} s_i + \hat{Y}, \, \hat{Z} \right] \right\} \tag{76.7}$$

at each s_i.

Iterating

Once we have the pairs $\{s_i, c_i\}$, the endogenous asset grid is obtained by $a_i = c_i + s_i$.

Also, we held $z \in \mathsf{Z}$ in the discussion above so we can pair it with a_i.

An approximation of the policy $(a, z) \mapsto \sigma(a, z)$ can be obtained by interpolating $\{a_i, c_i\}$ at each z.

In what follows, we use linear interpolation.

76.4 Implementation

Here's the model as a `NamedTuple`.

```
class IFP(NamedTuple):
    """
    A NamedTuple that stores primitives for the income fluctuation
    problem, using JAX.
    """
    γ: float
    β: float
    P: jnp.ndarray
    a_r: float
    b_r: float
    a_y: float
    b_y: float
    s_grid: jnp.ndarray
    η_draws: jnp.ndarray
    ζ_draws: jnp.ndarray


def create_ifp(
        γ=1.5,                      # Utility parameter
        β=0.96,                     # Discount factor
        P=jnp.array([(0.9, 0.1),    # Default Markov chain for Z
                     (0.1, 0.9)]),
        a_r=0.16,                   # Volatility term in R shock
        b_r=0.0,                    # Mean shift R shock
        a_y=0.2,                    # Volatility term in Y shock
        b_y=0.5,                    # Mean shift Y shock
        shock_draw_size=100,        # For Monte Carlo
        grid_max=100,               # Exogenous grid max
        grid_size=100,              # Exogenous grid size
        seed=1234                   # Random seed
    ):
```

(continues on next page)

(continued from previous page)

```
    """
    Create an instance of IFP with the given parameters.

    """
    # Test stability assuming {R_t} is IID and ln R ~ N(b_r, a_r)
    ER = np.exp(b_r + a_r**2 / 2)
    assert β * ER < 1, "Stability condition failed."

    # Generate random draws using JAX
    key = jax.random.key(seed)
    subkey1, subkey2 = jax.random.split(key)
    η_draws = jax.random.normal(subkey1, (shock_draw_size,))
    ζ_draws = jax.random.normal(subkey2, (shock_draw_size,))
    s_grid = jnp.linspace(0, grid_max, grid_size)

    return IFP(
        γ, β, P, a_r, b_r, a_y, b_y, s_grid, η_draws, ζ_draws
    )


def u_prime(c, γ):
    """Marginal utility"""
    return c**(-γ)

def u_prime_inv(c, γ):
    """Inverse of marginal utility"""
    return c**(-1/γ)

def R(z, ζ, a_r, b_r):
    """Gross return on assets"""
    return jnp.exp(a_r * ζ + b_r)

def Y(z, η, a_y, b_y):
    """Labor income"""
    return jnp.exp(a_y * η + (z * b_y))
```

Here's the Coleman-Reffett operator using JAX:

```
def K(
        c_in: jnp.array,   # c_in[i, z] = consumption at a_in[i, z]
        a_in: jnp.array,   # a_in[i, z] is an asset grid
        ifp: IFP
    ):
    """
    The Coleman--Reffett operator for the income fluctuation problem,
    using the endogenous grid method with JAX.

    """

    # Extract parameters from ifp
    γ, β, P, a_r, b_r, a_y, b_y, s_grid, η_draws, ζ_draws = ifp
    n = len(P)

    def compute_expectation(s, z):
        def inner_expectation(z_hat):
            def compute_term(η, ζ):
                R_hat = R(z_hat, ζ, a_r, b_r)
```

(continues on next page)

(continued from previous page)

```
                Y_hat = Y(z_hat, η, a_y, b_y)
                a_val = R_hat * s + Y_hat
                # Interpolate consumption
                c_interp = jnp.interp(a_val, a_in[:, z_hat], c_in[:, z_hat])
                mu = u_prime(c_interp, γ)
                return R_hat * mu
            # Vectorize over all shock combinations
            η_grid, ζ_grid = jnp.meshgrid(η_draws, ζ_draws, indexing='ij')
            terms = vmap(vmap(compute_term))(η_grid, ζ_grid)
            return P[z, z_hat] * jnp.mean(terms)
        # Sum over z_hat states
        Ez = jnp.sum(vmap(inner_expectation)(jnp.arange(n)))
        return u_prime_inv(β * Ez, γ)

    # Vectorize over s_grid and z
    compute_exp_v1 = vmap(compute_expectation, in_axes=(None, 0))
    compute_exp_v2 = vmap(compute_exp_v1,      in_axes=(0, None))
    c_out = compute_exp_v2(s_grid, jnp.arange(n))
    # Calculate endogenous asset grid
    a_out = s_grid[:, None] + c_out
    # Fix consumption-asset pair at (0, 0)
    c_out = c_out.at[0, :].set(0)
    a_out = a_out.at[0, :].set(0)

    return c_out, a_out
```

The next function solves for an approximation of the optimal consumption policy via time iteration using JAX:

```
@jax.jit
def solve_model(
        ifp: IFP,
        c_init: jnp.ndarray,  # Initial guess of σ on grid endogenous grid
        a_init: jnp.ndarray,  # Initial endogenous grid
        tol: float = 1e-5,
        max_iter: int = 1000
    ) -> jnp.ndarray:
    " Solve the model using time iteration with EGM. "

    def condition(loop_state):
        c_in, a_in, i, error = loop_state
        return (error > tol) & (i < max_iter)

    def body(loop_state):
        c_in, a_in, i, error = loop_state
        c_out, a_out = K(c_in, a_in, ifp)
        error = jnp.max(jnp.abs(c_out - c_in))
        i += 1
        return c_out, a_out, i, error

    i, error = 0, tol + 1
    initial_state = (c_init, a_init, i, error)
    final_loop_state = jax.lax.while_loop(condition, body, initial_state)
    c_out, a_out, i, error = final_loop_state

    return c_out, a_out
```

Now we can create an instance and solve the model using JAX:

```
ifp = create_ifp()
```

Set up the initial condition:

```
# Initial guess of σ = consume all assets
k = len(ifp.s_grid)
n = len(ifp.P)
σ_init = jnp.empty((k, n))
for z in range(n):
    σ_init = σ_init.at[:, z].set(ifp.s_grid)
a_init = σ_init.copy()
```

Let's generate an approximation solution with JAX:

```
σ_star, a_star = solve_model(ifp, σ_init, a_init)
```

Let's try it again with a timer.

```
with qe.Timer(precision=8):
    σ_star, a_star = solve_model(ifp, σ_init, a_init)
    σ_star.block_until_ready()
```

```
0.45794129 seconds elapsed
```

76.5 Simulation

Let's return to the default model and study the stationary distribution of assets.

Our plan is to run a large number of households forward for T periods and then histogram the cross-sectional distribution of assets.

Set `num_households=50_000, T=500`.

First we write a function to run a single household forward in time and record the final value of assets.

The function takes a solution pair `c_vec` and `a_vec`, understanding them as representing an optimal policy associated with a given model `ifp`

```
def simulate_household(
        key, a_0, z_idx_0, c_vec, a_vec, ifp, T
    ):
    """
    Simulates a single household for T periods to approximate the stationary
    distribution of assets.

    - key is the state of the random number generator
    - ifp is an instance of IFP
    - c_vec, a_vec are the optimal consumption policy, endogenous grid for ifp

    """
    # Extract parameters from ifp
    γ, β, P, a_r, b_r, a_y, b_y, s_grid, η_draws, ζ_draws = ifp
    n_z = len(P)

    # Create interpolation function for consumption policy
```

(continues on next page)

(continued from previous page)

```
    σ = lambda a, z_idx: jnp.interp(a, a_vec[:, z_idx], c_vec[:, z_idx])

    # Simulate forward T periods
    def update(t, state):
        a, z_idx = state
        # Draw next shock z' from P[z, z']
        current_key = jax.random.fold_in(key, 3*t)
        z_next_idx = jax.random.choice(current_key, n_z, p=P[z_idx]).astype(jnp.int32)
        # Draw η shock for income
        η_key = jax.random.fold_in(key, 3*t + 1)
        η = jax.random.normal(η_key)
        # Draw ζ shock for return
        ζ_key = jax.random.fold_in(key, 3*t + 2)
        ζ = jax.random.normal(ζ_key)
        # Compute stochastic return
        R_next = R(z_next_idx, ζ, a_r, b_r)
        # Compute income
        Y_next = Y(z_next_idx, η, a_y, b_y)
        # Update assets: a' = R' * (a - c) + Y'
        a_next = R_next * (a - σ(a, z_idx)) + Y_next
        # Return updated state
        return a_next, z_next_idx

    initial_state = a_0, z_idx_0
    final_state = jax.lax.fori_loop(0, T, update, initial_state)
    a_final, _ = final_state
    return a_final
```

Now we write a function to simulate many households in parallel.

```
@partial(jax.jit, static_argnums=(3, 4, 5))
def compute_asset_stationary(
        c_vec, a_vec, ifp, num_households=50_000, T=500, seed=1234
    ):
    """
    Simulates num_households households for T periods to approximate
    the stationary distribution of assets.

    Returns the final cross-section of asset holdings.

    - ifp is an instance of IFP
    - c_vec, a_vec are the optimal consumption policy and endogenous grid.

    """
    # Extract parameters from ifp
    γ, β, P, a_r, b_r, a_y, b_y, s_grid, η_draws, ζ_draws = ifp

    # Start with assets = savings_grid_max / 2
    a_0_vector = jnp.full(num_households, s_grid[-1] / 2)
    # Initialize the exogenous state of each household
    z_idx_0_vector = jnp.zeros(num_households).astype(jnp.int32)

    # Vectorize over many households
    key = jax.random.key(seed)
    keys = jax.random.split(key, num_households)
    # Vectorize simulate_household in (key, a_0, z_idx_0)
    sim_all_households = jax.vmap(
```

(continues on next page)

(continued from previous page)

```
        simulate_household, in_axes=(0, 0, 0, None, None, None, None)
    )
    assets = sim_all_households(keys, a_0_vector, z_idx_0_vector, c_vec, a_vec, ifp,␣
↪T)

    return jnp.array(assets)
```

We'll need some inequality measures for visualization, so let's define them first:

```
def gini_coefficient(x):
    """
    Compute the Gini coefficient for array x.

    """
    x = jnp.asarray(x)
    n = len(x)
    x_sorted = jnp.sort(x)
    # Compute Gini coefficient
    cumsum = jnp.cumsum(x_sorted)
    a = (2 * jnp.sum((jnp.arange(1, n+1)) * x_sorted)) / (n * cumsum[-1])
    return a - (n + 1) / n


def top_share(
        x: jnp.array,    # array of wealth values
        p: float=0.01    # fraction of top households (default 0.01 for top 1%)
    ):
    """
    Compute the share of total wealth held by the top p fraction of households.

    """
    x = jnp.asarray(x)
    x_sorted = jnp.sort(x)
    # Number of households in top p%
    n_top = int(jnp.ceil(len(x) * p))
    # Wealth held by top p%
    wealth_top = jnp.sum(x_sorted[-n_top:])
    # Total wealth
    wealth_total = jnp.sum(x_sorted)
    return wealth_top / wealth_total
```

Now we call the function, generate the asset distribution and visualize it:

```
ifp = create_ifp()
# Extract parameters for initialization
s_grid = ifp.s_grid
n_z = len(ifp.P)
a_init = s_grid[:, None] * jnp.ones(n_z)
c_init = a_init
c_vec, a_vec = solve_model(ifp, c_init, a_init)
assets = compute_asset_stationary(c_vec, a_vec, ifp, num_households=200_000)

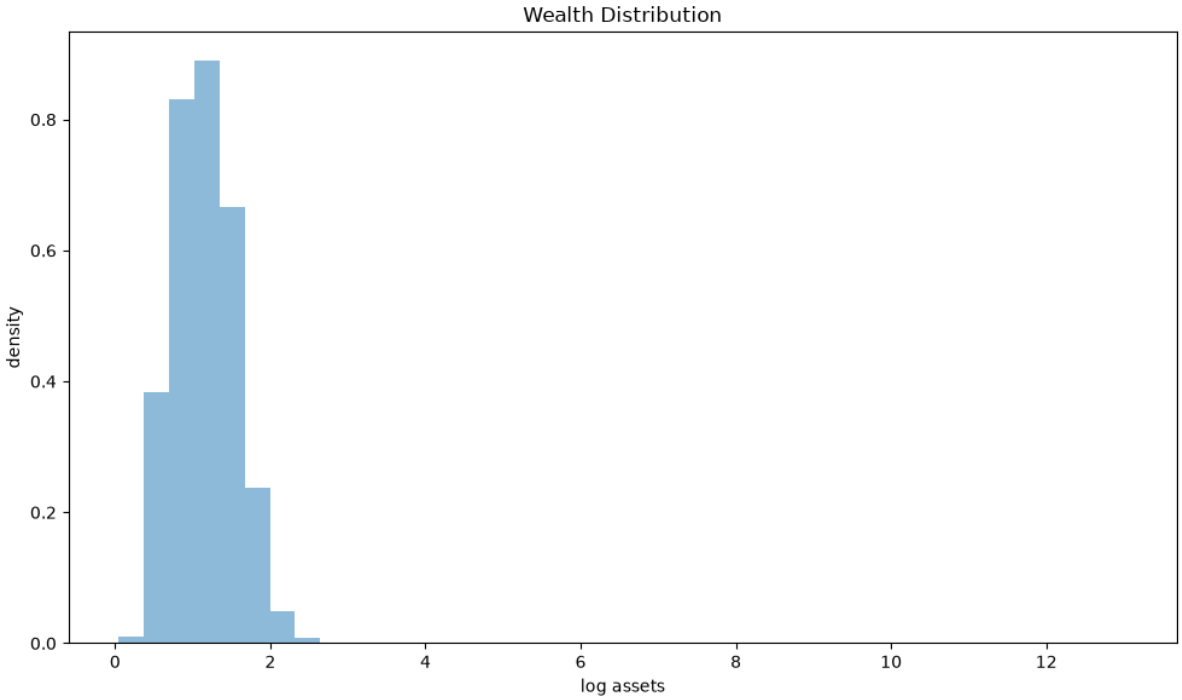
# Compute Gini coefficient for the plot
gini_plot = gini_coefficient(assets)

# Plot histogram of log wealth
fig, ax = plt.subplots(figsize=(10, 6))
```

(continues on next page)

(continued from previous page)

```
ax.hist(jnp.log(assets), bins=40, alpha=0.5, density=True)
ax.set(xlabel='log assets', ylabel='density', title="Wealth Distribution")
plt.tight_layout()
plt.show()
```



The histogram shows the distribution of log wealth.

Bearing in mind that we are looking at log values, the histogram suggests a long right tail of the distribution.

Below we examine this in more detail.

76.6 Wealth Inequality

Lets' look at wealth inequality by computing some standard measures of this phenomenon.

We will also examine how inequality varies with the interest rate.

76.6.1 Measuring Inequality

Let's print the Gini coefficient and the top 1% wealth share from our simulation:

```
gini = gini_coefficient(assets)
top1 = top_share(assets, p=0.01)

print(f"Gini coefficient: {gini:.4f}")
print(f"Top 1% wealth share: {top1:.4f}")
```

```
Gini coefficient: 0.7847
Top 1% wealth share: 0.7306
```

Recent numbers suggest that

- the Gini coefficient for wealth in the US is around 0.8
- the top 1% wealth share is over 0.3

Our model with stochastic returns generates a Gini coefficient close to the empirical value, demonstrating that capital income risk is an important factor in wealth inequality.

The top 1% wealth share is, however, too large.

Our model needs proper calibration and additional work – we set these tasks aside for now.

76.7 Exercises

Exercise 76.7.1

Plot how the Gini coefficient varies with the volatility of returns on assets.

Specifically, compute the Gini coefficient for values of `a_r` ranging from 0.10 to 0.16 (use at least 5 different values) and plot the results.

What does this tell you about the relationship between capital income risk and wealth inequality?

Solution

We loop over different values of `a_r`, solve the model for each, simulate the wealth distribution, and compute the Gini coefficient.

```
# Range of a_r values to explore
a_r_vals = np.linspace(0.10, 0.16, 5)
gini_vals = []

print("Computing Gini coefficients for different return volatilities...\n")

for a_r in a_r_vals:
    print(f"a_r = {a_r:.3f}...", end=" ")

    # Create model with this a_r value
    ifp_temp = create_ifp(a_r=a_r, grid_max=100)

    # Solve the model
    s_grid_temp = ifp_temp.s_grid
    n_z_temp = len(ifp_temp.P)
    a_init_temp = s_grid_temp[:, None] * jnp.ones(n_z_temp)
    c_init_temp = a_init_temp
    c_vec_temp, a_vec_temp = solve_model(
        ifp_temp, c_init_temp, a_init_temp
    )

    # Simulate households
    assets_temp = compute_asset_stationary(
        c_vec_temp, a_vec_temp, ifp_temp, num_households=200_000
    )

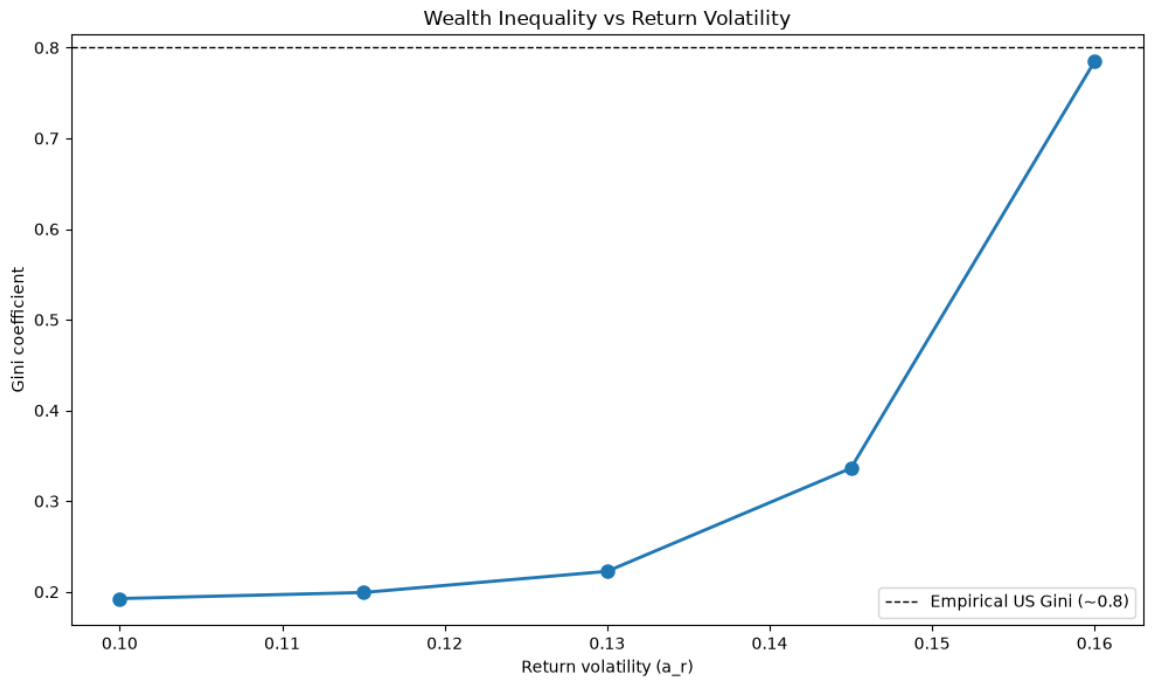
    # Compute Gini coefficient
```

```
    gini_temp = gini_coefficient(assets_temp)
    gini_vals.append(gini_temp)
    print(f"Gini = {gini_temp:.4f}")

# Plot the results
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(a_r_vals, gini_vals, 'o-', linewidth=2, markersize=8)
ax.set(xlabel='Return volatility (a_r)',
       ylabel='Gini coefficient',
       title='Wealth Inequality vs Return Volatility')
ax.axhline(y=0.8, color='k', linestyle='--', linewidth=1,
           label='Empirical US Gini (~0.8)')
ax.legend()
plt.tight_layout()
plt.show()
```

```
Computing Gini coefficients for different return volatilities...

a_r = 0.100...
Gini = 0.1925
a_r = 0.115...
Gini = 0.1992
a_r = 0.130...
Gini = 0.2225
a_r = 0.145...
Gini = 0.3362
a_r = 0.160...
Gini = 0.7847
```



The plot shows that wealth inequality (measured by the Gini coefficient) increases with return volatility.

This demonstrates that capital income risk is a key driver of wealth inequality.

When returns are more volatile, lucky households who experience sequences of high returns accumulate substantially

more wealth than unlucky households, leading to greater inequality in the wealth distribution.

Exercise 76.7.2

Plot how the Gini coefficient varies with the volatility of labor income.

Specifically, compute the Gini coefficient for values of `a_y` ranging from 0.125 to 0.20 and plot the results. Set `a_r=0.10` for this exercise.

What does this tell you about the relationship between labor income risk and wealth inequality? Can we achieve the same rise in inequality by varying labor income volatility as we can by varying return volatility?

Solution

We loop over different values of `a_y`, solve the model for each, simulate the wealth distribution, and compute the Gini coefficient.

```
# Range of a_y values to explore
a_y_vals = np.linspace(0.125, 0.20, 5)
gini_vals_y = []

print("Computing Gini coefficients for different labor income volatilities...\n")

for a_y in a_y_vals:
    print(f"a_y = {a_y:.3f}...", end=" ")

    # Create model with this a_y value and a_r=0.10
    ifp_temp = create_ifp(a_y=a_y, a_r=0.10, grid_max=100)

    # Solve the model
    s_grid_temp = ifp_temp.s_grid
    n_z_temp = len(ifp_temp.P)
    a_init_temp = s_grid_temp[:, None] * jnp.ones(n_z_temp)
    c_init_temp = a_init_temp
    c_vec_temp, a_vec_temp = solve_model(
        ifp_temp, c_init_temp, a_init_temp
    )

    # Simulate households
    assets_temp = compute_asset_stationary(
        c_vec_temp, a_vec_temp, ifp_temp, num_households=200_000
    )

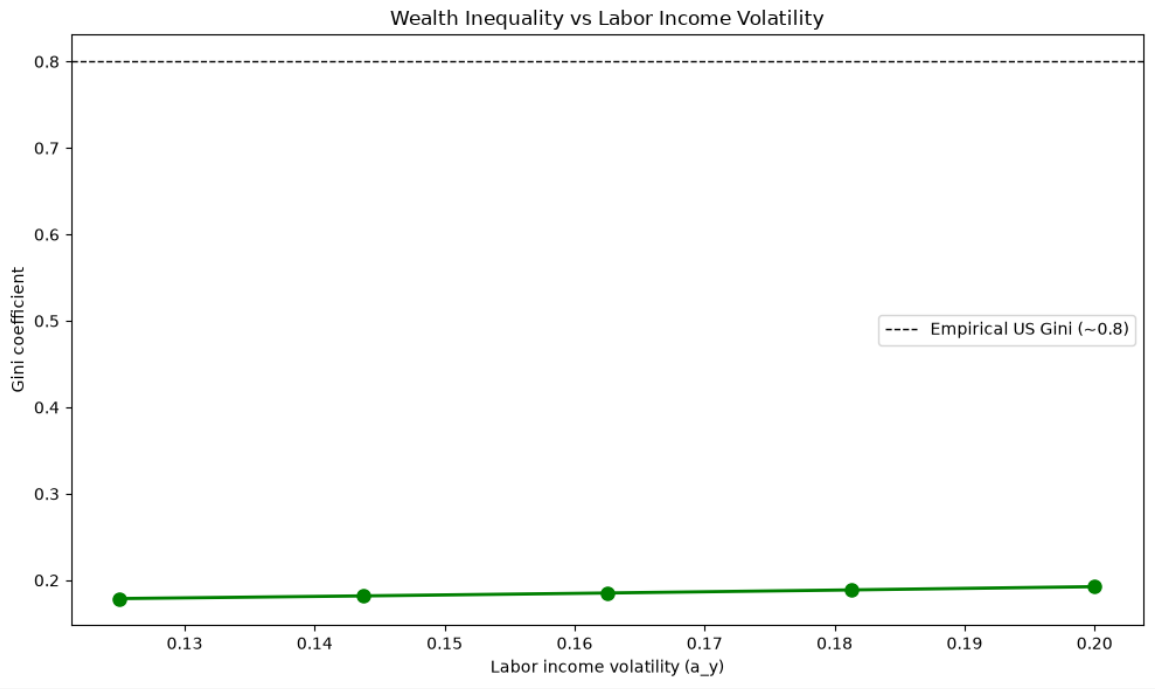
    # Compute Gini coefficient
    gini_temp = gini_coefficient(assets_temp)
    gini_vals_y.append(gini_temp)
    print(f"Gini = {gini_temp:.4f}")

# Plot the results
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(a_y_vals, gini_vals_y, 'o-', linewidth=2, markersize=8, color='green')
ax.set(xlabel='Labor income volatility (a_y)',
       ylabel='Gini coefficient',
       title='Wealth Inequality vs Labor Income Volatility')
ax.axhline(y=0.8, color='k', linestyle='--', linewidth=1,
```

```
            label='Empirical US Gini (~0.8)')
ax.legend()
plt.tight_layout()
plt.show()
```

```
Computing Gini coefficients for different labor income volatilities...

a_y = 0.125...
Gini = 0.1787
a_y = 0.144...
Gini = 0.1818
a_y = 0.163...
Gini = 0.1852
a_y = 0.181...
Gini = 0.1887
a_y = 0.200...
Gini = 0.1925
```



The plot shows that wealth inequality increases with labor income volatility, but the effect is much weaker than the effect of return volatility.

Comparing the two exercises:

- When return volatility (`a_r`) varies from 0.10 to 0.16, the Gini coefficient rises dramatically from around 0.20 to 0.79
- When labor income volatility (`a_y`) varies from 0.125 to 0.20, a similar amount in percentage terms, the Gini coefficient increases but by a much smaller amount

This suggests that capital income risk is a more important driver of wealth inequality than labor income risk.

The intuition is that wealth accumulation compounds over time: households who experience favorable returns on their assets can reinvest those returns, leading to exponential growth.

In contrast, labor income shocks, while they affect current consumption and savings, do not have the same compound-

ing effect on wealth accumulation.

Part XI

LQ Control

CHAPTER

SEVENTYSEVEN

LQ CONTROL: FOUNDATIONS

Contents

- *LQ Control: Foundations*
  - *Overview*
  - *Introduction*
  - *Optimality – finite horizon*
  - *Implementation*
  - *Extensions and comments*
  - *Further applications*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

77.1 Overview

Linear quadratic (LQ) control refers to a class of dynamic optimization problems that have found applications in almost every scientific field.

This lecture provides an introduction to LQ control and its economic applications.

As we will see, LQ systems have a simple structure that makes them an excellent workhorse for a wide variety of economic problems.

Moreover, while the linear-quadratic structure is restrictive, it is in fact far more flexible than it may appear initially.

These themes appear repeatedly below.

Mathematically, LQ control problems are closely related to *the Kalman filter*

- Recursive formulations of linear-quadratic control problems and Kalman filtering problems both involve matrix Riccati equations.
- Classical formulations of linear control and linear filtering problems make use of similar matrix decompositions (see for example Classical Control with Linear Algebra and Classical Prediction and Filtering With Linear Algebra).

In reading what follows, it will be useful to have some familiarity with

- matrix manipulations
- vectors of random variables
- dynamic programming and the Bellman equation (see for example Shortest Paths and *Optimal Savings III: Stochastic Returns*)

For additional reading on LQ control, see, for example,

- [Ljungqvist and Sargent, 2018], chapter 5
- [Hansen and Sargent, 2008], chapter 4
- [Hernandez-Lerma and Lasserre, 1996], section 3.5

In order to focus on computation, we leave longer proofs to these sources (while trying to provide as much intuition as possible).

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
from quantecon import LQ
from typing import NamedTuple
```

77.2 Introduction

The "linear" part of LQ is a linear law of motion for the state, while the "quadratic" part refers to preferences.

Let's begin with the former, move on to the latter, and then put them together into an optimization problem.

77.2.1 The law of motion

Let x_t be a vector describing the state of some economic system.

Suppose that x_t follows a linear law of motion given by

$$x_{t+1} = Ax_t + Bu_t + Cw_{t+1}, \qquad t = 0, 1, 2, \ldots \tag{77.1}$$

Here

- u_t is a "control" vector, incorporating choices available to a decision-maker confronting the current state x_t
- $\{w_t\}$ is a sequence of uncorrelated zero mean shocks satisfying $\mathbb{E} w_t w_t^\top = I$, where the right-hand side is the identity matrix

Regarding the dimensions

- x_t is $n \times 1$, A is $n \times n$
- u_t is $k \times 1$, B is $n \times k$
- w_t is $j \times 1$, C is $n \times j$

Example 1

Consider a household budget constraint given by

$$a_{t+1} + c_t = (1 + r)a_t + y_t$$

Here a_t is assets, r is a fixed interest rate, c_t is current consumption, and y_t is current non-financial income.

If we suppose that $\{y_t\}$ is serially uncorrelated and $N(0, \sigma^2)$, then, taking $\{w_t\}$ to be standard normal, we can write the system as

$$a_{t+1} = (1 + r)a_t - c_t + \sigma w_{t+1}$$

This is clearly a special case of (77.1), with assets being the state and consumption being the control.

Example 2

One unrealistic feature of the previous model is that non-financial income has a zero mean and is often negative.

This can easily be overcome by adding a sufficiently large mean.

Hence in this example, we take $y_t = \sigma w_{t+1} + \mu$ for some positive real number μ.

Another alteration that's useful to introduce (we'll see why soon) is to change the control variable from consumption to the deviation of consumption from some "ideal" quantity $\bar{c}$.

(Most parameterizations will be such that $\bar{c}$ is large relative to the amount of consumption that is attainable in each period, and hence the household wants to increase consumption.)

For this reason, we now take our control to be $u_t := c_t - \bar{c}$.

In terms of these variables, the budget constraint $a_{t+1} = (1 + r)a_t - c_t + y_t$ becomes

$$a_{t+1} = (1 + r)a_t - u_t - \bar{c} + \sigma w_{t+1} + \mu \tag{77.2}$$

How can we write this new system in the form of equation (77.1)?

If, as in the previous example, we take a_t as the state, then we run into a problem: the law of motion contains some constant terms on the right-hand side.

This means that we are dealing with an *affine* function, not a linear one (recall *this discussion*).

Fortunately, we can easily circumvent this problem by adding an extra state variable.

In particular, if we write

$$\begin{bmatrix} a_{t+1} \\ 1 \end{bmatrix} = \begin{bmatrix} 1 + r & -\bar{c} + \mu \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a_t \\ 1 \end{bmatrix} + \begin{bmatrix} -1 \\ 0 \end{bmatrix} u_t + \begin{bmatrix} \sigma \\ 0 \end{bmatrix} w_{t+1} \tag{77.3}$$

then the first row is equivalent to (77.2).

Moreover, the model is now linear and can be written in the form of (77.1) by setting

$$x_t := \begin{bmatrix} a_t \\ 1 \end{bmatrix}, \quad A := \begin{bmatrix} 1 + r & -\bar{c} + \mu \\ 0 & 1 \end{bmatrix}, \quad B := \begin{bmatrix} -1 \\ 0 \end{bmatrix}, \quad C := \begin{bmatrix} \sigma \\ 0 \end{bmatrix} \tag{77.4}$$

In effect, we've bought ourselves linearity by adding another state.

77.2.2 Preferences

In the LQ model, the aim is to minimize flow of losses, where time-t loss is given by the quadratic expression

$$x_t^\top R x_t + u_t^\top Q u_t \tag{77.5}$$

The entries in R and Q are chosen to match the specific problem being studied.

Here

- R is assumed to be $n \times n$, symmetric and nonnegative definite.
- Q is assumed to be $k \times k$, symmetric and positive definite.

Note

In fact, for many economic problems, the definiteness conditions on R and Q can be relaxed. It is sufficient that certain submatrices of R and Q be nonnegative definite. See [Hansen and Sargent, 2008] for details.

Note also that this notation is not universal: some authors let Q be the matrix associated with the state and R the matrix associated with the control.

Example 1

A very simple example that satisfies these assumptions is to take R and Q to be identity matrices so that current loss is

$$x_t^\top I x_t + u_t^\top I u_t = \|x_t\|^2 + \|u_t\|^2$$

Thus, for both the state and the control, loss is measured as squared distance from the origin.

(In fact, the general case (77.5) can also be understood in this way, but with R and Q identifying other – non-Euclidean – notions of "distance" from the zero vector.)

Intuitively, we can often think of the state x_t as representing deviation from a target, such as

- deviation of inflation from some target level
- deviation of a firm's capital stock from some desired quantity

The aim is to put the state close to the target, while using controls parsimoniously.

Example 2

In the household problem *studied above*, setting $R = 0$ and $Q = 1$ yields preferences

$$x_t^\top R x_t + u_t^\top Q u_t = u_t^2 = (c_t - \bar{c})^2$$

Under this specification, the household's current loss is the squared deviation of consumption from the ideal level $\bar{c}$.

77.3 Optimality – finite horizon

Let's now be precise about the optimization problem we wish to consider, and look at how to solve it.

77.3.1 The objective

We will begin with the finite horizon case, with terminal time $T \in \mathbb{N}$.

In this case, the aim is to choose a sequence of controls $\{u_0, \dots, u_{T-1}\}$ to minimize the objective

$$\mathbb{E}\left\{\sum_{t=0}^{T-1} \beta^t (x_t^\top R x_t + u_t^\top Q u_t) + \beta^T x_T^\top R_f x_T\right\} \tag{77.6}$$

subject to the law of motion (77.1) and initial state x_0.

The new objects introduced here are β and the matrix R_f.

The scalar β is the discount factor, while $x^\top R_f x$ gives terminal loss associated with state x.

Comments:

- We assume R_f to be $n \times n$, symmetric and nonnegative definite.
- We allow $\beta = 1$, and hence include the undiscounted case.
- x_0 may itself be random, in which case we require it to be independent of the shock sequence $\{w_1, \dots, w_T\}$.

77.3.2 Information

There's one constraint we've neglected to mention so far, which is that the decision-maker who solves this LQ problem knows only the present and the past, not the future.

To clarify this point, consider the sequence of controls $\{u_0, \dots, u_{T-1}\}$.

When choosing these controls, the decision-maker is permitted to take into account the effects of the shocks $\{w_1, \dots, w_T\}$ on the system.

However, it is typically assumed — and will be assumed here — that the time-t control u_t can be made with knowledge of past and present shocks only.

The fancy measure-theoretic way of saying this is that u_t must be measurable with respect to the σ-algebra generated by $x_0, w_1, w_2, \dots, w_t$.

This is in fact equivalent to stating that u_t can be written in the form $u_t = g_t(x_0, w_1, w_2, \dots, w_t)$ for some Borel measurable function g_t.

(Just about every function that's useful for applications is Borel measurable, so, for the purposes of intuition, you can read that last phrase as "for some function g_t")

Now note that x_t will ultimately depend on the realizations of $x_0, w_1, w_2, \dots, w_t$.

In fact, it turns out that x_t summarizes all the information about these historical shocks that the decision-maker needs to set controls optimally.

More precisely, it can be shown that any optimal control u_t can always be written as a function of the current state alone.

Hence in what follows we restrict attention to control policies (i.e., functions) of the form $u_t = g_t(x_t)$.

Actually, the preceding discussion applies to all standard dynamic programming problems.

What's special about the LQ case is that – as we shall soon see — the optimal u_t turns out to be a linear function of x_t.

77.3.3 Solution

To solve the finite horizon LQ problem we can use a dynamic programming strategy based on backward induction that is conceptually similar to the approach adopted in Shortest Paths.

For reasons that will soon become clear, we first introduce the notation $J_T(x) = x^\top R_f x$.

Now consider the problem of the decision-maker in the second to last period.

In particular, let the time be $T-1$, and suppose that the state is x_{T-1}.

The decision-maker must trade-off current and (discounted) final losses, and hence solves

$$\min_u \{x_{T-1}^\top R x_{T-1} + u^\top Q u + \beta \, \mathbb{E} J_T(A x_{T-1} + B u + C w_T)\}$$

We use u instead of u_{T-1} here to simplify notation.

At this stage, it is convenient to define the function

$$J_{T-1}(x) = \min_u \{x^\top R x + u^\top Q u + \beta \, \mathbb{E} J_T(A x + B u + C w_T)\} \tag{77.7}$$

The function J_{T-1} will be called the $T-1$ value function, and $J_{T-1}(x)$ can be thought of as representing total "loss-to-go" from state x at time $T-1$ when the decision-maker behaves optimally.

Now let's step back to $T-2$.

For a decision-maker at $T-2$, the value $J_{T-1}(x)$ plays a role analogous to that played by the terminal loss $J_T(x) = x^\top R_f x$ for the decision-maker at $T-1$.

That is, $J_{T-1}(x)$ summarizes the future loss associated with moving to state x.

The decision-maker chooses her control u to trade off current loss against future loss, where

- the next period state is $x_{T-1} = A x_{T-2} + B u + C w_{T-1}$, and hence depends on the choice of current control.
- the "cost" of landing in state x_{T-1} is $J_{T-1}(x_{T-1})$.

Her problem is therefore

$$\min_u \{x_{T-2}^\top R x_{T-2} + u^\top Q u + \beta \, \mathbb{E} J_{T-1}(A x_{T-2} + B u + C w_{T-1})\}$$

Letting

$$J_{T-2}(x) = \min_u \{x^\top R x + u^\top Q u + \beta \, \mathbb{E} J_{T-1}(A x + B u + C w_{T-1})\}$$

the pattern for backward induction is now clear.

In particular, we define a sequence of value functions $\{J_0, \ldots, J_T\}$ via

$$J_{t-1}(x) = \min_u \{x^\top R x + u^\top Q u + \beta \, \mathbb{E} J_t(A x + B u + C w_t)\} \quad \text{and} \quad J_T(x) = x^\top R_f x$$

The first equality is the Bellman equation from dynamic programming theory specialized to the finite horizon LQ problem.

Now that we have $\{J_0, \ldots, J_T\}$, we can obtain the optimal controls.

As a first step, let's find out what the value functions look like.

It turns out that every J_t has the form $J_t(x) = x^\top P_t x + d_t$ where P_t is a $n \times n$ matrix and d_t is a constant.

We can show this by induction, starting from $P_T := R_f$ and $d_T = 0$.

Using this notation, (77.7) becomes

$$J_{T-1}(x) = \min_u \{x^\top R x + u^\top Q u + \beta \, \mathbb{E} (A x + B u + C w_T)^\top P_T (A x + B u + C w_T)\} \tag{77.8}$$

To obtain the minimizer, we can take the derivative of the r.h.s. with respect to u and set it equal to zero.

Applying the relevant rules of *matrix calculus*, this gives

$$u = -(Q + \beta B^\top P_T B)^{-1} \beta B^\top P_T A x \tag{77.9}$$

Plugging this back into (77.8) and rearranging yields

$$J_{T-1}(x) = x^\top P_{T-1} x + d_{T-1}$$

where

$$P_{T-1} = R - \beta^2 A^\top P_T B (Q + \beta B^\top P_T B)^{-1} B^\top P_T A + \beta A^\top P_T A \tag{77.10}$$

and

$$d_{T-1} := \beta \, \text{trace}(C^\top P_T C) \tag{77.11}$$

(The algebra is a good exercise — we'll leave it up to you.)

If we continue working backwards in this manner, it soon becomes clear that $J_t(x) = x^\top P_t x + d_t$ as claimed, where P_t and d_t satisfy the recursions

$$P_{t-1} = R - \beta^2 A^\top P_t B (Q + \beta B^\top P_t B)^{-1} B^\top P_t A + \beta A^\top P_t A \quad \text{with} \quad P_T = R_f \tag{77.12}$$

and

$$d_{t-1} = \beta (d_t + \text{trace}(C^\top P_t C)) \quad \text{with} \quad d_T = 0 \tag{77.13}$$

Recalling (77.9), the minimizers from these backward steps are

$$u_t = -F_t x_t \quad \text{where} \quad F_t := (Q + \beta B^\top P_{t+1} B)^{-1} \beta B^\top P_{t+1} A \tag{77.14}$$

These are the linear optimal control policies we *discussed above*.

In particular, the sequence of controls given by (77.14) and (77.1) solves our finite horizon LQ problem.

Rephrasing this more precisely, the sequence $\{u_0, \ldots, u_{T-1}\}$ given by

$$u_t = -F_t x_t \quad \text{with} \quad x_{t+1} = (A - BF_t) x_t + C w_{t+1} \tag{77.15}$$

for $t = 0, \ldots, T-1$ attains the minimum of (77.6) subject to our constraints.

77.4 Implementation

We will use code from lqcontrol.py in QuantEcon.py to solve finite and infinite horizon linear quadratic control problems.

In the module, the various updating, simulation and fixed point methods are wrapped in a class called `LQ`, which includes

- Instance data:
  - The required parameters Q, R, A, B and optional parameters C, β, T, R_f, N specifying a given LQ model
    - set T and R_f to `None` in the infinite horizon case
    - set `C = None` (or zero) in the deterministic case
  - the value function and policy data

    * d_t, P_t, F_t in the finite horizon case
    * d, P, F in the infinite horizon case

- Methods:
  - `update_values` — shifts d_t, P_t, F_t to their $t-1$ values via (77.12), (77.13) and (77.14)
  - `stationary_values` — computes P, d, F in the infinite horizon case
  - `compute_sequence` —- simulates the dynamics of x_t, u_t, w_t given x_0 and assuming standard normal shocks

77.4.1 An application

Early Keynesian models assumed that households have a constant marginal propensity to consume from current income.

Data contradicted the constancy of the marginal propensity to consume.

In response, Milton Friedman, Franco Modigliani and others built models based on a consumer's preference for an intertemporally smooth consumption stream.

(See, for example, [Friedman, 1956] or [Modigliani and Brumberg, 1954].)

One property of those models is that households purchase and sell financial assets to make consumption streams smoother than income streams.

The household savings problem *outlined above* captures these ideas.

The optimization problem for the household is to choose a consumption sequence in order to minimize

$$\mathbb{E}\left\{\sum_{t=0}^{T-1}\beta^t(c_t-\bar{c})^2+\beta^T q a_T^2\right\} \tag{77.16}$$

subject to the sequence of budget constraints $a_{t+1}=(1+r)a_t-c_t+y_t,\ t\geq 0$.

Here q is a large positive constant, the role of which is to induce the consumer to target zero debt at the end of her life.

(Without such a constraint, the optimal choice is to choose $c_t=\bar{c}$ in each period, letting assets adjust accordingly.)

As before we set $y_t=\sigma w_{t+1}+\mu$ and $u_t:=c_t-\bar{c}$, after which the constraint can be written as in (77.2).

We saw how this constraint could be manipulated into the LQ formulation $x_{t+1}=Ax_t+Bu_t+Cw_{t+1}$ by setting $x_t=(a_t\ 1)^\top$ and using the definitions in (77.4).

To match with this state and control, the objective function (77.16) can be written in the form of (77.6) by choosing

$$Q:=1,\quad R:=\begin{bmatrix}0&0\\0&0\end{bmatrix},\quad \text{and}\quad R_f:=\begin{bmatrix}q&0\\0&0\end{bmatrix}$$

Now that the problem is expressed in LQ form, we can proceed to the solution by applying (77.12) and (77.14).

After generating shocks $\{w_1,\dots,w_T\}$, the dynamics for assets and consumption can be simulated via (77.15).

The following figure was computed using $r=0.05, \beta=1/(1+r), \bar{c}=2, \mu=1, \sigma=0.25, T=45$ and $q=10^6$.

The shocks $\{w_t\}$ were taken to be IID and standard normal.

```
class LQModel(NamedTuple):
    r: float              # interest rate
    β: float              # discount factor
    T: int                # time horizon
    c_bar: float          # target consumption level
```

(continues on next page)

(continued from previous page)

```
    σ: float               # income shock standard deviation
    μ: float               # mean income level
    q: float               # terminal penalty weight
    Q: float               # control penalty matrix
    R: np.ndarray          # state penalty matrix
    Rf: np.ndarray         # terminal state penalty matrix
    A: np.ndarray          # state transition matrix
    B: np.ndarray          # control matrix
    C: np.ndarray          # shock matrix


def create_lq_model(r=0.05,
                    T=45,
                    c_bar=2,
                    σ=0.25,
                    μ=1,
                    q=1e6):
    β = 1 / (1 + r)

    # Formulate as an LQ problem
    Q = 1
    R = np.zeros((2, 2))
    Rf = np.zeros((2, 2))
    Rf[0, 0] = q
    A = np.array([[1 + r, -c_bar + μ],
                  [0,      1]])
    B = np.array([[-1],
                  [ 0]])
    C = np.array([[σ],
                  [0]])

    return LQModel(r=r, β=β, T=T, c_bar=c_bar, σ=σ, μ=μ,
                   q=q, Q=Q, R=R, Rf=Rf, A=A, B=B, C=C)

def simulate_and_plot(model):
    # Unpack model
    r, β, T, c_bar, σ, μ, q, Q, R, Rf, A, B, C = model

    # Compute solutions and simulate
    lq = LQ(Q, R, A, B, C, beta=β, T=T, Rf=Rf)
    x0 = (0, 1)
    xp, up, wp = lq.compute_sequence(x0)

    # Convert back to assets, consumption and income
    assets = xp[0, :]            # a_t
    c = up.flatten() + c_bar     # c_t
    income = σ * wp[0, 1:] + μ   # y_t

    # Plot results
    n_rows = 2
    fig, axes = plt.subplots(n_rows, 1, figsize=(12, 10))

    plt.subplots_adjust(hspace=0.5)

    bbox = (0., 1.02, 1., .102)
    legend_args = {'bbox_to_anchor': bbox, 'loc': 3, 'mode': 'expand'}
    p_args = {'lw': 2, 'alpha': 0.7}
```

(continues on next page)

(continued from previous page)

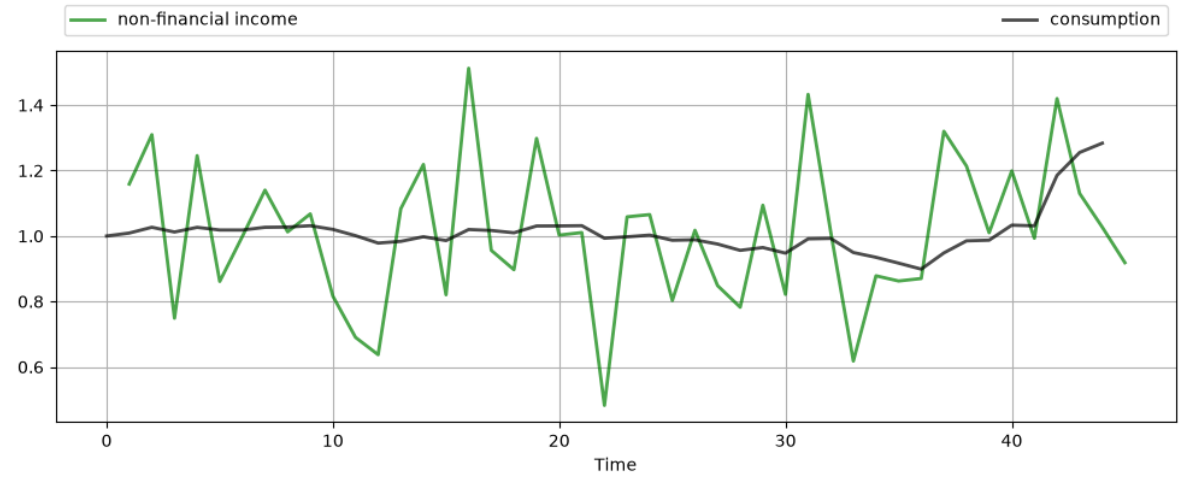
```
    axes[0].plot(range(1, T+1), income, 'g-', label="non-financial income",
                 **p_args)
    axes[0].plot(range(T), c, 'k-', label="consumption", **p_args)

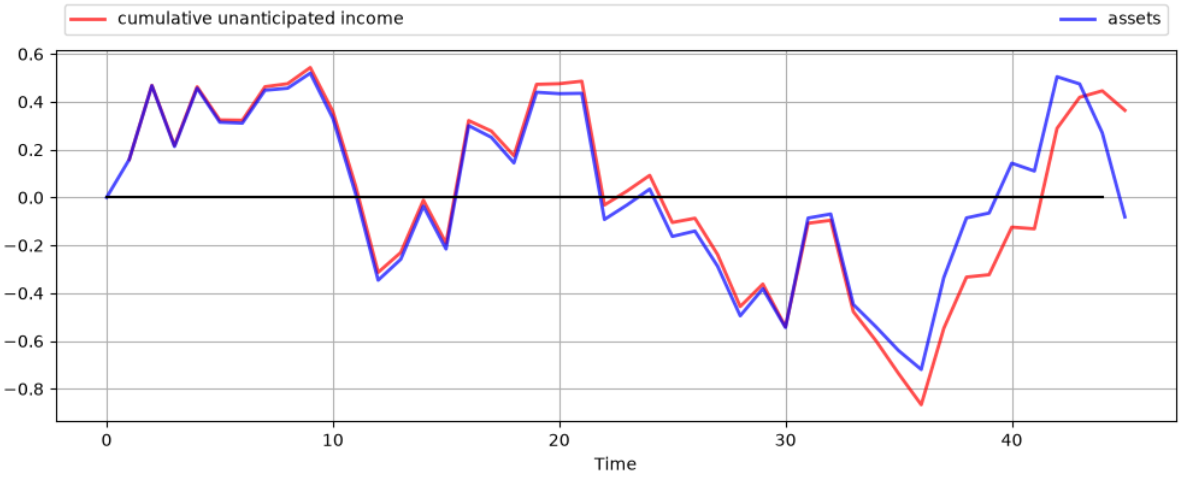
    axes[1].plot(range(1, T+1), np.cumsum(income - μ), 'r-',
                 label="cumulative unanticipated income", **p_args)
    axes[1].plot(range(T+1), assets, 'b-', label="assets", **p_args)
    axes[1].plot(range(T), np.zeros(T), 'k-')

    for ax in axes:
        ax.grid()
        ax.set_xlabel('Time')
        ax.legend(ncol=2, **legend_args)

    plt.show()

# Create model instance and simulate
model = create_lq_model()
simulate_and_plot(model)
```





The top panel shows the time path of consumption c_t and income y_t in the simulation.

As anticipated by the discussion on consumption smoothing, the time path of consumption is much smoother than that for income.

(But note that consumption becomes more irregular towards the end of life, when the zero final asset requirement impinges more on consumption choices.)

The second panel in the figure shows that the time path of assets a_t is closely correlated with cumulative unanticipated income, where the latter is defined as

$$z_t := \sum_{j=0}^{t} \sigma w_t$$

A key message is that unanticipated windfall gains are saved rather than consumed, while unanticipated negative shocks are met by reducing assets.

(Again, this relationship breaks down towards the end of life due to the zero final asset requirement.)

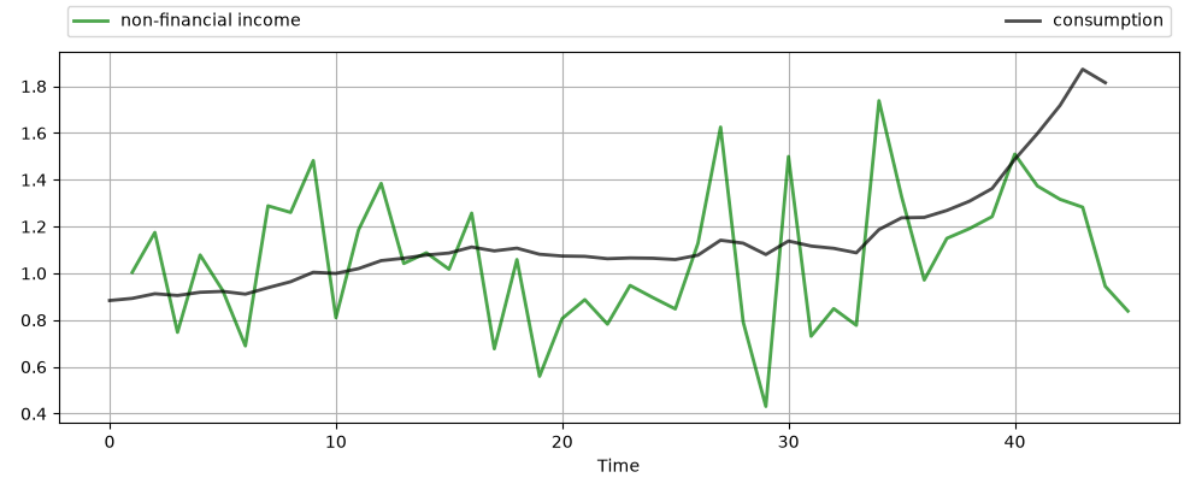
These results are relatively robust to changes in parameters.

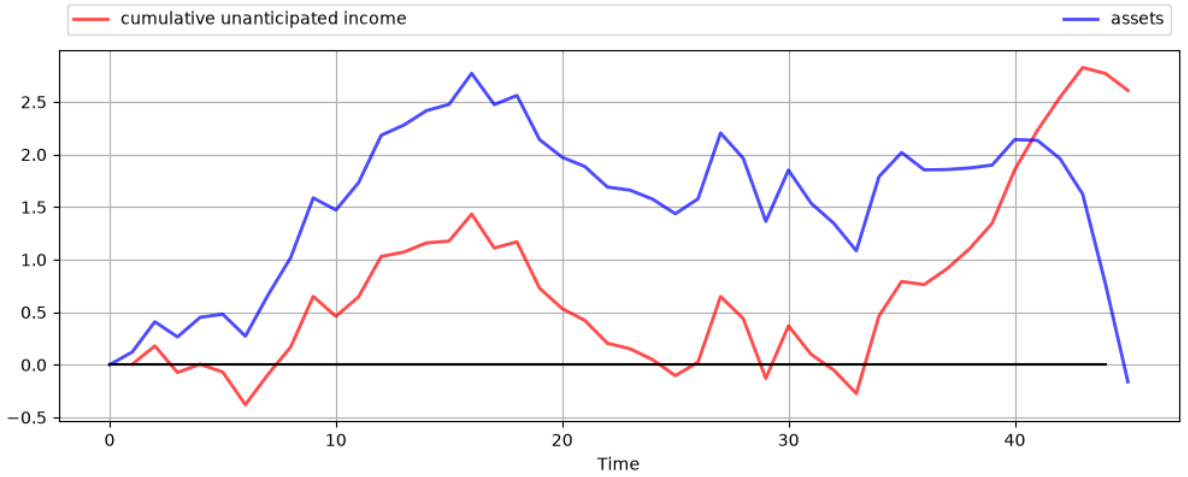
For example, let's increase β from $1/(1+r) \approx 0.952$ to 0.96 while keeping other parameters fixed.

This consumer is slightly more patient than the last one, and hence puts relatively more weight on later consumption values.

Since `LQModel` is a `NamedTuple`, we can obtain the new model by copying the old one with a single field changed.

```
simulate_and_plot(model._replace(β=0.96))
```





We now have a slowly rising consumption stream and a hump-shaped build-up of assets in the middle periods to fund rising consumption.

However, the essential features are the same: consumption is smooth relative to income, and assets are strongly positively correlated with cumulative unanticipated income.

77.5 Extensions and comments

Let's now consider a number of standard extensions to the LQ problem treated above.

77.5.1 Time-varying parameters

In some settings, it can be desirable to allow A, B, C, R and Q to depend on t.

For the sake of simplicity, we've chosen not to treat this extension in our implementation given below.

However, the loss of generality is not as large as you might first imagine.

In fact, we can tackle many models with time-varying parameters by suitable choice of state variables.

One illustration is given *below*.

For further examples and a more systematic treatment, see [Hansen and Sargent, 2013], section 2.4.

77.5.2 Adding a cross-product term

In some LQ problems, preferences include a cross-product term $u_t^\top N x_t$, so that the objective function becomes

$$\mathbb{E}\left\{\sum_{t=0}^{T-1}\beta^t(x_t^\top R x_t + u_t^\top Q u_t + 2u_t^\top N x_t) + \beta^T x_T^\top R_f x_T\right\} \tag{77.17}$$

Our results extend to this case in a straightforward way.

The sequence $\{P_t\}$ from (77.12) becomes

$$P_{t-1} = R - (\beta B^\top P_t A + N)^\top (Q + \beta B^\top P_t B)^{-1}(\beta B^\top P_t A + N) + \beta A^\top P_t A \quad \text{with} \quad P_T = R_f \tag{77.18}$$

The policies in (77.14) are modified to

$$u_t = -F_t x_t \quad \text{where} \quad F_t := (Q + \beta B^\top P_{t+1} B)^{-1}(\beta B^\top P_{t+1} A + N) \tag{77.19}$$

The sequence $\{d_t\}$ is unchanged from (77.13).

We leave interested readers to confirm these results (the calculations are long but not overly difficult).

77.5.3 Infinite horizon

Finally, we consider the infinite horizon case, with *cross-product term*, unchanged dynamics and objective function given by

$$\mathbb{E}\left\{\sum_{t=0}^{\infty}\beta^t(x_t^\top R x_t + u_t^\top Q u_t + 2u_t^\top N x_t)\right\} \tag{77.20}$$

In the infinite horizon case, optimal policies can depend on time only if time itself is a component of the state vector x_t.

In other words, there exists a fixed matrix F such that $u_t = -Fx_t$ for all t.

That decision rules are constant over time is intuitive — after all, the decision-maker faces the same infinite horizon at every stage, with only the current state changing.

Not surprisingly, P and d are also constant.

The stationary matrix P is the solution to the discrete-time algebraic Riccati equation.

$$P = R - (\beta B^\top P A + N)^\top (Q + \beta B^\top P B)^{-1}(\beta B^\top P A + N) + \beta A^\top P A \tag{77.21}$$

Equation (77.21) is also called the *LQ Bellman equation*, and the map that sends a given P into the right-hand side of (77.21) is called the *LQ Bellman operator*.

The stationary optimal policy for this model is

$$u = -Fx \quad \text{where} \quad F = (Q + \beta B^\top P B)^{-1}(\beta B^\top P A + N) \tag{77.22}$$

The sequence $\{d_t\}$ from (77.13) is replaced by the constant value

$$d := \text{trace}(C^\top P C)\frac{\beta}{1-\beta} \tag{77.23}$$

The state evolves according to the time-homogeneous process $x_{t+1} = (A - BF)x_t + Cw_{t+1}$.

An example infinite horizon problem is treated *below*.

77.5.4 Certainty equivalence

Linear quadratic control problems of the class discussed above have the property of *certainty equivalence*.

By this, we mean that the optimal policy F is not affected by the parameters in C, which specify the shock process.

This can be confirmed by inspecting (77.22) or (77.19).

It follows that we can ignore uncertainty when solving for optimal behavior, and plug it back in when examining optimal state dynamics.

77.6 Further applications

77.6.1 Application 1: Age-dependent income process

Previously we studied a permanent income model that generated consumption smoothing.

One unrealistic feature of that model is the assumption that the mean of the random income process does not depend on the consumer's age.

A more realistic income profile is one that rises in early working life, peaks towards the middle and maybe declines toward the end of working life and falls more during retirement.

In this section, we will model this rise and fall as a symmetric inverted "U" using a polynomial in age.

As before, the consumer seeks to minimize

$$\mathbb{E}\left\{\sum_{t=0}^{T-1}\beta^t(c_t - \bar{c})^2 + \beta^T q a_T^2\right\} \tag{77.24}$$

subject to $a_{t+1} = (1+r)a_t - c_t + y_t,\ t \geq 0$.

For income we now take $y_t = p(t) + \sigma w_{t+1}$ where $p(t) := m_0 + m_1 t + m_2 t^2$.

(In *the next section* we employ some tricks to implement a more sophisticated model.)

The coefficients m_0, m_1, m_2 are chosen such that $p(0) = 0, p(T/2) = \mu$, and $p(T) = 0$.

You can confirm that the specification $m_0 = 0, m_1 = T\mu/(T/2)^2, m_2 = -\mu/(T/2)^2$ satisfies these constraints.

To put this into an LQ setting, consider the budget constraint, which becomes

$$a_{t+1} = (1+r)a_t - u_t - \bar{c} + m_1 t + m_2 t^2 + \sigma w_{t+1} \tag{77.25}$$

The fact that a_{t+1} is a linear function of $(a_t, 1, t, t^2)$ suggests taking these four variables as the state vector x_t.

Once a good choice of state and control (recall $u_t = c_t - \bar{c}$) has been made, the remaining specifications fall into place relatively easily.

Thus, for the dynamics we set

$$x_t := \begin{bmatrix} a_t \\ 1 \\ t \\ t^2 \end{bmatrix}, \quad A := \begin{bmatrix} 1+r & -\bar{c} & m_1 & m_2 \\ 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 0 \\ 0 & 1 & 2 & 1 \end{bmatrix}, \quad B := \begin{bmatrix} -1 \\ 0 \\ 0 \\ 0 \end{bmatrix}, \quad C := \begin{bmatrix} \sigma \\ 0 \\ 0 \\ 0 \end{bmatrix} \tag{77.26}$$

If you expand the expression $x_{t+1} = Ax_t + Bu_t + Cw_{t+1}$ using this specification, you will find that assets follow (77.25) as desired and that the other state variables also update appropriately.

To implement preference specification (77.24) we take

$$Q := 1, \quad R := \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad R_f := \begin{bmatrix} q & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \tag{77.27}$$

The next figure shows a simulation of consumption and assets computed using the `compute_sequence` method of `lqcontrol.py` with initial assets set to zero.

Once again, smooth consumption is a dominant feature of the sample paths.

The asset path exhibits dynamics consistent with standard life cycle theory.

Exercise 77.7.1 gives the full set of parameters used here and asks you to replicate the figure.

77.6.2 Application 2: A permanent income model with retirement

In the *previous application*, we generated income dynamics with an inverted U shape using polynomials and placed them in an LQ framework.

It is arguably the case that this income process still contains unrealistic features.

A more common earning profile is where

1. income grows over working life, fluctuating around an increasing trend, with growth flattening off in later years
2. retirement follows, with lower but relatively stable (non-financial) income

Letting K be the retirement date, we can express these income dynamics by

$$y_t = \begin{cases} p(t) + \sigma w_{t+1} & \text{if } t \leq K \\ s & \text{otherwise} \end{cases} \tag{77.28}$$

Here

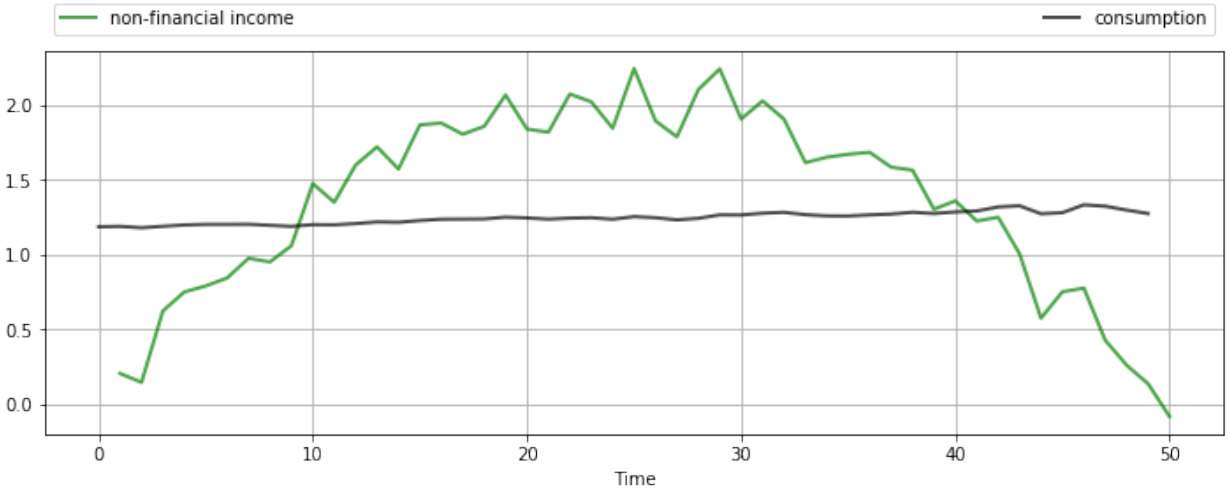
- $p(t) := m_1 t + m_2 t^2$ with the coefficients m_1, m_2 chosen such that $p(K) = \mu$ and $p(0) = p(2K) = 0$
- s is retirement income

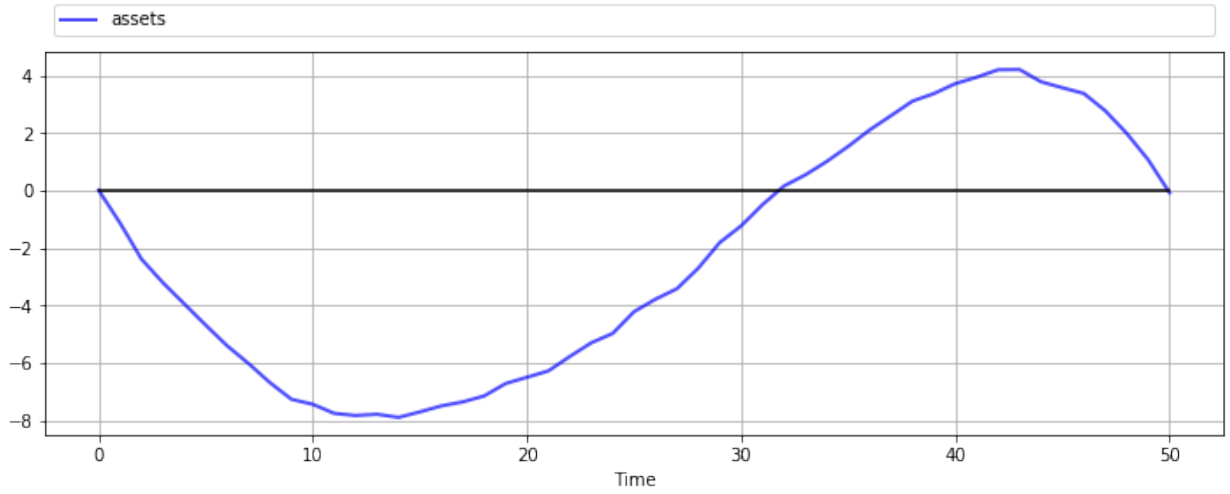
We suppose that preferences are unchanged and given by (77.16).

The budget constraint is also unchanged and given by $a_{t+1} = (1+r)a_t - c_t + y_t$.

Our aim is to solve this problem and simulate paths using the LQ techniques described in this lecture.

In fact, this is a nontrivial problem, as the kink in the dynamics (77.28) at K makes it very difficult to express the law of motion as a fixed-coefficient linear system.


non-financial income
consumption
2.0
1.5
1.0
0.5
0.0
0
10
20
30
40
50
Time


assets
4
2
0
−2
−4
−6
−8
0
10
20
30
40
50
Time

However, we can still use our LQ methods here by suitably linking two-component LQ problems.

These two LQ problems describe the consumer's behavior during her working life (`lq_working`) and retirement (`lq_retired`).

(This is possible because, in the two separate periods of life, the respective income processes [polynomial trend and constant] each fit the LQ framework.)

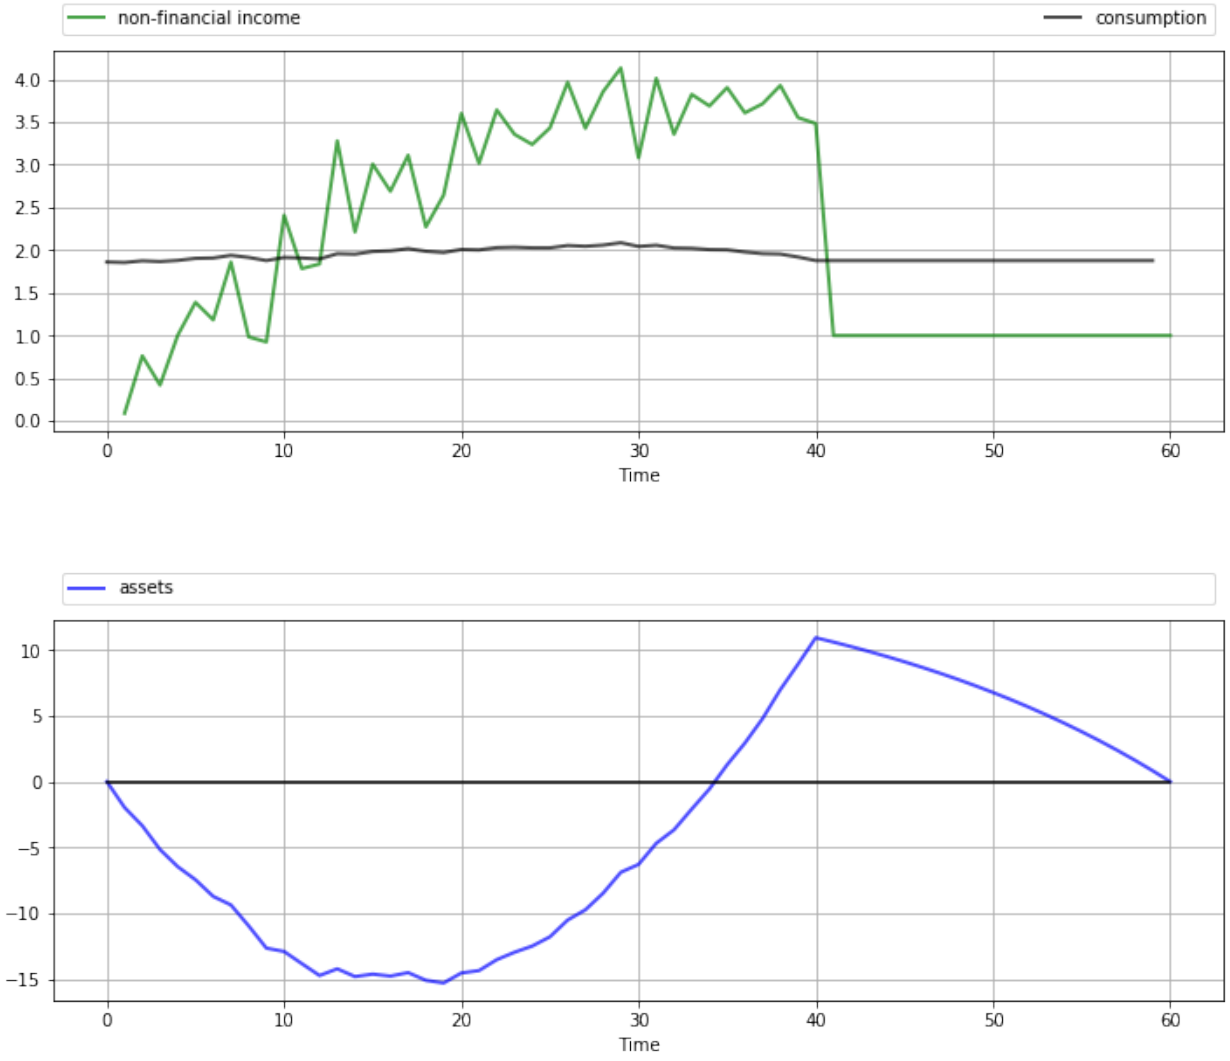
The basic idea is that although the whole problem is not a single time-invariant LQ problem, it is still a dynamic programming problem, and hence we can use appropriate Bellman equations at every stage.

Based on this logic, we can

1. solve `lq_retired` by the usual backward induction procedure, iterating back to the start of retirement.
2. take the start-of-retirement value function generated by this process, and use it as the terminal condition R_f to feed into the `lq_working` specification.
3. solve `lq_working` by backward induction from this choice of R_f, iterating back to the start of working life.

This process gives the entire life-time sequence of value functions and optimal policies.

The next figure shows one simulation based on this procedure.



The full set of parameters used in the simulation is discussed in Exercise 77.7.2, where you are asked to replicate the figure.

Once again, the dominant feature observable in the simulation is consumption smoothing.

The asset path fits well with standard life cycle theory, with dissaving early in life followed by later saving.

Assets peak at retirement and subsequently decline.

77.6.3 Application 3: Monopoly with adjustment costs

Consider a monopolist facing stochastic inverse demand function

$$p_t = a_0 - a_1 q_t + d_t$$

Here q_t is output, and the demand shock d_t follows

$$d_{t+1} = \rho d_t + \sigma w_{t+1}$$

where $\{w_t\}$ is IID and standard normal.

The monopolist maximizes the expected discounted sum of present and future profits

$$\mathbb{E}\left\{\sum_{t=0}^{\infty} \beta^t \pi_t\right\} \quad \text{where} \quad \pi_t := p_t q_t - c q_t - \gamma (q_{t+1} - q_t)^2 \tag{77.29}$$

Here

- $\gamma (q_{t+1} - q_t)^2$ represents adjustment costs
- c is average cost of production

This can be formulated as an LQ problem and then solved and simulated, but first let's study the problem and try to get some intuition.

One way to start thinking about the problem is to consider what would happen if $\gamma = 0$.

Without adjustment costs there is no intertemporal trade-off, so the monopolist will choose output to maximize current profit in each period.

It's not difficult to show that profit-maximizing output is

$$\bar{q}_t := \frac{a_0 - c + d_t}{2a_1}$$

In light of this discussion, what we might expect for general γ is that

- if γ is close to zero, then q_t will track the time path of $\bar{q}_t$ relatively closely.
- if γ is larger, then q_t will be smoother than $\bar{q}_t$, as the monopolist seeks to avoid adjustment costs.

This intuition turns out to be correct.

The following figures show simulations produced by solving the corresponding LQ problem.

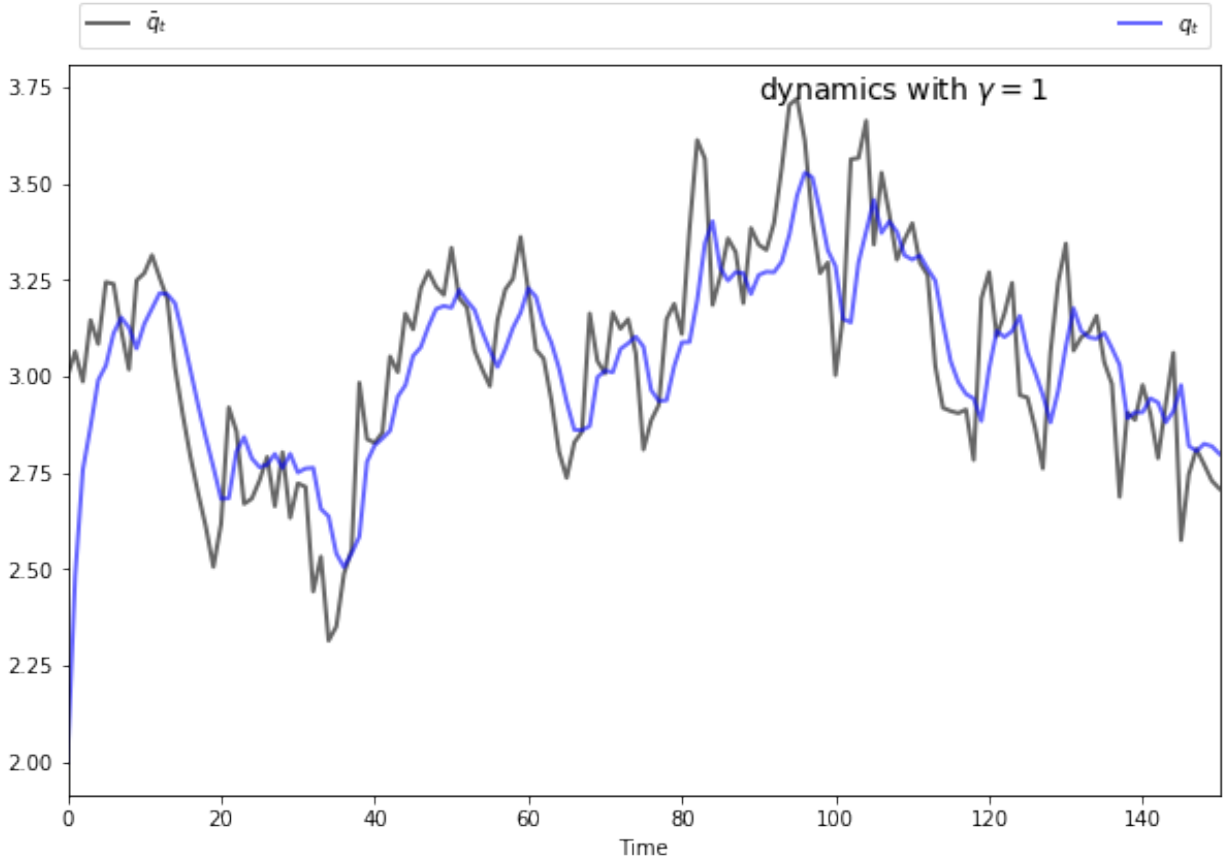
The only difference in parameters across the figures is the size of γ

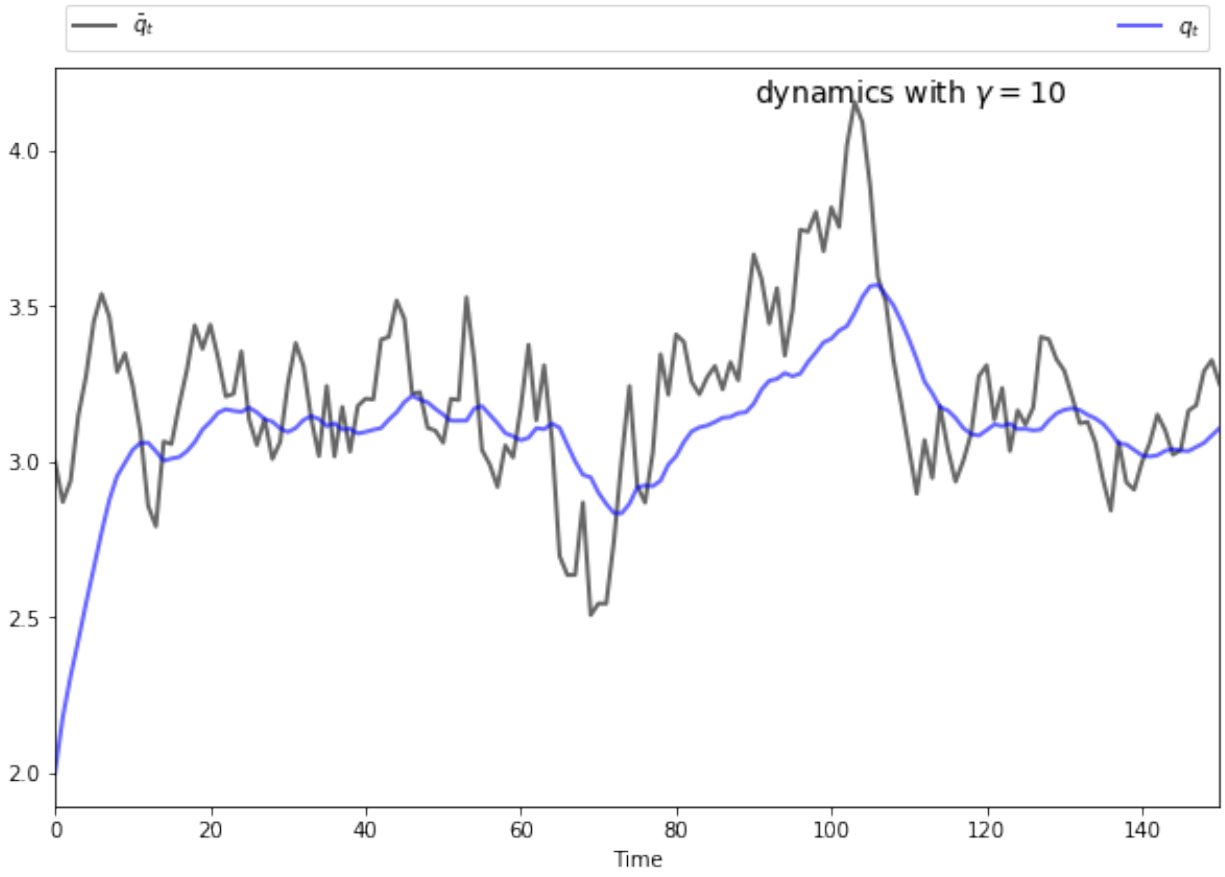
To produce these figures we converted the monopolist problem into an LQ problem.

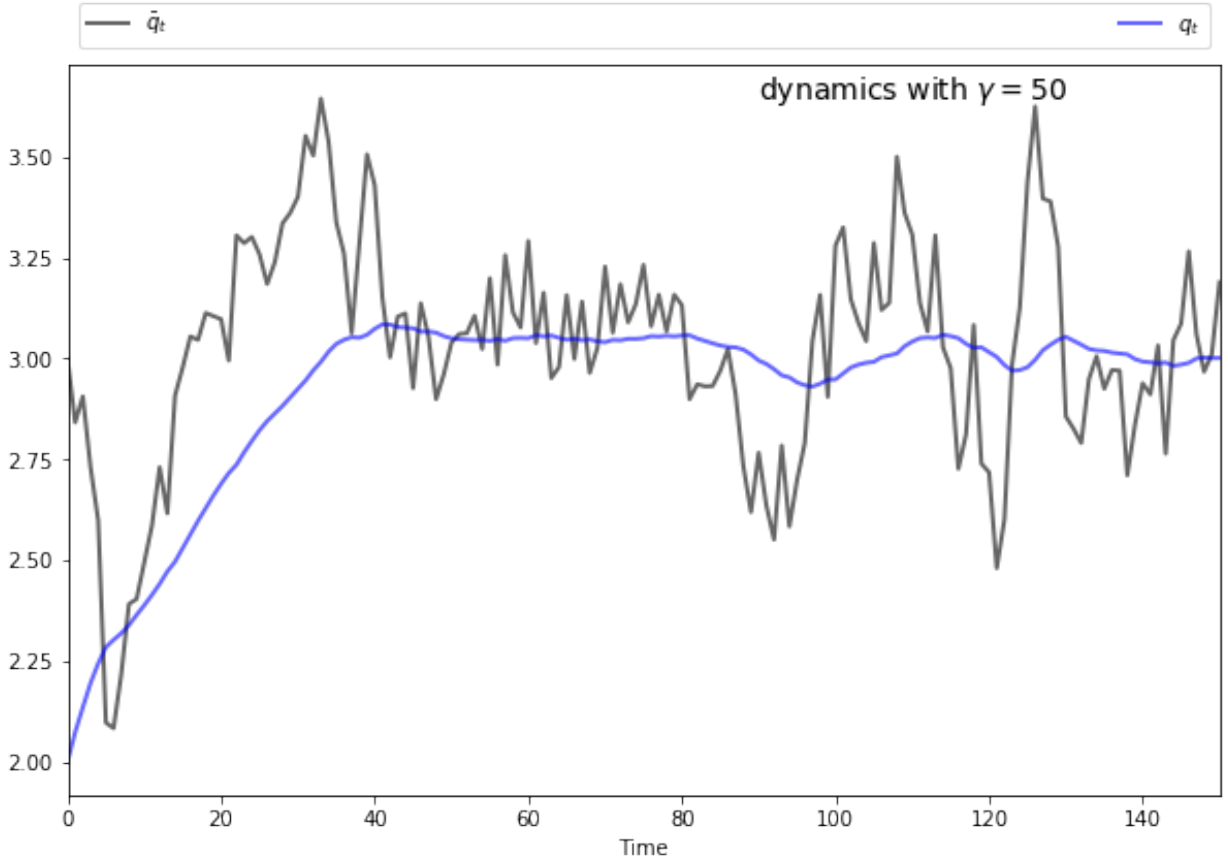
The key to this conversion is to choose the right state — which can be a bit of an art.

Here we take $x_t = (\bar{q}_t \;\; q_t \;\; 1)^\top$, while the control is chosen as $u_t = q_{t+1} - q_t$.

We also manipulated the profit function slightly.


$\bar{q}_t$
q_t
dynamics with $\gamma = 1$
3.75
3.50
3.25
3.00
2.75
2.50
2.25
2.00
0
20
40
60
80
100
120
140
Time


$\tilde{q}_t$
q_t
dynamics with $\gamma = 10$
4.0
3.5
3.0
2.5
2.0
0
20
40
60
80
100
120
140
Time


$\bar{q}_t$
q_t
dynamics with $\gamma = 50$
3.50
3.25
3.00
2.75
2.50
2.25
2.00
0
20
40
60
80
100
120
140
Time

In (77.29), current profits are $\pi_t := p_t q_t - c q_t - \gamma (q_{t+1} - q_t)^2$.

Let's now replace π_t in (77.29) with $\hat{\pi}_t := \pi_t - a_1 \bar{q}_t^2$.

This makes no difference to the solution, since $a_1 \bar{q}_t^2$ does not depend on the controls.

(In fact, we are just adding a constant term to (77.29), and optimizers are not affected by constant terms.)

The reason for making this substitution is that, as you will be able to verify, $\hat{\pi}_t$ reduces to the simple quadratic

$$\hat{\pi}_t = -a_1 (q_t - \bar{q}_t)^2 - \gamma u_t^2$$

After negation to convert to a minimization problem, the objective becomes

$$\min \mathbb{E} \sum_{t=0}^{\infty} \beta^t \left\{ a_1 (q_t - \bar{q}_t)^2 + \gamma u_t^2 \right\} \tag{77.30}$$

It's now relatively straightforward to find R and Q such that (77.30) can be written as (77.20).

Furthermore, the matrices A, B and C from (77.1) can be found by writing down the dynamics of each element of the state.

Exercise 77.7.3 asks you to complete this process, and reproduce the preceding figures.

77.7 Exercises

Exercise 77.7.1

Replicate the figure with polynomial income *shown above*.

The parameters are $r = 0.05, \beta = 1/(1+r), \bar{c} = 1.5, \mu = 2, \sigma = 0.15, T = 50$ and $q = 10^4$.

Solution

Here's one solution.

We use some fancy plot commands to get a certain style – feel free to use simpler ones.

The model is an LQ permanent income / life-cycle model with hump-shaped income

$$y_t = m_1 t + m_2 t^2 + \sigma w_{t+1}$$

where $\{w_t\}$ is IID $N(0,1)$ and the coefficients m_1 and m_2 are chosen so that $p(t) = m_1 t + m_2 t^2$ has an inverted U shape with

- $p(0) = 0, p(T/2) = \mu$, and
- $p(T) = 0$

```
# Model parameters
r = 0.05
β = 1/(1 + r)
T = 50
c_bar = 1.5
σ = 0.15
μ = 2
q = 1e4
```

```
m1 = T * (μ/(T/2)**2)
m2 = -(μ/(T/2)**2)

# Formulate as an LQ problem
Q = 1
R = np.zeros((4, 4))
Rf = np.zeros((4, 4))
Rf[0, 0] = q
A = [[1 + r, -c_bar, m1, m2],
     [0,           1,  0,  0],
     [0,           1,  1,  0],
     [0,           1,  2,  1]]
B = [[-1],
     [ 0],
     [ 0],
     [ 0]]
C = [[σ],
     [0],
     [0],
     [0]]

# Compute solutions and simulate
lq = LQ(Q, R, A, B, C, beta=β, T=T, Rf=Rf)
x0 = (0, 1, 0, 0)
xp, up, wp = lq.compute_sequence(x0)

# Convert results back to assets, consumption and income
ap = xp[0, :]               # Assets
c = up.flatten() + c_bar    # Consumption
time = np.arange(1, T+1)
income = σ * wp[0, 1:] + m1 * time + m2 * time**2  # Income


# Plot results
n_rows = 2
fig, axes = plt.subplots(n_rows, 1, figsize=(12, 10))

plt.subplots_adjust(hspace=0.5)

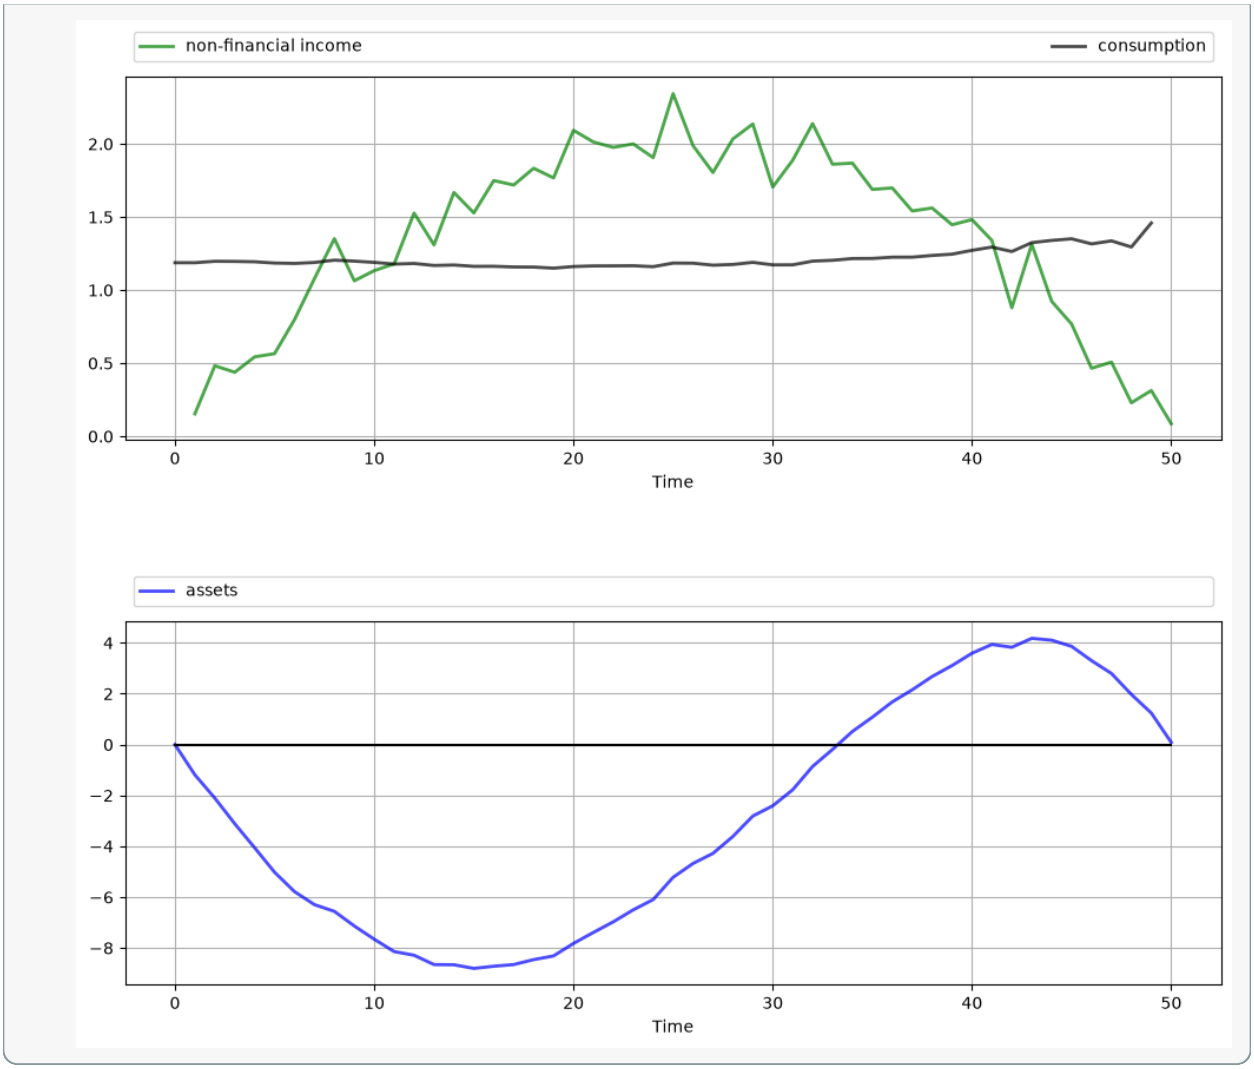
bbox = (0., 1.02, 1., .102)
legend_args = {'bbox_to_anchor': bbox, 'loc': 3, 'mode': 'expand'}
p_args = {'lw': 2, 'alpha': 0.7}

axes[0].plot(range(1, T+1), income, 'g-', label="non-financial income",
             **p_args)
axes[0].plot(range(T), c, 'k-', label="consumption", **p_args)

axes[1].plot(range(T+1), ap.flatten(), 'b-', label="assets", **p_args)
axes[1].plot(range(T+1), np.zeros(T+1), 'k-')

for ax in axes:
    ax.grid()
    ax.set_xlabel('Time')
    ax.legend(ncol=2, **legend_args)

plt.show()
```



Exercise 77.7.2

Replicate the figure on work and retirement *shown above*.

The parameters are $r = 0.05, \beta = 1/(1+r), \bar{c} = 4, \mu = 4, \sigma = 0.35, K = 40, T = 60, s = 1$ and $q = 10^4$.

To understand the overall procedure, carefully read the section containing that figure.

Hint

First, in order to make our approach work, we must ensure that both LQ problems have the same state variables and control.
As with previous applications, the control can be set to $u_t = c_t - \bar{c}$.
For `lq_working`, x_t, A, B, C can be chosen as in (77.26).

- Recall that m_1, m_2 are chosen so that $p(K) = \mu$ and $p(2K) = 0$.

For `lq_retired`, use the same definition of x_t and u_t, but modify A, B, C to correspond to constant income $y_t = s$.
For `lq_retired`, set preferences as in (77.27).
For `lq_working`, preferences are the same, except that R_f should be replaced by the final value function that emerges from iterating `lq_retired` back to the start of retirement.
With some careful footwork, the simulation can be generated by patching together the simulations from these two separate models.

Solution

This is a permanent income / life-cycle model with polynomial growth in income over working life followed by a fixed retirement income.

The model is solved by combining two LQ programming problems as described in the lecture.

```
# Model parameters
r = 0.05
β = 1/(1 + r)
T = 60
K = 40
c_bar = 4
σ = 0.35
μ = 4
q = 1e4
s = 1
m1 = 2 * μ/K
m2 = -μ/K**2

# Formulate LQ problem 1 (retirement)
Q = 1
R = np.zeros((4, 4))
Rf = np.zeros((4, 4))
Rf[0, 0] = q
A = [[1 + r, s - c_bar, 0, 0],
     [0,               1, 0, 0],
     [0,               1, 1, 0],
     [0,               1, 2, 1]]
B = [[-1],
     [ 0],
     [ 0],
     [ 0]]
C = [[0],
     [0],
     [0],
     [0]]

# Initialize LQ instance for retired agent
lq_retired = LQ(Q, R, A, B, C, beta=β, T=T-K, Rf=Rf)
```

```
# Iterate back to start of retirement, record final value function
for i in range(T-K):
    lq_retired.update_values()
Rf2 = lq_retired.P

# Formulate LQ problem 2 (working life)
R = np.zeros((4, 4))
A = [[1 + r, -c_bar, m1, m2],
     [0,           1,  0,  0],
     [0,           1,  1,  0],
     [0,           1,  2,  1]]
B = [[-1],
     [ 0],
     [ 0],
     [ 0]]
C = [[σ],
     [0],
     [0],
     [0]]

# Set up working life LQ instance with terminal Rf from lq_retired
lq_working = LQ(Q, R, A, B, C, beta=β, T=K, Rf=Rf2)

# Simulate working state / control paths
x0 = (0, 1, 0, 0)
xp_w, up_w, wp_w = lq_working.compute_sequence(x0)
# Simulate retirement paths (note the initial condition)
xp_r, up_r, wp_r = lq_retired.compute_sequence(xp_w[:, K])

# Convert results back to assets, consumption and income
xp = np.column_stack((xp_w, xp_r[:, 1:]))
assets = xp[0, :]                    # Assets

up = np.column_stack((up_w, up_r))
c = up.flatten() + c_bar             # Consumption

time = np.arange(1, K+1)
income_w = σ * wp_w[0, 1:K+1] + m1 * time + m2 * time**2  # Income
income_r = np.full(T-K, s)
income = np.concatenate((income_w, income_r))

# Plot results
n_rows = 2
fig, axes = plt.subplots(n_rows, 1, figsize=(12, 10))

plt.subplots_adjust(hspace=0.5)

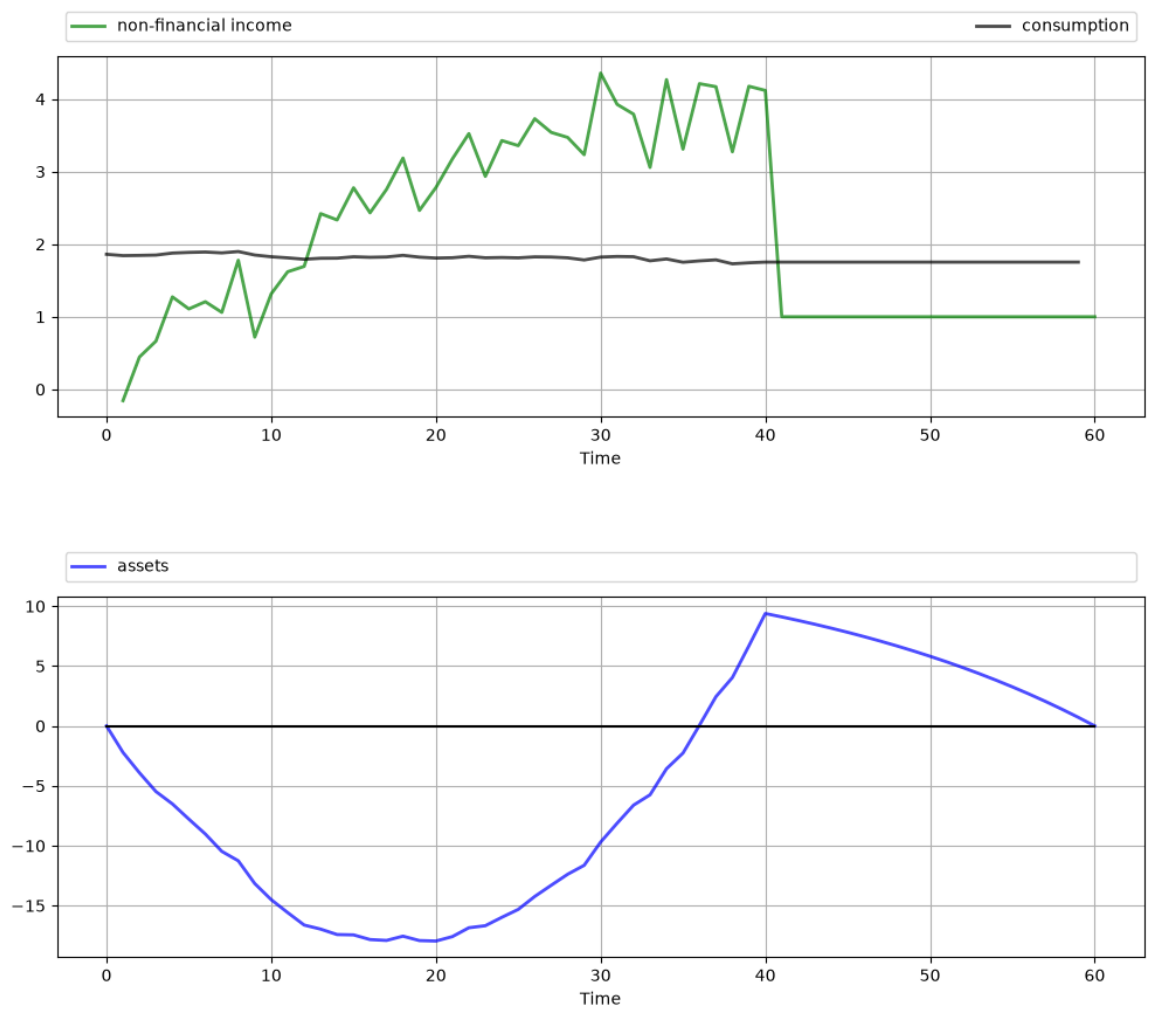
bbox = (0., 1.02, 1., .102)
legend_args = {'bbox_to_anchor': bbox, 'loc': 3, 'mode': 'expand'}
p_args = {'lw': 2, 'alpha': 0.7}

axes[0].plot(range(1, T+1), income, 'g-', label="non-financial income",
             **p_args)
axes[0].plot(range(T), c, 'k-', label="consumption", **p_args)

axes[1].plot(range(T+1), assets, 'b-', label="assets", **p_args)
axes[1].plot(range(T+1), np.zeros(T+1), 'k-')
```

```
for ax in axes:
    ax.grid()
    ax.set_xlabel('Time')
    ax.legend(ncol=2, **legend_args)

plt.show()
```



Exercise 77.7.3

Reproduce the figures from the monopolist application *given above*.

For parameters, use $a_0 = 5, a_1 = 0.5, \sigma = 0.15, \rho = 0.9, \beta = 0.95$ and $c = 2$, while γ varies between 1 and 50 (see figures).

Solution

The first task is to find the matrices A, B, C, Q, R that define the LQ problem.

Recall that $x_t = (\bar{q}_t \;\; q_t \;\; 1)^\top$, while $u_t = q_{t+1} - q_t$.

Letting $m_0 := (a_0 - c)/2a_1$ and $m_1 := 1/2a_1$, we can write $\bar{q}_t = m_0 + m_1 d_t$, and then, with some manipulation

$$\bar{q}_{t+1} = m_0(1-\rho) + \rho\bar{q}_t + m_1\sigma w_{t+1}$$

By our definition of u_t, the dynamics of q_t are $q_{t+1} = q_t + u_t$.

Using these facts you should be able to build the correct A, B, C matrices (and then check them against those found in the solution code below).

Suitable R, Q matrices can be found by inspecting the objective function, which we repeat here for convenience:

$$\min \mathbb{E}\left\{\sum_{t=0}^{\infty} \beta^t a_1(q_t - \bar{q}_t)^2 + \gamma u_t^2\right\}$$

Our solution code is

```
# Model parameters
a0 = 5
a1 = 0.5
σ = 0.15
ρ = 0.9
γ = 1
β = 0.95
c = 2
T = 120

# Useful constants
m0 = (a0-c)/(2 * a1)
m1 = 1/(2 * a1)

# Formulate LQ problem
Q = γ
R = [[ a1, -a1,  0],
     [-a1,  a1,  0],
     [  0,   0,  0]]
A = [[ρ, 0, m0 * (1 - ρ)],
     [0, 1,           0],
     [0, 0,           1]]

B = [[0],
     [1],
     [0]]
C = [[m1 * σ],
     [     0],
     [     0]]

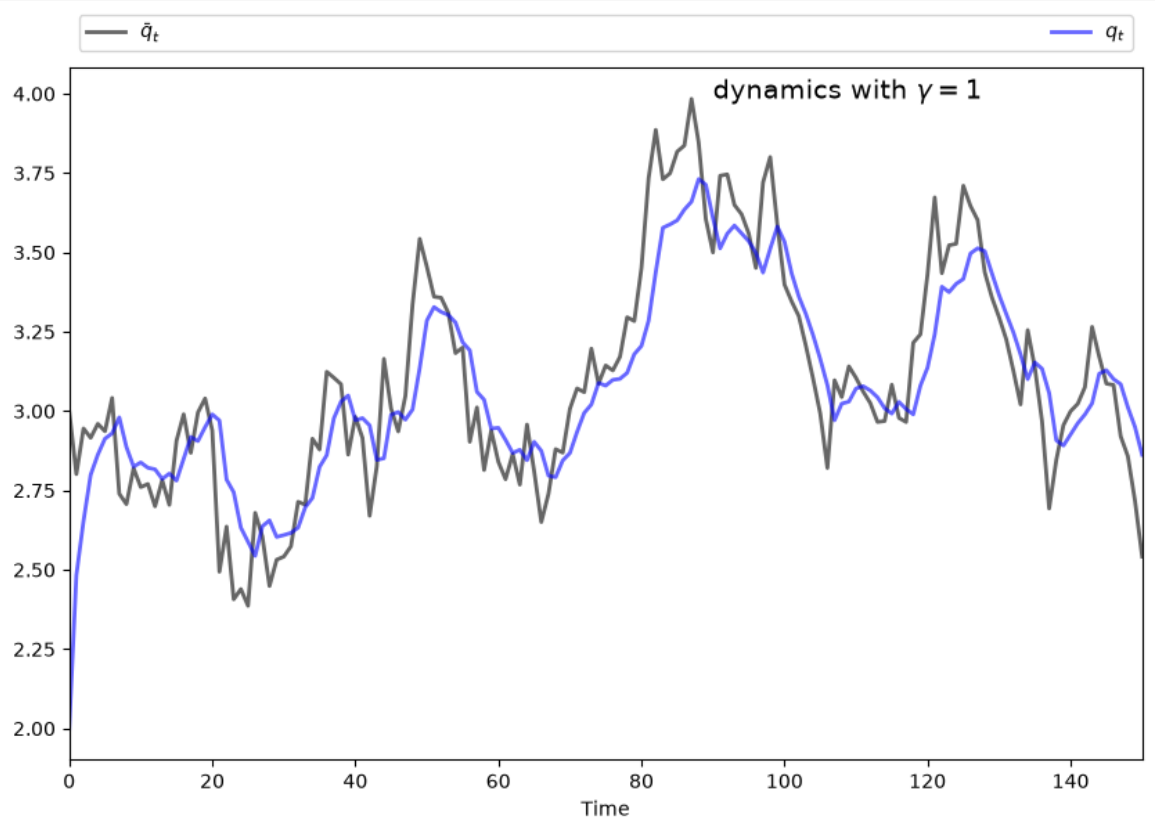
lq = LQ(Q, R, A, B, C=C, beta=β)

# Simulate state / control paths
x0 = (m0, 2, 1)
xp, up, wp = lq.compute_sequence(x0, ts_length=150)
q_bar = xp[0, :]
q = xp[1, :]

# Plot simulation results
fig, ax = plt.subplots(figsize=(10, 6.5))
```

```
# Some fancy plotting stuff -- simplify if you prefer
bbox = (0., 1.01, 1., .101)
legend_args = {'bbox_to_anchor': bbox, 'loc': 3, 'mode': 'expand'}
p_args = {'lw': 2, 'alpha': 0.6}

time = range(len(q))
ax.set(xlabel='Time', xlim=(0, max(time)))
ax.plot(time, q_bar, 'k-', lw=2, alpha=0.6, label=r'$\bar q_t$')
ax.plot(time, q, 'b-', lw=2, alpha=0.6, label='$q_t$')
ax.legend(ncol=2, **legend_args)
s = fr'dynamics with $\gamma = {γ}$'
ax.text(max(time) * 0.6, 1 * q_bar.max(), s, fontsize=14)
plt.show()
```



CHAPTER

SEVENTYEIGHT

LAGRANGIAN FOR LQ CONTROL

```
!pip install quantecon
```

```
import numpy as np
from quantecon import LQ
from scipy.linalg import schur
```

78.1 Overview

This is a sequel to this lecture *linear quadratic dynamic programming*

It can also be regarded as presenting **invariant subspace** techniques that extend ones that we encountered earlier in this lecture *stability in linear rational expectations models*

We present a Lagrangian formulation of an infinite horizon linear quadratic undiscounted dynamic programming problem.

Such a problem is also sometimes called an optimal linear regulator problem.

A Lagrangian formulation

- carries insights about connections between stability and optimality
- is the basis for fast algorithms for solving Riccati equations
- opens the way to constructing solutions of dynamic systems that don't come directly from an intertemporal optimization problem

A key tool in this lecture is the concept of an $n \times n$ **symplectic** matrix.

A symplectic matrix has eigenvalues that occur in **reciprocal pairs**, meaning that if $\lambda_i \in (-1, 1)$ is an eigenvalue, then so is λ_i^{-1}.

This reciprocal pairs property of the eigenvalues of a matrix is a tell-tale sign that the matrix describes the joint dynamics of a system of equations describing the **states** and **costates** that constitute first-order necessary conditions for solving an undiscounted linear-quadratic infinite-horizon optimization problem.

The symplectic matrix that will interest us describes the first-order dynamics of **state** and **co-state** vectors of an optimally controlled system.

In focusing on eigenvalues and eigenvectors of this matrix, we capitalize on an analysis of **invariant subspaces.**

These invariant subspace formulations of LQ dynamic programming problems provide a bridge between recursive (i.e., dynamic programming) formulations and classical formulations of linear control and linear filtering problems that make use of related matrix decompositions (see for example this lecture and this lecture).

While most of this lecture focuses on undiscounted problems, later sections describe handy ways of transforming discounted problems to undiscounted ones.

The techniques in this lecture will prove useful when we study Stackelberg and Ramsey problem in this lecture.

78.2 Undiscounted LQ DP Problem

The problem is to choose a sequence of controls $\{u_t\}_{t=0}^{\infty}$ to maximize the criterion

$$-\sum_{t=0}^{\infty}\{x_t' R x_t + u_t' Q u_t\}$$

subject to $x_{t+1} = Ax_t + Bu_t$, where x_0 is a given initial state vector.

Here x_t is an $(n \times 1)$ vector of state variables, u_t is a $(k \times 1)$ vector of controls, R is a positive semidefinite symmetric matrix, Q is a positive definite symmetric matrix, A is an $(n \times n)$ matrix, and B is an $(n \times k)$ matrix.

The optimal value function turns out to be quadratic, $V(x) = -x'Px$, where P is a positive semidefinite symmetric matrix.

Using the transition law to eliminate next period's state, the Bellman equation becomes

$$-x'Px = \max_u\{-x'Rx - u'Qu - (Ax + Bu)'P(Ax + Bu)\} \tag{78.1}$$

The first-order necessary conditions for the maximum problem on the right side of equation (78.1) are

Note

We use the following rules for differentiating quadratic and bilinear matrix forms: $\frac{\partial x'Ax}{\partial x} = (A + A')x$; $\frac{\partial y'Bz}{\partial y} = Bz$, $\frac{\partial y'Bz}{\partial z} = B'y$.

$$(Q + B'PB)u = -B'PAx,$$

which implies that an optimal decision rule for u is

$$u = -(Q + B'PB)^{-1}B'PAx$$

or

$$u = -Fx,$$

where

$$F = (Q + B'PB)^{-1}B'PA.$$

Substituting $u = -(Q + B'PB)^{-1}B'PAx$ into the right side of equation (78.1) and rearranging gives

$$P = R + A'PA - A'PB(Q + B'PB)^{-1}B'PA. \tag{78.2}$$

Equation (78.2) is called an **algebraic matrix Riccati** equation.

There are multiple solutions of equation (78.2).

But only one of them is positive definite.

The positive define solution is associated with the maximum of our problem.

It expresses the matrix P as an implicit function of the matrices R, Q, A, B.

Notice that the **gradient of the value function** is

$$\frac{\partial V(x)}{\partial x} = -2Px \tag{78.3}$$

We shall use fact (78.3) later.

78.3 Lagrangian

For the undiscounted optimal linear regulator problem, form the Lagrangian

$$L = -\sum_{t=0}^{\infty}\left\{x_t'Rx_t + u_t'Qu_t + 2\mu_{t+1}'[Ax_t + Bu_t - x_{t+1}]\right\} \tag{78.4}$$

where $2\mu_{t+1}$ is a vector of Lagrange multipliers on the time t transition law $x_{t+1} = Ax_t + Bu_t$.

(We put the 2 in front of μ_{t+1} to make things match up nicely with equation (78.3).)

First-order conditions for maximization with respect to $\{u_t, x_{t+1}\}_{t=0}^{\infty}$ are

$$\begin{aligned} 2Qu_t + 2B'\mu_{t+1} &= 0 \ , \ t \geq 0 \\ \mu_t &= Rx_t + A'\mu_{t+1} \ , \ t \geq 1. \end{aligned} \tag{78.5}$$

Define μ_0 to be a vector of shadow prices of x_0 and apply an envelope condition to (78.4) to deduce that

$$\mu_0 = Rx_0 + A'\mu_1,$$

which is a time $t = 0$ counterpart to the second equation of system (78.5).

An important fact is that

$$\mu_{t+1} = Px_{t+1} \tag{78.6}$$

where P is a positive define matrix that solves the algebraic Riccati equation (78.2).

Thus, from equations (78.3) and (78.6), $-2\mu_t$ is the gradient of the value function with respect to x_t.

The Lagrange multiplier vector μ_t is often called the **costate** vector that corresponds to the **state** vector x_t.

It is useful to proceed with the following steps:

- solve the first equation of (78.5) for u_t in terms of μ_{t+1}.
- substitute the result into the law of motion $x_{t+1} = Ax_t + Bu_t$.
- arrange the resulting equation and the second equation of (78.5) into the form

$$L \begin{bmatrix} x_{t+1} \\ \mu_{t+1} \end{bmatrix} = N \begin{bmatrix} x_t \\ \mu_t \end{bmatrix} \ , \ t \geq 0, \tag{78.7}$$

where

$$L = \begin{bmatrix} I & BQ^{-1}B' \\ 0 & A' \end{bmatrix}, \quad N = \begin{bmatrix} A & 0 \\ -R & I \end{bmatrix}.$$

When L is of full rank (i.e., when A is of full rank), we can write system (78.7) as

$$\begin{bmatrix} x_{t+1} \\ \mu_{t+1} \end{bmatrix} = M \begin{bmatrix} x_t \\ \mu_t \end{bmatrix} \tag{78.8}$$

where

$$M \equiv L^{-1}N = \begin{bmatrix} A + BQ^{-1}B'A'^{-1}R & -BQ^{-1}B'A'^{-1} \\ -A'^{-1}R & A'^{-1} \end{bmatrix}. \tag{78.9}$$

78.4 State-Costate Dynamics

We seek to solve the difference equation system (78.8) for a sequence $\{x_t\}_{t=0}^{\infty}$ that satisfies

- an initial condition for x_0
- a terminal condition $\lim_{t\to+\infty} x_t = 0$

This terminal condition reflects our desire for a **stable** solution, one that does not diverge as $t \to \infty$.

We inherit our wish for stability of the $\{x_t\}$ sequence from a desire to maximize

$$-\sum_{t=0}^{\infty}[x_t'Rx_t + u_t'Qu_t],$$

which requires that $x_t'Rx_t$ converge to zero as $t \to +\infty$.

78.5 Reciprocal Pairs Property

To proceed, we study properties of the $(2n \times 2n)$ matrix M defined in (78.9).

It helps to introduce a $(2n \times 2n)$ matrix

$$J = \begin{bmatrix} 0 & -I_n \\ I_n & 0 \end{bmatrix}.$$

The rank of J is $2n$.

Definition: A matrix M is called **symplectic** if

$$MJM' = J. \tag{78.10}$$

Salient properties of symplectic matrices that are readily verified include:

- If M is symplectic, then M^2 is symplectic
- The determinant of a symplectic, then $\det(M) = 1$

It can be verified directly that M in equation (78.9) is symplectic.

It follows from equation (78.10) and from the fact $J^{-1} = J' = -J$ that for any symplectic matrix M,

$$M' = J^{-1}M^{-1}J. \tag{78.11}$$

Equation (78.11) states that M' is related to the inverse of M by a **similarity transformation**.

For square matrices, recall that

- similar matrices share eigenvalues
- eigenvalues of the inverse of a matrix are inverses of eigenvalues of the matrix
- a matrix and its transpose share eigenvalues

It then follows from equation (78.11) that the eigenvalues of M occur in reciprocal pairs: if λ is an eigenvalue of M, so is λ^{-1}.

Write equation (78.8) as

$$y_{t+1} = My_t \tag{78.12}$$

where $y_t = \begin{bmatrix} x_t \\ \mu_t \end{bmatrix}$.

Consider a **triangularization** of M

$$V^{-1}MV = \begin{bmatrix} W_{11} & W_{12} \\ 0 & W_{22} \end{bmatrix} \tag{78.13}$$

where

- each block on the right side is $(n \times n)$
- V is nonsingular
- all eigenvalues of W_{22} exceed 1 in modulus
- all eigenvalues of W_{11} are less than 1 in modulus

78.6 Schur decomposition

The **Schur decomposition** and the **eigenvalue decomposition** are two decompositions of the form (78.13).

Write equation (78.12) as

$$y_{t+1} = VWV^{-1}y_t. \tag{78.14}$$

A solution of equation (78.14) for arbitrary initial condition y_0 is evidently

$$y_t = V \begin{bmatrix} W_{11}^t & W_{12,t} \\ 0 & W_{22}^t \end{bmatrix} V^{-1}y_0 \tag{78.15}$$

where $W_{12,t} = W_{12}$ for $t = 1$ and for $t \geq 2$ obeys the recursion

$$W_{12,t} = W_{11}^{t-1}W_{12,t-1} + W_{12,t-1}W_{22}^{t-1}$$

and where W_{ii}^t is W_{ii} raised to the tth power.

Write equation (78.15) as

$$\begin{bmatrix} y_{1t}^* \\ y_{2t}^* \end{bmatrix} = \begin{bmatrix} W_{11}^t & W_{12,t} \\ 0 & W_{22}^t \end{bmatrix} \begin{bmatrix} y_{10}^* \\ y_{20}^* \end{bmatrix}$$

where $y_t^* = V^{-1}y_t$, and in particular where

$$y_{2t}^* = V^{21}x_t + V^{22}\mu_t, \tag{78.16}$$

and where V^{ij} denotes the (i, j) piece of the partitioned V^{-1} matrix.

Because W_{22} is an unstable matrix, y_t^* will diverge unless $y_{20}^* = 0$.

Let V^{ij} denote the (i, j) piece of the partitioned V^{-1} matrix.

To attain stability, we must impose $y_{20}^* = 0$, which from equation (78.16) implies

$$V^{21}x_0 + V^{22}\mu_0 = 0$$

or

$$\mu_0 = -(V^{22})^{-1}V^{21}x_0.$$

This equation replicates itself over time in the sense that it implies

$$\mu_t = -(V^{22})^{-1} V^{21} x_t.$$

But notice that because $(V^{21}\ V^{22})$ is the second row block of the inverse of V, it follows that

$$(V^{21}\ V^{22}) \begin{bmatrix} V_{11} \\ V_{21} \end{bmatrix} = 0$$

which implies

$$V^{21} V_{11} + V^{22} V_{21} = 0.$$

Therefore,

$$-(V^{22})^{-1} V^{21} = V_{21} V_{11}^{-1}.$$

So we can write

$$\mu_0 = V_{21} V_{11}^{-1} x_0$$

and

$$\mu_t = V_{21} V_{11}^{-1} x_t.$$

However, we know that $\mu_t = P x_t$, where P occurs in the matrix that solves the Riccati equation.

Thus, the preceding argument establishes that

$$P = V_{21} V_{11}^{-1}. \tag{78.17}$$

Remarkably, formula (78.17) provides us with a computationally efficient way of computing the positive definite matrix P that solves the algebraic Riccati equation (78.2) that emerges from dynamic programming.

This same method can be applied to compute the solution of any system of the form (78.8) if a solution exists, even if eigenvalues of M fail to occur in reciprocal pairs.

The method will typically work so long as the eigenvalues of M split half inside and half outside the unit circle.

Systems in which eigenvalues (properly adjusted for discounting) fail to occur in reciprocal pairs arise when the system being solved is an equilibrium of a model in which there are distortions that prevent there being any optimum problem that the equilibrium solves. See [Ljungqvist and Sargent, 2018], ch 12.

78.7 Application

Here we demonstrate the computation with the deterministic permanent-income example from this *LQ Control: Foundations*.

Because that model is discounted, we apply the invariant-subspace method to the equivalent *undiscounted* system obtained from the transformed matrices $\hat{A} = \beta^{1/2} A$ and $\hat{B} = \beta^{1/2} B$.

```
# Model parameters
r = 0.05
β = 1 / (1 + r)
c_bar = 2
μ = 1
```

(continues on next page)

(continued from previous page)

```
# Formulate as an LQ problem
Q = np.array([[1]])
R = np.zeros((2, 2))
A = [[1 + r, -c_bar + μ],
     [0,               1]]
B = [[-1],
     [0]]

# Construct an LQ instance
lq = LQ(Q, R, A, B, beta=β)
```

Given matrices A, B, Q, R, we can then compute L, N, and $M = L^{-1}N$.

```
def construct_LNM(A, B, Q, R):

    n = A.shape[0]

    # construct L and N
    L = np.zeros((2*n, 2*n))
    L[:n, :n] = np.eye(n)
    L[:n, n:] = B @ np.linalg.inv(Q) @ B.T
    L[n:, n:] = A.T

    N = np.zeros((2*n, 2*n))
    N[:n, :n] = A
    N[n:, :n] = -R
    N[n:, n:] = np.eye(n)

    # compute M
    M = np.linalg.inv(L) @ N

    return L, N, M
```

```
A_bar = lq.A * lq.beta ** (1/2)
B_bar = lq.B * lq.beta ** (1/2)

L, N, M = construct_LNM(A_bar, B_bar, lq.Q, lq.R)
```

```
M
```

```
array([[ 1.02469508, -0.97590007, -0.92942864,  0.        ],
       [ 0.        ,  0.97590007,  0.        ,  0.        ],
       [ 0.        ,  0.        ,  0.97590007,  0.        ],
       [ 0.        ,  0.        ,  0.97590007,  1.02469508]])
```

Let's verify that M is symplectic.

```
n = lq.n
J = np.zeros((2*n, 2*n))
J[n:, :n] = np.eye(n)
J[:n, n:] = -np.eye(n)

M @ J @ M.T - J
```

```
array([[-5.15703240e-17,  0.00000000e+00,  0.00000000e+00,
         9.31818882e-18],
       [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
         1.11022302e-16],
       [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
         0.00000000e+00],
       [-9.31818882e-18, -1.11022302e-16,  0.00000000e+00,
         0.00000000e+00]])
```

We can compute the eigenvalues of M using `np.linalg.eigvals`, arranged in ascending order.

```
eigvals = sorted(np.linalg.eigvals(M), key=lambda z: (abs(z), z.real, z.imag))
eigvals
```

```
[np.float64(0.9759000729485331),
 np.float64(0.9759000729485332),
 np.float64(1.0246950765959597),
 np.float64(1.02469507659596)]
```

When we apply Schur decomposition such that $M = VWV^{-1}$, we want

- the upper left block of W, W_{11}, to have all of its eigenvalues less than 1 in modulus, and
- the lower right block W_{22} to have eigenvalues that exceed 1 in modulus.

To get what we want, let's define a sorting function that tells `scipy.schur` to sort the corresponding eigenvalues with modulus smaller than 1 to the upper left.

```
tol = 1e-10

def sort_fun(x):
    "Sort the eigenvalues with modulus smaller than 1 to the top-left."
    return abs(x) < 1 - tol

W, V, stable_dim = schur(M, sort=sort_fun)
stable_dim
```

```
2
```

```
W
```

```
array([[ 9.75900073e-01,  1.06779938e-16,  9.77000404e-01,
        -9.27110864e-01],
       [ 0.00000000e+00,  9.75900073e-01, -2.31777716e-03,
         9.77000404e-01],
       [ 0.00000000e+00,  0.00000000e+00,  1.02469508e+00,
        -2.12887870e-16],
       [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
         1.02469508e+00]])
```

```
V
```

```
array([[ 0.99875234, -0.00236938,  0.04988138,  0.        ],
       [ 0.04993762,  0.0473876 , -0.99762751,  0.        ],
       [ 0.        , -0.04988138, -0.00236938,  0.99875234],
       [ 0.        ,  0.99762751,  0.0473876 ,  0.04993762]])
```

We can check the modulus of eigenvalues of W_{11} and W_{22}.

Since they are both triangular matrices, eigenvalues are the diagonal elements.

```
# W11
np.diag(W[:n, :n])
```

```
array([0.97590007, 0.97590007])
```

```
# W22
np.diag(W[n:, n:])
```

```
array([1.02469508, 1.02469508])
```

The following functions wrap M matrix construction, Schur decomposition, and stability-imposing computation of P.

```
def stable_solution(M, verbose=True):
    """
    Given a system of linear difference equations

        y' = |a b| y
        x' = |c d| x

    which is potentially unstable, find the solution
    by imposing stability.

    Parameter
    ---------
    M : np.ndarray(float)
        The matrix represents the linear difference equations system.
    """
    n = M.shape[0] // 2
    tol = 1e-10

    def sort_fun(x):
        "Sort the eigenvalues with modulus smaller than 1 to the top-left."
        return abs(x) < 1 - tol

    W, V, stable_dim = schur(M, sort=sort_fun)
    if stable_dim != n:
        raise ValueError(
    f"Expected {n} stable eigenvalues inside the unit circle, found {stable_dim}."
    )
    if verbose:
        print('eigenvalues:\n')
        print('    W11: {}'.format(np.diag(W[:n, :n])))
        print('    W22: {}'.format(np.diag(W[n:, n:])))

    # compute V21 V11^{-1} without forming the inverse explicitly
    P = np.linalg.solve(V[:n, :n].T, V[n:, :n].T).T

    return W, V, P

def stationary_P(lq, verbose=True):
    """
    Computes the matrix :math:`P` that represent the value function
```

(continues on next page)

(continued from previous page)

```
        V(x) = x' P x

    in the infinite horizon case. Computation is via imposing stability
    on the solution path and using Schur decomposition.

    Parameters
    ----------
    lq : qe.LQ
        QuantEcon class for analyzing linear quadratic optimal control
        problems of infinite horizon form.

    Returns
    -------
    P : array_like(float)
        P matrix in the value function representation.
    """

    Q = lq.Q
    R = lq.R
    A = lq.A * lq.beta ** (1/2)
    B = lq.B * lq.beta ** (1/2)

    n, k = lq.n, lq.k

    L, N, M = construct_LNM(A, B, Q, R)
    W, V, P = stable_solution(M, verbose=verbose)

    return P
```

```
# compute P
stationary_P(lq)
```

```
eigenvalues:

    W11: [0.97590007 0.97590007]
    W22: [1.02469508 1.02469508]
```

```
array([[ 0.0525, -1.05  ],
       [-1.05  , 21.    ]])
```

Note that the matrix P computed in this way is close to what we get from the routine in quantecon that solves an algebraic Riccati equation by iterating to convergence on a Riccati difference equation.

The small difference comes from computational errors and will decrease as we increase the maximum number of iterations or decrease the tolerance for convergence.

```
lq.stationary_values()
```

```
(array([[ 0.0525, -1.05  ],
        [-1.05  , 21.    ]]),
 array([[-0.05,  1.  ]]),
 np.float64(0.0))
```

Using a Schur decomposition is much more efficient.

```
%%timeit
stationary_P(lq, verbose=False)
```

```
97.3 µs ± 65.8 ns per loop (mean ± std. dev. of 7 runs, 10,000 loops each)
```

```
%%timeit
lq.stationary_values()
```

```
1.81 ms ± 1.21 µs per loop (mean ± std. dev. of 7 runs, 1,000 loops each)
```

78.8 Other Applications

The preceding approach to imposing stability on a system of potentially unstable linear difference equations is not limited to linear quadratic dynamic optimization problems.

For example, the same method is used in our Stability in Linear Rational Expectations Models lecture.

Let's try to solve the model described in that lecture by applying the `stable_solution` function defined in this lecture above.

```
def construct_H(ρ, λ, δ):
    "contruct matrix H given parameters."

    H = np.empty((2, 2))
    H[0, :] = ρ,δ
    H[1, :] = - (1 - λ) / λ, 1 / λ

    return H

H = construct_H(ρ=.9, λ=.5, δ=0)
```

```
W, V, P = stable_solution(H)
P
```

```
eigenvalues:

    W11: [0.9]
    W22: [2.]
```

```
array([[0.90909091]])
```

78.9 Discounted Problems

78.9.1 Transforming States and Controls to Eliminate Discounting

A pair of useful transformations allows us to convert a discounted problem into an undiscounted one.

Thus, suppose that we have a discounted problem with objective

$$-\sum_{t=0}^{\infty} \beta^t \left\{ x_t' R x_t + u_t' Q u_t \right\}$$

and that the state transition equation is again $x_{t+1} = Ax_t + Bu_t$.

Define the transformed state and control variables

- $\hat{x}_t = \beta^{\frac{t}{2}} x_t$
- $\hat{u}_t = \beta^{\frac{t}{2}} u_t$

and the transformed transition equation matrices

- $\hat{A} = \beta^{\frac{1}{2}} A$
- $\hat{B} = \beta^{\frac{1}{2}} B$

so that the adjusted state and control variables obey the transition law

$$\hat{x}_{t+1} = \hat{A}\hat{x}_t + \hat{B}\hat{u}_t.$$

Then a discounted optimal control problem defined by A, B, R, Q, β having optimal policy characterized by P, F is associated with an equivalent undiscounted problem defined by $\hat{A}, \hat{B}, Q, R$ having optimal policy characterized by $\hat{F}, \hat{P}$ that satisfy the following equations:

$$\hat{F} = (Q + B'\hat{P}B)^{-1}\hat{B}'P\hat{A}$$

and

$$\hat{P} = R + \hat{A}'P\hat{A} - \hat{A}'P\hat{B}(Q + B'\hat{P}\hat{B})^{-1}\hat{B}'P\hat{A}$$

It follows immediately from the definitions of $\hat{A}, \hat{B}$ that $\hat{F} = F$ and $\hat{P} = P$.

By exploiting these transformations, we can solve a discounted problem by solving an associated undiscounted problem.

In particular, we can first transform a discounted LQ problem to an undiscounted one and then solve that discounted optimal regulator problem using the Lagrangian and invariant subspace methods described above.

For example, when $\beta = \frac{1}{1+r}$, we can solve for P with $\hat{A} = \beta^{1/2}A$ and $\hat{B} = \beta^{1/2}B$.

These settings are adopted by default in the function `stationary_P` defined above.

```
stationary_P(lq)
```

```
eigenvalues:

    W11: [0.97590007 0.97590007]
    W22: [1.02469508 1.02469508]
```

```
array([[ 0.0525, -1.05  ],
       [-1.05  , 21.    ]])
```

We can verify that the solution agrees with one that comes from applying the routine `LQ.stationary_values` in the quantecon package.

```
lq.stationary_values()
```

```
(array([[ 0.0525, -1.05  ],
        [-1.05  , 21.    ]]),
 array([[-0.05,  1.  ]]),
 np.float64(0.0))
```

78.9.2 Lagrangian for Discounted Problem

For several purposes, it is useful explicitly briefly to describe a Lagrangian for a discounted problem.

Thus, for the discounted optimal linear regulator problem, form the Lagrangian

$$L = -\sum_{t=0}^{\infty} \beta^t \left\{ x_t' R x_t + u_t' Q u_t + 2\beta \mu_{t+1}' [A x_t + B u_t - x_{t+1}] \right\} \tag{78.18}$$

where $2\mu_{t+1}$ is a vector of Lagrange multipliers on the state vector x_{t+1}.

First-order conditions for maximization with respect to $\{u_t, x_{t+1}\}_{t=0}^{\infty}$ are

$$\begin{aligned} 2Qu_t + 2\beta B' \mu_{t+1} = 0 \ , \ t \geq 0 \\ \mu_t = Rx_t + \beta A' \mu_{t+1} \ , \ t \geq 1. \end{aligned} \tag{78.19}$$

Define $2\mu_0$ to be the vector of shadow prices of x_0 and apply an envelope condition to (78.18) to deduce that

$$\mu_0 = Rx_0 + \beta A' \mu_1,$$

which is a time $t = 0$ counterpart to the second equation of system (78.19).

Proceeding as we did above with the undiscounted system (78.5), we can rearrange the first-order conditions into the system

$$\begin{bmatrix} I & \beta B Q^{-1} B' \\ 0 & \beta A' \end{bmatrix} \begin{bmatrix} x_{t+1} \\ \mu_{t+1} \end{bmatrix} = \begin{bmatrix} A & 0 \\ -R & I \end{bmatrix} \begin{bmatrix} x_t \\ \mu_t \end{bmatrix} \tag{78.20}$$

which in the special case that $\beta = 1$ agrees with equation (78.5), as expected.

By staring at system (78.20), we can infer identities that shed light on the structure of optimal linear regulator problems, some of which will be useful in this lecture when we apply and extend the methods of this lecture to study Stackelberg and Ramsey problems.

First, note that the first block of equation system (78.20) asserts that when $\mu_{t+1} = P x_{t+1}$, then

$$(I + \beta Q^{-1} B' P B P) x_{t+1} = A x_t,$$

which can be rearranged to sbe

$$x_{t+1} = (I + \beta B Q^{-1} B' P)^{-1} A x_t.$$

This expression for the optimal closed loop dynamics of the state must agree with an alternative expression that we had derived with dynamic programming, namely,

$$x_{t+1} = (A - BF) x_t.$$

But using

$$F = \beta (Q + \beta B' P B)^{-1} B' P A \tag{78.21}$$

it follows that

$$A - BF = (I - \beta B (Q + \beta B' P B)^{-1} B' P) A.$$

Thus, our two expressions for the closed loop dynamics agree if and only if

$$(I + \beta B Q^{-1} B' P)^{-1} = (I - \beta B (Q + \beta B' P B)^{-1} B' P). \tag{78.22}$$

Matrix equation (78.22) can be verified by applying a partitioned inverse formula.

Note

Just use the formula $(a-bd^{-1}c)^{-1} = a^{-1}+a^{-1}b(d-ca^{-1}b)^{-1}ca^{-1}$ for appropriate choices of the matrices a, b, c, d.

Next, note that for *any* fixed F for which eigenvalues of $A-BF$ are less than $\frac{1}{\beta}$ in modulus, the value function associated with using this rule forever is $-x_0\tilde{P}x_0$ where $\tilde{P}$ obeys the following matrix equation:

$$\tilde{P} = (R + F'QF) + \beta(A - BF)'P(A - BF). \tag{78.23}$$

Evidently, $\tilde{P} = P$ only when F obeys formula (78.21).

Next, note that the second equation of system (78.20) implies the "forward looking" equation for the Lagrange multiplier

$$\mu_t = Rx_t + \beta A'\mu_{t+1}$$

whose solution is

$$\mu_t = Px_t,$$

where

$$P = R + \beta A'P(A - BF) \tag{78.24}$$

where we must require that F obeys equation (78.21).

Equations (78.23) and (78.24) provide different perspectives on the optimal value function.

CHAPTER

SEVENTYNINE

ELIMINATING CROSS PRODUCTS

79.1 Overview

This lecture describes formulas for eliminating

- cross products between states and control in linear-quadratic dynamic programming problems
- covariances between state and measurement noises in Kalman filtering problems

For a linear-quadratic dynamic programming problem, the idea involves these steps

- transform states and controls in a way that leads to an equivalent problem with no cross-products between transformed states and controls
- solve the transformed problem using standard formulas for problems with no cross-products between states and controls presented in this lecture *Linear Control: Foundations*
- transform the optimal decision rule for the altered problem into the optimal decision rule for the original problem with cross-products between states and controls

79.2 Undiscounted Dynamic Programming Problem

Here is a nonstochastic undiscounted LQ dynamic programming with cross products between states and controls in the objective function.

The problem is defined by the 5-tuple of matrices (A, B, R, Q, H) where R and Q are positive definite symmetric matrices and $A \sim m \times m, B \sim m \times k, Q \sim k \times k, R \sim m \times m$ and $H \sim k \times m$.

The problem is to choose $\{x_{t+1}, u_t\}_{t=0}^{\infty}$ to maximize

$$-\sum_{t=0}^{\infty}(x_t' R x_t + u_t' Q u_t + 2 u_t H x_t)$$

subject to the linear constraints

$$x_{t+1} = A x_t + B u_t, \quad t \geq 0$$

where x_0 is a given initial condition.

The solution to this undiscounted infinite-horizon problem is a time-invariant feedback rule

$$u_t = -F x_t$$

where

$$F = -(Q + B'PB)^{-1}B'PA$$

and $P \sim m \times m$ is a positive definite solution of the algebraic matrix Riccati equation

$$P = R + A'PA - (A'PB + H')(Q + B'PB)^{-1}(B'PA + H).$$

It can be verified that an **equivalent** problem without cross-products between states and controls is defined by a 4-tuple of matrices : (A^*, B, R^*, Q).

That the omitted matrix $H = 0$ indicates that there are no cross products between states and controls in the equivalent problem.

The matrices (A^*, B, R^*, Q) defining the equivalent problem and the value function, policy function matrices P, F^* that solve it are related to the matrices (A, B, R, Q, H) defining the original problem and the value function, policy function matrices P, F that solve the original problem by

$$\begin{aligned}
A^* &= A - BQ^{-1}H, \\
R^* &= R - H'Q^{-1}H, \\
P &= R^* + A^{*'}PA - (A^{*'}PB)(Q + B'PB)^{-1}B'PA^*, \\
F^* &= (Q + B'PB)^{-1}B'PA^*, \\
F &= F^* + Q^{-1}H.
\end{aligned}$$

79.3 Kalman Filter

The **duality** that prevails between a linear-quadratic optimal control and a Kalman filtering problem means that there is an analogous transformation that allows us to transform a Kalman filtering problem with non-zero covariance matrix between between shocks to states and shocks to measurements to an equivalent Kalman filtering problem with zero covariance between shocks to states and measurments.

Let's look at the appropriate transformations.

First, let's recall the Kalman filter with covariance between noises to states and measurements.

The hidden Markov model is

$$\begin{aligned}
x_{t+1} &= Ax_t + Bw_{t+1}, \\
z_{t+1} &= Dx_t + Fw_{t+1},
\end{aligned}$$

where $A \sim m \times m, B \sim m \times p$ and $D \sim k \times m, F \sim k \times p$, and w_{t+1} is the time $t+1$ component of a sequence of i.i.d. $p \times 1$ normally distibuted random vectors with mean vector zero and covariance matrix equal to a $p \times p$ identity matrix.

Thus, x_t is $m \times 1$ and z_t is $k \times 1$.

The Kalman filtering formulas are

$$\begin{aligned}
K(\Sigma_t) &= (A\Sigma_t D' + BF')(D\Sigma_t D' + FF')^{-1}, \\
\Sigma_{t+1} &= A\Sigma_t A' + BB' - (A\Sigma_t D' + BF')(D\Sigma_t D' + FF')^{-1}(D\Sigma_t A' + FB').
\end{aligned}$$

(eq:Kalman102)

Define tranformed matrices

$$\begin{aligned}
A^* &= A - BF'(FF')^{-1}D, \\
B^*B^{*'} &= BB' - BF'(FF')^{-1}FB'.
\end{aligned}$$

79.3.1 Algorithm

A consequence of formulas {eq}`eq:Kalman102} is that we can use the following algorithm to solve Kalman filtering problems that involve non zero covariances between state and signal noises.

First, compute Σ, K^* using the ordinary Kalman filtering formula with $BF' = 0$, i.e., with zero covariance matrix between random shocks to states and random shocks to measurements.

That is, compute K^* and Σ that satisfy

$$
\begin{aligned}
K^* &= (A^*\Sigma D')(D\Sigma D' + FF')^{-1} \\
\Sigma &= A^*\Sigma A^{*'} + B^*B^{*'} - (A^*\Sigma D')(D\Sigma D' + FF')^{-1}(D\Sigma A^{*'}).
\end{aligned}
$$

The Kalman gain for the original problem **with non-zero covariance** between shocks to states and measurements is then

$$
K = K^* + BF'(FF')^{-1},
$$

The state reconstruction covariance matrix Σ for the original problem equals the state reconstrution covariance matrix for the transformed problem.

79.4 Duality table

Here is a handy table to remember how the Kalman filter and dynamic program are related.

Dynamic Program	Kalman Filter
A	A'
B	D'
H	FB'
Q	FF'
R	BB'
F	K'
P	Σ

CHAPTER

EIGHTY

THE PERMANENT INCOME MODEL

Contents

- *The Permanent Income Model*
  - *Overview*
  - *The Savings Problem*
  - *Alternative Representations*
  - *Two Classic Examples*
  - *Further Reading*
  - *Appendix: The Euler Equation*

80.1 Overview

This lecture describes a rational expectations version of the famous permanent income model of Milton Friedman [Friedman, 1956].

Robert Hall cast Friedman's model within a linear-quadratic setting [Hall, 1978].

Like Hall, we formulate an infinite-horizon linear-quadratic savings problem.

We use the model as a vehicle for illustrating

- alternative formulations of the *state* of a dynamic system
- the idea of *cointegration*
- impulse response functions
- the idea that changes in consumption are useful as predictors of movements in income

Background readings on the linear-quadratic-Gaussian permanent income model are Hall's [Hall, 1978] and chapter 2 of [Ljungqvist and Sargent, 2018].

Let's start with some imports

```
import matplotlib.pyplot as plt
import numpy as np
import random
from numba import jit
```

80.2 The Savings Problem

In this section, we state and solve the savings and consumption problem faced by the consumer.

80.2.1 Preliminaries

We use a class of stochastic processes called martingales.

A discrete-time martingale is a stochastic process (i.e., a sequence of random variables) $\{X_t\}$ with finite mean at each t and satisfying

$$\mathbb{E}_t[X_{t+1}] = X_t, \qquad t = 0, 1, 2, ...$$

Here $\mathbb{E}_t := \mathbb{E}[\cdot \mid \mathcal{F}_t]$ is a conditional mathematical expectation conditional on the time t *information set* $\mathcal{F}_t$.

The latter is just a collection of random variables that the modeler declares to be visible at t.

- When not explicitly defined, it is usually understood that $\mathcal{F}_t = \{X_t, X_{t-1}, ..., X_0\}$.

Martingales have the feature that the history of past outcomes provides no predictive power for changes between current and future outcomes.

For example, the current wealth of a gambler engaged in a "fair game" has this property.

One common class of martingales is the family of *random walks*.

A **random walk** is a stochastic process $\{X_t\}$ that satisfies

$$X_{t+1} = X_t + w_{t+1}$$

for some IID zero mean *innovation* sequence $\{w_t\}$.

Evidently, X_t can also be expressed as

$$X_t = \sum_{j=1}^{t} w_j + X_0$$

Not every martingale arises as a random walk (see, for example, Wald's martingale).

80.2.2 The Decision Problem

A consumer has preferences over consumption streams that are ordered by the utility functional

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t u(c_t) \right] \tag{80.1}$$

where

- $\mathbb{E}_t$ is the mathematical expectation conditioned on the consumer's time t information
- c_t is time t consumption
- u is a strictly concave one-period utility function
- $\beta \in (0, 1)$ is a discount factor

The consumer maximizes (80.1) by choosing a consumption, borrowing plan $\{c_t, b_{t+1}\}_{t=0}^{\infty}$ subject to the sequence of budget constraints

$$c_t + b_t = \frac{1}{1+r} b_{t+1} + y_t \quad t \geq 0 \tag{80.2}$$

Here

- y_t is an exogenous endowment process.
- $r > 0$ is a time-invariant risk-free net interest rate.
- b_t is one-period risk-free debt maturing at t.

The consumer also faces initial conditions b_0 and y_0, which can be fixed or random.

80.2.3 Assumptions

For the remainder of this lecture, we follow Friedman and Hall in assuming that $(1+r)^{-1} = \beta$.

Regarding the endowment process, we assume it has the *state-space representation*

$$\begin{aligned} z_{t+1} &= A z_t + C w_{t+1} \\ y_t &= U z_t \end{aligned} \tag{80.3}$$

where

- $\{w_t\}$ is an IID vector process with $\mathbb{E} w_t = 0$ and $\mathbb{E} w_t w_t' = I$.
- The *spectral radius* of A satisfies $\rho(A) < \sqrt{1/\beta}$.
- U is a selection vector that pins down y_t as a particular linear combination of components of z_t.

The restriction on $\rho(A)$ prevents income from growing so fast that discounted geometric sums of some quadratic forms to be described below become infinite.

Regarding preferences, we assume the quadratic utility function

$$u(c_t) = -(c_t - \gamma)^2$$

where γ is a bliss level of consumption.

Note

Along with this quadratic utility specification, we allow consumption to be negative. However, by choosing parameters appropriately, we can make the probability that the model generates negative consumption paths over finite time horizons as low as desired.

Finally, we impose the *no Ponzi scheme* condition

$$\mathbb{E}_0 \left[\sum_{t=0}^{\infty} \beta^t b_t^2 \right] < \infty \tag{80.4}$$

This condition rules out an always-borrow scheme that would allow the consumer to enjoy bliss consumption forever.

80.2.4 First-Order Conditions

First-order conditions for maximizing (80.1) subject to (80.2) are

$$\mathbb{E}_t[u'(c_{t+1})] = u'(c_t), \qquad t = 0, 1, \ldots \tag{80.5}$$

These optimality conditions are also known as *Euler equations*.

If you're not sure where they come from, you can find a proof sketch in the *appendix*.

With our quadratic preference specification, (80.5) has the striking implication that consumption follows a martingale:

$$\mathbb{E}_t[c_{t+1}] = c_t \tag{80.6}$$

(In fact, quadratic preferences are *necessary* for this conclusion[1].)

One way to interpret (80.6) is that consumption will change only when "new information" about permanent income is revealed.

These ideas will be clarified below.

80.2.5 The Optimal Decision Rule

Now let's deduce the optimal decision rule[2].

Note

One way to solve the consumer's problem is to apply *dynamic programming* as in *this lecture*. We do this later. But first we use an alternative approach that is revealing and shows the work that dynamic programming does for us behind the scenes.

In doing so, we need to combine

1. the optimality condition (80.6)
2. the period-by-period budget constraint (80.2), and
3. the boundary condition (80.4)

To accomplish this, observe first that (80.4) implies $\lim_{t\to\infty} \beta^{\frac{t}{2}} b_{t+1} = 0$.

Using this restriction on the debt path and solving (80.2) forward yields

$$b_t = \sum_{j=0}^{\infty} \beta^j (y_{t+j} - c_{t+j}) \tag{80.7}$$

Take conditional expectations on both sides of (80.7) and use the martingale property of consumption and the *law of iterated expectations* to deduce

$$b_t = \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t[y_{t+j}] - \frac{c_t}{1-\beta} \tag{80.8}$$

[1] A linear marginal utility is essential for deriving (80.6) from (80.5). Suppose instead that we had imposed the following more standard assumptions on the utility function: $u'(c) > 0, u''(c) < 0, u'''(c) > 0$ and required that $c \geq 0$. The Euler equation remains (80.5). But the fact that $u''' < 0$ implies via Jensen's inequality that $\mathbb{E}_t[u'(c_{t+1})] > u'(\mathbb{E}_t[c_{t+1}])$. This inequality together with (80.5) implies that $\mathbb{E}_t[c_{t+1}] > c_t$ (consumption is said to be a 'submartingale'), so that consumption stochastically diverges to $+\infty$. The consumer's savings also diverge to $+\infty$.

[2] An optimal decision rule is a map from the current state into current actions—in this case, consumption.

Expressed in terms of c_t we get

$$c_t = (1-\beta)\left[\sum_{j=0}^{\infty}\beta^j \mathbb{E}_t[y_{t+j}] - b_t\right] = \frac{r}{1+r}\left[\sum_{j=0}^{\infty}\beta^j \mathbb{E}_t[y_{t+j}] - b_t\right] \tag{80.9}$$

where the last equality uses $(1+r)\beta = 1$.

These last two equations assert that consumption equals *economic income*

- **financial wealth** equals $-b_t$
- **non-financial wealth** equals $\sum_{j=0}^{\infty}\beta^j \mathbb{E}_t[y_{t+j}]$
- **total wealth** equals the sum of financial and non-financial wealth
- a **marginal propensity to consume out of total wealth** equals the interest factor $\frac{r}{1+r}$
- **economic income** equals
  - a constant marginal propensity to consume times the sum of non-financial wealth and financial wealth
  - the amount the consumer can consume while leaving its wealth intact

Responding to the State

The *state* vector confronting the consumer at t is $\begin{bmatrix} b_t & z_t \end{bmatrix}$.

Here

- z_t is an *exogenous* component, unaffected by consumer behavior.
- b_t is an *endogenous* component (since it depends on the decision rule).

Note that z_t contains all variables useful for forecasting the consumer's future endowment.

It is plausible that current decisions c_t and b_{t+1} should be expressible as functions of z_t and b_t.

This is indeed the case.

In fact, from *this discussion*, we see that

$$\sum_{j=0}^{\infty}\beta^j \mathbb{E}_t[y_{t+j}] = \mathbb{E}_t\left[\sum_{j=0}^{\infty}\beta^j y_{t+j}\right] = U(I-\beta A)^{-1} z_t$$

Combining this with (80.9) gives

$$c_t = \frac{r}{1+r}\left[U(I-\beta A)^{-1} z_t - b_t\right] \tag{80.10}$$

Using this equality to eliminate c_t in the budget constraint (80.2) gives

$$\begin{aligned} b_{t+1} &= (1+r)(b_t + c_t - y_t) \\ &= (1+r)b_t + r[U(I-\beta A)^{-1} z_t - b_t] - (1+r)U z_t \\ &= b_t + U[r(I-\beta A)^{-1} - (1+r)I]z_t \\ &= b_t + U(I-\beta A)^{-1}(A-I)z_t \end{aligned}$$

To get from the second last to the last expression in this chain of equalities is not trivial.

A key is to use the fact that $(1+r)\beta = 1$ and $(I-\beta A)^{-1} = \sum_{j=0}^{\infty}\beta^j A^j$.

We've now successfully written c_t and b_{t+1} as functions of b_t and z_t.

A State-Space Representation

We can summarize our dynamics in the form of a linear state-space system governing consumption, debt and income:

$$\begin{aligned} z_{t+1} &= Az_t + Cw_{t+1} \\ b_{t+1} &= b_t + U[(I-\beta A)^{-1}(A-I)]z_t \\ y_t &= Uz_t \\ c_t &= (1-\beta)[U(I-\beta A)^{-1}z_t - b_t] \end{aligned} \tag{80.11}$$

To write this more succinctly, let

$$x_t = \begin{bmatrix} z_t \\ b_t \end{bmatrix}, \quad \tilde{A} = \begin{bmatrix} A & 0 \\ U(I-\beta A)^{-1}(A-I) & 1 \end{bmatrix}, \quad \tilde{C} = \begin{bmatrix} C \\ 0 \end{bmatrix}$$

and

$$\tilde{U} = \begin{bmatrix} U & 0 \\ (1-\beta)U(I-\beta A)^{-1} & -(1-\beta) \end{bmatrix}, \quad \tilde{y}_t = \begin{bmatrix} y_t \\ c_t \end{bmatrix}$$

Then we can express equation (80.11) as

$$\begin{aligned} x_{t+1} &= \tilde{A}x_t + \tilde{C}w_{t+1} \\ \tilde{y}_t &= \tilde{U}x_t \end{aligned} \tag{80.12}$$

We can use the following formulas from *linear state space models* to compute population mean $\mu_t = \mathbb{E}x_t$ and covariance $\Sigma_t := \mathbb{E}[(x_t - \mu_t)(x_t - \mu_t)']$

$$\mu_{t+1} = \tilde{A}\mu_t \quad \text{with} \quad \mu_0 \text{ given} \tag{80.13}$$

$$\Sigma_{t+1} = \tilde{A}\Sigma_t\tilde{A}' + \tilde{C}\tilde{C}' \quad \text{with} \quad \Sigma_0 \text{ given} \tag{80.14}$$

We can then compute the mean and covariance of $\tilde{y}_t$ from

$$\begin{aligned} \mu_{y,t} &= \tilde{U}\mu_t \\ \Sigma_{y,t} &= \tilde{U}\Sigma_t\tilde{U}' \end{aligned} \tag{80.15}$$

A Simple Example with IID Income

To gain some preliminary intuition on the implications of (80.11), let's look at a highly stylized example where income is just IID.

(Later examples will investigate more realistic income streams.)

In particular, let $\{w_t\}_{t=1}^{\infty}$ be IID and scalar standard normal, and let

$$z_t = \begin{bmatrix} z_t^1 \\ 1 \end{bmatrix}, \quad A = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix}, \quad U = \begin{bmatrix} 1 & \mu \end{bmatrix}, \quad C = \begin{bmatrix} \sigma \\ 0 \end{bmatrix}$$

Finally, let $b_0 = z_0^1 = 0$.

Under these assumptions, we have $y_t = \mu + \sigma w_t \sim N(\mu, \sigma^2)$.

Further, if you work through the state space representation, you will see that

$$b_t = -\sigma\sum_{j=1}^{t-1} w_j$$

$$c_t = \mu + (1-\beta)\sigma\sum_{j=1}^{t} w_j$$

Thus, income is IID and debt and consumption are both Gaussian random walks.

Defining assets as $-b_t$, we see that assets are just the cumulative sum of unanticipated incomes prior to the present date.

The next figure shows a typical realization with $r = 0.05$, $\mu = 1$, and $\sigma = 0.15$

```
r = 0.05
β = 1 / (1 + r)
σ = 0.15
μ = 1
T = 60

@jit
def time_path(T, rng):
    w = rng.standard_normal(T+1)  # w_0, w_1, ..., w_T
    w[0] = 0
    b = np.zeros(T+1)
    for t in range(1, T+1):
        b[t] = w[1:t].sum()
    b = -σ * b
    c = μ + (1 - β) * (σ * w - b)
    return w, b, c

rng = np.random.default_rng()
w, b, c = time_path(T, rng)

fig, ax = plt.subplots(figsize=(10, 6))

ax.plot(μ + σ * w, 'g-', label="Non-financial income")
ax.plot(c, 'k-', label="Consumption")
ax.plot( b, 'b-', label="Debt")
ax.legend(ncol=3, mode='expand', bbox_to_anchor=(0., 1.02, 1., .102))
ax.grid()
ax.set_xlabel('Time')

plt.show()
```



Observe that consumption is considerably smoother than income.

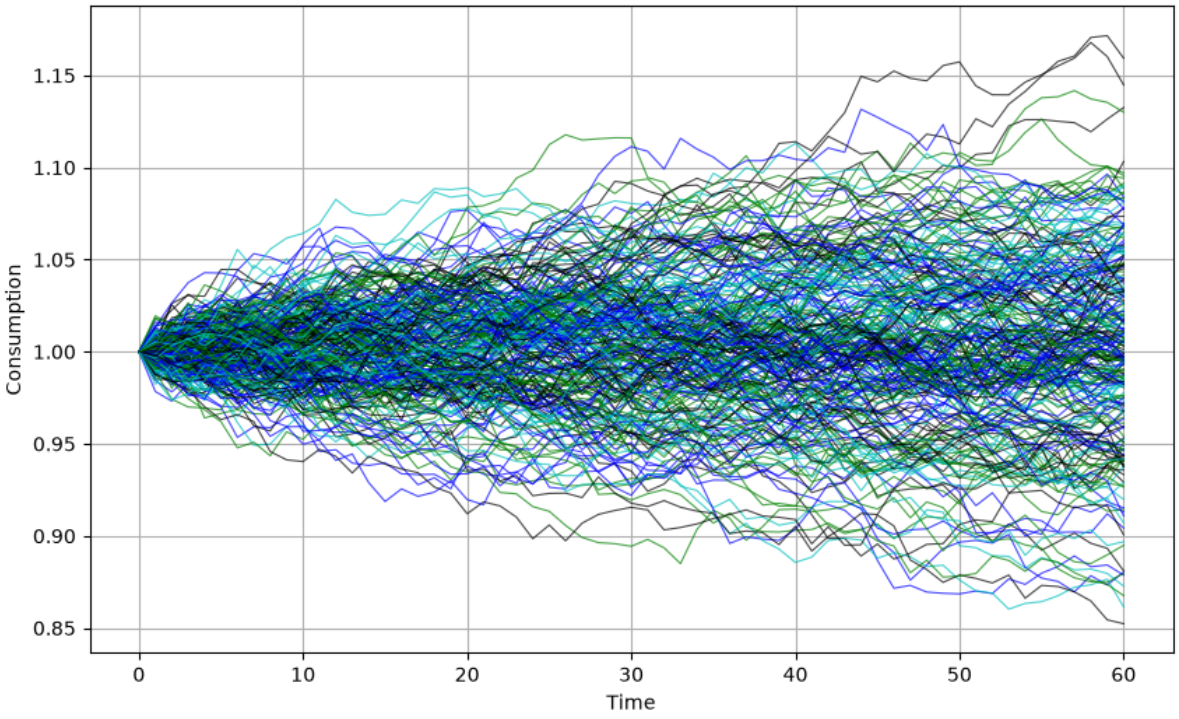
The figure below shows the consumption paths of 250 consumers with independent income streams

```
fig, ax = plt.subplots(figsize=(10, 6))

b_sum = np.zeros(T+1)
for i in range(250):
    w, b, c = time_path(T, rng)   # Generate new time path
    rcolor = random.choice(('c', 'g', 'b', 'k'))
    ax.plot(c, color=rcolor, lw=0.8, alpha=0.7)

ax.grid()
ax.set(xlabel='Time', ylabel='Consumption')

plt.show()
```



80.3 Alternative Representations

In this section, we shed more light on the evolution of savings, debt and consumption by representing their dynamics in several different ways.

80.3.1 Hall's Representation

Hall [Hall, 1978] suggested an insightful way to summarize the implications of LQ permanent income theory.

First, to represent the solution for b_t, shift (80.9) forward one period and eliminate b_{t+1} by using (80.2) to obtain

$$c_{t+1} = (1-\beta)\sum_{j=0}^{\infty}\beta^j \mathbb{E}_{t+1}[y_{t+j+1}] - (1-\beta)\left[\beta^{-1}(c_t + b_t - y_t)\right]$$

If we add and subtract $\beta^{-1}(1-\beta)\sum_{j=0}^{\infty}\beta^j \mathbb{E}_t y_{t+j}$ from the right side of the preceding equation and rearrange, we obtain

$$c_{t+1} - c_t = (1-\beta)\sum_{j=0}^{\infty}\beta^j\left\{\mathbb{E}_{t+1}[y_{t+j+1}] - \mathbb{E}_t[y_{t+j+1}]\right\} \tag{80.16}$$

The right side is the time $t+1$ *innovation to the expected present value* of the endowment process $\{y_t\}$.

We can represent the optimal decision rule for (c_t, b_{t+1}) in the form of (80.16) and (80.8), which we repeat:

$$b_t = \sum_{j=0}^{\infty}\beta^j \mathbb{E}_t[y_{t+j}] - \frac{1}{1-\beta}c_t \tag{80.17}$$

Equation (80.17) asserts that the consumer's debt due at t equals the expected present value of its endowment minus the expected present value of its consumption stream.

A high debt thus indicates a large expected present value of surpluses $y_t - c_t$.

Recalling again our discussion on *forecasting geometric sums*, we have

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j y_{t+j} = U(I - \beta A)^{-1} z_t$$
$$\mathbb{E}_{t+1} \sum_{j=0}^{\infty} \beta^j y_{t+j+1} = U(I - \beta A)^{-1} z_{t+1}$$
$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j y_{t+j+1} = U(I - \beta A)^{-1} A z_t$$

Using these formulas together with (80.3) and substituting into (80.16) and (80.17) gives the following representation for the consumer's optimum decision rule:

$$\begin{aligned} c_{t+1} &= c_t + (1-\beta) U(I - \beta A)^{-1} C w_{t+1} \\ b_t &= U(I - \beta A)^{-1} z_t - \frac{1}{1-\beta} c_t \\ y_t &= U z_t \\ z_{t+1} &= A z_t + C w_{t+1} \end{aligned} \tag{80.18}$$

Representation (80.18) makes clear that

- The state can be taken as (c_t, z_t).
  - The endogenous part is c_t and the exogenous part is z_t.
  - Debt b_t has disappeared as a component of the state because it is encoded in c_t.
- Consumption is a random walk with innovation $(1-\beta) U(I - \beta A)^{-1} C w_{t+1}$.
  - This is a more explicit representation of the martingale result in (80.6).

80.3.2 Cointegration

Representation (80.18) reveals that the joint process $\{c_t, b_t\}$ possesses the property that Engle and Granger [Engle and Granger, 1987] called cointegration.

Cointegration is a tool that allows us to apply powerful results from the theory of stationary stochastic processes to (certain transformations of) nonstationary models.

To apply cointegration in the present context, suppose that z_t is asymptotically stationary[3].

Despite this, both c_t and b_t will be non-stationary because they have unit roots (see (80.11) for b_t).

Nevertheless, there is a linear combination of c_t, b_t that *is* asymptotically stationary.

In particular, from the second equality in (80.18) we have

$$(1-\beta) b_t + c_t = (1-\beta) U(I - \beta A)^{-1} z_t \tag{80.19}$$

Hence the linear combination $(1-\beta) b_t + c_t$ is asymptotically stationary.

Accordingly, Granger and Engle would call $\begin{bmatrix} (1-\beta) & 1 \end{bmatrix}$ a **cointegrating vector** for the state.

When applied to the nonstationary vector process $\begin{bmatrix} b_t & c_t \end{bmatrix}'$, it yields a process that is asymptotically stationary.

[3] This would be the case if, for example, the *spectral radius* of A is strictly less than one.

Equation (80.19) can be rearranged to take the form

$$(1-\beta)b_t + c_t = (1-\beta)\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j y_{t+j} \tag{80.20}$$

Equation (80.20) asserts that the *cointegrating residual* on the left side equals the conditional expectation of the geometric sum of future incomes on the right[4].

80.3.3 Cross-Sectional Implications

Consider again (80.18), this time in light of our discussion of distribution dynamics in the *lecture on linear systems*.

The dynamics of c_t are given by

$$c_{t+1} = c_t + (1-\beta)U(I-\beta A)^{-1}Cw_{t+1} \tag{80.21}$$

or

$$c_t = c_0 + \sum_{j=1}^{t} \hat{w}_j \quad \text{for} \quad \hat{w}_{t+1} := (1-\beta)U(I-\beta A)^{-1}Cw_{t+1}$$

The unit root affecting c_t causes the time t variance of c_t to grow linearly with t.

In particular, since $\{\hat{w}_t\}$ is IID, we have

$$\text{Var}[c_t] = \text{Var}[c_0] + t\,\hat{\sigma}^2 \tag{80.22}$$

where

$$\hat{\sigma}^2 := (1-\beta)^2 U(I-\beta A)^{-1}CC'(I-\beta A')^{-1}U'$$

When $\hat{\sigma} > 0$, $\{c_t\}$ has no asymptotic distribution.

Let's consider what this means for a cross-section of ex-ante identical consumers born at time 0.

Let the distribution of c_0 represent the cross-section of initial consumption values.

Equation (80.22) tells us that the variance of c_t increases over time at a rate proportional to t.

A number of different studies have investigated this prediction and found some support for it (see, e.g., [Deaton and Paxson, 1994], [Storesletten *et al.*, 2004]).

80.3.4 Impulse Response Functions

Impulse response functions measure responses to various impulses (i.e., temporary shocks).

The impulse response function of $\{c_t\}$ to the innovation $\{w_t\}$ is a box.

In particular, the response of c_{t+j} to a unit increase in the innovation w_{t+1} is $(1-\beta)U(I-\beta A)^{-1}C$ for all $j \geq 1$.

[4] See [Campbell and Shiller, 1988], [Lettau and Ludvigson, 2001], [Lettau and Ludvigson, 2004] for interesting applications of related ideas.

80.3.5 Moving Average Representation

It's useful to express the innovation to the expected present value of the endowment process in terms of a moving average representation for income y_t.

The endowment process defined by (80.3) has the moving average representation

$$y_{t+1} = d(L)w_{t+1} \tag{80.23}$$

where

- $d(L) = \sum_{j=0}^{\infty} d_j L^j$ for some sequence d_j, where L is the lag operator[5]
- at time t, the consumer has an information set[6] $w^t = [w_t, w_{t-1}, ...]$

Notice that

$$y_{t+j} - \mathbb{E}_t[y_{t+j}] = d_0 w_{t+j} + d_1 w_{t+j-1} + \cdots + d_{j-1} w_{t+1}$$

It follows that

$$\mathbb{E}_{t+1}[y_{t+j}] - \mathbb{E}_t[y_{t+j}] = d_{j-1} w_{t+1} \tag{80.24}$$

Using (80.24) in (80.16) gives

$$c_{t+1} - c_t = (1-\beta) d(\beta) w_{t+1} \tag{80.25}$$

The object $d(\beta)$ is the **present value of the moving average coefficients** in the representation for the endowment process y_t.

80.4 Two Classic Examples

We illustrate some of the preceding ideas with two examples.

In both examples, the endowment follows the process $y_t = z_{1t} + z_{2t}$ where

$$\begin{bmatrix} z_{1t+1} \\ z_{2t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} z_{1t} \\ z_{2t} \end{bmatrix} + \begin{bmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{bmatrix} \begin{bmatrix} w_{1t+1} \\ w_{2t+1} \end{bmatrix}$$

Here

- w_{t+1} is an IID 2×1 process distributed as $N(0, I)$.
- z_{1t} is a permanent component of y_t.
- z_{2t} is a purely transitory component of y_t.

80.4.1 Example 1

Assume as before that the consumer observes the state z_t at time t.

In view of (80.18) we have

$$c_{t+1} - c_t = \sigma_1 w_{1t+1} + (1-\beta)\sigma_2 w_{2t+1} \tag{80.26}$$

Formula (80.26) shows how an increment $\sigma_1 w_{1t+1}$ to the permanent component of income z_{1t+1} leads to

[5] Representation (80.3) implies that $d(L) = U(I - AL)^{-1}C$.

[6] A moving average representation for a process y_t is said to be **fundamental** if the linear space spanned by y^t is equal to the linear space spanned by w^t. A time-invariant innovations representation, attained via the Kalman filter, is by construction fundamental.

- a permanent one-for-one increase in consumption and
- no increase in savings $-b_{t+1}$

But the purely transitory component of income $\sigma_2 w_{2t+1}$ leads to a permanent increment in consumption by a fraction $1-\beta$ of transitory income.

The remaining fraction β is saved, leading to a permanent increment in $-b_{t+1}$.

Application of the formula for debt in (80.11) to this example shows that

$$b_{t+1} - b_t = -z_{2t} = -\sigma_2 w_{2t} \tag{80.27}$$

This confirms that none of $\sigma_1 w_{1t}$ is saved, while all of $\sigma_2 w_{2t}$ is saved.

The next figure displays impulse-response functions that illustrates these very different reactions to transitory and permanent income shocks.

```
r = 0.05
β = 1 / (1 + r)
S = 5   # Impulse date
σ1 = σ2 = 0.15

@jit
def time_path(T, permanent=False):
    "Time path of consumption and debt given shock sequence"
    w1 = np.zeros(T+1)
    w2 = np.zeros(T+1)
    b = np.zeros(T+1)
    c = np.zeros(T+1)
    if permanent:
        w1[S+1] = 1.0
    else:
        w2[S+1] = 1.0
    for t in range(1, T):
        b[t+1] = b[t] - σ2 * w2[t]
        c[t+1] = c[t] + σ1 * w1[t+1] + (1 - β) * σ2 * w2[t+1]
    return b, c

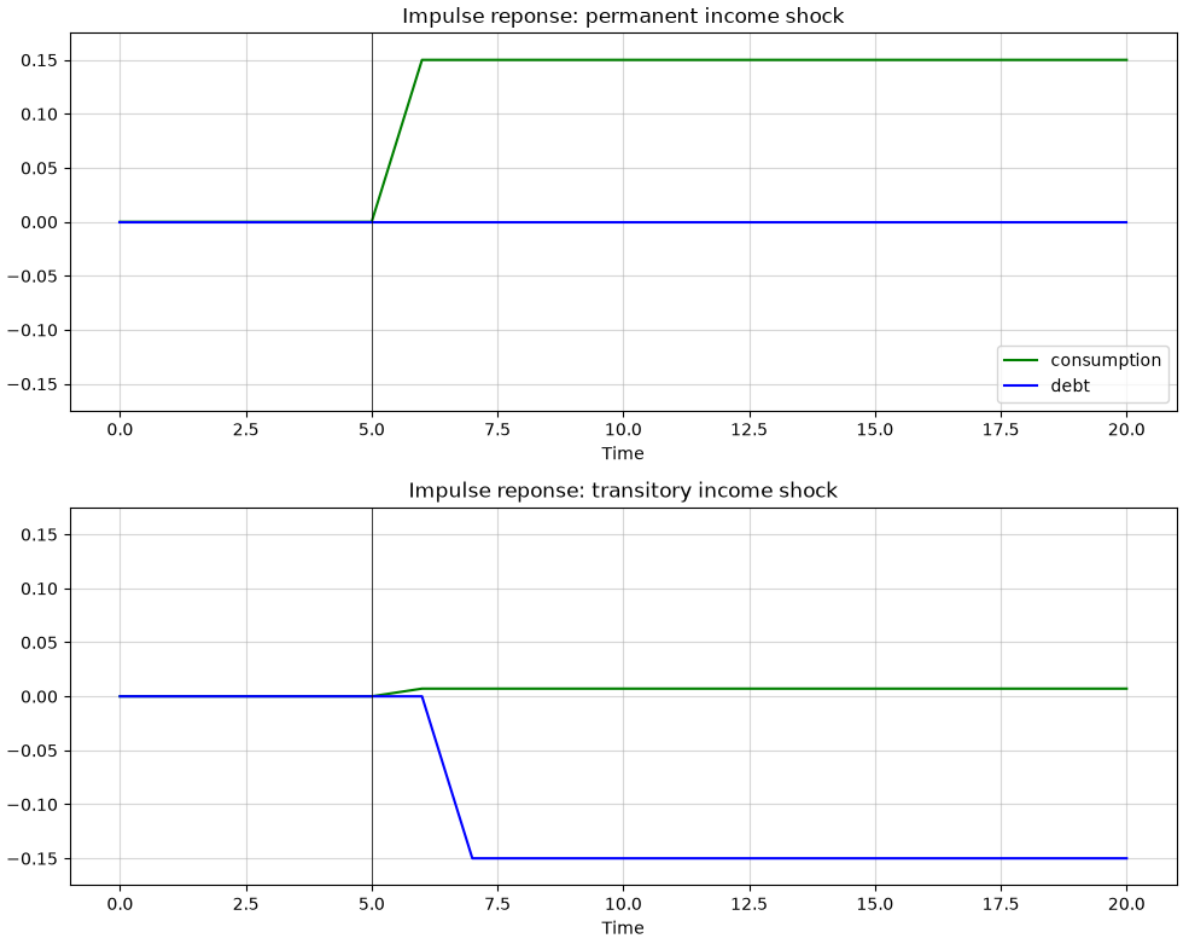

fig, axes = plt.subplots(2, 1, figsize=(10, 8))
titles = ['permanent', 'transitory']

L = 0.175

for ax, truefalse, title in zip(axes, (True, False), titles):
    b, c = time_path(T=20, permanent=truefalse)
    ax.set_title(f'Impulse reponse: {title} income shock')
    ax.plot(c, 'g-', label="consumption")
    ax.plot(b, 'b-', label="debt")
    ax.plot((S, S), (-L, L), 'k-', lw=0.5)
    ax.grid(alpha=0.5)
    ax.set(xlabel=r'Time', ylim=(-L, L))

axes[0].legend(loc='lower right')

plt.tight_layout()
plt.show()
```



Notice how the permanent income shock provokes no change in assets $-b_{t+1}$ and an immediate permanent change in consumption equal to the permanent increment in non-financial income.

In contrast, notice how most of a transitory income shock is saved and only a small amount is saved.

The box-like impulse responses of consumption to both types of shock reflect the random walk property of the optimal consumption decision.

80.4.2 Example 2

Assume now that at time t the consumer observes y_t, and its history up to t, but not z_t.

Under this assumption, it is appropriate to use an *innovation representation* to form A, C, U in (80.18).

The discussion in sections 2.9.1 and 2.11.3 of [Ljungqvist and Sargent, 2018] shows that the pertinent state space representation for y_t is

$$\begin{bmatrix} y_{t+1} \\ a_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & -(1-K) \\ 0 & 0 \end{bmatrix} \begin{bmatrix} y_t \\ a_t \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \end{bmatrix} a_{t+1}$$

$$y_t = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} y_t \\ a_t \end{bmatrix}$$

where

- $K :=$ the stationary Kalman gain

- $a_t := y_t - E[y_t \mid y_{t-1}, \ldots, y_0]$

In the same discussion in [Ljungqvist and Sargent, 2018] it is shown that $K \in [0, 1]$ and that K increases as σ_1/σ_2 does.

In other words, K increases as the ratio of the standard deviation of the permanent shock to that of the transitory shock increases.

Please see *first look at the Kalman filter*.

Applying formulas (80.18) implies

$$c_{t+1} - c_t = [1 - \beta(1 - K)]a_{t+1} \tag{80.28}$$

where the endowment process can now be represented in terms of the univariate innovation to y_t as

$$y_{t+1} - y_t = a_{t+1} - (1 - K)a_t \tag{80.29}$$

Equation (80.29) indicates that the consumer regards

- fraction K of an innovation a_{t+1} to y_{t+1} as *permanent*
- fraction $1 - K$ as purely transitory

The consumer permanently increases his consumption by the full amount of his estimate of the permanent part of a_{t+1}, but by only $(1 - \beta)$ times his estimate of the purely transitory part of a_{t+1}.

Therefore, in total, he permanently increments his consumption by a fraction $K + (1 - \beta)(1 - K) = 1 - \beta(1 - K)$ of a_{t+1}.

He saves the remaining fraction $\beta(1 - K)$.

According to equation (80.29), the first difference of income is a first-order moving average.

Equation (80.28) asserts that the first difference of consumption is IID.

Application of formula to this example shows that

$$b_{t+1} - b_t = (K - 1)a_t \tag{80.30}$$

This indicates how the fraction K of the innovation to y_t that is regarded as permanent influences the fraction of the innovation that is saved.

80.5 Further Reading

The model described above significantly changed how economists think about consumption.

While Hall's model does a remarkably good job as a first approximation to consumption data, it's widely believed that it doesn't capture important aspects of some consumption/savings data.

For example, liquidity constraints and precautionary savings appear to be present sometimes.

Further discussion can be found in, e.g., [Hall and Mishkin, 1982], [Parker, 1999], [Deaton, 1991], [Carroll, 2001].

80.6 Appendix: The Euler Equation

Where does the first-order condition (80.5) come from?

Here we'll give a proof for the two-period case, which is representative of the general argument.

The finite horizon equivalent of the no-Ponzi condition is that the agent cannot end her life in debt, so $b_2 = 0$.

From the budget constraint (80.2) we then have

$$c_0 = \frac{b_1}{1+r} - b_0 + y_0 \quad \text{and} \quad c_1 = y_1 - b_1$$

Here b_0 and y_0 are given constants.

Substituting these constraints into our two-period objective $u(c_0) + \beta \mathbb{E}_0[u(c_1)]$ gives

$$\max_{b_1} \left\{ u\left(\frac{b_1}{R} - b_0 + y_0\right) + \beta \, \mathbb{E}_0[u(y_1 - b_1)] \right\}$$

You will be able to verify that the first-order condition is

$$u'(c_0) = \beta R \, \mathbb{E}_0[u'(c_1)]$$

Using $\beta R = 1$ gives (80.5) in the two-period case.

The proof for the general case is similar.

CHAPTER

EIGHTYONE

PERMANENT INCOME II: LQ TECHNIQUES

Contents

- *Permanent Income II: LQ Techniques*
  - *Overview*
  - *Setup*
  - *The LQ Approach*
  - *Implementation*
  - *Two Example Economies*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

81.1 Overview

This lecture continues our analysis of the linear-quadratic (LQ) permanent income model of savings and consumption.

As we saw in our *previous lecture* on this topic, Robert Hall [Hall, 1978] used the LQ permanent income model to restrict and interpret intertemporal comovements of nondurable consumption, nonfinancial income, and financial wealth.

For example, we saw how the model asserts that for any covariance stationary process for nonfinancial income

- consumption is a random walk
- financial wealth has a unit root and is cointegrated with consumption

Other applications use the same LQ framework.

For example, a model isomorphic to the LQ permanent income model has been used by Robert Barro [Barro, 1979] to interpret intertemporal comovements of a government's tax collections, its expenditures net of debt service, and its public debt.

This isomorphism means that in analyzing the LQ permanent income model, we are in effect also analyzing the Barro tax smoothing model.

It is just a matter of appropriately relabeling the variables in Hall's model.

In this lecture, we'll

- show how the solution to the LQ permanent income model can be obtained using LQ control methods.
- represent the model as a linear state space system as in *this lecture*.
- apply QuantEcon's LinearStateSpace class to characterize statistical features of the consumer's optimal consumption and borrowing plans.

We'll then use these characterizations to construct a simple model of cross-section wealth and consumption dynamics in the spirit of Truman Bewley [Bewley, 1986].

(Later we'll study other Bewley models—see *this lecture*.)

The model will prove useful for illustrating concepts such as

- stationarity
- ergodicity
- ensemble moments and cross-section observations

Let's start with some imports:

```
import matplotlib.pyplot as plt
import quantecon as qe
import numpy as np
import scipy.linalg as la
```

81.2 Setup

Let's recall the basic features of the model discussed in the *permanent income model*.

Consumer preferences are ordered by

$$E_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{81.1}$$

where $u(c) = -(c - \gamma)^2$.

The consumer maximizes (81.1) by choosing a consumption, borrowing plan $\{c_t, b_{t+1}\}_{t=0}^{\infty}$ subject to the sequence of budget constraints

$$c_t + b_t = \frac{1}{1+r} b_{t+1} + y_t, \quad t \geq 0 \tag{81.2}$$

and the no-Ponzi condition

$$E_0 \sum_{t=0}^{\infty} \beta^t b_t^2 < \infty \tag{81.3}$$

The interpretation of all variables and parameters are the same as in the *previous lecture*.

We continue to assume that $(1+r)\beta = 1$.

The dynamics of $\{y_t\}$ again follow the linear state space model

$$\begin{aligned} z_{t+1} &= A z_t + C w_{t+1} \\ y_t &= U z_t \end{aligned} \tag{81.4}$$

The restrictions on the shock process and parameters are the same as in our *previous lecture*.

81.2.1 Digression on a Useful Isomorphism

The LQ permanent income model of consumption is mathematically isomorphic with a version of Barro's [Barro, 1979] model of tax smoothing.

In the LQ permanent income model

- the household faces an exogenous process of nonfinancial income
- the household wants to smooth consumption across states and time

In the Barro tax smoothing model

- a government faces an exogenous sequence of government purchases (net of interest payments on its debt)
- a government wants to smooth tax collections across states and time

If we set

- T_t, total tax collections in Barro's model to consumption c_t in the LQ permanent income model.
- G_t, exogenous government expenditures in Barro's model to nonfinancial income y_t in the permanent income model.
- B_t, government risk-free one-period assets falling due in Barro's model to risk-free one-period consumer debt b_t falling due in the LQ permanent income model.
- R, the gross rate of return on risk-free one-period government debt in Barro's model to the gross rate of return $1 + r$ on financial assets in the permanent income model of consumption.

then the two models are mathematically equivalent.

All characterizations of a $\{c_t, y_t, b_t\}$ in the LQ permanent income model automatically apply to a $\{T_t, G_t, B_t\}$ process in the Barro model of tax smoothing.

See consumption and tax smoothing models for further exploitation of an isomorphism between consumption and tax smoothing models.

81.2.2 A Specification of the Nonfinancial Income Process

For the purposes of this lecture, let's assume $\{y_t\}$ is a second-order univariate autoregressive process:

$$y_{t+1} = \alpha + \rho_1 y_t + \rho_2 y_{t-1} + \sigma w_{t+1}$$

We can map this into the linear state space framework in (81.4), as discussed in our lecture on *linear models*.

To do so we take

$$z_t = \begin{bmatrix} 1 \\ y_t \\ y_{t-1} \end{bmatrix}, \quad A = \begin{bmatrix} 1 & 0 & 0 \\ \alpha & \rho_1 & \rho_2 \\ 0 & 1 & 0 \end{bmatrix}, \quad C = \begin{bmatrix} 0 \\ \sigma \\ 0 \end{bmatrix}, \quad \text{and} \quad U = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

81.3 The LQ Approach

Previously we solved the permanent income model by solving a system of linear expectational difference equations subject to two boundary conditions.

Here we solve the same model using *LQ methods* based on dynamic programming.

After confirming that answers produced by the two methods agree, we apply QuantEcon's LinearStateSpace class to illustrate features of the model.

Why solve a model in two distinct ways?

Because by doing so we gather insights about the structure of the model.

Our earlier approach based on solving a system of expectational difference equations brought to the fore the role of the consumer's expectations about future nonfinancial income.

On the other hand, formulating the model in terms of an LQ dynamic programming problem reminds us that

- finding the state (of a dynamic programming problem) is an art, and
- iterations on a Bellman equation implicitly jointly solve both a forecasting problem and a control problem

81.3.1 The LQ Problem

Recall from our *lecture on LQ theory* that the optimal linear regulator problem is to choose a decision rule for u_t to minimize

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \{x_t' R x_t + u_t' Q u_t\},$$

subject to x_0 given and the law of motion

$$x_{t+1} = \tilde{A} x_t + \tilde{B} u_t + \tilde{C} w_{t+1}, \qquad t \geq 0, \tag{81.5}$$

where w_{t+1} is IID with mean vector zero and $\mathbb{E} w_t w_t' = I$.

The tildes in $\tilde{A}, \tilde{B}, \tilde{C}$ are to avoid clashing with notation in (81.4).

The value function for this problem is $v(x) = -x' P x - d$, where

- P is the unique positive semidefinite solution of the *corresponding matrix Riccati equation*.
- The scalar d is given by $d = \beta(1-\beta)^{-1} \text{trace}(P \tilde{C} \tilde{C}')$.

The optimal policy is $u_t = -F x_t$, where $F := \beta(Q + \beta \tilde{B}' P \tilde{B})^{-1} \tilde{B}' P \tilde{A}$.

Under an optimal decision rule F, the state vector x_t evolves according to $x_{t+1} = (\tilde{A} - \tilde{B}F) x_t + \tilde{C} w_{t+1}$.

81.3.2 Mapping into the LQ Framework

To map into the LQ framework, we'll use

$$x_t := \begin{bmatrix} z_t \\ b_t \end{bmatrix} = \begin{bmatrix} 1 \\ y_t \\ y_{t-1} \\ b_t \end{bmatrix}$$

as the state vector and $u_t := c_t - \gamma$ as the control.

With this notation and $U_\gamma := \begin{bmatrix}\gamma & 0 & 0\end{bmatrix}$, we can write the state dynamics as in (81.5) when

$$\tilde{A} := \begin{bmatrix} A & 0 \\ (1+r)(U_\gamma - U) & 1+r \end{bmatrix} \quad \tilde{B} := \begin{bmatrix} 0 \\ 1+r \end{bmatrix} \quad \text{and} \quad \tilde{C} := \begin{bmatrix} C \\ 0 \end{bmatrix} w_{t+1}$$

Please confirm for yourself that, with these definitions, the LQ dynamics (81.5) match the dynamics of z_t and b_t described above.

To map utility into the quadratic form $x_t' R x_t + u_t' Q u_t$ we can set

- $Q := 1$ (remember that we are minimizing) and
- $R :=$ a 4×4 matrix of zeros

However, there is one problem remaining.

We have no direct way to capture the non-recursive restriction (81.3) on the debt sequence $\{b_t\}$ from within the LQ framework.

To try to enforce it, we're going to use a trick: put a small penalty on b_t^2 in the criterion function.

In the present setting, this means adding a small entry $\epsilon > 0$ in the $(4, 4)$ position of R.

That will induce a (hopefully) small approximation error in the decision rule.

We'll check whether it really is small numerically soon.

81.4 Implementation

Let's write some code to solve the model.

One comment before we start is that the bliss level of consumption γ in the utility function has no effect on the optimal decision rule.

We saw this in the previous lecture *permanent income*.

The reason is that it drops out of the Euler equation for consumption.

In what follows we set it equal to unity.

81.4.1 The Exogenous Nonfinancial Income Process

First, we create the objects for the optimal linear regulator

```
# Set parameters
α, β, ρ1, ρ2, σ = 10.0, 0.95, 0.9, 0.0, 1.0

R = 1 / β
A = np.array([[1., 0., 0.],
              [α,  ρ1, ρ2],
              [0., 1., 0.]])
C = np.array([[0.], [σ], [0.]])
G = np.array([[0., 1., 0.]])

# Form LinearStateSpace system and pull off steady state moments
μ_z0 = np.array([[1.0], [0.0], [0.0]])
Σ_z0 = np.zeros((3, 3))
Lz = qe.LinearStateSpace(A, C, G, mu_0=μ_z0, Sigma_0=Σ_z0)
μ_z, μ_y, Σ_z, Σ_y, Σ_yx = Lz.stationary_distributions()
```

(continues on next page)

(continued from previous page)

```
# Mean vector of state for the savings problem
mxo = np.vstack([μ_z, 0.0])

# Create stationary covariance matrix of x -- start everyone off at b=0
a1 = np.zeros((3, 1))
aa = np.hstack([Σ_z, a1])
bb = np.zeros((1, 4))
sxo = np.vstack([aa, bb])

# These choices will initialize the state vector of an individual at zero
# debt and the ergodic distribution of the endowment process. Use these to
# create the Bewley economy.
mxbewley = mxo
sxbewley = sxo
```

The next step is to create the matrices for the LQ system

```
A12 = np.zeros((3,1))
ALQ_l = np.hstack([A, A12])
ALQ_r = np.array([[0, -R, 0, R]])
ALQ = np.vstack([ALQ_l, ALQ_r])

RLQ = np.array([[0., 0., 0., 0.],
                [0., 0., 0., 0.],
                [0., 0., 0., 0.],
                [0., 0., 0., 1e-9]])

QLQ = np.array([1.0])
BLQ = np.array([0., 0., 0., R]).reshape(4,1)
CLQ = np.array([0., σ, 0., 0.]).reshape(4,1)
β_LQ = β
```

Let's print these out and have a look at them

```
print(f"A = \n {ALQ}")
print(f"B = \n {BLQ}")
print(f"R = \n {RLQ}")
print(f"Q = \n {QLQ}")
```

```
A =
 [[ 1.          0.          0.          0.        ]
 [10.          0.9         0.          0.        ]
 [ 0.          1.          0.          0.        ]
 [ 0.         -1.05263158  0.          1.05263158]]
B =
 [[0.        ]
 [0.        ]
 [0.        ]
 [1.05263158]]
R =
 [[0.e+00 0.e+00 0.e+00 0.e+00]
 [0.e+00 0.e+00 0.e+00 0.e+00]
 [0.e+00 0.e+00 0.e+00 0.e+00]
 [0.e+00 0.e+00 0.e+00 1.e-09]]
Q =
 [1.]
```

Now create the appropriate instance of an LQ model

```
lqpi = qe.LQ(QLQ, RLQ, ALQ, BLQ, C=CLQ, beta=β_LQ)
```

We'll save the implied optimal policy function soon compare them with what we get by employing an alternative solution method

```
P, F, d = lqpi.stationary_values()  # Compute value function and decision rule
ABF = ALQ - BLQ @ F  #  Form closed loop system
```

81.4.2 Comparison with the Difference Equation Approach

In our *first lecture* on the infinite horizon permanent income problem we used a different solution method.

The method was based around

- deducing the Euler equations that are the first-order conditions with respect to consumption and savings.
- using the budget constraints and boundary condition to complete a system of expectational linear difference equations.
- solving those equations to obtain the solution.

Expressed in state space notation, the solution took the form

$$
\begin{aligned}
z_{t+1} &= Az_t + Cw_{t+1} \\
b_{t+1} &= b_t + U[(I - \beta A)^{-1}(A - I)]z_t \\
y_t &= Uz_t \\
c_t &= (1 - \beta)[U(I - \beta A)^{-1}z_t - b_t]
\end{aligned}
$$

Now we'll apply the formulas in this system

```
# Use the above formulas to create the optimal policies for b_{t+1} and c_t
b_pol = G @ la.inv(np.eye(3, 3) - β * A) @ (A - np.eye(3, 3))
c_pol = (1 - β) * G @ la.inv(np.eye(3, 3) - β * A)

# Create the A matrix for a LinearStateSpace instance
A_LSS1 = np.vstack([A, b_pol])
A_LSS2 = np.eye(4, 1, -3)
A_LSS = np.hstack([A_LSS1, A_LSS2])

# Create the C matrix for LSS methods
C_LSS = np.vstack([C, np.zeros(1)])

# Create the G matrix for LSS methods
G_LSS1 = np.vstack([G, c_pol])
G_LSS2 = np.vstack([np.zeros(1), -(1 - β)])
G_LSS = np.hstack([G_LSS1, G_LSS2])

# Use the following values to start everyone off at b=0, initial incomes zero
μ_0 = np.array([1., 0., 0., 0.])
Σ_0 = np.zeros((4, 4))
```

`A_LSS` calculated as we have here should equal `ABF` calculated above using the LQ model

```
ABF - A_LSS
```

```
array([[ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
         0.00000000e+00],
       [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
         0.00000000e+00],
       [ 0.00000000e+00,  0.00000000e+00,  0.00000000e+00,
         0.00000000e+00],
       [-9.51248364e-06,  9.51247768e-08,  0.00000000e+00,
        -1.99999900e-08]])
```

Now compare pertinent elements of `c_pol` and `F`

```
print(c_pol, "\n", -F)
```

```
[[65.51724138  0.34482759  0.        ]]
 [[ 6.55172323e+01  3.44827677e-01 -0.00000000e+00 -5.00000190e-02]]
```

We have verified that the two methods give the same solution.

Now let's create instances of the LinearStateSpace class and use it to do some interesting experiments.

To do this, we'll use the outcomes from our second method.

81.5 Two Example Economies

In the spirit of Bewley models [Bewley, 1986], we'll generate panels of consumers.

The examples differ only in the initial states with which we endow the consumers.

All other parameter values are kept the same in the two examples

- In the first example, all consumers begin with zero nonfinancial income and zero debt.
  - The consumers are thus *ex-ante* identical.
- In the second example, while all begin with zero debt, we draw their initial income levels from the invariant distribution of financial income.
  - Consumers are *ex-ante* heterogeneous.

In the first example, consumers' nonfinancial income paths display pronounced transients early in the sample

- these will affect outcomes in striking ways

Those transient effects will not be present in the second example.

We use methods affiliated with the LinearStateSpace class to simulate the model.

81.5.1 First Set of Initial Conditions

We generate 25 paths of the exogenous non-financial income process and the associated optimal consumption and debt paths.

In the first set of graphs, darker lines depict a particular sample path, while the lighter lines describe 24 other paths.

A second graph plots a collection of simulations against the population distribution that we extract from the `LinearStateSpace` instance `LSS`.

Comparing sample paths with population distributions at each date t is a useful exercise—see *our discussion* of the laws of large numbers

```
lss = qe.LinearStateSpace(A_LSS, C_LSS, G_LSS, mu_0=μ_0, Sigma_0=Σ_0)
```

81.5.2 Population and Sample Panels

In the code below, we use the LinearStateSpace class to

- compute and plot population quantiles of the distributions of consumption and debt for a population of consumers.
- simulate a group of 25 consumers and plot sample paths on the same graph as the population distribution.

```
def income_consumption_debt_series(A, C, G, μ_0, Σ_0, T=150, npaths=25):
    """
    This function takes initial conditions (μ_0, Σ_0) and uses the
    LinearStateSpace class from QuantEcon to  simulate an economy
    npaths times for T periods. It then uses that information to
    generate some graphs related to the discussion below.
    """
    lss = qe.LinearStateSpace(A, C, G, mu_0=μ_0, Sigma_0=Σ_0)

    # Simulation/Moment Parameters
    moment_generator = lss.moment_sequence()

    # Simulate various paths
    bsim = np.empty((npaths, T))
    csim = np.empty((npaths, T))
    ysim = np.empty((npaths, T))

    for i in range(npaths):
        sims = lss.simulate(T)
        bsim[i, :] = sims[0][-1, :]
        csim[i, :] = sims[1][1, :]
        ysim[i, :] = sims[1][0, :]

    # Get the moments
    cons_mean = np.empty(T)
    cons_var = np.empty(T)
    debt_mean = np.empty(T)
    debt_var = np.empty(T)
    for t in range(T):
        μ_x, μ_y, Σ_x, Σ_y = next(moment_generator)
        cons_mean[t], cons_var[t] = μ_y[1,0], Σ_y[1, 1]
        debt_mean[t], debt_var[t] = μ_x[3,0], Σ_x[3, 3]

    return bsim, csim, ysim, cons_mean, cons_var, debt_mean, debt_var

def consumption_income_debt_figure(bsim, csim, ysim):

    # Get T
    T =  bsim.shape[1]

    # Create the first figure
    fig, ax = plt.subplots(2, 1, figsize=(10, 8))
    xvals = np.arange(T)

    # Plot consumption and income
    ax[0].plot(csim[0, :], label="c", color="b")
```

(continues on next page)

(continued from previous page)

```
    ax[0].plot(ysim[0, :], label="y", color="g")
    ax[0].plot(csim.T, alpha=.1, color="b")
    ax[0].plot(ysim.T, alpha=.1, color="g")
    ax[0].legend(loc=4)
    ax[0].set(title="Nonfinancial Income, Consumption, and Debt",
              xlabel="t", ylabel="y and c")

    # Plot debt
    ax[1].plot(bsim[0, :], label="b", color="r")
    ax[1].plot(bsim.T, alpha=.1, color="r")
    ax[1].legend(loc=4)
    ax[1].set(xlabel="t", ylabel="debt")

    fig.tight_layout()
    return fig

def consumption_debt_fanchart(csim, cons_mean, cons_var,
                              bsim, debt_mean, debt_var):
    # Get T
    T =  bsim.shape[1]

    # Create percentiles of cross-section distributions
    cmean = np.mean(cons_mean)
    c90 = 1.65 * np.sqrt(cons_var)
    c95 = 1.96 * np.sqrt(cons_var)
    c_perc_95p, c_perc_95m = cons_mean + c95, cons_mean - c95
    c_perc_90p, c_perc_90m = cons_mean + c90, cons_mean - c90

    # Create percentiles of cross-section distributions
    dmean = np.mean(debt_mean)
    d90 = 1.65 * np.sqrt(debt_var)
    d95 = 1.96 * np.sqrt(debt_var)
    d_perc_95p, d_perc_95m = debt_mean + d95, debt_mean - d95
    d_perc_90p, d_perc_90m = debt_mean + d90, debt_mean - d90


    # Create second figure
    fig, ax = plt.subplots(2, 1, figsize=(10, 8))
    xvals = np.arange(T)

    # Consumption fan
    ax[0].plot(xvals, cons_mean, color="k")
    ax[0].plot(csim.T, color="k", alpha=.25)
    ax[0].fill_between(xvals, c_perc_95m, c_perc_95p, alpha=.25, color="b")
    ax[0].fill_between(xvals, c_perc_90m, c_perc_90p, alpha=.25, color="r")
    ax[0].set(title="Consumption/Debt over time",
              ylim=(cmean-15, cmean+15), ylabel="consumption")

    # Debt fan
    ax[1].plot(xvals, debt_mean, color="k")
    ax[1].plot(bsim.T, color="k", alpha=.25)
    ax[1].fill_between(xvals, d_perc_95m, d_perc_95p, alpha=.25, color="b")
    ax[1].fill_between(xvals, d_perc_90m, d_perc_90p, alpha=.25, color="r")
    ax[1].set(xlabel="t", ylabel="debt")

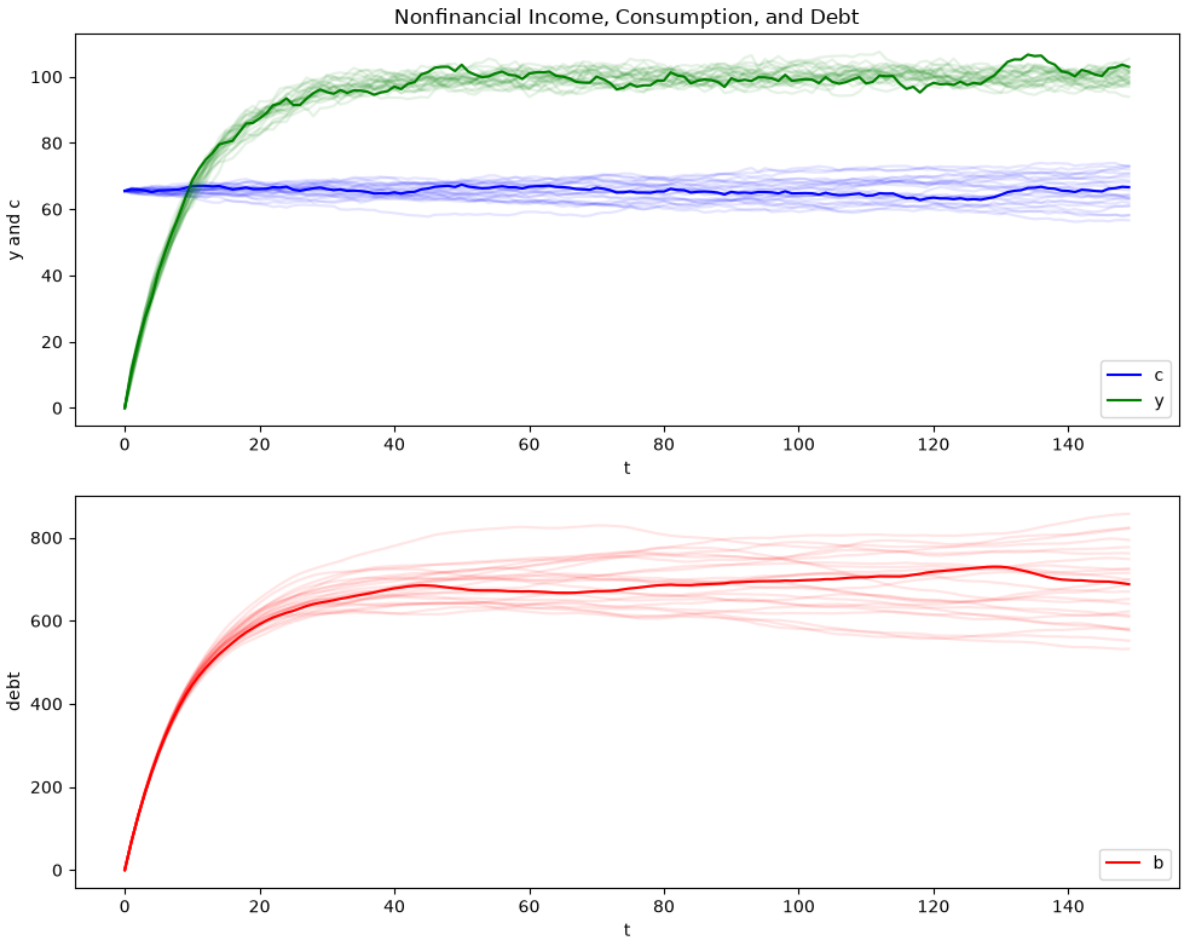
    fig.tight_layout()
    return fig
```

Now let's create figures with initial conditions of zero for y_0 and b_0

```
out = income_consumption_debt_series(A_LSS, C_LSS, G_LSS, μ_0, Σ_0)
bsim0, csim0, ysim0 = out[:3]
cons_mean0, cons_var0, debt_mean0, debt_var0 = out[3:]

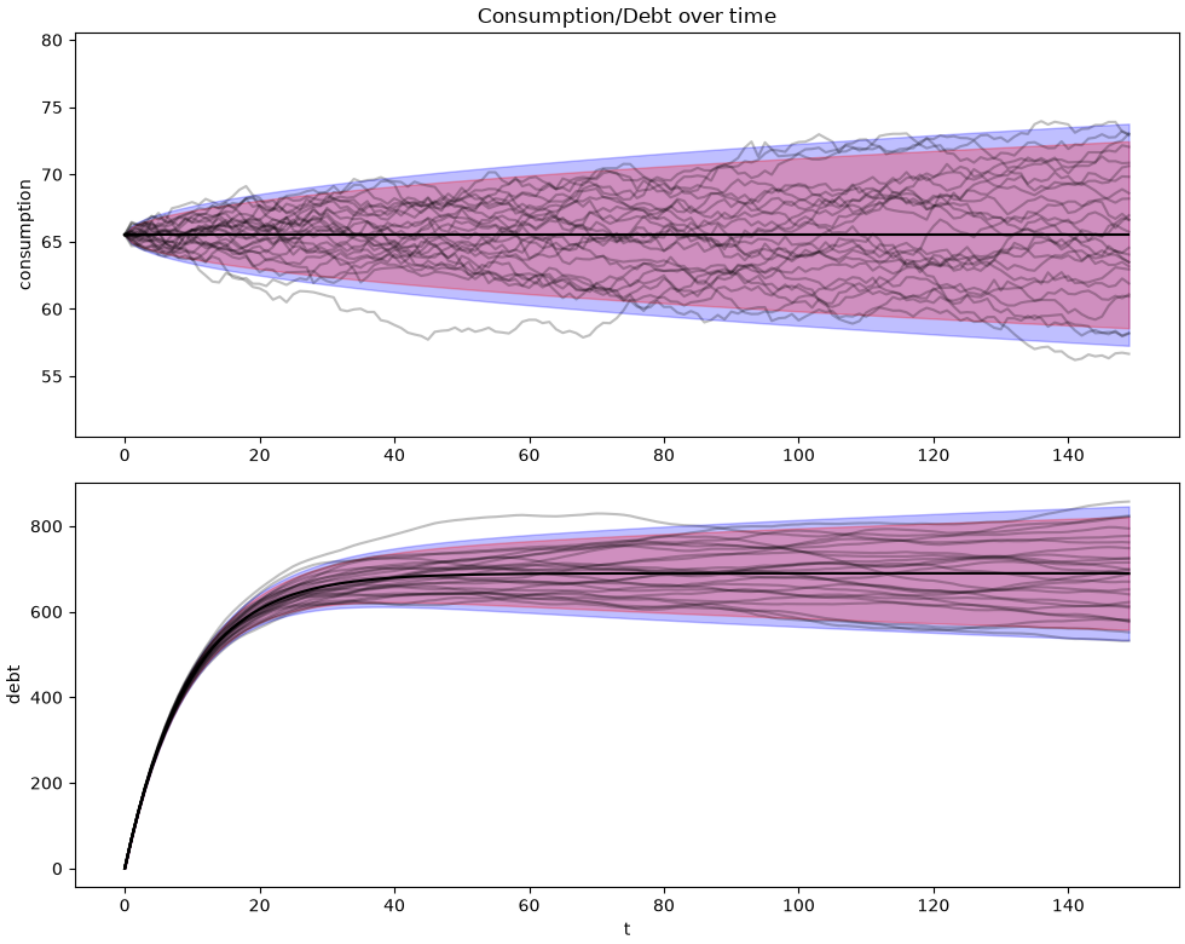
consumption_income_debt_figure(bsim0, csim0, ysim0)

plt.show()
```




```
consumption_debt_fanchart(csim0, cons_mean0, cons_var0,
                          bsim0, debt_mean0, debt_var0)

plt.show()
```



Here is what is going on in the above graphs.

For our simulation, we have set initial conditions $b_0 = y_{-1} = y_{-2} = 0$.

Because $y_{-1} = y_{-2} = 0$, nonfinancial income y_t starts far below its stationary mean $\mu_{y,\infty}$ and rises early in each simulation.

Recall from the *previous lecture* that we can represent the optimal decision rule for consumption in terms of the **cointegrating relationship**

$$(1-\beta)b_t + c_t = (1-\beta)E_t \sum_{j=0}^{\infty} \beta^j y_{t+j} \tag{81.6}$$

So at time 0 we have

$$c_0 = (1-\beta)E_0 \sum_{t=0}^{\infty} \beta^j y_t$$

This tells us that consumption starts at the income that would be paid by an annuity whose value equals the expected discounted value of nonfinancial income at time $t = 0$.

To support that level of consumption, the consumer borrows a lot early and consequently builds up substantial debt.

In fact, he or she incurs so much debt that eventually, in the stochastic steady state, he consumes less each period than his nonfinancial income.

He uses the gap between consumption and nonfinancial income mostly to service the interest payments due on his debt.

Thus, when we look at the panel of debt in the accompanying graph, we see that this is a group of *ex-ante* identical people each of whom starts with zero debt.

All of them accumulate debt in anticipation of rising nonfinancial income.

They expect their nonfinancial income to rise toward the invariant distribution of income, a consequence of our having started them at $y_{-1} = y_{-2} = 0$.

Cointegration Residual

The following figure plots realizations of the left side of (81.6), which, *as discussed in our last lecture*, is called the **cointegrating residual**.

As mentioned above, the right side can be thought of as an annuity payment on the expected present value of future income $E_t \sum_{j=0}^{\infty} \beta^j y_{t+j}$.

Early along a realization, c_t is approximately constant while $(1-\beta)b_t$ and $(1-\beta)E_t \sum_{j=0}^{\infty} \beta^j y_{t+j}$ both rise markedly as the household's present value of income and borrowing rise pretty much together.

This example illustrates the following point: the definition of cointegration implies that the cointegrating residual is *asymptotically* covariance stationary, not *covariance stationary*.

The cointegrating residual for the specification with zero income and zero debt initially has a notable transient component that dominates its behavior early in the sample.

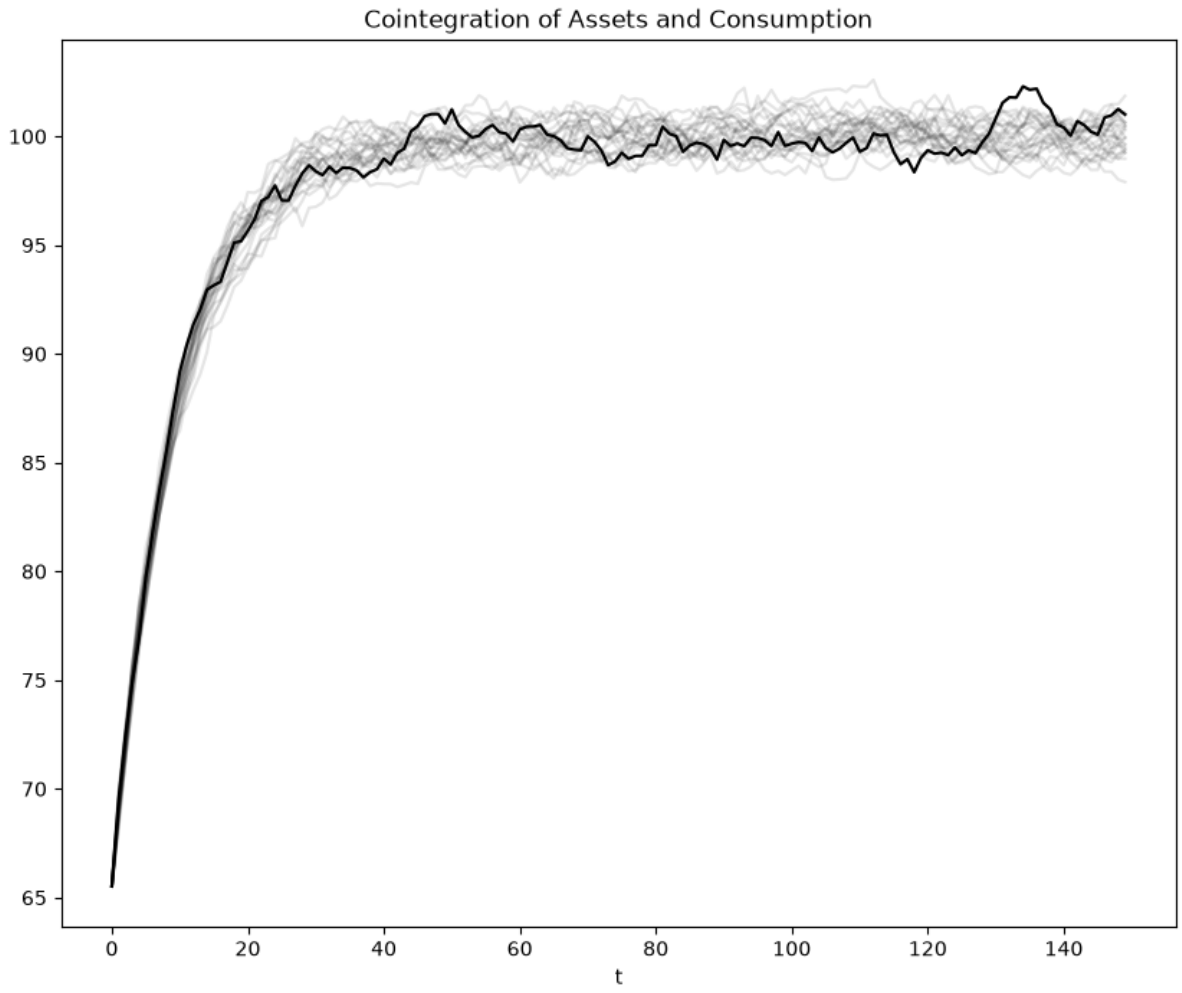
By altering initial conditions, we shall remove this transient in our second example to be presented below

```
def cointegration_figure(bsim, csim):
    """
    Plots the cointegration
    """
    # Create figure
    fig, ax = plt.subplots(figsize=(10, 8))
    ax.plot((1 - β) * bsim[0, :] + csim[0, :], color="k")
    ax.plot((1 - β) * bsim.T + csim.T, color="k", alpha=.1)

    ax.set(title="Cointegration of Assets and Consumption", xlabel="t")

    return fig
```

```
cointegration_figure(bsim0, csim0)
plt.show()
```



81.5.3 A "Borrowers and Lenders" Closed Economy

When we set $y_{-1} = y_{-2} = 0$ and $b_0 = 0$ in the preceding exercise, we make debt "head north" early in the sample.

Average debt in the cross-section rises and approaches the asymptote.

We can regard these as outcomes of a "small open economy" that borrows from abroad at the fixed gross interest rate $R = r + 1$ in anticipation of rising incomes.

So with the economic primitives set as above, the economy converges to a steady state in which there is an excess aggregate supply of risk-free loans at a gross interest rate of R.

This excess supply is filled by "foreigner lenders" willing to make those loans.

We can use virtually the same code to rig a "poor man's Bewley [Bewley, 1986] model" in the following way

- as before, we start everyone at $b_0 = 0$.
- But instead of starting everyone at $y_{-1} = y_{-2} = 0$, we draw $\begin{bmatrix} y_{-1} \\ y_{-2} \end{bmatrix}$ from the invariant distribution of the $\{y_t\}$ process.

This rigs a closed economy in which people are borrowing and lending with each other at a gross risk-free interest rate of $R = \beta^{-1}$.

Across the group of people being analyzed, risk-free loans are in zero excess supply.

We have arranged primitives so that $R = \beta^{-1}$ clears the market for risk-free loans at zero aggregate excess supply.

So the risk-free loans are being made from one person to another within our closed set of agents.

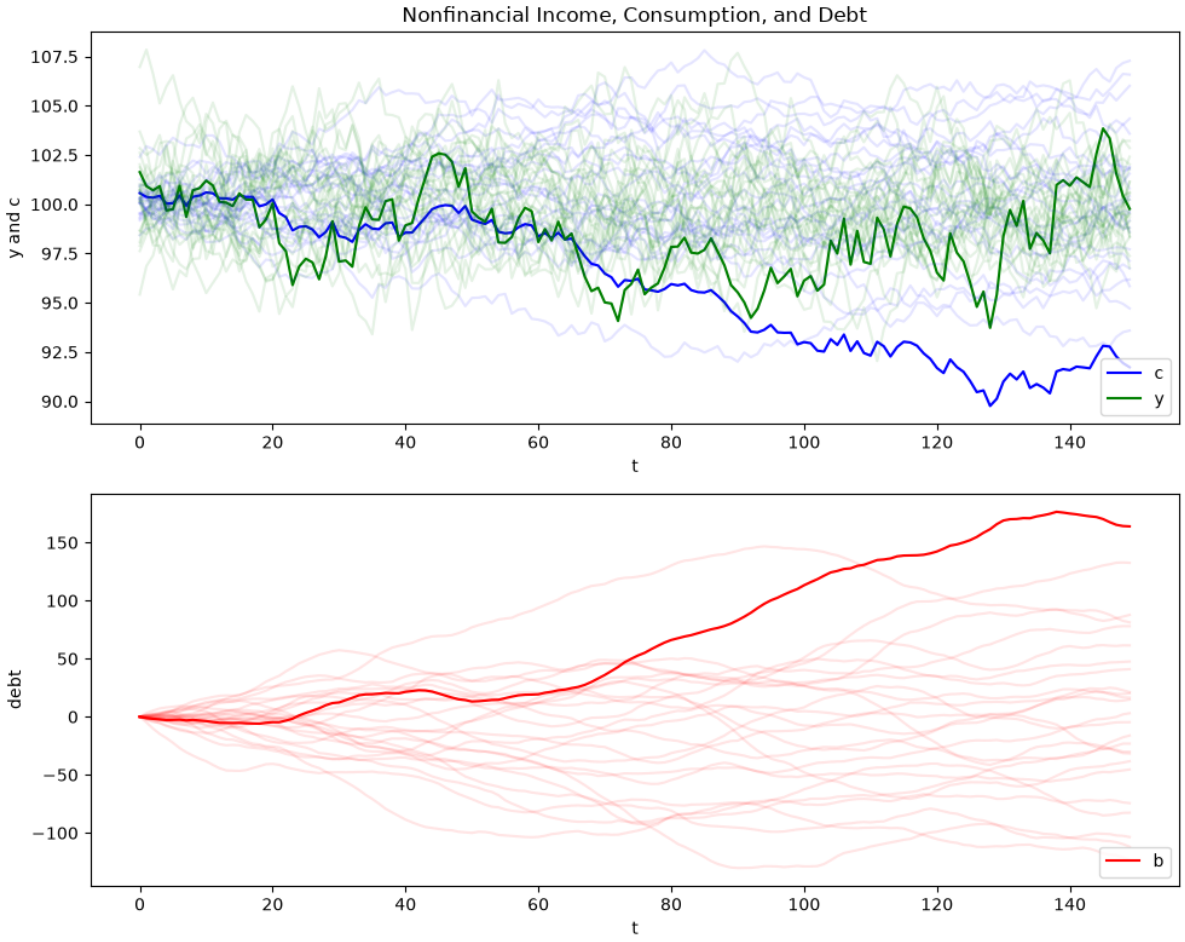
There is no need for foreigners to lend to our group.

Let's have a look at the corresponding figures

```
out = income_consumption_debt_series(A_LSS, C_LSS, G_LSS, mxbewley, sxbewley)
bsimb, csimb, ysimb = out[:3]
cons_meanb, cons_varb, debt_meanb, debt_varb = out[3:]

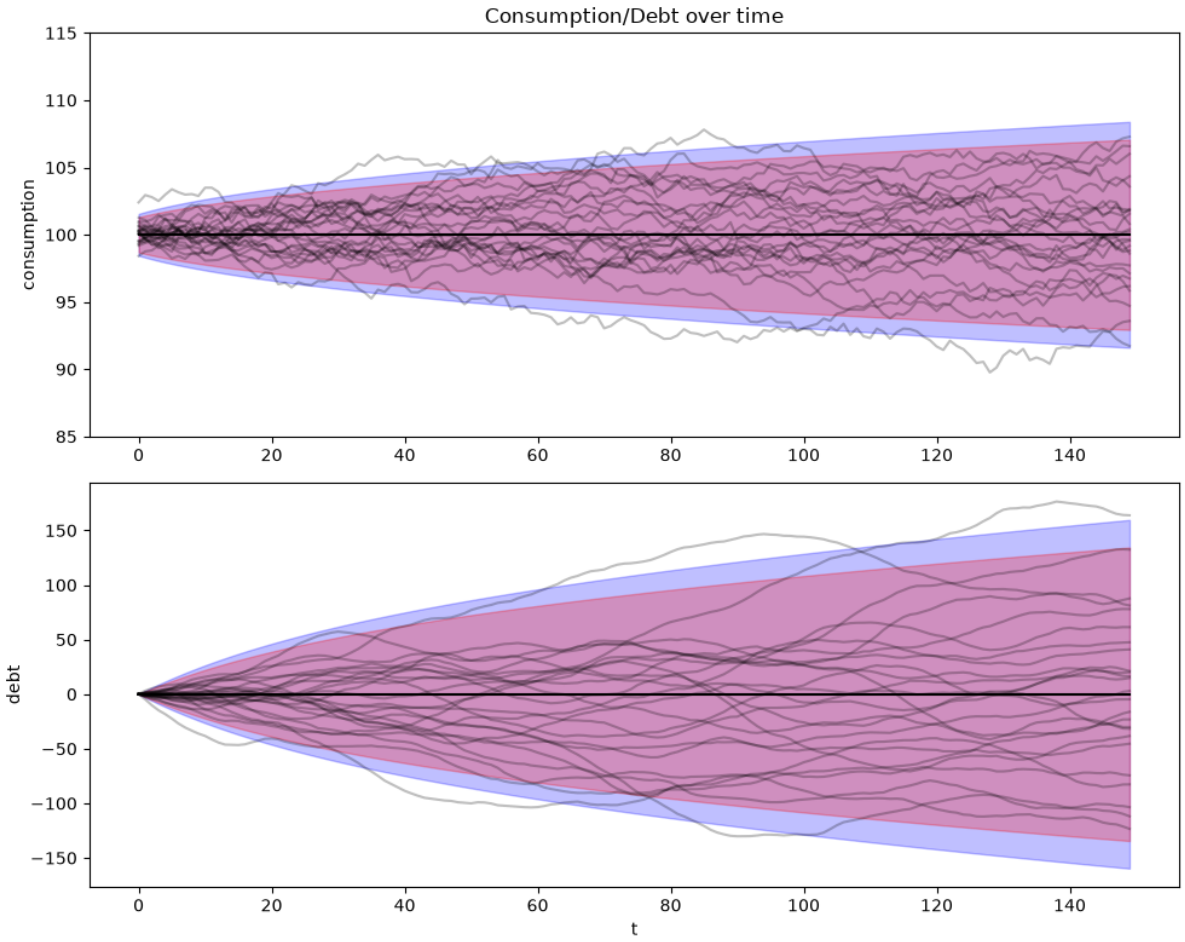
consumption_income_debt_figure(bsimb, csimb, ysimb)

plt.show()
```




```
consumption_debt_fanchart(csimb, cons_meanb, cons_varb,
                          bsimb, debt_meanb, debt_varb)

plt.show()
```



The graphs confirm the following outcomes:

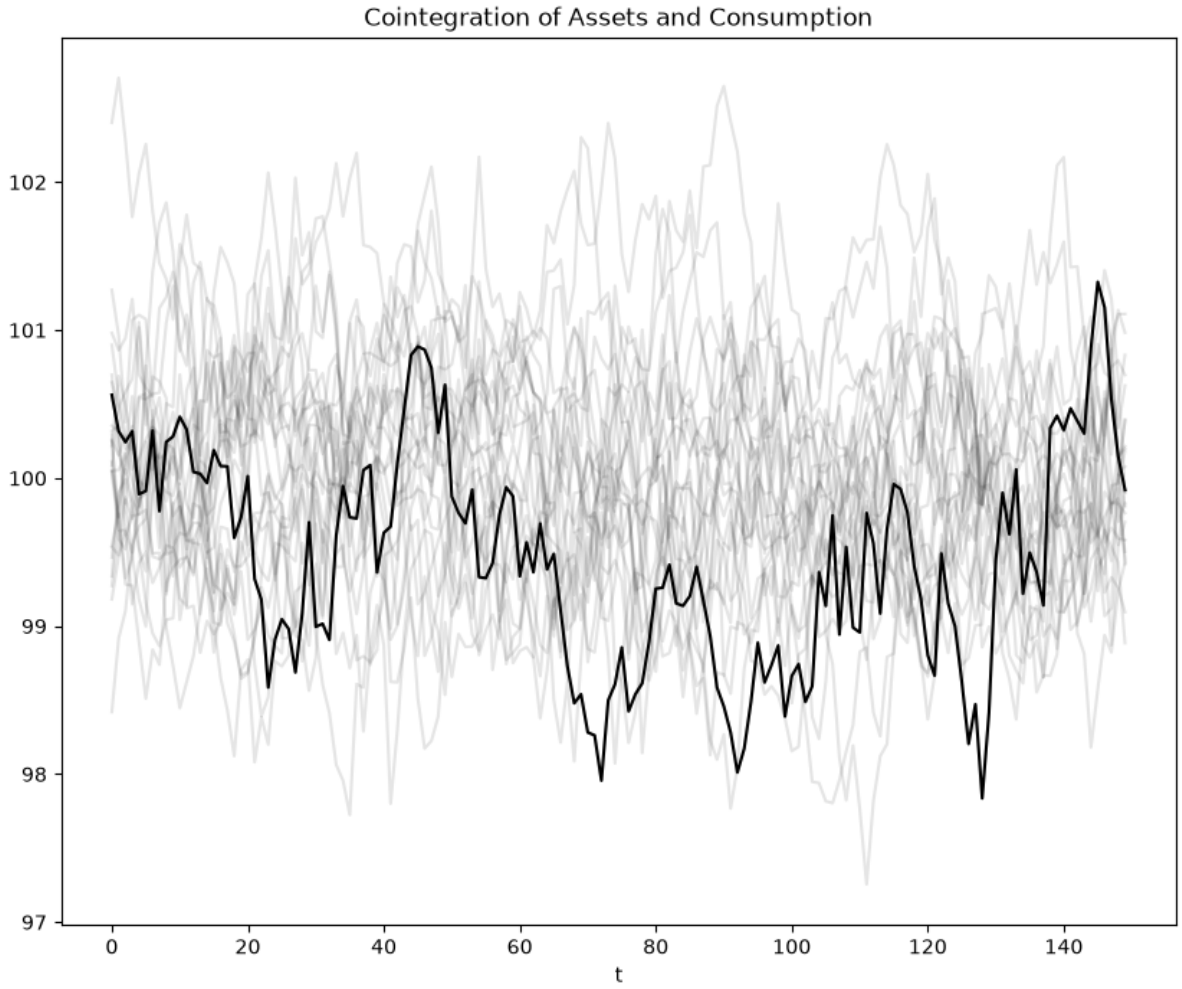
- As before, the consumption distribution spreads out over time.

But now there is some initial dispersion because there is *ex-ante* heterogeneity in the initial draws of $\begin{bmatrix} y_{-1} \\ y_{-2} \end{bmatrix}$.

- As before, the cross-section distribution of debt spreads out over time.
- Unlike before, the average level of debt stays at zero, confirming that this is a closed borrower-and-lender economy.
- Now the cointegrating residual seems stationary, and not just asymptotically stationary.

Let's have a look at the cointegration figure

```
cointegration_figure(bsimb, csimb)
plt.show()
```


Cointegration of Assets and Consumption
102
101
100
99
98
97
0
20
40
60
80
100
120
140
t

CHAPTER

EIGHTYTWO

ROBUST PERMANENT INCOME AND PRICING

Contents

- *Robust Permanent Income and Pricing*
  - *Overview*
  - *Risk-sensitive recursive preferences*
  - *A preference for robustness*
  - *The permanent income economy*
  - *Observational equivalence*
  - *Estimation*
  - *Asset pricing and the market price of risk*
  - *A risk-sensitive regulator*
  - *Exercises*
  - *Related lectures*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install --upgrade quantecon
```

82.1 Overview

This lecture studies the model of Hansen *et al.* [1999].

The paper asks a simple question with surprising consequences.

What happens to a classic permanent income consumer when, instead of trusting a single probability model of their income, they *fear that their model is misspecified* and want decisions that work well across a *family* of nearby models?

Such a consumer is called a **robust** decision maker.

The central findings are:

- A preference for robustness is hidden inside the quantity implications of the ordinary permanent income model.
- Robustness and risk-sensitivity are two interpretations of the *same* decision rules, with a single parameter σ governing both.

- Concern about *small* amounts of model misspecification can show up as *large* market-based measures of risk aversion.
- The consumption and savings data alone cannot identify the robustness parameter: the model is **observationally equivalent** to a standard permanent income model with a lower discount factor.
- But asset prices, in particular the **market price of risk**, *can* be used to pin robustness down.

We will learn about

- risk-sensitive recursive preferences and the operator $\mathcal{R}_t$
- how a malevolent "second agent" implements a preference for robustness through a two-player zero-sum game
- the link between robustness and **Knightian uncertainty** in the sense of Gilboa and Schmeidler [1989] and Epstein and Wang [1994]
- an **observational equivalence** result that we will reproduce numerically
- how a small worst-case distortion of conditional means translates almost one-for-one into a market price of risk

This lecture builds on ideas in *The Permanent Income Model*, *Permanent Income II: LQ Techniques*, and *LQ Control: Foundations*.

The robustness machinery here is developed at book length in Hansen and Sargent [2008], extended in Anderson *et al.* [2003], and reinterpreted through detection-error probabilities in Barillas *et al.* [2009].

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import brentq
from scipy.stats import norm
import quantecon as qe
```

82.2 Risk-sensitive recursive preferences

The theory rests on a recursive linear-quadratic optimization problem with a twist in how continuation utility is aggregated.

The state evolves as

$$x_{t+1} = Ax_t + Bi_t + Cw_{t+1}, \tag{82.1}$$

where i_t is a control vector, x_t is the state vector, and w_{t+1} is an IID Gaussian vector with $\mathbb{E} w_{t+1} = 0$ and $\mathbb{E} w_{t+1} w_{t+1}^\top = I$.

The one-period return is

$$u(i_t, x_t) = -i_t^\top Q i_t - x_t^\top R x_t, \tag{82.2}$$

with Q positive definite and R positive semidefinite.

Following Epstein and Zin [1989], Weil [1989], and Hansen and Sargent [1995], intertemporal preferences are induced by the recursion

$$U_t = u(i_t, x_t) + \beta \, \mathcal{R}_t(U_{t+1}), \tag{82.3}$$

where the **risk-sensitivity operator** is

$$\mathcal{R}_t(U_{t+1}) \equiv \frac{2}{\sigma} \log \mathbb{E}\left[\exp\left(\frac{\sigma U_{t+1}}{2} \right) \Big| J_t \right]. \tag{82.4}$$

Here J_t is the information available at t.

When $\sigma = 0$ we set $\mathcal{R}_t \equiv \mathbb{E}(\cdot \mid J_t)$ and recover the usual von Neumann–Morgenstern, state-additive specification.

When $\sigma \neq 0$ the operator $\mathcal{R}_t$ applies an additional risk adjustment *over and above* the one coming from the curvature of u.

Values of $\sigma < 0$ correspond to *more* aversion to risk than the von Neumann–Morgenstern benchmark, which is the case studied throughout the paper.

Note

The exponential-of-utility form in (82.4) originates in the *risk-sensitive control* literature started by Jacobson [1973] and extended by Whittle [1981] and Whittle [1990].

Hansen *et al.* [1999] give it an economic reinterpretation as a *preference for robustness*.

82.2.1 The operator under Gaussian uncertainty

The operator $\mathcal{R}_t$ has a transparent closed form when continuation utility is Gaussian.

Suppose $U_{t+1} \sim N(\mu, s^2)$ conditional on J_t. Using the Gaussian moment generating function $\mathbb{E}[\exp(aU_{t+1})] = \exp(a\mu + \frac{1}{2}a^2 s^2)$ with $a = \sigma/2$,

$$\mathcal{R}_t(U_{t+1}) = \frac{2}{\sigma} \log \mathbb{E}\left[\exp\left(\frac{\sigma U_{t+1}}{2}\right) \middle| J_t\right] = \frac{2}{\sigma}\left(\frac{\sigma}{2}\mu + \frac{\sigma^2}{8}s^2\right) = \mu + \frac{\sigma}{4}s^2. \tag{82.5}$$

For $\sigma < 0$ this is *below* the conditional mean μ: the decision maker evaluates uncertain prospects *pessimistically*, and the penalty grows with the conditional variance s^2.

This certainty equivalent has a revealing decomposition.

The expectation in (82.4) re-weights outcomes by $\exp(\sigma U_{t+1}/2)$; for a Gaussian this **exponential tilting** produces a new normal density with the *same* variance s^2 but a mean shifted to $\mu + \frac{\sigma}{2}s^2$.

The operator value $\mu + \frac{\sigma}{4}s^2$ lies *halfway* between the original mean μ and this worst-case mean: it equals the worst-case expected utility $\mu + \frac{\sigma}{2}s^2$ *plus* the relative-entropy penalty $-\frac{\sigma}{4}s^2$ that restrains the distortion.

Note

The two coefficients describe different objects.

The **worst-case mean** of U_{t+1} shifts by $\frac{\sigma}{2}s^2$, while the **operator value** (certainty equivalent) shifts by $\frac{\sigma}{4}s^2$.

Both are correct; the smaller shift of $\mathcal{R}_t$ reflects the entropy cost the malevolent player pays for the distortion.

A self-contained derivation of (82.5) is requested in Exercise 82.9.1.

Let's visualize both facts, with the certainty equivalent on the left and the worst-case (tilted) density of continuation utility on the right.

```
μ, s = 0.0, 1.0           # conditional mean and std of continuation utility

def R_operator(μ, s, σ):
    "Risk-sensitive operator for a Gaussian U ~ N(μ, s^2)."
    if σ == 0:
        return μ
```

(continues on next page)

(continued from previous page)

```
    return μ + σ * s**2 / 4

σ_vals = np.linspace(-1.5, 0.0, 200)
R_vals = [R_operator(μ, s, sg) for sg in σ_vals]

fig, axes = plt.subplots(1, 2, figsize=(11, 4))

axes[0].plot(σ_vals, R_vals, lw=2)
axes[0].axhline(μ, color='k', ls='--', lw=1, label=r'$E[U_{t+1}]=\mu$')
axes[0].set_xlabel(r'risk-sensitivity $\sigma$')
axes[0].set_ylabel(r'$\mathcal{R}_t(U_{t+1})$')
axes[0].set_title('certainty equivalent')
axes[0].legend()

# Tilted density
grid = np.linspace(μ - 4*s, μ + 4*s, 400)
axes[1].plot(grid, norm.pdf(grid, μ, s), lw=2,
             label=r'reference $N(\mu, s^2)$')
for σ in [-0.5, -1.0]:
    shift = σ * s**2 / 2      # worst-case mean shift of U_{t+1}
    axes[1].plot(grid, norm.pdf(grid, μ + shift, s), lw=2, ls='--',
                 label=f'worst-case, $\\sigma={σ}$')
    axes[1].axvline(μ + σ * s**2 / 4, color='C3', lw=0.8, ls=':')
axes[1].set_xlabel(r'continuation utility $U_{t+1}$')
axes[1].set_ylabel('density')
axes[1].set_title('worst-case density')
axes[1].legend()
fig.tight_layout()
plt.show()
```

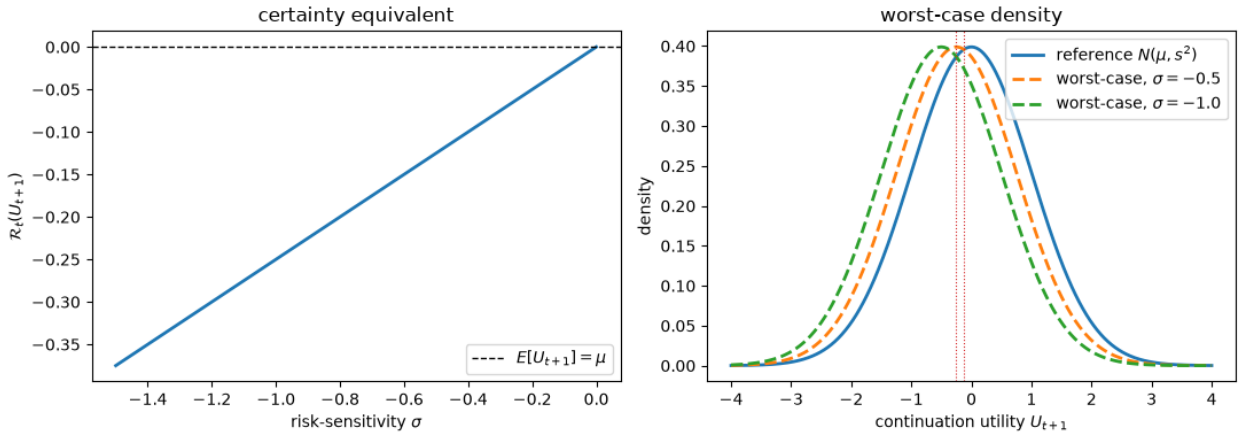


Fig. 82.1: Risk-sensitive operator and worst-case density

The left panel shows the certainty equivalent $\mathcal{R}_t$ sliding below the mean as σ becomes more negative.

The right panel shows the associated worst-case density of continuation utility: a robust agent behaves *as if* U_{t+1} were drawn from a pessimistically re-centered distribution (mean $\mu + \frac{\sigma}{2}s^2$, dashed), while the operator value $\mathcal{R}_t = \mu + \frac{\sigma}{4}s^2$ (dotted) sits halfway between it and the reference mean μ.

82.3 A preference for robustness

The pessimistic tilt in the right panel above is not just an analogy.

Hansen *et al.* [1999] show that the risk-sensitive problem is the value function of a **two-player zero-sum game**.

In this game, one player chooses the control $\{i_t\}$ while a second, malevolent player chooses a distortion $\{v_t\}$ to the conditional mean of the shocks.

The distorted law of motion is

$$x_{t+1} = Ax_t + Bi_t + C(w_{t+1} + v_t). \tag{82.6}$$

The minimizing player would like to push the state in painful directions, but is restrained by a penalty on the size of the distortion.

With $-1/\sigma \geq 0$ acting as a Lagrange multiplier on a constraint that bounds the distortion sequence, the Markov perfect equilibrium has value function

$$\tilde{W}(x) = \inf_v \sup_i \left\{ -i^\top Q i - x^\top R x + \beta \left[-\frac{1}{\sigma} v^\top v + \mathbb{E}\, \tilde{W}(Ax + Bi + C(w + v)) \right] \right\}. \tag{82.7}$$

Because $\sigma < 0$ makes $-1/\sigma > 0$, the term $-\frac{1}{\sigma} v^\top v$ *penalizes* the malevolent player for large distortions.

A smaller $|1/\sigma|$ (more negative σ) means a cheaper distortion budget and hence a larger family of models the agent guards against, reflecting a *stronger* preference for robustness.

This is exactly the max-min expected utility structure of Gilboa and Schmeidler [1989]: the agent's "nominal model" sets $v_t = 0$, but they entertain a whole family of alternatives indexed by $\{v_t\}$ and act against the worst case.

Following Epstein and Wang [1994], the non-uniqueness of the implied probability measures is a form of **Knightian uncertainty**.

The robust and risk-sensitive problems share the same value function matrix Ω and the same decision rule $i_t = -Fx_t$; they differ only in interpretation.

82.4 The permanent income economy

We now specialize to a habit-persistence version of the permanent income model.

A planner orders consumption streams $\{c_t\}$ through a service stream $\{s_t\}$ using the recursion

$$U_t = -(s_t - b_t)^2 + \beta\, \mathcal{R}_t(U_{t+1}), \tag{82.8}$$

where $\{b_t\}$ is an exogenous preference (bliss-point) shock.

Services are produced from consumption via the household technology

$$s_t = (1 + \lambda)c_t - \lambda h_{t-1}, \tag{82.9}$$

$$h_t = \delta_h h_{t-1} + (1 - \delta_h)c_t, \tag{82.10}$$

with $\lambda > 0$ and $\delta_h \in (0, 1)$.

Here h_t is a geometric average of current and past consumption, so (82.9) makes services depend *negatively* on a weighted average of past consumption, which is the **habit persistence**.

There is a linear production technology

$$c_t + i_t = \gamma k_{t-1} + d_t,$$

and capital accumulates as $k_t = \delta_k k_{t-1} + i_t$, where $\{d_t\}$ is an exogenous endowment.

Combining,

$$c_t + k_t = (\delta_k + \gamma) k_{t-1} + d_t, \qquad R \equiv \delta_k + \gamma, \tag{82.11}$$

so R is the gross physical return on capital, which in a decentralized economy equals the gross **risk-free rate**.

The endowment and preference shocks are driven by a common linear state,

$$z_{t+1} = A_{22} z_t + C_2 w_{t+1}, \qquad d_t = U_d z_t, \quad b_t = U_b z_t. \tag{82.12}$$

This whole economy is a special case of the control problem (82.1)–(82.3): stack h_{t-1}, k_{t-1}, and z_t into the state x_t and let the control be $i_t = s_t - b_t$.

82.4.1 The $\sigma = 0$ benchmark and the martingale

To build intuition, set $\sigma = 0$ and impose the permanent income restriction $\beta R = 1$, as in Hall [1978].

The first-order conditions then imply that the marginal utility of consumption services is a **martingale**,

$$\mathbb{E}_t \, \mu_{c,t+1} = \mu_{c,t}, \tag{82.13}$$

and that $\mu_{s,t}$ inherits the representation

$$\mu_{s,t} = \mu_{s,t-1} + v^\top w_t \tag{82.14}$$

for some loading vector v.

Equation (82.13) is the classic statement that, under $\beta R = 1$, consumption responds only to *news*, so it is a random walk.

This is the result that Hall [1978] and Campbell [1987] tested on aggregate U.S. data.

The scalar

$$\theta^2 \equiv v^\top v$$

measures the variance of the innovation to the marginal-utility martingale (82.14).

It will be the one summary statistic of the benchmark economy that we need below.

> **Note**
>
> Under a rational-expectations reading, the benchmark $\sigma = 0$ permanent income model has *no precautionary savings*, as emphasized by Zeldes.
>
> Introducing robustness ($\sigma < 0$) revives a precautionary motive: the consumer guards against worst-case mistakes in the conditional means of shocks.

82.5 Observational equivalence

Here is the paper's first headline result.

Proposition 82.5.1 (Observational Equivalence)

Fix all parameters except β and σ. Suppose $\beta R = 1$. There exists $\underline{\sigma} < 0$ such that the optimal consumption–investment plan for $\sigma = 0$ is *also* optimal for any $\sigma \in (\underline{\sigma}, 0)$, provided the discount factor is lowered to a value $\hat{\beta}(\sigma)$ that varies directly with σ.

In words: as far as the *quantities* $\{c_t, k_t\}$ are concerned, the robust ($\sigma < 0$) permanent income model is indistinguishable from the standard ($\sigma = 0$) one with a smaller discount factor.

Increasing the preference for robustness stimulates a precautionary motive for saving; lowering β makes saving less attractive; along a particular locus the two effects exactly cancel.

The proof is constructive and delivers an explicit locus of observationally equivalent $(\sigma, \hat{\beta})$ pairs. Define

$$\Omega(\beta) = \frac{\beta - 1 + \sigma\theta^2 + \sqrt{(\beta - 1 + \sigma\theta^2)^2 + 4\sigma\theta^2}}{-2\sigma\theta^2}, \tag{82.15}$$

$$\hat{\zeta}(\beta) = 1 + \frac{\theta^2 \sigma\, \Omega(\beta)}{1 - \sigma\theta^2 \Omega(\beta)}. \tag{82.16}$$

The equivalent discount factor $\hat{\beta}$ solves

$$\hat{\beta}\, R\, \hat{\zeta}(\hat{\beta}) = 1. \tag{82.17}$$

The lower bound $\underline{\sigma}$ is the most negative σ for which the square root in (82.15) stays real.

Let's reproduce the locus, a version of Figure 1 in Hansen *et al.* [1999].

```
β_bench = 0.9971          # benchmark discount factor (annual rate ~2.5%)
Rf = 1 / β_bench          # gross risk-free return fixed by β R = 1
θ2 = 0.01                # variance of the marginal-utility innovation, v^T v

def Ω_scalar(β, σ, θ2):
    disc = (β - 1 + σ * θ2)**2 + 4 * σ * θ2
    if disc < 0:
        return np.nan        # below σ-underbar: no real solution
    return (β - 1 + σ * θ2 + np.sqrt(disc)) / (-2 * σ * θ2)

def ζ_hat(β, σ, θ2):
    Ω = Ω_scalar(β, σ, θ2)
    return 1 + (θ2 * σ * Ω) / (1 - σ * θ2 * Ω)

def β_hat(σ, θ2, Rf):
    "Discount factor that makes σ observationally equivalent to σ=0."
    if σ == 0:
        return 1 / Rf
    f = lambda b: b * Rf * ζ_hat(b, σ, θ2) - 1
    return brentq(f, 0.95, 1 / Rf - 1e-12)

σ_vals = np.linspace(-1.2e-4, 0.0, 200)
β_vals = np.array([β_hat(sg, θ2, Rf) for sg in σ_vals])

fig, ax = plt.subplots()
ax.plot(β_vals, σ_vals, lw=2)
ax.set_xlabel(r'discount factor $\hat\beta$')
ax.set_ylabel(r'risk-sensitivity $\sigma$')
ax.axhline(0, color='k', lw=0.8, ls='--')
fig.tight_layout()
plt.show()
```

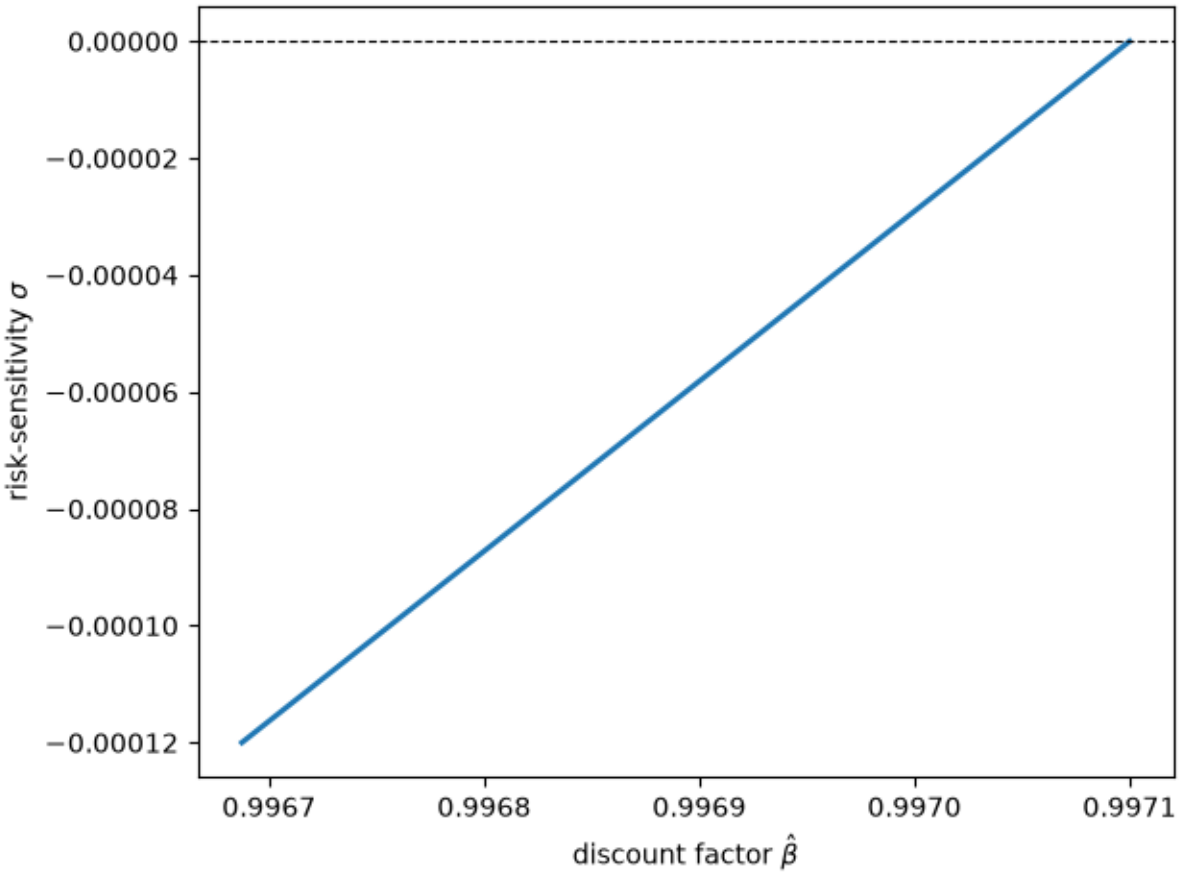


Fig. 82.2: Equivalent $(\sigma, \hat{\beta})$ pairs

Every point on this curve generates *exactly the same* consumption and investment data.

Moving down the curve (more negative σ, i.e. a stronger preference for robustness) requires a *lower* discount factor $\hat{\beta}$ to keep quantities unchanged.

This is why consumption and savings data alone cannot tell us how much the consumer fears model misspecification.

82.6 Estimation

Hansen *et al.* [1999] turns the observational-equivalence result into an empirical strategy.

Because the quantity data cannot pin down σ, the authors first estimate the $\sigma = 0$ version of the model, conditioning the likelihood *only on consumption and investment*, and then use the locus of *Proposition 82.5.1* to trace out the family of $(\sigma, \hat{\beta})$ pairs consistent with those estimates.

Asset prices (the next section) break the tie.

82.6.1 The data and the likelihood

The model is fit to U.S. post-war quarterly data, 1970:I–1996:III.

- *Consumption* is measured as nondurables plus services.
- *Investment* is measured as durables plus gross private investment.

Both series are deflated by the deterministic growth factor 1.0033^t, so the model is fit to *detrended* data.

The likelihood is Gaussian, built recursively (a Kalman filter), with the unobserved part of the initial state estimated using the methods of Hansen and Sargent.

82.6.2 Specification

The preference shock is a constant, $b_t = \mu_b$, fixed at $\mu_b = 32$; recall from the discussion of (82.11) that the *level* of b_t does not affect the decision rules, only prices.

The endowment is the sum of a *persistent* and a *transitory* component, each a second-order autoregression driven by orthogonal shocks,

$$(1 - \phi_1 L)(1 - \phi_2 L)\, d_t^* = c_{d^*}\, w_t^{d^*}, \tag{82.18}$$

$$(1 - \alpha_1 L)(1 - \alpha_2 L)\, \hat{d}_t = c_{\hat{d}}\, w_t^{\hat{d}}, \tag{82.19}$$

with $d_t = \mu_d + d_t^* + \hat{d}_t$.

A likelihood comparison (a gain from AR(1) to AR(2) but not beyond) led the authors to a second-order specification for the transitory component.

The four parameters governing the endogenous dynamics are $(\gamma, \delta_k, \beta, \lambda)$.

The depreciation factor is set to $\delta_k = 0.975$, and the permanent-income restriction $\beta R = 1$ (confirmed by the unrestricted estimates) is imposed with $\beta = 0.9971$, implying a 2.5% annual real interest rate after the growth adjustment.

The maximum-likelihood estimates (with habit persistence) are reproduced below, a version of Table 2 in Hansen *et al.* [1999].

Parameter	Symbol	Estimate
Discount factor	β	0.997
Habit depreciation	δ_h	0.682
Habit weight	λ	2.443
Transitory AR roots	α_1, α_2	0.813, 0.189
Persistent AR roots	ϕ_1, ϕ_2	0.998, 0.704
Endowment mean	μ_d	13.710
Transitory shock scale	$c_{\hat{d}}$	0.155
Persistent shock scale	c_{d^*}	0.108

The single most striking estimate is the autoregressive root $\phi_1 = 0.998$ of the persistent endowment component, a number all but indistinguishable from a unit root.

82.6.3 Impulse responses and the permanent income mechanism

The persistence of a shock is what determines how strongly consumption reacts to it.

Under the permanent income logic (with $\beta R = 1$ and, for transparency, *no* habit), consumption jumps on impact by the **annuity value** of the change in wealth and is a martingale thereafter,

$$\Delta c = \left(1 - \frac{1}{R}\right) \sum_{j \geq 0} R^{-j} \, \psi_j, \tag{82.20}$$

where $\psi_j = \partial d_{t+j} / \partial w_t$ is the endowment's own impulse response.

A near-permanent shock has a large present value and moves consumption a lot; a transitory shock has a small present value and barely moves it.

Let's compute the endowment responses (82.18)–(82.19) and the implied consumption responses.

```
def ar2_irf(r1, r2, c, H):
    "Impulse response of (1 - r1 L)(1 - r2 L) x = c w to a unit w shock."
    ψ = np.zeros(H)
    ψ[0] = c
    if H > 1:
        ψ[1] = (r1 + r2) * ψ[0]
    for j in range(2, H):
        ψ[j] = (r1 + r2) * ψ[j-1] - r1 * r2 * ψ[j-2]
    return ψ

H = 50
disc = Rf**(-np.arange(H))    # Rf = 1/β from the previous cell

ψ_p = ar2_irf(0.998, 0.704, 0.108, H)          # persistent endowment d*
ψ_t = ar2_irf(0.813, 0.189, 0.155, H)          # transitory endowment d_hat

# Flat consumption response
dc_p = (1 - 1/Rf) * np.sum(disc * ψ_p)
dc_t = (1 - 1/Rf) * np.sum(disc * ψ_t)

print(f"persistent shock: consumption responds to "
      f"{100*dc_p/ψ_p[0]:.0f}% of the impact")
print(f"transitory shock: consumption responds to "
      f"{100*dc_t/ψ_t[0]:.0f}% of the impact")
```

(continues on next page)

(continued from previous page)

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4), sharey=True)
for ax, ψ, dc, shock_label in [
        (axes[0], ψ_p, dc_p, 'persistent shock $d^{*}$'),
        (axes[1], ψ_t, dc_t, 'transitory shock $\\hat d$')]:
    ax.plot(ψ, lw=2, label=fr'{shock_label}: endowment response $\psi_j$')
    ax.axhline(dc, color='C3', ls='--', lw=2,
               label=r'consumption response $\Delta c$')
    ax.set_title(shock_label)
    ax.set_xlabel('horizon $j$ (quarters)')
    ax.legend()
axes[0].set_ylabel('response')
fig.tight_layout()
plt.show()
```

```
persistent shock: consumption responds to 41% of the impact
transitory shock: consumption responds to 2% of the impact
```

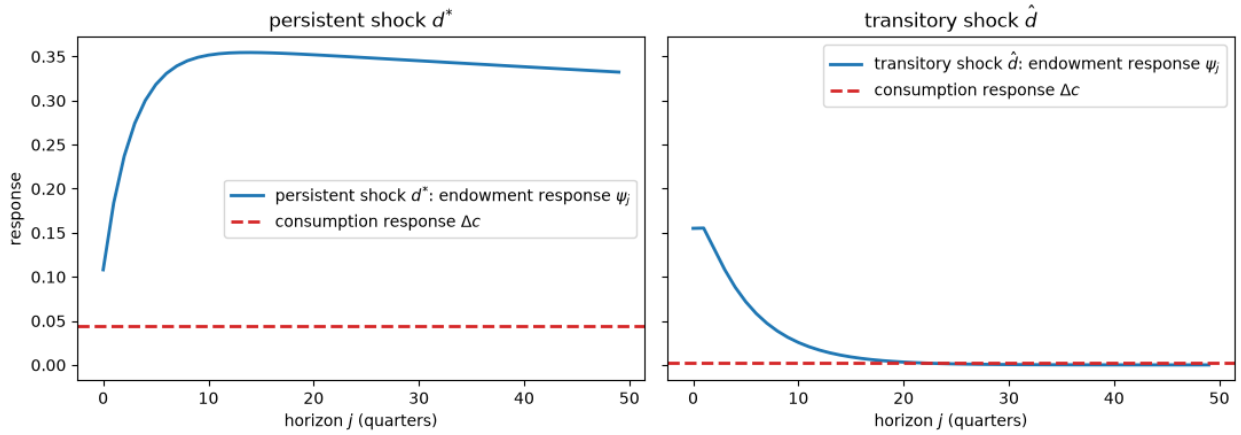


Fig. 82.3: Endowment and consumption impulse responses

The contrast is the heart of the permanent income hypothesis tested by Hall [1978] and Campbell [1987].

Consumption tracks a large fraction of the persistent shock, whose near-unit root makes it almost permanent income, but only a sliver of the transitory shock, the bulk of which is saved and shows up as investment.

Note

With habit persistence ($\lambda > 0$) the consumption responses are no longer flat: they become hump-shaped, because services rather than consumption obey the martingale logic.

The estimated $\lambda = 2.443$ and $\delta_h = 0.682$ imply economically important habit effects, and a likelihood-ratio comparison strongly rejects $\lambda = 0$.

Hansen *et al.* [1999] compare these magnitudes with the habit estimates in the time-nonseparable preference literature.

82.7 Asset pricing and the market price of risk

Hansen *et al.* [1999] shows how the observationally-equivalent pairs that look identical in quantity data have *different* implications for asset prices.

82.7.1 Decentralization

Following Lucas [1978], we regard the robust (or risk-sensitive) planning solution as the allocation of a competitive economy populated by a large number of identical agents who trade securities.

Equilibrium prices are the **shadow prices** that leave each agent content to consume the planner's allocation, treating it as an endowment process.

The equilibrium law of motion for the state is

$$x_{t+1} = A^0 x_t + C w_{t+1}, \tag{82.21}$$

and the value function at the optimum is $U_t^e = x_t^\top \Omega x_t + \rho$.

To support the robust allocation, prices must be computed using the *same* pessimistic, distorted beliefs that rationalize the planner's choices.

This is where risk-sensitivity ($\sigma < 0$) leaves its fingerprint on prices even though, by observational equivalence, it leaves no trace in quantities.

82.7.2 The twisting operator and distorted beliefs

Pricing a claim to next period's *utility* is trivial under the von Neumann–Morgenstern specification but nontrivial under risk-sensitivity.

The key object is the **twisting operator**

$$\mathcal{T}_t U_{t+1} \equiv \frac{\mathbb{E}(V_{t+1} U_{t+1} \mid J_t)}{\mathbb{E}(V_{t+1} \mid J_t)}, \qquad V_{t+1} \equiv \exp\left(\frac{\sigma U_{t+1}^e}{2}\right), \tag{82.22}$$

which re-weights outcomes by the exponential of equilibrium continuation utility.

It satisfies the subgradient inequality

$$\mathcal{R}_t(U_{t+1}) - \mathcal{R}_t(U_{t+1}^e) \leq \mathcal{T}_t U_{t+1} - \mathcal{T}_t U_{t+1}^e, \tag{82.23}$$

so $\mathcal{T}_t$ behaves like a *distorted* conditional expectation, exactly the change of measure used to price derivative claims in Epstein and Wang [1994].

Concretely, $\mathcal{T}_t$ is the ordinary conditional expectation under a **distorted transition law**

$$x_{t+1} = \hat{A} x_t + \hat{C} w_{t+1}, \qquad \hat{C}\hat{C}^\top = C(I - \sigma C^\top \Omega C)^{-1} C^\top, \tag{82.24}$$

with $\hat{A}$ given by (82.30)-style risk corrections.

Because $\sigma < 0$ and Ω is negative semidefinite, $(I - \sigma C^\top \Omega C)^{-1}$ exceeds the identity: the pricing measure assigns a *pessimistically shifted conditional mean and* an *inflated conditional variance* to next period's state.

These two distortions are precisely what generate risk premia.

82.7.3 Multi-period claims and the one-period stochastic discount factor

Prices of streams are built by iterating the operator.

Define $\mathcal{S}_{t,\tau} = \mathcal{T}_t \mathcal{T}_{t+1} \cdots \mathcal{T}_{t+\tau-1}$. The time-$t$ price of a claim to the consumption stream $\{c_{t+\tau}\}$ is then a discounted sum of twisted marginal utilities, and the one-period security with payoff p_{t+1} is priced as

$$q_t = \mathcal{T}_t \left\{ \beta \frac{\mathcal{M}^c_{t+1}}{\mathcal{M}^c_t} \, p_{t+1} \right\} = \mathbb{E}\left(m_{t+1,t} \, p_{t+1} \mid J_t \right), \tag{82.25}$$

where $\mathcal{M}^c_t$ is the marginal utility of consumption and $m_{t+1,t}$ is the one-period **stochastic discount factor** (intertemporal marginal rate of substitution).

Under risk-sensitivity, $m_{t+1,t}$ factors into two pieces,

$$m_{t+1,t} = m^f_{t+1,t} \, m^r_{t+1,t}, \tag{82.26}$$

where

$$m^f_{t+1,t} = \beta \frac{\mathcal{M}^c_{t+1}}{\mathcal{M}^c_t}$$

is the *familiar* intertemporal marginal rate of substitution (the only term present when $\sigma = 0$), and

$$m^r_{t+1,t} = \frac{\exp(\sigma U^e_{t+1}/2)}{\mathbb{E}[\exp(\sigma U^e_{t+1}/2) \mid J_t]} \tag{82.27}$$

is a multiplicative adjustment with conditional mean one.

The factor (82.27) is the source of the extra risk premia.

82.7.4 The market price of risk

Under the *robustness* interpretation, the same multiplicative factor equals a **likelihood ratio** between the worst-case and reference shock densities,

$$m^u_{t+1,t} = \frac{\exp[-(w_{t+1} - \hat{v}_t)^\top (w_{t+1} - \hat{v}_t)/2]}{\exp(-w^\top_{t+1} w_{t+1}/2)}, \tag{82.28}$$

where $\hat{v}_t$ is the worst-case conditional-mean distortion chosen by the malevolent player.

A direct calculation gives

$$\mathbb{E}_t(m^u_{t+1,t}) = 1, \qquad \mathbb{E}_t[(m^u_{t+1,t})^2] = \exp(\hat{v}_t^\top \hat{v}_t),$$

so that, for small distortions,

$$\mathrm{std}_t(m^u_{t+1,t}) = \left[\exp(\hat{v}_t^\top \hat{v}_t) - 1 \right]^{1/2} \approx |\hat{v}_t|. \tag{82.29}$$

The **market price of risk** is the maximal Sharpe ratio attainable, equal to $\mathrm{std}_t(m_{t+1,t})/\mathbb{E}_t(m_{t+1,t})$ along the efficient frontier (the Hansen and Jagannathan [1991] bound).

It is therefore approximately equal to the *magnitude of the worst-case distortion* $|\hat{v}_t|$.

This is the paper's punchline: a conditional-mean misspecification of $x\%$ of a unit-norm direction raises the market price of risk by roughly $x/100$.

A *small*, statistically-hard-to-detect doubt about the model can generate the *large* price of risk seen in the data.

Let's check the key identity (82.29) by Monte Carlo.

```
rng = np.random.default_rng(12345)

def mpr_check(v_hat, n=2_000_000):
    """
    Simulate the worst-case likelihood ratio m^u and compare its
    conditional standard deviation to |v_hat|.
    """
    k = len(v_hat)
    w = rng.standard_normal((n, k))
    # log likelihood ratio of N(v_hat, I) relative to N(0, I)
    log_mu = w @ v_hat - 0.5 * v_hat @ v_hat
    mu = np.exp(log_mu)
    return mu.mean(), mu.std(), np.linalg.norm(v_hat)

print(f"{'|v_hat|':>10}{'E[m^u]':>12}{'std(m^u)':>12}{'approx |v_hat|':>16}")
for scale in [0.05, 0.10, 0.20]:
    v_hat = np.array([scale, 0.0])      # distortion in one direction
    mean, std, norm_v = mpr_check(v_hat)
    print(f"{norm_v:10.3f}{mean:12.4f}{std:12.4f}{norm_v:16.3f}")
```

```
   |v_hat|      E[m^u]    std(m^u)  approx |v_hat|
     0.050      1.0000      0.0500           0.050
     0.100      0.9999      0.1003           0.100
     0.200      0.9997      0.2022           0.200
```

The simulated conditional mean of m^u is one, and its conditional standard deviation tracks $|\hat{v}_t|$ closely, confirming (82.29).

A 10% distortion delivers a market price of risk near 0.10.

82.8 A risk-sensitive regulator

To see the robust decision rules and worst-case shocks concretely, we solve the recursive risk-sensitive control problem (82.1)–(82.3) directly.

Guess a value function $W(x) = x^\top \Omega x + \rho$ with Ω negative semidefinite. The risk-sensitive operator (82.4) acting on this quadratic form introduces the **risk adjustment**

$$\mathcal{D}(\Omega) = \Omega + \sigma \Omega C (I - \sigma C^\top \Omega C)^{-1} C^\top \Omega, \tag{82.30}$$

so that, replacing Ω by $\mathcal{D}(\Omega)$, the Bellman equation becomes an ordinary linear-quadratic one.

Iterating

$$F = (Q - \beta B^\top \mathcal{D} B)^{-1} (N - \beta B^\top \mathcal{D} A),$$

$$\Omega \leftarrow -R - F^\top Q F + (F^\top N + N^\top F) + \beta (A - BF)^\top \mathcal{D} (A - BF),$$

to a fixed point yields the optimal rule $i_t = -F x_t$.

The worst-case mean distortion is then linear in the state, $\hat{v}_t = G x_t$, with

$$G = \sigma (I - \sigma C^\top \Omega C)^{-1} C^\top \Omega (A - BF). \tag{82.31}$$

When $\sigma = 0$ we have $\mathcal{D}(\Omega) = \Omega$ and $G = 0$, recovering the standard regulator.

```
def solve_rslq(A, B, C, Q, R, β, σ, N=None,
               tol=1e-12, max_iter=100_000):
    """
    Solve the recursive risk-sensitive LQ problem

        U_t = -(x^T R x + i^T Q i + 2 i^T N x) + β R_t(U_{t+1})
        x_{t+1} = A x_t + B i_t + C w_{t+1}

    Returns the feedback rule F (i = -F x), the value matrix Ω,
    the closed-loop matrix A - B F, and the worst-case loading G (v = G x).
    """
    A, B, C, Q, R = map(np.atleast_2d, (A, B, C, Q, R))
    n, kw = A.shape[0], C.shape[1]
    if N is None:
        N = np.zeros((B.shape[1], n))
    Ω = -np.eye(n)                # negative-definite start
    Iw = np.eye(kw)
    for it in range(max_iter):
        M = Iw - σ * C.T @ Ω @ C
        D = Ω + σ * Ω @ C @ np.linalg.solve(M, C.T @ Ω)
        F = np.linalg.solve(Q - β * B.T @ D @ B, N - β * B.T @ D @ A)
        Acl = A - B @ F
        Ω_new = (-R - F.T @ Q @ F + (F.T @ N + N.T @ F)
                 + β * Acl.T @ D @ Acl)
        if np.max(np.abs(Ω_new - Ω)) < tol:
            Ω = Ω_new
            break
        Ω = Ω_new
    M = Iw - σ * C.T @ Ω @ C
    G = σ * np.linalg.solve(M, C.T @ Ω @ (A - B @ F))
    return F, Ω, A - B @ F, G
```

We first verify that at $\sigma = 0$ our solver reproduces QuantEcon's ordinary LQ regulator.

```
# Stable regulator
A = np.array([[0.9, 0.0],
              [0.0, 0.8]])
B = np.array([[1.0],
              [0.0]])
C = np.array([[0.3],
              [0.2]])
Q = np.array([[1.0]])
R = np.eye(2)
β = 0.95

# QuantEcon ordinary LQ
lq = qe.LQ(Q, R, A, B, C=C, beta=β)
P, F_qe, d = lq.stationary_values()

# Sigma-zero check
F0, Ω0, Acl0, G0 = solve_rslq(A, B, C, Q, R, β, σ=0.0)

print("QuantEcon LQ feedback rule F :", F_qe.flatten())
print("solve_rslq  feedback rule F  :", F0.flatten())
print("max |difference|             :", np.max(np.abs(F0 - F_qe)))
```

```
QuantEcon LQ feedback rule F : [0.52479699 0.        ]
```

(continues on next page)

(continued from previous page)

```
solve_rslq  feedback rule F  : [0.52479699 0.        ]
max |difference|             : 1.1102230246251565e-16
```

The two agree to machine precision.

Now we crank up the preference for robustness and inspect how the control rule and the worst-case shock respond.

```
σ_vals = [0.0, -0.3, -0.6]

print(f"{'σ':>7}{'F[0]':>10}{'F[1]':>10}{'G[0]':>10}{'G[1]':>10}")
for σ in σ_vals:
    F, Q, Acl, G = solve_rslq(A, B, C, Q, R, β, σ)
    print(f"{σ:7.2f}{F[0,0]:10.4f}{F[0,1]:10.4f}"
          f"{G[0,0]:10.4f}{G[0,1]:10.4f}")
```

```
      σ      F[0]      F[1]      G[0]      G[1]
   0.00    0.5248    0.0000   -0.0000   -0.0000
  -0.30    0.5357    0.0340    0.0532    0.1370
  -0.60    0.5490    0.0803    0.1154    0.3169
```

As σ becomes more negative:

- the feedback gain F grows, since the robust agent reacts *more aggressively* to the state, fearing the worst-case shock will amplify deviations;
- the worst-case loading G moves away from zero, as the malevolent player pushes the shock in the direction that hurts most.

Finally, let's see the worst-case distortion in action by simulating the controlled state under the *reference* model while displaying the conditional mean distortion $\hat{v}_t = Gx_t$ the robust agent is guarding against.

```
def simulate(A, B, C, F, G, T=60, seed=0):
    rng = np.random.default_rng(seed)
    n, kw = A.shape[0], C.shape[1]
    x = np.zeros((T + 1, n))
    v = np.zeros((T, kw))
    x[0] = np.array([1.0, 1.0])          # initial deviation
    for t in range(T):
        v[t] = (G @ x[t]).flatten()
        w = rng.standard_normal(kw)
        x[t + 1] = A @ x[t] + (B @ (-F @ x[t])).flatten() + (C @ w).flatten()
    return x, v

fig, axes = plt.subplots(1, 2, figsize=(11, 4))

for σ, color in zip([0.0, -0.6], ['C0', 'C3']):
    F, Q, Acl, G = solve_rslq(A, B, C, Q, R, β, σ)
    x, v = simulate(A, B, C, F, G, seed=42)
    axes[0].plot(x[:, 0], color=color, lw=2, label=f'$\\sigma={σ}$')
    axes[1].plot(v[:, 0], color=color, lw=2, label=f'$\\sigma={σ}$')

axes[0].set_xlabel('$t$')
axes[0].set_ylabel(r'controlled state $x_{1,t}$')
axes[0].axhline(0, color='k', lw=0.8, ls='--')
axes[0].set_title('controlled state')
axes[0].legend()
```

(continues on next page)

(continued from previous page)

```
axes[1].set_xlabel('$t$')
axes[1].set_ylabel(r'worst-case mean distortion $\hat v_{1,t}$')
axes[1].axhline(0, color='k', lw=0.8, ls='--')
axes[1].set_title('mean distortion')
axes[1].legend()
fig.tight_layout()
plt.show()
```

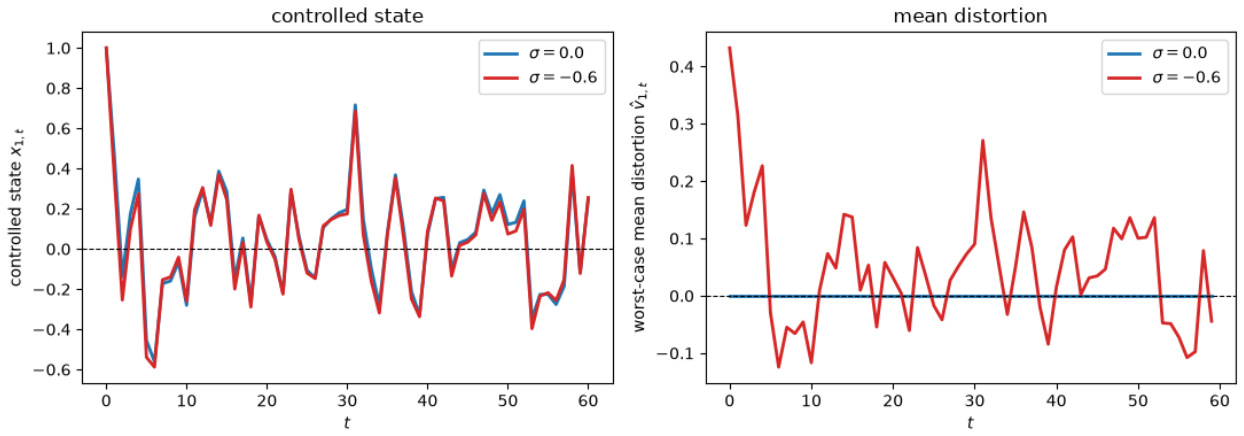


Fig. 82.4: Controlled state and distortion

For $\sigma = 0$ the distortion is identically zero, since the agent fully trusts the model.

For $\sigma < 0$ the robust agent's decisions are shaped by a nonzero worst-case distortion that feeds back on the state, exactly the mechanism that, through (82.29), inflates the market price of risk.

82.9 Exercises

Exercise 82.9.1

The Gaussian formula (82.5) says that for $U_{t+1} \sim N(\mu, s^2)$,

$$\mathcal{R}_t(U_{t+1}) = \mu + \frac{\sigma}{4} s^2.$$

Derive this result directly from the definition (82.4).

Hint: use the moment generating function of a normal random variable, $\mathbb{E}[\exp(aU)] = \exp(a\mu + a^2 s^2/2)$.

Solution

Here is one solution:

Start from the definition with $a = \sigma/2$:

$$\mathcal{R}_t(U_{t+1}) = \frac{2}{\sigma}\log\mathbb{E}\Big[\exp\Big(\frac{\sigma}{2}U_{t+1}\Big)\Big].$$

Since $U_{t+1} \sim N(\mu, s^2)$, the moment generating function gives

$$\mathbb{E}\Big[\exp\Big(\frac{\sigma}{2}U_{t+1}\Big)\Big] = \exp\Big(\frac{\sigma}{2}\mu + \frac{1}{2}\frac{\sigma^2}{4}s^2\Big).$$

Taking logs and multiplying by $2/\sigma$,

$$\mathcal{R}_t(U_{t+1}) = \frac{2}{\sigma}\left(\frac{\sigma}{2}\mu + \frac{\sigma^2}{8}s^2\right) = \mu + \frac{\sigma}{4}s^2.$$

Letting $\sigma \to 0$ recovers $\mathcal{R}_t(U_{t+1}) = \mu = \mathbb{E}_t U_{t+1}$, as expected.

Exercise 82.9.2

The observational-equivalence locus has a left endpoint $\underline{\sigma}$: the most negative σ for which (82.15) has a real solution.

Using the code from the lecture, find $\underline{\sigma}$ numerically for `theta2 = 0.01` and for `theta2 = 0.02`, and explain why a larger θ^2 shrinks the admissible range of σ.

Hint: the square root is real when the discriminant $(\beta - 1 + \sigma\theta^2)^2 + 4\sigma\theta^2 \geq 0$, evaluated at the relevant $\hat{\beta}$.

Solution

Here is one solution:

The boundary $\underline{\sigma}$ is the most negative σ at which a valid $\hat{\beta}$ can still be found.

We scan σ downward and stop when `beta_hat` can no longer return a real solution.

```
def σ_underbar(θ2, Rf, grid=np.linspace(-1e-6, -5e-4, 5000)):
    last_ok = 0.0
    for sg in grid:
        try:
            b = β_hat(sg, θ2, Rf)
            if np.isnan(Ω_scalar(b, sg, θ2)):
                break
            last_ok = sg
        except ValueError:
            break
    return last_ok

for θ2 in [0.01, 0.02]:
    sb = σ_underbar(θ2, Rf)
    print(f"θ2 = {θ2}:  σ_underbar ≈ {sb:.3e}")
```

```
θ2 = 0.01:  σ_underbar ≈ -2.105e-04
θ2 = 0.02:  σ_underbar ≈ -1.052e-04
```

A larger θ^2 means the marginal-utility martingale (82.14) carries a bigger innovation variance, so each unit of $|\sigma|$ generates a larger risk adjustment.

The discriminant in (82.15), which contains the term $4\sigma\theta^2 < 0$, turns negative at a *smaller* $|\sigma|$.

Hence the admissible range $(\underline{\sigma}, 0)$ shrinks as θ^2 grows.

Exercise 82.9.3

The market-price-of-risk approximation (82.29) states that $\text{std}_t(m^u_{t+1,t}) = [\exp(\hat{v}_t^\top \hat{v}_t) - 1]^{1/2} \approx |\hat{v}_t|$.

For the risk-sensitive regulator solved in the lecture (with A, B, C, Q, R, β as given there), compute the *exact* market price of risk $[\exp(\hat{v}_t^\top \hat{v}_t) - 1]^{1/2}$ as a function of σ, evaluated at the state $x = (1, 1)^\top$, and plot it together with the linear approximation $|\hat{v}_t|$.

Comment on the range of σ over which the approximation is accurate.

Solution

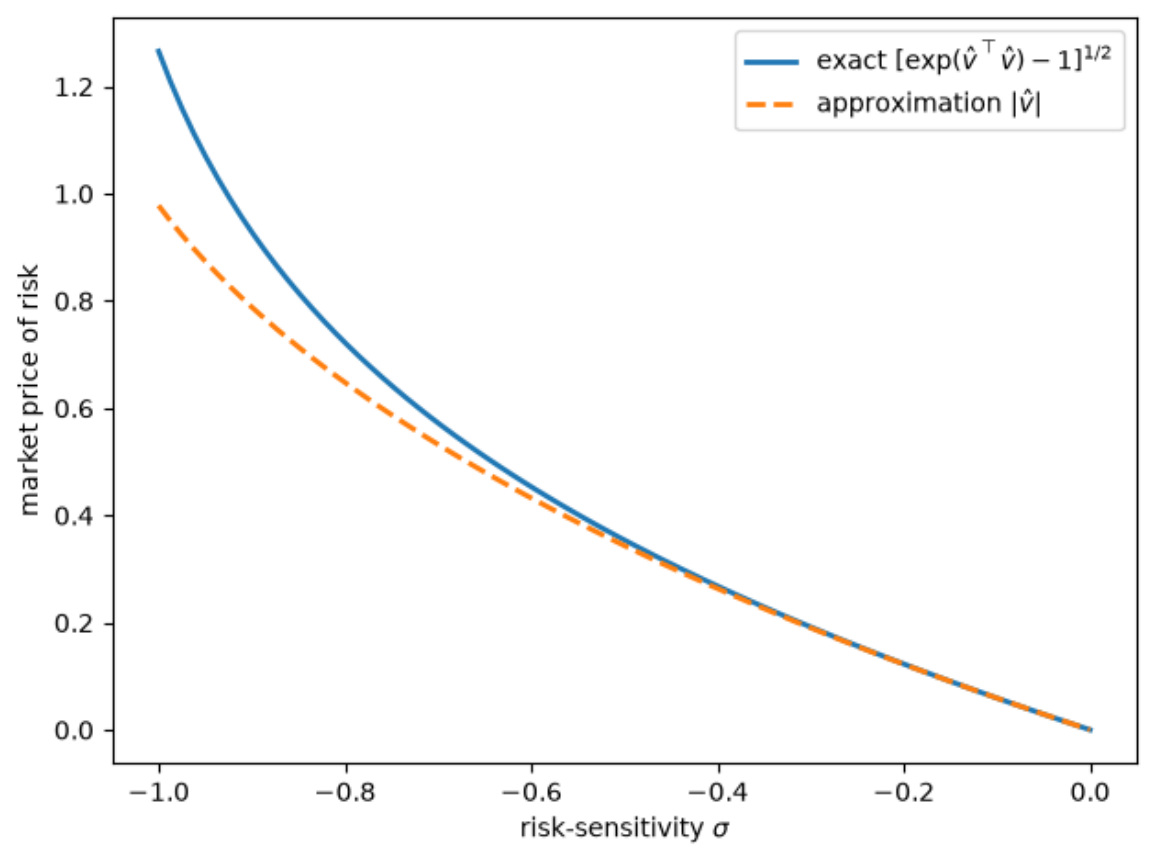
Here is one solution:

We solve the regulator for a grid of σ values, evaluate $\hat{v}_t = Gx$ at $x = (1, 1)^\top$, and compare the exact and approximate market prices of risk.

```
x_eval = np.array([1.0, 1.0])
σ_grid = np.linspace(-1.0, 0.0, 80)

exact, approx = [], []
for σ in σ_grid:
    F, Ω, Acl, G = solve_rslq(A, B, C, Q, R, β, σ)
    v_hat = (G @ x_eval).flatten()
    nv2 = v_hat @ v_hat
    exact.append(np.sqrt(np.exp(nv2) - 1))
    approx.append(np.sqrt(nv2))            # |v_hat|

fig, ax = plt.subplots()
ax.plot(σ_grid, exact, lw=2,
        label=r'exact $[\exp(\hat v^\top \hat v)-1]^{1/2}$')
ax.plot(σ_grid, approx, lw=2, ls='--', label=r'approximation $|\hat v|$')
ax.set_xlabel(r'risk-sensitivity $\sigma$')
ax.set_ylabel('market price of risk')
ax.legend()
fig.tight_layout()
plt.show()
```



The two curves are nearly indistinguishable for small $|\hat{v}_t|$ (i.e. σ close to zero) because $\exp(z) - 1 \approx z$ when z is small.

They separate only when the preference for robustness, and hence the worst-case distortion, becomes large.

This is precisely the regime Hansen *et al.* [1999] emphasize: *small*, hard-to-detect distortions map almost linearly into the market price of risk.

82.10 Related lectures

This lecture connects to several others in the QuantEcon collection.

The underlying consumption-smoothing economics is developed in *The Permanent Income Model* and *Permanent Income II: LQ Techniques*, and the linear-quadratic control machinery is laid out in *LQ Control: Foundations*.

The reinterpretation of the market price of risk as a *price of model uncertainty*, calibrated through detection-error probabilities, is the subject of Barillas *et al.* [2009]; the broader semigroup treatment of robustness, pricing, and model detection is in Anderson *et al.* [2003].

Both build directly on Hansen *et al.* [1999].

CHAPTER

EIGHTYTHREE

CERTAINTY EQUIVALENCE

Contents

- *Certainty Equivalence*
  - *Overview*
  - *A central problem of empirical economics*
  - *A formal setup*
  - *Estimated rules are not enough*
  - *An optimization problem*
  - *A linear-quadratic DP problem and certainty equivalence*
  - *A trouble with ad hoc expectations*
  - *Exercises*
  - *Concluding remarks*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

```
import numpy as np
import matplotlib.pyplot as plt
from quantecon import LQ
```

83.1 Overview

Simon [1956] and Theil [1957] established a celebrated *certainty equivalence* (CE) property for linear-quadratic (LQ) dynamic programming problems.

Their result justifies a convenient two-step algorithm:

1. **Optimize** under perfect foresight (treat future exogenous variables as known).
2. **Forecast** – substitute optimal forecasts for the unknown future values.

The striking insight is that these two steps are completely separable.

The decision rule that emerges from step 1 is *identical* to the decision rule for the original stochastic problem once optimal forecasts are substituted in step 2.

The decision rule does not depend on the variance of the shocks, but the *level* of the optimal value function *does*.

After describing the structure of the certainty equivalence property in detail, this lecture describes its role in rational expectations modeling.

We do so by drawing heavily on the introduction to Lucas and Sargent [1981].

In addition to learning the certainty equivalence principle, this lecture describes troubles with pre-rational expectations econometric policy evaluation procedures described by Lucas [1976].

Note

That volume collected early papers on rational expectations modeling and econometrics.

83.2 A central problem of empirical economics

To set the stage, Lucas and Sargent [1981] stated the central question for empirical economics that had been posed by Leonid Hurwicz ([Hurwicz, 1962]):

- Given observations on an agent's behavior in a particular economic environment, what can we infer about how that behavior *would have differed* had the environment been altered?

Note

Hurwicz formulates a notion of 'causality' as a context-specific concept that he casts in terms of a well posed decision problem.

This is the problem of policy-invariant structural inference in the following setting.

- Observations emerged under one environment or 'regime'
- We want to predict behavior under another 'regime'
- Unless we understand *why* agents behaved as they did in the historical regime, i.e., their purposes, we can't predict their behavior under the constraints they face in the new regime.

To confront the problem that Hurwicz had posed, Lucas and Sargent [1981] formulated the following decision framework.

83.3 A formal setup

Consider a single decision maker whose situation at date t is fully described by two state variables (x_t, z_t).

The environment $z_t \in S_1$ is selected by "nature" and evolves exogenously according to

$$z_{t+1} = f(z_t, \epsilon_{t+1}), \tag{83.1}$$

where the innovations $\{\epsilon_t\}$ are IID draws from a fixed CDF $\Phi(\cdot) : \mathcal{E} \to [0, 1]$.

The function $f : S_1 \times \mathcal{E} \to S_1$ is called the **decision maker's environment**.

The endogenous state $x_t \in S_2$ is under partial control of the agent.

Each period the agent selects an action $u_t \in U$.

A fixed technology $g : S_1 \times S_2 \times U \to S_2$ governs the transition

$$x_{t+1} = g(z_t,\, x_t,\, u_t). \tag{83.2}$$

The decision rule $h : S_1 \times S_2 \to U$ maps the agent's current situation into an action:

$$u_t = h(z_t,\, x_t). \tag{83.3}$$

The econometrician observes (some or all of) the process $\{z_t, x_t, u_t\}$, the joint motion of which is determined by (83.1), (83.2), and (83.3).

83.4 Estimated rules are not enough

Suppose we have estimated f, g, and h from a long time series generated under a fixed environment f_0.

This gives us $h_0 = T(f_0)$, where T is the (unknown) functional mapping environments into optimal decision rules.

But this single estimate, however precise, *reveals nothing* about how $T(f)$ varies with f.

Policy evaluation requires knowledge of the entire map $f \mapsto T(f)$.

Under an environment change $f_0 \to f_1$, agents will in general revise their decision rules $h_0 \to h_1 = T(f_1)$, rendering the estimated rule h_0 invalid for forecasting behavior under f_1.

Lucas [1976] and the introduction to Lucas and Sargent [1981] conclude that the only nonexperimental path forward is to recover the **return function** V from which h is derived as the solution to an optimization problem, and then re-solve that problem under the counterfactual environment f_1.

83.5 An optimization problem

Assume the agent selects h to maximize the expected discounted sum of current-period returns $V : S_1 \times S_2 \times U \to \mathbb{R}$:

$$E_0\left\{\sum_{t=0}^{\infty} \beta^t\, V(z_t,\, x_t,\, u_t)\right\}, \qquad 0 < \beta < 1, \tag{83.4}$$

given initial conditions (z_0, x_0), the environment f, and the technology g.

Here $E_0\{\cdot\}$ denotes expectation conditional on (z_0, x_0) with respect to the distribution of $\{z_1, z_2, \ldots\}$ induced by (83.1).

In principle, knowledge of V (together with g and f) allows one to compute $h = T(f)$ theoretically and hence to trace out $T(f)$ for any counterfactual f.

The essential question is whether V can itself be recovered from observations on $\{f, g, h\}$.

> **Note**
>
> The decision rule is in general a functional $h = T(f, g, V)$. The dependence on g and V is suppressed in the main text but made explicit when needed.

83.6 A linear-quadratic DP problem and certainty equivalence

Progress beyond the level of generality of the previous section requires restricting the primitives.

A productive restriction, exploited in the papers collected in Lucas and Sargent [1981], imposes *quadratic* V and *linear* g, which forces h to be linear.

As part of its computational tractability, this specialization delivers a striking structural result:

- the **certainty equivalence** theorem of Simon [1956] and Theil [1957].

83.6.1 Decomposition of h

Under quadratic V and linear g, the optimal decision rule h decomposes into two components applied in sequence.

Step 1 – Forecasting. Define the infinite sequence of optimal point forecasts of all current and future states of nature:

$$\tilde{z}_t \;=\; (z_t,\;\; {}_{t+1}z_t^e,\;\; {}_{t+2}z_t^e,\;\ldots) \;\in\; S_1^\infty, \tag{83.5}$$

where ${}_{t+j}z_t^e$ denotes the least-mean-squared-error forecast of z_{t+j} formed at time t.

The optimal forecast sequence is a (generally nonlinear) function of the current state:

$$\tilde{z}_t = h_2(z_t). \tag{83.6}$$

The function $h_2 : S_1 \to S_1^\infty$ depends entirely on the environment (f, Φ) and is obtained as the solution to a **pure forecasting problem**, with no reference to preferences or technology.

Step 2 – Optimization. Given the forecast sequence $\tilde{z}_t$, the optimal action is a **linear** function of $\tilde{z}_t$ and x_t:

$$u_t = h_1(\tilde{z}_t,\, x_t). \tag{83.7}$$

The function $h_1 : S_1^\infty \times S_2 \to U$ depends entirely on preferences (V) and technology (g) but **not** on the stochastic environment (f, Φ).

The ultimate decision rule is therefore the **composite**:

$$\boxed{h(z_t, x_t) \;=\; h_1[h_2(z_t),\; x_t].} \tag{83.8}$$

83.6.2 The separation principle

(83.8) embodies a clean **separation** of the two sources of dependence in h:

Component	Depends on	Independent of
h_1 (optimization)	V, g	f, Φ
h_2 (forecasting)	f, Φ	V, g

Since policy analysis concerns changes in f, and since h_1 is invariant to f, the policy analyst need only re-solve the forecasting problem $h_2 = S(f)$ under the new environment, keeping h_1 fixed.

The relationship of original interest, $h = T(f)$, then follows directly from (83.8).

83.6.3 Certainty equivalence and perfect foresight

The name "certainty equivalence" reflects a further implication of the LQ structure: the function h_1 can be derived as if the agent **knew the future path** $z_{t+1}, z_{t+2}, \ldots$ **with certainty** – i.e., by solving the deterministic problem in which $\tilde{z}_t$ is treated as the realized path rather than a forecast.

Randomness of the environment affects actions only through the forecast $\tilde{z}_t$; conditional on $\tilde{z}_t$, the optimization problem is deterministic.

This means the LQ problem decouples into:

- **Dynamic optimization under perfect foresight** – solve for h_1 from (V, g) by treating $\tilde{z}_t$ as known, yielding a standard deterministic LQ regulator problem independent of the environment (f, Φ).
- **Optimal linear prediction** – solve for $h_2 = S(f)$ from (f, Φ) using least-squares forecasting theory, which reduces to a standard Kalman/Wiener prediction formula when f is itself linear.

83.6.4 Cross-equation restrictions

A hallmark of the rational expectations hypothesis as it appears in this framework is that it ties together what would otherwise be free parameters in different equations.

The requirement that $\tilde{z}_t = h_2(z_t) = S(f)(z_t)$ – i.e., that agents' forecasts be *optimal* with respect to the *actual* law of motion f – imposes **cross-equation restrictions** between the parameters of the forecasting rule h_2 and the parameters of the environment f.

These restrictions, rather than any conditions on distributed lags within a single equation, are the operative empirical content of rational expectations.

Note

This is the message of Lucas [1976] and Sargent [1981].

The following code verifies the CE principle numerically.

We consider a simple scalar LQ problem:

$$y_{t+1} = a\, y_t + b\, u_t + \sigma\, \varepsilon_{t+1}, \qquad r(y_t, u_t) = -(q\, y_t^2 + r\, u_t^2)$$

and vary the noise standard deviation σ across a wide range.

The CE theorem predicts that:

- the **policy gain** F (the coefficient in $u_t = -F y_t$) is independent of σ, and
- the **value constant** d (the additive term in $V(y) = -y^\top P y - d$) grows with σ.

```
a, b_coeff = 0.9, 1.0
q, r = 1.0, 1.0
β = 0.95

A = np.array([[a]])
B = np.array([[b_coeff]])
R_mat = np.array([[q]])          # state cost
Q_mat = np.array([[r]])          # control cost

σ_vals = np.linspace(0.0, 3.0, 80)
```

(continues on next page)

(continued from previous page)

```
F_vals, d_vals = [], []

for σ in σ_vals:
    C = np.array([[σ]])
    lq = LQ(Q_mat, R_mat, A, B, C=C, beta=β)
    P, F, d = lq.stationary_values()
    F_vals.append(float(F[0, 0]))
    d_vals.append(float(d))

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].plot(σ_vals, F_vals, lw=2)
axes[0].set_xlabel('noise level $\\sigma$')
axes[0].set_ylabel('policy gain $F$')
axes[0].set_title('policy gain')
axes[0].set_ylim(0, 2 * max(F_vals) + 0.1)

axes[1].plot(σ_vals, d_vals, lw=2, color='C1')
axes[1].set_xlabel('noise level $\\sigma$')
axes[1].set_ylabel('value constant $d$')
axes[1].set_title('value constant')
fig.tight_layout()
plt.show()
```

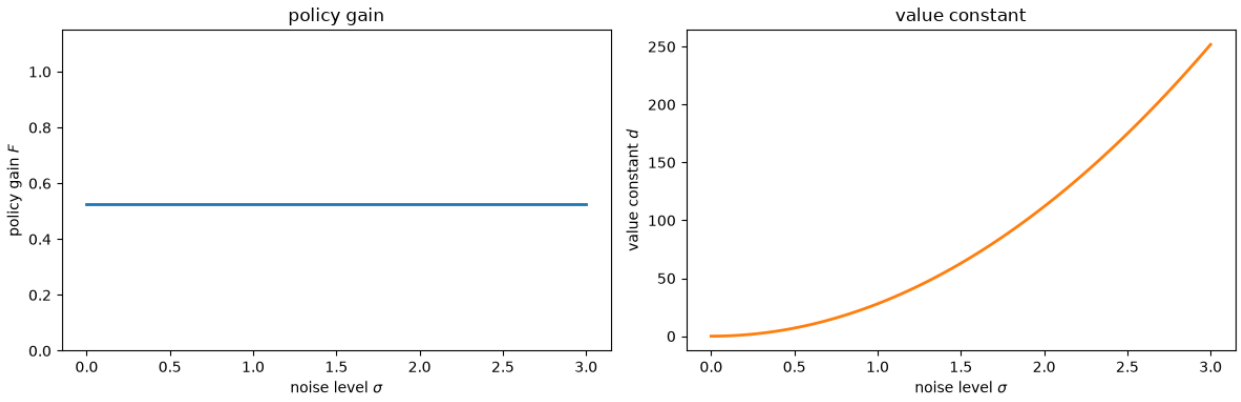


Fig. 83.1: CE: policy does not depend on noise

As the plot confirms, F (the policy gain) is *flat* across all noise levels, while the value constant d increases monotonically with σ.

This is the CE principle in action: uncertainty changes the value of the problem but not the optimal decision rule.

83.7 A trouble with ad hoc expectations

Prior practice, exemplified by the adaptive expectations mechanisms of Friedman [1956] and Cagan [1956], directly postulated a particular form of (83.6):

$$\theta_t^e = \lambda \sum_{i=0}^{\infty} (1-\lambda)^i \, \theta_{t-i}, \qquad 0 < \lambda < 1, \tag{83.9}$$

treating the coefficient λ as a free parameter to be estimated from data, with no reference to the underlying environment f.

The deficiency is not that (83.9) is a distributed lag – linear forecasting rules are perfectly acceptable simplifications.

The deficiency is that the **coefficients** of the distributed lag are left unrestricted by theory.

The mapping $h_2 = S(f)$ shows that optimal forecasting coefficients are *determined* by f: when f changes, h_2 changes, and so does h.

An estimated λ calibrated under f_0 is therefore non-structural and will give incorrect predictions whenever f is altered.

This is the econometric content of the critique delivered by Muth [1960].

Rational expectations equates the subjective distribution that agents use in forming $\tilde{z}_t$ to the objective distribution f that actually generates the data, thereby closing the model and eliminating free parameters in h_2.

83.8 Exercises

Exercise 83.8.1

Using the scalar LQ setup in the code cell above (with $a = 0.9$, $b = 1$, $q = r = 1$, $\beta = 0.95$), verify numerically that the value constant d satisfies $d \propto \sigma^2$.

Hint: From the CE analysis, the value constant satisfies $d = \frac{\beta}{1-\beta} \operatorname{tr}(C^\top P C)$, and since $C = \sigma$ in the scalar case, this gives $d = \frac{\beta}{1-\beta} P \, \sigma^2$.

Confirm that a plot of d against σ^2 is linear and compute the theoretical slope $\frac{\beta}{1-\beta} P$.

Solution

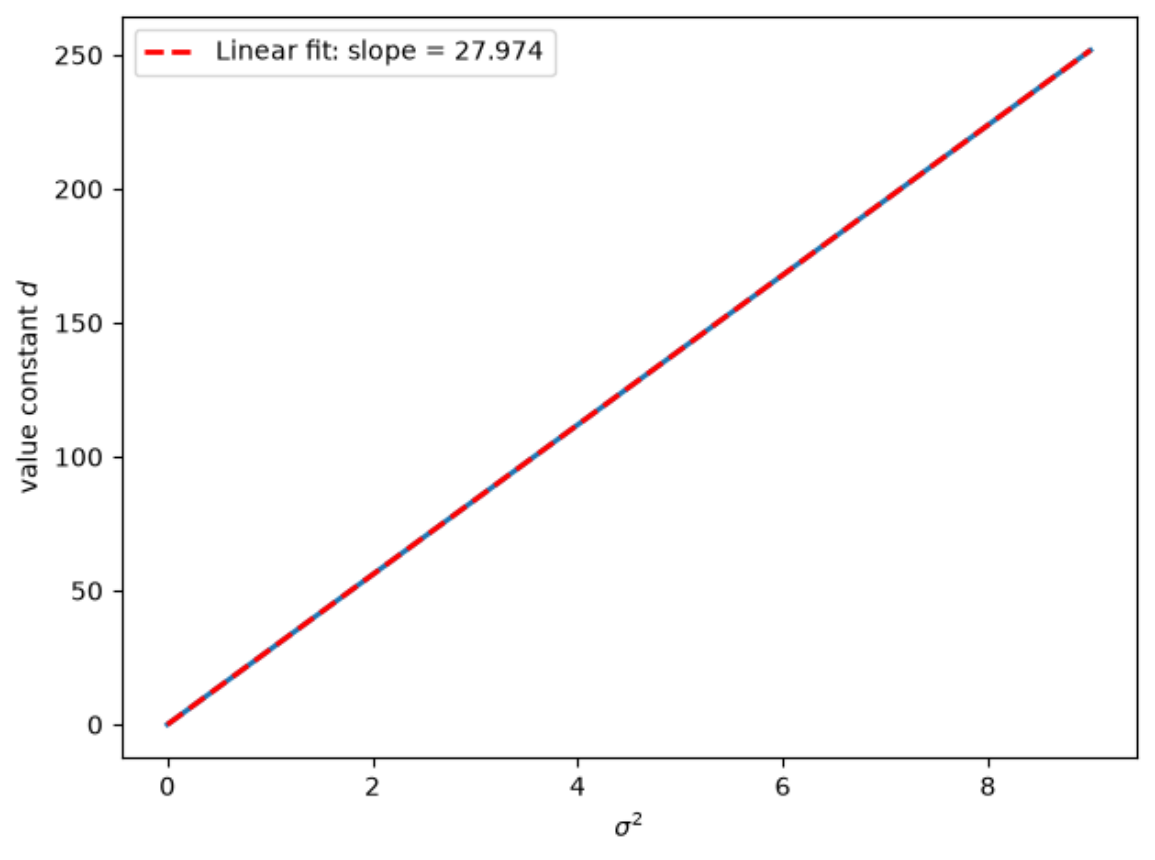
```
σ_sq_vals = σ_vals ** 2

fig, ax = plt.subplots()
ax.plot(σ_sq_vals, d_vals, lw=2)
ax.set_xlabel('$\\sigma^2$')
ax.set_ylabel('value constant $d$')

coeffs = np.polyfit(σ_sq_vals, d_vals, 1)
ax.plot(σ_sq_vals, np.polyval(coeffs, σ_sq_vals),
        'r--', lw=2, label=f'Linear fit: slope = {coeffs[0]:.3f}')
ax.legend()
fig.tight_layout()
plt.show()

P_scalar = float(LQ(Q_mat, R_mat, A, B, C=np.zeros((1, 1)),
```

```
                    beta=β).stationary_values()[0].item())
theoretical_slope = β / (1 - β) * P_scalar
print(f"Empirical slope:    {coeffs[0]:.4f}")
print(f"Theoretical slope β/(1-β)*P = {theoretical_slope:.4f}")
```




```
Empirical slope:    27.9740
Theoretical slope β/(1-β)*P = 27.9740
```

The slope is indeed $\frac{\beta}{1-\beta}P$, confirming the analytic formula.

The value matrix P is determined entirely by preferences and technology, not by the noise level – a direct consequence of the certainty equivalence principle.

83.9 Concluding remarks

This sequel *certainty equivalence and model uncertainty* describes how to extend the certainty equivalence principle to a linear-quadratic setting in which a decision maker distrusts the transition dynamics specified in his baseline model.

CHAPTER

EIGHTYFOUR

CERTAINTY EQUIVALENCE AND MODEL UNCERTAINTY

Contents

- *Certainty Equivalence and Model Uncertainty*
  - *Overview*
  - *Recap: ordinary certainty equivalence*
  - *Model uncertainty and robustness*
  - *Value function under robustness*
  - *Risk-sensitive preferences*
  - *Application: permanent income model*
  - *Summary*
  - *Exercises*

84.1 Overview

This is a sequel to *this lecture on certainty equivalence* that established an important *certainty equivalence* (CE) property for linear-quadratic (LQ) dynamic programming problems.

The property justifies a two-step algorithm for computing optimal decision rules:

1. *Optimize* under perfect foresight (treat future exogenous variables as known).
2. *Forecast* – substitute optimal forecasts for the unknown future values.

This lecture extends the certainty equivalence property in two directions motivated by [Hansen and Sargent, 2004]:

- *Model uncertainty and robustness.* What happens when the decision maker does not trust his model? A remarkable version of CE survives, but now the "forecasting" step uses a *distorted* probability distribution that the decision maker deliberately tilts against himself in order to achieve robustness.
- *Risk-sensitive preferences.* A mathematically equivalent reformulation interprets the same decision rules through recursive risk-sensitive preferences.

  The robustness parameter θ and the risk-sensitivity parameter σ are linked by $\theta = -\sigma^{-1}$.

We illustrate all three settings – ordinary CE, robust CE, and the permanent income application – with Python code using `quantecon`.

84.1.1 Model features

- Linear transition laws and quadratic objectives (LQ framework).
- Ordinary CE: optimal policy independent of noise variance.
- Robust CE: distorted forecasts replace baseline model forecasts; policy function depends on θ.
- Permanent income application: Hall's martingale, precautionary savings under robustness, and observational equivalence between robustness and patience.

This lecture draws on [Hansen and Sargent, 2004] and [Hansen and Sargent, 2008].

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

We use the following imports:

```
import numpy as np
import matplotlib.pyplot as plt
from quantecon import LQ, RBLQ
```

84.2 Recap: ordinary certainty equivalence

The *companion lecture* established the CE property in detail.

Here we collect only the elements needed for the robustness extension below.

The state vector $y_t = \begin{bmatrix} x_t \\ z_t \end{bmatrix}$ has an exogenous component z_t with transition law

$$z_{t+1} = f(z_t,\, \epsilon_{t+1}) \tag{84.1}$$

and an endogenous component x_t obeying

$$x_{t+1} = g(x_t,\, z_t,\, u_t). \tag{84.2}$$

Under the LQ assumption (quadratic return $r(y,u) = -y^\top Q y - u^\top R u$, linear f and g, Gaussian shocks), the optimal decision rule h decomposes as $u_t = h_1(x_t,\, h_2 \cdot z_t)$ where h_1 solves a nonstochastic control problem and h_2 solves an optimal forecasting problem.

The optimal value function is $V(y_0) = -y_0^\top P\, y_0 - p$ where, writing $z_{t+1} = f_1 z_t + f_2 \epsilon_{t+1}$:

- P is the fixed point of an operator $T(P; r, g, f_1)$ that does *not* involve the volatility matrix f_2, so neither P nor the decision rule h depends on the noise loadings.
- The constant $p = \beta/(1-\beta)\,\mathrm{tr}(f_2^\top P f_2)$ grows with volatility.

Uncertainty lowers the value (larger p) but does not alter behaviour.

The following code sets up a scalar LQ problem and confirms that the policy gain F is invariant to the noise level σ while d grows with it.

```
a, b_coeff = 0.9, 1.0
q, r = 1.0, 1.0
β = 0.95

A = np.array([[a]])
```

(continues on next page)

(continued from previous page)

```
B = np.array([[b_coeff]])
Q_mat = np.array([[q]])          # state cost
R_mat = np.array([[r]])          # control cost

σ_vals = np.linspace(0.0, 3.0, 80)
F_vals, d_vals = [], []

for σ in σ_vals:
    C = np.array([[σ]])
    lq = LQ(R_mat, Q_mat, A, B, C=C, beta=β)
    P, F, d = lq.stationary_values()
    F_vals.append(float(F[0, 0]))
    d_vals.append(float(d))

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].plot(σ_vals, F_vals, lw=2)
axes[0].set_xlabel('noise level $\\sigma$')
axes[0].set_ylabel('policy gain $F$')
axes[0].set_ylim(0, 2 * max(F_vals) + 0.1)

axes[1].plot(σ_vals, d_vals, lw=2, color='darkorange')
axes[1].set_xlabel('noise level $\\sigma$')
axes[1].set_ylabel('value constant $d$')

plt.tight_layout()
plt.show()
```

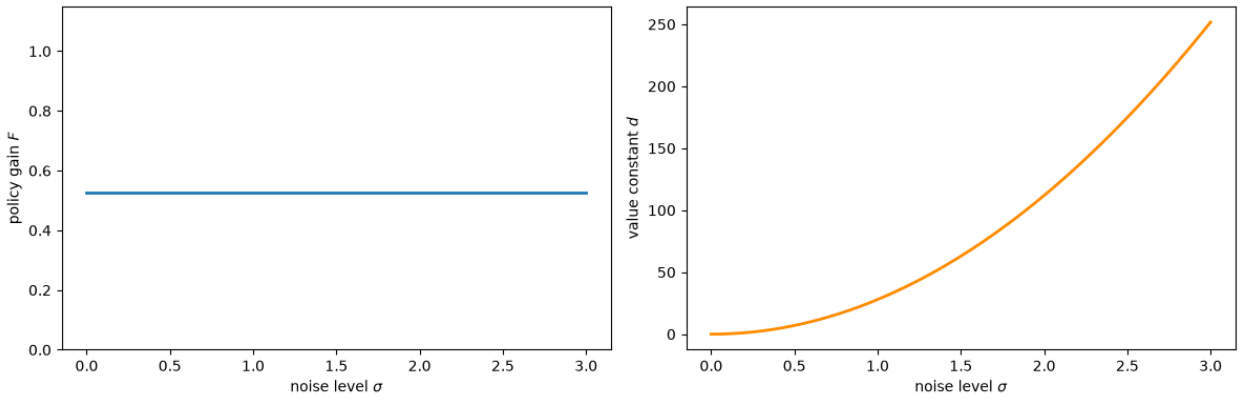


Fig. 84.1: CE principle – policy vs. value

84.3 Model uncertainty and robustness

84.3.1 Setup and the multiplier problem

The decision maker in Simon and Theil's setting knows his model exactly – he has no doubt about the transition law (84.1).

Now suppose he suspects that the true data-generating process is

$$z_{t+1} = f(z_t,\ \epsilon_{t+1} + w_{t+1}) \tag{84.3}$$

where $w_{t+1} = \omega_t(x^t, z^t)$ is a misspecification term chosen by an adversarial "nature."

The decision maker believes his approximating model is a good approximation in the sense that

$$\hat{\mathbb{E}}\left[\sum_{t=0}^{\infty} \beta^t\, w_{t+1}^\top w_{t+1} \,\Big|\, y_0\right] \le \eta_0, \tag{84.4}$$

where η_0 parametrises the tolerated misspecification budget and $\hat{\mathbb{E}}$ is the expectation under the distorted law (84.3).

To construct a *robust* decision rule the decision maker solves the **multiplier problem** – a two-player zero-sum dynamic game:

$$\min_{\{w_{t+1}\}} \max_{\{u_t\}} \hat{\mathbb{E}}\left[\sum_{t=0}^{\infty} \beta^t \Big\{ r(y_t, u_t) + \theta\beta\, w_{t+1}^\top w_{t+1} \Big\} \,\Big|\, y_0\right] \tag{84.5}$$

where $\theta > 0$ penalises large distortions.

A larger θ shrinks the feasible misspecification set; as $\theta \to \infty$ the problem reduces to ordinary LQ.

The Markov perfect equilibrium of (84.5) delivers a *robust* rule $u_t = h(x_t, z_t)$ together with a worst-case distortion process $w_{t+1} = W(x_t, z_t)$.

84.3.2 Stackelberg timing and the modified CE

The Markov perfect equilibrium *conceals* a form of CE.

To reveal it, Hansen and Sargent [2001] impose a **Stackelberg timing protocol**: at time 0, the *minimising* player commits once and for all to a plan $\{w_{t+1}\}$, after which the *maximising* player chooses u_t sequentially.

This makes the minimiser the Stackelberg leader.

To describe the leader's committed plan, introduce "big-letter" state variables (X_t, Z_t) (same dimensions as (x_t, z_t)) that encode the leader's pre-committed strategy:

$$\begin{aligned} w_{t+1} &= W(X_t, Z_t), \\ X_{t+1} &= g(X_t, Z_t,\ h(X_t, Z_t)), \\ Z_{t+1} &= f(Z_t,\ W(X_t, Z_t) + \epsilon_{t+1}). \end{aligned} \tag{84.6}$$

Summarised with $Y_t = \begin{bmatrix} X_t \\ Z_t \end{bmatrix}$:

$$Y_{t+1} = MY_t + N\epsilon_{t+1}, \qquad w_{t+1} = W(Y_t). \tag{84.7}$$

The maximising player then faces an *ordinary* dynamic programming problem subject to his own dynamics (84.2), the distorted z-law (84.3), and the exogenous process (84.7).

His optimal rule takes the form

$$u_t = \tilde{H}(x_t, z_t, Y_t). \tag{84.8}$$

Başar and Bernhard [2008] and Hansen and Sargent [2008] establish that at equilibrium (with "big K = little k" imposed) this collapses to

$$\tilde{H}(X_t, Z_t, Y_t) = h(Y_t), \tag{84.9}$$

the *same* rule as the Markov perfect equilibrium of (84.5).

84.3.3 Modified separation principle

The Stackelberg timing permits an Euler-equation approach.

The two-step algorithm becomes:

The first step is unchanged: solve the same nonstochastic control problem as before, with $\mathbf{z}_t = (z_t, z_{t+1}, ...)$ treated as known, giving $u_t = h_1(x_t, \mathbf{z}_t)$.

The second step is modified: form forecasts using the *distorted* law of motion (84.7). By the linearity and Gaussianity of the system,

$$\hat{\mathbb{E}}[\mathbf{z}_t \mid z^t, Y^t] = \hat{h}_2 \begin{bmatrix} z_t \\ Y_t \end{bmatrix} \tag{84.10}$$

where $\hat{\mathbb{E}}$ uses the distorted model.

Substituting (84.10) into h_1 and imposing $Y_t = y_t$ gives the robust rule

$$u_t = h_1\big(x_t,\ \hat{h}_2 \cdot y_t\big) = h(x_t, z_t). \tag{84.11}$$

This is the modified CE: *step 1 is identical to the non-robust case*; only step 2 changes, using distorted rather than rational forecasts.

In contrast to ordinary CE, the robust policy *does* change as θ varies.

As $\theta \to \infty$ (no robustness) the robust policy converges to the standard LQ policy.

```
σ_fixed = 1.0
C_fixed = np.array([[σ_fixed]])

lq_std = LQ(R_mat, Q_mat, A, B, C=C_fixed, beta=β)
P_std, F_std_arr, d_std = lq_std.stationary_values()
F_standard = float(F_std_arr[0, 0])
P_standard = float(P_std[0, 0])

θ_vals = np.linspace(2.0, 30.0, 120)   # restrict attention to a numerically stable␣
↪range
F_rob_vals, P_rob_vals = [], []

for θ in θ_vals:
    rblq = RBLQ(R_mat, Q_mat, A, B, C_fixed, β, θ)
    F_rob, K_rob, P_rob = rblq.robust_rule()
    F_rob_vals.append(float(F_rob[0, 0]))
    P_rob_vals.append(float(P_rob[0, 0]))

fig, axes = plt.subplots(1, 2, figsize=(12, 4))
```

(continues on next page)

(continued from previous page)

```
axes[0].plot(θ_vals, F_rob_vals, lw=2, label='Robust $F(\\theta)$')
axes[0].axhline(F_standard, color='r', linestyle='--', lw=2,
                label=f'Standard LQ ($F = {F_standard:.3f}$)')
axes[0].set_xlabel('robustness parameter $\\theta$')
axes[0].set_ylabel('policy gain $F$')
axes[0].legend()

axes[1].plot(θ_vals, P_rob_vals, lw=2, color='purple',
             label='Robust $P(\\theta)$')
axes[1].axhline(P_standard, color='r', linestyle='--', lw=2,
                label=f'Standard LQ ($P = {P_standard:.3f}$)')
axes[1].set_xlabel('robustness parameter $\\theta$')
axes[1].set_ylabel('value matrix $P$')
axes[1].legend()

plt.tight_layout()
plt.show()
```

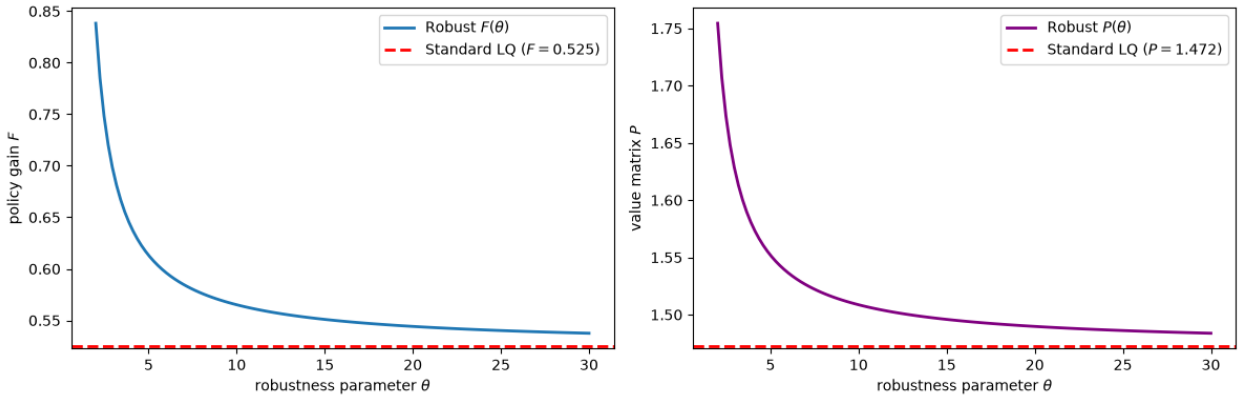


Fig. 84.2: Robust policy varies with θ

Observe that for small θ (strong preference for robustness) both F and P deviate substantially from their non-robust counterparts, converging to the standard values as $\theta \to \infty$.

This contrasts sharply with ordinary CE: under robustness, *both the policy gain and the value matrix depend on the robustness parameter θ and the noise-loading matrix C*.

84.4 Value function under robustness

Under a preference for robustness, the optimised value of (84.5) is again quadratic,

$$V(y_0) = -y_0^\top P\, y_0 - p, \tag{84.12}$$

but now *both P and p* depend on the volatility parameter f_2.

Specifically, P is the fixed point of the composite operator $T \circ \mathcal{D}$ where T is the same Bellman operator as in the non-robust case and $\mathcal{D}$ is the **distortion operator**:

$$\mathcal{D}(P) = \mathcal{D}(P;\, f_2,\, \theta). \tag{84.13}$$

Given the fixed point $P = T(\mathcal{D}(P))$, the constant is

$$p = p(P;\, f_2,\, \beta,\, \theta). \tag{84.14}$$

Despite P now depending on f_2, a form of CE still prevails: the same decision rule (84.11) also emerges from the *nonstochastic* game that maximises (84.5) subject to (84.2) and

$$z_{t+1} = f(z_t,\, w_{t+1}), \tag{84.15}$$

i.e., setting $\epsilon_{t+1} \equiv 0$.

The presence of randomness lowers the value (the constant p) but does not change the decision rule.

84.5 Risk-sensitive preferences

Building on Jacobson [1973] and Whittle [1990], Hansen and Sargent [2004] showed that the same decision rules can be reinterpreted through **risk-sensitive preferences**.

Suppose the decision maker *fully trusts* his model

$$y_{t+1} = A\, y_t + B\, u_t + C\, \epsilon_{t+1} \tag{84.16}$$

but evaluates stochastic processes according to the recursion

$$U_t = r(y_t, u_t) + \beta\, \mathcal{R}_t(U_{t+1}) \tag{84.17}$$

where the *risk-adjusted* continuation operator is

$$\mathcal{R}_t(U_{t+1}) = \frac{2}{\sigma} \log \mathbb{E}\left[\exp\left(\frac{\sigma U_{t+1}}{2} \right) \,\middle|\, y^t \right], \qquad \sigma \le 0. \tag{84.18}$$

When $\sigma = 0$, L'Hôpital's rule recovers the standard expectation operator.

When $\sigma < 0$, $\mathcal{R}_t$ penalises right-tail risk in the continuation utility U_{t+1}.

For a candidate quadratic continuation value $U^e_{t+1} = -y^\top_{t+1} \Omega\, y_{t+1} - \rho$, let $\hat{y}_{t+1} \equiv A y_t + B u_t$ denote the conditional mean of y_{t+1}.

Evaluating $\mathcal{R}_t$ via the log-moment-generating function of the Gaussian distribution yields

$$\mathcal{R}_t U^e_{t+1} = -\hat{y}^\top_{t+1} \mathcal{D}(\Omega)\, \hat{y}_{t+1} - \hat{\rho} \tag{84.19}$$

where $\mathcal{D}$ is the *same* distortion operator as in (84.13) with $\theta = -\sigma^{-1}$, and $\hat{\rho}$ is the corresponding scalar adjustment term.

Consequently, the risk-sensitive Bellman equation has the *same* fixed point P as the robust control problem, and therefore the *same decision rule* $u_t = -F y_t$.

> **Key equivalence:** robust control with parameter θ and risk-sensitive control with parameter $\sigma = -\theta^{-1}$ produce identical decision rules.

84.6 Application: permanent income model

We now illustrate all of the above in a concrete linear-quadratic permanent income model.

84.6.1 Model setup

A consumer receives an exogenous endowment process $\{z_t\}$ and allocates it between consumption c_t and savings x_t to maximise

$$-\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t (c_t - b)^2, \qquad \beta \in (0,1) \tag{84.20}$$

where b is a bliss level of consumption.

Defining the **marginal utility of consumption** $\mu_{ct} \equiv b - c_t$ (the control), the budget constraint and endowment process are

$$x_{t+1} = R\, x_t + z_t - b + \mu_{ct} \tag{84.21}$$

$$z_{t+1} = \mu_d(1-\rho) + \rho\, z_t + c_d(\epsilon_{t+1} + w_{t+1}) \tag{84.22}$$

where $R > 1$ is the gross return on savings, $|\rho| < 1$, and w_{t+1} is an optional shock-mean distortion representing model misspecification.

After absorbing the constants $-b$ and $\mu_d(1-\rho)$ by augmenting the state vector, or equivalently by working with deviations from steady state, setting $w_{t+1} \equiv 0$ and taking $Q = 0$ (return depends only on the control μ_{ct}) and $R_{\text{ctrl}} = 1$ puts this in the standard LQ form

$$y_t = \begin{bmatrix} x_t \\ z_t \end{bmatrix}, \quad A = \begin{bmatrix} R & 1 \\ 0 & \rho \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad C = \begin{bmatrix} 0 \\ c_d \end{bmatrix}. \tag{84.23}$$

In the numerical code below we add a negligible `1e-8 * I` regularisation to the state-cost matrix to keep the Riccati computation well conditioned in Hall's unit-root case $\beta R = 1$.

We calibrate to parameters estimated by Hansen *et al.* [1999] from post-WWII U.S. data:

```
β_hat = 0.9971
R_rate = 1.0 / β_hat          # β*R = 1 (Hall's case)
ρ     = 0.9992
c_d   = 5.5819
σ_rs  = -2e-7                  # σ_hat < 0
θ_pi  = -1.0 / σ_rs           # θ = -1/σ_hat

A_pi = np.array([[R_rate, 1.0],
                 [0.0,    ρ]])
B_pi = np.array([[1.0],
                 [0.0]])
C_pi = np.array([[0.0],
                 [c_d]])
Q_pi = 1e-8 * np.eye(2)       # regularise for β*R = 1
R_pi = np.array([[1.0]])
```

84.6.2 Without robustness: Hall's martingale

Setting $\sigma = 0$ (no preference for robustness), the consumer's Euler equation is

$$\mathbb{E}_t[\mu_{c,t+1}] = (\beta R)^{-1} \mu_{ct}. \tag{84.24}$$

With $\beta R = 1$ (Hall's case), this is $\mathbb{E}_t[\mu_{c,t+1}] = \mu_{ct}$, i.e., the **marginal utility of consumption is a martingale** – equivalently, consumption follows a random walk.

The optimal policy is $\mu_{ct} = -F y_t$ where, from the solved-forward Euler equation, $F = [(R-1),\ (R-1)/(R-\rho)]$.

The resulting closed-loop projection onto the one-dimensional direction of μ_{ct} gives the scalar AR(1) representation

$$\mu_{c,t+1} = \varphi\, \mu_{ct} + \nu\, \epsilon_{t+1}. \tag{84.25}$$

```
F_pi     = np.array([[(R_rate - 1.0), (R_rate - 1.0) / (R_rate - ρ)]])
A_cl_std = A_pi - B_pi @ F_pi

ψ_std = 1.0 / (β_hat * R_rate)
ν_std = (R_rate - 1.0) * c_d / (R_rate - ρ)

print(f"ψ = {ψ_std:.6f}, ν = {ν_std:.4f}")
```

```
ψ = 1.000000, ν = 4.3777
```

84.6.3 With robustness: precautionary savings

Under a preference for robustness ($\sigma < 0$, $\theta < \infty$), the consumer uses distorted forecasts $\hat{\mathbb{E}}_t[\cdot]$ evaluated under the worst-case model.

The consumption rule takes the certainty-equivalent form

$$\mu_{ct} = -(1 - R^{-2}\beta^{-1})\left(R\, x_t + \hat{\mathbb{E}}_t\left[\sum_{j=0}^{\infty} R^{-j}(z_{t+j} - b)\right]\right) \tag{84.26}$$

where h_1 – the first step of the CE algorithm – is *identical* to the non-robust case.

Only the expectations operator changes.

The resulting AR(1) dynamics for μ_{ct} become:

$$\mu_{c,t+1} = \tilde{\varphi}\, \mu_{ct} + \tilde{\nu}\, \epsilon_{t+1} \tag{84.27}$$

with $\tilde{\varphi} < 1$, implying $\mathbb{E}_t[c_{t+1}] > c_t$ under the approximating model – a form of **precautionary saving**.

The observational equivalence formula (84.28) (derived below) immediately gives the robust AR(1) coefficient: $\tilde{\varphi} = 1/(\tilde{\beta} R)$ where $\tilde{\beta} = \tilde{\beta}(\sigma)$.

The innovation scale $\tilde{\nu}$ follows from the robust permanent income formula with the distorted persistence; Hansen *et al.* [1999] report $\tilde{\nu} \approx 8.0473$ for their calibration.

```
def beta_tilde(σ, β_hat_val, α_sq_val):
    """Observational-equivalence locus: β_tilde(σ)."""
    denom = 2.0 * (1.0 + σ * α_sq_val)
    numer = β_hat_val * (1.0 + β_hat_val)
    disc  = 1.0 - 4.0 * β_hat_val * (1.0 + σ * α_sq_val) / \
            (1.0 + β_hat_val) ** 2
    return (numer / denom) * (1.0 + np.sqrt(np.maximum(disc, 0.0)))

ν_rob = 8.0473
α_sq  = ν_rob ** 2
bt    = beta_tilde(σ_rs, β_hat, α_sq)
ψ_rob = 1.0 / (bt * R_rate)

print(f"β_tilde = {bt:.5f}, ψ_tilde = {ψ_rob:.4f}, ν_tilde = {ν_rob:.4f}")
```

```
β_tilde = 0.99953, ψ_tilde = 0.9976, ν_tilde = 8.0473
```

```
rng = np.random.default_rng(0)
T_sim = 100

def simulate_ar1(ψ, ν, shocks, mu0=0.0):
    path = np.empty(len(shocks) + 1)
    path[0] = mu0
    for t, ε in enumerate(shocks, start=1):
        path[t] = ψ * path[t-1] + ν * ε
    return path

shock_path = rng.standard_normal(T_sim - 1)
mu0_init = 10.0
mu_std_path = simulate_ar1(ψ_std, ν_std, shock_path, mu0=mu0_init)
mu_rob_path = simulate_ar1(ψ_rob, ν_rob, shock_path, mu0=mu0_init)

fig, axes = plt.subplots(2, 1, figsize=(11, 6), sharex=True)
t_grid = np.arange(T_sim)

axes[0].plot(t_grid, mu_std_path, lw=2, label=f'$\\mu_{{ct}}$ (standard, $\\varphi={ψ_
↪std:.4f}$)')
axes[0].axhline(0, color='k', lw=0.8, linestyle='--')
axes[0].set_ylabel('$\\mu_{ct}$')
axes[0].legend(loc='upper right')

axes[1].plot(t_grid, mu_rob_path, lw=2, color='darkorange',
             label=f'$\\mu_{{ct}}$ (robust, $\\tilde{{\\varphi}}={ψ_rob:.4f}$)')
axes[1].axhline(0, color='k', lw=0.8, linestyle='--')
axes[1].set_xlabel('period $t$')
axes[1].set_ylabel('$\\mu_{ct}$')
axes[1].legend(loc='upper right')

plt.tight_layout()
plt.show()
```

84.6.4 Observational equivalence: robustness acts like patience

A key insight of Hansen and Sargent [2001] is that, in the permanent income model, a preference for robustness ($\sigma < 0$) is *observationally equivalent* to an increase in the discount factor from $\hat{\beta}$ to a larger value $\tilde{\beta}(\sigma)$, with σ set back to zero.

The equivalence locus is given by

$$\tilde{\beta}(\sigma) = \frac{\hat{\beta}(1+\hat{\beta})}{2(1+\sigma\alpha^2)}\left[1 + \sqrt{1 - \frac{4\hat{\beta}(1+\sigma\alpha^2)}{(1+\hat{\beta})^2}}\right] \tag{84.28}$$

where $\alpha^2 = \tilde{\nu}^2$ is the squared innovation loading in the robust AR(1) representation (84.27).

```
σ_range = np.linspace(-3e-7, 0.0, 200)
bt_vals = [beta_tilde(s, β_hat, α_sq) for s in σ_range]
bt_check = beta_tilde(σ_rs, β_hat, α_sq)

fig, ax = plt.subplots(figsize=(9, 5))
```

(continues on next page)

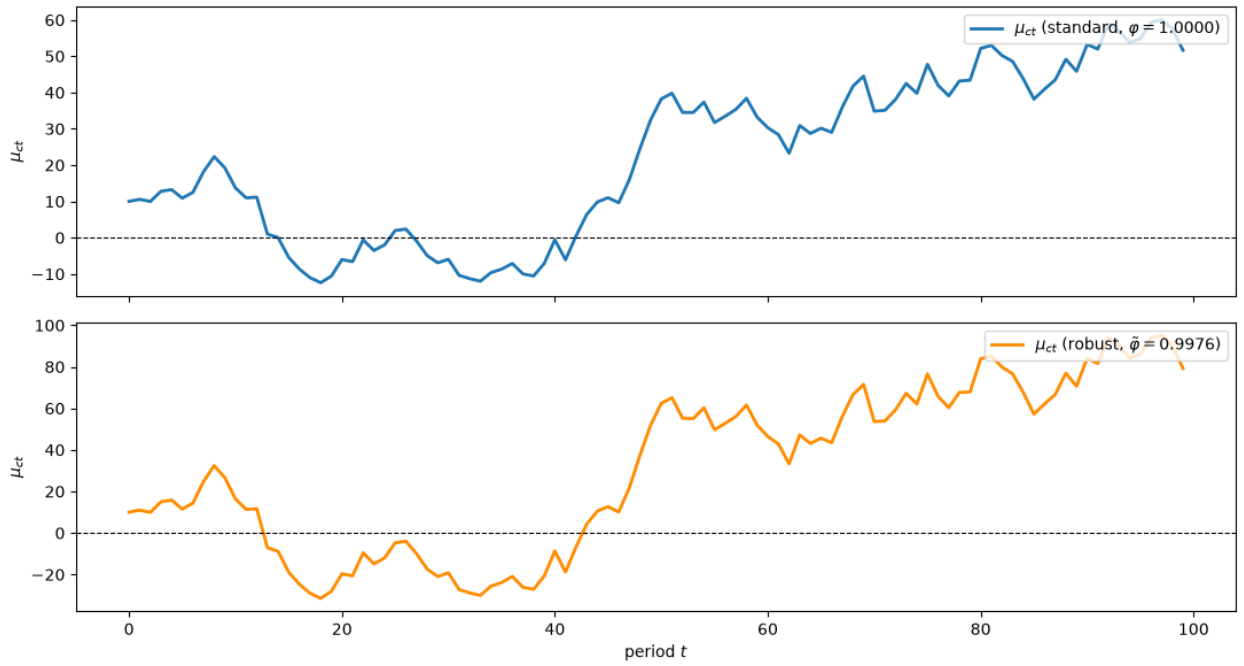


Fig. 84.3: Standard vs robust consumption paths

(continued from previous page)

```
ax.plot(-σ_range * 1e7, bt_vals, lw=2, color='steelblue',
        label='$\\tilde{\\beta}(\\sigma)$')
ax.axhline(β_hat, color='r', linestyle='--', lw=2,
           label=f'$\\hat{{\\beta}} = {β_hat}$')
ax.scatter([-σ_rs * 1e7], [bt_check], zorder=5, color='darkorange', s=80,
           label=f'$(\\hat{{\\sigma}},\\, \\tilde{{\\beta}}) '
                 f'= ({σ_rs:.0e},\\, {bt_check:.4f})$')
ax.set_xlabel('risk sensitivity $-\\sigma$ ($\\times 10^{-7}$)')
ax.set_ylabel('observationally equivalent discount factor $\\tilde{\\beta}$')
ax.legend()
plt.tight_layout()
plt.show()
```

The plot confirms the paper's key finding: *activating a preference for robustness is observationally equivalent – for consumption and saving behaviour – to increasing the discount factor*.

However, Hansen *et al.* [1999] show that the two parametrisations do *not* imply the same asset prices.

This happens because a preference for robustness generates different state-prices through the $\mathcal{D}(P)$ matrix that enters the stochastic discount factor.

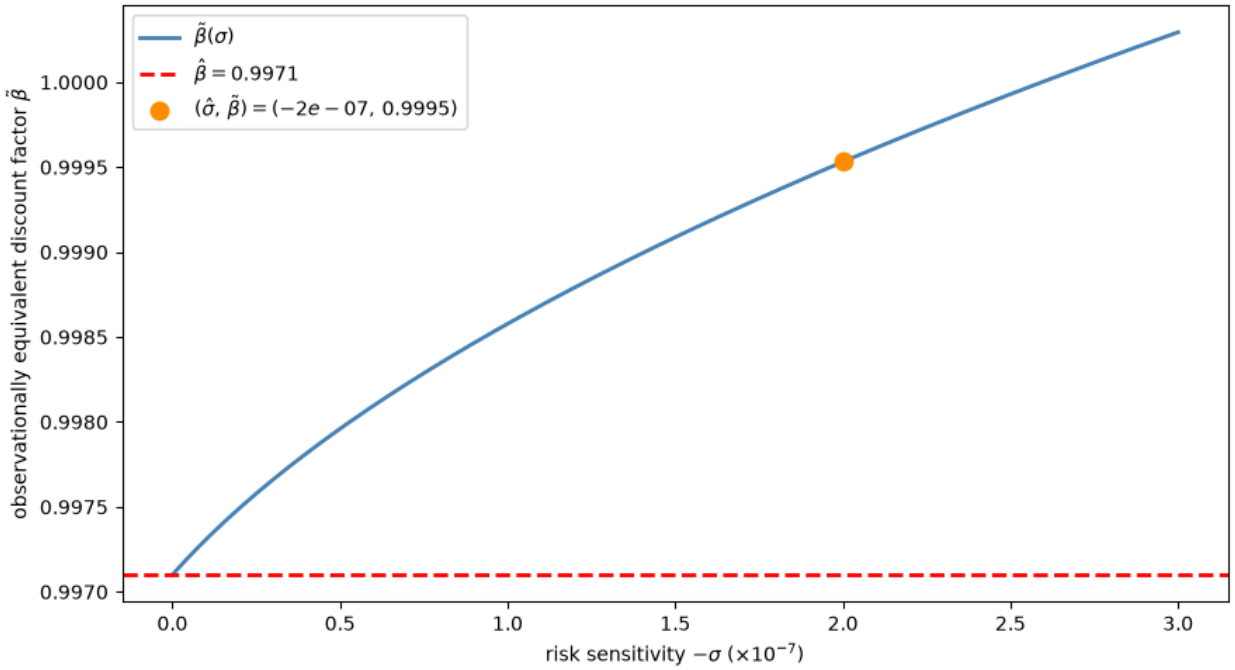


Fig. 84.4: Observational equivalence locus

84.7 Summary

The table below condenses the main results:

Setting	Policy depends on noise?	Forecasts used	CE survives?
Simon–Theil (ordinary LQ)	No	Rational	Yes
Robust control (multiplier)	Yes (P changes with f_2 and θ)	Distorted (worst-case)	Yes (modified)
Risk-sensitive preferences	Yes (same as robust)	Distorted (same)	Yes (same)

In all three cases, the decision maker can be described as following a two-step procedure: first solve a nonstochastic control problem, then form beliefs.

The difference is in which beliefs are formed in the second step.

84.8 Exercises

Exercise 84.8.1

Show numerically that as $\theta \to \infty$ the robust policy $F(\theta)$ converges to the standard LQ policy F_{std} and that the rate of convergence is of order $1/\theta$. Plot $|F(\theta) - F_{\text{std}}|$ against $1/\theta$ on a log–log scale.

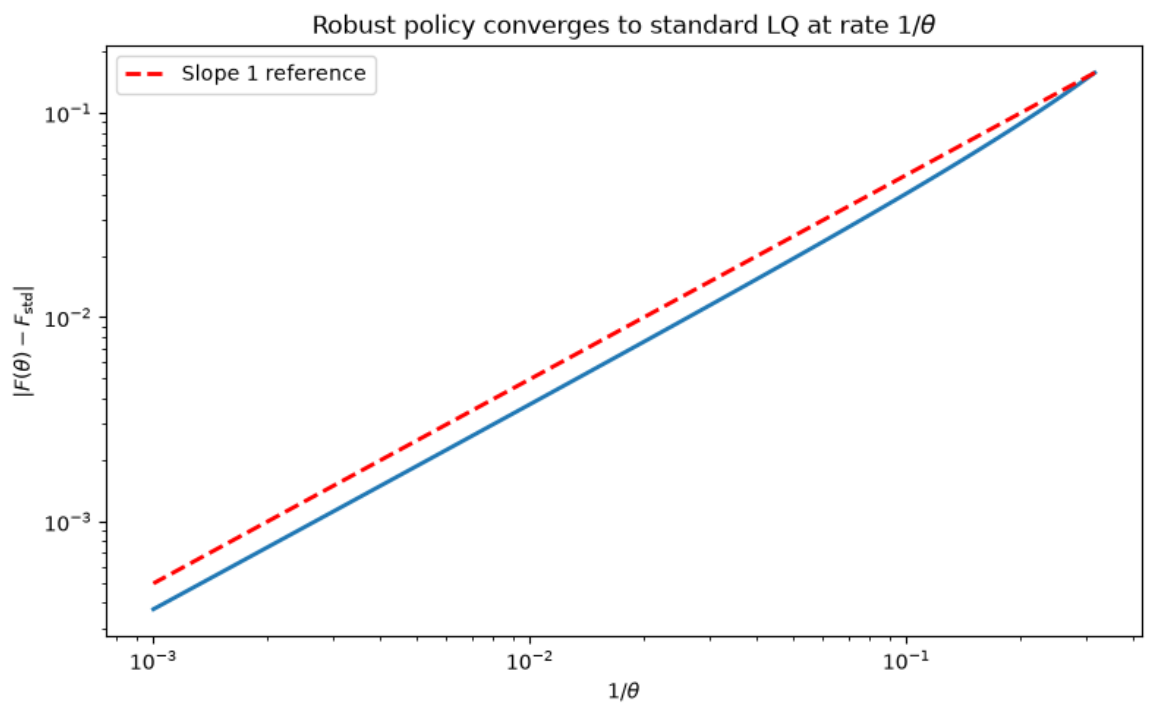
Solution

```
θ_large = np.logspace(0.5, 3.0, 100)
gap_vals = []

for θ in θ_large:
    rblq = RBLQ(R_mat, Q_mat, A, B, C_fixed, β, θ)
    F_r, _, _ = rblq.robust_rule()
    gap_vals.append(abs(float(F_r[0, 0]) - F_standard))

fig, ax = plt.subplots(figsize=(8, 5))
ax.loglog(1.0 / θ_large, gap_vals, lw=2)
ax.set_xlabel('$1/\\theta$')
ax.set_ylabel('$|F(\\theta) - F_{\\mathrm{std}}|$')
ax.set_title('Robust policy converges to standard LQ at rate $1/\\theta$')

x_ref = 1.0 / θ_large
ax.loglog(x_ref, x_ref * gap_vals[0] / x_ref[0],
          'r--', lw=2, label='Slope 1 reference')
ax.legend()
plt.tight_layout()
plt.show()
```



The log–log plot reveals an approximately linear relationship, confirming $O(1/\theta)$ convergence.

Exercise 84.8.2

Pick three values $\sigma_i < 0$ and verify numerically that the robust permanent income model with $(\sigma_i, \hat{\beta})$ produces the same policy matrix F as a suitably chosen non-robust model with $(0, \tilde{\beta}_i)$.

To find $\tilde{\beta}_i$, extract the AR(1) coefficient φ_i for μ_{ct} from the robust closed-loop dynamics and set $\tilde{\beta}_i = 1/(\varphi_i R)$.

Show that $\tilde{\beta}_i > \hat{\beta}$ in every case, confirming that robustness acts like increased patience.

Solution

For each σ_i we solve the robust problem with `RBLQ` and extract the AR(1) coefficient φ for μ_{ct} from the closed-loop dynamics $A_{\text{cl}} = A - BF_{\text{rob}}$.

If F is a left eigenvector of A_{cl} with eigenvalue φ, then $\mu_{ct} = -Fy_t$ satisfies $\mu_{c,t+1} = \varphi\,\mu_{ct} + \nu\,\epsilon_{t+1}$.

Setting $\tilde{\beta} = 1/(\varphi R)$ and solving a standard (non-robust) LQ problem with discount factor $\tilde{\beta}$ should reproduce F.

```
σ_trio = np.array([-5e-8, -1e-7, -2e-7])

for s in σ_trio:
    # Robust model: (σ, β_hat)
    θ_val = -1.0 / s
    rblq = RBLQ(R_pi, Q_pi, A_pi, B_pi, C_pi, β_hat, θ_val)
    F_rob, K_rob, P_rob = rblq.robust_rule()

    # Extract ψ from closed-loop under the approximating model
    A_cl = A_pi - B_pi @ F_rob
    ψ_rob = float((F_rob @ A_cl)[0, 1] / F_rob[0, 1])

    # Implied discount factor
    bt = 1.0 / (ψ_rob * R_rate)

    # Non-robust model with β_tilde
    lq_nr = LQ(R_pi, Q_pi, A_pi, B_pi, C=C_pi, beta=bt)
    P_nr, F_nr, d_nr = lq_nr.stationary_values()

    print(f"σ = {s:.1e},  θ = {θ_val:.1e},  β = {bt:.6f} (> β = {β_hat})")
    print(f"  ψ_rob = {ψ_rob:.8f}")
    print(f"  F_robust  = [{F_rob[0,0]:.6f}, {F_rob[0,1]:.6f}]")
    print(f"  F_non-rob = [{F_nr[0,0]:.6f}, {F_nr[0,1]:.6f}]")
    print(f"  |F_rob - F_nr| = {np.max(np.abs(F_rob - F_nr)):.2e}")
    print(f"  K (worst-case distortion): [{K_rob[0,0]:.2e}, {K_rob[0,1]:.2e}]")
    print()
```

```
σ = -5.0e-08,  θ = 2.0e+07,  β = 0.997472 (> β = 0.9971)
  ψ_rob = 0.99962721
  F_robust  = [0.003284, 0.884886]
  F_non-rob = [0.003284, 0.884894]
  |F_rob - F_nr| = 8.14e-06
  K (worst-case distortion): [2.48e-07, 6.68e-05]

σ = -1.0e-07,  θ = 1.0e+07,  β = 0.997953 (> β = 0.9971)
  ψ_rob = 0.99914484
  F_robust  = [0.003766, 1.014865]
  F_non-rob = [0.003766, 1.014863]
  |F_rob - F_nr| = 2.05e-06
  K (worst-case distortion): [5.68e-07, 1.53e-04]

σ = -2.0e-07,  θ = 5.0e+06,  β = 0.999522 (> β = 0.9971)
  ψ_rob = 0.99757687
  F_robust  = [0.005332, 1.437585]
  F_non-rob = [0.005333, 1.437465]
  |F_rob - F_nr| = 1.20e-04
  K (worst-case distortion): [1.01e-06, 2.72e-04]
```

The policy matrices F match to high precision, confirming observational equivalence for consumption and saving decisions.

In every case $\tilde{\beta} > \hat{\beta}$: a preference for robustness makes the agent behave as if he were more patient.

The non-zero worst-case distortion K in the robust model has no analogue in the non-robust model.

As Hansen *et al.* [1999] show, this is why the two parametrisations imply different asset prices even though saving plans coincide.

CHAPTER

EIGHTYFIVE

THE LQ PERMANENT INCOME MODEL

Contents

- *The LQ Permanent Income Model*
  - *Overview*
  - *The standard LQ permanent income model*
  - *Exercises*

85.1 Overview

This lecture studies consumption smoothing in a linear-quadratic (LQ) permanent income model.

It presents a rational-expectations version of the permanent income theories of Friedman [1956] and Hall [1978].

Throughout, we set $\beta R = 1$, so that the baseline consumer's subjective discount factor equals the bond price.

The model is useful for studying

- impulse response functions
- alternative state-space representations of the optimal decision rule
- cointegration of consumption and assets

We derive the consumer's optimal consumption function, present two state-space representations of the optimal decision rule, and illustrate them with two classic examples.

This is the first of four lectures on the LQ permanent income model.

The two sequels build directly on the tools developed here.

- *Consumption Smoothing with Incomplete and Complete Markets* studies the cross-section behavior of consumption and embeds the single consumer in closed economies with incomplete and complete markets.
- *Robust Consumption Smoothing and Precautionary Savings* studies a consumer who distrusts his model of income and engages in precautionary savings.
- *A Robust LQ Bewley Model* combines the two, building a Bewley economy whose consumers differ in how much they distrust their income model.

Let's begin with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.linalg import solve, inv, solve_discrete_lyapunov
from scipy.stats import norm
```

85.2 The standard LQ permanent income model

85.2.1 Setup

A consumer has preferences over consumption streams ordered by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t) \tag{85.1}$$

where $\mathbb{E}_t$ is a mathematical expectation conditioned on the consumer's time-t information, c_t is time-t consumption, $u(c)$ is a strictly concave one-period utility function, and $\beta \in (0,1)$ is a discount factor.

The consumer maximizes (85.1) by choosing a plan $\{c_t, b_{t+1}\}_{t=0}^{\infty}$ subject to the sequence of budget constraints

$$c_t + b_t = R^{-1} b_{t+1} + y_t, \quad t \geq 0 \tag{85.2}$$

where $\{y_t\}$ is an exogenous stationary endowment process, R is a constant gross risk-free interest rate, b_t is a one-period risk-free bond maturing at t, and b_0 is a given initial condition.

Note

For $t \geq 1$, b_t is chosen at time $t-1$.

The bond $b_t > 0$ represents debt owed by the consumer at the start of period t.

We assume $R^{-1} = \beta$.

The endowment or non-financial income process has the state-space representation

$$\begin{aligned} z_{t+1} &= \check{A} z_t + \check{C} w_{t+1} \\ y_t &= \check{G} z_t \end{aligned} \tag{85.3}$$

where w_{t+1} is IID with mean zero and identity covariance matrix, $\check{A}$ is a stable matrix (eigenvalues strictly less than one in modulus), and $\check{G}$ is a row vector.

The state confronting the household at t is $\begin{bmatrix} b_t & z_t^\top \end{bmatrix}^\top$, where b_t is its one-period debt due at the start of period t and z_t contains all variables useful for forecasting its future endowment.

To make the problem linear-quadratic, we adopt the **quadratic utility function**

$$u(c_t) = -\tfrac{1}{2}(c_t - \gamma)^2$$

where $\gamma > 0$ is a bliss level of consumption.

We allow c_t to be negative (a producer rather than a consumer).

We impose a **transversality condition**

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t b_t^2 < +\infty, \tag{85.4}$$

which rules out Ponzi schemes.

85.2.2 Euler equation and certainty equivalence

With quadratic utility, the first-order conditions for the consumer's problem imply that

$$\mathbb{E}_t c_{t+1} = c_t \tag{85.5}$$

Note

Equation (85.5) says that consumption is a martingale.

This is the key implication of the LQ permanent income model.

It contrasts with models that have convex marginal utility ($u''' > 0$), where consumption is instead a submartingale.

Because the consumer maximizes a quadratic objective subject to a linear transition equation, the problem satisfies a **certainty-equivalence** property.

This implies that we can find the optimal plan by

- first solving the problem while pretending to have perfect foresight; this lets us express c_t as a function of b_t and the continuation sequence $\{y_{t+j}\}_{j=0}^{\infty}$
- then simply replace $\{y_{t+j}\}_{j=0}^{\infty}$ with $\{\mathbb{E}_t y_{t+j}\}_{j=0}^{\infty}$.

85.2.3 The optimal consumption function

Solving the budget constraint (85.2) forward, imposing the transversality condition, and taking conditional expectations gives

$$b_t = \sum_{j=0}^{\infty} \beta^j \mathbb{E}_t y_{t+j} - \frac{1}{1-\beta} c_t \tag{85.6}$$

Rearranging yields the **consumption function**

$$c_t = (1-\beta)\left[\sum_{j=0}^{\infty} \beta^j \mathbb{E}_t y_{t+j} - b_t\right] \tag{85.7}$$

Equivalently, with net interest rate r defined by $\beta = 1/(1+r)$,

$$c_t = \frac{r}{1+r}\left[\sum_{j=0}^{\infty} \beta^j \mathbb{E}_t y_{t+j} - b_t\right] \tag{85.8}$$

Evidently, consumption at t equals $r/(1+r)$ times total wealth, where total wealth is the sum of human wealth $\sum_{j=0}^{\infty} \beta^j \mathbb{E}_t y_{t+j}$ and financial wealth $-b_t$.

Using state-space representation (85.3) to evaluate the geometric sum of expected future endowments,

$$\sum_{j=0}^{\infty} \beta^j \mathbb{E}_t y_{t+j} = \check{G}(I - \beta \check{A})^{-1} z_t, \tag{85.9}$$

we obtain

$$c_t = (1-\beta)\left[\check{G}(I - \beta \check{A})^{-1} z_t - b_t\right] \tag{85.10}$$

This expresses c_t as a function of the state $[b_t,\ z_t^\top]^\top$ that confronts the household.

85.2.4 Representation 1: state (b_t, z_t)

Combining the endowment law of motion with the optimal debt dynamics (derived by substituting (85.10) into (85.2)) gives the following representation:

$$\begin{aligned} z_{t+1} &= \check{A}\, z_t + \check{C}\, w_{t+1} \\ b_{t+1} &= b_t + \check{G}[(I - \beta\check{A})^{-1}(\check{A} - I)]z_t \\ y_t &= \check{G}\, z_t \\ c_t &= (1-\beta)\left[\check{G}(I - \beta\check{A})^{-1} z_t - b_t\right] \end{aligned} \tag{85.11}$$

In this representation the **exogenous** state is z_t and the **endogenous** state is b_t.

We turn now to an alternative representation.

85.2.5 Representation 2: state (c_t, z_t)

Hall [1978] showed that the LQ permanent income model implies a representation in which the state consists of current consumption c_t and the exogenous endowment state z_t.

In this representation, b_t becomes an outcome rather than a state variable.

Shifting (85.7) forward, eliminating b_{t+1} via (85.2), and rearranging yields

$$c_{t+1} - c_t = (1-\beta)\sum_{j=0}^{\infty} \beta^j (\mathbb{E}_{t+1} y_{t+j+1} - \mathbb{E}_t y_{t+j+1}) \tag{85.12}$$

The right-hand side is $(1-\beta)$ times the time-$(t+1)$ **innovation** to the expected present value of the endowment stream.

Suppose the endowment has the (Wold) moving-average representation

$$y_{t+1} = d(L)\, w_{t+1}, \qquad d(L) = \sum_{j=0}^{\infty} d_j L^j \tag{85.13}$$

where $d(L) = \check{G}(I - \check{A}L)^{-1}\check{C}$.

Then

$$\mathbb{E}_{t+1} y_{t+j} - \mathbb{E}_t y_{t+j} = d_{j-1}\, w_{t+1} \tag{85.14}$$

Substituting (85.14) into (85.12) gives the key result

$$c_{t+1} - c_t = (1-\beta)\, d(\beta)\, w_{t+1}. \tag{85.15}$$

Here, $d(\beta) = \check{G}(I - \beta\check{A})^{-1}\check{C}$ is the **present value of the (Wold) moving-average coefficients**.

Thus, consumption is a **random walk** with innovation $(1-\beta)d(\beta)w_{t+1}$.

Combining (85.15) and (85.6) gives

$$\begin{aligned} c_{t+1} &= c_t + (1-\beta)\,\check{G}(I - \beta\check{A})^{-1}\check{C}\, w_{t+1} \\ b_t &= \check{G}(I - \beta\check{A})^{-1} z_t - \frac{1}{1-\beta}\, c_t \\ y_t &= \check{G}\, z_t \\ z_{t+1} &= \check{A}\, z_t + \check{C}\, w_{t+1} \end{aligned} \tag{85.16}$$

This representation reveals several important features of the optimal decision rule:

1. **State**: The state consists of the endogenous component c_t and the exogenous component z_t, with financial assets b_t encoded in c_t rather than carried as a separate state.
2. **Random walk**: Consumption is a random walk with innovation $(1-\beta)d(\beta)w_{t+1}$, which confirms that the Euler equation (85.5) is built into the solution and implies that consumption has no asymptotic stationary distribution.
3. **Box impulse response**: For all $j \geq 1$, the response of c_{t+j} to the innovation w_{t+1} is the constant $(1-\beta)d(\beta)$, giving a "box-shaped" impulse response.
4. **Cointegration**: Both c_t and b_t are nonstationary (unit-root processes), but the linear combination $(1-\beta)b_t + c_t$ is stationary.

   From (85.6),

$$(1-\beta)b_t + c_t = (1-\beta)\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j y_{t+j} \tag{85.17}$$

The left side is the cointegrating residual.

85.2.6 Debt dynamics

Subtracting (85.16) (equation for b_t) at time t from the same equation at time $t+1$ and substituting gives

$$b_{t+1} - b_t = \check{G}(I - \beta\check{A})^{-1}(\check{A} - I)\, z_t \tag{85.18}$$

This shows that b_{t+1} is **predetermined** at time t as a function of z_t alone.

Solving backward from any t, b_t depends on the entire history $z^{t-1} = [z_{t-1}, \ldots, z_0]$ and the initial condition b_0.

Such **history dependence** is a hallmark of a consumption plan in various incomplete-markets economies.

85.2.7 Two classic examples

We illustrate formulas (85.16) with two examples.

In both, the endowment is $y_t = z_{1t} + z_{2t}$, where

$$\begin{pmatrix} z_{1,t+1} \\ z_{2,t+1} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} z_{1t} \\ z_{2t} \end{pmatrix} + \begin{pmatrix} \sigma_1 & 0 \\ 0 & \sigma_2 \end{pmatrix} \begin{pmatrix} w_{1,t+1} \\ w_{2,t+1} \end{pmatrix} \tag{85.19}$$

Here z_{1t} is a **permanent** component of y_t and z_{2t} is a **purely transitory** component.

In the full-information example, the consumer observes the state z_t at time t, so he can reconstruct w_{t+1} from z_{t+1} and z_t.

Applying (85.16):

$$c_{t+1} - c_t = \sigma_1 w_{1,t+1} + (1-\beta)\,\sigma_2\, w_{2,t+1} \tag{85.20}$$

A unit increment to the permanent component z_{1t} raises consumption *one-for-one* permanently and causes **zero net saving**.

A unit increment to the purely transitory component raises consumption by only the fraction $(1-\beta)$ permanently, while the remaining fraction β is saved.

From (85.18):

$$b_{t+1} - b_t = -z_{2t} = -\sigma_2 w_{2t} \tag{85.21}$$

confirming that none of the permanent shock is saved, while all of the transitory shock is saved.

In the incomplete-information (Muth model) example, the consumer observes y_t and its history, but not z_{1t} and z_{2t} separately.

The appropriate approach uses an **innovations representation** derived by the Kalman filter.

At the Kalman filter steady state, the **Kalman gain** $K \in [0, 1]$ satisfies

$$K = \frac{\Sigma}{\Sigma + \sigma_2^2}, \qquad \Sigma = \frac{\sigma_1^2 + \sqrt{\sigma_1^4 + 4\sigma_1^2\sigma_2^2}}{2} \tag{85.22}$$

where K increases with the ratio σ_1^2/σ_2^2 (the variance of the permanent shock relative to the transitory shock).

The innovations representation expresses the endowment as an ARMA(1,1) in its own innovation $a_t = y_t - \mathbb{E}[y_t \mid y^{t-1}]$ (the one-step-ahead forecast error):

$$y_{t+1} = y_t - (1 - K)\, a_t + a_{t+1} \tag{85.23}$$

Here the coefficient $-(1 - K)$ on the lagged innovation reflects that only the fraction K of last period's surprise was treated as permanent; the remainder mean-reverts.

The scalar a_t is IID with variance $\Sigma + \sigma_2^2$.

Applying (85.16) to this innovation representation:

$$c_{t+1} - c_t = [1 - \beta(1 - K)]\, a_{t+1} \tag{85.24}$$

The consumer regards a fraction K of the innovation a_{t+1} as permanent and fraction $1 - K$ as transitory.

He permanently increments consumption by $K + (1 - \beta)(1 - K) = 1 - \beta(1 - K)$ of a_{t+1} and saves the remaining fraction $\beta(1 - K)$.

The first difference of income obeys a first-order moving average:

$$y_{t+1} - y_t = a_{t+1} - (1 - K)\, a_t \tag{85.25}$$

By contrast, the first difference of consumption is IID by (85.24).

85.2.8 Implementation

```
# Parameters
β = 0.95       # discount factor (so R = 1/β)
σ1 = 0.15      # std of permanent shock
σ2 = 0.30      # std of transitory shock

# Example 1: full information
A_check = np.array([[1.0, 0.0],
                    [0.0, 0.0]])
C_check = np.array([[σ1, 0.0],
                    [0.0, σ2]])
G_check = np.array([[1.0, 1.0]])

# Key matrix M = G(I - βA)^{-1}
IbA = np.eye(2) - β * A_check
M = G_check @ inv(IbA)   # shape (1, 2)

# Consumption impulse responses
h = (1 - β) * M @ C_check          # shape (1, 2)
irf_perm_ex1 = h[0, 0] / σ1       # response per unit std of permanent shock
```

(continues on next page)

(continued from previous page)

```
irf_trans_ex1 = h[0, 1] / σ2       # response per unit std of transitory shock

print("Example 1 (full information)")
print(f"  IRF c to permanent shock  (normalised): {irf_perm_ex1:.4f}   "
      f"(theory: 1.0)")
print(f"  IRF c to transitory shock (normalised): {irf_trans_ex1:.4f}   "
      f"(theory: {1-β:.4f})")
```

```
Example 1 (full information)
  IRF c to permanent shock  (normalised): 1.0000   (theory: 1.0)
  IRF c to transitory shock (normalised): 0.0500   (theory: 0.0500)
```

```
# Example 2: partial information
Σ = (σ1**2 + np.sqrt(σ1**4 + 4 * σ1**2 * σ2**2)) / 2
K = Σ / (Σ + σ2**2)

print("Example 2 (partial information)")
print(f"  Steady-state Kalman gain K = {K:.4f}")
print(f"  IRF c to unit innovation a_{{t+1}}: {1 - β*(1-K):.4f}")
print(f"  Fraction of innovation treated as permanent (K): {K:.4f}")
print(f"  Fraction saved: β(1-K) = {β*(1-K):.4f}")
```

```
Example 2 (partial information)
  Steady-state Kalman gain K = 0.3904
  IRF c to unit innovation a_{t+1}: 0.4209
  Fraction of innovation treated as permanent (K): 0.3904
  Fraction saved: β(1-K) = 0.5791
```

```
# Compare impulse responses
T = 30
irf_c_ex1_perm = np.ones(T) * irf_perm_ex1 * σ1
irf_c_ex1_trans = np.ones(T) * irf_trans_ex1 * σ2

irf_c_ex2 = np.ones(T) * (1 - β * (1 - K))   # per unit innovation a_t

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].axhline(0, color='k', linewidth=0.8)
axes[0].step(range(T), irf_c_ex1_perm, where='post',
             label='permanent shock ($z_1$)', color='C0', lw=2)
axes[0].step(range(T), irf_c_ex1_trans, where='post',
             label='transitory shock ($z_2$)', color='C3',
             linestyle='--', lw=2)
axes[0].set_xlabel('periods after shock')
axes[0].set_ylabel('response of $c$')
axes[0].set_title('full information')
axes[0].legend()

axes[1].axhline(0, color='k', linewidth=0.8)
axes[1].step(range(T), irf_c_ex2, where='post',
             label=f'unit innovation $a_{{t+1}}$ (K = {K:.2f})',
             color='C4', lw=2)
axes[1].set_xlabel('periods after shock')
axes[1].set_ylabel('response of $c$')
axes[1].set_title('partial information')
```

(continues on next page)

(continued from previous page)

```
axes[1].legend()
fig.tight_layout()
plt.show()
```

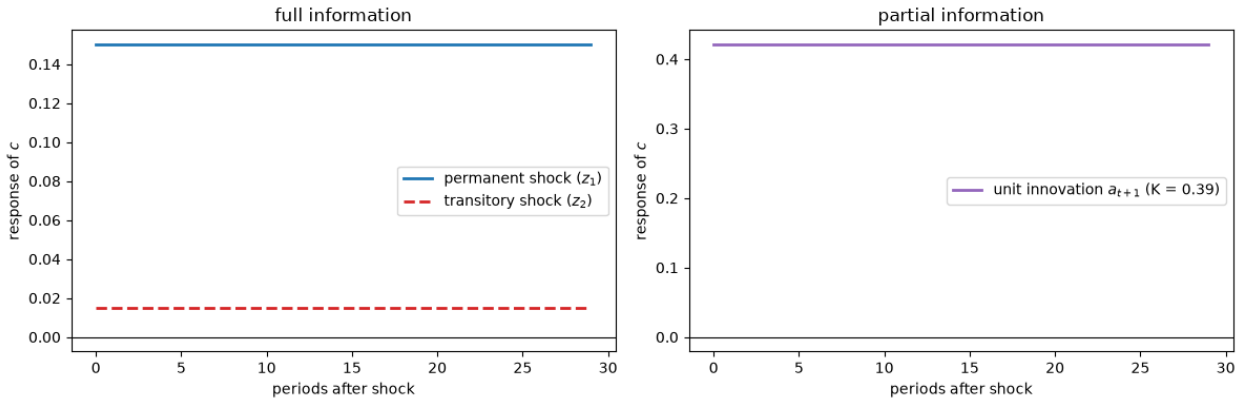


Fig. 85.1: Consumption impulse responses

Note

The impulse responses have the "box" shape characteristic of the LQ permanent income model: once a shock occurs, consumption shifts permanently to a new level and stays there.

The two representations and examples developed here are the foundation for the two sequel lectures.

Consumption Smoothing with Incomplete and Complete Markets uses the (c_t, z_t) representation (85.16) to study how the cross-section distribution of consumption evolves in closed economies with incomplete and complete markets.

Robust Consumption Smoothing and Precautionary Savings studies a consumer who distrusts the endowment process (85.3) and engages in precautionary savings.

A Robust LQ Bewley Model puts such consumers into a Bewley economy and shows that their concerns about misspecification leave no trace in quantity data.

85.3 Exercises

Exercise 85.3.1

This exercise verifies the response of consumption and debt to permanent and transitory shocks in the full-information two-factor model.

Consider the model of (85.19) with the calibration used above.

Suppose the economy starts at $z_0 = 0$ and $b_0 = 0$ and is hit by a single shock at $t = 1$: either a unit permanent shock $w_1 = (1, 0)^\top$ or a unit transitory shock $w_1 = (0, 1)^\top$, with no further shocks thereafter.

1. Using representation (85.11), compute and plot the paths of consumption c_t and debt b_t following each shock.

2. Confirm that the permanent shock raises consumption permanently by σ_1 and induces no saving, while the transitory shock raises consumption permanently by $(1-\beta)\sigma_2$ and is otherwise saved.

Solution

Here is one solution:

We iterate representation (85.11) forward with a single shock at $t = 1$.

```
def impulse_response(shock, T=20):
    """Paths of c and b after a single shock at t=1 (z_0 = b_0 = 0)."""
    I2 = np.eye(2)
    b_coef = M @ (A_check - I2)       # coefficient on z_t in the b law

    z = np.zeros((T + 1, 2))
    b = np.zeros(T + 1)

    z[1] = C_check @ shock            # shock realized at t = 1
    for t in range(1, T):
        z[t + 1] = A_check @ z[t]
    for t in range(T):
        b[t + 1] = b[t] + (b_coef @ z[t]).item()

    c = np.array([((1 - β) * (M @ z[t] - b[t])).item() for t in range(T + 1)])
    return c, b

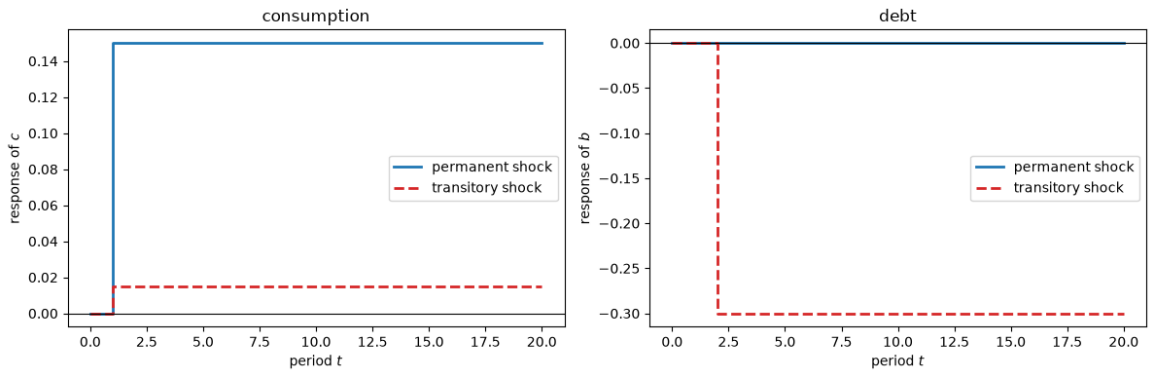

c_perm, b_perm = impulse_response(np.array([1.0, 0.0]))
c_tran, b_tran = impulse_response(np.array([0.0, 1.0]))

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].step(range(len(c_perm)), c_perm, where='post',
             label='permanent shock', color='C0', lw=2)
axes[0].step(range(len(c_tran)), c_tran, where='post',
             label='transitory shock', color='C3', linestyle='--', lw=2)
axes[0].axhline(0, color='k', lw=0.8)
axes[0].set_xlabel('period $t$')
axes[0].set_ylabel('response of $c$')
axes[0].set_title('consumption')
axes[0].legend()

axes[1].step(range(len(b_perm)), b_perm, where='post',
             label='permanent shock', color='C0', lw=2)
axes[1].step(range(len(b_tran)), b_tran, where='post',
             label='transitory shock', color='C3', linestyle='--', lw=2)
axes[1].axhline(0, color='k', lw=0.8)
axes[1].set_xlabel('period $t$')
axes[1].set_ylabel('response of $b$')
axes[1].set_title('debt')
axes[1].legend()
fig.tight_layout()
plt.show()

print(f"permanent shock:  Δc = {c_perm[-1]:.4f}  (theory σ1 = {σ1:.4f}),"
      f"  Δb = {b_perm[-1]:.4f}")
print(f"transitory shock: Δc = {c_tran[-1]:.4f}  "
      f"(theory (1-β)σ2 = {(1-β)*σ2:.4f}),  Δb = {b_tran[-1]:.4f}")
```




```
permanent shock:  Δc = 0.1500  (theory σ1 = 0.1500),  Δb = 0.0000
transitory shock: Δc = 0.0150  (theory (1-β)σ2 = 0.0150),  Δb = -0.3000
```

The permanent shock lifts consumption by σ_1 and leaves debt at zero: the shock is fully capitalised, so there is no net saving.

The transitory shock lifts consumption by only $(1-\beta)\sigma_2$; the consumer saves the remainder, so debt falls to $-\sigma_2$ (assets accumulate).

Exercise 85.3.2

This exercise illustrates the **cointegration** of consumption and debt described in (85.17).

The cointegration result requires a *stationary* endowment, so we replace the two-factor process with a scalar AR(1),

$$z_{t+1} = \rho z_t + \sigma_\varepsilon w_{t+1}, \qquad y_t = z_t,$$

with $\rho = 0.7$ and $\sigma_\varepsilon = 0.5$.

1. Simulate a long path of c_t and b_t using representation (85.11).
2. Verify that c_t and b_t each inherit a unit root (they wander), while the cointegrating residual $(1-\beta)b_t + c_t$ is stationary and equals $(1-\beta)\check{G}(I-\beta\check{A})^{-1}z_t$.

Solution

Here is one solution:

```
ρ, σε = 0.7, 0.5
A_ar = np.array([[ρ]])
C_ar = np.array([[σε]])
G_ar = np.array([[1.0]])

M_ar = G_ar @ inv(np.eye(1) - β * A_ar)      # G(I - βA)^{-1}
b_coef_ar = M_ar @ (A_ar - np.eye(1))        # coefficient on z_t in b law

rng = np.random.default_rng(0)
T = 400
z = np.zeros((T + 1, 1))
b = np.zeros(T + 1)
for t in range(T):
    z[t + 1] = A_ar @ z[t] + C_ar @ rng.standard_normal(1)
```

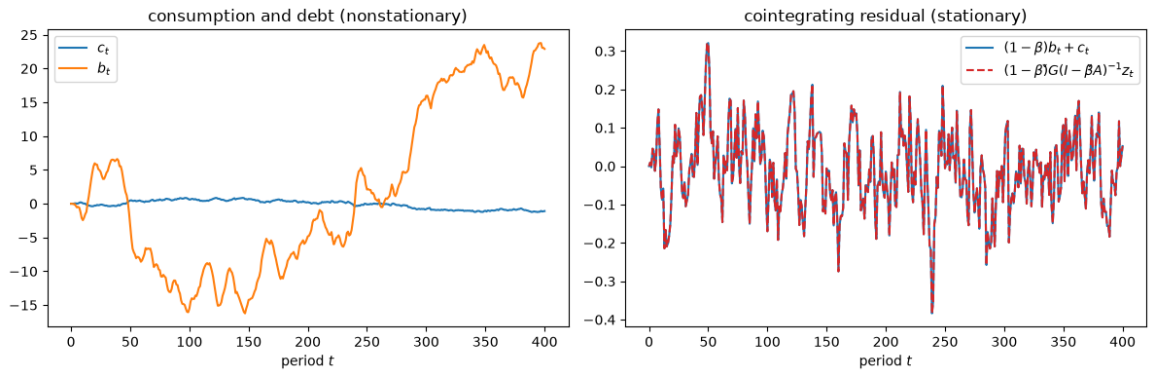
```
    b[t + 1] = b[t] + (b_coef_ar @ z[t]).item()

c = np.array([((1 - β) * (M_ar @ z[t] - b[t])).item() for t in range(T + 1)])
residual = (1 - β) * b + c
theory = np.array([((1 - β) * (M_ar @ z[t])).item() for t in range(T + 1)])

fig, axes = plt.subplots(1, 2, figsize=(12, 4))
axes[0].plot(c, label='$c_t$', lw=1.5)
axes[0].plot(b, label='$b_t$', lw=1.5)
axes[0].set_xlabel('period $t$')
axes[0].set_title('consumption and debt (nonstationary)')
axes[0].legend()

axes[1].plot(residual, label=r'$(1-\beta)b_t + c_t$', lw=1.5, color='C0')
axes[1].plot(theory, label=r'$(1-\beta)\check{G}(I-\beta\check{A})^{-1}z_t$',
             lw=1.5, linestyle='--', color='C3')
axes[1].set_xlabel('period $t$')
axes[1].set_title('cointegrating residual (stationary)')
axes[1].legend()
fig.tight_layout()
plt.show()

print(f"max |residual - theory| = {np.max(np.abs(residual - theory)):.2e}")
```




```
max |residual - theory| = 2.36e-16
```

Both c_t and b_t inherit the unit root that the random-walk consumption rule builds into the solution, so they drift without settling down.

Their linear combination $(1-\beta)b_t + c_t$ coincides with $(1-\beta)\check{G}(I-\beta\check{A})^{-1}z_t$ (up to floating-point error), which is a stationary function of the stationary state z_t.

Exercise 85.3.3

This exercise explores how the consumer's information problem shapes his response to income surprises in the Muth model.

Recall from (85.24) that the permanent consumption response to a unit income innovation a_{t+1} is $1-\beta(1-K)$, where the Kalman gain K in (85.22) depends on the ratio σ_1/σ_2.

1. Plot the response $1-\beta(1-K)$ as a function of the ratio σ_1/σ_2, holding σ_2 fixed.
2. Explain the two limiting cases $\sigma_1/\sigma_2 \to 0$ and $\sigma_1/\sigma_2 \to \infty$, and relate them to the two shocks of the full-

information model.

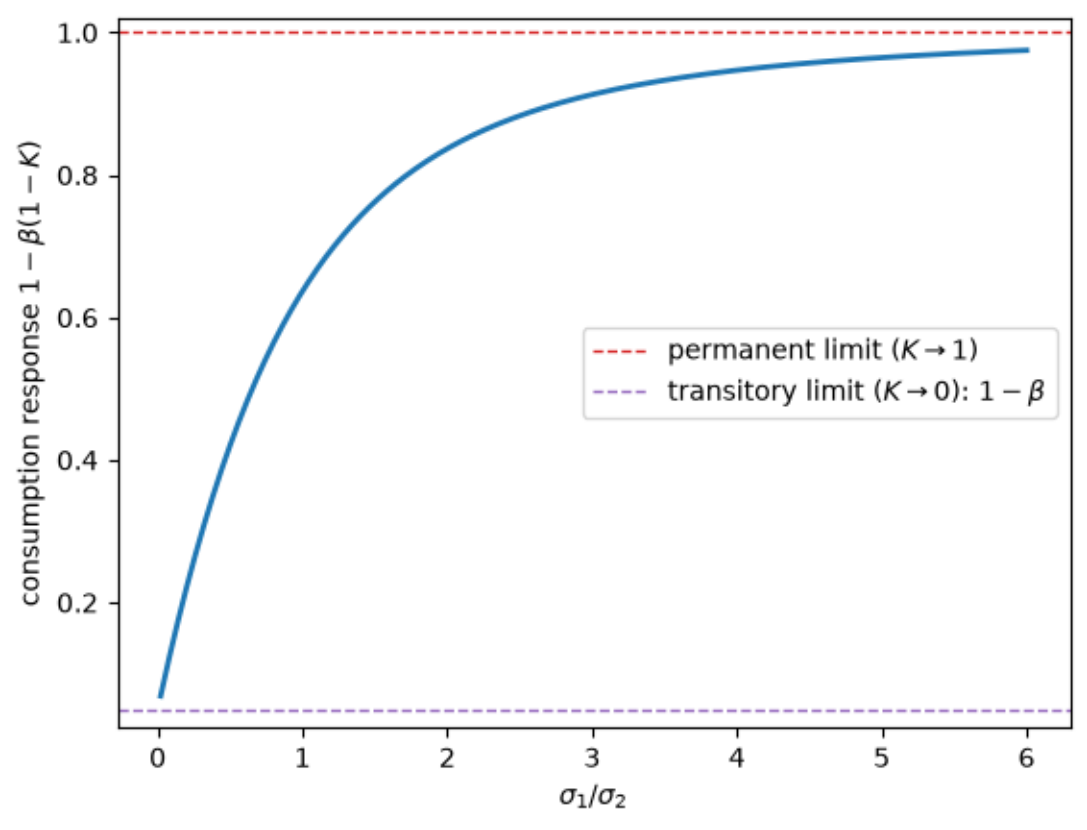
Solution

Here is one solution:

```
ratios = np.linspace(0.02, 6.0, 300)
σ2_fixed = 0.30

responses = []
for r in ratios:
    s1 = r * σ2_fixed
    Σ = (s1**2 + np.sqrt(s1**4 + 4 * s1**2 * σ2_fixed**2)) / 2
    K = Σ / (Σ + σ2_fixed**2)
    responses.append(1 - β * (1 - K))

fig, ax = plt.subplots()
ax.plot(ratios, responses, lw=2, color='C0')
ax.axhline(1.0, color='C3', linestyle='--', lw=1,
           label=r'permanent limit ($K\to1$)')
ax.axhline(1 - β, color='C4', linestyle='--', lw=1,
           label=r'transitory limit ($K\to0$): $1-\beta$')
ax.set_xlabel(r'$\sigma_1/\sigma_2$')
ax.set_ylabel(r'consumption response $1-\beta(1-K)$')
ax.legend()
plt.show()
```



As $\sigma_1/\sigma_2 \to 0$ the endowment is dominated by transitory noise, so $K \to 0$: the consumer treats each innovation as transitory and raises consumption by only $1-\beta$, exactly as for the purely transitory shock in the full-information model.

As $\sigma_1/\sigma_2 \to \infty$ the endowment is dominated by permanent shocks, so $K \to 1$: the consumer treats each innovation as permanent and raises consumption one-for-one, exactly as for the permanent shock.

For intermediate ratios the consumer optimally splits each surprise, capitalising the fraction K he attributes to the permanent component and saving the rest.

CHAPTER

EIGHTYSIX

CONSUMPTION SMOOTHING WITH INCOMPLETE AND COMPLETE MARKETS

Contents

- *Consumption Smoothing with Incomplete and Complete Markets*
  - *Overview*
  - *A brief review*
  - *Spreading consumption cross sections*
  - *A borrowers and lenders economy*
  - *Consumption smoothing with complete markets*
  - *Exercises*

86.1 Overview

This lecture studies how the cross-section distribution of consumption evolves when many consumers each solve the LQ permanent income problem.

It is the second of four lectures on the LQ permanent income model and builds directly on *The LQ Permanent Income Model*.

We first show that the unit root in individual consumption causes the cross-section variance of consumption to grow linearly with time.

We then embed the individual consumer in a **closed economy** with a continuum of consumers, following Bewley [1977], and show how the gross interest rate $R = \beta^{-1}$ emerges as an equilibrium outcome.

Finally, we replace the single risk-free bond with a complete set of Arrow securities and show how complete markets deliver a time-invariant cross-section distribution of consumption.

The third lecture, *Robust Consumption Smoothing and Precautionary Savings*, relaxes the assumption that the consumer fully trusts his income model.

The fourth, *A Robust LQ Bewley Model*, returns to the Bewley economy of this lecture and populates it with consumers who differ in how much they trust that model.

Let's begin with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.linalg import inv
```

86.2 A brief review

We recall the essentials from *The LQ Permanent Income Model*.

A consumer with quadratic utility, discount factor β, and access to a risk-free bond with gross return $R = \beta^{-1}$ faces the endowment process

$$\begin{aligned} z_{t+1} &= \check{A}\, z_t + \check{C}\, w_{t+1} \\ y_t &= \check{G}\, z_t \end{aligned} \tag{86.1}$$

where w_{t+1} is IID with mean zero and identity covariance matrix.

The optimal consumption function expresses consumption as $r/(1+r) = (1-\beta)$ times total wealth,

$$c_t = (1-\beta)\left[\check{G}(I - \beta\check{A})^{-1} z_t - b_t\right] \tag{86.2}$$

The model has a state-space representation in which the state is current consumption c_t and the exogenous endowment state z_t:

$$\begin{aligned} c_{t+1} &= c_t + (1-\beta)\,\check{G}(I - \beta\check{A})^{-1}\check{C}\, w_{t+1} \\ b_t &= \check{G}(I - \beta\check{A})^{-1} z_t - \frac{1}{1-\beta}\, c_t \\ y_t &= \check{G}\, z_t \\ z_{t+1} &= \check{A}\, z_t + \check{C}\, w_{t+1} \end{aligned} \tag{86.3}$$

Consumption is a **random walk**: its first difference is the IID innovation $h\, w_{t+1}$, where $h = (1-\beta)\,\check{G}(I - \beta\check{A})^{-1}\check{C}$.

Throughout we use the two-factor endowment $y_t = z_{1t} + z_{2t}$ from *The LQ Permanent Income Model*, with z_{1t} a permanent component and z_{2t} a purely transitory component, so that $\check{A} = \mathrm{diag}(1, 0)$ and $\check{C} = \mathrm{diag}(\sigma_1, \sigma_2)$.

The following cell reproduces the calibration and the key matrices.

```
# Parameters
β = 0.95        # discount factor (so R = 1/β)
σ1 = 0.15       # std of permanent shock
σ2 = 0.30       # std of transitory shock

# Two-factor endowment
A_check = np.array([[1.0, 0.0],
                    [0.0, 0.0]])
C_check = np.array([[σ1, 0.0],
                    [0.0, σ2]])
G_check = np.array([[1.0, 1.0]])

# Key matrix M = G(I - βA)^{-1}
IbA = np.eye(2) - β * A_check
M = G_check @ inv(IbA)    # shape (1, 2)
```

86.3 Spreading consumption cross sections

The unit root in consumption (representation (86.3)) causes a **cross-section variance** of consumption to grow linearly with time.

Consider a continuum of *ex ante* identical households born at $t = 0$.

All households i share the same preferences.

They all face a stochastic process for non-financial income of the same form

$$y_t^i = g(w_t^i, w_{t-1}^i, \ldots, w_0^i; z_0^i), \tag{86.4}$$

While all consumers face the same g process, they have different, statistically independent realizations of the **idiosyncratic** shock sequences $\{w_t^i\}_{t=0}^{\infty}$.

Let all households start from the same initial conditions $c_0^i = c_0$ and z_0^i.

From (86.3), household i's consumption follows

$$c_{t+1}^i = c_t^i + h\, w_{t+1}^i, \qquad h = (1-\beta)\,\check{G}(I - \beta\check{A})^{-1}\check{C}$$

Since $\{w_t^i\}$ realizations are independent across agents,

$$\mathbb{E}_0(c_t^i - c_0^i)^2 = t\, hh^\top \tag{86.5}$$

In the two-factor model, h is a 1×2 row vector so $hh^\top$ is a positive scalar equal to $\sigma_1^2 + (1-\beta)^2\sigma_2^2$.

The cross-section variance of consumption grows like t.

```
# Simulate cross-section spreading
rng = np.random.default_rng(42)
N = 5000        # number of agents
T_sim = 80          # number of periods

h_vec = (1 - β) * (M @ C_check)    # shape (1, 2), then flatten
h_vec = h_vec.flatten()             # h = [h1, h2]

c = np.zeros((N, T_sim + 1))        # consumption paths
# initialise all agents at c_0 = 0 (demeaned)
for t in range(T_sim):
    eps = rng.standard_normal((N, 2))     # N draws of 2D shock
    dc = eps @ h_vec                # shape (N,)
    c[:, t+1] = c[:, t] + dc

# Cross-section variance at each date
var_c = np.var(c, axis=0)
theory = np.arange(T_sim + 1) * np.dot(h_vec, h_vec)

fig, ax = plt.subplots()
ax.plot(var_c, label='simulated cross-section variance', lw=2)
ax.plot(theory, label=r'theoretical: $t \cdot h h^\top$',
        linestyle='--', color='C3', lw=2)
ax.set_xlabel('period $t$')
ax.set_ylabel('cross-section variance of $c$')
ax.legend()
plt.show()
```

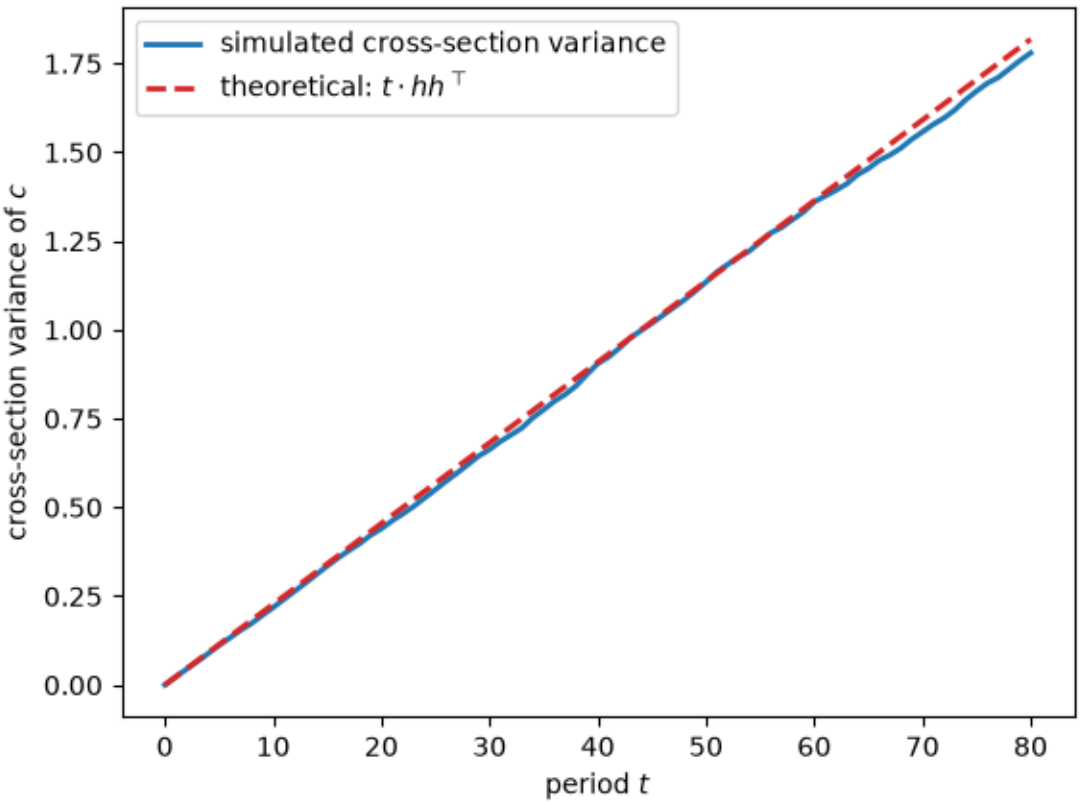


Fig. 86.1: Spreading consumption cross sections

86.4 A borrowers and lenders economy

Up to now we have set $R = \beta^{-1}$ and taken it as determined outside the model ("small open economy").

Following ideas of Bewley [1977], we can construct a **closed economy** in which $R = \beta^{-1}$ is an **equilibrium outcome**.

A continuum of measure one of consumers, indexed by $i \in [0, 1]$, trade a risk-free one-period bond with price β.

All consumers have the same preferences and the same stochastic income process (86.4), but face **idiosyncratic** non-financial income shock process realizations.

Initial bond positions are zero: $b_0^i = 0$ for all i.

Initial endowment states z_0^i are independent draws from a common initial distribution.

Because the permanent component z_{1t} has a unit root, process (86.1) has no stationary distribution, so in the simulation below we draw the permanent component $z_{10}^i \sim N(0, 1)$ and draw the transitory component $z_{20}^i \sim N(0, \sigma_2^2)$ from its stationary distribution.

From (86.2), with $b_0^i = 0$, agent i's time-0 consumption is

$$c_0^i = (1 - \beta)\, \check{G}(I - \beta \check{A})^{-1} z_0^i \tag{86.6}$$

For $t \geq 1$, from (86.3):

$$c_{t+1}^i = c_t^i + h\, w_{t+1}^i, \qquad h = (1 - \beta)\, \check{G}(I - \beta \check{A})^{-1} \check{C} \tag{86.7}$$

Let Y denote the stationary mean of the cross-section average of non-financial income.

Integrating (86.6) over all agents:

$$\int_0^1 c_0^i \, di = (1 - \beta) \sum_{j=0}^{\infty} \beta^j \mathbb{E}_0 \int_0^1 y_j^i \, di = Y \tag{86.8}$$

because the continuum of idiosyncratic shocks averages to zero.

For future periods, integrating (86.7):

$$\int_0^1 c_{t+1}^i \, di = \int_0^1 c_t^i \, di + h \underbrace{\int_0^1 w_{t+1}^i \, di}_{=0} = Y$$

The goods market clears at every date at *constant* aggregate consumption equal to Y.

The bond market clears at zero net supply each period.

Thus $R = \beta^{-1}$ is an equilibrium outcome.

While the cross-section mean of consumption is constant, the cross-section *variance* grows without bound according to (86.5).

Initial differences in endowment draws z_0^i create permanent differences in consumption levels.

```
# Verify Bewley market clearing via simulation
# Online mean and variance avoid storing all paths.
rng = np.random.default_rng(0)
N_bew = 10000    # number of agents
T_bew = 60

# Draw initial states for the simulation.
```

(continues on next page)

(continued from previous page)

```
z0_i = rng.standard_normal((N_bew, 2)) * np.array([1.0, σ2])
c0_i = ((1 - β) * (M @ z0_i.T)).flatten()   # shape (N_bew,)

# Propagate consumption across agents.
mean_c = np.zeros(T_bew + 1)
var_c2 = np.zeros(T_bew + 1)
mean_c[0] = c0_i.mean()
var_c2[0] = c0_i.var()

c_now = c0_i.copy()
for t in range(T_bew):
    eps = rng.standard_normal((N_bew, 2))
    c_now = c_now + eps @ h_vec
    mean_c[t + 1] = c_now.mean()
    var_c2[t + 1] = c_now.var()

# Reuse initial consumption below.
c_bew_t0 = c0_i

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].plot(mean_c, lw=2, color='C0')
axes[0].axhline(mean_c[0], linestyle='--', color='C3', label='initial mean')
axes[0].set_xlabel('period $t$')
axes[0].set_ylabel('mean consumption')
axes[0].set_title('mean')
axes[0].legend()

axes[1].plot(var_c2, lw=2, color='C0', label='simulated variance')
axes[1].set_xlabel('period $t$')
axes[1].set_ylabel('variance of consumption')
axes[1].set_title('variance')
axes[1].legend()
fig.tight_layout()
plt.show()
```

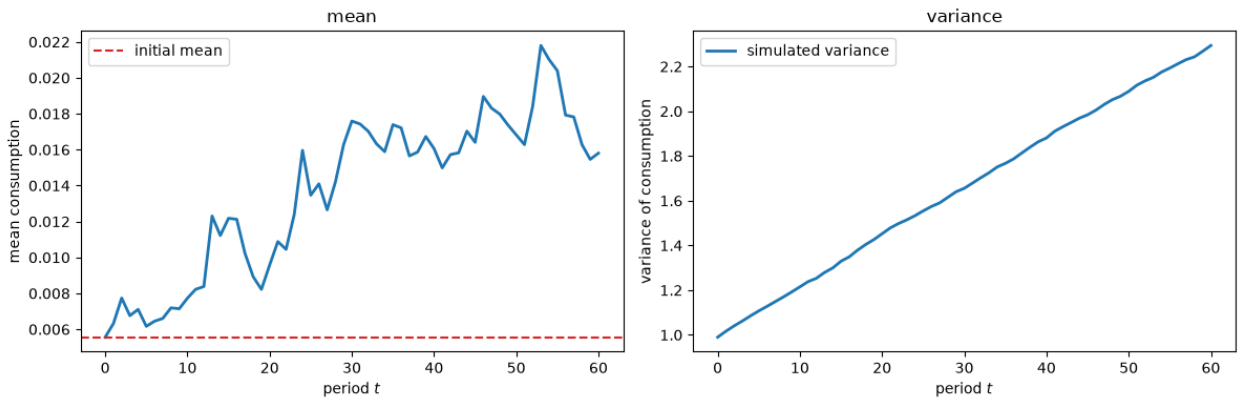


Fig. 86.2: Bewley economy cross-section moments

Because each consumer dislikes variation of consumption over time, each consumer would prefer a completely smoothed stream $c_t^i = c_0^i$ for all t.

Such an allocation is feasible because the cross-section average of income is constant.

The next section describes a complete-markets allocation that supports this outcome.

86.5 Consumption smoothing with complete markets

We replace the single bond with a **complete set of Arrow securities**.

The budget constraint becomes

$$c_t + b_{t-1}(z_t) = \int q(z_{t+1}|z_t)\, b_t(z_{t+1})\, dz_{t+1} + y_t \tag{86.9}$$

where $q(z_{t+1}|z_t)$ is the pricing kernel for one-period state-contingent claims and $b_t(z_{t+1})$ is the household's portfolio of Arrow securities chosen at t.

We guess (and verify) that the equilibrium pricing kernel is

$$q(z_{t+1}|z_t) = \beta\, \phi(z_{t+1}|z_t) \tag{86.10}$$

where $\phi(z_{t+1}|z_t)$ is the transition density of z.

This kernel prices a one-period risk-free bond at β, so $R = \beta^{-1}$, consistent with the incomplete-markets equilibrium.

We conjecture that the equilibrium delivers each consumer i a *constant* consumption level:

$$c_t^i = \bar{c}^i = c_0^i, \quad \forall\, t \geq 0 \tag{86.11}$$

where $c_0^i = (1-\beta)\, \check{G}(I - \beta \check{A})^{-1} z_0^i$ is the consumer's time-0 consumption in the incomplete-markets economy.

The state-contingent debt that supports constant consumption is

$$b_{t-1}(z_t) = \check{G}(I - \beta \check{A})^{-1} z_t - \frac{1}{1-\beta}\, \bar{c}^i \;\equiv\; b(z_t, \bar{c}^i) \tag{86.12}$$

Note that indebtedness depends only on the current Markov state z_t, *not* on the history of earlier states.

This absence of history dependence reflects the **complete risk sharing** attained under complete markets.

Substituting the pricing kernel (86.10) and the portfolio conjecture (86.12) into the budget constraint (86.9) and using the law of iterated expectations confirms that the budget constraint simplifies to $c_t = \bar{c}^i$ in every state and period.

Under complete markets, the cross-section distribution of consumption is **time-invariant**.

Consumer i's rank in the consumption distribution is fixed forever.

A lucky initial draw z_0^i manifests itself as perpetually high consumption $\bar{c}^i$ and lower indebtedness $b(z_t^i, \bar{c}^i)$ across all future states.

This outcome contrasts with what happens in the incomplete-markets Bewley economy, where the cross-section variance of consumption grows without bound.

```
# Complete and incomplete consumption distributions
rng = np.random.default_rng(1)
N_cm = 5000
T_cm = 50

# Initial consumption draws (same as Bewley economy)
c0_cm = c_bew_t0[:N_cm]
```

(continues on next page)

(continued from previous page)

```
# Incomplete markets: consumption evolves (random walk)
c_inc = np.zeros((N_cm, T_cm + 1))
c_inc[:, 0] = c0_cm
for t in range(T_cm):
    eps = rng.standard_normal((N_cm, 2))
    c_inc[:, t+1] = c_inc[:, t] + eps @ h_vec

# Complete markets: consumption stays constant
c_comp = np.tile(c0_cm[:, np.newaxis], T_cm + 1)

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

for t_plot, color in zip([0, 10, 30, 50], ['C0', 'C1', 'C3', 'C4']):
    axes[0].hist(c_inc[:, t_plot], bins=60, alpha=0.4,
                 label=f't = {t_plot}', color=color, density=True)
axes[0].set_xlabel('$c$')
axes[0].set_ylabel('density, incomplete markets')
axes[0].set_title('incomplete markets')
axes[0].legend(fontsize=9)

for t_plot, color in zip([0, 10, 30, 50], ['C0', 'C1', 'C3', 'C4']):
    axes[1].hist(c_comp[:, t_plot], bins=60, alpha=0.4,
                 label=f't = {t_plot}', color=color, density=True)
axes[1].set_xlabel('$c$')
axes[1].set_ylabel('density, complete markets')
axes[1].set_title('complete markets')
axes[1].legend(fontsize=9)
fig.tight_layout()
plt.show()
```

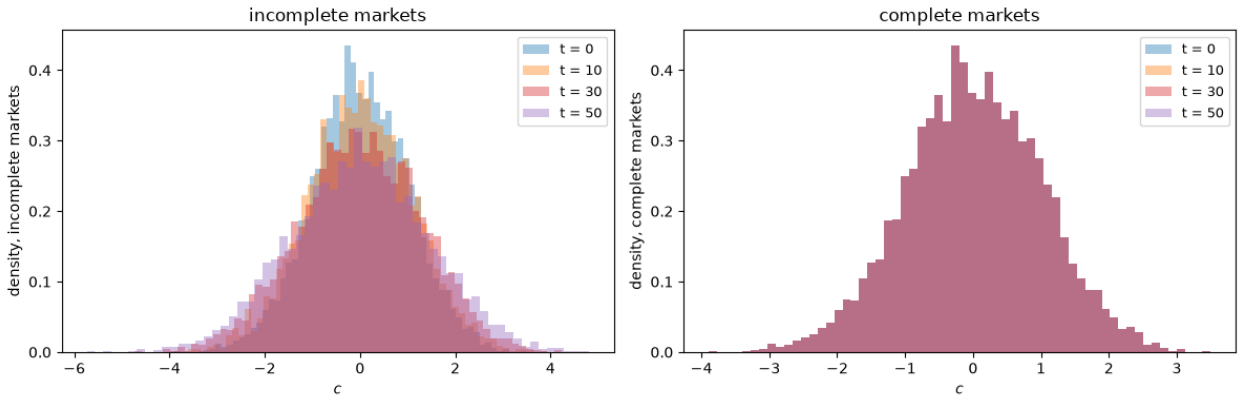


Fig. 86.3: Cross section distributions with incomplete and complete markets

Note

Under **complete markets** the histogram stays the same across all t (distributions coincide perfectly), while under **incomplete markets** distributions spread out over time.

So far the consumer fully trusts his stochastic income model.

In *Robust Consumption Smoothing and Precautionary Savings* we relax that assumption and let the consumer seek decision rules that are robust to plausible misspecifications.

A Robust LQ Bewley Model then returns to the Bewley economy studied here and shows that a cross-section of consumers with different concerns about misspecification can reproduce exactly the equilibrium of this lecture.

The optimal robust rule takes the same form as the rule above, but under a distorted model of the income process that looks more persistent than the approximating one.

86.6 Exercises

Exercise 86.6.1

This exercise studies how patience governs the speed at which the cross-section of consumption spreads out.

From (86.5), the cross-section variance of consumption grows by $hh^\top = \sigma_1^2 + (1-\beta)^2\sigma_2^2$ per period.

1. Compute this per-period growth rate for $\beta \in \{0.90, 0.95, 0.99\}$ and report the permanent and transitory contributions separately.
2. Confirm by simulation that the cross-section variance grows linearly at the predicted rate for $\beta = 0.95$.
3. Explain why increasing β slows the spreading due to transitory shocks but leaves the contribution of permanent shocks unchanged.

Solution

Here is one solution:

```
for b in (0.90, 0.95, 0.99):
    perm = σ1**2
    tran = (1 - b)**2 * σ2**2
    print(f"β = {b}:  growth = {perm + tran:.5f}  "
          f"(permanent {perm:.5f}, transitory {tran:.5f})")
```

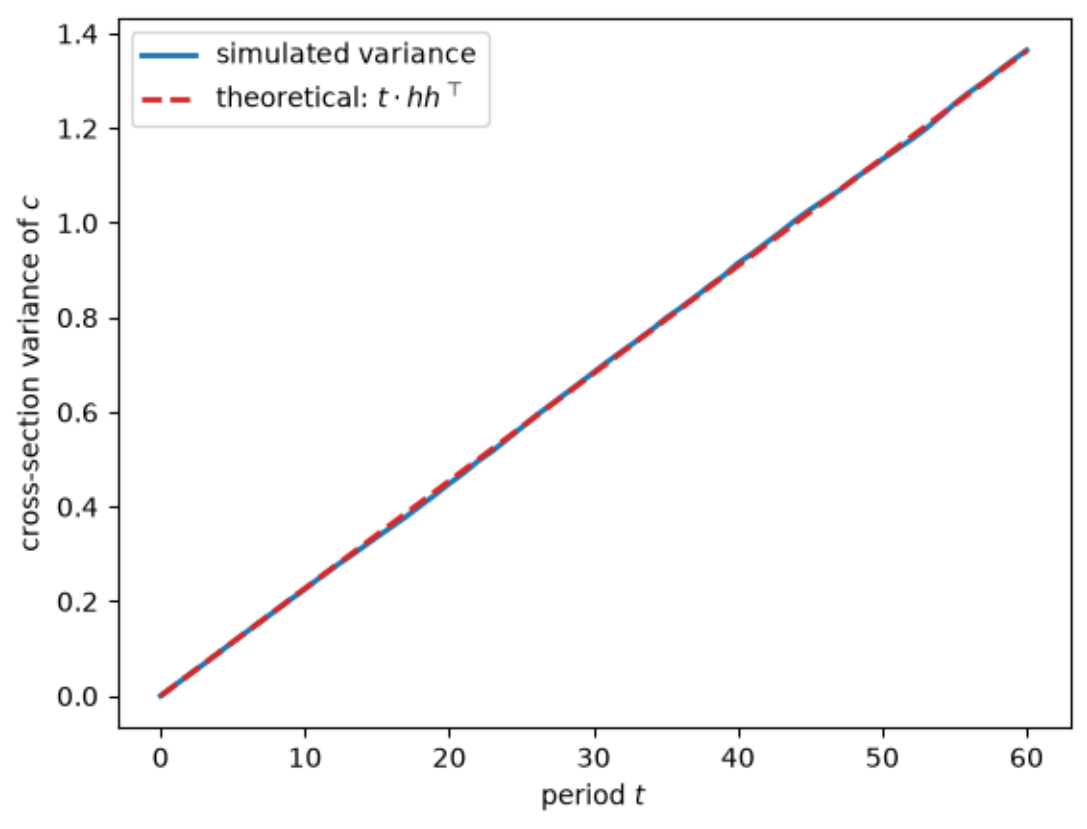
```
β = 0.9:  growth = 0.02340  (permanent 0.02250, transitory 0.00090)
β = 0.95:  growth = 0.02272  (permanent 0.02250, transitory 0.00023)
β = 0.99:  growth = 0.02251  (permanent 0.02250, transitory 0.00001)
```

```
rng = np.random.default_rng(7)
N, T_sim = 20000, 60
hh = float(h_vec @ h_vec)

c = np.zeros((N, T_sim + 1))
for t in range(T_sim):
    c[:, t + 1] = c[:, t] + rng.standard_normal((N, 2)) @ h_vec

var_c = c.var(axis=0)
theory = np.arange(T_sim + 1) * hh

fig, ax = plt.subplots()
ax.plot(var_c, lw=2, label='simulated variance')
ax.plot(theory, lw=2, linestyle='--', color='C3',
        label=r'theoretical: $t\cdot h h^\top$')
ax.set_xlabel('period $t$')
ax.set_ylabel('cross-section variance of $c$')
ax.legend()
plt.show()
```



The permanent contribution σ_1^2 does not depend on β because a permanent shock is capitalised one-for-one into consumption regardless of patience.

The transitory contribution $(1-\beta)^2\sigma_2^2$ shrinks as $\beta \to 1$ because a more patient consumer smooths a transitory shock over a longer horizon, moving consumption by only the small annuity value $(1-\beta)$.

Exercise 86.6.2

This exercise contrasts the cross-section variance of consumption under incomplete and complete markets.

Start all consumers from the initial consumption draws `c_bew_t0` computed above.

1. Compute the cross-section variance at each date $t = 0, 1, \ldots, 60$ under incomplete markets (consumption is a random walk) and under complete markets (consumption is constant).
2. Plot the two variance paths on the same axes and comment on the difference.

Solution

Here is one solution:

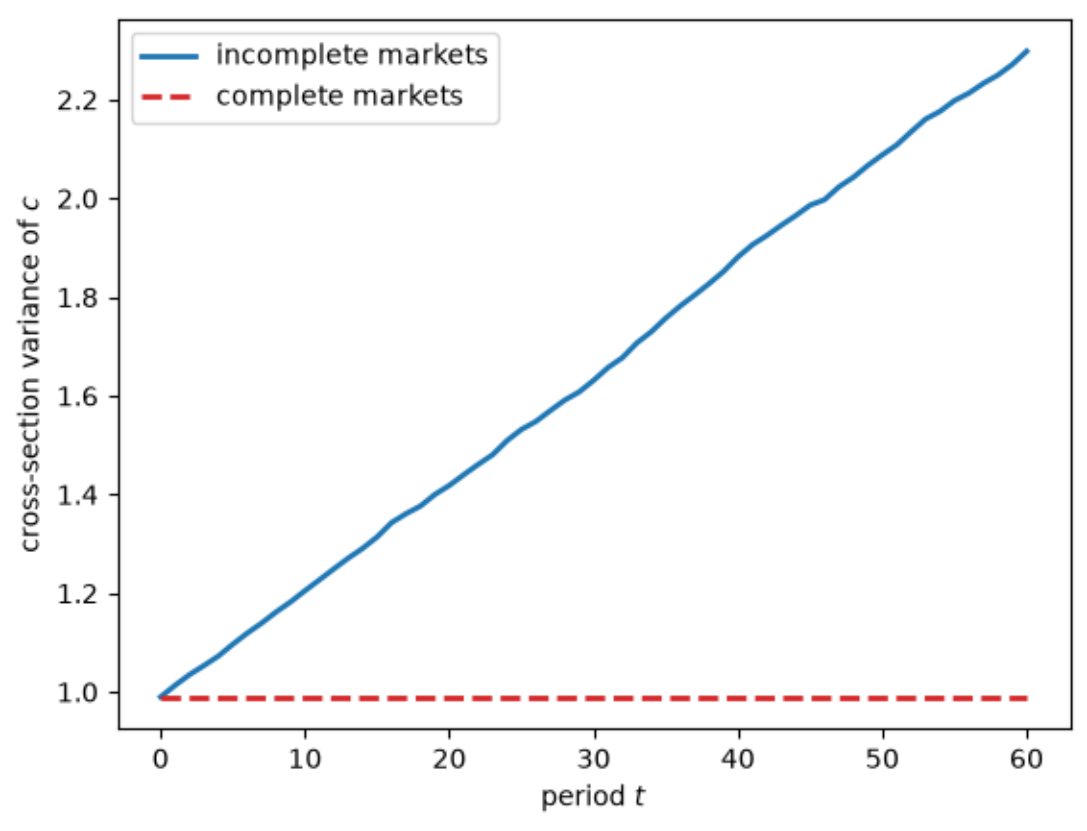
```
rng = np.random.default_rng(11)
T = 60
c0 = c_bew_t0

# Incomplete markets: random-walk consumption
```

```
c_inc = np.zeros((len(c0), T + 1))
c_inc[:, 0] = c0
for t in range(T):
    c_inc[:, t + 1] = c_inc[:, t] + rng.standard_normal((len(c0), 2)) @ h_vec

var_inc = c_inc.var(axis=0)
var_comp = np.full(T + 1, c0.var())   # complete markets: constant consumption

fig, ax = plt.subplots()
ax.plot(var_inc, lw=2, label='incomplete markets')
ax.plot(var_comp, lw=2, linestyle='--', color='C3', label='complete markets')
ax.set_xlabel('period $t$')
ax.set_ylabel('cross-section variance of $c$')
ax.legend()
plt.show()
```



Under incomplete markets the variance rises linearly without bound: each consumer accumulates an independent random walk of consumption innovations.

Under complete markets the variance is flat: each consumer locks in a constant consumption level $\bar{c}^i$, so the cross-section distribution never changes.

Exercise 86.6.3

This exercise shows how complete markets freeze each consumer's *rank* in the consumption distribution.

Under complete markets consumption is constant, so a consumer's position in the cross-section distribution never

changes; under incomplete markets the random walk gradually scrambles ranks.

1. Simulate consumption paths under both market structures, starting from `c_bew_t0`.
2. Compute the Spearman rank correlation between consumption at $t = 0$ and at later dates t under each market structure, and plot it against t.
3. Interpret the result.

Solution

Here is one solution:

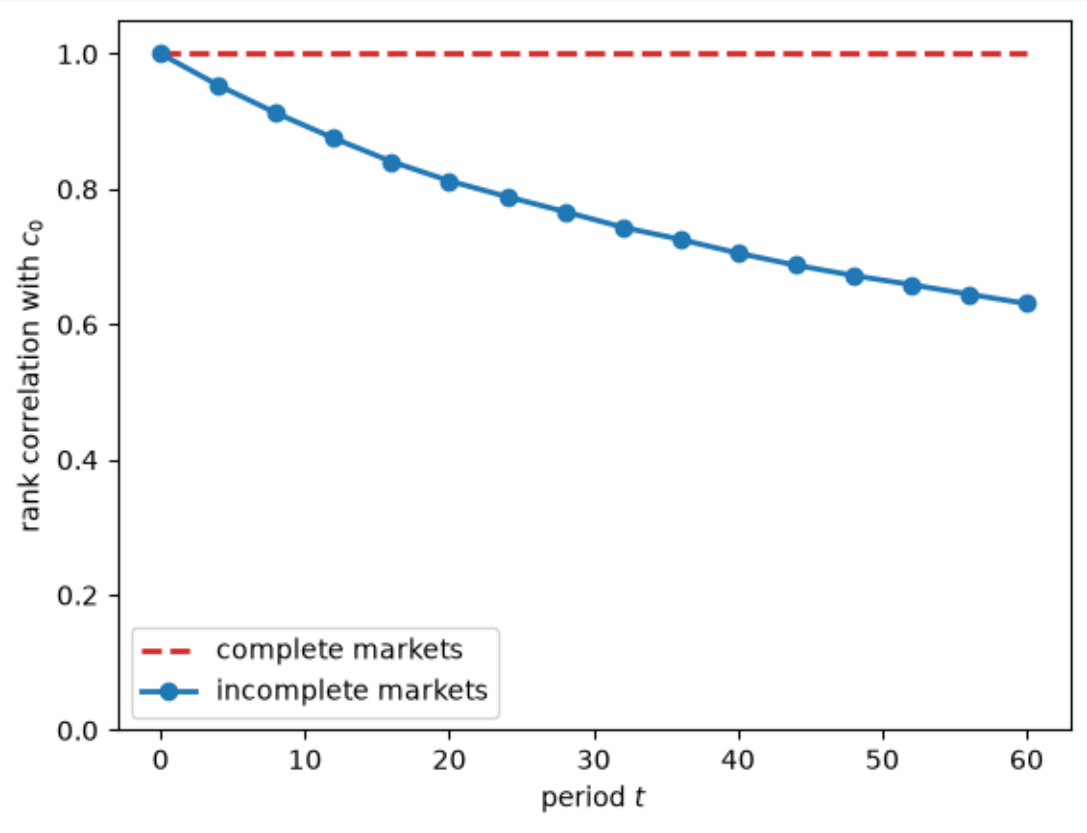
```
from scipy.stats import spearmanr

rng = np.random.default_rng(3)
T = 60
c0 = c_bew_t0

c_inc = np.zeros((len(c0), T + 1))
c_inc[:, 0] = c0
for t in range(T):
    c_inc[:, t + 1] = c_inc[:, t] + rng.standard_normal((len(c0), 2)) @ h_vec

dates = np.arange(0, T + 1, 4)
rank_inc = [spearmanr(c_inc[:, 0], c_inc[:, t]).statistic for t in dates]
rank_comp = [1.0 for _ in dates]   # constant consumption ⇒ ranks fixed

fig, ax = plt.subplots()
ax.plot(dates, rank_comp, lw=2, linestyle='--', color='C3',
        label='complete markets')
ax.plot(dates, rank_inc, lw=2, marker='o', color='C0',
        label='incomplete markets')
ax.set_xlabel('period $t$')
ax.set_ylabel(r'rank correlation with $c_0$')
ax.set_ylim(0, 1.05)
ax.legend()
plt.show()
```



Under complete markets the rank correlation stays exactly one: each consumer keeps the rank determined by the initial draw z_0^i forever.

Under incomplete markets the rank correlation decays toward zero as the accumulated random walk reshuffles who is rich and who is poor, even though every consumer faces the same income process.

CHAPTER

EIGHTYSEVEN

ROBUST CONSUMPTION SMOOTHING AND PRECAUTIONARY SAVINGS

Contents

- *Robust Consumption Smoothing and Precautionary Savings*
  - *Overview*
  - *A brief review*
  - *A robust permanent income model*
  - *Observational equivalence*
  - *The scalar worst-case model*
  - *Frequency domain interpretation*
  - *Detection error probabilities*
  - *Robustness of decision rules*
  - *Concluding remarks*
  - *Exercises*
  - *Related lectures*

87.1 Overview

This lecture studies a robust version of the LQ permanent income model due to Hansen *et al.* [1999] and Hansen and Sargent [2008].

A consumer who distrusts his specification of the labor income process engages in a form of precautionary savings.

This is the third of four lectures on the LQ permanent income model.

It builds on *The LQ Permanent Income Model*, which develops the standard model, and *Consumption Smoothing with Incomplete and Complete Markets*, which studies its cross-section and market-structure implications.

The sequel, *A Robust LQ Bewley Model*, uses the results developed here to build a Bewley economy populated by consumers who differ in how much they distrust their income model.

Our description of the model with concerns about robustness includes

- how, for quantities, a concern for robustness is observationally equivalent to an increase in impatience
- how the worst-case model that the consumer uses to shape his decision rule distorts the baseline model's endowment process toward greater persistence
- a **breakdown point** beyond which the robust control problem ceases to have a solution
- a frequency-domain representation of the effects of concerns about misspecification of the endowment process
- a detection-error-probability characterization of the amount of model uncertainty

A recurring theme is that a single scalar α^2, the variance of the innovation to the consumer's marginal utility, summarises everything about the endowment process that matters for robustness.

Let's begin with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
```

87.2 A brief review

We recall the essentials from *The LQ Permanent Income Model* and *Consumption Smoothing with Incomplete and Complete Markets*.

87.2.1 Notation

Because a robust decision maker guards against distortions to the mean of a shock, we need separate symbols for the shock and for the distortion.

We therefore adopt the following conventions, which differ in three places from the two preceding lectures.

Note

- w_{t+1} is the baseline IID shock, as in *The LQ Permanent Income Model*, and v_{t+1} is a **distortion** to its conditional mean, as in *Robust Permanent Income and Pricing*.
- $\sigma \leq 0$ is the **robustness parameter**. The standard deviations of the two endowment shocks, written σ_1 and σ_2 in the preceding lectures, are renamed η_1 and η_2 here so that σ is free.
- a_t denotes the consumer's **net assets**, equal to minus the debt b_t of *The LQ Permanent Income Model*. This frees b_t for the preference shifter of Hansen *et al.* [1999].

87.2.2 The model

A consumer with quadratic utility and discount factor β faces the endowment process

$$
\begin{aligned}
z_{t+1} &= \check{A}\, z_t + \check{C}\, w_{t+1} \\
y_t &= \check{G}\, z_t
\end{aligned}
\tag{87.1}
$$

The optimal decision rule has a state-space representation in which the state is current consumption c_t and the exogenous endowment state z_t:

$$\begin{aligned} c_{t+1} &= c_t + (1-\beta)\,\check{G}(I-\beta\check{A})^{-1}\check{C}\,w_{t+1} \\ a_t &= \frac{1}{1-\beta}\,c_t - \check{G}(I-\beta\check{A})^{-1}z_t \\ y_t &= \check{G}\,z_t \\ z_{t+1} &= \check{A}\,z_t + \check{C}\,w_{t+1} \end{aligned} \tag{87.2}$$

We again use the two-factor endowment $y_t = z_{1t} + z_{2t}$,

$$\begin{pmatrix} z_{1,t+1} \\ z_{2,t+1} \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} z_{1t} \\ z_{2t} \end{pmatrix} + \begin{pmatrix} \eta_1 & 0 \\ 0 & \eta_2 \end{pmatrix} \begin{pmatrix} w_{1,t+1} \\ w_{2,t+1} \end{pmatrix} \tag{87.3}$$

with z_{1t} a permanent component and z_{2t} a purely transitory component.

87.2.3 The consumption innovation

One scalar built from (87.2) will do all of the work below.

The first line of (87.2) says that consumption is a random walk whose innovation is $h\,w_{t+1}$, where

$$h = (1-\beta)\,\check{G}(I-\beta\check{A})^{-1}\check{C} \tag{87.4}$$

Define α^2 to be the variance of that innovation,

$$\alpha^2 = hh^\top = (1-\beta)^2\,\check{G}(I-\beta\check{A})^{-1}\check{C}\check{C}^\top(I-\beta\check{A}^\top)^{-1}\check{G}^\top \tag{87.5}$$

For the two-factor endowment (87.3) we have $\check{A} = \mathrm{diag}(1,0)$, $\check{C} = \mathrm{diag}(\eta_1,\eta_2)$ and $\check{G} = \begin{pmatrix} 1 & 1 \end{pmatrix}$, so that $(I-\beta\check{A})^{-1} = \mathrm{diag}((1-\beta)^{-1},1)$ and

$$h = \begin{pmatrix} \eta_1 & (1-\beta)\eta_2 \end{pmatrix}, \qquad \alpha^2 = \eta_1^2 + (1-\beta)^2\,\eta_2^2 \tag{87.6}$$

The permanent shock variance η_1^2 enters with coefficient 1 because a unit permanent shock is *fully* capitalised into consumption.

The transitory shock variance η_2^2 enters with the small coefficient $(1-\beta)^2$ because only its annuity value is consumed.

This scalar does triple duty across the three lectures of this suite.

Note

α^2 is simultaneously

- the variance of the consumption innovation in *The LQ Permanent Income Model*,
- the rate at which the cross-section variance of consumption grows with age in *Consumption Smoothing with Incomplete and Complete Markets*, and
- the quantity that, multiplied by σ, governs every robustness result in this lecture.

Robust Permanent Income and Pricing writes the same object as θ^2.

The following cell fixes the calibration used below.

```
β = 0.95        # discount factor, so R = 1/β
η1 = 0.15       # std of permanent shock
η2 = 0.30       # std of transitory shock

R = 1 / β
α2 = η1**2 + (1 - β)**2 * η2**2
α = np.sqrt(α2)

print(f"α^2 = {α2:.6f}")
print(f"  permanent  η1^2          = {η1**2:.6f} "
      f"({100 * η1**2 / α2:5.1f}% of α^2)")
print(f"  transitory (1-β)^2 η2^2 = {(1 - β)**2 * η2**2:.6f} "
      f"({100 * (1 - β)**2 * η2**2 / α2:5.1f}% of α^2)")
```

```
α^2 = 0.022725
  permanent  η1^2          = 0.022500 ( 99.0% of α^2)
  transitory (1-β)^2 η2^2 = 0.000225 (  1.0% of α^2)
```

Permanent shocks account for almost all of α^2 in this calibration.

87.3 A robust permanent income model

87.3.1 Robustness and precautionary savings

We now study a consumer who *distrusts* his specification of the stochastic process governing his labor income.

The model is due to Hansen *et al.* [1999] (HST), who estimated it on US quarterly consumption and investment data.

For a fuller treatment of the HST model and its asset-pricing implications, see *Robust Permanent Income and Pricing*.

A consumer who fears model misspecification engages in a form of **precautionary savings** that is distinct from the usual precautionary motive, which requires a convex marginal utility.

Here, the precautionary motive arises because the consumer wants to protect against misspecification of the **conditional means** of income shocks, and it operates even with quadratic preferences.

HST showed an important **observational equivalence** result: for quantities (c_t, i_t) alone, a concern for robustness is indistinguishable from an increase in impatience, that is, a decrease in β.

We develop this result carefully below.

87.3.2 The HST model

HST's model features a planner with preferences over consumption streams $\{c_t\}$, mediated through **service streams** $\{s_t\}$.

Let b be a preference shifter, or utility bliss point.

The **Bellman equation for the robust planner** is

$$-x^\top P x - p = \sup_c \inf_{v^*} \left\{ -(s-b)^2 + \beta\left(\theta\,(v^*)^\top v^* - \mathbb{E}\,(x^*)^\top P x^* - p\right) \right\} \tag{87.7}$$

subject to the household technology, capital accumulation, endowment dynamics, and the state law:

$$\begin{aligned} s &= (1+\lambda)c - \lambda h \\ h^* &= \delta_h h + (1-\delta_h)c \\ k^* &= \delta_k k + i \\ c + i &= \gamma k + d \\ \begin{pmatrix} d \\ b \end{pmatrix} &= Uz \\ z^* &= A_{22} z + C_2(w^* + v^*) \end{aligned} \tag{87.8}$$

Here * denotes the next-period value; c is consumption; s is the scalar service measure; h is a habit stock; k is the capital stock; i is investment; d is an endowment shock; b is a **preference shock**; γ is the marginal product of capital; $w^* \sim N(0, I)$ is the baseline shock; and v^* is a **distortion** to the conditional mean of w^* chosen by a minimizing agent.

The penalty parameter θ governs the consumer's concern about robustness.

A large θ makes distortions expensive and so restrains the minimizing agent.

We use the transformation

$$\sigma = -\theta^{-1} \leq 0, \qquad \theta \in (0, \infty] \tag{87.9}$$

so that $\sigma = 0$, equivalently $\theta = \infty$, corresponds to no robustness concern, and $\sigma < 0$ to an increasing concern.

When $\lambda > 0$ and $\delta_h \in (0, 1)$, the technology (87.8) accommodates **habit persistence** or durability, and the stock h_t is a geometric weighted average of current and past consumption.

Equation $c_t + k_t = Rk_{t-1} + d_t$ with $R = \delta_k + \gamma$ combines capital accumulation with a linear production technology, so R is the physical gross return on capital.

Let $x_t^\top = [h_{t-1},\, k_{t-1},\, z_t^\top]$.

The state transition equation is

$$x_{t+1} = A\, x_t + B\, u_t + C(w_{t+1} + v_{t+1}) \tag{87.10}$$

where $u_t = c_t$ and v_{t+1} is the distortion to the conditional mean of w_{t+1}.

HST estimated the model on US quarterly data from 1970Q1 to 1996Q3, using nondurables plus services for consumption and durable consumption plus gross private investment for investment.

They imposed $\beta R = 1$ and $\delta_k = 0.975$, so γ is pinned down once β is estimated.

Two of their preference estimates are worth recording.

Parameter	Habit	No habit
β	0.997	0.997
δ_h	0.682	—
λ	2.443	0
$2 \times \log L$	779.05	762.55

At a quarterly frequency, $\beta = 0.997$ implies an annual real interest rate of $\beta^{-4} - 1 \approx 1.2\%$.

The remaining estimated parameters govern the exogenous d_t and b_t processes and are reported in Appendix A of Hansen *et al.* [1999].

87.3.3 Solution when $\sigma = 0$

When $\sigma = 0$ the objective reduces to

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \{-(s_t - b_t)^2\} \tag{87.11}$$

Forming a Lagrangian and deriving first-order conditions yields

$$\begin{aligned} \mu_{st} &= b_t - s_t \\ \mu_{ct} &= (1+\lambda)\mu_{st} + (1-\delta_h)\mu_{ht} \\ \mu_{ht} &= \beta \mathbb{E}_t[\delta_h \mu_{h,t+1} - \lambda \mu_{s,t+1}] \\ \mu_{ct} &= \beta R \, \mathbb{E}_t \mu_{c,t+1} \end{aligned} \tag{87.12}$$

Here μ_{st} is the **marginal valuation of consumption services**, which summarises the endogenous state variables h_{t-1} and k_{t-1}.

The last line of (87.12) implies $\mathbb{E}_t \mu_{c,t+1} = (\beta R)^{-1} \mu_{ct}$, so μ_{st} is a martingale when $\beta R = 1$:

$$\mu_{st} = \mu_{s,t-1} + \nu^\top w_t \tag{87.13}$$

for some vector ν.

Solving forward and substituting gives

$$\mu_{st} = \Psi_1 k_{t-1} + \Psi_2 h_{t-1} + \Psi_3 \sum_{j=0}^{\infty} R^{-j} \mathbb{E}_t b_{t+j} + \Psi_4 \sum_{j=0}^{\infty} R^{-j} \mathbb{E}_t d_{t+j} \tag{87.14}$$

where

$$\Psi_1 = -(1+\lambda) R (1 - R^{-2}\beta^{-1}) \left[\frac{1 - R^{-1}\tilde{\delta}_h}{1 - R^{-1}\tilde{\delta}_h + \lambda(1 - \tilde{\delta}_h)} \right], \quad \Psi_4 = R^{-1}\Psi_1 \tag{87.15}$$

and $\tilde{\delta}_h = (\delta_h + \lambda)/(1+\lambda)$.

In the widely-studied special case $\lambda = \delta_h = 0$, we have $s_t = c_t$ and $\mu_{st} = b_t - c_t$, and the marginal propensity to consume out of **non-human wealth** $R k_{t-1}$ equals that out of **human wealth** $\sum_{j=0}^{\infty} R^{-j} \mathbb{E}_t d_{t+j}$, a well-known feature of the LQ model.

The formula for μ_{st} can be written as $\mu_{st} = M_s x_t$ where x_t follows (87.10).

It follows that

$$\nu^\top = M_s C, \qquad \alpha = \sqrt{\nu^\top \nu} = \sqrt{M_s C C^\top M_s^\top} \tag{87.16}$$

This α is the same scalar we met in (87.5).

To see why, set $\lambda = \delta_h = 0$ and hold b_t fixed, so that $\mu_{st} = b - c_t$ and the innovation to μ_{st} is minus the innovation to c_t.

Hence $\nu^\top = -h$ and $\alpha^2 = \nu^\top \nu = h h^\top$, exactly as in (87.5).

The sign is irrelevant because only α^2 ever appears.

87.4 Observational equivalence

HST state an observational-equivalence theorem.

Theorem 87.4.1 (Observational equivalence, I)

Fix all parameters except (σ, β) and suppose $\beta R = 1$ when $\sigma = 0$.

There exists $\underline{\sigma} < 0$ such that for any $\sigma \in (\underline{\sigma}, 0)$, the optimal consumption-investment plan for $(0, \beta)$ is also chosen by a robust decision maker with parameters $(\sigma, \hat{\beta}(\sigma))$, where

$$\hat{\beta}(\sigma) = \frac{1}{R} + \frac{\sigma\alpha^2}{R-1} = \beta + \frac{\sigma\alpha^2\beta}{1-\beta} \tag{87.17}$$

and $\hat{\beta}(\sigma) < \beta$.

The second equality in (87.17) uses $R = \beta^{-1}$ and will be the form we use in computations.

Since $R > 1$ and $\alpha^2 > 0$, a more negative σ, meaning a stronger robustness concern, lowers $\hat{\beta}$.

A robust consumer wants to save more because his alter ego, a utility-minimizing agent, makes future income look worse than the approximating model predicts.

A lower discount factor makes a consumer less patient and therefore reduces saving.

When these two forces are balanced according to (87.17), consumption plans are identical across $(\sigma, \hat{\beta}(\sigma))$ pairs.

Proof. When $\beta R = 1$ and $\sigma = 0$, the marginal utility μ_{st} obeys the martingale

$$\mu_{st} = \mu_{s,t-1} + \alpha \tilde{w}_t \tag{87.18}$$

where $\tilde{w}_t$ is scalar IID with mean zero and unit variance.

Activating a concern about robustness, $\sigma < 0$, leads the utility-minimizing alter ego to set

$$\tilde{v}_t = K(\sigma, \hat{\beta})\, \mu_{s,t-1} \tag{87.19}$$

making the worst-case model for μ_{st}

$$\mu_{st} = \zeta\, \mu_{s,t-1} + \alpha\, \tilde{w}_t, \qquad \zeta \equiv 1 + \alpha\, K(\sigma, \hat{\beta}) \tag{87.20}$$

For the allocation to remain the same, we require the robust Euler equation $\hat{\beta} R\, \hat{\mathbb{E}}_t \mu_{s,t+1} = \mu_{st}$ to hold under the worst-case model, which gives

$$\zeta = (\hat{\beta} R)^{-1} \tag{87.21}$$

The minimizing agent's Bellman equation, a pure forecasting problem, yields

$$\zeta = \frac{1}{1 - \sigma \alpha^2 P(\hat{\beta})} \tag{87.22}$$

where $P(\hat{\beta})$ solves the scalar Bellman equation

$$P(\hat{\beta}) = \frac{\hat{\beta} - 1 + \sigma\alpha^2 + \sqrt{(\hat{\beta} - 1 + \sigma\alpha^2)^2 + 4\sigma\alpha^2}}{-2\sigma\alpha^2} \tag{87.23}$$

Solving (87.21)-(87.23) for $\hat{\beta}$ gives exactly (87.17).

Equation (87.17) is the useful numerical object because it gives a straight-line map from the robustness parameter to the observationally equivalent discount factor.

87.4.1 Precautionary savings interpretation

The consumer's concern about model misspecification activates the precautionary savings motive that underlies the observational-equivalence theorem.

A concern about robustness makes the consumer save *more*.

Decreasing β makes the consumer save *less*.

The observational-equivalence theorem says that these two forces can be made to offset each other exactly.

In the special case $\lambda = \delta_h = 0$, $s_t = c_t$ and the consumption rule is

$$c_t = (1 - R^{-2}\beta^{-1}) \left[R k_{t-1} + \mathbb{E}_t \sum_{j=0}^{\infty} R^{-j} d_{t+j} \right] + \left(\frac{(R\beta)^{-1} - 1}{R - 1} \right) b \tag{87.24}$$

The **marginal propensity to consume** out of non-human wealth $R k_{t-1}$ *equals* that out of human wealth $\mathbb{E}_t \sum R^{-j} d_{t+j}$.

This equal-propensity property is a hallmark of the LQ model and *persists* when a concern for robustness is present, in contrast to usual precautionary-savings models with convex marginal utility.

Theorem 87.4.1 says that with $\sigma < 0$, the observationally equivalent $\hat{\beta}$ satisfies $\hat{\beta} < \beta$.

If the starting point has $\beta R = 1$, then $\hat{\beta} R < 1$.

For a non-robust consumer with discount factor $\hat{\beta}$ at the same interest rate, the Euler equation implies $\mathbb{E}_t c_{t+1} < c_t$, so expected consumption declines over time.

This downward drift is the impatience offset in *Theorem 87.4.1*.

It cancels the robust consumer's precautionary-savings motive, leaving the consumption and investment quantities unchanged.

The classical precautionary motive arises because

$$u'''(c) > 0 \;\Rightarrow\; \mathbb{E}_t u'(c_{t+1}) > u'(\mathbb{E}_t c_{t+1}) \;\Rightarrow\; \mathbb{E}_t c_{t+1} > c_t \tag{87.25}$$

This channel requires *convexity of marginal utility* and is absent with quadratic preferences.

In contrast, the robustness-based precautionary motive operates through distortions of **conditional means** of shocks, shifting the first moment of the innovation to non-financial income.

87.4.2 Observational equivalence and distorted expectations

The observational-equivalence result can be interpreted using a **Stackelberg multiplier game**.

After the minimizing agent has committed to a distortion process $\{v_{t+1}\}$, the maximizing consumer faces the following worst-case law of motion for the state X_t:

$$\begin{aligned} X_{t+1} &= (A - BF(\sigma, \hat{\beta}) + CK(\sigma, \hat{\beta}))X_t + C\, w_{t+1} \\ \begin{pmatrix} b_t \\ d_t \end{pmatrix} &= S X_t \end{aligned} \tag{87.26}$$

A robust consumer forms expectations of future income using the **distorted transition matrix** $A - BF + CK$ rather than the approximating transition matrix $A - BF$.

The distorted expectations operator $\hat{\mathbb{E}}_t$ satisfies

$$\hat{\mathbb{E}}_t X_{t+j} = (A - BF(\sigma, \hat{\beta}) + CK(\sigma, \hat{\beta}))^j X_t \tag{87.27}$$

Observational equivalence requires that the modified human-wealth formula

$$\hat{\Psi}_4 \sum_{j=0}^{\infty} R^{-j} \hat{\mathbb{E}}_t d_{t+j} \tag{87.28}$$

equals its benchmark counterpart $\Psi_4 \sum_{j=0}^{\infty} R^{-j} \mathbb{E}_t d_{t+j}$.

This is achieved by a mutual adjustment of the coefficients $\hat{\Psi}_j$ through $\hat{\beta}$ and of the distorted expectation operator $\hat{\mathbb{E}}_t$ through σ.

The worst-case eigenvalue of $A - BF + CK$ exceeds that of $A - BF$ in modulus, so the worst-case distortions make the income process *more persistent* than under the approximating model.

This is the precautionary motive in state-space form: the minimizing agent makes future income look more risky by introducing low-frequency persistence.

In the scalar reduction of the next section, this eigenvalue is ζ, and we verify that $\zeta > 1$ while the approximating model has a unit root.

87.4.3 Another observational equivalence result

Theorem 87.4.2 (Observational equivalence, II)

Fix all parameters except (σ, β) and consider a consumption-investment allocation for $(\hat{\sigma}, \hat{\beta})$ where $\hat{\beta}R = 1$ and $\hat{\sigma} < 0$.

Then there exists $\tilde{\beta} > \hat{\beta}$ such that the $(\hat{\sigma}, \hat{\beta})$ allocation also solves the $(0, \tilde{\beta})$ problem.

Theorem 87.4.1 showed that starting from a benchmark with $\beta R = 1$, activating robustness is equivalent to *reducing* β.

Theorem 87.4.2 goes in the opposite direction: the effects of activating a concern for robustness from a starting point with $\beta R = 1$ are replicated by *increasing* β while setting $\sigma = 0$.

In other words, when $\beta R = 1$, a concern for robustness operates like an *increase* in the discount factor, pushing $\beta R > 1$ and imparting an *upward drift* to the expected consumption profile.

Proof. With $\hat{\beta}R = 1$ and $\hat{\sigma} < 0$, the robust Euler equation implies

$$\hat{\mathbb{E}}_t \mu_{c,t+1} = \mu_{ct} \tag{87.29}$$

One seeks $\tilde{\beta} > \hat{\beta}$ and $\sigma = 0$ such that the same allocation solves the non-robust problem with discount factor $\tilde{\beta}$.

The key step is to observe that the worst-case distortion $K(\hat{\sigma}, \hat{\beta})$ introduces a drift in the marginal utility process that is equivalent to the drift produced by raising the discount factor above $\hat{\beta}$.

Equating the two drifts and solving the scalar Bellman equation for K yields

$$\tilde{\beta}(\hat{\sigma}) = \frac{\hat{\beta}(1+\hat{\beta})}{2(1+\hat{\sigma}\alpha^2)} \left[1 + \sqrt{1 - 4\hat{\beta}\,\frac{1+\hat{\sigma}\alpha^2}{(1+\hat{\beta})^2}} \right] \tag{87.30}$$

Setting $\hat{\sigma} = 0$ makes the square root equal $(1-\hat{\beta})/(1+\hat{\beta})$, so that $\tilde{\beta} = \hat{\beta}$.

Since $1 + \hat{\sigma}\alpha^2$ decreases as $\hat{\sigma}$ falls below zero, both the prefactor and the square root increase, so $\tilde{\beta} > \hat{\beta}$ whenever $\hat{\sigma} < 0$.

87.4.4 Comparing the two loci

Both (87.17) and (87.30) are closed forms, so we can plot them directly.

We start from a benchmark with $\beta R = 1$.

```
σ_grid = np.linspace(0.0, -0.16, 60)

# locus I: non-robust agent fixed at βR=1, report the robust twin's β̂(σ)
β_hat = β + σ_grid * α2 * β / (1 - β)

# locus II: robust agent fixed at βR=1, report the non-robust twin's β̃(σ̂)
q = 1 + σ_grid * α2
β_tilde = β * (1 + β) / (2 * q) * (1 + np.sqrt(1 - 4 * β * q / (1 + β)**2))
```

(continues on next page)

(continued from previous page)

```
fig, ax = plt.subplots()
ax.plot(-σ_grid, β_hat, lw=2, color='C3',
        label=r'locus I: robust twin $\hat\beta(\sigma) < \beta$')
ax.plot(-σ_grid, β_tilde, lw=2, color='C0',
        label=r'locus II: non-robust twin $\tilde\beta(\hat\sigma) > \beta$')
ax.axhline(β, color='k', linestyle=':', lw=1,
           label=r'benchmark $\beta$ ($\beta R = 1$)')
ax.set_xlabel(r'robustness concern $|\sigma|$')
ax.set_ylabel('discount factor of the equivalent agent')
ax.legend()
plt.show()
```

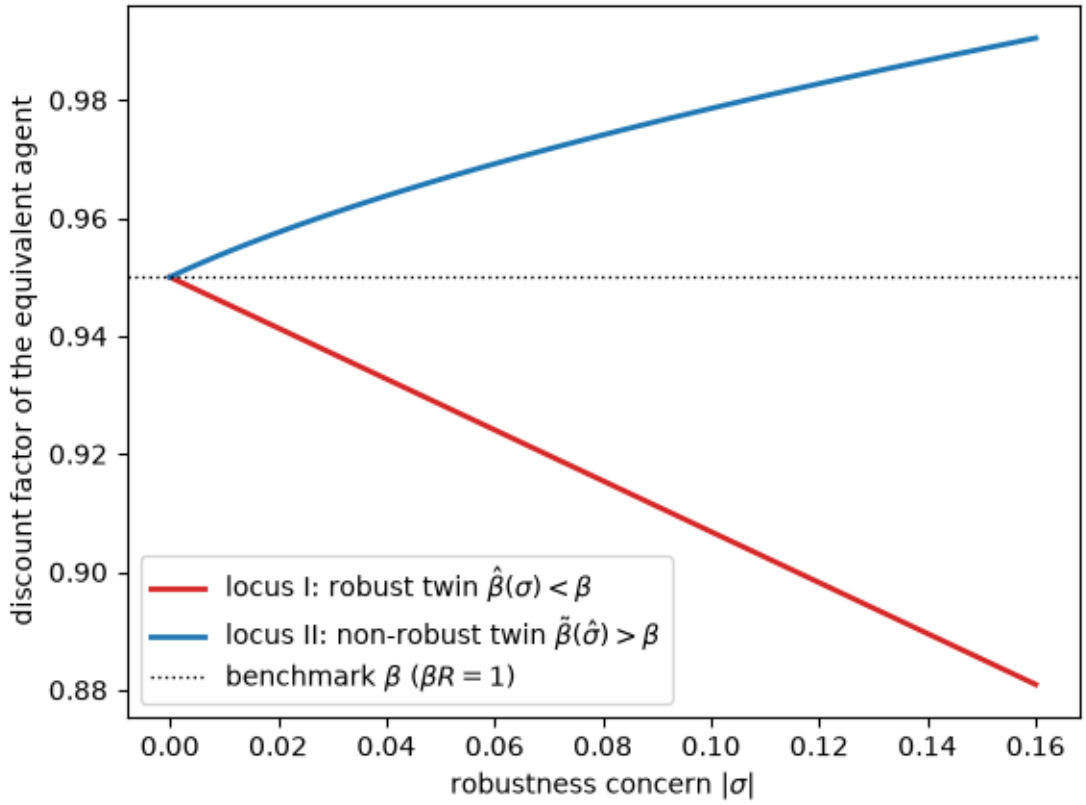


Fig. 87.1: Two observational-equivalence experiments. Locus I holds the *non-robust* agent fixed at $\beta R = 1$ and reports the *robust* twin's discount factor $\hat\beta(\sigma)$; locus II holds the *robust* agent fixed at $\beta R = 1$ and reports the *non-robust* twin's discount factor $\tilde\beta(\hat\sigma)$.

Both loci pass through the benchmark β at $\sigma = 0$ and separate as the robustness concern grows.

The key to reading Fig. 87.1 is that the two loci hold *different* agents fixed, so the discount factor plotted on the vertical axis refers to a different agent on each curve.

Locus I, from *Theorem 87.4.1*, holds the **non-robust** agent fixed at the benchmark $(\sigma = 0, \beta)$ with $\beta R = 1$ and reports the discount factor $\hat\beta(\sigma) < \beta$ of the **robust** agent that mimics it.

This is the sense in which HST call a concern for robustness observationally equivalent to a *lower* discount factor: because robustness already makes the agent save more, its discount factor must be lowered to hold the allocation at the benchmark.

Because the non-robust benchmark has $\beta R = 1$, its optimal consumption is a martingale, $\mathbb{E}_t c_{t+1} = c_t$.

The robust twin chooses the identical consumption process, so it too satisfies $\mathbb{E}_t c_{t+1} = c_t$.

The lower $\hat{\beta}$, which has $\hat{\beta} R < 1$, would on its own impart a downward drift, but the robust agent's precautionary saving offsets it exactly, leaving expected consumption flat.

Locus II, from *Theorem 87.4.2*, instead holds the **robust** agent fixed at $(\hat{\sigma}, \beta)$ with $\beta R = 1$ and reports the discount factor $\tilde{\beta}(\hat{\sigma}) > \beta$ of the **non-robust** agent that mimics it.

Here there is no impatience offset, so the common allocation inherits the robust agent's precautionary *upward* drift, which the non-robust twin reproduces through $\tilde{\beta} R > 1$.

The two experiments encode the *same* economics: a concern for robustness adds precautionary saving that acts like extra patience.

They differ only in which agent is anchored at $\beta R = 1$, and hence in whether the common saving motive shows up as an exactly-offsetting impatience adjustment, as in locus I where expected consumption is flat, or as an upward drift in expected consumption, as in locus II.

87.5 The scalar worst-case model

The proof of *Theorem 87.4.1* reduced the robust problem to a scalar forecasting problem in the marginal utility μ_{st}.

That scalar problem can be solved in closed form, which is worth doing because it makes the worst-case dynamics, the breakdown point, and the frequency-domain results all transparent.

87.5.1 A closed-form solution

Combining (87.21) with (87.17) and $R = \beta^{-1}$ gives the worst-case persistence directly:

$$\zeta(\sigma) = \frac{1}{\hat{\beta}(\sigma) R} = \frac{\beta}{\hat{\beta}(\sigma)} = \left[1 + \frac{\sigma \alpha^2}{1 - \beta}\right]^{-1} \tag{87.31}$$

This is the central formula of the lecture.

The worst-case persistence of marginal utility is exactly the ratio of the two discount factors.

Since $\hat{\beta}(\sigma) < \beta$ for $\sigma < 0$, we have $\zeta(\sigma) > 1$: the approximating model for μ_{st} has a unit root, and the worst-case model is mildly explosive.

That is the scalar counterpart of the statement in (87.26) that the worst-case transition matrix has a larger eigenvalue than the approximating one.

We can also solve (87.23) explicitly.

Write $u = \sigma \alpha^2$ and $\delta = 1 - \beta$, so that (87.17) reads $\hat{\beta} - 1 + u = (u - \delta^2)/\delta$.

The discriminant in (87.23) is then a **perfect square**:

$$(\hat{\beta} - 1 + u)^2 + 4u = \frac{(u - \delta^2)^2}{\delta^2} + 4u = \left(\frac{u + \delta^2}{\delta}\right)^2 \tag{87.32}$$

so the two roots of (87.23) are available in closed form:

$$P = -\frac{1}{1 - \beta} \qquad \text{and} \qquad P = \frac{1 - \beta}{\sigma \alpha^2} \tag{87.33}$$

Substituting into (87.22), the first root reproduces (87.31) while the second gives the constant $\zeta = R$, which violates the Euler equation (87.21) except at the single point where the two roots coincide.

So the economically relevant root is the constant $P = -(1-\beta)^{-1}$, independent of σ.

Note

Selecting the root numerically, by taking whichever of the two is closer to the target $(\hat{\beta}R)^{-1}$, is treacherous.

Because (87.32) is a perfect square, the square root is $|u+\delta^2|/\delta$, and the *sign* in front of it that selects $P = -(1-\beta)^{-1}$ flips as u crosses $-\delta^2$.

A solver that picks the root by a distance criterion will silently switch branches there.

87.5.2 The breakdown point

Theorem 87.4.1 asserts the existence of a lower bound $\underline{\sigma} < 0$ without saying what it is.

The scalar model tells us.

The minimizing agent's problem has a finite value only if the discounted worst-case state is square summable, that is, only if $\hat{\beta}\,\zeta^2 < 1$.

Using $\hat{\beta} = \beta/\zeta$ from (87.31), this condition is $\beta\zeta < 1$, or equivalently $\zeta < R$.

Substituting (87.31) and solving gives the **breakdown point**

$$\underline{\sigma} = -\frac{(1-\beta)^2}{\alpha^2} \tag{87.34}$$

At $\sigma = \underline{\sigma}$ three things happen at once.

- The discriminant (87.32) has a double root, so the two roots in (87.33) coincide.
- The worst-case persistence reaches $\zeta = R$, so $\hat{\beta}\zeta^2 = 1$ exactly.
- The observationally equivalent discount factor reaches $\hat{\beta} = \beta^2$.

For $\sigma < \underline{\sigma}$ the robust control problem has no solution, and any numerical answer reported there is meaningless.

We therefore restrict every plot below to $\sigma \in (\underline{\sigma}, 0]$.

87.5.3 Verifying the closed form

The following cell solves the quadratic (87.23) numerically and checks it against the closed forms (87.31) and (87.33).

```
def worst_case_persistence(σ, β, α2):
    """
    Worst-case persistence ζ(σ) of marginal utility on the
    observational-equivalence locus, from eq:rcs-zeta.
    """
    return 1 / (1 + σ * α2 / (1 - β))


def solve_scalar_riccati(σ, β, α2):
    """
    Solve the scalar Bellman equation eq:rcs-riccati by brute force and
    return both roots together with the implied persistence ζ = 1/(1-σα²P).
```

(continues on next page)

(continued from previous page)

```
    """
    β_hat = β + σ * α2 * β / (1 - β)
    u = σ * α2
    disc = (β_hat - 1 + u)**2 + 4 * u
    roots = np.array([(β_hat - 1 + u + s * np.sqrt(disc)) / (-2 * u)
                      for s in (1.0, -1.0)])
    return roots, 1 / (1 - u * roots)


σ_lo = -(1 - β)**2 / α2               # breakdown point, eq:rcs-breakdown
print(f"breakdown point  σ̲  = {σ_lo:.6f}")
print(f"there            β  = {β + σ_lo * α2 * β / (1 - β):.6f} "
      f"(β² = {β**2:.6f})")
print(f"                 ζ   = {worst_case_persistence(σ_lo, β, α2):.6f} "
      f"(R = {R:.6f})")

print(f"\n{'σ':>10}{'P (numerical roots)':>28}{'-1/(1-β)':>12}"
      f"{'ζ (num)':>12}{'ζ (closed)':>12}")
for σ in [-0.02, -0.05, -0.09, -0.105]:
    roots, ζs = solve_scalar_riccati(σ, β, α2)
    keep = np.argmin(np.abs(roots + 1 / (1 - β)))
    print(f"{σ:10.3f}{str(np.round(roots, 4)):>28}{-1 / (1 - β):12.4f}"
          f"{ζs[keep]:12.6f}{worst_case_persistence(σ, β, α2):12.6f}")
```

```
breakdown point  σ̲  = -0.110011
there            β  = 0.902500 (β² = 0.902500)
                 ζ   = 1.052632 (R = 1.052632)

         σ          P (numerical roots)    -1/(1-β)      ζ (num)   ζ (closed)
    -0.020          [ -20.    -110.011]    -20.0000    1.009173    1.009173
    -0.050          [-20.     -44.0044]    -20.0000    1.023253    1.023253
    -0.090          [-20.     -24.4469]    -20.0000    1.042650    1.042650
    -0.105          [-20.     -20.9545]    -20.0000    1.050114    1.050114
```

One root is pinned at $-(1-\beta)^{-1} = -20$ for every σ, exactly as (87.33) predicts, and the implied ζ agrees with the closed form to displayed precision.

Notice also that the two roots approach each other as σ falls toward $\underline{\sigma} \approx -0.11$.

The next figure plots the worst-case impulse response ζ^h over the admissible range.

```
horizons = np.arange(31)

fig, ax = plt.subplots()
for frac in [0.0, 0.4, 0.8, 0.98]:
    σ = frac * σ_lo
    ζ = worst_case_persistence(σ, β, α2)
    ax.plot(horizons, ζ**horizons, lw=2,
            label=rf'$\sigma={σ:.3f}$, $\zeta={ζ:.3f}$')

ax.set_xlabel('horizon $h$')
ax.set_ylabel(r'response of $\mu_{s,t+h}$')
ax.legend()
plt.show()
```

Fig. 87.2 shows that as σ falls, the minimizing agent converts the unit root of the approximating model into mild explo-

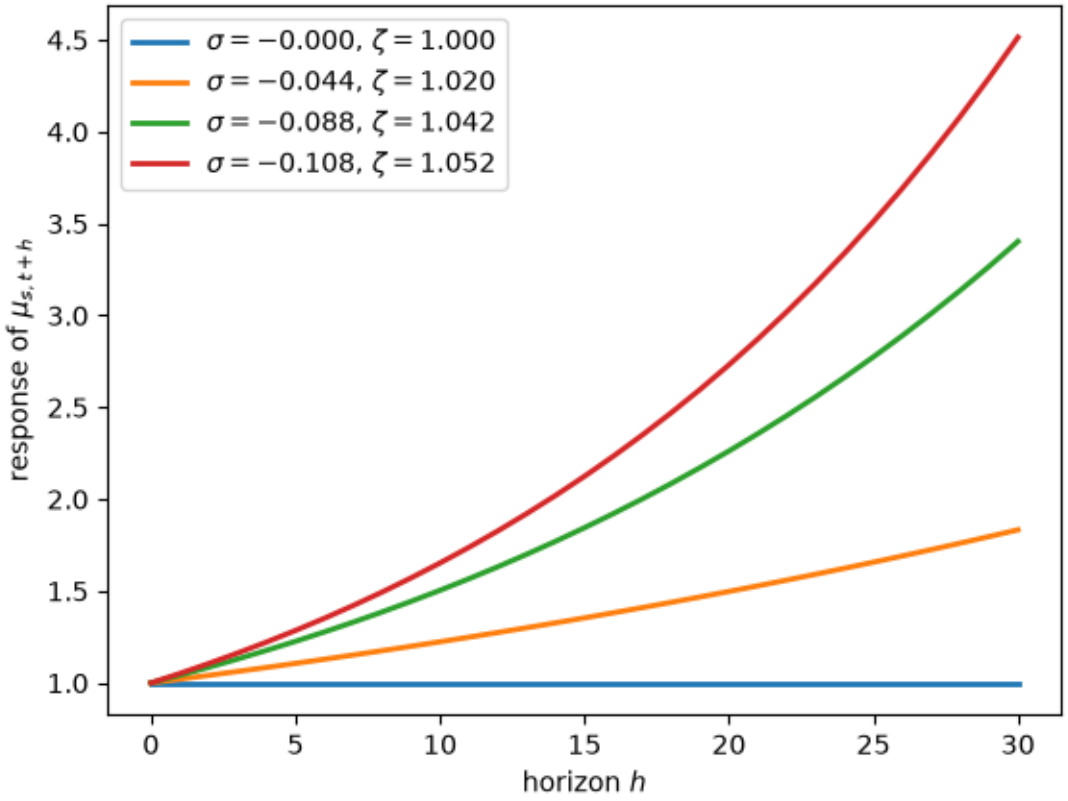


Fig. 87.2: Worst-case impulse response of marginal utility. The approximating model has a unit root; the worst-case model is increasingly explosive as σ falls toward the breakdown point $\underline{\sigma}$.

siveness.

At the breakdown point the response would grow at the gross interest rate R, so that the discounted worst-case state ceases to be square summable.

87.6 Frequency domain interpretation

The LQ permanent income framework has a natural frequency-domain interpretation.

The consumer's concave utility makes him dislike **high-frequency** fluctuations in consumption, which he smooths by adjusting savings.

High-frequency fluctuations are easy to smooth, so the consumer is automatically robust to misspecification of high-frequency features of the income process.

Low-frequency fluctuations are harder to smooth because they are more persistent.

In the frequency-domain notation of HST, the transfer function from shocks w_t to the target $s_t - b_t$ is $G(\cdot)$, and the frequency decomposition of the H_2 criterion is

$$H_2 = -\frac{1}{2\pi}\int_{-\pi}^{\pi} \operatorname{trace}\!\left[G(\sqrt{\beta}\,e^{i\omega})^\top G(\sqrt{\beta}\,e^{i\omega})\right] d\omega \tag{87.35}$$

The evaluation at $\sqrt{\beta}\,e^{i\omega}$ rather than at $e^{i\omega}$ is essential and not merely conventional.

Both the approximating model and the worst-case model for μ_{st} are non-stationary, the first with a unit root and the second explosive, so neither has an ordinary spectral density.

Discounting by $\sqrt{\beta}$ is exactly what makes the object in (87.35) finite.

In the scalar model of *The scalar worst-case model*, the target is $s_t - b_t = -\mu_{st}$, and from (87.20) we can read off the transfer function and the discounted spectral density

$$G(z) = \frac{\alpha}{1-\zeta z}, \qquad S(\omega;\sigma) = \left|G(\sqrt{\hat\beta}\,e^{i\omega})\right|^2 = \frac{\alpha^2}{\left|1-\zeta\sqrt{\hat\beta}\,e^{i\omega}\right|^2} \tag{87.36}$$

This is finite precisely when $\zeta\sqrt{\hat\beta} < 1$, which is the condition $\hat\beta\zeta^2 < 1$ that defines the breakdown point (87.34).

So the frequency-domain object and the breakdown point are two views of the same restriction.

At $\sigma = 0$ we have $\zeta = 1$ and $\hat\beta = \beta$, and (87.36) reduces to the discounted spectral density of a random walk.

```
ω = np.linspace(0, np.pi, 400)


def spectrum(σ, β, α2):
    ζ = worst_case_persistence(σ, β, α2)
    β_hat = β / ζ
    return α2 / np.abs(1 - ζ * np.sqrt(β_hat) * np.exp(1j * ω))**2


S0 = spectrum(0.0, β, α2)

fig, axes = plt.subplots(1, 2, figsize=(11, 4))

for frac in [0.0, 0.4, 0.8, 0.95]:
    σ = frac * σ_lo
```

(continues on next page)

(continued from previous page)

```
    S = spectrum(σ, β, α2)
    axes[0].plot(ω, S, lw=2, label=rf'$\sigma={σ:.3f}$')
    axes[1].plot(ω, S / S0, lw=2, label=rf'$\sigma={σ:.3f}$')

axes[0].set_yscale('log')
axes[0].set_xlabel(r'frequency $\omega$')
axes[0].set_ylabel(r'$S(\omega;\sigma)$')
axes[0].set_title('discounted spectral density')
axes[0].legend()

axes[1].set_yscale('log')
axes[1].set_xlabel(r'frequency $\omega$')
axes[1].set_ylabel(r'$S(\omega;\sigma)\,/\,S(\omega;0)$')
axes[1].set_title('relative to the approximating model')
axes[1].legend()

fig.tight_layout()
plt.show()
```

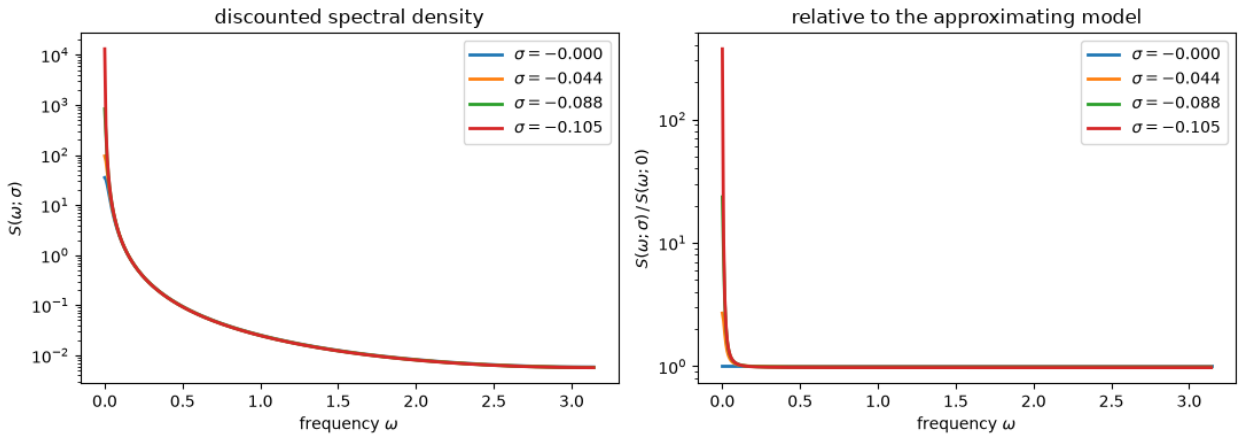


Fig. 87.3: Left: the discounted spectral density of the robust consumer's target. Right: the same densities relative to the approximating model. The minimizing agent concentrates its distortion at low frequencies, where the permanent income consumer is least able to smooth.

The left panel of Fig. 87.3 shows that $S(\omega;\sigma)$ is largest at $\omega \approx 0$, where the consumer's welfare is most sensitive to income variability.

The right panel makes the key point: the ratio $S(\omega;\sigma)/S(\omega;0)$ is sharply decreasing in ω.

Recognizing where the consumer is vulnerable, the minimizing agent concentrates the worst-case distortions at low frequencies, and it does so more aggressively as $|\sigma|$ grows.

The peak at $\omega = 0$ equals $\alpha^2/(1-\sqrt{\beta\zeta})^2$ and diverges as $\sigma \to \underline{\sigma}$, which is another way of seeing the breakdown point.

Because the distortion is $v_t = K\mu_{s,t-1}$ by (87.19), the spectral density of the distortion process is K^2 times the density of $\mu_{s,t-1}$, so it inherits the same low-frequency concentration.

87.7 Detection error probabilities

A natural way to discipline the choice of σ is to ask: **how difficult would it be to distinguish the approximating model from the worst-case model statistically?**

For a sample of length T, one can use a **log-likelihood ratio test** to compare the two hypotheses.

The **detection error probability** (DEP) is the probability of making the wrong decision using the log-likelihood ratio statistic when one does not know which model generated the data:

$$\text{DEP}(\sigma) = \frac{1}{2}\big[\mathbb{P}\{\text{prefer approximating} \mid \text{worst-case is true}\} + \mathbb{P}\{\text{prefer worst-case} \mid \text{approximating is true}\}\big] \tag{87.37}$$

When $\sigma = 0$ the two models are identical and DEP $= 0.5$.

As $|\sigma|$ increases the models diverge and the DEP falls toward zero.

In the scalar model the two hypotheses are fully explicit:

$$\text{approximating:} \quad \mu_{t+1} = \mu_t + \alpha w_{t+1}, \qquad \text{worst-case:} \quad \mu_{t+1} = \zeta(\sigma)\,\mu_t + \alpha w_{t+1} \tag{87.38}$$

Both have Gaussian innovations with the same variance α^2, so the log-likelihood ratio is a difference of sums of squares.

```
def simulate_paths(ζ, α, T, n_paths, seed):
    "Simulate n_paths draws of μ_{t+1} = ζ μ_t + α w_{t+1} from μ_0 = 0."
    rng = np.random.default_rng(seed)
    paths = np.zeros((n_paths, T + 1))
    shocks = rng.standard_normal((n_paths, T))
    for t in range(T):
        paths[:, t + 1] = ζ * paths[:, t] + α * shocks[:, t]
    return paths


def log_likelihood_ratio(paths, ζ, α):
    "Return log p_worst(path) - log p_approx(path)."
    lag, lead = paths[:, :-1], paths[:, 1:]
    return 0.5 * (np.sum(((lead - lag) / α)**2, axis=1)
                  - np.sum(((lead - ζ * lag) / α)**2, axis=1))


def detection_error_probability(ζ, α, T=40, n_paths=10_000, seed=1234):
    "Finite-sample DEP for the approximating and worst-case scalar laws."
    if np.isclose(ζ, 1.0):
        return 0.5
    approx = simulate_paths(1.0, α, T, n_paths, seed)
    worst = simulate_paths(ζ, α, T, n_paths, seed + 1)
    return 0.5 * (np.mean(log_likelihood_ratio(worst, ζ, α) < 0)
                  + np.mean(log_likelihood_ratio(approx, ζ, α) > 0))
```

We use $T = 40$, which is ten years of quarterly data.

Before plotting, it is worth noting a property that makes the DEP the right way to report a robustness concern.

The parameter σ is *not* scale free: it always enters through the product $\sigma\alpha^2$, and α^2 has the units of consumption squared.

Doubling the units in which consumption is measured therefore changes the numerical value of σ that represents a given concern for robustness.

The DEP has no such problem, because α cancels from the likelihood ratio in (87.38).

```
ζ_test = worst_case_persistence(0.6 * σ_lo, β, α2)
for scale in [0.5, 1.0, 2.0]:
    dep = detection_error_probability(ζ_test, scale * α)
    print(f"α scaled by {scale:>4}:  DEP = {dep:.4f}")
```

```
α scaled by  0.5:  DEP = 0.3334
α scaled by  1.0:  DEP = 0.3334
α scaled by  2.0:  DEP = 0.3334
```

The DEP depends only on ζ and on the sample size T.

```
σ_vals = np.linspace(0.0, 0.999 * σ_lo, 31)
ζ_vals = worst_case_persistence(σ_vals, β, α2)

fig, ax = plt.subplots()
for T, color in [(40, 'C0'), (160, 'C2')]:
    dep_vals = np.array([detection_error_probability(ζ, α, T=T)
                         for ζ in ζ_vals])
    ax.plot(-σ_vals, dep_vals, lw=2, color=color, label=f'$T = {T}$')

ax.axhline(0.2, color='C3', linestyle='--', lw=1.2, label='DEP = 0.2')
ax.axvline(-σ_lo, color='k', linestyle=':', lw=1,
           label='breakdown point')
ax.set_xlabel(r'robustness concern $|\sigma|$')
ax.set_ylabel('detection error probability')
ax.set_ylim(0.0, 0.52)
ax.legend()
plt.show()
```

Note

HST suggested that a DEP above 0.2 is "plausible", meaning the models are still hard enough to distinguish statistically that a concern for robustness is warranted.

Values of σ with DEP ≥ 0.2 therefore define a set of plausible worst-case models.

Fig. 87.4 shows that the two disciplines interact in an interesting way.

At $T = 40$ the DEP curve reaches the 0.2 threshold essentially exactly at the breakdown point.

In other words, with ten years of quarterly data, every robustness concern that this model can represent at all is also statistically plausible: the mathematical limit binds first, and the statistical one is slack.

```
for frac in [0.6, 0.9, 0.999]:
    ζ = worst_case_persistence(frac * σ_lo, β, α2)
    print(f"σ = {frac * σ_lo:.5f}  (= {frac:.3f} σ̲):  "
          f"DEP = {detection_error_probability(ζ, α):.4f}")
```

```
σ = -0.06601  (= 0.600 σ̲):  DEP = 0.3334
σ = -0.09901  (= 0.900 σ̲):  DEP = 0.2329
σ = -0.10990  (= 0.999 σ̲):  DEP = 0.2034
```

That near-coincidence is a property of this calibration and of $T = 40$, not a theorem.

At $T = 160$ the ordering reverses, and statistical detectability binds well before the breakdown point.

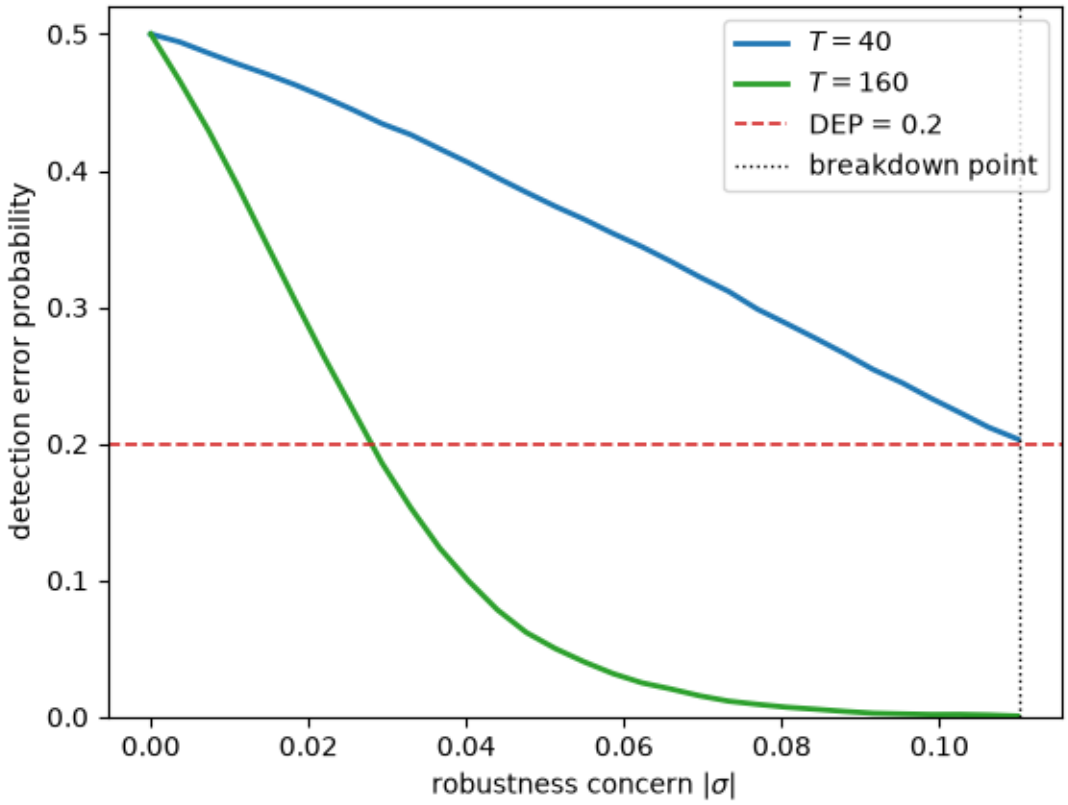


Fig. 87.4: Finite-sample detection error probability against the robustness concern, over the admissible range $(\underline{\sigma}, 0]$, for two sample lengths. Values below the dashed line are the ones HST regard as implausible, because the approximating and worst-case models would then be too easy to tell apart.

```
σ_report = 0.6 * σ_lo
ζ_report = worst_case_persistence(σ_report, β, α2)
print(f"at σ = {σ_report:.4f}  (ζ = {ζ_report:.4f}):")
for T in [20, 40, 80, 160]:
    print(f"  T = {T:>3}:  DEP = "
          f"{detection_error_probability(ζ_report, α, T=T):.4f}")
```

```
at σ = -0.0660  (ζ = 1.0309):
  T =  20:  DEP = 0.4147
  T =  40:  DEP = 0.3334
  T =  80:  DEP = 0.1507
  T = 160:  DEP = 0.0203
```

The lesson is that the plausible amount of model uncertainty depends on how much data the econometrician is imagined to have.

87.8 Robustness of decision rules

A natural follow-up question is whether robust decision rules perform better than the non-robust rule when the data are in fact generated by a distorted model.

Define the **payoff** when the decision rule is designed for robustness parameter σ^r and the data are generated by the distorted model associated with σ^d:

$$\pi(\sigma^d;\sigma^r) = -\mathbb{E}_{0,\sigma^d}\sum_{t=0}^{\infty}\beta^t\, x_t^\top H(\sigma^r)^\top H(\sigma^r)\, x_t \tag{87.39}$$

where the state evolves under decision rule $F(\sigma^r)$ and worst-case shocks $K(\sigma^d)$:

$$x_{t+1} = (A - BF(\sigma^r) + CK(\sigma^d))x_t + C\, w_{t+1} \tag{87.40}$$

There is an important caveat about *which* family of rules to compare, and it follows directly from *Theorem 87.4.1*.

Along the observational-equivalence locus $(\sigma, \hat{\beta}(\sigma))$, every agent chooses the *same* decision rule.

So $F(\sigma^r)$ does not vary along the locus, and $\pi(\sigma^d;\sigma^r)$ is constant in σ^r there.

A meaningful comparison of rules must therefore move *off* the locus, for example by holding β fixed at the benchmark and varying σ^r alone.

That experiment requires solving the full HST matrix problem rather than the scalar reduction of *The scalar worst-case model*, because once we leave the locus the marginal-utility process no longer summarises the decision rule.

Robust Permanent Income and Pricing carries out that computation with the QuantEcon `LQ` and robust-control routines.

87.9 Concluding remarks

A concern for model misspecification, parameterised by $\sigma = -\theta^{-1} \le 0$, alters the permanent income model in ways that are subtle rather than dramatic.

A concern for robustness generates a precautionary savings motive even under quadratic preferences, by distorting the conditional means of income shocks.

The distorted worst-case model makes the income process **more persistent**, shifting power toward the low frequencies where the permanent income consumer is most vulnerable.

The observational equivalence theorem *Theorem 87.4.1* shows that for quantities (c_t, i_t) alone, a concern for robustness is indistinguishable from a reduction in β.

The reverse theorem *Theorem 87.4.2* shows that, starting from $\beta R = 1$, robustness is observationally equivalent to an *increase* in β, which imparts an upward drift to expected consumption.

Two disciplines bound how large a robustness concern can sensibly be.

The breakdown point (87.34) is a hard mathematical limit beyond which no solution exists.

Detection error probabilities provide a softer and scale-free statistical discipline: choose $|\sigma|$ small enough that the approximating and worst-case models remain difficult to distinguish.

The observationally equivalent $(\sigma, \hat{\beta})$ pairs **do** have different implications for asset prices, a point pursued in *Robust Permanent Income and Pricing*.

They also have different implications for the *beliefs* that agents hold, which is the subject of *A Robust LQ Bewley Model*.

87.10 Exercises

Exercise 87.10.1

This exercise derives the breakdown point (87.34).

1. Using (87.31), show that $\hat{\beta}(\sigma)\,\zeta(\sigma)^2 = \beta\,\zeta(\sigma)$.
2. The minimizing agent's problem has a finite value only if $\hat{\beta}\zeta^2 < 1$. Use part 1 to show that this is equivalent to $\zeta < R$, and hence that $\underline{\sigma} = -(1-\beta)^2/\alpha^2$.
3. Verify that $\hat{\beta}(\underline{\sigma}) = \beta^2$.
4. Explain why the breakdown point moves toward zero when the endowment becomes more volatile.

Solution

Here is one solution.

1. From (87.31), $\zeta = \beta/\hat{\beta}$, so $\hat{\beta} = \beta/\zeta$ and therefore $\hat{\beta}\zeta^2 = (\beta/\zeta)\zeta^2 = \beta\zeta$.
2. By part 1 the condition $\hat{\beta}\zeta^2 < 1$ is $\beta\zeta < 1$, that is, $\zeta < \beta^{-1} = R$.

   Substituting the closed form (87.31) gives

   $$\left[1 + \frac{\sigma\alpha^2}{1-\beta}\right]^{-1} < \frac{1}{\beta} \iff 1 + \frac{\sigma\alpha^2}{1-\beta} > \beta \iff \sigma\alpha^2 > -(1-\beta)^2,$$

   which is $\sigma > -(1-\beta)^2/\alpha^2 = \underline{\sigma}$.
3. At $\sigma = \underline{\sigma}$ we have $\sigma\alpha^2 = -(1-\beta)^2$, so $\zeta = [1-(1-\beta)]^{-1} = \beta^{-1} = R$ and $\hat{\beta} = \beta/\zeta = \beta^2$.
4. A more volatile endowment raises α^2, and $\underline{\sigma} = -(1-\beta)^2/\alpha^2$ is therefore closer to zero.

   The economics is that σ only ever matters through $\sigma\alpha^2$, so when the consumer faces more income risk a *smaller* $|\sigma|$ already buys a given amount of pessimism.

Exercise 87.10.2

This exercise explains why a naive numerical solver for (87.23) misbehaves.

1. Verify algebraically that on the locus (87.17) the discriminant of (87.23) equals $\left[(\sigma\alpha^2 + (1-\beta)^2)/(1-\beta)\right]^2$.
2. Conclude that the square root equals $|\sigma\alpha^2 + (1-\beta)^2|/(1-\beta)$ and hence that the two roots are those in (87.33).
3. Write code that, for a grid of σ in $(\underline{\sigma}, 0)$, records which *sign* in front of the square root in (87.23) yields the economically relevant root $P = -(1-\beta)^{-1}$.

   Confirm that the answer flips at $\sigma = \underline{\sigma}$.

Solution

Here is one solution.

1. With $u = \sigma\alpha^2$ and $\delta = 1 - \beta$, equation (87.17) gives $\hat{\beta} - 1 = -\delta + u/\delta \cdot \beta/\beta$, more directly $\hat{\beta} - 1 + u = (u - \delta^2)/\delta$.

   Hence the discriminant is

$$\frac{(u-\delta^2)^2}{\delta^2} + 4u = \frac{(u-\delta^2)^2 + 4u\delta^2}{\delta^2} = \frac{(u+\delta^2)^2}{\delta^2}.$$

2. Taking the square root gives $|u + \delta^2|/\delta$, and substituting the two signs into (87.23) gives $P = -1/\delta$ and $P = \delta/u$.
3. The sign flips exactly where $u + \delta^2$ changes sign, that is at $u = -\delta^2$, which is $\sigma = \underline{\sigma}$.

```
σ_test = np.linspace(-0.01, 1.4 * σ_lo, 40)
signs = []
for σ in σ_test:
    roots, _ = solve_scalar_riccati(σ, β, α2)
    signs.append('+' if np.argmin(np.abs(roots + 1 / (1 - β))) == 0 else '-')

print(''.join(signs))
flip = next(i for i in range(1, len(signs)) if signs[i] != signs[i - 1])
print(f"sign flips between σ = {σ_test[flip - 1]:.5f} "
      f"and σ = {σ_test[flip]:.5f}")
print(f"breakdown point σ = {σ_lo:.5f}")
```

```
++++++++++++++++++++++++++++++------------
sign flips between σ = -0.10970 and σ = -0.11340
breakdown point σ = -0.11001
```

Exercise 87.10.3

This exercise calibrates σ using the detection error probability.

Write a bisection that finds the σ at which $\text{DEP}(\sigma) = 0.2$, and report the associated $\hat{\beta}$, ζ, and the ratio $\sigma/\underline{\sigma}$.

Run it for $T = 40$ and $T = 160$.

Take care: the target need not be attained on the admissible range $(\underline{\sigma}, 0]$, and your code should say so rather than silently return an endpoint.

Solution

The DEP is decreasing in $|\sigma|$, so a bisection works, but we must first check that the target lies within reach.

```
def σ_for_target_dep(target, T, β, α2, tol=1e-5):
    """
    Find σ ∈ (σ̲, 0) with DEP(σ) = target by bisection.

    Returns None if the DEP never falls to the target on the admissible
    range, which happens when the breakdown point binds before statistical
    detectability does.
    """
    α_loc = np.sqrt(α2)
    lo, hi = 0.999 * (-(1 - β)**2 / α2), 0.0     # lo is the more negative end

    def dep_at(σ):
        return detection_error_probability(
            worst_case_persistence(σ, β, α2), α_loc, T=T)

    if dep_at(lo) > target:
        return None

    while hi - lo > tol:
        mid = 0.5 * (lo + hi)
        if dep_at(mid) < target:
            lo = mid                # too easy to detect, move toward zero
        else:
            hi = mid
    return 0.5 * (lo + hi)


for T in [40, 160]:
    σ_star = σ_for_target_dep(0.2, T, β, α2)
    if σ_star is None:
        print(f"T = {T:>3}:  DEP stays above 0.2 on the whole admissible "
              f"range; the breakdown point binds first")
    else:
        ζ_star = worst_case_persistence(σ_star, β, α2)
        print(f"T = {T:>3}:  σ = {σ_star:.5f}   β = {β / ζ_star:.5f}   "
              f"ζ = {ζ_star:.5f}   σ/σ̲ = {σ_star / σ_lo:.3f}")
```

```
T =  40:  DEP stays above 0.2 on the whole admissible range; the breakdown␣
↪point binds first
T = 160:  σ = -0.02806   β = 0.93788   ζ = 1.01292   σ/σ̲ = 0.255
```

A longer sample makes the two models easier to distinguish, so the σ that keeps the DEP at 0.2 moves closer to zero.

At $T = 40$ no admissible σ has a DEP as low as 0.2, so the breakdown point is the binding constraint.

At $T = 160$ statistical detectability binds first, at about a quarter of the way to the breakdown point.

87.11 Related lectures

- *The LQ Permanent Income Model* develops the standard LQ permanent income model used here.
- *Consumption Smoothing with Incomplete and Complete Markets* studies the cross-section of consumption and the market structures that support it.
- *A Robust LQ Bewley Model* applies the observational-equivalence results of this lecture to build a Bewley economy with heterogeneous concerns about misspecification.
- *Robust Permanent Income and Pricing* treats risk-sensitive preferences, estimation, and asset pricing in the HST model.

CHAPTER

EIGHTYEIGHT

A ROBUST LQ BEWLEY MODEL

Contents

- *A Robust LQ Bewley Model*
  - *Overview*
  - *Mapping the Bewley economy into the HST framework*
  - *The robustness scalar for this economy*
  - *The Bewley observational-equivalence locus*
  - *Equilibrium with heterogeneous types*
  - *Where the agents genuinely differ*
  - *Computation*
  - *Concluding remarks*
  - *Exercises*
  - *Related lectures*

88.1 Overview

This lecture embeds the Bewley economy of *Consumption Smoothing with Incomplete and Complete Markets* in the robust permanent income framework of *Robust Consumption Smoothing and Precautionary Savings*.

It is the last of four lectures on the LQ permanent income model.

The result is a family of economies in which consumers disagree about the model generating their income, yet behave identically.

Using the observational-equivalence theorem of Hansen *et al.* [1999], we show

- how a continuum of consumers i can differ in their robustness parameters $\sigma_i \leq 0$ and their discount factors β_i, provided each pair (σ_i, β_i) lies on an observational-equivalence locus
- how every such consumer chooses the **same consumption-saving rule** as a benchmark $(\sigma = 0, \beta)$ agent who fully trusts the endowment process
- how the equilibrium interest rate $R = \beta^{-1}$ and all aggregate and cross-section dynamics therefore coincide with those of the plain-vanilla Bewley model

- how, despite all of this, distinct (σ_i, β_i) agents hold genuinely different subjective models of their non-financial income

The economy is a pure endowment economy, so there is no physical capital and investment plays no role.

We read *Robust Consumption Smoothing and Precautionary Savings* as a prerequisite and carry over its notation.

> **Note**
>
> As in *Robust Consumption Smoothing and Precautionary Savings*, w_{t+1} is the baseline shock and v_{t+1} a distortion to its conditional mean, $\sigma \leq 0$ is the robustness parameter, η_1 and η_2 are the standard deviations of the permanent and transitory endowment shocks, and a_t denotes net assets.

Let's begin with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
```

88.2 Mapping the Bewley economy into the HST framework

We specialise the robust model of *Robust Consumption Smoothing and Precautionary Savings* to $\lambda = \delta_h = 0$, so there are no habits and no durable goods, and to a pure endowment economy with no physical capital, $k_t = 0$.

In this case services equal consumption, $s_t = c_t$.

The only traded security is a one-period risk-free bond, and a_t denotes the household's net asset position, so that positive a_t is wealth.

The endowment process follows the state-space representation

$$
\begin{aligned}
z_{t+1} &= \check{A}\, z_t + \check{C}\, w_{t+1} \\
y_t &= \check{G}\, z_t
\end{aligned}
\tag{88.1}
$$

with the two-factor specification $y_t = z_{1t} + z_{2t}$, $\check{A} = \mathrm{diag}(1, 0)$ and $\check{C} = \mathrm{diag}(\eta_1, \eta_2)$.

The household's augmented state vector is $x_t = [a_t,\ z_t^\top]^\top$, and the law of motion of *Robust Consumption Smoothing and Precautionary Savings* specialises to

$$
\begin{pmatrix} a_{t+1} \\ z_{t+1} \end{pmatrix}
= \underbrace{\begin{pmatrix} R & R\check{G} \\ 0 & \check{A} \end{pmatrix}}_{A}
\begin{pmatrix} a_t \\ z_t \end{pmatrix}
+ \underbrace{\begin{pmatrix} -R \\ 0 \end{pmatrix}}_{B} c_t
+ \underbrace{\begin{pmatrix} 0 \\ \check{C} \end{pmatrix}}_{C} (w_{t+1} + v_{t+1})
\tag{88.2}
$$

The objective is $\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t [-(c_t - b)^2]$, which is the HST criterion with $\sigma = 0$ and a constant bliss level $b_t \equiv b$.

The period return is written without a factor of $\frac{1}{2}$, and that choice is not cosmetic: it fixes the scale of the robustness parameter σ used throughout, since rescaling the return rescales σ by the same factor.

The robust Bellman equation with $\sigma = 0$ therefore reduces exactly to the LQ problem of *The LQ Permanent Income Model*, confirming that the HST framework nests the Bewley model.

88.3 The robustness scalar for this economy

Everything about the endowment process that matters for robustness is summarised by the scalar α^2 derived in *Robust Consumption Smoothing and Precautionary Savings*.

For the two-factor endowment it is

$$\alpha^2 = \eta_1^2 + (1-\beta)^2\,\eta_2^2 \tag{88.3}$$

This is the variance of the consumption innovation $h\,w_{t+1}$, where $h = (1-\beta)\check{G}(I-\beta\check{A})^{-1}\check{C} = \begin{pmatrix}\eta_1 & (1-\beta)\eta_2\end{pmatrix}$.

In the present setting α^2 has a second, equally concrete meaning that we met in *Consumption Smoothing with Incomplete and Complete Markets*.

Because individual consumption is a random walk with innovation variance α^2, the cross-section variance of consumption among agents of age t who started from a common initial condition is exactly $t\,\alpha^2$.

So the same scalar that governs how much a consumer's robustness concern bites also governs how fast the Bewley cross-section fans out.

88.4 The Bewley observational-equivalence locus

Applying the observational-equivalence theorem *Theorem 87.4.1* of *Robust Consumption Smoothing and Precautionary Savings* at the equilibrium interest rate $R = \beta^{-1}$ gives the **Bewley observational-equivalence locus**

$$\hat{\beta}(\sigma) = \beta + \frac{\sigma\,\alpha^2\,\beta}{1-\beta} \tag{88.4}$$

For $\sigma < 0$ we have $\hat{\beta}(\sigma) < \beta$.

An agent with the pair $(\sigma, \hat{\beta}(\sigma))$ is more concerned about model misspecification, because σ is lower, but also more impatient, because $\hat{\beta}$ is lower.

The two forces cancel exactly, leaving the consumption decision rule unchanged.

The locus is admissible only down to the breakdown point of *Robust Consumption Smoothing and Precautionary Savings*,

$$\underline{\sigma} = -\frac{(1-\beta)^2}{\alpha^2}, \qquad \text{at which} \qquad \hat{\beta}(\underline{\sigma}) = \beta^2 \tag{88.5}$$

Below $\underline{\sigma}$ the individual robust control problem has no solution, so there is no economy to describe.

The bound has a transparent reading in terms of the worst-case dynamics derived below.

Since $\zeta(\sigma) = \beta/\hat{\beta}(\sigma)$ is the growth rate that agent σ fears for its own marginal utility, the breakdown point is exactly the σ at which

$$\zeta(\underline{\sigma}) = \frac{1}{\beta} = R, \qquad \text{equivalently} \qquad \beta\,\zeta(\underline{\sigma}) = 1 \tag{88.6}$$

So $\underline{\sigma}$ is the robustness level at which the feared growth in marginal utility just reaches the gross interest rate, and the agent's discounted worst-case objective ceases to converge.

Any stronger concern for robustness would have the agent guarding against a future it cannot value.

88.5 Equilibrium with heterogeneous types

We can now populate the economy with a continuum of types that differ in their concern for robustness.

Proposition 88.5.1 (A robust Bewley equilibrium)

Let each agent i in the unit interval be indexed by a robustness parameter $\sigma_i \in (\underline{\sigma}, 0]$, distributed according to any distribution Φ, and let agent i have discount factor

$$\beta_i = \hat{\beta}(\sigma_i) = \beta + \frac{\sigma_i \, \alpha^2 \, \beta}{1 - \beta} \tag{88.7}$$

so that every pair (σ_i, β_i) lies on the locus (88.4).

Then

1. every agent's optimal consumption plan is identical to that of the plain-vanilla $(\sigma = 0, \beta)$ agent,
2. $R = \beta^{-1}$ is an equilibrium gross interest rate, independently of Φ, and
3. the aggregate and cross-section dynamics coincide with those of the benchmark Bewley economy of *Consumption Smoothing with Incomplete and Complete Markets*.

Proof. By *Theorem 87.4.1*, an agent with parameters $(\sigma_i, \hat{\beta}(\sigma_i))$ facing gross interest rate $R = \beta^{-1}$ chooses the same consumption-saving rule as the benchmark $(0, \beta)$ agent.

This holds agent by agent and does not require the σ_i to be equal, which gives part 1.

Since all individual rules coincide with the benchmark rule, the goods-market clearing condition $\int c_t^i \, di = Y$ and the bond-market condition $\int a_t^i \, di = 0$ are the benchmark conditions, so they are satisfied at $R = \beta^{-1}$ for exactly the reason given in *Consumption Smoothing with Incomplete and Complete Markets*.

Because market clearing never refers to Φ, neither does the rate that clears the market, which gives part 2.

The argument is a verification: the locus (88.4) is itself constructed at $R = \beta^{-1}$, so what we have shown is that this rate reproduces itself as an equilibrium for any Φ, not that no other rate could.

Part 3 follows because aggregate and cross-section objects are integrals of individual paths, and the individual paths are the benchmark paths.

The distribution Φ of robustness types is therefore completely unidentified by quantity data.

An econometrician who observes $\{c_t^i, a_t^i\}$ for every agent and every date cannot tell whether the economy is populated entirely by $\sigma_i = 0$ agents, entirely by σ_i near $\underline{\sigma}$ agents, or by any mixture.

One feature of this equilibrium deserves comment, because it runs against a familiar result.

In a model with heterogeneous discount factors and a common interest rate, the most patient type ordinarily comes to hold all of the wealth, and the long-run distribution degenerates.

Here every type with $\sigma_i < 0$ has $\beta_i R < 1$ and so is impatient at the market rate, yet no type decumulates relative to any other.

The reason is that impatience and the precautionary motive are offset at every date, not merely on average, so asset paths as well as consumption paths coincide across types.

Robustness type is therefore uncorrelated with wealth at every horizon, and the usual sorting force is exactly neutralized rather than merely slowed.

88.6 Where the agents genuinely differ

Agents on the locus are indistinguishable in what they *do* but not in what they *believe*.

An agent with $\sigma_i < 0$ applies a worst-case distortion $v^i_{t+1} = K(\sigma_i, \beta_i)\, \mu^i_{s,t}$ to its conditional expectations, while an agent with $\sigma_i = 0$ takes the approximating model at face value.

From *Robust Consumption Smoothing and Precautionary Savings*, the worst-case law for agent i's marginal utility is

$$\mu^i_{s,t+1} = \zeta_i\, \mu^i_{st} + \alpha\, w_{t+1}, \qquad \zeta_i = \frac{\beta}{\beta_i} = \left[1 + \frac{\sigma_i \alpha^2}{1-\beta}\right]^{-1} \geq 1 \tag{88.8}$$

with equality only for the fully trusting type $\sigma_i = 0$.

With $\lambda = \delta_h = 0$ and a constant bliss point we have $\mu_{st} = b - c_t$, so agent i's **worst-case expected consumption path** is

$$\hat{\mathbb{E}}_t\, c^i_{t+h} = b - \zeta_i^h\, (b - c_t) \tag{88.9}$$

Under the approximating model, by contrast, consumption is a martingale, $\mathbb{E}_t c_{t+h} = c_t$ for every h.

Equation (88.9) says that a robust agent below its bliss point expects, under its worst-case model, that consumption will *drift away* from bliss at the geometric rate ζ_i.

The more robust the agent, the faster the drift it guards against, and the more precautionary saving it does.

That extra saving is exactly offset by the lower β_i, which is why the realized path is the same for all types.

88.7 Computation

We use the calibration of the preceding lectures.

```
β = 0.95         # benchmark discount factor, R = 1/β
η1 = 0.15        # std of permanent shock
η2 = 0.30        # std of transitory shock
b = 1.0          # bliss point

R = 1 / β
h = np.array([η1, (1 - β) * η2])       # consumption innovation loadings
α2 = h @ h
α = np.sqrt(α2)
σ_lo = -(1 - β)**2 / α2                # breakdown point, eq:rbew-breakdown

print(f"α^2 = {α2:.6f}   α = {α:.6f}")
print(f"breakdown σ̲ = {σ_lo:.6f},  where β = {β + σ_lo * α2 * β / (1 - β):.6f}"
      f"  (β² = {β**2:.6f})")
```

```
α^2 = 0.022725   α = 0.150748
breakdown σ̲ = -0.110011,  where β = 0.902500  (β² = 0.902500)
```

Next we build a set of types spread across the admissible range and record what distinguishes them.

We reuse the detection error probability of *Robust Consumption Smoothing and Precautionary Savings* to report how hard each type's worst-case model would be to detect in a sample of $T = 40$ quarters.

```
def worst_case_persistence(σ, β, α2):
    "Worst-case persistence ζ(σ) of marginal utility, eq:rbew-zeta."
    return 1 / (1 + σ * α2 / (1 - β))


def simulate_paths(ζ, α, T, n_paths, seed):
    "Simulate n_paths draws of μ_{t+1} = ζ μ_t + α w_{t+1} from μ_0 = 0."
    rng = np.random.default_rng(seed)
    paths = np.zeros((n_paths, T + 1))
    shocks = rng.standard_normal((n_paths, T))
    for t in range(T):
        paths[:, t + 1] = ζ * paths[:, t] + α * shocks[:, t]
    return paths


def log_likelihood_ratio(paths, ζ, α):
    "Return log p_worst(path) - log p_approx(path)."
    lag, lead = paths[:, :-1], paths[:, 1:]
    return 0.5 * (np.sum(((lead - lag) / α)**2, axis=1)
                  - np.sum(((lead - ζ * lag) / α)**2, axis=1))


def detection_error_probability(ζ, α, T=40, n_paths=10_000, seed=1234):
    "Finite-sample DEP for the approximating and worst-case scalar laws."
    if np.isclose(ζ, 1.0):
        return 0.5
    approx = simulate_paths(1.0, α, T, n_paths, seed)
    worst = simulate_paths(ζ, α, T, n_paths, seed + 1)
    return 0.5 * (np.mean(log_likelihood_ratio(worst, ζ, α) < 0)
                  + np.mean(log_likelihood_ratio(approx, ζ, α) > 0))
```

```
σ_types = np.array([0.0, 0.3, 0.6, 0.9]) * σ_lo
β_types = β + σ_types * α2 * β / (1 - β)
ζ_types = worst_case_persistence(σ_types, β, α2)
dep_types = np.array([detection_error_probability(ζ, α) for ζ in ζ_types])

print(f"{'σ_i':>10}{'β_i':>10}{'ζ_i':>10}{'DEP':>8}")
for σ_i, β_i, ζ_i, dep in zip(σ_types, β_types, ζ_types, dep_types):
    print(f"{σ_i:10.4f}{β_i:10.4f}{ζ_i:10.4f}{dep:8.3f}")
```

```
       σ_i       β_i       ζ_i     DEP
   -0.0000    0.9500    1.0000   0.500
   -0.0330    0.9357    1.0152   0.426
   -0.0660    0.9215    1.0309   0.333
   -0.0990    0.9072    1.0471   0.233
```

These four types differ substantially in patience and in the pessimism of their worst-case model.

We now confirm that they nonetheless behave identically.

Proposition 88.5.1 asserts that each type, solving *its own* problem at (σ_i, β_i), arrives at the benchmark decision rule.

Testing that claim means solving each type's robust problem separately and comparing the rules that come out.

Assuming the common rule and then reporting that the resulting paths coincide would establish nothing.

We reuse the risk-sensitive LQ solver of *Robust Permanent Income and Pricing*, renaming its state-cost argument to `Rc` because R is the gross interest rate here.

The period return below is $-(c_t - b)^2$, matching the normalization of σ fixed above.

```
def solve_rslq(A, B, C, Q, Rc, β, σ, N=None, tol=1e-12, max_iter=50_000):
    "Risk-sensitive LQ regulator; returns F in the rule c = -F x."
    A, B, C, Q, Rc = map(np.atleast_2d, (A, B, C, Q, Rc))
    n, kw = A.shape[0], C.shape[1]
    if N is None:
        N = np.zeros((B.shape[1], n))
    Ω, Iw = -np.eye(n), np.eye(kw)
    for _ in range(max_iter):
        M = Iw - σ * C.T @ Ω @ C
        D = Ω + σ * Ω @ C @ np.linalg.solve(M, C.T @ Ω)
        F = np.linalg.solve(Q - β * B.T @ D @ B, N - β * B.T @ D @ A)
        Acl = A - B @ F
        Ω_new = -Rc - F.T @ Q @ F + (F.T @ N + N.T @ F) + β * Acl.T @ D @ Acl
        if np.max(np.abs(Ω_new - Ω)) < tol:
            return F
        Ω = Ω_new
    raise RuntimeError('risk-sensitive Riccati iteration did not converge')
```

Write the agent's problem with state $x_t = \begin{pmatrix} 1 & a_t & z_{1t} & z_{2t} \end{pmatrix}'$ and control c_t, matching the timing of (88.2).

The constant carries the bliss point, and every agent faces the same market rate $R = \beta^{-1}$; only β_i and σ_i differ across types.

```
def bewley_lq(b, η1, η2, R):
    "State-space matrices for the agent's problem, period return -(c-b)^2."
    A_x = np.array([[1, 0, 0, 0],
                    [0, R, R, R],
                    [0, 0, 1, 0],
                    [0, 0, 0, 0]], float)
    B_x = np.array([[0.], [-R], [0.], [0.]])
    C_x = np.array([[0, 0], [0, 0], [η1, 0], [0, η2]], float)
    Q_x = np.array([[1.0]])
    Rc_x = np.zeros((4, 4))
    Rc_x[0, 0] = b**2
    N_x = np.array([[-b, 0, 0, 0]], float)
    return A_x, B_x, C_x, Q_x, Rc_x, N_x


A_x, B_x, C_x, Q_x, Rc_x, N_x = bewley_lq(b, η1, η2, R)
F_bench = solve_rslq(A_x, B_x, C_x, Q_x, Rc_x, β, 0.0, N_x)

print("benchmark rule  c = "
      + " + ".join(f"{v:.4f}·{n}" for v, n in
                   zip(-F_bench.ravel(), ["1", "a", "z1", "z2"])))
print(f"implied consumption innovation  {np.round((-F_bench @ C_x).ravel(), 6)}")
print(f"analytic h                      {np.round(h, 6)}")
```

```
benchmark rule  c = 0.0000·1 + 0.0500·a + 1.0000·z1 + 0.0500·z2
implied consumption innovation  [0.15  0.015]
analytic h                      [0.15  0.015]
```

The benchmark rule reproduces the analytic innovation loadings h, which checks that the state-space setup is the one the algebra describes.

Now solve each type's own problem and compare.

```
F_types = [solve_rslq(A_x, B_x, C_x, Q_x, Rc_x, β_i, σ_i, N_x)
           for σ_i, β_i in zip(σ_types, β_types)]

print(f"{'σ_i':>10}{'β_i':>10}{'max |F_i - F_bench|':>22}")
for σ_i, β_i, F_i in zip(σ_types, β_types, F_types):
    print(f"{σ_i:10.4f}{β_i:10.4f}{np.max(np.abs(F_i - F_bench)):22.2e}")
```

```
       σ_i       β_i   max |F_i - F_bench|
   -0.0000    0.9500              0.00e+00
   -0.0330    0.9357              1.27e-12
   -0.0660    0.9215              2.32e-12
   -0.0990    0.9072              9.43e-12
```

Every coefficient of every type's rule agrees with the benchmark to eleven decimal places or better, which is *Proposition 88.5.1* part 1.

To see that this test has power, move each discount factor one percent off the locus, keeping σ_i fixed, and solve again.

```
print(f"{'σ_i':>10}{'on locus':>12}{'+1% off':>12}{'-1% off':>12}")
for σ_i, β_i in zip(σ_types[1:], β_types[1:]):
    devs = []
    for mult in (1.0, 1.01, 0.99):
        F_off = solve_rslq(A_x, B_x, C_x, Q_x, Rc_x, β_i * mult, σ_i, N_x)
        devs.append(np.max(np.abs(F_off - F_bench)))
    print(f"{σ_i:10.4f}" + "".join(f"{d:12.1e}" for d in devs))
```

```
       σ_i    on locus     +1% off     -1% off
   -0.0330     1.3e-12     2.6e-01     2.7e-01
   -0.0660     2.3e-12     4.5e-01     4.7e-01
   -0.0990     9.4e-12     1.7e+00     1.0e+00
```

A one percent departure from the locus moves the rule in the first decimal place, so the agreement above is not an artifact of the comparison.

Finally, simulate each type using **its own** solved rule, with common shocks.

```
T = 60
rng = np.random.default_rng(42)
shocks = rng.standard_normal((T, 2))          # common shocks for all types

c_paths = np.zeros((len(σ_types), T + 1))
for i, F_i in enumerate(F_types):
    x = np.array([1.0, 0.0, 0.0, 0.0])        # [1, a_t, z_1, z_2]
    for t in range(T + 1):
        c_paths[i, t] = -(F_i @ x).item()
        if t < T:
            x = A_x @ x + B_x.ravel() * c_paths[i, t] + C_x @ shocks[t]

print("max absolute difference across types:"
      f" {np.abs(c_paths - c_paths[0]).max():.2e}")
print("max deviation from the random walk with innovation h:"
      f" {np.abs(c_paths[0] - np.concatenate([[0], np.cumsum(shocks @ h)])).max():.2e}
 ↪")
```

```
max absolute difference across types: 1.12e-11
max deviation from the random walk with innovation h: 5.78e-14
```

The paths coincide, and each reproduces the random walk with innovation h that the algebra predicts.

The next figure contrasts what the types do with what they believe.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))

for i, σ_i in enumerate(σ_types):
    axes[0].plot(c_paths[i], lw=3 - 0.6 * i, alpha=0.9,
                 label=rf'$\sigma_i={σ_i:.3f}$')
axes[0].set_xlabel('$t$')
axes[0].set_ylabel('$c_t$')
axes[0].set_title('realized consumption')
axes[0].legend()

horizons = np.arange(41)
c_now = c_paths[0, 20]
axes[1].axhline(c_now, color='k', linestyle=':', lw=1.2,
                label='approximating model')
for σ_i, ζ_i in zip(σ_types, ζ_types):
    if σ_i == 0.0:
        continue
    axes[1].plot(horizons, b - ζ_i**horizons * (b - c_now), lw=2,
                 label=rf'worst case, $\sigma_i={σ_i:.3f}$')
axes[1].set_xlabel('horizon $h$')
axes[1].set_ylabel(r'$\hat{\mathbb{E}}_t\,c_{t+h}$')
axes[1].set_title('expected consumption under each belief')
axes[1].legend()

fig.tight_layout()
plt.show()
```

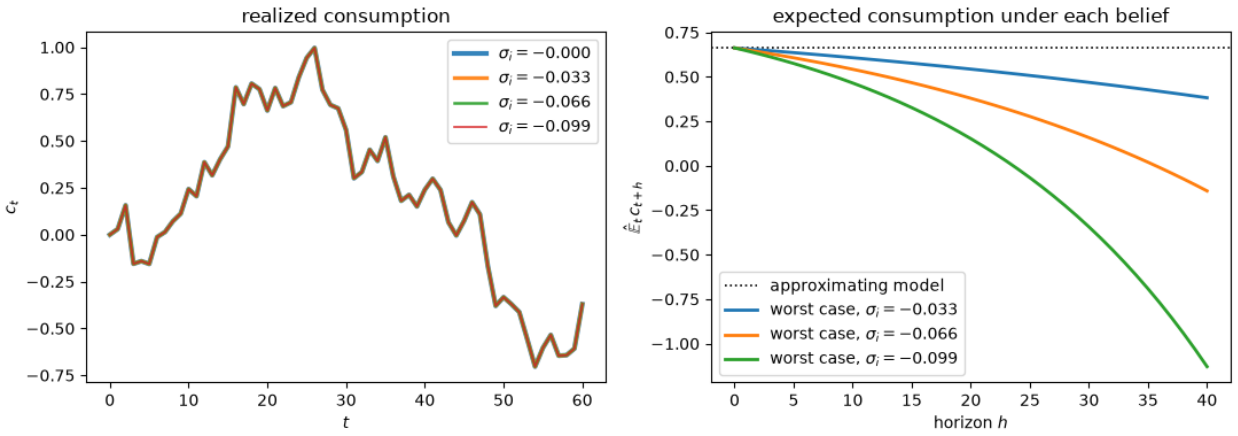


Fig. 88.1: Same actions, different beliefs. Left: realized consumption paths for four robustness types facing common shocks; the curves lie exactly on top of one another. Right: each type's worst-case expected consumption path from a common date, against the flat martingale forecast of the approximating model.

The left panel of Fig. 88.1 shows four curves drawn on top of one another.

The right panel shows that the same four agents expect very different futures.

The $\sigma_i = 0$ agent expects consumption to stay where it is.

Every robust agent guards against a future in which consumption drifts away from the bliss point, and the drift is faster the more robust the agent.

Finally we check part 3 of *Proposition 88.5.1*, that the cross-section behaves as in the benchmark Bewley economy.

We simulate a large population spread over the admissible range of types, giving each agent its own shocks and solving each type's own problem, and compare the cross-section variance of consumption to the benchmark prediction $t\,\alpha^2$.

```
n_agents, T_pop = 20_000, 40
rng = np.random.default_rng(1234)

# a population spread over the admissible range, each solving its own problem
σ_pop = np.linspace(0.99 * σ_lo, 0.0, 12)
β_pop = β + σ_pop * α2 * β / (1 - β)
F_pop = [solve_rslq(A_x, B_x, C_x, Q_x, Rc_x, β_j, σ_j, N_x)
         for σ_j, β_j in zip(σ_pop, β_pop)]

type_of = rng.integers(0, len(σ_pop), size=n_agents)
pop_shocks = rng.standard_normal((n_agents, T_pop, 2))
c_pop = np.zeros((n_agents, T_pop + 1))

for j, F_j in enumerate(F_pop):
    m = type_of == j
    x = np.zeros((m.sum(), 4))
    x[:, 0] = 1.0                                   # the constant
    for t in range(T_pop + 1):
        c_pop[m, t] = -(x @ F_j.ravel())
        if t < T_pop:
            x = (x @ A_x.T + np.outer(c_pop[m, t], B_x.ravel())
                 + pop_shocks[m, t] @ C_x.T)

print(f"{'t':>5}{'cross-section var':>20}{'t·α²':>12}")
for t in [10, 20, 30, 40]:
    print(f"{t:5d}{c_pop[:, t].var():20.5f}{t * α2:12.5f}")
print(f"\ncorrelation between σ_i and c_T: "
      f"{np.corrcoef(σ_pop[type_of], c_pop[:, T_pop])[0, 1]:+.4f}")
```

```
    t   cross-section var        t·α²
   10             0.22836     0.22725
   20             0.46220     0.45450
   30             0.68968     0.68175
   40             0.91503     0.90900

correlation between σ_i and c_T: -0.0002
```

The cross-section variance grows linearly at rate α^2, exactly as in *Consumption Smoothing with Incomplete and Complete Markets*.

Robustness type and consumption are uncorrelated, so the distribution of types leaves no trace in the data even though each agent has genuinely solved a different problem.

88.8 Concluding remarks

We embedded the single-agent robust permanent income model in a Bewley equilibrium with a continuum of agents who differ in how much they distrust their income process.

Provided each agent's pair (σ_i, β_i) lies on the observational-equivalence locus (88.4), every agent chooses the benchmark consumption rule.

The equilibrium interest rate $R = \beta^{-1}$, the aggregate dynamics, and the linear growth of the cross-section variance of consumption are all inherited unchanged from the plain-vanilla Bewley model of *Consumption Smoothing with Incomplete and Complete Markets*.

This is a strong non-identification result.

Quantity data pin down the decision rule, but they cannot decompose it into a degree of impatience and a degree of concern about misspecification.

What the agents do *not* share is their view of the world: each robust type acts as if its consumption were about to drift away from bliss at its own rate ζ_i.

Two routes lead out of this observational equivalence.

One is to look at asset prices, which do distinguish the $(\sigma, \hat{\beta})$ pairs, and which are studied in *Robust Permanent Income and Pricing*.

The other is to bound the plausible range of σ statistically, using the detection error probabilities of *Robust Consumption Smoothing and Precautionary Savings*.

88.9 Exercises

Exercise 88.9.1

This exercise asks you to carry out the translation from the benchmark Bewley economy into HST notation.

Specialise the robust-control setup to the no-habit, no-capital LQ Bewley environment, so $\lambda = \delta_h = 0$ and $k_t = 0$, and let the endowment follow the two-factor model.

1. Write the household state as $x_t = [a_t, z_t^\top]^\top$, where a_t is net assets, and derive the matrices (A, B, C) in (88.2).
2. Show that when $\sigma = 0$ the Bellman problem coincides with the LQ permanent-income problem of *The LQ Permanent Income Model*.
3. HST define $\alpha^2 = \nu^\top \nu$ with $\nu^\top = M_s C$, where $\mu_{st} = M_s x_t$. Compute M_s for this economy and verify that this route delivers the same α^2 as (88.3).

Solution

Here is one solution.

1. With budget law $a_{t+1} = R(a_t + y_t - c_t)$, $y_t = \check{G} z_t$ and $z_{t+1} = \check{A} z_t + \check{C} w_{t+1}$, stacking gives

$$\begin{pmatrix} a_{t+1} \\ z_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} R & R\check{G} \\ 0 & \check{A} \end{pmatrix}}_{A} \begin{pmatrix} a_t \\ z_t \end{pmatrix} + \underbrace{\begin{pmatrix} -R \\ 0 \end{pmatrix}}_{B} c_t + \underbrace{\begin{pmatrix} 0 \\ \check{C} \end{pmatrix}}_{C} w_{t+1}.$$

The sign of B is negative because higher c_t reduces asset accumulation.

2. At $\sigma = 0$ the minimizing agent is absent, the distortion term drops out of the Bellman equation, and the objective is $\mathbb{E}_0 \sum \beta^t[-(c_t - b)^2]$ subject to a linear law of motion.

That is precisely the LQ permanent-income problem.

3. With $\lambda = \delta_h = 0$ and constant bliss, $\mu_{st} = b - c_t$, and the optimal rule is

$$c_t = (1-\beta)[a_t + \check{G}(I - \beta\check{A})^{-1} z_t] + \text{constant},$$

so $M_s = -(1-\beta)\begin{pmatrix} 1 & \check{G}(I-\beta\check{A})^{-1}\end{pmatrix}$ up to the constant.

Since $C = [0;\ \check{C}]$, the first column of M_s is annihilated and

$$\nu^\top = M_s C = -(1-\beta)\check{G}(I-\beta\check{A})^{-1}\check{C} = -h.$$

Hence $\alpha^2 = \nu^\top \nu = hh^\top = \eta_1^2 + (1-\beta)^2\eta_2^2$, matching (88.3).

The minus sign is immaterial because only α^2 appears.

Exercise 88.9.2

This exercise works through the equilibrium logic of *Proposition 88.5.1*.

Fix a benchmark pair $(\beta, \sigma = 0)$ with $R = \beta^{-1}$ and let a unit interval of consumers be indexed by i, with type $\sigma_i \in (\underline{\sigma}, 0]$ and discount factor $\beta_i = \hat{\beta}(\sigma_i)$ from (88.4).

1. Use *Theorem 87.4.1* of *Robust Consumption Smoothing and Precautionary Savings* to show that each type has the same consumption rule as the benchmark $(\beta, 0)$ agent.
2. Show that goods- and bond-market clearing imply the same equilibrium interest rate $R = \beta^{-1}$ as in the plain-vanilla Bewley model, whatever the distribution of types.
3. Explain why agents can be observationally equivalent in quantities while holding different worst-case subjective models.
4. Suppose instead that agents share a common β but differ in σ_i, so that they are *not* on the locus. Explain why $R = \beta^{-1}$ would then generally fail to clear the bond market.

Solution

Here is one solution.

1. *Theorem 87.4.1* says that if (σ_i, β_i) satisfies $\beta_i = \beta + \sigma_i\alpha^2\beta/(1-\beta)$, then type i chooses the same decision rule as the benchmark agent, so all types share the policy function $c_t = \mathcal{C}(a_t, z_t)$.
2. Since all individual rules coincide with the benchmark rule, aggregating over i reproduces the benchmark market-clearing conditions, which hold at $R = \beta^{-1}$.

The type distribution never enters, so it cannot affect the equilibrium rate.

3. Observational equivalence is a statement about quantities generated by optimal rules.

The minimizing feedback $K(\sigma_i, \beta_i)$ still differs across types, so the agents attach different worst-case conditional means to the same shock process while making identical choices.

4. Off the locus the impatience offset is missing.

An agent with $\sigma_i < 0$ and discount factor β has a precautionary motive that is not cancelled, so it wants to save more

than the benchmark agent at $R = \beta^{-1}$.

With positive net demand for bonds in the aggregate, the equilibrium interest rate must fall below β^{-1} to clear the market, and the equilibrium then depends on the whole distribution of types.

Exercise 88.9.3

This exercise separates quantities from beliefs for two individual agents.

Consider agents a and b with $\sigma^a < \sigma^b \leq 0$, both on the locus (88.4).

1. Show that the two agents have the same consumption innovation $h\, w_{t+1}$.
2. Show that if they start from the same (a_t, z_t) and observe the same shock w_{t+1}, their next-period consumption and assets coincide.
3. Using (88.9), compute the ratio of the two agents' worst-case forecasts of $b - c_{t+h}$ and show that it grows geometrically in h.
4. Summarise what is and is not identified by data on quantities alone.

Solution

Here is one solution.

1. Both pairs lie on (88.4), so by *Theorem 87.4.1* both use the benchmark rule and hence the same innovation vector h.
2. With a common state and a common shock, both agents apply the same policy function and the same law of motion, so $c^a_{t+1} = c^b_{t+1}$ and $a^a_{t+1} = a^b_{t+1}$.
3. From (88.9), $\widehat{\mathbb{E}}_t(b - c^j_{t+h}) = \zeta_j^h (b - c_t)$, so the ratio is $(\zeta_a/\zeta_b)^h$.

Since $\sigma^a < \sigma^b$ implies $\zeta_a > \zeta_b$, the ratio grows geometrically: the two agents' beliefs diverge without bound as the horizon lengthens, even though their actions never differ at all.

4. Quantities identify the equilibrium decision rule, and hence the single combination of parameters that appears in it.

They do not identify the decomposition of that rule into impatience β_i and robustness σ_i along the locus.

Exercise 88.9.4

This exercise asks how much belief heterogeneity is statistically plausible.

Restrict attention to types whose worst-case model has a detection error probability of at least 0.25 in a sample of $T = 40$.

1. Find the most robust admissible type $\sigma^{\min}$ by bisection.
2. For that type, report β_i, ζ_i, and the horizon h at which its worst-case forecast of $b - c_{t+h}$ is twice the approximating model's forecast.
3. Repeat part 1 with $T = 160$ and comment on what a longer sample does to the plausible amount of belief heterogeneity.

Solution

Here is one solution.

The approximating model forecasts $b - c_{t+h} = b - c_t$ for every h, while the worst-case forecast is $\zeta_i^h(b - c_t)$, so the doubling horizon solves $\zeta_i^h = 2$.

```
def σ_for_target_dep(target, T, β, α2, tol=1e-5):
    """
    Find σ ∈ (σ̲, 0) with DEP(σ) = target by bisection.

    Returns None if the DEP never falls to the target on the admissible
    range, in which case the breakdown point is the binding constraint.
    """
    α_loc = np.sqrt(α2)
    lo, hi = 0.999 * (-(1 - β)**2 / α2), 0.0

    def dep_at(σ):
        return detection_error_probability(
            worst_case_persistence(σ, β, α2), α_loc, T=T)

    if dep_at(lo) > target:
        return None

    while hi - lo > tol:
        mid = 0.5 * (lo + hi)
        if dep_at(mid) < target:
            lo = mid
        else:
            hi = mid
    return 0.5 * (lo + hi)


for T in [40, 160]:
    σ_min = σ_for_target_dep(0.25, T, β, α2)
    if σ_min is None:
        σ_min = 0.999 * σ_lo        # breakdown binds before detectability
        note = ' (breakdown point binds)'
    else:
        note = ''
    ζ_min = worst_case_persistence(σ_min, β, α2)
    print(f"T = {T:>3}:  σ_min = {σ_min:.5f}   β_i = {β / ζ_min:.5f}   "
          f"ζ_i = {ζ_min:.5f}   doubling horizon = "
          f"{np.log(2) / np.log(ζ_min):.1f} quarters{note}")
```

```
T =  40:  σ_min = -0.09313   β_i = 0.90979   ζ_i = 1.04420   doubling horizon =␣
↪16.0 quarters
T = 160:  σ_min = -0.02322   β_i = 0.93997   ζ_i = 1.01067   doubling horizon =␣
↪65.3 quarters
```

At $T = 40$ statistical detectability binds before the breakdown point does, so the most robust plausible type is interior.

At $T = 160$ it binds far sooner and the plausible set of types shrinks sharply toward $\sigma = 0$.

The doubling horizon lengthens correspondingly: with more data, only agents whose pessimism accumulates slowly remain statistically credible.

Note

The detection error probability is estimated by simulation, so a target close to the value attained at $\underline{\sigma}$ makes the answer sensitive to the random seed.
At $T = 40$ the DEP at the breakdown point is almost exactly 0.2, which is why the target here is 0.25.

88.10 Related lectures

- *The LQ Permanent Income Model* develops the standard LQ permanent income model.
- *Consumption Smoothing with Incomplete and Complete Markets* builds the benchmark Bewley economy whose equilibrium is reproduced here.
- *Robust Consumption Smoothing and Precautionary Savings* derives the observational-equivalence theorems, the breakdown point, and the detection error probabilities used in this lecture.
- *Robust Permanent Income and Pricing* shows how asset prices break the observational equivalence.

CHAPTER

EIGHTYNINE

PRODUCTION SMOOTHING VIA INVENTORIES

Contents

- *Production Smoothing via Inventories*
  - *Overview*
  - *Example 1*
  - *Inventories Not Useful*
  - *Inventories Useful but are Hardwired to be Zero Always*
  - *Example 2*
  - *Example 3*
  - *Example 4*
  - *Example 5*
  - *Example 6*
  - *Exercises*

In addition to what's in Anaconda, this lecture employs the following library:

```
!pip install quantecon
```

89.1 Overview

This lecture can be viewed as an application of this *quantecon lecture* about linear quadratic control theory.

It formulates a discounted dynamic program for a firm that chooses a production schedule to balance

- minimizing costs of production across time, against
- keeping costs of holding inventories low

In the tradition of a classic book by Holt, Modigliani, Muth, and Simon [Holt *et al.*, 1960], we simplify the firm's problem by formulating it as a linear quadratic discounted dynamic programming problem of the type studied in this *quantecon lecture*.

Because its costs of production are increasing and quadratic in production, the firm holds inventories as a buffer stock in order to smooth production across time, provided that holding inventories is not too costly.

But the firm also wants to make its sales out of existing inventories, a preference that we represent by a cost that is quadratic in the difference between sales in a period and the firm's beginning of period inventories.

We compute examples designed to indicate how the firm optimally smooths production while keeping inventories close to sales.

To introduce components of the model, let

- S_t be sales at time t
- Q_t be production at time t
- I_t be inventories at the beginning of time t
- $\beta \in (0,1)$ be a discount factor
- $c(Q_t) = c_1 Q_t + c_2 Q_t^2$, be a cost of production function, where $c_1 > 0, c_2 > 0$, be an inventory cost function
- $d(I_t, S_t) = d_1 I_t + d_2 (S_t - I_t)^2$, where $d_1 > 0, d_2 > 0$, be a cost-of-holding-inventories function, consisting of two components:
  - a cost $d_1 I_t$ of carrying inventories, and
  - a cost $d_2 (S_t - I_t)^2$ of having inventories deviate from sales
- $p_t = a_0 - a_1 S_t + v_t$ be an inverse demand function for a firm's product, where $a_0 > 0, a_1 > 0$ and v_t is a demand shock at time t
- $\pi_t = p_t S_t - c(Q_t) - d(I_t, S_t)$ be the firm's profits at time t
- $\sum_{t=0}^{\infty} \beta^t \pi_t$ be the present value of the firm's profits at time 0
- $I_{t+1} = I_t + Q_t - S_t$ be the law of motion of inventories
- $z_{t+1} = A_{22} z_t + C_2 \epsilon_{t+1}$ be a law of motion for an exogenous state vector z_t that contains time t information useful for predicting the demand shock v_t
- $v_t = G z_t$ link the demand shock to the information set z_t
- the constant 1 be the first component of z_t

To map our problem into a linear-quadratic discounted dynamic programming problem (also known as an optimal linear regulator), we define the **state** vector at time t as

$$x_t = \begin{bmatrix} I_t \\ z_t \end{bmatrix}$$

and the **control** vector as

$$u_t = \begin{bmatrix} Q_t \\ S_t \end{bmatrix}$$

The law of motion for the state vector x_t is evidently

$$\begin{bmatrix} I_{t+1} \\ z_t \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & A_{22} \end{bmatrix} \begin{bmatrix} I_t \\ z_t \end{bmatrix} + \begin{bmatrix} 1 & -1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} Q_t \\ S_t \end{bmatrix} + \begin{bmatrix} 0 \\ C_2 \end{bmatrix} \epsilon_{t+1}$$

or

$$x_{t+1} = A x_t + B u_t + C \epsilon_{t+1}$$

(At this point, we ask that you please forgive us for using Q_t to be the firm's production at time t, while below we use Q as the matrix in the quadratic form $u_t' Q u_t$ that appears in the firm's one-period profit function)

We can express the firm's profit as a function of states and controls as

$$\pi_t = -(x_t' R x_t + u_t' Q u_t + 2 u_t' N x_t)$$

To form the matrices R, Q, N in an LQ dynamic programming problem, we note that the firm's profits at time t function can be expressed

$$
\begin{aligned}
\pi_t =& p_t S_t - c\left(Q_t\right) - d\left(I_t, S_t\right) \\
=& \left(a_0 - a_1 S_t + v_t\right) S_t - c_1 Q_t - c_2 Q_t^2 - d_1 I_t - d_2 \left(S_t - I_t\right)^2 \\
=& a_0 S_t - a_1 S_t^2 + G z_t S_t - c_1 Q_t - c_2 Q_t^2 - d_1 I_t - d_2 S_t^2 - d_2 I_t^2 + 2 d_2 S_t I_t \\
=& - \left(\underbrace{d_1 I_t + d_2 I_t^2}_{x_t' R x_t} + \underbrace{a_1 S_t^2 + d_2 S_t^2 + c_2 Q_t^2}_{u_t' Q u_t} \underbrace{- a_0 S_t - G z_t S_t + c_1 Q_t - 2 d_2 S_t I_t}_{2 u_t' N x_t} \right) \\
=& - \Bigg(\begin{bmatrix} I_t & z_t' \end{bmatrix} \underbrace{\begin{bmatrix} d_2 & \frac{d_1}{2} S_c \\ \frac{d_1}{2} S_c' & 0 \end{bmatrix}}_{\equiv R} \begin{bmatrix} I_t \\ z_t \end{bmatrix} + \begin{bmatrix} Q_t & S_t \end{bmatrix} \underbrace{\begin{bmatrix} c_2 & 0 \\ 0 & a_1 + d_2 \end{bmatrix}}_{\equiv Q} \begin{bmatrix} Q_t \\ S_t \end{bmatrix} + 2 \begin{bmatrix} Q_t & S_t \end{bmatrix} \underbrace{\begin{bmatrix} 0 & \frac{c_1}{2} S_c \\ -d_2 & -\frac{a_0}{2} S_c - \frac{G}{2} \end{bmatrix}}_{\equiv N} \begin{bmatrix} I \\ z \end{bmatrix}
\end{aligned}
$$

where $S_c = [1, 0]$.

Remark on notation: The notation for cross product term in the QuantEcon library is N.

The firms' optimum decision rule takes the form

$$u_t = -F x_t$$

and the evolution of the state under the optimal decision rule is

$$x_{t+1} = (A - BF) x_t + C \epsilon_{t+1}$$

The firm chooses a decision rule for u_t that maximizes

$$E_0 \sum_{t=0}^{\infty} \beta^t \pi_t$$

subject to a given x_0.

This is a stochastic discounted LQ dynamic program.

Here is code for computing an optimal decision rule and for analyzing its consequences.

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
```

```
class SmoothingExample:
    """
    Class for constructing, solving, and plotting results for
    inventories and sales smoothing problem.
    """

    def __init__(self,
                 β=0.96,            # Discount factor
                 c1=1,              # Cost-of-production
                 c2=1,
                 d1=1,              # Cost-of-holding inventories
                 d2=1,
                 a0=10,             # Inverse demand function
                 a1=1,
                 A22=[[1,   0],     # z process
```

(continues on next page)

(continued from previous page)

```
                    [1, 0.9]],
              C2=[[0], [1]],
              G=[0, 1]):

        self.β = β
        self.c1, self.c2 = c1, c2
        self.d1, self.d2 = d1, d2
        self.a0, self.a1 = a0, a1
        self.A22 = np.atleast_2d(A22)
        self.C2 = np.atleast_2d(C2)
        self.G = np.atleast_2d(G)

        # Dimensions
        k, j = self.C2.shape        # Dimensions for randomness part
        n = k + 1                   # Number of states
        m = 2                       # Number of controls

        Sc = np.zeros(k)
        Sc[0] = 1

        # Construct matrices of transition law
        A = np.zeros((n, n))
        A[0, 0] = 1
        A[1:, 1:] = self.A22

        B = np.zeros((n, m))
        B[0, :] = 1, -1

        C = np.zeros((n, j))
        C[1:, :] = self.C2

        self.A, self.B, self.C = A, B, C

        # Construct matrices of one period profit function
        R = np.zeros((n, n))
        R[0, 0] = d2
        R[1:, 0] = d1 / 2 * Sc
        R[0, 1:] = d1 / 2 * Sc

        Q = np.zeros((m, m))
        Q[0, 0] = c2
        Q[1, 1] = a1 + d2

        N = np.zeros((m, n))
        N[1, 0] = - d2
        N[0, 1:] = c1 / 2 * Sc
        N[1, 1:] = - a0 / 2 * Sc - self.G / 2

        self.R, self.Q, self.N = R, Q, N

        # Construct LQ instance
        self.LQ = qe.LQ(Q, R, A, B, C, N, beta=β)
        self.LQ.stationary_values()

    def simulate(self, x0, T=100):

        c1, c2 = self.c1, self.c2
```

(continues on next page)

(continued from previous page)

```
        d1, d2 = self.d1, self.d2
        a0, a1 = self.a0, self.a1
        G = self.G

        x_path, u_path, w_path = self.LQ.compute_sequence(x0, ts_length=T)

        I_path = x_path[0, :-1]
        z_path = x_path[1:, :-1]
        ν_path = (G @ z_path)[0, :]

        Q_path = u_path[0, :]
        S_path = u_path[1, :]

        revenue = (a0 - a1 * S_path + ν_path) * S_path
        cost_production = c1 * Q_path + c2 * Q_path ** 2
        cost_inventories = d1 * I_path + d2 * (S_path - I_path) ** 2

        Q_no_inventory = (a0 + ν_path - c1) / (2 * (a1 + c2))
        Q_hardwired = (a0 + ν_path - c1) / (2 * (a1 + c2 + d2))

        fig, ax = plt.subplots(2, 2, figsize=(15, 10))

        ax[0, 0].plot(range(T), I_path, label="inventories")
        ax[0, 0].plot(range(T), S_path, label="sales")
        ax[0, 0].plot(range(T), Q_path, label="production")
        ax[0, 0].legend(loc=1)
        ax[0, 0].set_title("inventories, sales, and production")

        ax[0, 1].plot(range(T), (Q_path - S_path), color='b')
        ax[0, 1].set_ylabel("change in inventories", color='b')
        span = max(abs(Q_path - S_path))
        ax[0, 1].set_ylim(0-span*1.1, 0+span*1.1)
        ax[0, 1].set_title("demand shock and change in inventories")

        ax1_ = ax[0, 1].twinx()
        ax1_.plot(range(T), ν_path, color='r')
        ax1_.set_ylabel("demand shock", color='r')
        span = max(abs(ν_path))
        ax1_.set_ylim(0-span*1.1, 0+span*1.1)

        ax1_.plot([0, T], [0, 0], '--', color='k')

        ax[1, 0].plot(range(T), revenue, label="revenue")
        ax[1, 0].plot(range(T), cost_production, label="cost_production")
        ax[1, 0].plot(range(T), cost_inventories, label="cost_inventories")
        ax[1, 0].legend(loc=1)
        ax[1, 0].set_title("profits decomposition")

        ax[1, 1].plot(range(T), Q_path, label="production")
        ax[1, 1].plot(range(T), Q_hardwired, label='production when  $I_t$ \
            forced to be zero')
        ax[1, 1].plot(range(T), Q_no_inventory, label='production when \
            inventories not useful')
        ax[1, 1].legend(loc=1)
        ax[1, 1].set_title('three production concepts')

        plt.show()
```

Notice that the above code sets parameters at the following default values

- discount factor $\beta = 0.96$,
- inverse demand function: $a0 = 10, a1 = 1$
- cost of production $c1 = 1, c2 = 1$
- costs of holding inventories $d1 = 1, d2 = 1$

In the examples below, we alter some or all of these parameter values.

89.2 Example 1

In this example, the demand shock follows AR(1) process:

$$\nu_t = \alpha + \rho \nu_{t-1} + \epsilon_t,$$

which implies

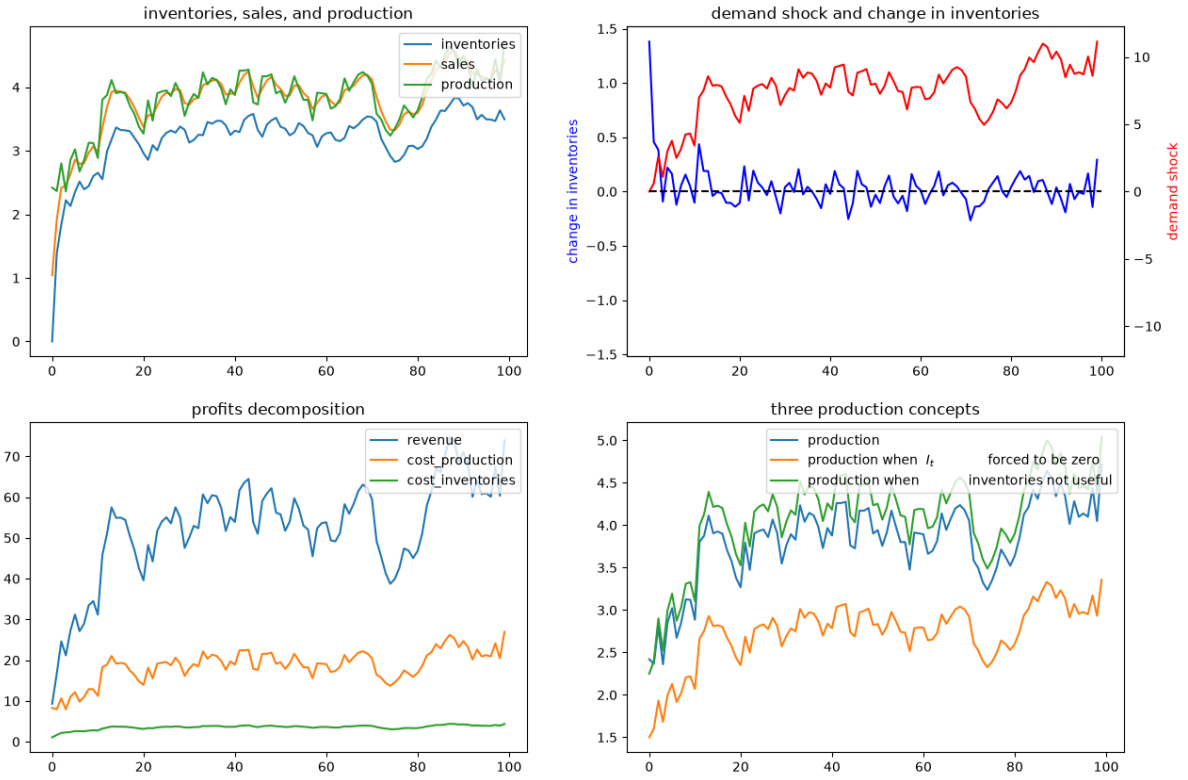
$$z_{t+1} = \begin{bmatrix} 1 \\ v_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \alpha & \rho \end{bmatrix} \underbrace{\begin{bmatrix} 1 \\ v_t \end{bmatrix}}_{z_t} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} \epsilon_{t+1}.$$

We set $\alpha = 1$ and $\rho = 0.9$, their default values.

We'll calculate and display outcomes, then discuss them below the pertinent figures.

```
ex1 = SmoothingExample()

x0 = [0, 1, 0]
ex1.simulate(x0)
```



The figures above illustrate various features of an optimal production plan.

Starting from zero inventories, the firm builds up a stock of inventories and uses them to smooth costly production in the face of demand shocks.

Optimal decisions evidently respond to demand shocks.

Inventories are always less than sales, so some sales come from current production, a consequence of the cost, $d_1 I_t$ of holding inventories.

The lower right panel shows differences between optimal production and two alternative production concepts that come from altering the firm's cost structure – i.e., its technology.

These two concepts correspond to these distinct altered firm problems.

- a setting in which inventories are not needed
- a setting in which they are needed but we arbitrarily prevent the firm from holding inventories by forcing it to set $I_t = 0$ always

We use these two alternative production concepts in order to shed light on the baseline model.

89.3 Inventories Not Useful

Let's turn first to the setting in which inventories aren't needed.

In this problem, the firm forms an output plan that maximizes the expected value of

$$\sum_{t=0}^{\infty} \beta^t \{p_t Q_t - C(Q_t)\}$$

It turns out that the optimal plan for Q_t for this problem also solves a sequence of static problems $\max_{Q_t} \{p_t Q_t - c(Q_t)\}$.

When inventories aren't required or used, sales always equal production.

This simplifies the problem and the optimal no-inventory production maximizes the expected value of

$$\sum_{t=0}^{\infty} \beta^t \left\{p_t Q_t - C\left(Q_t\right)\right\}.$$

The optimum decision rule is

$$Q_t^{ni} = \frac{a_0 + \nu_t - c_1}{c_2 + a_1}.$$

89.4 Inventories Useful but are Hardwired to be Zero Always

Next, we turn to a distinct problem in which inventories are useful – meaning that there are costs of $d_2(I_t - S_t)^2$ associated with having sales not equal to inventories – but we arbitrarily impose on the firm the costly restriction that it never hold inventories.

Here the firm's maximization problem is

$$\max_{\{I_t, Q_t, S_t\}} \sum_{t=0}^{\infty} \beta^t \left\{p_t S_t - C\left(Q_t\right) - d\left(I_t, S_t\right)\right\}$$

subject to the restrictions that $I_t = 0$ for all t and that $I_{t+1} = I_t + Q_t - S_t$.

The restriction that $I_t = 0$ implies that $Q_t = S_t$ and that the maximization problem reduces to

$$\max_{Q_t} \sum_{t=0}^{\infty} \beta^t \left\{ p_t Q_t - C\left(Q_t\right) - d\left(0, Q_t\right) \right\}$$

Here the optimal production plan is

$$Q_t^h = \frac{a_0 + \nu_t - c_1}{c_2 + a_1 + d_2}.$$

We introduce this I_t **is hardwired to zero** specification in order to shed light on the role that inventories play by comparing outcomes with those under our two other versions of the problem.

The bottom right panel displays a production path for the original problem that we are interested in (the blue line) as well with an optimal production path for the model in which inventories are not useful (the green path) and also for the model in which, although inventories are useful, they are hardwired to zero and the firm pays cost $d(0, Q_t)$ for not setting sales $S_t = Q_t$ equal to zero (the orange line).

Notice that it is typically optimal for the firm to produce more when inventories aren't useful. Here there is no requirement to sell out of inventories and no costs from having sales deviate from inventories.

But "typical" does not mean "always".

Thus, if we look closely, we notice that for small t, the green "production when inventories aren't useful" line in the lower right panel is below optimal production in the original model.

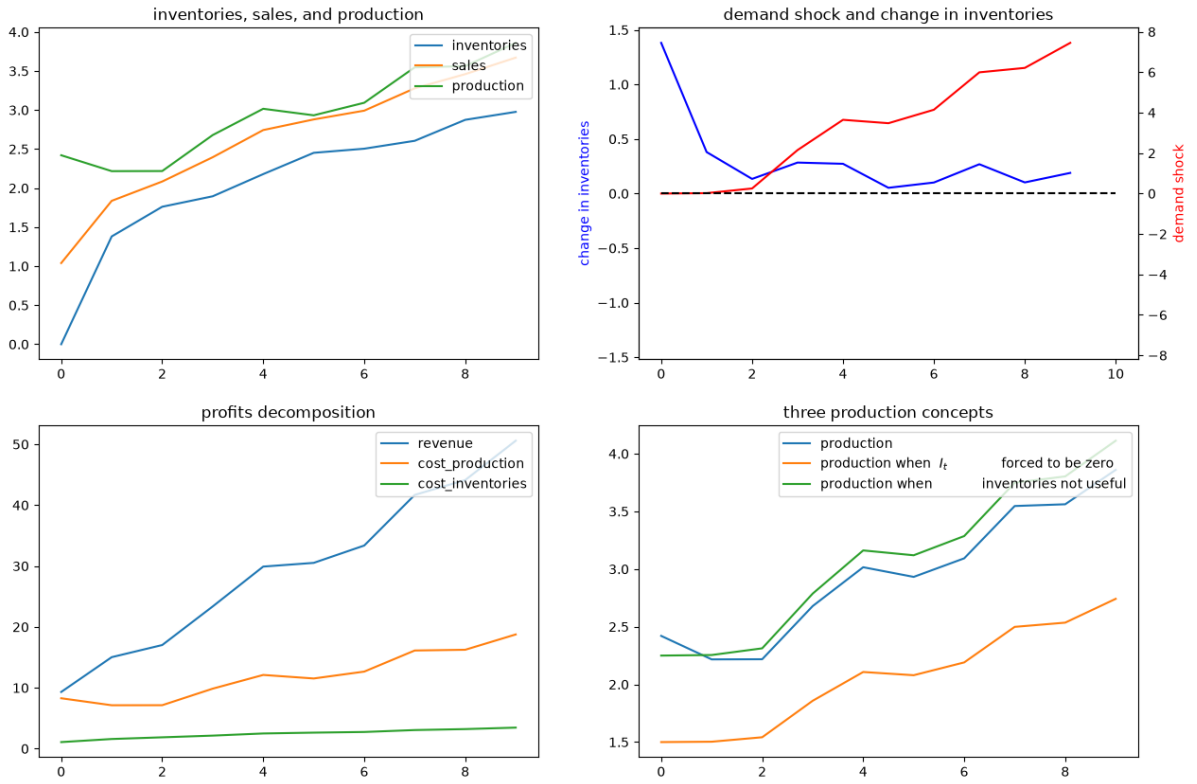
High optimal production in the original model early on occurs because the firm wants to accumulate inventories quickly in order to acquire high inventories for use in later periods.

But how the green line compares to the blue line early on depends on the evolution of the demand shock, as we will see in a deterministically seasonal demand shock example to be analyzed below.

In that example, the original firm optimally accumulates inventories slowly because the next positive demand shock is in the distant future.

To make the green-blue model production comparison easier to see, let's confine the graphs to the first 10 periods:

```
ex1.simulate(x0, T=10)
```



89.5 Example 2

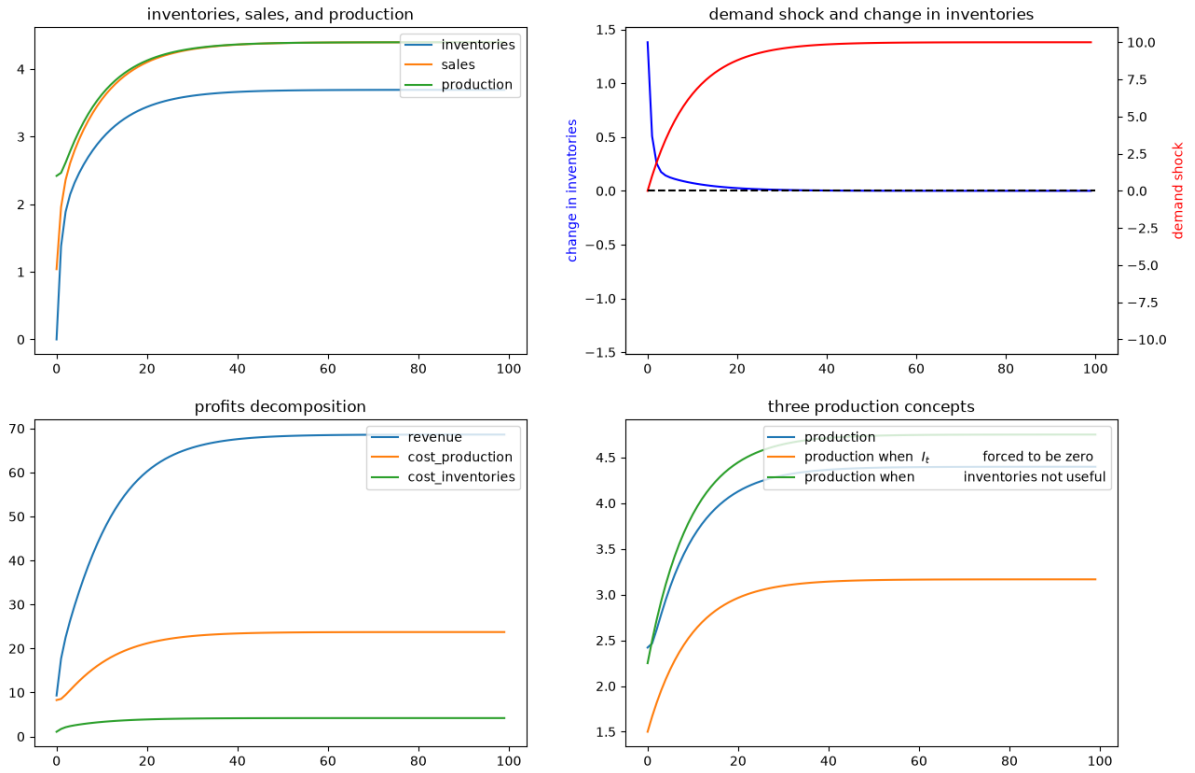
Next, we shut down randomness in demand and assume that the demand shock ν_t follows a deterministic path:

$$\nu_t = \alpha + \rho \nu_{t-1}$$

Again, we'll compute and display outcomes in some figures

```
ex2 = SmoothingExample(C2=[[0], [0]])

x0 = [0, 1, 0]
ex2.simulate(x0)
```



89.6 Example 3

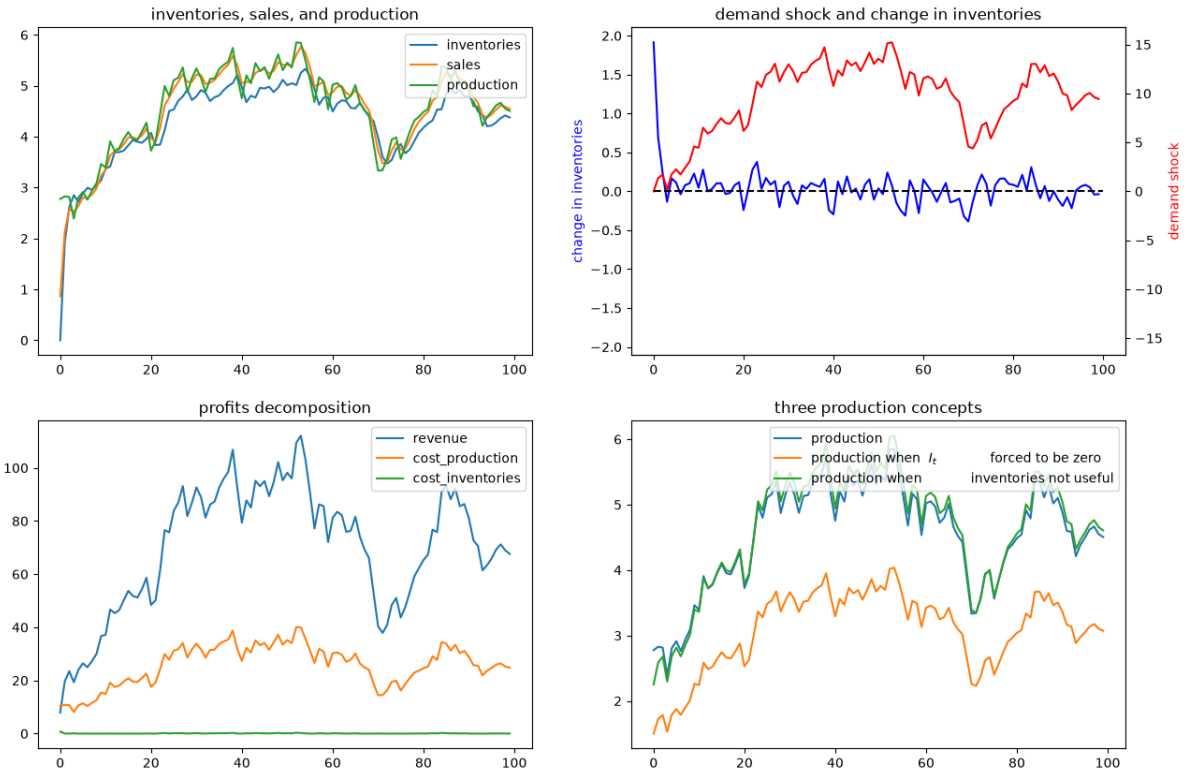
Now we'll put randomness back into the demand shock process and also assume that there are zero costs of holding inventories.

In particular, we'll look at a situation in which $d_1 = 0$ but $d_2 > 0$.

Now it becomes optimal to set sales approximately equal to inventories and to use inventories to smooth production quite well, as the following figures confirm

```
ex3 = SmoothingExample(d1=0)

x0 = [0, 1, 0]
ex3.simulate(x0)
```



89.7 Example 4

To bring out some features of the optimal policy that are related to some technical issues in linear control theory, we'll now temporarily assume that it is costless to hold inventories.

When we completely shut down the cost of holding inventories by setting $d_1 = 0$ and $d_2 = 0$, something absurd happens (because the Bellman equation is opportunistic and very smart).

(Technically, we have set parameters that end up violating conditions needed to assure **stability** of the optimally controlled state.)

The firm finds it optimal to set $Q_t \equiv Q^* = \frac{-c_1}{2c_2}$, an output level that sets the costs of production to zero (when $c_1 > 0$, as it is with our default settings, then it is optimal to set production negative, whatever that means!).

Recall the law of motion for inventories

$$I_{t+1} = I_t + Q_t - S_t$$

So when $d_1 = d_2 = 0$ so that the firm finds it optimal to set $Q_t = \frac{-c_1}{2c_2}$ for all t, then

$$I_{t+1} - I_t = \frac{-c_1}{2c_2} - S_t < 0$$

for almost all values of S_t under our default parameters that keep demand positive almost all of the time.

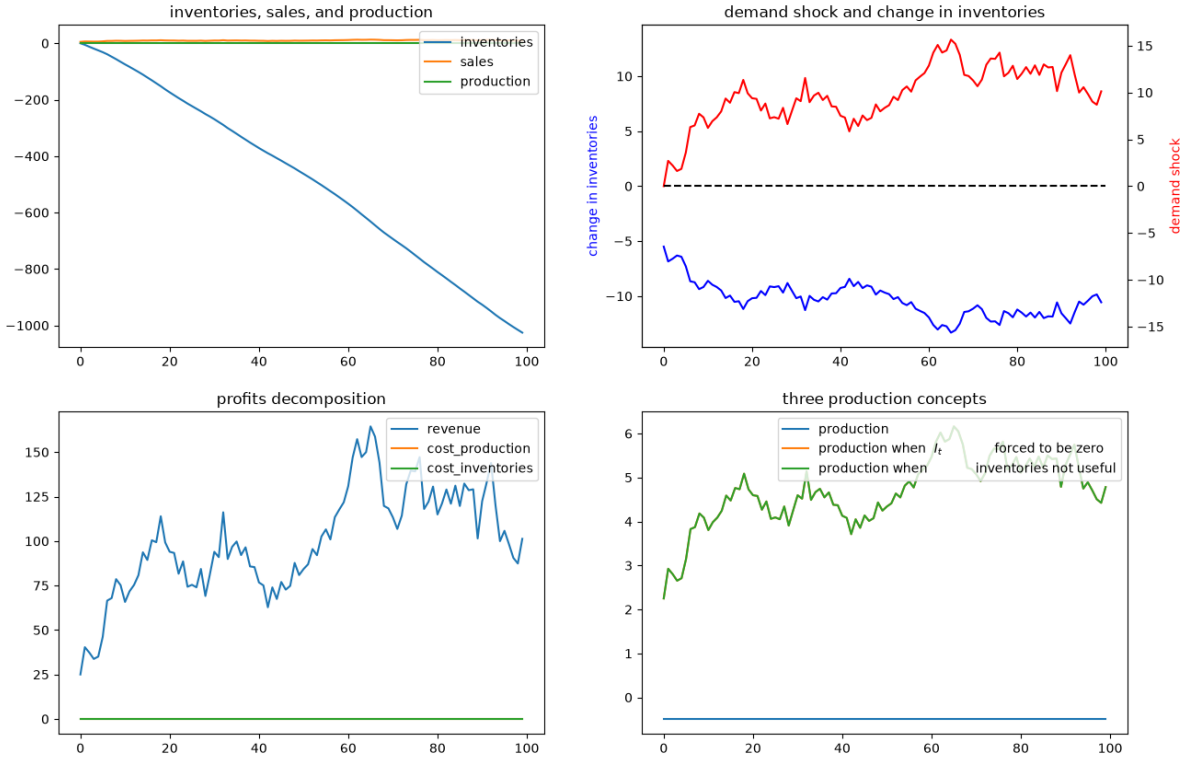
The dynamic program instructs the firm to set production costs to zero and to **run a Ponzi scheme** by running inventories down forever.

(We can interpret this as the firm somehow **going short in** or **borrowing** inventories)

The following figures confirm that inventories head south without limit

```
ex4 = SmoothingExample(d1=0, d2=0)

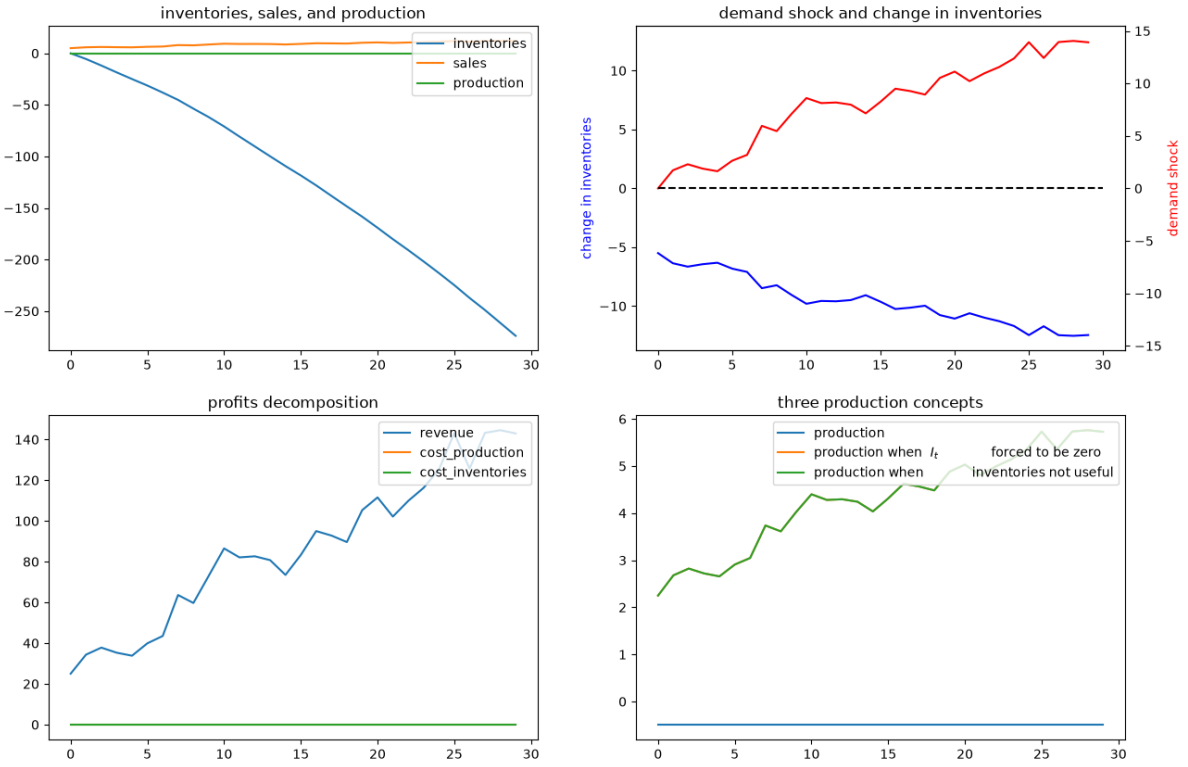
x0 = [0, 1, 0]
ex4.simulate(x0)
```



Let's shorten the time span displayed in order to highlight what is going on.

We'll set the horizon $T = 30$ with the following code

```
# shorter period
ex4.simulate(x0, T=30)
```



89.8 Example 5

Now we'll assume that the demand shock that follows a linear time trend

$$v_t = b + at, a > 0, b > 0$$

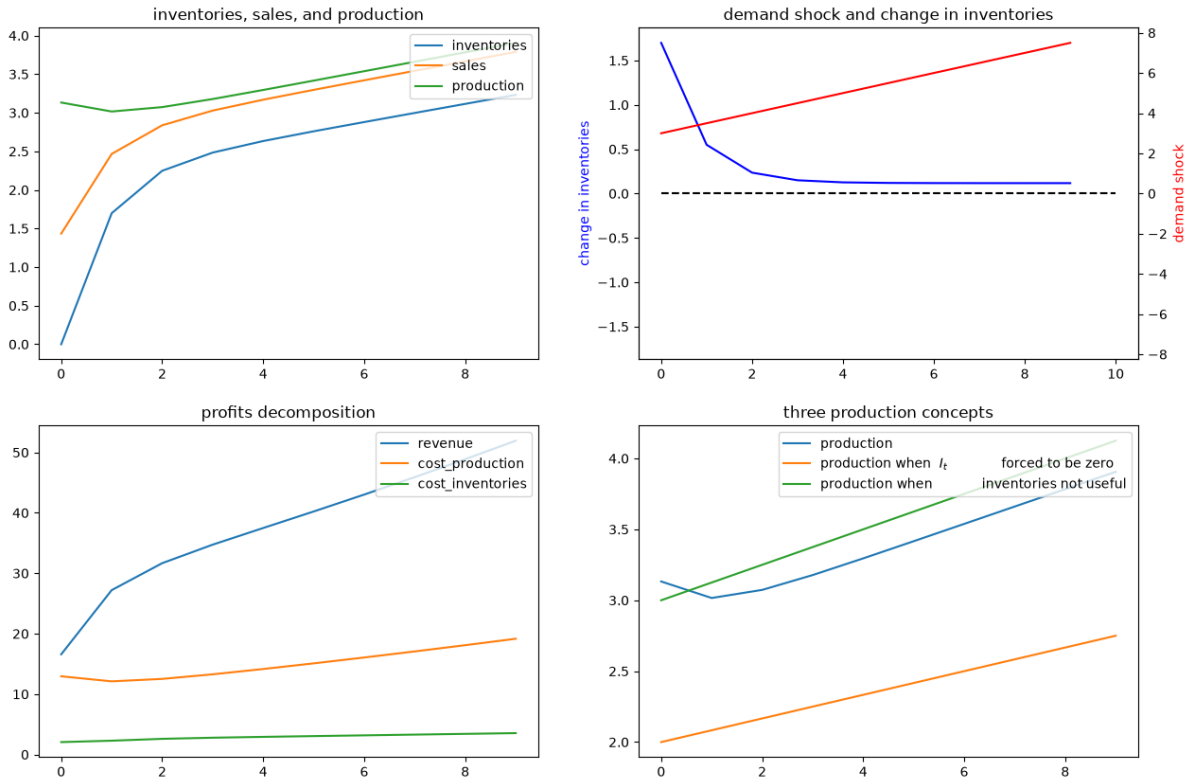
To represent this, we set $C_2 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$ and

$$A_{22} = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, x_0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, G = \begin{bmatrix} b & a \end{bmatrix}$$

```
# Set parameters
a = 0.5
b = 3.
```

```
ex5 = SmoothingExample(A22=[[1, 0], [1, 1]], C2=[[0], [0]], G=[b, a])

x0 = [0, 1, 0] # set the initial inventory as 0
ex5.simulate(x0, T=10)
```



89.9 Example 6

Now we'll assume a deterministically seasonal demand shock.

To represent this we'll set

$$A_{22} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \end{bmatrix}, C_2 = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}, G' = \begin{bmatrix} b \\ a \\ 0 \\ 0 \\ 0 \end{bmatrix}$$

where $a > 0, b > 0$ and

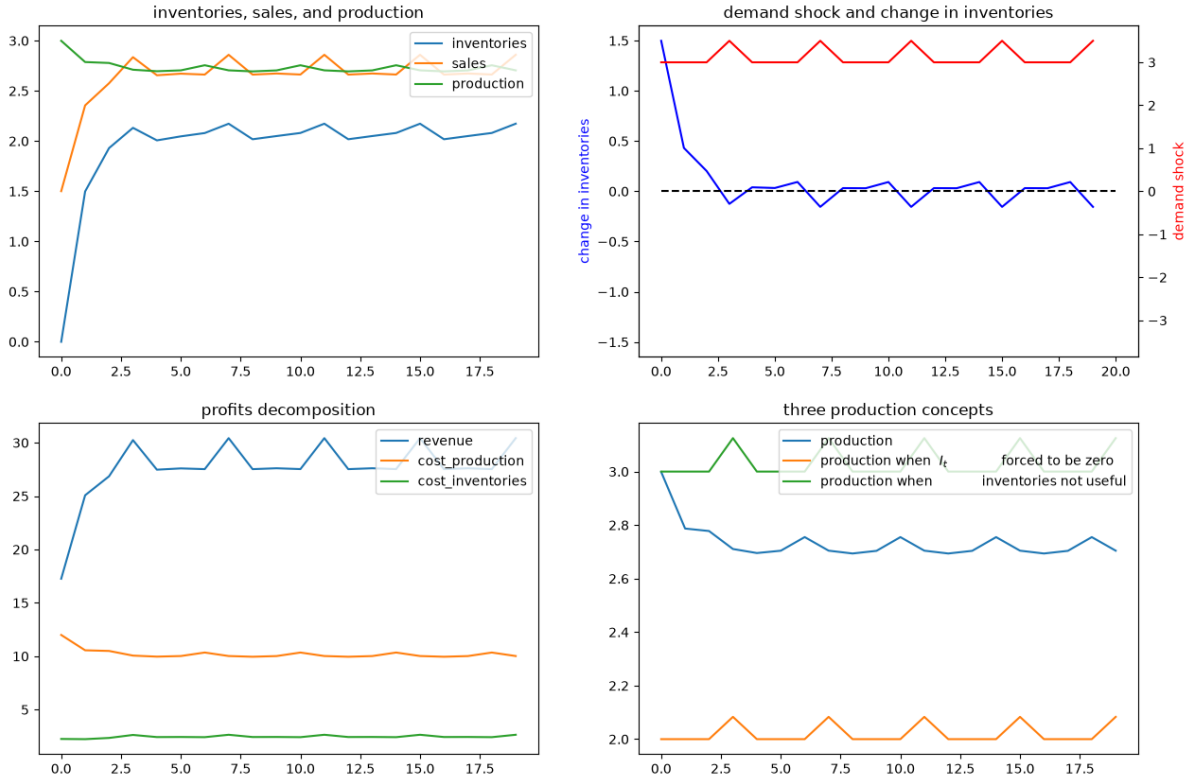
$$x_0 = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 0 \\ 0 \end{bmatrix}$$

```
ex6 = SmoothingExample(A22=[[1, 0, 0, 0, 0],
                           [0, 0, 0, 0, 1],
                           [0, 1, 0, 0, 0],
                           [0, 0, 1, 0, 0],
                           [0, 0, 0, 1, 0]],
                       C2=[[0], [0], [0], [0], [0]],
                       G=[b, a, 0, 0, 0])
```

(continues on next page)

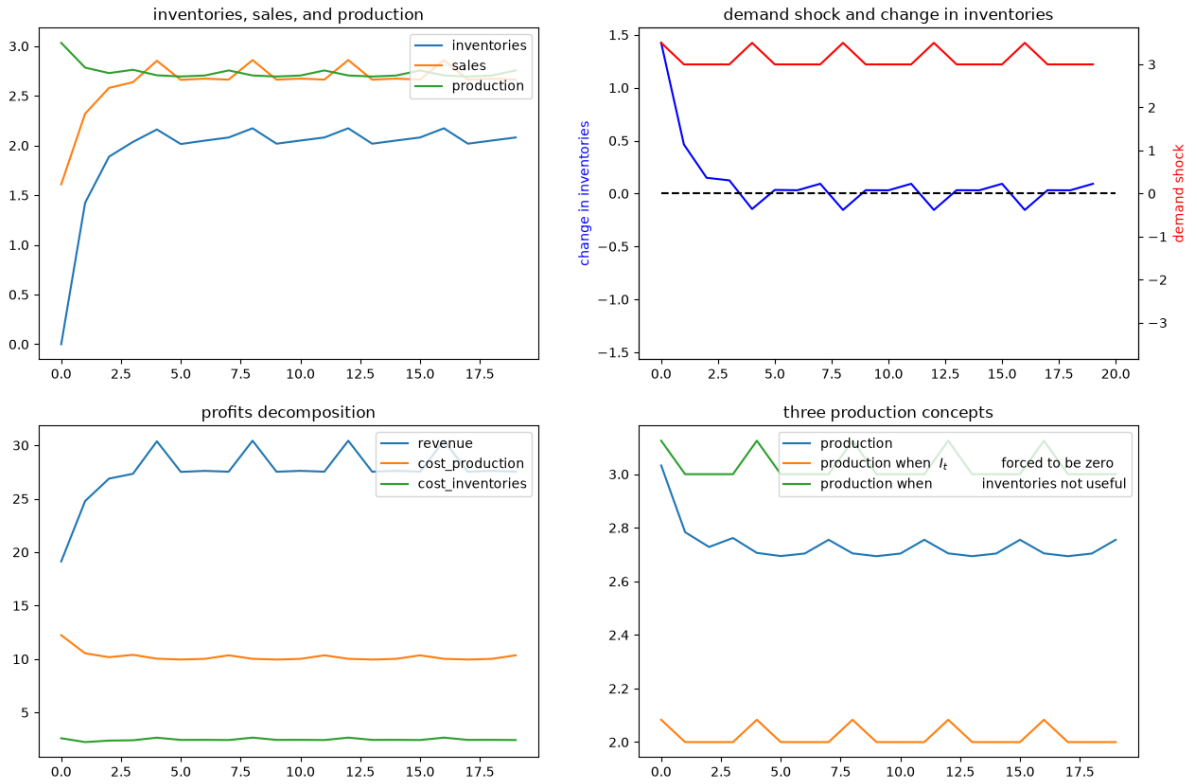
(continued from previous page)

```
x00 = [0, 1, 0, 1, 0, 0] # Set the initial inventory as 0
ex6.simulate(x00, T=20)
```

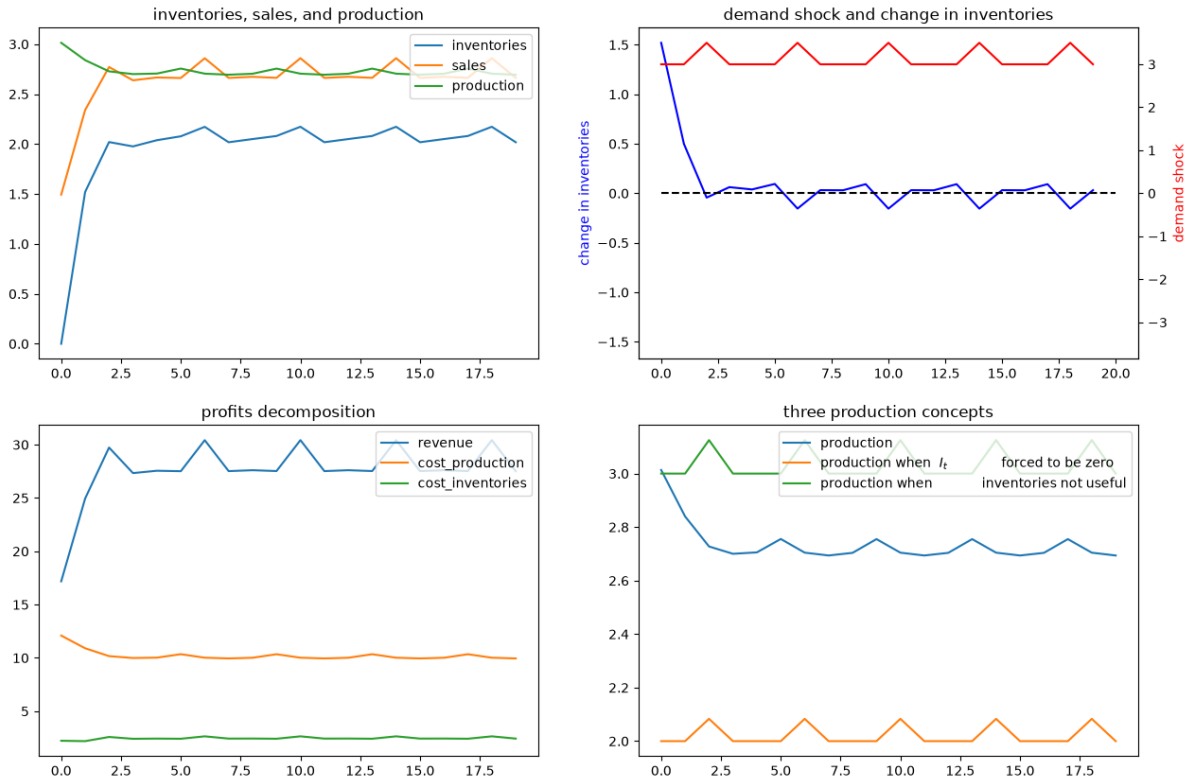


Now we'll generate some more examples that differ simply from the initial **season** of the year in which we begin the demand shock

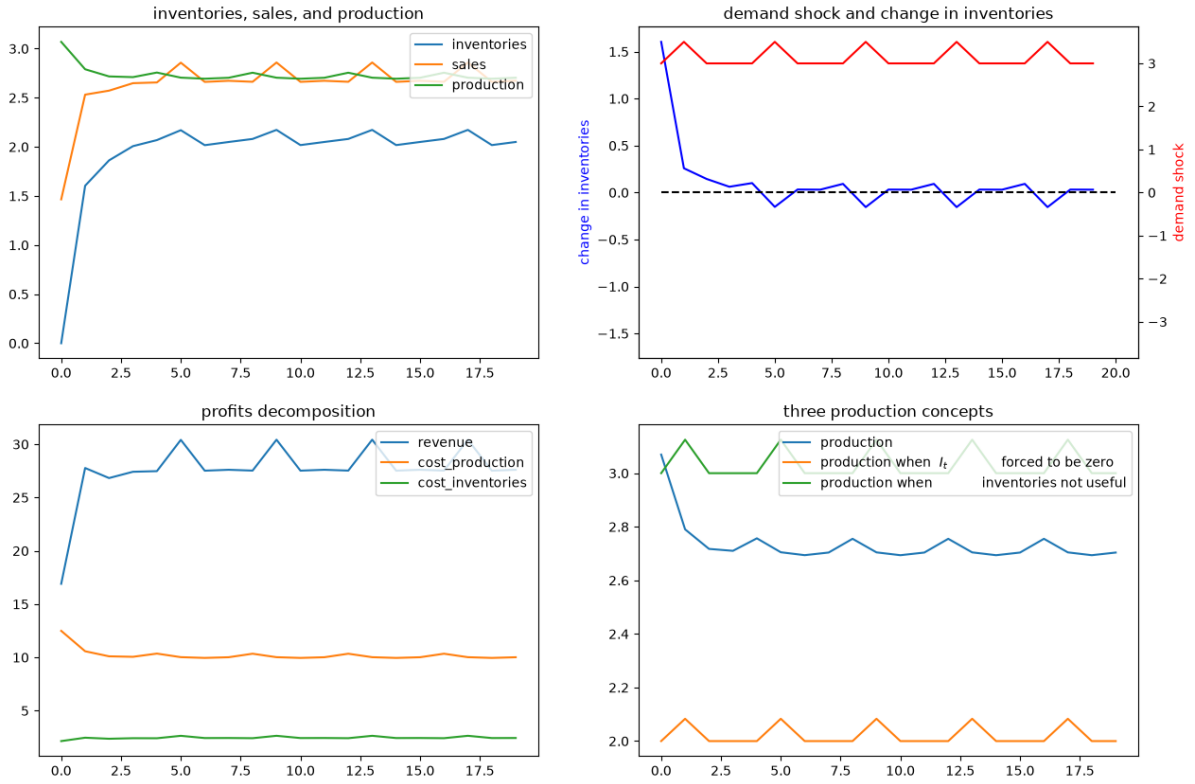
```
x01 = [0, 1, 1, 0, 0, 0]
ex6.simulate(x01, T=20)
```




```
x02 = [0, 1, 0, 0, 1, 0]
ex6.simulate(x02, T=20)
```



```
x03 = [0, 1, 0, 0, 0, 1]
ex6.simulate(x03, T=20)
```



89.10 Exercises

Please try to analyze some inventory sales smoothing problems using the `SmoothingExample` class.

Exercise 89.10.1

Assume that the demand shock follows AR(2) process below:

$$\nu_t = \alpha + \rho_1 \nu_{t-1} + \rho_2 \nu_{t-2} + \epsilon_t.$$

where $\alpha = 1$, $\rho_1 = 1.2$, and $\rho_2 = -0.3$. You need to construct $A22$, C, and G matrices properly and then to input them as the keyword arguments of `SmoothingExample` class. Simulate paths starting from the initial condition $x_0 = [0, 1, 0, 0]'$.

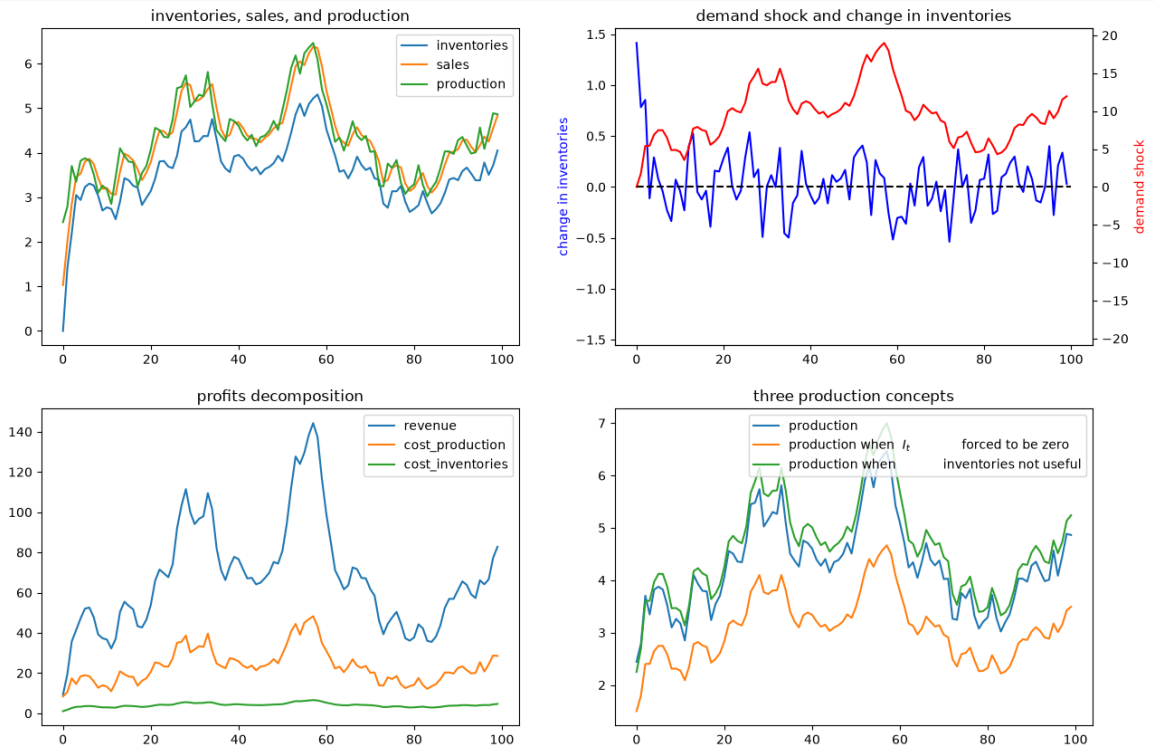
After this, try to construct a very similar `SmoothingExample` with the same demand shock process but exclude the randomness ϵ_t. Compute the stationary states $\bar{x}$ by simulating for a long period. Then try to add shocks with different magnitude to $\bar{\nu}_t$ and simulate paths. You should see how firms respond differently by staring at the production plans.

Solution

```
# set parameters
a = 1
```

```
ρ1 = 1.2
ρ2 = -.3
# construct matrices
A22 =[[1,  0,  0],
      [1, ρ1, ρ2],
      [0,  1,  0]]
C2 = [[0], [1], [0]]
G = [0, 1, 0]
ex1 = SmoothingExample(A22=A22, C2=C2, G=G)

x0 = [0, 1, 0, 0] # initial condition
ex1.simulate(x0)
```




```
# now silence the noise
ex1_no_noise = SmoothingExample(A22=A22, C2=[[0], [0], [0]], G=G)

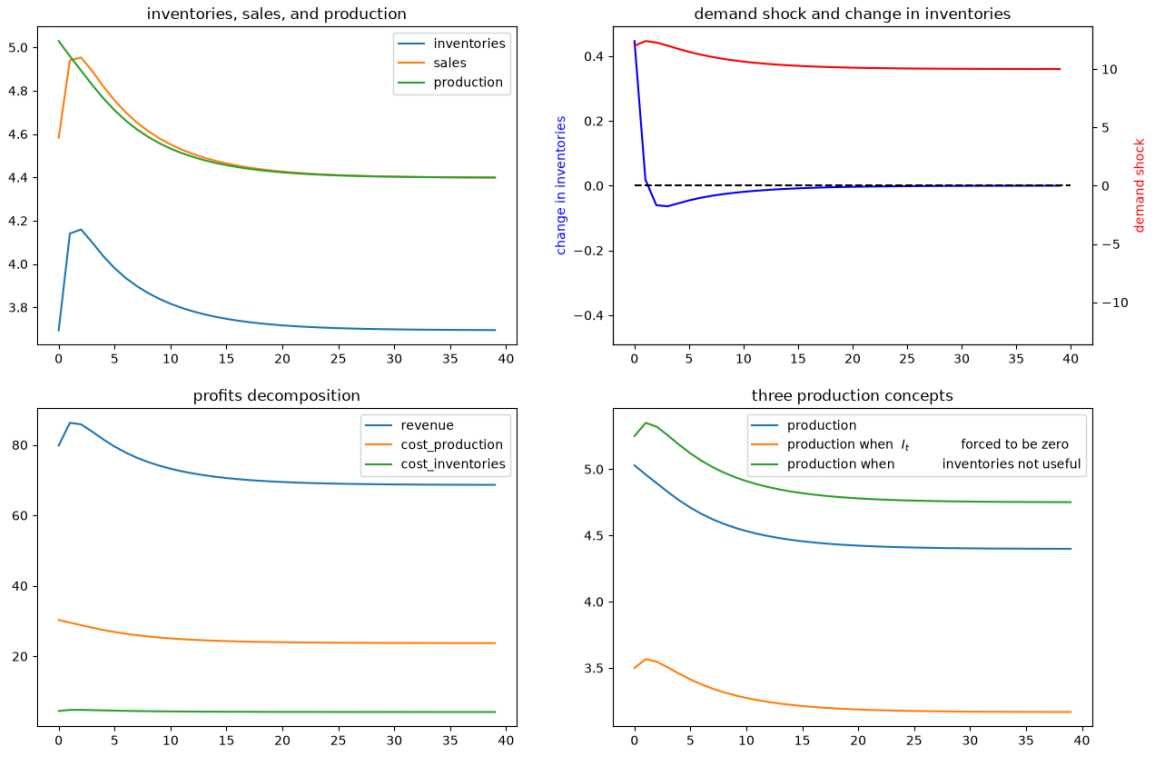
# initial condition
x0 = [0, 1, 0, 0]

# compute stationary states
x_bar = ex1_no_noise.LQ.compute_sequence(x0, ts_length=250)[0][:, -1]
x_bar
```

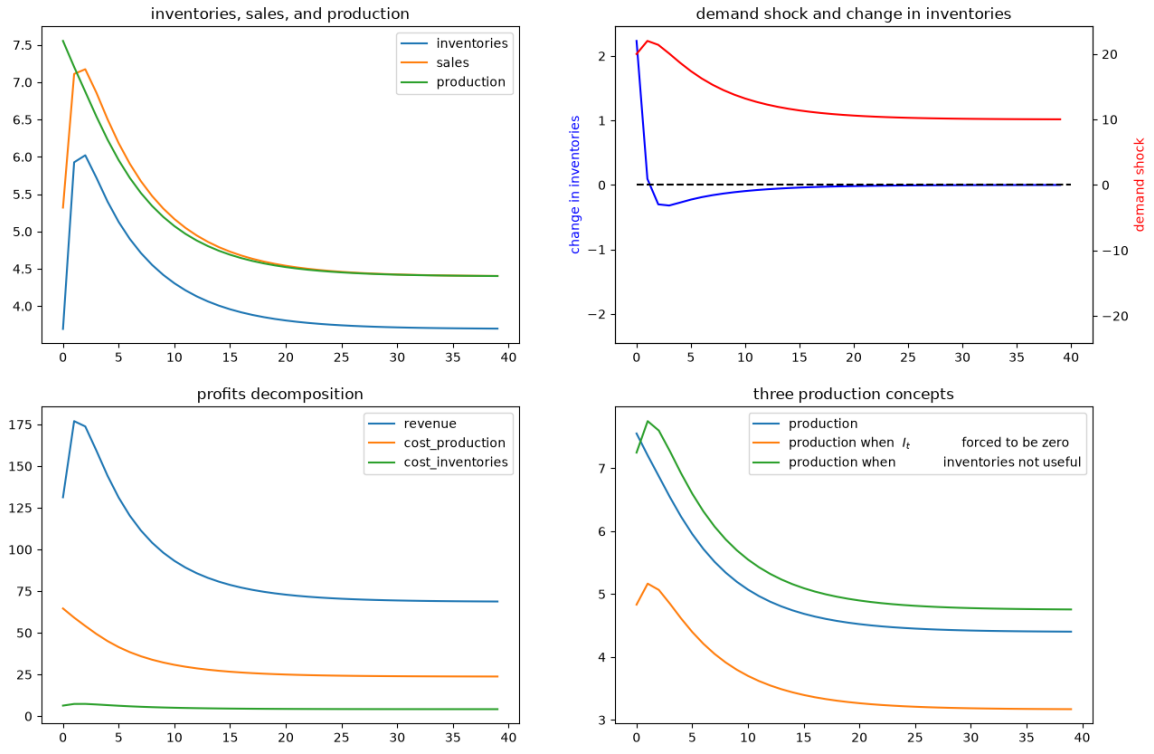
```
array([ 3.69387755,  1.        , 10.        , 10.        ])
```

In the following, we add small and large shocks to $\bar{\nu}_t$ and compare how firm responds differently in quantity. As the shock is not very persistent under the parameterization we are using, we focus on a short period response.

```
T = 40
# small shock
x_bar1 = x_bar.copy()
x_bar1[2] += 2
ex1_no_noise.simulate(x_bar1, T=T)
```




```
# large shock
x_bar1 = x_bar.copy()
x_bar1[2] += 10
ex1_no_noise.simulate(x_bar1, T=T)
```



Exercise 89.10.2

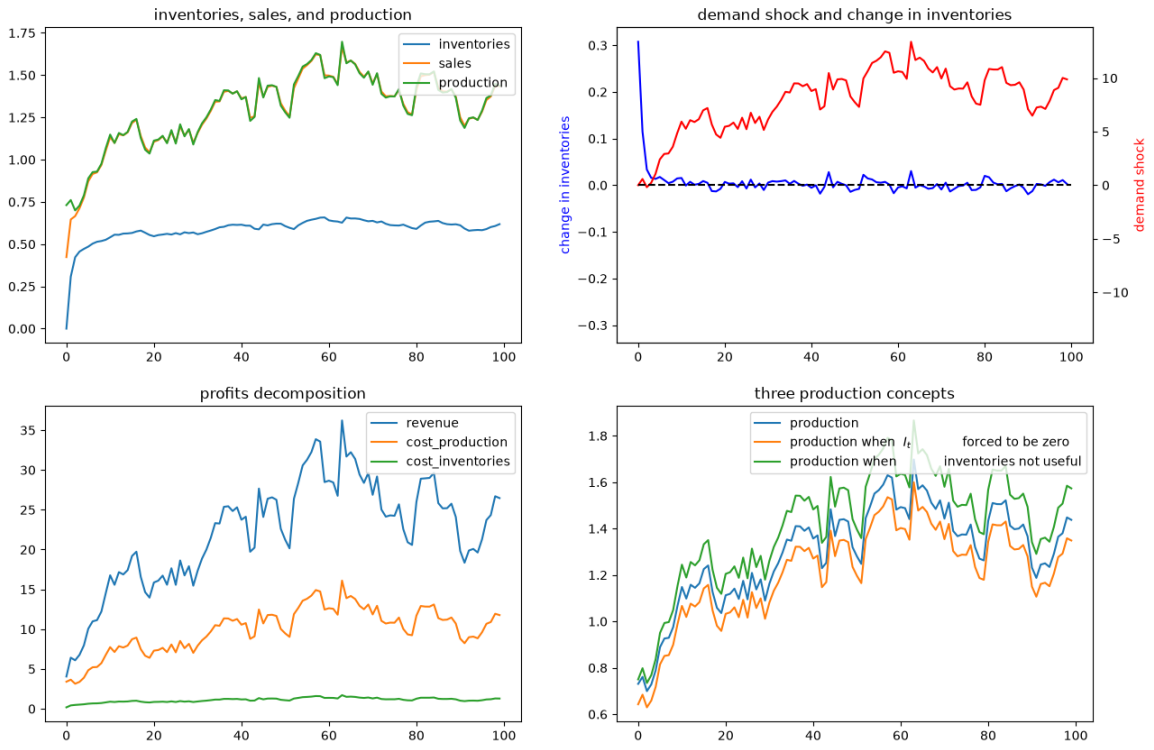
Change parameters of $C(Q_t)$ and $d(I_t, S_t)$.

1. Make production more costly, by setting $c_2 = 5$.
2. Increase the cost of having inventories deviate from sales, by setting $d_2 = 5$.

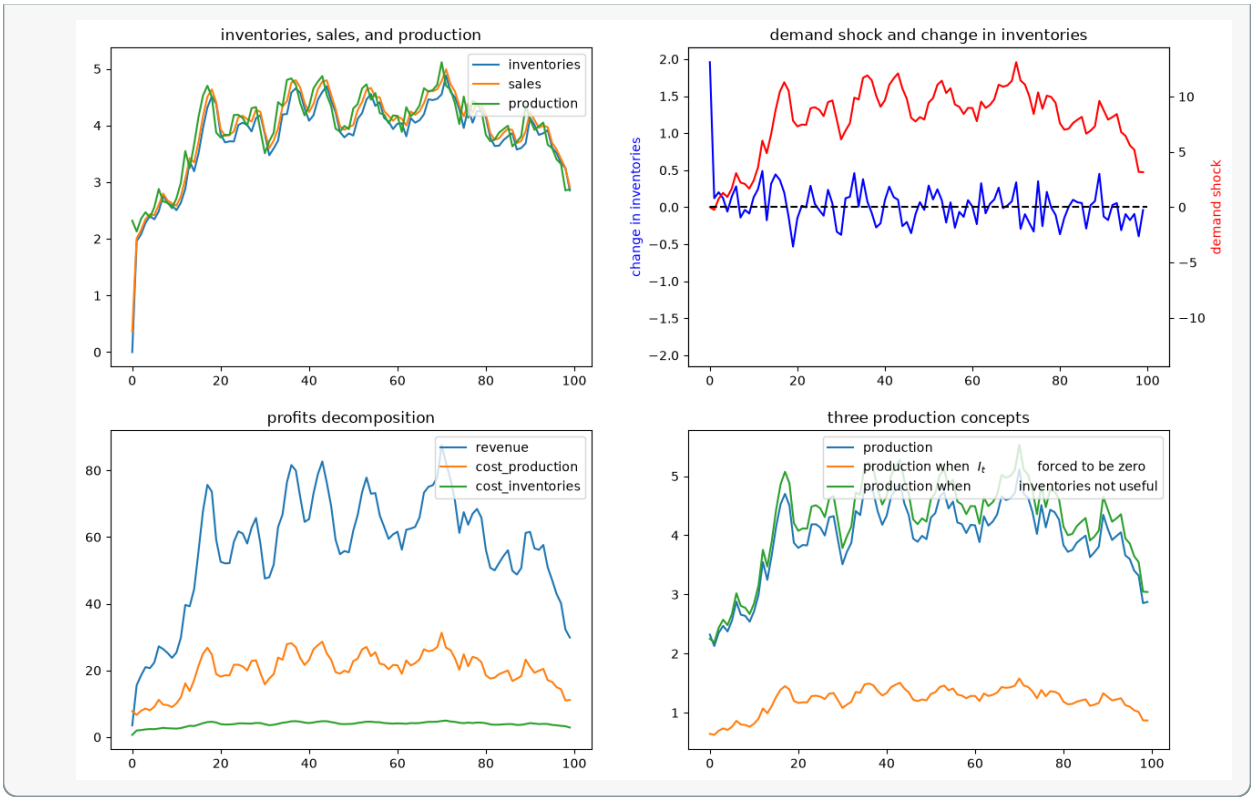
Solution

```
x0 = [0, 1, 0]
```

```
SmoothingExample(c2=5).simulate(x0)
```




```
SmoothingExample(d2=5).simulate(x0)
```


inventories, sales, and production
inventories
sales
production
demand shock and change in inventories
change in inventories
demand shock
profits decomposition
revenue
cost_production
cost_inventories
three production concepts
production
production when I_t
forced to be zero
production when
inventories not useful

Part XII

Bounded Rationality in Macroeconomics

CHAPTER

NINETY

A PECULIAR DEFINITION OF BOUNDED RATIONALITY

Contents

- *A Peculiar Definition of Bounded Rationality*
  - *Overview*
  - *Rational expectations as a fixed point*
  - *One recursion, many learners*
  - *Money and prices*
  - *Two currencies*
  - *Where this leaves us*
  - *Exercises*

90.1 Overview

This lecture opens a series based on Sargent [1993] about Bounded Rationality in Macroeconomics, from one point of view in the late 1980s and early 1990s.

Those were years when the Soviet Union ended and formerly Warsaw Pact countries were rearranging their political economies.

For those countries, those were thus truly *regime changes* in the technical sense of modern dynamic macroeconomics.

As *What the program is good for* explains, rational expectations models handle such changes less comfortably than their users usually suppose.

Two generations of work on economic dynamics — in game theory, macroeconomics, and general equilibrium theory — had produced theories that embraced the rational expectations assumption, models that were designed to understand settings in which people face recurrent situations they have lived through many times before.

The transitions across regimes underway in Eastern Europe in the early 1990s were not like that.

People there were confronted with unprecedented opportunities, new and ill-defined rules, and a daily struggle to figure out the mechanism that would eventually govern trade and production.

Economists with good models of a market economy had ample *equilibrium theories* describing how a system behaves once it has fully adjusted to a new and coherent set of rules and expectations.

They knew much less about the dynamics of *transitions* from a Soviet system to a market economy.

They might hold prejudices and anecdotes about how to manage such a transition, but no empirically confirmed formal theory of it.

Against this background, some economists ventured into what Sims [1980] called the "wilderness" of irrational expectations and bounded rationality.

The aim was partly to build theories of transition dynamics, partly to understand the properties of equilibrium dynamics themselves, and partly to study systems that never settle down.

This series follows Sargent [1993] a little way into that wilderness.

90.1.1 The knowledge that rational expectations imputes

To see what bounded rationality retreats *from*, start with rational expectations, which imposes **two** requirements:

1. **Individual rationality** — each artificial agent's behavior maximizes an objective function subject to perceived constraints.
2. **Mutual consistency** — the constraints perceived by everybody in the system agree with one another and with the outcomes that their decisions jointly generate.

The second requirement is Muth's *rational expectations* assumption [Muth, 1961].

In an economy one person's decisions are part of another person's constraints, so consistency requires each person to hold correct beliefs about everyone else's decisions, decision processes, and beliefs.

Consistency is also what gives rational expectations its power: without some restriction on perceptions, a model in which behavior depends on arbitrary assumptions about subjective beliefs can produce almost any outcome at all.

But look at what that requirement imputes to people once a model is taken to data.

The agents inside a rational expectations model evaluate their Euler equations using *equilibrium* probability distributions.

Those are the very distributions that the econometrician studying them is still struggling to estimate.

The agents, in other words, have somehow already solved the inference problem that the economist is only part way through.

90.1.2 Sargent's formulation of bounded rationality

Sargent's **bounded rationality** program keeps individual rationality and retreats from mutual consistency in a particular way, chosen to remove the asymmetry just described between the agents in a model and the econometrician who studies them.

Sargent's version of a bounded rationality program is (Sargent [1993], ch. 1):

> I interpret a proposal to build models with 'boundedly rational' agents as a call to retreat from the second piece of rational expectations (mutual consistency of perceptions) by expelling rational agents from our model environments and replacing them with 'artificially intelligent' agents who behave like econometricians. These 'econometricians' theorize, estimate, and adapt in attempting to learn about probability distributions which, under rational expectations, they already know.

The agents are made more like the people who build the models: they gather data, form theories, estimate, and adapt.

Note

Why "peculiar"?

For Herbert Simon, who coined the term, bounded rationality meant limits on people's capacity to gather and process information, so that they *satisfice* rather than optimize.

Sargent's agents are bounded in a different way: they compute like econometricians, but they do not know the equilibrium probability distributions.

The book credits Simon with the aim of removing the asymmetry between agents and econometrician, but Simon rejected this way of doing it.

After he saw the manuscript of the book, Simon wrote Sargent a letter objecting to the formulation and recommending that Sargent not call it "bounded rationality".

Simon particularly disliked making the agents inside the model act like econometricians, a pretense that he found preposterous.

Rational expectations and bounded rationality share a birthplace: in the 1950s Muth and Simon were colleagues and coauthors at the Carnegie Institute of Technology, now Carnegie Mellon University [Holt *et al.*, 1960].

Sargent's proposal makes work harder for the model builder, not easier.

Withdrawing the assumption of a commonly understood environment means we must put something in its place, and there are many plausible somethings to choose among (Sargent [1993], ch. 1):

> This area is wilderness because the researcher faces so many choices after he decides to forgo the discipline provided by equilibrium theorizing. The commitment to equilibrium theorizing made many choices for him by requiring that people be modelled as optimal decision-makers within a commonly understood environment. When we withdraw the assumption of a commonly understood environment, we have to replace it with something, and there are so many plausible possibilities.

90.1.3 What the program is good for

Chapter 2 of Sargent [1993] lists four payoffs that might justify the effort.

1. **Equilibrium selection.** When a model has many equilibria, adaptive agents may converge to only some of them, turning learning into a device for selecting among rational expectations equilibria; where the equilibrium is unique, agents who learn their way to it lend it a plausibility it lacked as a mere assumption.
2. **Computation.** Learning algorithms double as algorithms for computing equilibria too complicated to solve by hand, and the literatures that supply the agents' "brains" also supply economists with new methods of optimization and estimation.
3. **Transition dynamics and regime changes.** Adaptive dynamics hold out the hope of a theory of the out-of-equilibrium adjustment that the Eastern European reformers had to manage blind.
4. **New sources of dynamics.** Where rational expectations makes a sharp prediction that the data reject, adaptive agents can supply what is missing; the book's leading example is the no-trade theorem of finance, under which no trade can occur in equilibrium.

The third payoff answers a criticism that the book levels at standard practice.

A common use of rational expectations models is to compute one equilibrium under a "pre-reform" policy rule and another under a "post-reform" rule, and to read the difference as the effect of the reform.

That comparison is itself inconsistent with rational expectations: each equilibrium assumes a government committed forever to its rule, so the comparison predicts what would happen if the government did what the agents were certain it never would.

The book illustrates with the abrupt ends of the German and Austrian hyperinflations, where the shortcut seems to cost little, and with two stories about the onset of the Great Depression, each of which works only by backing off rational

expectations at some point.

A theory of how people learn their way into a new regime would fill that gap.

That promise is the least fulfilled, as the series will acknowledge; selection and computation are the surer payoffs.

This lecture sets up the selection problem in its sharpest form: models with **too many equilibria**.

When a rational expectations model has a continuum of equilibria, the physical description of the economy plus the equilibrium concept fail to pin down what happens.

Something else must choose, and a plausible account of how people grope toward equilibrium is a natural candidate for that something.

We build two monetary examples that later lectures return to, both with a continuum of rational expectations equilibria:

- a quantity-theory model of money and prices in which the price level is determined only up to an arbitrary **bubble** term, and
- a two-currency version in which the *exchange rate is completely unrestricted*.

Along the way we set up the machinery — the fixed-point view of equilibrium, the relaxation algorithm, adaptive expectations, Muth's inverse optimal-prediction problem, and the stochastic-approximation recursion whose associated differential equation governs where adaptive agents end up — that later lectures use to expel the rational agents and put adaptive ones in their place.

90.1.4 The rest of the series

- *Adaptive Agents in an Overlapping Generations Monetary Economy* — adaptive households in Samuelson's overlapping generations monetary model.
  - Least squares learning selects the low-inflation equilibrium that the rational expectations dynamics reject, and laboratory subjects go the same way.
  - A government learning a Phillips curve closes the lecture.
- *Learning, Approximation, and Equilibrium Computation* — what an agent must give up when the state is continuous and a separate response for each contingency is out of the question, which brings in approximate equilibria and the resemblance between learning algorithms and equilibrium computation algorithms.
- *Exchange Rate Indeterminacy, Learning, and Experiments* — the two-currency model of this lecture with adaptive agents, where learning pins the exchange rate down only by making it depend on history, while a genetic-algorithm economy instead produces volatility that never dies.
- *Networks, Genetic Algorithms, and Classifier Systems* — a catalogue of candidate brains, from Holland and the connectionists: perceptrons, associative memories, genetic algorithms, classifier systems.
- *Money as a Medium of Exchange among Artificially Intelligent Agents* — populations of classifier systems that discover, from scratch, which good will serve as money.
- *1993 Prospects for Bounded Rationality in Macroeconomics* — the 1993 ledger of what the program achieved and what it did not, and a postscript on the three decades since.

The lectures are ordered by how much the modeller tells the agents, a gradient of *prompting*:

- Bray's least squares learners, met *below*, know everything about their environment except one coefficient.
- The Robbins–Monro agents of *Adaptive Agents in an Overlapping Generations Monetary Economy* know their utility function but not the distribution of the returns they face.
- The agents of *Learning, Approximation, and Equilibrium Computation* must also choose a functional form for what they learn.

- The Newton learners of *Exchange Rate Indeterminacy, Learning, and Experiments* know the gradient of their own payoff; its genetic-algorithm strings, like the classifier systems of *Networks, Genetic Algorithms, and Classifier Systems* and *Money as a Medium of Exchange among Artificially Intelligent Agents*, see only a payoff signal.

1993 Prospects for Bounded Rationality in Macroeconomics then enters the prompting on the debit side of the ledger.

90.1.5 Notation across the series

Each lecture follows the notation of its sources, so some symbols change jobs from lecture to lecture.

symbol	in this lecture	later in the series
λ	relaxation parameter (gain); growth rate of the bubble	portfolio share (*Exchange Rate Indeterminacy, Learning, and Experiments*); trade decision (*Money as a Medium of Exchange among Artificially Intelligent Agents*)
γ	curvature of cost; coefficient on money in the price level	learning gain (*Adaptive Agents in an Overlapping Generations Monetary Economy* – *Networks, Genetic Algorithms, and Classifier Systems*); consumption decision (*Money as a Medium of Exchange among Artificially Intelligent Agents*)
θ	Muth's moving-average coefficient; generic belief parameters	Phillips-curve slope (*Adaptive Agents in an Overlapping Generations Monetary Economy*); approximation coefficients (*Learning, Approximation, and Equilibrium Computation*)
w_1, u	endowments are $2w_1, 2w_2$	endowments are w_1, w_2 (*Adaptive Agents in an Overlapping Generations Monetary Economy* – *Exchange Rate Indeterminacy, Learning, and Experiments*); weights (*Networks, Genetic Algorithms, and Classifier Systems*)
π	—	gross inflation (*Adaptive Agents in an Overlapping Generations Monetary Economy*); choice and holding probabilities (*Networks, Genetic Algorithms, and Classifier Systems*, *Money as a Medium of Exchange among Artificially Intelligent Agents*)
s	—	saving (*Adaptive Agents in an Overlapping Generations Monetary Economy* – *Exchange Rate Indeterminacy, Learning, and Experiments*); network state (*Networks, Genetic Algorithms, and Classifier Systems*); storage cost (*Money as a Medium of Exchange among Artificially Intelligent Agents*)
N	(the number of firms is n)	cohort size (*Adaptive Agents in an Overlapping Generations Monetary Economy*, *Exchange Rate Indeterminacy, Learning, and Experiments*); population or pattern size (*Exchange Rate Indeterminacy, Learning, and Experiments*, *Networks, Genetic Algorithms, and Classifier Systems*)

From *Adaptive Agents in an Overlapping Generations Monetary Economy* on, R is reserved for gross returns and M for estimated moment or curvature matrices.

Adaptive expectations means Cagan's and Friedman's geometric distributed lag below, while *adaptive agents* means agents who revise any rule — a forecast, a regression coefficient, the strength of a rule of thumb — in light of experience.

The learning gain is written λ_t here, because γ is taken, and γ_t from *Adaptive Agents in an Overlapping Generations Monetary Economy* on.

Let's start with some imports.

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from scipy.signal import lfilter
```

90.2 Rational expectations as a fixed point

90.2.1 A static market

Take a competitive industry with a large number n of identical firms.

Each firm chooses output x to maximize

$$\Pi(x, p) = px - c(x),$$

where c is an increasing, convex cost function, and price is determined by a downward-sloping inverse demand curve

$$p = p(nX)$$

in which X is the output of the *average* firm.

Each firm is a price-taker and an X-taker: it takes p as given and sets marginal cost equal to price, $p = c'(x)$.

Write the solution as $x = g(p)$.

Substituting the demand curve gives the **best-response map**

$$x = g(p(nX)) \equiv h(X), \tag{90.1}$$

which sends a conjectured industry average X to the individual output an optimizing firm would choose against it.

That is the first component of rational expectations, and it is all that individual rationality delivers.

The second component — consistency — requires that what each firm chooses coincides with what it assumed the average firm was choosing:

$$X = h(X). \tag{90.2}$$

A static rational expectations equilibrium is a **fixed point of the best-response map**.

For a concrete example, take a quadratic cost function $c(x) = \frac{\gamma}{2}x^2$ and a linear inverse demand curve $p = a - bnX$.

Then $p = c'(x) = \gamma x$ gives $x = p/\gamma$, so

$$h(X) = \frac{a - bnX}{\gamma}, \qquad \text{with fixed point} \qquad X^* = \frac{a}{\gamma + bn}.$$

```
a, b, n, γ = 10.0, 0.3, 5, 1.0

def h(X):
    "Best response of an individual firm to an industry average X."
    return (a - b * n * X) / γ

X_star = a / (γ + b * n)
print(f"X*     = {X_star}")
print(f"h(X*)  = {h(X_star)}")
print(f"h'(X)  = {-b * n / γ}")
```

```
X*     = 4.0
h(X*)  = 4.0
h'(X)  = -1.5
```

90.2.2 Computing the fixed point

Now suppose we want to *find* X^* without solving for it in closed form.

A natural starting point is the **relaxation algorithm**: carry an estimate X_k^* of the equilibrium, compute the best response to it, and move part of the way toward that best response,

$$X_k^* = X_{k-1}^* + \lambda(h(X_{k-1}^*) - X_{k-1}^*), \tag{90.3}$$

where $\lambda \in (0, 1]$ is a **relaxation parameter**.

With $\lambda = 1$ this is simple iteration on h, the classic cobweb.

```
def relax(λ, X0=1.0, n_iter=12):
    "Iterate the relaxation algorithm, returning the whole path."
    path = np.empty(n_iter + 1)
    path[0] = X0
    for k in range(n_iter):
        path[k + 1] = path[k] + λ * (h(path[k]) - path[k])
    return path


fig, axes = plt.subplots(1, 2, figsize=(11, 4))

axes[0].plot(relax(1.0), 'o-', ms=4, lw=1, color='C0',
             label=r"$\lambda = 1.0$ (cobweb)")
axes[0].set_title(r"$\lambda = 1$: the cobweb")
axes[0].set_ylabel("$X^*_k$")

for i, λ in enumerate((0.8, 0.5, 0.3)):
    axes[1].plot(relax(λ), 'o-', ms=4, lw=1, color=f'C{i+1}',
                 label=fr"$\lambda = {λ}$")
axes[1].set_title(r"$\lambda < 1$: relaxation")
axes[1].set_ylim(-1, 9)

for ax in axes:
    ax.axhline(X_star, color='k', lw=0.8, ls='--', label="$X^*$")
    ax.set_xlabel("$k$")
    ax.legend(frameon=False, fontsize=9)
plt.tight_layout()
plt.show()
```

The naive cobweb $\lambda = 1$ **diverges**, oscillating ever further away from the equilibrium it is trying to find; note the scale on the left panel.

Since h is affine, iteration (90.3) has multiplier $1 + \lambda(h' - 1)$, so it converges if and only if

$$\left|1 + \lambda(h' - 1)\right| < 1 \qquad \iff \qquad \lambda < \frac{2}{1 + bn/\gamma}.$$

```
λ_max = 2 / (1 + b * n / γ)
print(f"converges iff λ < {λ_max}")
print(f"λ = 1.0 : multiplier = {1 + 1.0 * (-b*n/γ - 1):+.2f}   (diverges)")
print(f"λ = 0.8 : multiplier = {1 + 0.8 * (-b*n/γ - 1):+.2f}   (period-2 cycle)")
print(f"λ = 0.5 : multiplier = {1 + 0.5 * (-b*n/γ - 1):+.2f}   (converges)")
```

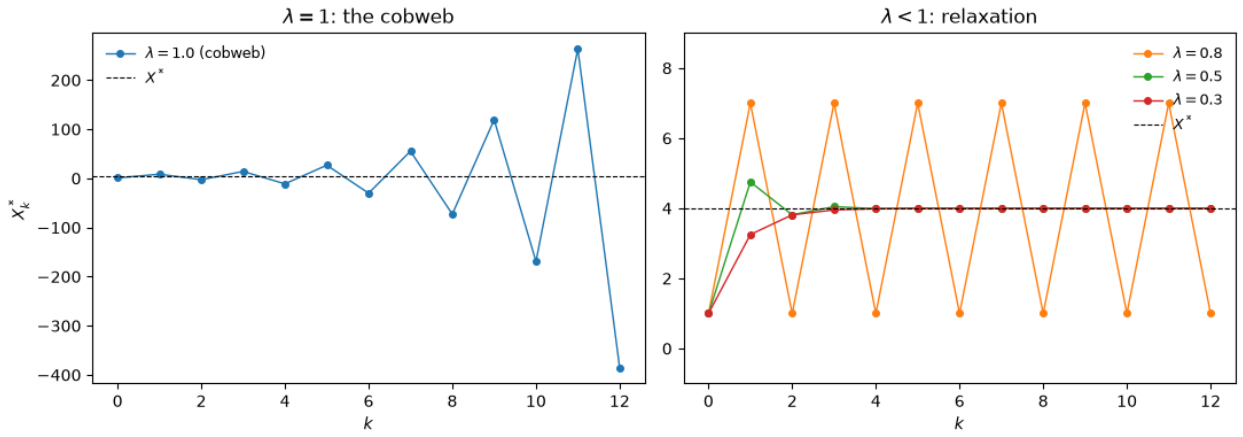


Fig. 90.1: Relaxation algorithm paths: the cobweb and three smaller relaxation parameters

```
converges iff λ < 0.8
λ = 1.0 : multiplier = -1.50   (diverges)
λ = 0.8 : multiplier = -1.00   (period-2 cycle)
λ = 0.5 : multiplier = -0.25   (converges)
```

At $\lambda = 0.8$ the multiplier is exactly -1: the scheme maps $X \mapsto 8 - X$ and so bounces forever between 1 and 7 without either converging or diverging, which is the zigzag that never dies down in the right panel.

Damping the adjustment enough makes the algorithm find the equilibrium; not damping it enough makes the algorithm chase its own tail.

Hold on to that observation.

A recurring theme of this series is that *how* you adapt determines where, and whether, you end up.

90.2.3 Adaptive expectations

Equation (90.3) was introduced as an algorithm running in iteration count k.

Reinterpret k as **calendar time** t, read X_t^* as the value people *expect*, and $X_t = h(X_{t-1}^*)$ as what actually happens, and it becomes a theory of expectation formation:

$$X_t^* = (1-\lambda)X_{t-1}^* + \lambda X_t = \lambda \sum_{j=0}^{\infty} (1-\lambda)^j X_{t-j}. \tag{90.4}$$

This is the **adaptive expectations** scheme that Cagan [Cagan, 1956] used to study hyperinflations and that Friedman [1957] used to study consumption.

Expectations are a geometrically declining distributed lag of past observations, with a single free parameter λ describing beliefs.

Cagan and Friedman took λ as a free parameter and left open the question of *why* anyone would form expectations this way.

90.2.4 Muth's inverse problem

Muth [1960] set out to eliminate λ as a free parameter by turning the question around.

Instead of asking what forecasts a given environment implies, he asked: **for what environment would exponential smoothing be the optimal forecast?**

His answer is that (90.4) is the least-squares forecast of X_{t+k} at every horizon k if and only if X_t follows

$$X_t = X_{t-1} + \epsilon_t - \theta \epsilon_{t-1}, \tag{90.5}$$

with $\{\epsilon_t\}$ a martingale difference sequence and $\theta \in [0, 1)$, and the smoothing weight tied to the moving-average coefficient by

$$\lambda = 1 - \theta \in (0, 1].$$

Note

Sargent writes both the smoothing weight in (90.4) and the moving-average coefficient in (90.5) as λ.

We use θ for the second to keep the relationship $\lambda = 1 - \theta$ visible; *Notation across the series* lists the other symbols that change jobs.

The forecasting rule inherits its one parameter from the stochastic process being forecast.

Let's confirm this numerically at horizon $k = 1$ (Exercise 90.7.4 checks a longer horizon): simulate (90.5), run exponential smoothing at a range of weights, and see which weight forecasts best.

```
def smoothing_mse(x, λ, k=1):
    "MSE of exponential smoothing with weight λ as a forecast k steps ahead."
    # f_i = λ x_i + (1 - λ) f_{i-1}, started at f_0 = x_0
    f, _ = lfilter([λ], [1, λ - 1], x, zi=[(1 - λ) * x[0]])
    return np.mean((x[k:] - f[:-k]) ** 2)


rng = np.random.default_rng(0)
T = 200_000
ε = rng.standard_normal(T + 1)            # one shock path, shared across all λ
grid = np.linspace(0.02, 0.98, 97)

fig, ax = plt.subplots(figsize=(7, 4))
for θ in (0.2, 0.5, 0.8):
    X = np.cumsum(ε[1:] - θ * ε[:-1])     # Muth's process
    mse = np.array([smoothing_mse(X, λ) for λ in grid])
    line, = ax.plot(grid, mse, lw=1.2, label=fr"$\theta = {θ}$")
    ax.axvline(1 - θ, color=line.get_color(), lw=0.8, ls='--')
    print(f"θ = {θ}:  best λ = {grid[mse.argmin()]:.2f},  1 - θ = {1 - θ:.2f},"
          f"  min MSE = {mse.min():.4f}")
ax.set_xlabel(r"smoothing weight $\lambda$")
ax.set_ylabel("forecast MSE")
ax.set_ylim(0.9, 3)
ax.legend(frameon=False)
plt.show()
```

```
θ = 0.2:  best λ = 0.80,  1 - θ = 0.80,  min MSE = 1.0024
θ = 0.5:  best λ = 0.50,  1 - θ = 0.50,  min MSE = 1.0024
θ = 0.8:  best λ = 0.20,  1 - θ = 0.20,  min MSE = 1.0024
```

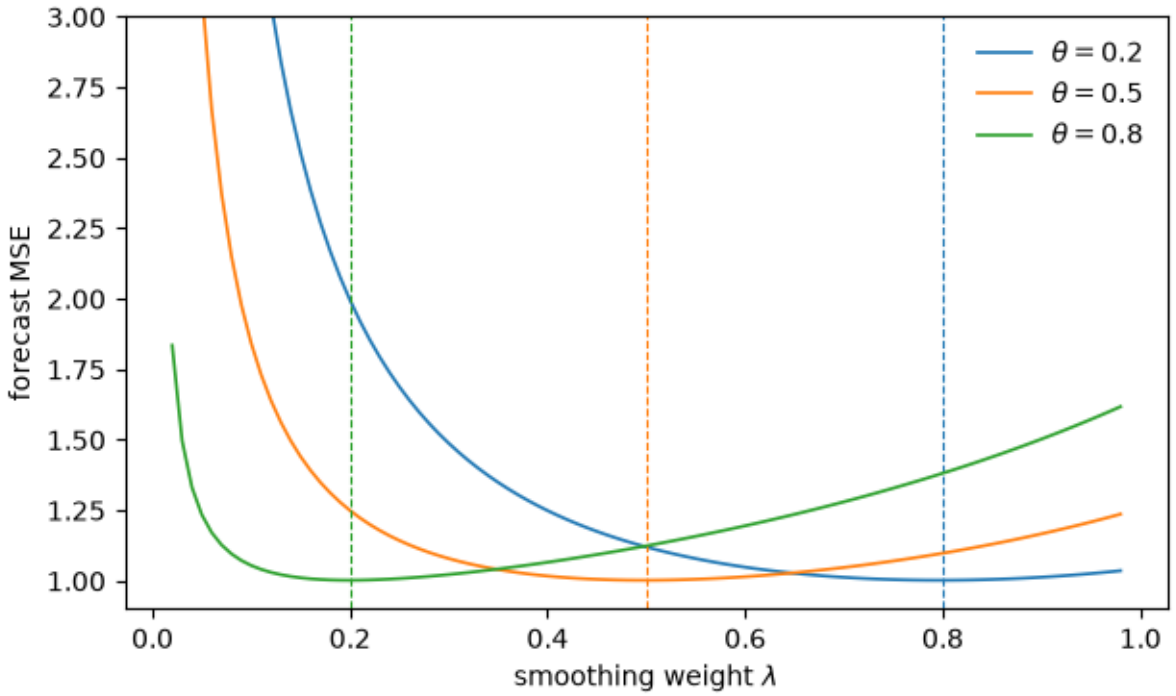


Fig. 90.2: Forecast MSE of exponential smoothing

Each curve bottoms out exactly at its dashed line $\lambda = 1 - \theta$, and the minimized mean squared error is $\mathbb{V}[\epsilon_t] = 1$ up to simulation noise — exponential smoothing at that weight *is* the conditional expectation, and nothing can beat it.

Muth's exercise was the first application of the rational expectations idea in the form that became standard: find the restrictions that link a forecasting scheme to the environment in which it is used.

It also pushed later researchers [Lucas and Prescott, 1971] to treat the **forecasting scheme itself** as the object in terms of which equilibrium is defined.

90.2.5 The dynamic analogue

That is exactly what happens when we move from static to dynamic models.

Let each agent choose a sequence rather than a single action, a contingency plan for x_t that maximizes

$$E \sum_{t=0}^{\infty} \beta^t \, \Pi(x_t, x_{t-1}, p_t),$$

taking as given a price equation $p_t = p(X_t, u_t)$ and a **perceived law of motion** for the aggregate state,

$$X_t = H(X_{t-1}, u_t),$$

where $\{u_t\}$ is IID and $\beta \in (0, 1)$ is a discount factor.

The Euler equation for this problem is

$$\Pi_1(x_t, x_{t-1}, p_t) + \beta \int \Pi_2(x_{t+1}, x_t, p_{t+1}) \, dF(p_{t+1} \mid X_t) = 0,$$

where Π_i is the derivative with respect to the ith argument and $F(\cdot \mid X_t)$ is the distribution of next period's price that the perceived law of motion H induces.

In a rational expectations equilibrium F is the *equilibrium* distribution, the very object that, as we saw at the start, the econometrician is still trying to estimate.

Solving the agent's dynamic program yields an individual decision rule $x_t = h(x_{t-1}, X_{t-1}, u_t)$.

Imposing that the representative agent is representative, $x_t = X_t$, delivers the **actual law of motion**

$$X_t = h(X_{t-1}, X_{t-1}, u_t) \equiv H^*(X_{t-1}, u_t),$$

and hence a map from perceived to actual laws of motion,

$$H^* = T(H). \tag{90.6}$$

A dynamic rational expectations equilibrium is a fixed point $H = T(H)$, the same idea as (90.2), but the fixed point now lives in a space of *functions* rather than a space of numbers.

The relaxation algorithm carries over unchanged,

$$H^*_k = H^*_{k-1} + \lambda(T(H^*_{k-1}) - H^*_{k-1}), \tag{90.7}$$

except that it now revises entire expectations-generating functions in response to the gap between what they predicted and what happened.

The econometric learners of this series are versions of (90.7) in which the revision is driven by data rather than by an exact evaluation of T, and in which the gain λ sometimes declines over time and sometimes stays constant.

The next section makes that family precise.

See also

Rational Expectations Equilibrium develops the T map in detail for a Lucas–Prescott industry model and computes its fixed point, and *Investment Under Uncertainty* studies the Lucas–Prescott model itself.

Least Squares Learning in Self-Referential Models studies what happens when agents estimate the perceived law of motion by least squares while living inside the system their estimates help determine.

90.3 One recursion, many learners

The relaxation algorithm is the simplest member of a family of recursions that runs through the whole series; to see the family, put some noise into the static market and let the firms learn.

90.3.1 From relaxation to learning: Bray's market

Suppose firms must commit to output before they see what the average firm produces.

Each firm enters period t with a belief X^*_{t-1} about average output, formed from data through $t-1$, and produces its best response $h(X^*_{t-1})$; a common supply shock u_t, IID with mean zero and variance σ_u^2, then determines what actually happens, and firms revise their belief toward it:

$$X_t = h(X^*_{t-1}) + u_t, \qquad X^*_t = X^*_{t-1} + \lambda_t(X_t - X^*_{t-1}). \tag{90.8}$$

Without the shock and with a constant gain $\lambda_t = \lambda$, this is the relaxation algorithm (90.3) run in calendar time, that is, adaptive expectations.

With the gain $\lambda_t = 1/t$, X_t^* is the sample mean of $X_1, \ldots, X_t$, the least squares estimate of a constant, and the model is, up to notation, that of Bray [1982], whose feedback coefficient is our h' (Bray writes it b, a letter that here denotes the slope of demand).

The belief error $e_t = X_t^* - X^*$ obeys

$$e_t = [1 + \lambda_t(h' - 1)]e_{t-1} + \lambda_t u_t,$$

so each period applies the relaxation multiplier at that period's gain.

A constant gain above the bound $2/(1 + bn/\gamma)$ makes beliefs explode, as the cobweb did; below the bound they do not explode, but each shock moves them by λu_t, so they wander forever in a band around X^*.

A declining gain eventually falls below any bound, and as $\lambda \to 0$ the stability condition $|1 + \lambda(h' - 1)| < 1$ collapses to $h' < 1$, Bray's condition for least squares learners to converge, which the cobweb market, with $h' = -bn/\gamma < 0$, satisfies however steep it is.

The next cell puts firms in this lecture's market, where $h' = -1.5$ and fixed gains above 0.8 explode, and compares three gains on one path of shocks.

```
def learn(gain, u, X0=1.0):
    "Beliefs X*_t of firms that update with gain(t) in the market with shocks u."
    X = np.empty(len(u) + 1)
    X[0] = X0
    for t in range(1, len(u) + 1):
        X_t = h(X[t-1]) + u[t-1]                    # what actually happens
        X[t] = X[t-1] + gain(t) * (X_t - X[t-1])    # belief revision
    return X


σ_u, T_b = 0.5, 2_000
u = σ_u * np.random.default_rng(1).standard_normal(T_b)
t_b = np.arange(1, T_b + 1)

fig, ax = plt.subplots(figsize=(8, 4))
for label, gain, lw in [(r"$\lambda_t = 0.9$", lambda t: 0.9, 1.0),
                        (r"$\lambda_t = 0.3$", lambda t: 0.3, 0.5),
                        (r"$\lambda_t = 1/t$", lambda t: 1 / t, 1.0)]:
    ax.plot(t_b, np.abs(learn(gain, u)[1:] - X_star), lw=lw, label=label)
ax.plot(t_b, σ_u / np.sqrt(t_b), 'k--', lw=0.8, label=r"$\sigma_u / \sqrt{t}$")
ax.set_xscale('log')
ax.set_yscale('log')
ax.set_ylim(1e-4, 1e3)
ax.set_xlabel("$t$")
ax.set_ylabel(r"$|X^*_t - X^*|$")
ax.legend(frameon=False, fontsize=9, ncol=2)
plt.show()
```

The fixed gain above the bound explodes, and the fixed gain below it hovers in a band whose width is set by the gain (Exercise 90.7.5 computes it).

Only the declining gain converges, with an error that shrinks like $t^{-1/2}$, the rate at which a sample mean settles down.

Least Squares Learning in Self-Referential Models analyzes Bray's model in detail, including the case $h' > 1$ in which least squares diverges, and *Optimal Misspecified Beliefs* revisits it with a constant gain.

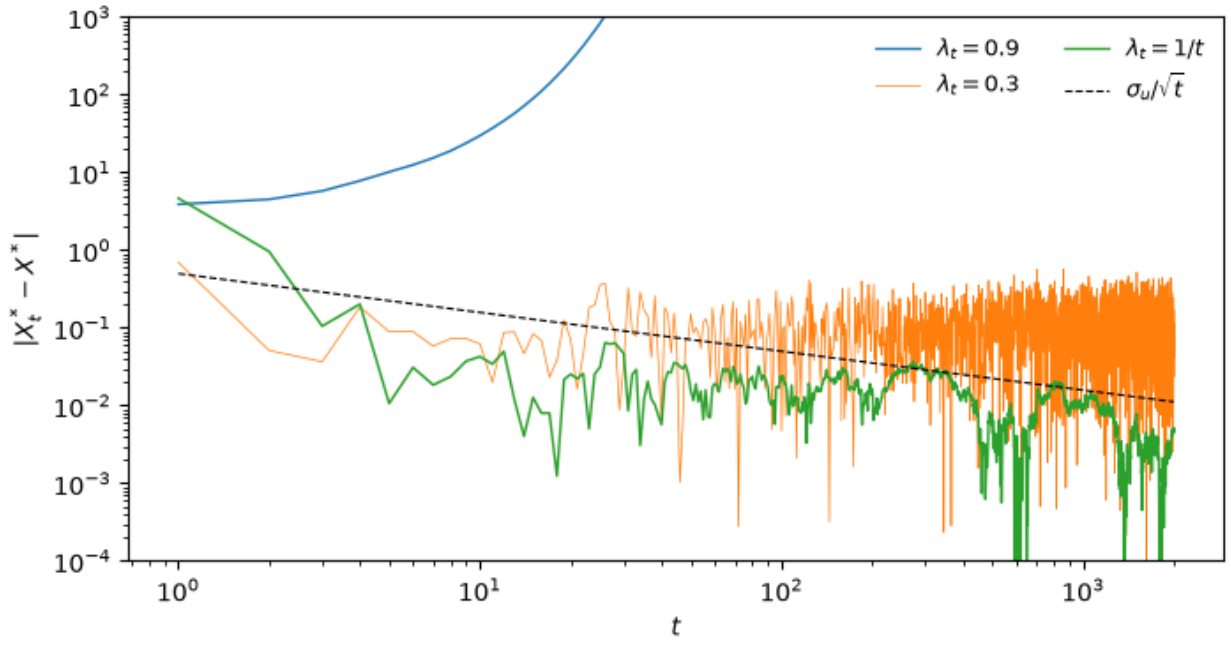


Fig. 90.3: Belief errors in Bray's market under fixed and declining gains

90.3.2 Stochastic approximation and E-stability

Most learners in this series update a vector of beliefs θ_t by a recursion of the form

$$\theta_t = \theta_{t-1} + \lambda_t \, Q(\theta_{t-1}, x_t), \tag{90.9}$$

where x_t is the latest data, λ_t is the gain, and Q is a direction of improvement, such as a forecast error or a gradient, sometimes premultiplied by the inverse of an estimated moment or curvature matrix, as in recursive least squares and Newton's method.

Equation (90.8) is the case $\theta = X^*$ and $Q = X_t - X^*_{t-1}$.

Robbins and Monro [1951] introduced such recursions to find a zero of a function observed only with noise, and Ljung [1977] showed that for small gains they follow, on average, the differential equation $d\theta/d\tau = E\,Q(\theta, x)$ in the rescaled time $\tau_t = \lambda_1 + \cdots + \lambda_t$, with the expectation taken holding θ fixed.

When agents estimate a perceived law of motion by least squares, the average update points from the perceived law toward the actual one, and the differential equation becomes

$$\frac{d\theta}{d\tau} = T(\theta) - \theta, \tag{90.10}$$

where T is the map (90.6) from perceived to actual laws of motion.

For Bray's market $T = h$, and (90.10) is $dX/d\tau = h(X) - X$, the equation that Sargent [1993] (ch. 2) already attaches to the relaxation algorithm.

A fixed point $\theta^* = T(\theta^*)$ is a rational expectations equilibrium; it is **E-stable** (expectationally stable) if it is a stable rest point of (90.10), which requires every eigenvalue of the Jacobian $DT(\theta^*)$ to have real part less than one.

What learners do near an E-stable equilibrium depends on the gain.

- With a **decreasing gain** ($\sum_t \lambda_t = \infty$ and $\sum_t \lambda_t^2 < \infty$, as for $\lambda_t = 1/t$), recursive least squares converges to an E-stable equilibrium, under regularity conditions that may include a projection facility keeping beliefs in bounds, and it converges to an E-unstable one with probability zero.

- With a **constant gain** in a stochastic environment, beliefs never settle: they fluctuate around the E-stable equilibrium in an ergodic distribution whose variance is proportional to the gain when the gain is small, with occasional large escapes [Cho *et al.*, 2002, Sargent, 1999, Williams, 2019].
- In a deterministic environment there is nothing to average away, and a small enough constant gain converges too, as the relaxation algorithm does.

E-stability is the right condition for least squares learning near an equilibrium, but it is not a theorem about every algorithm: stochastic-gradient learners can fail to converge to an E-stable equilibrium [Giannitsarou, 2005].

Population algorithms, such as the genetic algorithms and classifier systems of later lectures, lie outside the theory; drawing on Arifovic's 1991 dissertation, the book (ch. 5) reports that a genetic-algorithm population can sometimes learn its way to the equilibrium of Bray's model even when $h' > 1$.

The core of the theory predates the book, whose chapter 5 (footnote 73 and the section "An associated ordinary differential equation") uses (90.10) to study least squares learning.

The idea goes back to DeCanio [1979] and Evans [1985], Evans [1989], who used the T map to define expectational stability, and to Marcet and Sargent [1989] and Woodford [1990], who showed that it governs the convergence of least squares learning, using the method of Ljung [1977]; Evans and Honkapohja [2001] later systematized and generalized the theory.

Least Squares Learning in Self-Referential Models develops it, and *Adaptive Learning and Escape Dynamics* gives a compact primer on recursive algorithms, mean dynamics, and constant-gain escapes.

90.3.3 How the series uses it

lecture	learners and gains	what the differential equation explains
Adaptive Agents in an Overlapping Generations Monetary Economy	Robbins–Monro saving rules, declining gain	convergence to the stationary equilibrium
	least squares on inflation, weights growing with the price level (an effectively constant gain)	*the stability reversal*: $T' = \frac{1}{2} < 1$ at the low steady state $\pi = 2$, $T' = 2 > 1$ at the high one $\pi = 3$
	a government learning a Phillips curve, constant gain	recurrent escapes
Learning, Approximation, and Equilibrium Computation	online *parameterized expectations*, declining gain	slow convergence, at a rate set by the largest eigenvalue of DT
Exchange Rate Indeterminacy, Learning, and Experiments	Newton–Raphson portfolio rules, constant gain	a continuum of rest points with no drift along it, so the limit depends on where beliefs started
	a genetic algorithm	outside the theory
Networks, Genetic Algorithms, and Classifier Systems, *Money as a Medium of Exchange among Artificially Intelligent Agents*	classifier systems, $1/\tau$ averaging of rule strengths	rest points that value future payoffs as if the agents were very impatient

90.4 Money and prices

We now build the first of the two models that motivate everything that follows.

A representative money-holder chooses nominal balances m_t to carry from t to $t+1$ to maximize

$$\ln\left(2w_1 - \frac{m_t}{p_t}\right) + \ln\left(2w_2 + \frac{m_t}{p^*_{t+1}}\right), \qquad w_1 > w_2 > 0, \tag{90.11}$$

where p_t is the current price level and p^*_{t+1} is the price level expected next period.

The first term is consumption today, reduced by the real resources m_t/p_t given up to acquire currency; the second is consumption tomorrow, augmented by the goods m_t/p^*_{t+1} that the currency is expected to command.

Differentiating (90.11) with respect to m_t and rearranging gives the demand for money

$$\frac{m_t}{p_t} = w_1 - w_2\frac{p^*_{t+1}}{p_t}. \tag{90.12}$$

Real balances fall when currency is expected to lose value faster.

This is a version of the demand function Cagan [Cagan, 1956] used to study hyperinflations, and it also arises in Samuelson's [Samuelson, 1958] overlapping generations model.

Following the book, the endowments are written $2w_1$ and $2w_2$, so that the demand for real balances is $w_1 - w_2\pi$ with $\pi = p^*_{t+1}/p_t$.

From *Adaptive Agents in an Overlapping Generations Monetary Economy* on the endowments are w_1 and w_2, so the same household saves $(w_1 - w_2\pi)/2$.

To close the model we need a theory of p^*_{t+1}.

Suppose the money supply grows at a constant rate,

$$M_{t+1} = \mu M_t, \tag{90.13}$$

and that the household believes the price level is related to the money supply by

$$p_t = \gamma M_t + \lambda^t c, \tag{90.14}$$

for constants (γ, λ, c) that summarize its beliefs, with $\gamma > 0$, $\lambda > 0$ and $c \geq 0$ so that the price level stays positive.

Knowing μ, the household forecasts $p^*_{t+1} = \gamma\mu M_t + \lambda^{t+1}c$.

Substituting the forecast and (90.14) into (90.12) gives money demand as a function of the current money supply,

$$m_t = \gamma(w_1 - w_2\mu)M_t + \lambda^t(w_1 - w_2\lambda)c. \tag{90.15}$$

Notice a feature common to models in which expectations matter: the *demand* for money depends on its *supply*, because today's demand depends on tomorrow's expected price level, which is believed to depend on tomorrow's money supply.

90.4.1 Equilibrium

Setting demand equal to supply, $m_t = M_t$, turns (90.15) into a functional equation,

$$M_t = \gamma(w_1 - w_2\mu)M_t + \lambda^t(w_1 - w_2\lambda)c,$$

which must hold at every date.

Matching the two terms gives

$$\gamma = (w_1 - \mu w_2)^{-1}, \qquad \lambda = \frac{w_1}{w_2}, \qquad c \geq 0 \text{ arbitrary}, \tag{90.16}$$

so the equilibrium price level is

$$p_t = (w_1 - \mu w_2)^{-1} M_t + \left(\frac{w_1}{w_2}\right)^t c. \tag{90.17}$$

The parameters γ and λ are pinned down.

The constant c is not.

Every $c \geq 0$ is a rational expectations equilibrium, and expectations formed from (90.14) are always exactly right in each of them.

Let's verify that the residual really does vanish for any c we care to try.

```
w1, w2, μ, M0 = 2.0, 1.0, 1.5, 1.0
γ_m, λ_m = 1 / (w1 - μ * w2), w1 / w2

t = np.arange(12)
M = M0 * μ ** t

def price_path(c):
    "Equilibrium price level for bubble constant c."
    return γ_m * M + λ_m ** t * c

for c in (0.0, 0.5, 3.0, 25.0):
    p = price_path(c)
    p_star = γ_m * μ * M + λ_m ** (t + 1) * c      # forecast of next period's price
    m = w1 * p - w2 * p_star                       # money demand
    print(f"c = {c:5}:  max |demand - supply| = {np.max(np.abs(m - M)):.2e}")
```

```
c =   0.0:  max |demand - supply| = 0.00e+00
c =   0.5:  max |demand - supply| = 0.00e+00
c =   3.0:  max |demand - supply| = 0.00e+00
c =  25.0:  max |demand - supply| = 0.00e+00
```

Money demand equals money supply exactly, at every date, for every c.

90.4.2 The bubble

Since $w_1 > w_2$, the second term in (90.17) grows at the gross rate $\lambda = w_1/w_2 > 1$.

In the equilibrium with $c = 0$ the price level is proportional to the money supply: the quantity theory in its textbook form.

In every other equilibrium the price level carries a component that has nothing to do with the money supply and grows exponentially: a purely speculative **bubble**.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))

for c in (0.0, 0.5, 3.0):
    lab = f"$c = {c}$" + (" (quantity theory)" if c == 0 else "")
    axes[0].plot(t, price_path(c), 'o-', ms=3, lw=1, label=lab)
    axes[1].plot(t, M / price_path(c), 'o-', ms=3, lw=1, label=lab)
```

(continues on next page)

(continued from previous page)

```
axes[0].set_yscale('log')
axes[0].set_ylabel("$p_t$  (log scale)")
axes[1].set_ylabel("real balances  $M_t / p_t$")
axes[1].axhline(w1 - μ * w2, color='k', lw=0.8, ls='--')
axes[1].set_ylim(0, 0.6)
for ax in axes:
    ax.set_xlabel("$t$")
    ax.legend(frameon=False, fontsize=9)
plt.tight_layout()
plt.show()
```

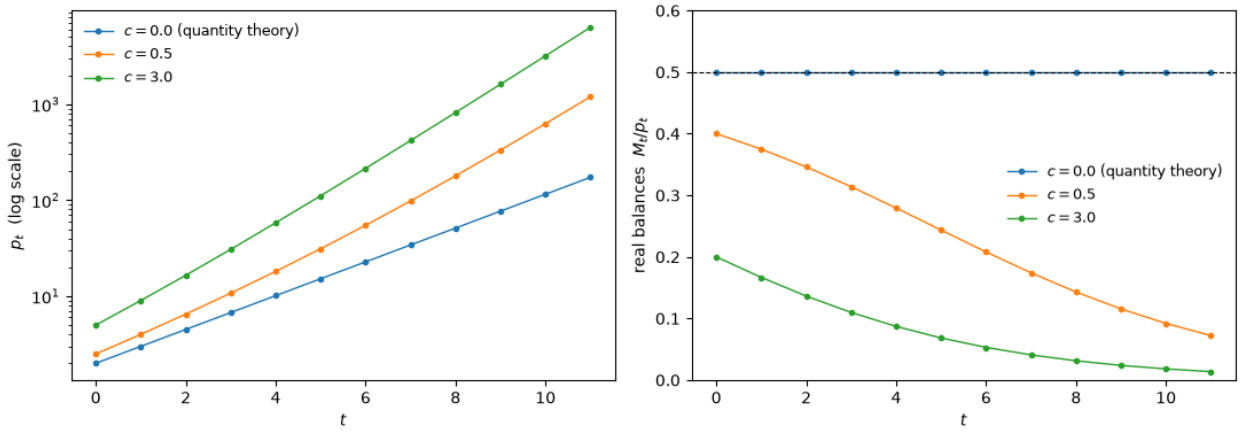


Fig. 90.4: Price levels and real balances under bubbles

The left panel shows the price level diverging further and further from the quantity-theory path as c rises.

The right panel shows what that does to the real value of the money stock: with $c = 0$ real balances are constant at $w_1 - \mu w_2$, while with $c > 0$ the bubble drives them steadily toward zero.

Along a bubble path the gross inflation rate climbs toward w_1/w_2, which from (90.12) is precisely the rate at which the demand for real balances vanishes.

The economy demonetizes itself, purely because everyone expects it to.

```
infl = pd.DataFrame(
    {f"c = {c}": price_path(c)[1:] / price_path(c)[:-1] for c in (0.0, 0.5, 3.0)},
    index=pd.Index(t[1:], name="t"),
).round(3)
infl
```

```
     c = 0.0  c = 0.5  c = 3.0
t
1        1.5    1.600    1.800
2        1.5    1.625    1.833
3        1.5    1.654    1.864
4        1.5    1.686    1.890
5        1.5    1.721    1.913
6        1.5    1.757    1.932
7        1.5    1.792    1.947
8        1.5    1.826    1.959
```

(continues on next page)

(continued from previous page)

```
9       1.5    1.857    1.969
10      1.5    1.885    1.976
11      1.5    1.908    1.982
```

There is nothing in the model — no preference, no technology, no policy — that says which c we are in.

As Sargent puts it, "Rational expectations alone is not a sufficiently restrictive principle to determine outcomes."

Adaptive Agents in an Overlapping Generations Monetary Economy returns to this household with a constant money supply ($\mu = 1$, which is a zero deficit there).

Its two stationary inflation rates are then $\pi = 1$, the bubbleless equilibrium here, and $\pi = w_1/w_2$, the limit of every bubble path.

Perfect-foresight dynamics lead away from $\pi = 1$, but that lecture shows that least squares learners converge to it: learning kills the bubble.

90.5 Two currencies

The indeterminacy gets worse when we allow more than one currency.

Versions of the result below are due to Russell Boyer (1971) and Kareken and Wallace [1981].

Keep (90.12) as the demand for currency *in total*, and suppose there are two fiat currencies in supplies M_{1t} and M_{2t} that are perfect substitutes as long as their rates of return are equal:

$$\frac{p^*_{1,t+1}}{p_{1t}} = \frac{p^*_{2,t+1}}{p_{2t}}. \tag{90.18}$$

This indifference about *which* currency to hold is what makes the exchange rate indeterminate.

Let people believe the price levels are given by

$$p_{1t} = \gamma_1 M_{1t} + \gamma_2 e M_{2t} + c\lambda^t, \qquad p_{2t} = e^{-1} p_{1t}, \tag{90.19}$$

where e is a constant exchange rate.

Requiring that the demand for currency, valued in units of currency 1, equal the total supply $M_{1t} + eM_{2t}$ gives

$$\gamma_1 = (w_1 - \mu_1 w_2)^{-1}, \quad \gamma_2 = (w_1 - \mu_2 w_2)^{-1}, \quad \lambda = \frac{w_1}{w_2}, \quad c \geq 0, \quad e \in [0, \infty). \tag{90.20}$$

These equations are remarkable for what they leave out.

The exchange rate e is *entirely unrestricted* — if the equations have a solution for one e, they have a solution for every other — and the formulas for γ_1 and γ_2 do not involve e at all.

Take the simplest case: two currencies in fixed supply, $\mu_1 = \mu_2 = 1$, and set $c = 0$.

Then $\gamma_1 = \gamma_2 = (w_1 - w_2)^{-1}$ and the price levels are constant.

```
w1, w2 = 2.0, 1.0
H1, H2 = 100.0, 120.0                  # fixed supplies of the two currencies
γ_e = 1 / (w1 - w2)

def two_currency(e):
    "Price levels and the real allocation at exchange rate e."
    p1 = γ_e * (H1 + e * H2)
```

(continues on next page)

(continued from previous page)

```
    p2 = p1 / e
    supply = H1 + e * H2              # total currency, in units of currency 1
    demand = (w1 - w2) * p1           # money demand, with p* = p since prices are␣
 ↪constant
    return p1, p2, supply / p1, demand - supply

pd.DataFrame(
    [two_currency(e) for e in (0.25, 0.5, 1.0, 2.0, 4.0)],
    index=pd.Index([0.25, 0.5, 1.0, 2.0, 4.0], name="e"),
    columns=["$p_1$", "$p_2$", "real balances", "excess demand"],
).round(4)
```

e	p_1	p_2	real balances	excess demand
0.25	130.0	520.0	1.0	0.0
0.50	160.0	320.0	1.0	0.0
1.00	220.0	220.0	1.0	0.0
2.00	340.0	170.0	1.0	0.0
4.00	580.0	145.0	1.0	0.0

Every row is an equilibrium.

The nominal price levels move around a great deal as e varies.

Since e is the value of a unit of currency 2 in units of currency 1, a larger e means currency 2 is worth more, which raises p_1 and lowers p_2.

But the last two columns tell the real story: total real balances are $w_1 - w_2 = 1$ regardless of e, and markets clear exactly.

The real allocation is identical in every one of these equilibria.

The model determines what people consume and how much purchasing power the currency stock commands; it says nothing whatever about the rate at which the two monies exchange.

This is a sharp version of a problem that has haunted international monetary theory.

The indeterminacy is not a matter of a few isolated equilibria: every $e \geq 0$ is one.

That e leaves the real allocation untouched does depend on equal money growth rates, as Exercise 90.7.3 shows.

90.6 Where this leaves us

We now have two models in which rational expectations is silent about something we would very much like to predict.

There are three ways to respond.

The first is to add restrictions to the environment until the equilibrium becomes unique.

The second is to declare the indeterminacy a genuine feature of the world and, in a tradition that runs back to Ricardo and Wicksell, to criticize the arrangements that cause it; the book calls a regime that leaves the exchange rate history-dependent "an ill-formed mechanism".

The third — the one this series pursues — is to ask what happens when we **replace the rational agents with adaptive ones** and watch where the system goes.

This is a substantive change, not a technicality.

An adaptive agent is not endowed with the equilibrium; it has beliefs, a rule for revising them, and initial conditions.

Those extra objects are exactly what the rational expectations equilibrium conditions failed to pin down, so a system of adaptive agents can select an outcome where rational expectations could not.

The rest of this series takes that idea seriously, and the results are mixed in an instructive way.

In the overlapping generations monetary economy of *Adaptive Agents in an Overlapping Generations Monetary Economy*, least squares learning selects the *opposite* equilibrium from the one the rational expectations dynamics converge to (with no deficit, the bubbleless one), and human experimental subjects side with the adaptive model.

In the two-currency model of *Exchange Rate Indeterminacy, Learning, and Experiments*, adaptive agents do pin the exchange rate down, but only by making it depend on where their beliefs started and on how their algorithm is tuned.

The rest points of the learning algorithm reproduce the indeterminacy exactly; what fixes the outcome is what the book calls *the dead hand of history*, and it grants that this "may seem a weak reed on which to base exchange rate determination."

Learning does not select an exchange rate so much as remember one.

In a Kiyotaki–Wright search economy, the classifier-system agents of *Money as a Medium of Exchange among Artificially Intelligent Agents* learn to use a medium of exchange.

But in *one of its economies*, where the speculative equilibrium is the unique stationary equilibrium for patient agents, they settle on the fundamental pattern of trade instead.

That is not learning selecting among equilibria; it is convergence to an outcome that is not an equilibrium of the economy the agents live in.

Marimon, McGrattan and Sargent suspected early myopia, and *that lecture shows why*: the way the classifier systems assign credit makes the agents effectively very impatient.

In the first two cases the algorithm supplies what the equilibrium concept did not; in the third it overrides the equilibrium concept.

Whether such outcomes are discoveries about economies or artifacts of the algorithm is the question the series keeps returning to, and *1993 Prospects for Bounded Rationality in Macroeconomics* renders a verdict.

90.7 Exercises

Exercise 90.7.1

The relaxation algorithm (90.3) converges for the static market if and only if $\lambda < 2/(1 + bn/\gamma)$.

The ratio bn/γ measures the slope of demand relative to the curvature of costs, so a market with steep demand and near-linear costs is one in which the naive cobweb ($\lambda = 1$) is badly behaved.

1. Verify the stability boundary numerically: for a grid of values of bn/γ, find the largest λ (on a fine grid) for which the algorithm converges, and compare with the analytical prediction.
2. Now let the relaxation parameter decline with the iteration count, $\lambda_k = 1/k$. Show that the algorithm converges to X^* for every slope, including slopes at which the cobweb and most fixed values of λ fail. Explain why, using the differential equation $dX/d\tau = h(X) - X$ of *One recursion, many learners*.

Solution

For part 1, we test convergence by iterating 400 times and asking whether the last step is smaller than 10^{-8}.

```
def converges(slope, λ, n_iter=400, tol=1e-8):
    """Does the relaxation algorithm converge when h'(X) = -slope?"""
```

```
    a_, X = 10.0, 1.0
    for _ in range(n_iter):
        X_new = X + λ * ((a_ - slope * X) - X)
        if not np.isfinite(X_new) or abs(X_new) > 1e12:
            return False
        X, X_prev = X_new, X
    return abs(X - X_prev) < tol


λ_grid = np.linspace(0.01, 1.5, 300)
rows = []
for slope in (0.5, 1.0, 1.5, 2.0, 4.0):
    ok = [λ for λ in λ_grid if converges(slope, λ)]
    rows.append((slope, max(ok) if ok else np.nan, 2 / (1 + slope)))

pd.DataFrame(rows, columns=["$bn/\\gamma$", "largest $\\lambda$ found",
                            "$2/(1 + bn/\\gamma)$"]).round(3)
```

	bn/γ	largest λ found	$2/(1 + bn/\gamma)$
0	0.5	1.296	1.333
1	1.0	0.972	1.000
2	1.5	0.777	0.800
3	2.0	0.648	0.667
4	4.0	0.389	0.400

The numerical boundary sits consistently a little *below* the analytical one, and that gap is not grid spacing; it is a real feature of the test.

As λ approaches the boundary the multiplier $m = 1 - \lambda(1 + bn/\gamma)$ approaches -1, so convergence becomes arbitrarily slow, and a test with a fixed iteration count and tolerance declares failure just before the true boundary.

The size of the gap is predictable.

Starting from $X_0 = 1$, the last step of the test has size $|X_{400} - X_{399}| = \lambda\,|9 - bn/\gamma|\,|m|^{399}$, so the test passes only while that is below 10^{-8}.

```
fine = np.linspace(0.01, 1.5, 1_000_001)
predicted = []
for slope in (0.5, 1.0, 1.5, 2.0, 4.0):
    m = np.abs(1 - fine * (1 + slope))
    with np.errstate(over='ignore'):
        passes = (m < 1) & (fine * abs(9 - slope) * m ** 399 < 1e-8)
    predicted.append(fine[passes].max())

pd.DataFrame({"$bn/\\gamma$": [0.5, 1.0, 1.5, 2.0, 4.0],
              "found": [r[1] for r in rows],
              "predicted by the test": predicted,
              "true boundary": [r[2] for r in rows]}).round(3)
```

	bn/γ	found	predicted by the test	true boundary
0	0.5	1.296	1.299	1.333
1	1.0	0.972	0.975	1.000
2	1.5	0.777	0.780	0.800
3	2.0	0.648	0.650	0.667
4	4.0	0.389	0.391	0.400

The predicted cutoffs lie just above the grid values found, within one grid step.

Note the first row: when $bn/\gamma < 1$ the boundary exceeds one, so the naive cobweb converges on its own and no damping is needed.

Damping is what buys convergence in the steep markets, and with a fixed λ the steeper the market the more damping is required.

For part 2, we run the declining-gain version from the same starting point and report the distance from X^* after 400 and 4000 iterations.

```
def relax_declining(slope, n_iter):
    "Distance from X* after n_iter steps of relaxation with λ_k = 1/k."
    a_, X = 10.0, 1.0
    for k in range(1, n_iter + 1):
        X = X + (1 / k) * ((a_ - slope * X) - X)
    return abs(X - a_ / (1 + slope))


slopes = (0.25, 0.5, 1.5, 4.5, 19.5)
pd.DataFrame({"$bn/\\gamma$": slopes,
              "fixed-λ bound": [2 / (1 + s_) for s_ in slopes],
              "error, k = 400": [relax_declining(s_, 400) for s_ in slopes],
              "error, k = 4000": [relax_declining(s_, 4000) for s_ in slopes]})
```

	bn/γ	fixed-λ bound	error, k = 400	error, k = 4000
0	0.25	1.600000	7.986362e-04	4.489482e-05
1	0.50	1.333333	2.000047e-04	6.319365e-06
2	1.50	0.800000	3.985628e-07	1.255046e-09
3	4.50	0.363636	6.861178e-14	1.088019e-14
4	19.50	0.097561	1.110223e-16	1.110223e-16

Every slope converges, including $bn/\gamma = 19.5$, where any fixed λ above 0.1 fails.

With $\lambda_k = 1/k$ the error is multiplied at step k by $1 - (1 + bn/\gamma)/k$; once $k > (1 + bn/\gamma)/2$ every step contracts, and the product of the multipliers behaves like $k^{-(1+bn/\gamma)}$, which goes to zero whatever the slope.

The differential equation says the same thing: in the rescaled time $\tau_k = 1 + \frac{1}{2} + \cdots + \frac{1}{k} \approx \log k$ the iteration tracks $dX/d\tau = h(X) - X$, whose solution $X(\tau) - X^* = e^{(h'-1)\tau}(X(0) - X^*) \approx k^{h'-1}(X(0) - X^*)$ is stable whenever $h' < 1$, which every downward-sloping h satisfies.

With a declining gain, then, a steeper market converges *faster*, the reverse of the fixed-gain case; the cost is that in flat markets the error shrinks only like a power of k (compare the two error columns in the first row), where a well-chosen fixed λ would shrink it geometrically.

Exercise 90.7.2

The equilibrium (90.16) was derived without checking that it makes economic sense.

Show that a monetary equilibrium requires $\mu < w_1/w_2$, by finding what goes wrong with the demand for real balances when money grows faster than that.

Solution

Along the $c = 0$ equilibrium, real balances are constant at $M_t/p_t = \gamma^{-1} = w_1 - \mu w_2$.

This is positive only when $\mu < w_1/w_2$.

If money grows faster, (90.16) still "solves" the functional equation, but it asks the household to hold a negative quantity of currency.

```
w1, w2 = 2.0, 1.0
μ_grid = np.array([0.5, 1.0, 1.5, 1.9, 2.0, 2.5])
```

```
with np.errstate(divide='ignore'):        # γ is infinite exactly at μ = w1/w2
    table = pd.DataFrame({
        "$\\mu$": μ_grid,
        "$\\gamma = (w_1 - \\mu w_2)^{-1}$": 1 / (w1 - μ_grid * w2),
        "real balances $w_1 - \\mu w_2$": w1 - μ_grid * w2,
        "monetary equilibrium?": np.where(μ_grid < w1 / w2, "yes", "no"),
    }).round(3)
table
```

```
   $\mu$  $\gamma = (w_1 - \mu w_2)^{-1}$  real balances $w_1 - \mu w_2$  \
0    0.5                           0.667                           1.5
1    1.0                           1.000                           1.0
2    1.5                           2.000                           0.5
3    1.9                          10.000                           0.1
4    2.0                             inf                           0.0
5    2.5                          -2.000                          -0.5

  monetary equilibrium?
0                   yes
1                   yes
2                   yes
3                   yes
4                    no
5                    no
```

At $\mu = w_1/w_2 = 2$ the demand for real balances hits zero and γ blows up; beyond it, both are negative.

The intuition runs through (90.12): the household holds currency only if the expected loss of purchasing power, p^*_{t+1}/p_t, is smaller than w_1/w_2.

Money growing at rate μ produces inflation at rate μ, so $\mu \geq w_1/w_2$ drives the demand for money to zero, the same boundary that the *bubble* equilibria approach asymptotically from below.

Exercise 90.7.3

In the two-currency example we set $\mu_1 = \mu_2 = 1$ and found that the real allocation was the same at every exchange rate.

That is special.

Repeat the calculation with $\mu_1 \neq \mu_2$ — say $\mu_1 = 1.0$ and $\mu_2 = 1.3$, with $w_1 = 2$, $w_2 = 1$, $M_{1,0} = M_{2,0} = 100$, and $c = 0$ — and compute total real balances at several exchange rates over the first several periods.

Does the choice of e still leave the real allocation untouched?

Solution

With $c = 0$, total real balances in units of currency 1 are $(M_{1t} + eM_{2t})/(\gamma_1 M_{1t} + \gamma_2 eM_{2t})$.

```
w1, w2 = 2.0, 1.0
μ1, μ2 = 1.0, 1.3
γ1, γ2 = 1 / (w1 - μ1 * w2), 1 / (w1 - μ2 * w2)
t = np.arange(10)
M1, M2 = 100.0 * μ1 ** t, 100.0 * μ2 ** t

def real_balances(e):
```

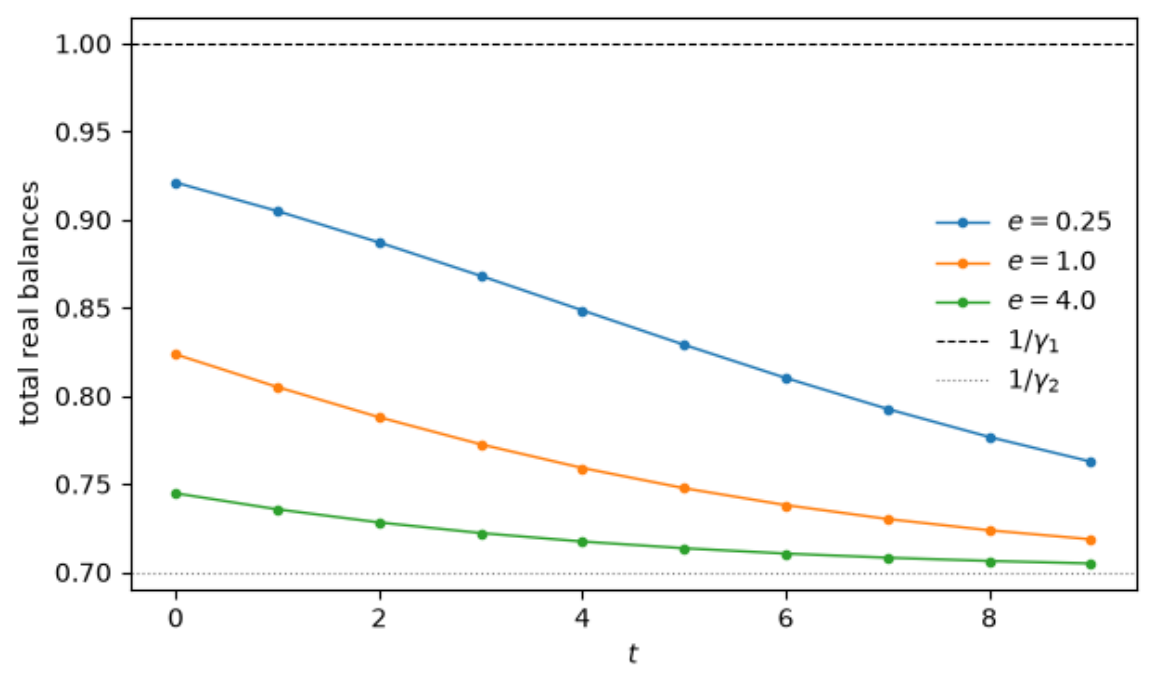
```
    p1 = y1 * M1 + y2 * e * M2
    return (M1 + e * M2) / p1

pd.DataFrame({f"e = {e}": real_balances(e) for e in (0.25, 1.0, 4.0)},
             index=pd.Index(t, name="t")).round(4)
```

t	e = 0.25	e = 1.0	e = 4.0
0	0.9211	0.8235	0.7447
1	0.9049	0.8050	0.7356
2	0.8871	0.7879	0.7282
3	0.8681	0.7725	0.7221
4	0.8485	0.7591	0.7173
5	0.8290	0.7476	0.7135
6	0.8101	0.7380	0.7105
7	0.7926	0.7301	0.7081
8	0.7767	0.7237	0.7063
9	0.7627	0.7186	0.7049

The same paths in a figure:

```
fig, ax = plt.subplots(figsize=(7, 4))
for e in (0.25, 1.0, 4.0):
    ax.plot(t, real_balances(e), 'o-', ms=3, lw=1, label=f"$e = {e}$")
ax.axhline(1 / y1, color='k', lw=0.8, ls='--', label=r"$1/\gamma_1$")
ax.axhline(1 / y2, color='gray', lw=0.8, ls=':', label=r"$1/\gamma_2$")
ax.set_xlabel("$t$")
ax.set_ylabel("total real balances")
ax.legend(frameon=False)
plt.show()
```



No.

With unequal money growth rates the exchange rate affects the real allocation, and the allocation is no longer even constant over time.

The reason is that $\gamma_1 \neq \gamma_2$: the two currencies are valued differently because they are expected to be diluted at different rates, so the *composition* of the currency stock matters, and e is what fixes that composition.

Since currency 2 grows faster, it comes to dominate the stock whatever e we choose, and real balances converge to $1/\gamma_2$ from wherever e starts them.

So the irrelevance of e for the real allocation is a knife-edge: it needs $\mu_1 = \mu_2$ (with any common growth rate μ, real balances are $w_1 - \mu w_2$ at every e).

The indeterminacy of e is not a knife-edge: every e in the table is still an equilibrium, but now the choice among them has real consequences.

Exercise 90.7.4

Muth's result says that exponential smoothing with $\lambda = 1 - \theta$ is the optimal forecast of X_{t+k} at *every* horizon k, but Fig. 90.2 checked only $k = 1$.

1. Show that for (90.5) the conditional expectation $E_t X_{t+k}$ is the same for every $k \geq 1$, that it equals the exponentially smoothed value with $\lambda = 1 - \theta$, and that the k-step forecast error has variance $1 + (k-1)(1-\theta)^2$ when $\mathbb{V}[\epsilon_t] = 1$.
2. Verify the result numerically at $k = 4$, reusing `smoothing_mse`.

Solution

Write $X_{t+k} = X_t + \sum_{j=1}^{k} (\epsilon_{t+j} - \theta \epsilon_{t+j-1})$.

Every shock dated after t has conditional mean zero, so $E_t X_{t+k} = X_t - \theta \epsilon_t \equiv f_t$ for every $k \geq 1$.

Since $\epsilon_t = X_t - f_{t-1}$, we have $f_t = X_t - \theta(X_t - f_{t-1}) = (1-\theta) X_t + \theta f_{t-1}$, which is (90.4) with $\lambda = 1 - \theta$.

The forecast error is

$$X_{t+k} - f_t = \epsilon_{t+k} + (1-\theta)(\epsilon_{t+1} + \cdots + \epsilon_{t+k-1}),$$

whose variance is $1 + (k-1)(1-\theta)^2$.

Here is the numerical check, on the same shocks as the figure.

```
k = 4
rows = []
for θ in (0.2, 0.5, 0.8):
    X = np.cumsum(ε[1:] - θ * ε[:-1])
    mse = np.array([smoothing_mse(X, λ, k) for λ in grid])
    rows.append((θ, grid[mse.argmin()], 1 - θ, mse.min(), 1 + (k - 1) * (1 - θ) **␣
↪2))

pd.DataFrame(rows, columns=["$\\theta$", "best $\\lambda$", "$1 - \\theta$",
                            "min MSE", "$1 + (k-1)(1-\\theta)^2$"]).round(3)
```

	θ	best λ	$1 - \theta$	min MSE	$1 + (k-1)(1-\theta)^2$
0	0.2	0.79	0.8	2.916	2.92
1	0.5	0.50	0.5	1.749	1.75
2	0.8	0.20	0.2	1.121	1.12

At horizon 4 the best weight is again $1 - \theta$, up to the grid and sampling noise, and the minimized mean squared error matches the formula.

The error variance grows with the horizon, but the forecast does not change: the same smoothed value is optimal at every horizon.

Exercise 90.7.5

In Bray's market (90.8) with a constant gain λ below the stability bound, beliefs never settle down.

1. Show that the belief error $e_t = X_t^* - X^*$ is a stationary first-order autoregression with variance

$$V(\lambda) = \frac{\lambda \sigma_u^2}{(1-h')(2-\lambda(1-h'))}.$$

2. Using `learn` from *From relaxation to learning: Bray's market*, simulate 100,000 periods for $\lambda \in \{0.025, 0.05, 0.1, 0.2, 0.4\}$ and compare the sample variance of e_t with $V(\lambda)$. How does the spread of beliefs scale with the gain?

Solution

From (90.8), $e_t = \rho e_{t-1} + \lambda u_t$ with $\rho = 1 - \lambda(1-h')$.

Below the bound $|\rho| < 1$, so e_t is a stationary autoregression with variance $\lambda^2\sigma_u^2/(1-\rho^2)$, and $1-\rho^2 = \lambda(1-h')(2-\lambda(1-h'))$ gives the formula.

The simulation starts beliefs at X^* and drops the first 1,000 periods.

```
u_long = σ_u * np.random.default_rng(2).standard_normal(100_000)
h_prime = -b * n / γ

rows = []
for λ in (0.025, 0.05, 0.1, 0.2, 0.4):
    e = learn(lambda t: λ, u_long, X0=X_star)[1_000:] - X_star
    V = λ * σ_u**2 / ((1 - h_prime) * (2 - λ * (1 - h_prime)))
    rows.append((λ, e.var(), V, e.var() / λ))

pd.DataFrame(rows, columns=["gain $\\lambda$", "sample variance", "$V(\\lambda)$",
                            "variance / $\\lambda$"]).round(5)
```

	gain λ	sample variance	$V(\lambda)$	variance / λ
0	0.025	0.00129	0.00129	0.05147
1	0.050	0.00266	0.00267	0.05318
2	0.100	0.00570	0.00571	0.05704
3	0.200	0.01332	0.01333	0.06661
4	0.400	0.03991	0.04000	0.09978

The sample variances match the formula.

For small gains $V(\lambda) \approx \lambda\sigma_u^2/(2(1-h'))$, so the variance of beliefs is proportional to the gain and their standard deviation to its square root; the ratio in the last column creeps up at larger gains because of the second factor in the denominator.

A constant-gain learner therefore converges only in distribution, to a cloud around X^* that collapses to a point only as $\lambda \to 0$.

Adaptive Learning and Escape Dynamics shows what such a cloud can do in a model with richer feedback: escape.

CHAPTER

NINETYONE

ADAPTIVE AGENTS IN AN OVERLAPPING GENERATIONS MONETARY ECONOMY

Contents

- *Adaptive Agents in an Overlapping Generations Monetary Economy*
  - *Overview*
  - *The environment*
  - *Part 1: a stochastic deficit*
  - *Part 2: a constant deficit and two steady states*
  - *A government learning the Phillips curve*
  - *Concluding remarks*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

91.1 Overview

A Peculiar Definition of Bounded Rationality studied monetary models with many rational expectations equilibria, among them a Cagan money-demand model whose equilibria include a continuum of bubbles.

This lecture works with a close relative: the overlapping generations model of fiat money of Samuelson [1958], in the version that Bryant, Wallace and others used to study inflationary finance, in which a government prints currency to pay for its expenditures.

We replace agents who forecast with the equilibrium law of motion for the price level by "adaptive" ones who do not.

We do this twice, for two different purposes.

Part 1 puts a *stochastic* government deficit into the model and asks whether successive generations, groping their way by trial and error, can somehow as a society converge to a rational expectations equilibrium.

They can.

Each generation observes what happened to its predecessors and adjusts their saving decision in a utility-improving direction, using a **Robbins–Monro** recursion.

Given enough time the economy converges to a stationary equilibrium that we compute independently, and of the two such equilibria it finds only the low-inflation one.

This exercise also shows how sensitive learning is to the complexity of what must be learned: a deficit state that occurs rarely is learned about slowly, because the observations arrive slowly.

Part 2 removes the randomness and turns to a sharper question.

With a *constant* deficit the model has **two** stationary equilibria, one with low inflation and one with high inflation.

The low-inflation equilibrium Pareto-dominates the other.

Under the rational expectations dynamics, the low-inflation equilibrium is **unstable** and the high-inflation one is stable: theory pushes the economy toward the bad outcome.

Under least squares learning, the stability is **reversed**.

The reversal is not a numerical accident: it follows from *E-stability*, the stability condition for learning that *A Peculiar Definition of Bounded Rationality* sets out in *One recursion, many learners*.

And when Marimon and Sunder [1993] ran the economy as a laboratory experiment with paid human subjects, the subjects behaved like the adaptive model, not like the rational expectations model.

That is the strongest evidence in Sargent [1993] that adaptive dynamics are doing real work as a selection device, tempered by a counterexample, due to Benjamin Bental, warning against concluding that adaptation reliably selects the *good* equilibrium.

We then close with a **third application**, in which the adaptive agent is a government learning a Phillips curve.

Least squares learning now converges to the model's unique self-confirming equilibrium, which has an inflation bias, while a government that keeps a constant gain makes recurrent **escapes** toward the zero-inflation Ramsey outcome, which is not an equilibrium at all.

That phenomenon seeded Sargent's later work on *The Conquest of American Inflation*, and the Phillips curve lectures that begin with *The Rise and Fall of U.S. Inflation* study it in depth.

Let's start with some imports.

```
import numpy as np
import pandas as pd
import quantecon as qe
import matplotlib.pyplot as plt
from scipy.optimize import fsolve
```

91.2 The environment

The economy consists of overlapping generations of two-period-lived agents.

At each date $t \geq 1$, N identical agents are born, endowed with w_1 units of a single consumption good when young and w_2 units when old.

A young agent's preferences over lifetime consumption (c_1, c_2) are ordered by the expected value of $u(c_1) + u(c_2)$.

Throughout we use logarithmic utility, $u(c) = \ln c$.

A government prints currency to finance expenditures G_t, subject to the budget constraint

$$G_t = \frac{H_t - H_{t-1}}{p_t}, \tag{91.1}$$

where H_t is the stock of currency carried by the young at t into $t+1$, and p_t is the price level.

Currency is the only store of value, so a young agent's saving s_t is entirely in the form of real balances, and market clearing requires

$$\frac{H_t}{p_t} = s_t N. \tag{91.2}$$

Write $R_t = p_t / p_{t+1}$ for the gross rate of return on currency between t and $t+1$.

Combining (91.1) and (91.2) gives a relation we will use constantly: the return on currency is determined entirely by saving behavior and the deficit,

$$R_t = \frac{N s_{t+1} - G_{t+1}}{N s_t}. \tag{91.3}$$

Currency holds its value only if the next generation is willing to absorb the outstanding stock plus the new issue.

Note

(91.3) is worth pausing on, because it is what makes this a *self-referential* system, the property studied at length in *Least Squares Learning in Self-Referential Models*.

The return that today's young earn on their saving depends on how much tomorrow's young choose to save, which in turn depends on what tomorrow's young expect to earn.

Nobody's beliefs are about an exogenous process.

91.3 Part 1: a stochastic deficit

Now let government expenditures follow a Markov chain with transition probabilities

$$P(i,j) = \text{Prob}\{G_{t+1} = \bar{G}_j \mid G_t = \bar{G}_i\},$$

over a finite list of levels $G = [\bar{G}_1, \dots, \bar{G}_n]$.

(Sargent writes $\pi(i,j)$; we reserve π for inflation.)

91.3.1 Stationary rational expectations equilibrium

In a stationary equilibrium, saving depends on the current deficit state, so it is described by an $n \times 1$ vector $s = [s_1, \dots, s_n]$, and the return on currency by an $n \times n$ matrix $R(i,j)$.

A young agent who observes $G_t = \bar{G}_i$ chooses s_i to maximize

$$V(s_i) = u(w_1 - s_i) + \sum_j u(w_2 + s_i R(i,j)) P(i,j),$$

with first-order condition

$$u'(w_1 - s_i) = \sum_{j=1}^{n} u'(w_2 + s_i R(i,j)) R(i,j) P(i,j). \tag{91.4}$$

Substituting (91.3) — which in the stationary equilibrium reads $R(i,j) = (s_j N - \bar{G}_j)/(s_i N)$ — turns (91.4) into a system of n nonlinear equations in the saving rates alone:

$$u'(w_1 - s_i) = \sum_{j=1}^{n} u'\left(w_2 + \frac{s_j N - \bar{G}_j}{N}\right) \cdot \frac{s_j N - \bar{G}_j}{s_i N} \cdot P(i,j). \tag{91.5}$$

A stationary equilibrium exists if and only if (91.5) has a solution with $s_i \in (0, w_1)$ for every i and with every return positive, that is, $s_j N > \bar{G}_j$ for every j.

The code below solves (91.5) from an initial guess and checks both conditions.

```
w1, w2, N = 20.0, 10.0, 1.0

def u_prime(c):
    return 1 / c                          # log utility

def equilibrium_residual(s, G, P):
    "Left minus right side of the equilibrium condition, one entry per deficit state."
    n = len(s)
    out = np.empty(n)
    for i in range(n):
        rhs = sum(u_prime(w2 + (s[j]*N - G[j])/N)
                  * (s[j]*N - G[j])/(s[i]*N)
                  * P[i, j] for j in range(n))
        out[i] = u_prime(w1 - s[i]) - rhs
    return out

def stationary_equilibrium(G, P, guess=4.0):
    "Solve the equilibrium condition from `guess`; check that the solution is an␣
 ↪equilibrium."
    s, _, ier, msg = fsolve(equilibrium_residual, np.full(len(G), guess),
                            args=(G, P), full_output=True)
    R = np.array([[(s[j]*N - G[j])/(s[i]*N) for j in range(len(s))]
                  for i in range(len(s))])
    assert ier == 1, msg
    assert np.all((s > 0) & (s < w1)) and np.all(R > 0), "not an equilibrium"
    return s, R
```

Following Sargent [1993], we study two economies that differ only in the process for government expenditures, with $w_1 = 20$, $w_2 = 10$ and $N = 1$, so that $\bar{G}$ is measured per young person.

In the first, the deficit is always zero and the chain is a fair coin.

```
G1, P1 = np.array([0.0, 0.0]), np.full((2, 2), 0.5)
s1, R1 = stationary_equilibrium(G1, P1)
print(f"saving rates      {np.round(s1, 4)}")
print(f"returns\n{np.round(R1, 4)}")
```

```
saving rates      [5. 5.]
returns
[[1. 1.]
 [1. 1.]]
```

With no deficit to finance, currency simply holds its value: agents save 5 in either state and $R \equiv 1$.

In the second economy the deficit is 0.8 in state 1 and zero in state 2, and the chain is persistent.

```
G2 = np.array([0.8, 0.0])
P2 = np.array([[0.75, 0.25],
               [0.50, 0.50]])
s2, R2 = stationary_equilibrium(G2, P2)
print(f"saving rates      {np.round(s2, 4)}      (Sargent reports 4.211, 4.364)")
print(f"returns\n{np.round(R2, 4)}")
print("Sargent reports  [[0.81, 1.0362], [0.7817, 1.00]]")
```

```
saving rates       [4.2111 4.3637]       (Sargent reports 4.211, 4.364)
returns
[[0.81   1.0362]
 [0.7817 1.    ]]
Sargent reports  [[0.81, 1.0362], [0.7817, 1.00]]
```

Both match the values reported in the text.

Notice the pattern in R: the return on currency is poor whenever the government is about to run a deficit (the first column), because new issue dilutes the outstanding stock.

The system (91.5) is nonlinear, and Sargent warns that in general it has more than one solution.

Economy 2 has a second stationary equilibrium, which the solver finds if we start it at low saving rates.

```
s2_hi, R2_hi = stationary_equilibrium(G2, P2, guess=1.2)    # the high-inflation␣
↪equilibrium
print(f"saving rates       {np.round(s2_hi, 4)}")
print(f"returns\n{np.round(R2_hi, 4)}")
```

```
saving rates       [1.1896 1.6343]
returns
[[0.3275 1.3738]
 [0.2384 1.    ]]
```

Saving is far lower, so each deficit must be financed by a large proportional issue of new currency, and the return on currency held into a deficit state falls from about 0.8 to between 0.24 and 0.33.

All its returns are positive, so this is a genuine stationary equilibrium, with much higher inflation.

The learning economy below never finds it, and Exercise 91.7.3 shows why.

91.3.2 Learning by successive generations

Now withdraw the equilibrium from the agents.

They are still assumed to know their own utility function and to remember the experience of earlier agents like themselves, but they are *not* told the distribution of returns, which is exactly the object that (91.4) requires.

Instead, each generation adjusts its predecessors' saving decision in the direction that *realized* experience suggests would have been better.

The realized lifetime utility of an agent who saved s and earned R is

$$U(s) = u(w_1 - s) + u(w_2 + sR),$$

with derivatives

$$U'(s) = -u'(w_1 - s) + u'(w_2 + sR)R, \qquad U''(s) = u''(w_1 - s) + u''(w_2 + sR)R^2.$$

The connection to equilibrium is that in a rational expectations equilibrium $V'(s_i) = E_t U'(s_i)$ and $V''(s_i) = E_t U''(s_i)$, where E_t conditions on $G_t = \bar{G}_i$.

So an agent who could compute $E_t U'$ would just set it to zero and be done.

Our agents cannot, and must estimate it from experience instead.

They use a **Robbins–Monro** algorithm, applied state by state.

Let τ_i count the number of times the deficit state $\bar{G}_i$ has been visited so far, and let γ_τ be a decreasing gain sequence.

The rules are

$$\begin{aligned} M(i,\tau_i+1) &= M(i,\tau_i) + \gamma_{\tau_i}(U''(s(i,\tau_i)) - M(i,\tau_i)), \\ s(i,\tau_i+1) &= s(i,\tau_i) - \gamma_{\tau_i} M(i,\tau_i+1)^{-1} U'(s(i,\tau_i)). \end{aligned} \tag{91.6}$$

M accumulates a running estimate of $E_i U''$, and the saving rule takes a Newton step against it.

The gain γ_τ plays the role of the relaxation parameter λ in *A Peculiar Definition of Bounded Rationality*: each observation moves the rule a fraction γ_τ of the way toward what that observation suggests.

In the relaxation algorithm the gain was constant; here it declines, as it does for Bray's least squares learners in that lecture, so each new observation moves the rule less than the one before.

Sargent sets $\gamma_\tau = 1/\tau$ and starts M at -1.

Our code uses $\gamma_\tau = 1/(\tau + 20)$, which keeps the first steps cautious, as if the initial rule had already been learned from 20 observations, and starts M at $U''(s,1)$, the curvature the initial rule would have if currency held its value.

The algorithm is set up so that, *if* the saving rules converge, $M(i,\tau_i) \to E_i U''$ and the limit solves $E_i U'(s_i) = 0$, which is exactly the equilibrium condition (91.4).

Whether they converge, and to which solution when (91.5) has several, is governed by the stability of an associated ordinary differential equation, as explained in *One recursion, many learners*.

Two features of the setup deserve comment.

Two classes of agent. To evaluate a saving decision we must wait until *two* periods of that agent's consumption are known, so the population is split into an "odd" and an "even" subsequence.

Odd agents reset their rule in odd periods and learn only from earlier odd agents; even agents likewise.

Sargent notes that this alternation resembles the designs of laboratory experiments.

Marimon and Sunder [1993], whose experiment appears in Part 2, drew their young subjects afresh at random each period, but the two-currency experiment of Arifovic [1996], discussed in *Exchange Rate Indeterminacy, Learning, and Experiments*, split its subjects into two fixed groups that were young in alternate periods, just like the odd and even classes here.

A projection facility. Nothing in (91.6) prevents a Newton step from pushing saving below the deficit.

The price level implied by (91.1) and (91.2), $p_t = H_{t-1}/(N s_t - G_t)$, would then fail to be positive, and the economy would cease to exist.

We therefore confine each state's saving rate to the interval $[\bar{G}_i/N + 0.4,\ w_1 - 0.4]$, a region where an equilibrium price level exists.

Devices of this kind are needed for the convergence theorems too; *Least Squares Learning in Self-Referential Models* discusses the projection facility in detail.

The simulation draws the whole path of deficit states first, then lets the two classes of agent learn along it.

```
def U_prime(s, R):
    return -1/(w1 - s) + R/(w2 + s*R)

def U_double(s, R):
    return -1/(w1 - s)**2 - R**2/(w2 + s*R)**2


def simulate(G, P, T=50_000, s0=None, seed=0, tau0=20, band=0.4, gain=None):
    """
    Overlapping generations learning a state-contingent saving rule.

    The gain is 1/(τ + tau0), where τ counts a class's visits to a state,
```

(continues on next page)

(continued from previous page)

```
    unless a constant `gain` is supplied.  Returns the history of saving
    rules, shape (T, 2, n) = (time, class, state), and the deficit states.
    """
    n = len(G)
    states = qe.MarkovChain(P).simulate(ts_length=T, init=0, random_state=seed)
    s = np.array(s0, float)
    M = np.array([[U_double(s[j, k], 1.0) for k in range(n)] for j in range(2)])
    τ = np.zeros((2, n), int)
    floor = G/N + band                  # projection facility: saving must cover the␣
↪deficit
    ceiling = w1 - band

    hist = np.empty((T, 2, n))
    prev = None

    for t in range(T):
        i = states[t]                   # today's deficit state
        j = t % 2                       # which class is young this period
        s_t = s[j, i]

        if prev is not None:            # last period's young now learn their return
            j_p, i_p, s_p = prev
            R = (N*s_t - G[i]) / (N*s_p)          # the return on currency
            τ[j_p, i_p] += 1
            γ = 1 / (τ[j_p, i_p] + tau0) if gain is None else gain
            M[j_p, i_p] += γ * (U_double(s_p, R) - M[j_p, i_p])
            s_new = s[j_p, i_p] - γ * U_prime(s_p, R) / M[j_p, i_p]
            s[j_p, i_p] = min(max(s_new, floor[i_p]), ceiling)

        prev = (j, i, s_t)
        hist[t] = s

    return hist, states
```

Note

The simulation never computes a price level.

Because (91.3) expresses the return on currency purely in terms of saving rates and the deficit, the *real* side of the model is self-contained.

This is more than a convenience: the nominal money stock in these economies grows without bound whenever the government runs a deficit, so a simulation that tracked H_t and p_t directly would overflow long before the learning converged.

91.3.3 Do they find it?

We start both classes of agent far from equilibrium, with asymmetric rules, and let each economy run for 50,000 periods.

```
start = [[8.0, 2.0],
         [2.0, 8.0]]                    # deliberately far from equilibrium, and␣
↪asymmetric

hist1, _ = simulate(G1, P1, s0=start)
hist2, _ = simulate(G2, P2, s0=start)

fig, axes = plt.subplots(1, 2, figsize=(11, 4), sharex=True)
for ax, hist, s_star, title in [(axes[0], hist1, s1, "Economy 1:  $\\bar G = (0, 0)$
↪"),
                                (axes[1], hist2, s2, "Economy 2:  $\\bar G = (0.8, 0)$
↪")]:
    for k, colour in enumerate(['C0', 'C1']):
        ax.plot(hist[:, 0, k], color=colour, lw=1, label=f"state {k+1}, even")
        ax.plot(hist[:, 1, k], color=colour, lw=1, ls=':', label=f"state {k+1}, odd")
        ax.axhline(s_star[k], color='k', lw=0.8, ls='--')
    ax.set_xscale('log')
    ax.set_xlabel("$t$  (log scale)")
    ax.set_title(title)
axes[0].set_ylabel("saving rate")
axes[0].legend(frameon=False, fontsize=8)
plt.tight_layout()
plt.show()
```

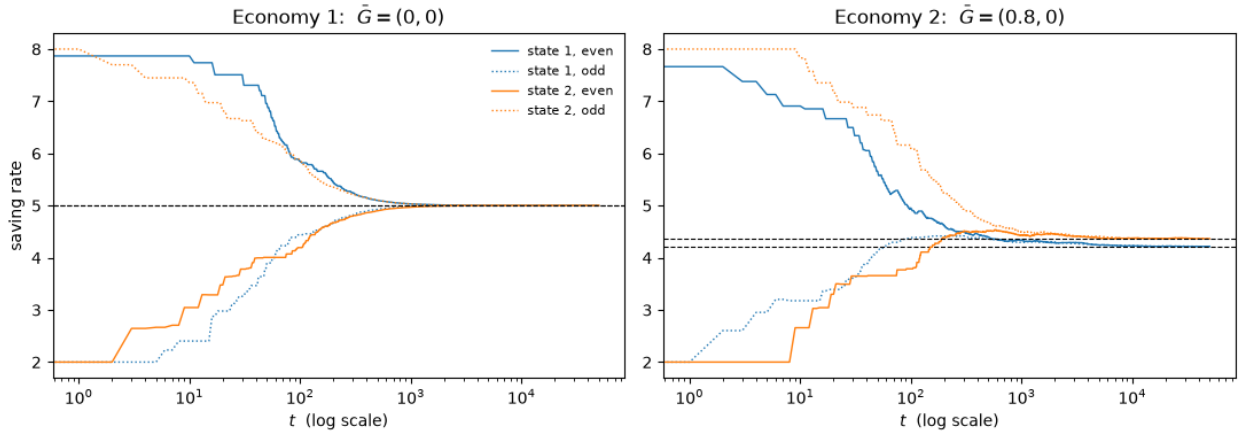


Fig. 91.1: Saving rates converge in both economies

Fig. 91.1 shows that both economies converge, and that both classes of agent converge to the same rule even though neither learns from the other.

The dashed black lines are the equilibrium saving rates computed earlier from (91.5), objects that no agent in the simulation has access to.

The table reports the learned rules, averaged over the two classes, at three horizons.

```
rows = []
for T in (5_000, 50_000, 200_000):
```

(continues on next page)

(continued from previous page)

```
    h1, _ = simulate(G1, P1, T=T, s0=start)
    h2, _ = simulate(G2, P2, T=T, s0=start)
    rows.append([f"learned, T = {T:,}", *h1[-1].mean(axis=0), *h2[-1].mean(axis=0)])
rows.append(["rational expectations", *s1, *s2])

pd.DataFrame(rows, columns=["", "econ 1 state 1", "econ 1 state 2",
                            "econ 2 state 1", "econ 2 state 2"]).set_index("").
↪round(4)
```

	econ 1 state 1	econ 1 state 2	econ 2 state 1	\
learned, T = 5,000	5.0011	5.0007	4.2429	
learned, T = 50,000	5.0001	5.0001	4.2170	
learned, T = 200,000	5.0000	5.0000	4.2156	
rational expectations	5.0000	5.0000	4.2111	

	econ 2 state 2
learned, T = 5,000	4.3930
learned, T = 50,000	4.3658
learned, T = 200,000	4.3693
rational expectations	4.3637

Convergence is slow, which is what a gain of order $1/\tau$ buys you, but it is convergence, and to the low-inflation equilibrium of economy 2.

91.3.4 How much do the agents have to be told?

The agents learn a *separate* saving rate for each level of the deficit: in effect they learn the policy function $s = f(G)$ **non-parametrically**, one number per state.

That is feasible here because there are only two states, and even so it has a cost.

A state that is rarely visited is learned about slowly, because its observations arrive slowly; Exercise 91.7.1 makes this concrete.

When the number of states is large, or the state is continuous, one number per state is hopeless, and agents must learn a parametric rule instead.

Learning, Approximation, and Equilibrium Computation takes up what that buys and what it costs.

91.4 Part 2: a constant deficit and two steady states

Now strip out the randomness.

Let the government finance a *constant* deficit G per young person, and let utility be $\ln c_1 + \ln c_2$ as before.

With perfect foresight over the price level, the young save

$$s_t = \frac{w_1 - w_2 \pi_t}{2}, \qquad \pi_t \equiv \frac{p_{t+1}}{p_t}, \tag{91.7}$$

where π_t is the gross inflation rate, the reciprocal of the return on currency.

Let $h_t = H_t/(N p_t)$ be real balances per young person.

The government's budget constraint in real terms is $h_t = h_{t-1}/\pi_{t-1} + G$, and equilibrium requires $h_t = s_t$.

Eliminating h between these gives an autonomous difference equation in the inflation rate,

$$\pi_{t+1} = g(\pi_t) \equiv A_1 - \frac{A_2}{\pi_t}, \qquad A_1 = \frac{w_1 + w_2 - 2G}{w_2}, \qquad A_2 = \frac{w_1}{w_2}. \tag{91.8}$$

Stationary equilibria are the fixed points of g, and since g is a hyperbola there are generically two of them.

We use the parameters of Marimon and Sunder [1993]'s "Economy 7C": $w_1 = 6$, $w_2 = 1$, $G = 1$.

```
w1_d, w2_d, G_d = 6.0, 1.0, 1.0
A1 = (w1_d + w2_d - 2*G_d) / w2_d
A2 = w1_d / w2_d

def g_map(π):
    return A1 - A2/π

π_low, π_high = np.sort(np.roots([1, -A1, A2]))
print(f"A1 = {A1},  A2 = {A2}")
print(f"stationary gross inflation rates: {π_low} and {π_high}")
print(f"   i.e. net inflation of {100*(π_low-1):.0f}% and {100*(π_high-1):.0f}% per␣
↪period")
print(f"\nslope of g at the low  rate: {A2/π_low**2:.4f}   -> unstable under RE␣
↪dynamics")
print(f"slope of g at the high rate: {A2/π_high**2:.4f}   -> stable under RE dynamics
↪")
```

```
A1 = 5.0,  A2 = 6.0
stationary gross inflation rates: 2.0 and 3.0
   i.e. net inflation of 100% and 200% per period

slope of g at the low  rate: 1.5000   -> unstable under RE dynamics
slope of g at the high rate: 0.6667   -> stable under RE dynamics
```

There it is: the *low*-inflation stationary equilibrium is **unstable** under the rational expectations dynamics, and the high-inflation one is stable.

The case $G = 0$ connects this model to the bubbles of *A Peculiar Definition of Bounded Rationality*.

With no deficit, $A_1 = 1 + w_1/w_2$, and the fixed points of g are $\pi = 1$ and $\pi = w_1/w_2$.

These are the bubbleless, quantity-theory equilibrium of *The bubble* with a constant money supply, and the inflation rate that every one of its bubble paths approaches.

The scales differ.

That lecture endows agents with $2w_1$ and $2w_2$, so its demand for real balances is $w_1 - w_2\pi$, while saving here is $(w_1 - w_2\pi)/2$.

It also fixes the growth rate μ of the money supply where we fix the deficit G, and the two models coincide when $\mu = 1$ and $G = 0$.

The stability analysis below therefore answers that lecture's question of which bubble we are in: least squares learners kill the bubble (Exercise 91.7.5).

```
π_grid = np.linspace(1.35, 4.2, 400)

fig, ax = plt.subplots(figsize=(6.5, 5.5))
ax.plot(π_grid, g_map(π_grid), lw=1.6, label=r"$g(\pi) = A_1 - A_2/\pi$")
ax.plot(π_grid, π_grid, 'k--', lw=0.8, label="45 degree line")
```

(continues on next page)

(continued from previous page)

```
# cobweb the RE dynamics away from the low steady state
π = π_low + 0.05
for _ in range(30):
    nxt = g_map(π)
    ax.plot([π, π], [π, nxt], color='C3', lw=0.9)
    ax.plot([π, nxt], [nxt, nxt], color='C3', lw=0.9)
    π = nxt

for p, name in [(π_low, "low"), (π_high, "high")]:
    ax.plot(p, p, 'ko', ms=6)
    ax.annotate(f"{name}\n$\\pi = {p:.0f}$", (p, p), textcoords="offset points",
                xytext=(10, -22), fontsize=9)
ax.set_xlabel(r"$\pi_t$")
ax.set_ylabel(r"$\pi_{t+1}$")
ax.legend(frameon=False, loc='upper left')
plt.show()
```

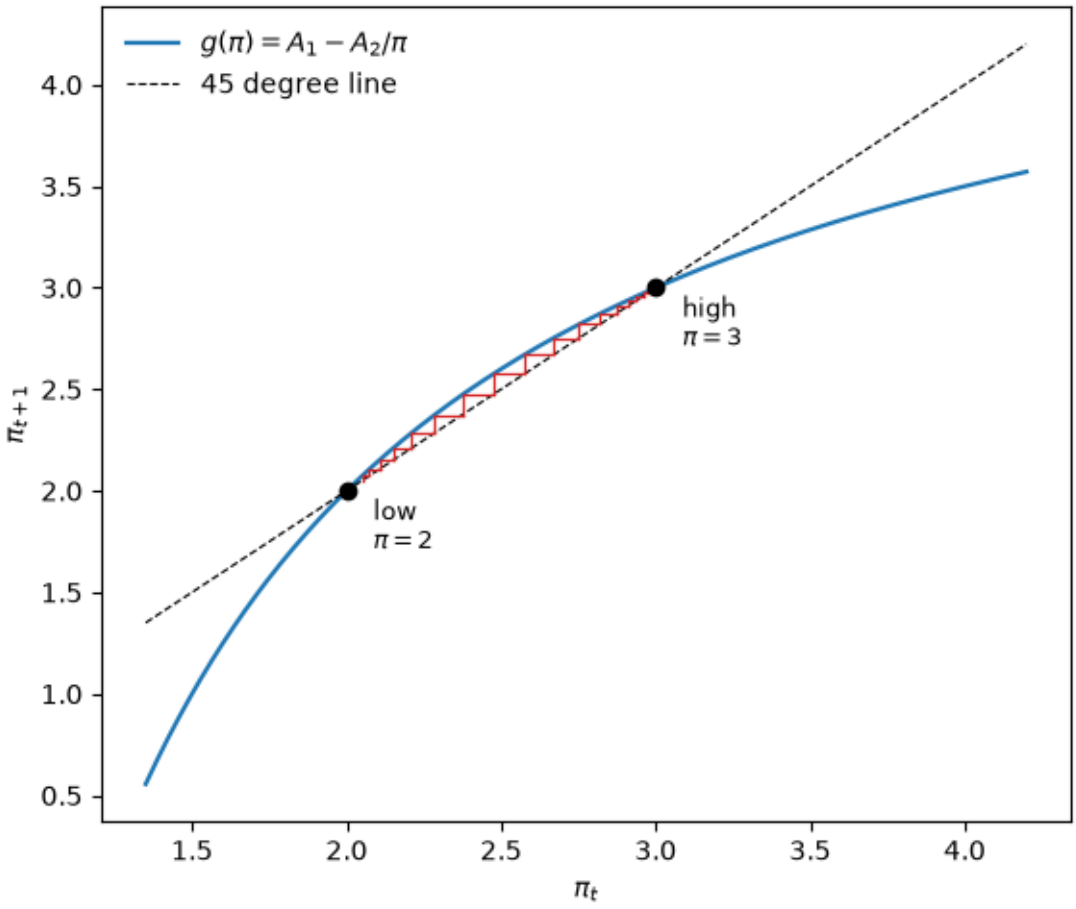


Fig. 91.2: Rational expectations inflation dynamics

In Fig. 91.2 the red staircase starts just above the low-inflation steady state and marches away from it, toward the high-inflation one.

The table follows rational expectations paths from four initial inflation rates.

```
def re_path(π0, T=25):
    path = [π0]
    for _ in range(T):
        nxt = g_map(path[-1])
        path.append(nxt if nxt > 0 else np.nan)
    return np.array(path)

pd.DataFrame({f"$\\pi_0 = {p}$": re_path(p) for p in (1.99, 2.01, 2.5, 4.0)}
             ).rename_axis("t").round(4).iloc[[0, 1, 2, 5, 10, 25]]
```

t	$\pi_0 = 1.99$	$\pi_0 = 2.01$	$\pi_0 = 2.5$	$\pi_0 = 4.0$
0	1.9900	2.0100	2.5000	4.0000
1	1.9849	2.0149	2.6000	3.5000
2	1.9772	2.0222	2.6923	3.2857
5	1.9187	2.0712	2.8836	3.0705
10	0.6693	2.3681	2.9830	3.0087
25	NaN	2.9961	3.0000	3.0000

Started fractionally below the low steady state the economy collapses (inflation computed from (91.8) turns negative, so no positive price level exists); started fractionally above it, the economy converges to $\pi = 3$.

This matters because of how the two equilibria rank.

The comparative statics of the low-inflation equilibrium are "classical": raising the deficit G lowers A_1, which raises the low stationary inflation rate.

The comparative statics of the high-inflation equilibrium are the reverse, because the economy is on the wrong side of an inflation-tax **Laffer curve**; see Exercise 91.7.4.

And the low-inflation equilibrium Pareto-dominates every other equilibrium of this model, stationary or not.

(The ranking relies on everyone in a generation being alike; Sargent notes that in versions with borrowers and lenders within a generation the equilibria are not Pareto-comparable.)

So the rational expectations dynamics repel the economy from the Pareto-optimal outcome, and most classical doctrine in monetary theory depends on selecting the equilibrium that those dynamics reject.

91.4.1 Least squares dynamics

Marcet and Sargent [1989] studied a version of this economy in which agents do not have perfect foresight but instead **run a regression**.

Each period they regress the price level on its own lagged value and use the fitted slope as their forecast of gross inflation,

$$\beta_t = \frac{\sum_{s<t} p_s p_{s-1}}{\sum_{s<t} p_{s-1}^2}, \qquad \pi_t^e = \beta_t,$$

and then save $s_t = (w_1 - w_2\beta_t)/2$ according to (91.7).

Since β_t uses data through $t-1$ only, there is no simultaneity: beliefs determine saving, saving determines the price level, and the new price level enters next period's regression.

Because $p_s = \pi_{s-1} p_{s-1}$, the estimator can be written

$$\beta_t = \frac{\sum_{s<t} \pi_{s-1}\, p_{s-1}^2}{\sum_{s<t} p_{s-1}^2},$$

a weighted average of *past realized inflation rates*, with weights proportional to p_{s-1}^2.

Those weights are far from equal.

Written recursively,

$$\beta_{t+1} = \beta_t + \gamma_t\,(\pi_{t-1} - \beta_t), \qquad \gamma_t = \frac{p_{t-1}^2}{\sum_{s\le t} p_{s-1}^2},$$

least squares is adaptive expectations whose gain γ_t is the newest observation's share of the weight.

During an inflation the price level grows geometrically, so that share does not shrink: when p grows at gross rate π, γ_t tends to $1 - 1/\pi^2$, which is 0.75 at $\pi = 2$.

So least squares in levels does not average a long history; during an inflation it behaves like adaptive expectations with a large gain.

The same form makes the recursion easy to compute.

The price level grows geometrically, so the two sums overflow quickly, but if we rescale both by the newest weight at each step everything stays of order one.

The code treats the initial belief π_0 as a single pseudo-observation.

```
def ls_dynamics(π_init, T=3000, G=G_d):
    """
    Least squares dynamics.  Carries the two regression sums rescaled by the
    newest p^2, so nothing overflows even though the price level does not stop
    growing.  Returns (beliefs, realized inflation, survived?).
    """
    s_prev = (w1_d - w2_d*π_init) / 2
    num, den, π_last = π_init, 1.0, π_init
    betas, infl = [], []

    for t in range(T):
        β = num / den
        s_t = (w1_d - w2_d*β) / 2
        if s_t <= G or not np.isfinite(s_t):
            return np.array(betas), np.array(infl), False   # no positive price level
        π_t = s_prev / (s_t - G)                            # realized gross inflation
        num = num/π_last**2 + π_t
        den = den/π_last**2 + 1.0
        π_last, s_prev = π_t, s_t
        betas.append(β)
        infl.append(π_t)

    return np.array(betas), np.array(infl), True


rows = []
for π0 in (0.5, 1.2, 1.5, 1.9, 2.1, 2.5, 2.9, 2.99, 3.01):
    b, infl, ok = ls_dynamics(π0)
    if ok:
        settled = np.argmax(np.abs(b - π_low) < 1e-3)
        rows.append([π0, "yes", b[-1], settled, np.nan])
    else:
        rows.append([π0, "no", np.nan, np.nan, len(b)])

pd.DataFrame(rows, columns=["initial belief $\\pi_0$", "survived", "belief $\\beta_T$
 ↪",
```

(continues on next page)

(continued from previous page)

```
                            "periods to within 0.001 of $\\pi_{low}$",
                            "period of collapse"]).round(5)
```

```
   initial belief $\pi_0$ survived  belief $\beta_T$  \
0                    0.50      yes               2.0
1                    1.20      yes               2.0
2                    1.50      yes               2.0
3                    1.90      yes               2.0
4                    2.10      yes               2.0
5                    2.50      yes               2.0
6                    2.90      yes               2.0
7                    2.99      yes               2.0
8                    3.01       no               NaN

   periods to within 0.001 of $\pi_{low}$  period of collapse
0                                    11.0                NaN
1                                    13.0                NaN
2                                    13.0                NaN
3                                     3.0                NaN
4                                     3.0                NaN
5                                    13.0                NaN
6                                    10.0                NaN
7                                    17.0                NaN
8                                     NaN                5.0
```

91.4.2 The stability reversal

Least squares learning converges to the **low**-inflation equilibrium from every initial belief below 3 in the table, within 20 periods, even from beliefs just below 3, the steady state that the rational expectations dynamics converge *to*.

A belief just above 3 ends the monetary economy: expected inflation rises until saving no longer covers the deficit.

Fig. 91.3 sets least squares paths beside rational expectations paths.

```
fig, ax = plt.subplots(figsize=(8, 4.5))

for π0 in (1.5, 2.5, 2.95):
    _, infl, _ = ls_dynamics(π0, T=60)
    ax.plot(infl, color='C0', lw=1.2,
            label="least squares dynamics" if π0 == 1.5 else None)

for π0 in (2.05, 2.5, 4.0):
    ax.plot(re_path(π0, 60), color='C3', lw=1.2, ls='--',
            label="rational expectations dynamics" if π0 == 2.05 else None)

ax.axhline(π_low, color='k', lw=0.8, ls=':')
ax.axhline(π_high, color='k', lw=0.8, ls=':')
ax.annotate("low", (52, π_low + 0.04), fontsize=9)
ax.annotate("high", (52, π_high + 0.04), fontsize=9)
ax.set_xlim(0, 60)
ax.set_ylim(1.2, 4.2)
ax.set_xlabel("$t$")
ax.set_ylabel("gross inflation")
ax.legend(frameon=False)
plt.show()
```

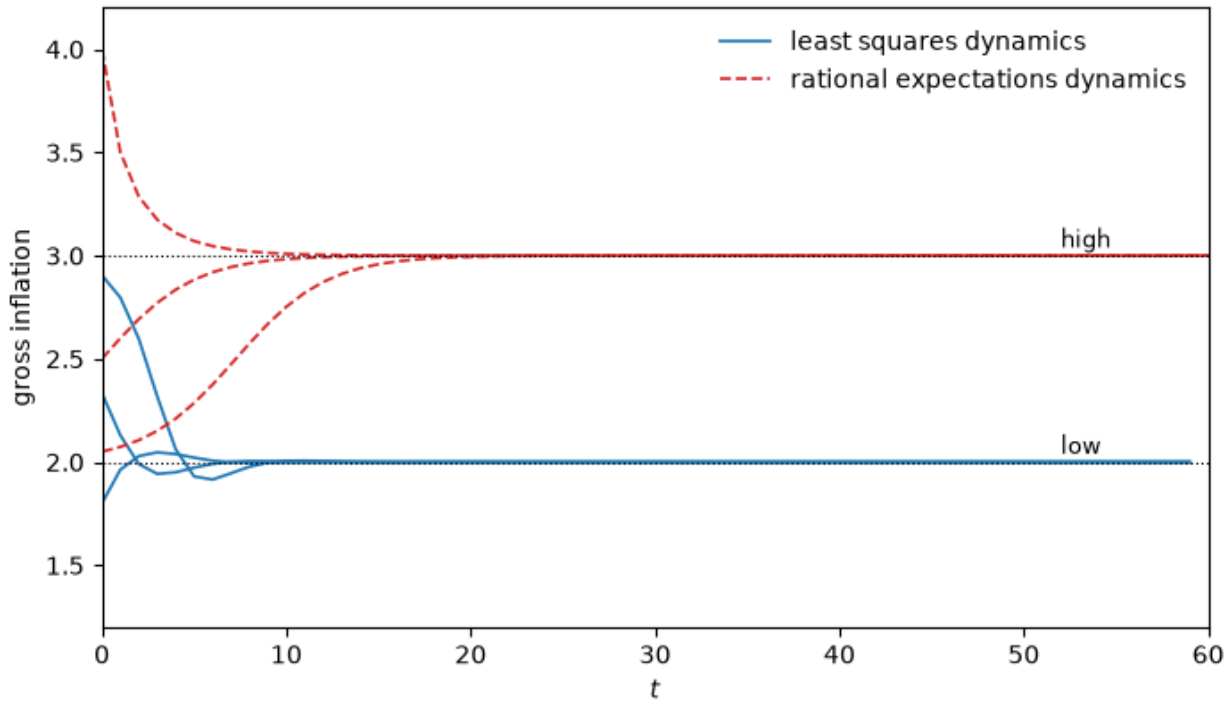


Fig. 91.3: Least squares and rational expectations inflation paths

The two families of paths run in opposite directions.

A bisection locates the edge of the least squares basin of attraction.

```
lo, hi = 2.5, 3.7
for _ in range(40):                    # bisect for the edge of the basin
    mid = (lo + hi) / 2
    lo, hi = (mid, hi) if ls_dynamics(mid)[2] else (lo, mid)
print(f"least squares dynamics survive for initial beliefs up to {lo:.6f}")
print(f"the rational-expectations-stable steady state is      {π_high}")
```

```
least squares dynamics survive for initial beliefs up to 3.000000
the rational-expectations-stable steady state is      3.0
```

The edge of the basin is π_{high} itself.

At exactly $\pi_0 = 3$, beliefs stay put: π_{high} is a rest point of the least squares dynamics, but an unstable one.

We never start at exactly 3, because in floating-point arithmetic a rounding error of order 10^{-16} eventually grows, and the economy then falls to π_{low} or collapses, depending on the sign of the error.

So the steady state that the rational expectations dynamics single out as *the* stable one is precisely the edge of the basin of attraction of the learning dynamics toward the other one.

Here the initial belief matters permanently, because it decides whether the monetary economy survives.

Why the reversal? E-stability

Fix a belief β and suppose agents held it forever.

Saving would be constant at $s(\beta) = (w_1 - w_2\beta)/2$, and the budget constraint $h_t = h_{t-1}/\pi_{t-1} + G$ with $h_t = s(\beta)$ would deliver constant actual inflation

$$T(\beta) = \frac{s(\beta)}{s(\beta) - G}. \tag{91.9}$$

T maps perceived inflation into actual inflation; it is this model's version of the T-map of *A Peculiar Definition of Bounded Rationality*, and its fixed points are the stationary equilibria.

Least squares moves beliefs toward the inflation they produce, so when the gain is small the average motion of beliefs follows the ordinary differential equation

$$\frac{d\beta}{d\tau} = T(\beta) - \beta.$$

A fixed point β^* is stable under this equation, or **E-stable**, if $T'(\beta^*) < 1$ (*One recursion, many learners*, and *Least Squares Learning in Self-Referential Models* for the general theory).

At Economy 7C's parameters $T(\beta) = (6 - \beta)/(4 - \beta)$, so $T'(\beta) = 2/(4 - \beta)^2$.

```
def T_map(β, G=G_d):
    "Actual gross inflation when the belief β is held forever."
    s = (w1_d - w2_d*β) / 2
    return s / (s - G)

def slope(f, x, h=1e-6):
    return (f(x + h) - f(x - h)) / (2*h)

for π_star, name in [(π_low, "low "), (π_high, "high")]:
    print(f"{name} steady state π = {π_star:.0f}:   T'(π) = {slope(T_map, π_star):.3f}
↪,"
          f"   g'(π) = {slope(g_map, π_star):.3f}")
```

```
low  steady state π = 2:   T'(π) = 0.500,   g'(π) = 1.500
high steady state π = 3:   T'(π) = 2.000,   g'(π) = 0.667
```

So $T'(2) = 1/2$ and $T'(3) = 2$: the low steady state is E-stable and the high one is not.

Every slope lies on the opposite side of one from the corresponding slope of g.

Under perfect foresight, a belief slightly above π_{low} is *validated and amplified*, because g has slope 1.5 there.

Under learning, the same belief lowers saving and so raises actual inflation, but only half as much, since T has slope $1/2$; realized inflation comes in *below* the belief, and the belief is revised back down.

At π_{high} it is the other way around.

This is not a coincidence of Economy 7C's parameters.

At every steady state with positive saving, $\operatorname{sign}(T' - 1) = -\operatorname{sign}(g' - 1)$, so whenever there are two steady states, the lower is E-stable and unstable under rational expectations, and the higher is the reverse; Exercise 91.7.6 asks you to prove it.

The reversal is a theorem about this model.

There is a second way to see why the two dynamics disagree.

Under perfect foresight nothing in the past pins down the initial inflation rate, so the rational-expectations-stable high steady state is approached by a continuum of equilibrium paths, while near the unstable low one the only equilibrium path is the stationary one.

Learning selects the steady state that is locally unique.

One caveat remains.

The differential equation describes the average motion of beliefs when each observation moves them only a little, and we saw above that least squares here gives the newest observation a weight near $1 - 1/\pi^2$, three quarters at π_{low}.

So we check the verdict directly.

Near a steady state π^* least squares behaves like adaptive expectations $\beta_{t+1} = \beta_t + \gamma(\pi_{t-1} - \beta_t)$ with $\gamma = 1 - 1/\pi^{*2}$.

Realized inflation $\pi_{t-1} = s(\beta_{t-1})/(s(\beta_t) - G)$ depends on two successive beliefs, so this is a second-order difference equation in β.

The code linearizes it at each steady state and reports its largest root in modulus, at the least squares gain and across all gains between 0 and 1.

```
def local_modulus(π_star, γ):
    "Largest root, in modulus, of the adaptive-expectations recursion linearized at π_
↪star."
    s = (w1_d - w2_d*π_star) / 2
    a = -(w2_d/2) / (s - G_d)                    # ∂π_{t-1} / ∂β_{t-1}
    b = (w2_d/2) * s / (s - G_d)**2              # ∂π_{t-1} / ∂β_t
    return np.max(np.abs(np.roots([1, -(1 - γ + γ*b), -γ*a])))

γ_grid = np.linspace(0.01, 1, 100)
for π_star, name in [(π_low, "low "), (π_high, "high")]:
    γ_ls = 1 - 1/π_star**2
    moduli = [local_modulus(π_star, γ) for γ in γ_grid]
    print(f"{name}: least squares gain {γ_ls:.3f}, largest root {local_modulus(π_star,
↪ γ_ls):.3f};"
          f"  over gains in (0, 1]: from {min(moduli):.3f} to {max(moduli):.3f}")
```

```
low : least squares gain 0.750, largest root 0.612;  over gains in (0, 1]: from 0.
↪500 to 0.995
high: least squares gain 0.889, largest root 2.409;  over gains in (0, 1]: from 1.
↪010 to 2.618
```

At π_{low} every root lies inside the unit circle, whatever the gain; at π_{high} some root lies outside it, whatever the gain.

The E-stability verdict survives the large gain.

It also explains why Bruno and Fischer [1990] found the same reversal in a closely related model using Friedman-style adaptive expectations instead of least squares.

Sargent nevertheless cites Bullard and Duffy [1993] as a warning not to overgeneralize from this result.

In overlapping generations economies whose agents live n periods, they found that for $n \geq 4$ least squares learning fails to converge locally to a rational expectations equilibrium.

91.4.3 Marimon and Sunder's experiment

Which set of dynamics describes people?

Marimon and Sunder [1993] built a laboratory economy to find out.

A fixed number of paying subjects were assigned roles as "young" and "old" in an implementation of exactly this model.

Each period the young submitted a schedule saying how much output they were willing to supply to the old for money at each price, and the experimenters cleared the market against the demand from the old and from the government.

Subjects temporarily on the sidelines played a **forecasting game** with its own cash prize, so the experiment produced a direct record of the price expectations driving the economy.

The session length was unknown to the subjects but governed by publicly announced rules that made the economy equivalent, from their point of view, to one that never ends.

Two findings matter here.

First, the experimental inflation paths were much better approximated by the least squares dynamics converging to the **low** stationary rate than by the rational expectations dynamics converging to the high one.

Sargent reports that this pattern is indicative of all of their results: *on average*, the experimental dynamics are much better approximated by the adaptive dynamics than by the rational expectations dynamics.

Second, the sideliners' forecast errors were comparable in size to those a least squares forecaster would have made using the same data.

The subjects were not doing something exotic, they were doing roughly what the adaptive model says.

The equilibrium that theory calls unstable is the one that both the algorithm and the humans select.

Note

Experimental evidence is not the only way to discriminate.

İmrohoroğlu [1993] estimated a rational expectations version of this model on German hyperinflation data and found the opposite: his estimated equilibrium slides along the *bad* side of the Laffer curve toward the high stationary rate, which is inconsistent with the equilibrium the adaptive dynamics select.

91.4.4 A warning

It is tempting to conclude that adaptive dynamics reliably select the Pareto-superior equilibrium.

Sargent offers a counterexample, shown to him by Benjamin Bental, built from a special case of Brock [1974]'s money-in-the-utility-function model.

An infinitely lived household with endowment y maximizes $\sum_t \delta^t (\ln c_t + \phi \ln(m_t/p_t))$ subject to a sequence of budget constraints, while the government finances a fixed expenditure $G > 0$ by printing currency.

(Sargent writes β and γ for the discount factor δ and the weight ϕ; we keep β for beliefs and γ for gains.)

Under the parameter restriction $\phi(y - G) > G$ the model has a unique stationary equilibrium, with gross inflation

$$\pi = \frac{(y - G)\phi - \delta G}{(y - G)\phi - G},$$

which has the classical property of rising with the deficit.

```
ψ_b, δ_b, y_b = 1.0, 1/1.05, 11.0

def brock_π(G):
    return ((y_b - G)*ψ_b - δ_b*G) / ((y_b - G)*ψ_b - G)

G_vals = (0.5, 1.0, 1.5, 2.0)
pd.DataFrame({
    "$G$": G_vals,
    "restriction $\\phi(y-G) > G$": [ψ_b*(y_b - G) > G for G in G_vals],
    "stationary $\\pi$": [brock_π(G) for G in G_vals],
    "return on currency $1/\\pi$": [1/brock_π(G) for G in G_vals],
}).round(5)
```

```
   $G$  restriction $\phi(y-G) > G$  stationary $\pi$  \
0  0.5                         True           1.00238
1  1.0                         True           1.00529
2  1.5                         True           1.00893
3  2.0                         True           1.01361

   return on currency $1/\pi$
0                     0.99762
1                     0.99474
2                     0.99115
3                     0.98658
```

But the model *also* has a continuum of non-stationary equilibria, indexed by initial price levels below the stationary one, in all of which the gross inflation rate converges to δ — that is, the economy ends in **deflation**, with the return on currency approaching $1/\delta = 1.05$ and real balances exploding.

These exist because the demand for currency implied by $\phi \ln(m/p)$ is so elastic with respect to the rate of return that the government can raise enough seigniorage to finance G even while paying interest on currency through deflation.

Sargent reports that least squares learning here selects the *classical stationary* equilibrium, just as in the overlapping generations model.

The T-map makes that easy to check.

With consumption constant at $y - G$, the household's first-order condition for money gives real balances $h(\pi^e) = \phi(y - G)\pi^e/(\pi^e - \delta)$ when expected inflation is π^e.

The budget constraint $h = h/\pi + G$ then turns a belief held forever into actual inflation $T(\pi^e) = h(\pi^e)/(h(\pi^e) - G)$.

```
def brock_T(πe, G):
    "Actual gross inflation when expected inflation πe is held forever."
    h = ψ_b * (y_b - G) * πe / (πe - δ_b)
    return h / (h - G)

G_b = 1.0
print(f"stationary π = {brock_π(G_b):.5f},   T'(π) = {slope(lambda x: brock_T(x, G_b),
 ↪ brock_π(G_b)):.3f}")
```

```
stationary π = 1.00529,   T'(π) = 0.095
```

The slope is far below one, so the classical stationary equilibrium is E-stable.

The difference from the overlapping generations model is the welfare ranking: in the Brock model, the non-stationary equilibria all Pareto-dominate the classical stationary one that learning picks.

So in both examples least squares dynamics select the *classical* equilibrium.

Whether that equilibrium is the good one is a separate question, and the answer depends on the model.

91.5 A government learning the Phillips curve

The examples so far put adaptive *households* into a monetary economy.

We close with an application in which the adaptive agent is the **government**, learning an econometric relationship that it acts on.

The model, due to Sims [1988] and estimated by Chung [1990], extends the Phillips curve example of Kydland and Prescott [1977].

It nests two stories about post-war U.S. inflation: a government misled by an apparently exploitable Phillips curve, and Phillips-curve fitters whose adaptive methods caught the adverse shift in time.

The Phillips curve lectures that begin with *The Rise and Fall of U.S. Inflation* develop this material in depth, with the same calibration; here we give a condensed version.

91.5.1 The model

Inflation is decomposed into an expected and an unexpected part,

$$\pi_t = g_{t-1} + \eta_t, \tag{91.10}$$

where g_{t-1} is the inflation the public expects (and which the government controls) and η_t is the public's forecast error, orthogonal to earlier information.

The *true* Phillips curve is the natural-rate one: only the **surprise** η_t moves unemployment:

$$U_t = U^* - \theta(\pi_t - g_{t-1}) + u_t = U^* - \theta\eta_t + u_t, \qquad \theta > 0. \tag{91.11}$$

The government does not know this.

It believes unemployment depends on the *level* of inflation, and fits the non-expectational regression

$$U_t = \alpha_0 + \alpha_1 \pi_t + \epsilon_t. \tag{91.12}$$

Each period it minimizes $\frac{1}{2}\mathbb{E}(U_t^2 + \pi_t^2)$ subject to its perceived model, which yields the myopic target

$$g_{t-1} = -\frac{\alpha_0\,\alpha_1}{1+\alpha_1^2}. \tag{91.13}$$

The system is **self-referential**: the government's beliefs α set its policy g through (91.13); the policy shapes the joint distribution of (π, U); and that distribution is what the government's regression (91.12) estimates.

> **Note**
>
> We follow Sargent [1993]'s notation, in which π_t is *net* inflation in percent.
>
> The table translates to *Escaping Nash Inflation* and *Priors, Escapes, and Learning Cycles*; *A note on notation* covers the rest of the Phillips curve lectures.

object	here	*Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*
inflation	π_t	π_n
government's target (systematic) inflation	g_{t-1}	x_n
natural rate of unemployment	U^*	u
inflation surprise	η_t	$\sigma_2 W_{2n}$
unemployment shock	u_t	$\sigma_1 W_{1n}$
belief: intercept, slope	α_0, α_1	γ_{-1}, γ_1 (escaping_nash); a, b (priors)
constant gain	γ	ε
equilibrium with the inflation bias	self-confirming (Kydland–Prescott's "consistent")	self-confirming (Nash)

```
U_star, θ_pc = 5.0, 1.0              # natural rate 5%, Phillips slope
σ_η, σ_u = 0.3, 0.3                  # inflation-surprise and unemployment shocks

def govt_target(α):
    "Government's myopic optimal target inflation."
    α0, α1 = α
    return -α0 * α1 / (1 + α1**2)
```

91.5.2 The self-confirming equilibrium

Kydland and Prescott's **consistent** equilibrium is what the Phillips curve lectures call a **self-confirming** equilibrium: a belief vector α whose induced policy generates data whose regression returns the same α.

Under a constant policy g, the surprise η_t is the only thing moving inflation around its mean, so the population regression of U on π has slope exactly $-\theta$ and intercept $U^* + \theta g$.

Solving that fixed point with the decision rule (91.13) gives Kydland and Prescott's time-consistent outcome,

$$\alpha_0 = (1 + \theta^2)U^*, \qquad \alpha_1 = -\theta, \qquad g = \theta U^*.$$

```
α_consistent = np.array([(1 + θ_pc**2) * U_star, -θ_pc])
g_consistent = θ_pc * U_star
print(f"self-confirming beliefs α = {α_consistent},  target inflation g = {g_
 ↪consistent}")
print(f"optimal (Ramsey) outcome:                    target inflation g = 0")
```

```
self-confirming beliefs α = [10. -1.],  target inflation g = 5.0
optimal (Ramsey) outcome:                    target inflation g = 0
```

This is the **inflation bias**.

The government inflates at rate θU^* and gets nothing for it: because only surprises move unemployment, average unemployment is U^* either way.

Had the government understood the true model (91.11), it would have chosen $g = 0$, the Ramsey outcome, with the same unemployment and none of the inflation.

The self-confirming equilibrium is unique, and the Ramsey outcome is not an equilibrium: to choose $g = 0$ the government would have to believe there is no tradeoff, but the data that $g = 0$ generates show the slope $-\theta$.

The self-confirming equilibrium is E-stable, so a government that believes its coefficients are constant and estimates them by least squares with a decreasing gain is drawn toward it asymptotically (*Self-Confirming Equilibria*, *Adaptive Learning and Escape Dynamics*); in finite samples a falling gain can also let an early escape stick, because it slows the return as well.

That is the natural-rate story: the government keeps inflating because its own policy keeps producing data that confirm the tradeoff it believes in.

91.5.3 Constant-gain learning and escapes

Sims [1988] and Chung [1990] made the government *less* sure that its coefficients are constant: it suspects they drift as a random walk and estimates them with a Kalman filter (see *A First Look at the Kalman Filter*), which discounts old data in favor of recent data.

Sims found two kinds of history.

Some stayed near the self-confirming equilibrium; others converged to a process that spends most of its time near the zero-inflation Ramsey outcome, because a government willing to entertain drifting coefficients learns, without a big inflation, that the tradeoff is poor.

Which kind emerges depends partly on parameters, especially the drift variances the government entertains.

Our government learns in a simpler way, by **constant-gain recursive least squares**, the version studied by Sargent [1999] and Cho *et al.* [2002], which approximates a Kalman filter with a particular prior about the drift [Sargent and Williams, 2005].

With $x_t = (1, \pi_t)^\top$ and a constant gain γ, it sets g_{t-1} by (91.13) using its estimate α_t, observes (π_t, U_t), and updates

$$M_t = M_{t-1} + \gamma(x_t x_t^\top - M_{t-1}), \qquad \alpha_{t+1} = \alpha_t + \gamma M_t^{-1} x_t (U_t - x_t^\top \alpha_t).$$

Unlike a $1/t$ gain, which eventually freezes the estimates, a constant gain keeps the government paying attention to what just happened.

We start it at the self-confirming beliefs, with $M_0 = I$ and $\gamma = 0.02$.

```
def constant_gain(T=40_000, gain=0.02, seed=1):
    """
    Government estimates its Phillips curve by constant-gain recursive least
    squares, starting from the self-confirming beliefs.  Returns the paths of
    target inflation g_{t-1}, inflation π_t and unemployment U_t.
    """
    rng = np.random.default_rng(seed)
    α = α_consistent.copy()
    M = np.eye(2)                                          # moment-matrix estimate
    g_path, π_path, U_path = np.empty(T), np.empty(T), np.empty(T)
    for t in range(T):
        g = govt_target(α)
        η = σ_η * rng.standard_normal()
        π = g + η
        U = U_star - θ_pc * η + σ_u * rng.standard_normal()    # true Phillips curve
        x = np.array([1.0, π])
        M = M + gain * (np.outer(x, x) - M)
        α = α + gain * np.linalg.solve(M, x * (U - x @ α))
        g_path[t], π_path[t], U_path[t] = g, π, U
    return g_path, π_path, U_path

g_path, _, _ = constant_gain()
```

(continues on next page)

(continued from previous page)

```
fig, ax = plt.subplots(figsize=(9, 4))
ax.plot(g_path, lw=0.5)
ax.axhline(g_consistent, color='C3', ls='--', lw=1, label="self-confirming␣
↪equilibrium (inflation bias)")
ax.axhline(0, color='C2', ls=':', lw=1, label="Ramsey / zero inflation")
ax.set_ylim(-1, 6.8)
ax.set_xlabel("$t$")
ax.set_ylabel("target inflation $g_t$")
ax.legend(frameon=False, loc='upper center', ncol=2)
plt.show()
```

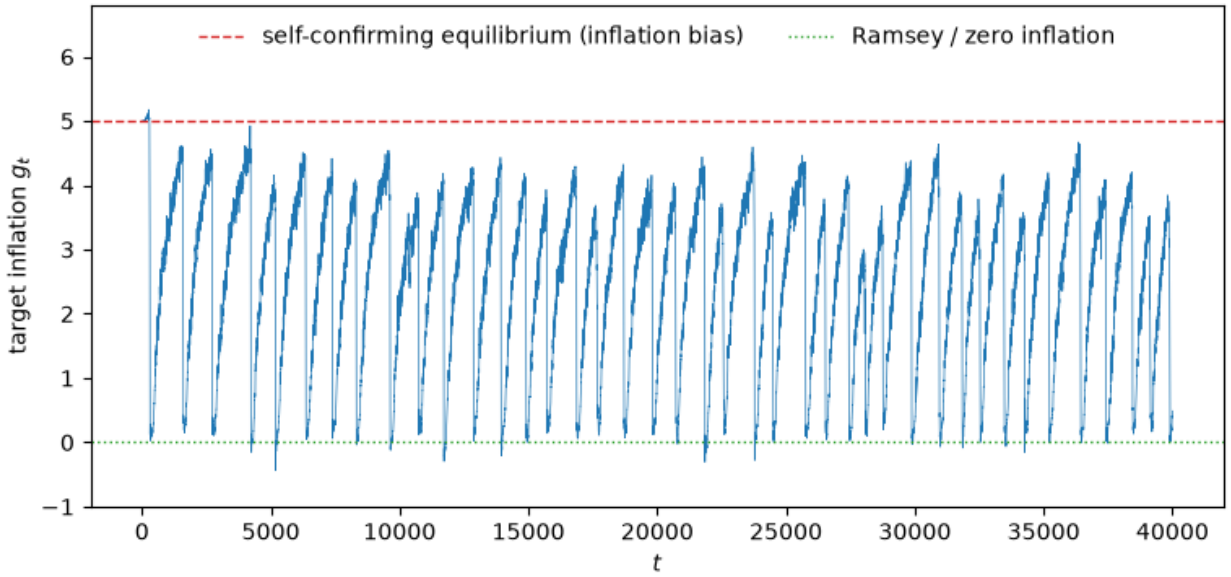


Fig. 91.4: Constant-gain learning and recurrent escapes toward zero inflation

Fig. 91.4 shows that the constant-gain government does not settle at the inflation bias.

It repeatedly and abruptly **escapes** toward zero inflation, then climbs back toward the inflation bias, and at this gain it usually escapes again before it gets there: a recurring sawtooth.

The escapes reinforce themselves.

A run of shocks that makes unemployment look unresponsive to inflation flattens the estimated slope α_1, and the government cuts inflation.

That predictable fall in inflation does not raise unemployment, since only surprises move it, so the government's own policy generates data that confirm the flatter slope and push inflation further toward zero.

This recurrent escape-and-return is the dynamics of Sargent [1999] and Cho *et al.* [2002], not Sims's, whose escaping histories stayed near the Ramsey outcome; *Priors, Escapes, and Learning Cycles* traces the difference to how the government sets up its Kalman filter.

Adaptive Learning and Escape Dynamics and *Escaping Nash Inflation* show how the frequency of escapes depends on the gain, and how Sargent [1999] turned them into a theory of how U.S. inflation was conquered in the early 1980s.

Chung went further and estimated the government's learning rule itself on post-war U.S. data; Exercise 91.7.7 asks you to solve a stripped-down version of his problem.

91.6 Concluding remarks

Part 1 showed that a system of adaptive agents can find a rational expectations equilibrium it was never told about.

It also showed that *which* equilibrium it finds is not up to the agents: economy 2 has two stationary equilibria, and learning finds only the low-inflation one.

And it made visible how much the complexity of the environment governs the speed of learning: a rare deficit state is learned about slowly.

Part 2 sharpened the selection result.

Where the model has two steady states, least squares learning picks out the one that the rational expectations dynamics reject, and laboratory subjects go the same way.

E-stability explains why: at every steady state with positive saving the slopes of the T-map and of the rational expectations map lie on opposite sides of one.

The same logic disposes of the bubbles of *A Peculiar Definition of Bounded Rationality* (Exercise 91.7.5).

That is the best case in Sargent [1993] for treating adaptive dynamics as an equilibrium selection device.

The Brock counterexample marks its limit.

In both models learning selects the classical equilibrium, which is the good one in one model and a bad one in the other.

The selection is a fact about the dynamics, not a welfare theorem, and *1993 Prospects for Bounded Rationality in Macroeconomics* returns to how much weight it can bear.

The Phillips curve application turned the *gain* into the object of interest.

With a decreasing gain the government settles at the inflation bias; with a constant gain it keeps escaping toward the Ramsey outcome, which is not an equilibrium at all.

The Phillips curve lectures that begin with *The Rise and Fall of U.S. Inflation* take this up in depth.

Learning, Approximation, and Equilibrium Computation takes up the question Part 1 left open: what an agent must give up when the state is continuous.

Exchange Rate Indeterminacy, Learning, and Experiments then returns to this overlapping generations economy with two currencies, where the rational expectations equilibria form a continuum and the limit of learning depends on where beliefs start.

Further ahead, *Money as a Medium of Exchange among Artificially Intelligent Agents* takes the selection question to a setting where the competing equilibria are not high and low inflation but different *monetary institutions*: which good becomes money.

91.7 Exercises

Exercise 91.7.1

Because agents learn a separate saving rate for each deficit state, they can only learn about a state as fast as that state occurs.

Compare two economies with $\bar{G} = (0.8, 0)$ that differ only in the transition matrix: one where the two states are equally likely, and one where the high-deficit state is rare.

For each, compute the stationary equilibrium, run the learning simulation for 30,000 periods, and report how far each state's learned saving rate is from its target along with how often that state was visited.

A single run can mislead, so repeat the comparison across 20 seeds.

Solution

For each chain we record, in every run, the error of each state's learned rule (averaged over the two classes) and the number of periods spent in each state.

```
G_ex = np.array([0.8, 0.0])
cases = {"equally likely": np.array([[0.5, 0.5], [0.5, 0.5]]),
         "high deficit rare": np.array([[0.5, 0.5], [0.03, 0.97]])}
seeds = range(20)

rows, ratios = [], {}
for name, P in cases.items():
    s_star, _ = stationary_equilibrium(G_ex, P)
    errors, visits = [], []
    for seed in seeds:
        hist, states = simulate(G_ex, P, T=30_000, s0=start, seed=seed)
        errors.append(np.abs(hist[-1].mean(axis=0) - s_star))
        visits.append(np.bincount(states, minlength=2))
    errors, visits = np.array(errors), np.array(visits)
    rmse = np.sqrt(np.mean(errors**2, axis=0))
    ratios[name] = errors[:, 0] / errors[:, 1]
    for k in (0, 1):
        rows.append([name, k+1, visits[:, k].mean(), s_star[k], rmse[k]])

display(pd.DataFrame(rows, columns=["chain", "state", "visits (mean over seeds)",
                                    "equilibrium $s_i$", "RMS error over seeds"]
                     ).set_index(["chain", "state"]).round(4))

for name, r in ratios.items():
    print(f"{name:18s}: error ratio state 1 / state 2 across seeds: "
          f"median {np.median(r):.1f}, range {r.min():.1f} to {r.max():.1f}, "
          f"state 1 worse in {np.sum(r > 1)} of {len(r)} seeds")
```

```
                          visits (mean over seeds)  equilibrium $s_i$  \
chain             state
equally likely    1                      15001.05             4.4965
                  2                      14998.95             4.4965
high deficit rare 1                       1707.65             4.6828
                  2                      28292.35             4.9646

                          RMS error over seeds
chain             state
equally likely    1                     0.0053
                  2                     0.0051
high deficit rare 1                     0.0108
                  2                     0.0015
```

```
equally likely    : error ratio state 1 / state 2 across seeds: median 0.8,␣
 ↪range 0.0 to 181.2, state 1 worse in 9 of 20 seeds
high deficit rare : error ratio state 1 / state 2 across seeds: median 9.1,␣
 ↪range 0.2 to 39.8, state 1 worse in 15 of 20 seeds
```

When the two states are equally likely the chain is IID, so both states face the same distribution of future returns and the equilibrium has $s_1 = s_2$.

Each state is visited about half the time, and the two rules are learned about equally well.

When the high-deficit state is rare it is visited only about 6 percent of the time, roughly 1,700 periods out of 30,000, and the root-mean-square error in its rule is several times that of the common state, even though both have had 30,000 periods of calendar time to settle.

The ratio is large in most runs but varies a good deal from seed to seed, which is why a single run is a poor basis for the comparison.

Sargent's comment is worth keeping in mind: in terms of *unconditional expected utility*, failing to learn what to do in a rarely visited state may cost very little.

The cost of slow learning is not proportional to the size of the error.

Note also that changing the transition matrix changes the equilibrium itself — beliefs about the persistence of deficits feed straight into (91.5) — so the comparison has to be against a separately recomputed target for each chain.

Exercise 91.7.2

In a stochastic environment, the declining gain $\gamma_\tau = 1/(\tau + 20)$ is what makes (91.6) converge.

An agent who suspected the environment might be drifting would use a **constant gain** instead, weighting recent experience more heavily forever.

Use the `gain` argument of `simulate` to give the agents a constant gain, and describe what happens to the limiting behavior of the saving rules.

Compare gains of 0.02 and 0.005 against the declining-gain benchmark, in both economy 1 and economy 2.

Solution

For each economy and gain we average the last 5,000 periods of the rules (over the two classes), and report the mean error relative to equilibrium and the standard deviation.

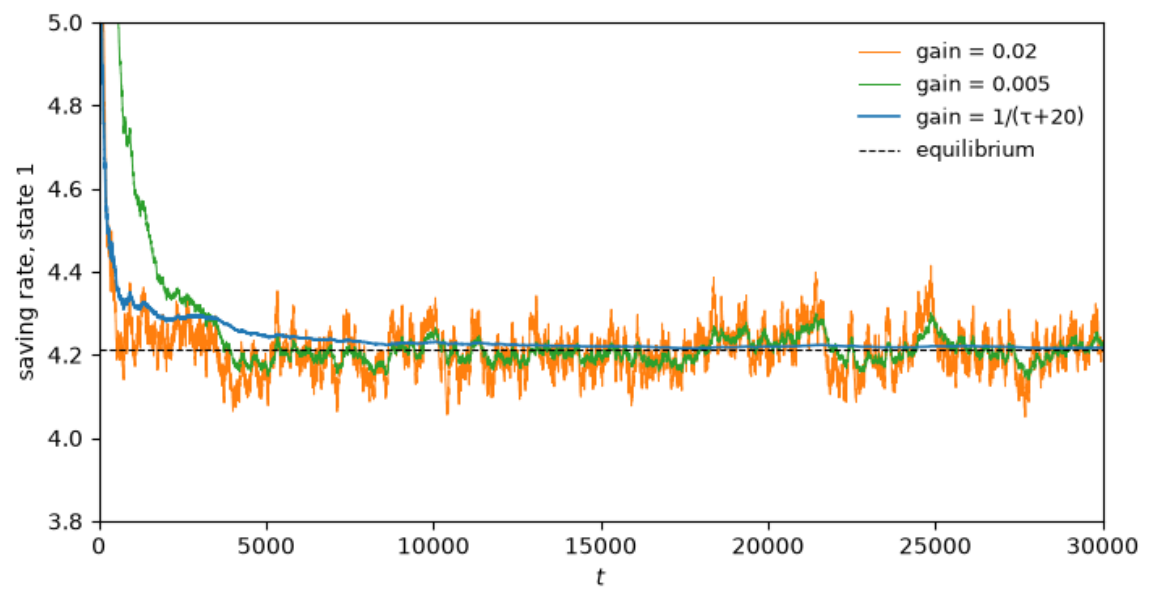
```
rows, paths = [], {}
for G, P, s_star, econ in [(G1, P1, s1, "economy 1"), (G2, P2, s2, "economy 2")]:
    for label, gain in [("0.02", 0.02), ("0.005", 0.005), ("1/(τ+20)", None)]:
        hist, _ = simulate(G, P, T=30_000, s0=start, gain=gain)
        paths[econ, label] = hist[:, 0, 0]
        tail = hist[-5_000:].mean(axis=1)        # average the two classes
        rows.append([econ, label,
                     *[f"{e:.1e}" for e in tail.mean(axis=0) - s_star],
                     *[f"{d:.1e}" for d in tail.std(axis=0)]])

pd.DataFrame(rows, columns=["economy", "gain", "mean error $s_1$", "mean error $s_2
↪$",
                            "s.d. $s_1$", "s.d. $s_2$"]).set_index(["economy",
↪"gain"])
```

		mean error s_1	mean error s_2	s.d. s_1	s.d. s_2
economy	gain				
economy 1	0.02	-8.9e-16	-8.9e-16	2.2e-14	2.2e-14
	0.005	-4.4e-15	-4.4e-15	8.3e-14	8.3e-14
	1/(τ+20)	1.7e-04	1.6e-04	9.3e-06	8.2e-06
economy 2	0.02	-2.9e-03	2.2e-02	3.7e-02	3.9e-02
	0.005	2.5e-03	2.1e-02	1.8e-02	1.6e-02
	1/(τ+20)	9.1e-03	1.0e-02	1.1e-03	1.8e-03

The figure follows the even class's rule for state 1 in economy 2.

```
fig, ax = plt.subplots(figsize=(8, 4))
for label, colour, lw in [("0.02", 'C1', 0.7), ("0.005", 'C2', 0.7), ("1/(τ+20)",
 ↪'C0', 1.2)]:
    ax.plot(paths["economy 2", label], color=colour, lw=lw, label=f"gain = {label}
 ↪")
ax.axhline(s2[0], color='k', lw=0.8, ls='--', label="equilibrium")
ax.set_xlim(0, 30_000)
ax.set_ylim(3.8, 5.0)
ax.set_xlabel("$t$")
ax.set_ylabel("saving rate, state 1")
ax.legend(frameon=False, fontsize=9)
plt.show()
```



In economy 2 the constant-gain rules do not converge.

They settle into a *neighborhood* of the equilibrium and then keep rattling around inside it forever, because each new observation is always given the same weight and so never stops moving the estimate.

Shrinking the gain tightens the band but does not eliminate it.

In economy 1, by contrast, the constant-gain rules converge to the equilibrium to machine precision, and faster than the declining-gain rule.

Nothing random hits the agents at that equilibrium: with no deficit, $R \equiv 1$ in every state, so once the rules are right every observation confirms them.

So the correct statement is conditional.

When there are intrinsic shocks at the equilibrium, only a gain that declines to zero delivers convergence to a point, which is why the convergence theorems for these systems require it.

Without such shocks a constant gain converges too, and quickly.

The table also shows the price of a declining gain in economy 2.

After 30,000 periods the $1/(\tau+20)$ rule is still about 0.01 above the equilibrium in both states, and since its standard deviation is several times smaller, that error is bias: the rule has not yet arrived.

The constant-gain rules, by contrast, wander around the equilibrium with standard deviations of about 0.02 to 0.04, slowly enough that even a 5,000-period average can sit a couple of hundredths away from it.

What the constant-gain agent buys in exchange for its permanent noise is speed and adaptability: if the deficit process were to change, the declining-gain agent would barely notice, while the constant-gain agent would track the change.

That trade-off reappears throughout this literature, and in the Phillips curve application above.

Exercise 91.7.3

Economy 2 has a second, high-inflation stationary equilibrium, `s2_hi`, computed above.

In the simulations of Fig. 91.1 the projection facility keeps state 1's saving above $\bar{G}_1 + 0.4 = 1.2$, which already rules it out, since its s_1 is below 1.2.

1. Lower the projection band to 0.01, start both classes at the high-inflation equilibrium, and run the learning economy for 20,000 periods with 20 different seeds. Where do the saving rules go, and how often does the projection facility bind?
2. Explain the result with the associated ordinary differential equation. When both classes use the rule s, the average Newton step in state i is $-E_i U'(s_i)/E_i U''(s_i)$, where E_i averages over next period's state with the returns $R(i,j) = (Ns_j - \bar{G}_j)/(Ns_i)$. Compute the eigenvalues of the Jacobian of this map at both equilibria.

Solution

We start every run exactly at the high-inflation equilibrium and record where the rules (averaged over the two classes) end up, and whether state 1's rule ever hits the floor $\bar{G}_1 + 0.01$.

```
band = 0.01
floor_1 = G2[0]/N + band
final, hit_floor = [], []
for seed in range(20):
    hist, _ = simulate(G2, P2, T=20_000, s0=np.tile(s2_hi, (2, 1)), seed=seed,␣
 ↪band=band)
    final.append(hist[-1].mean(axis=0))
    hit_floor.append(np.any(np.isclose(hist[:, :, 0], floor_1)))
final, hit_floor = np.array(final), np.array(hit_floor)
climbed = final[:, 0] > s2_hi[0] + 1

print(f"climbed toward the low-inflation equilibrium:  {climbed.sum()} of 20 runs,"
      f"  mean final rule {np.round(final[climbed].mean(axis=0), 3)}")
print(f"fell to the projection floor:                  {(~climbed).sum()} of 20␣
 ↪runs,"
      f"  mean final rule {np.round(final[~climbed].mean(axis=0), 3)}")
print(f"hit the floor at least once:                   {hit_floor.sum()} of 20 runs
 ↪")
print(f"low-inflation equilibrium:                     {np.round(s2, 3)}")
```

```
climbed toward the low-inflation equilibrium:  15 of 20 runs,  mean final rule␣
 ↪[4.096 4.254]
fell to the projection floor:                  5 of 20 runs,  mean final rule␣
 ↪[0.81  0.348]
hit the floor at least once:                   12 of 20 runs
low-inflation equilibrium:                     [4.211 4.364]
```

Learning never stays at the high-inflation equilibrium, even when it starts exactly there.

Most runs climb toward the low-inflation equilibrium, covering most of the distance in 20,000 periods and still climbing.

The rest go the other way: state 1's rule falls to the floor, where saving barely covers the deficit, and state 2's rule sinks far below its starting point.

Without the projection facility those economies would end, and many of the runs that eventually climb also touch the floor on the way.

The mean dynamics explain both outcomes.

```
def mean_newton_step(s, G, P):
    "Average Newton step -E_i U'/E_i U'' in each state when both classes use the␣
↪rule s."
    R = (N*s[None, :] - G[None, :]) / (N*s[:, None])
    EU1 = (P * U_prime(s[:, None], R)).sum(axis=1)
    EU2 = (P * U_double(s[:, None], R)).sum(axis=1)
    return -EU1 / EU2

def jacobian(f, x, h=1e-6):
    J = np.empty((len(x), len(x)))
    for k in range(len(x)):
        e = np.zeros(len(x))
        e[k] = h
        J[:, k] = (f(x + e) - f(x - e)) / (2*h)
    return J

for s_eq, name in [(s2, "low inflation "), (s2_hi, "high inflation")]:
    eig = np.linalg.eigvals(jacobian(lambda s: mean_newton_step(s, G2, P2), s_eq))
    print(f"{name} equilibrium {np.round(s_eq, 4)}:  eigenvalues {np.round(np.
↪sort(eig.real), 3)}")
```

```
low inflation  equilibrium [4.2111 4.3637]:  eigenvalues [-1.948 -0.795]
high inflation equilibrium [1.1896 1.6343]:  eigenvalues [-3.696  3.263]
```

At the low-inflation equilibrium both eigenvalues are negative, so it is a stable rest point of the associated differential equation.

At the high-inflation equilibrium one eigenvalue is positive: it is a saddle, and the noise in realized returns pushes the rules off it along the unstable direction, which leads up toward the low-inflation equilibrium on one side and down toward collapse on the other.

This is Part 2's stability reversal in a stochastic economy: as with π_{high} there, the high-inflation equilibrium is the boundary between convergence to the low-inflation equilibrium and the end of the monetary economy.

Learning by itself excludes the high-inflation equilibrium; the projection facility matters on the downside, where it keeps collapsing economies alive.

Exercise 91.7.4

The claim that the high-inflation steady state sits on "the wrong side of a Laffer curve" can be made precise.

In a stationary equilibrium with gross inflation π, real balances are $h = (w_1 - w_2\pi)/2$ and the government collects seigniorage $h(1 - 1/\pi)$.

Plot seigniorage as a function of π, mark the two stationary equilibria, and use the picture to explain why raising G

raises the low stationary inflation rate but *lowers* the high one, and what happens when G gets too large.

Solution

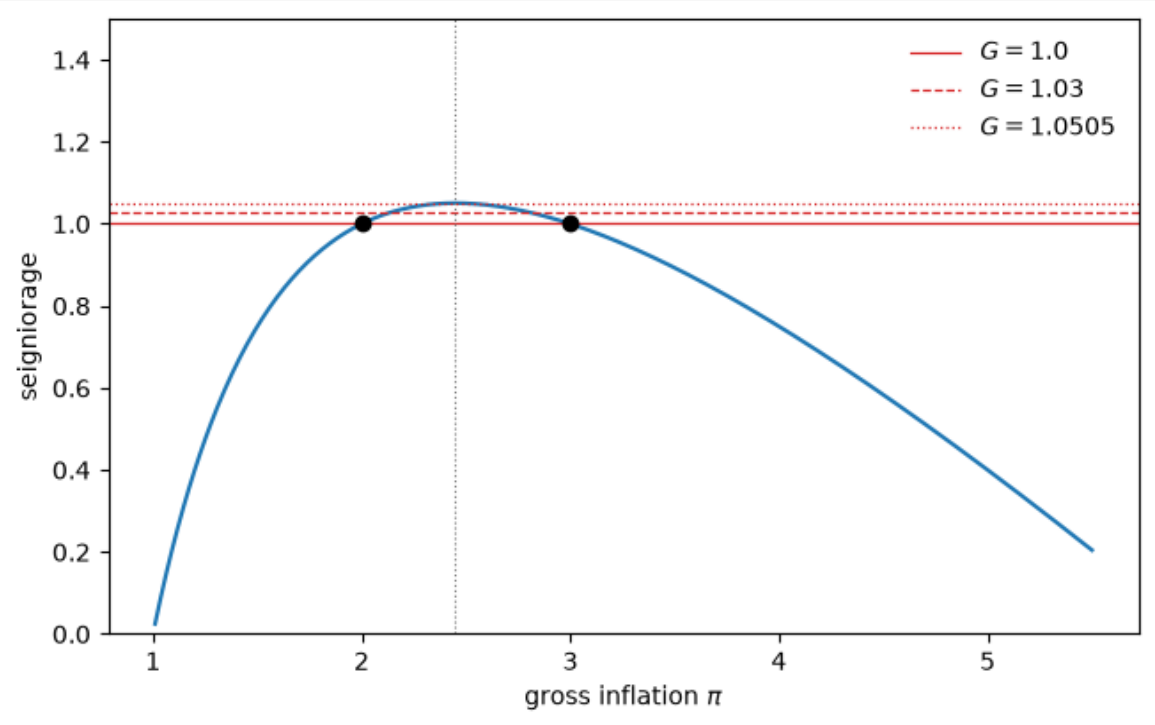
Here is the inflation-tax Laffer curve for Economy 7C.

```
def seigniorage(π):
    return (w1_d - w2_d*π)/2 * (1 - 1/π)

π_grid = np.linspace(1.01, 5.5, 500)
peak = π_grid[np.argmax(seigniorage(π_grid))]

fig, ax = plt.subplots(figsize=(7.5, 4.5))
ax.plot(π_grid, seigniorage(π_grid), lw=1.6)
for G_try, style in [(1.0, '-'), (1.03, '--'), (1.0505, ':')]:
    ax.axhline(G_try, color='C3', lw=0.9, ls=style, label=f"$G = {G_try}$")
ax.plot([π_low, π_high], [seigniorage(π_low), seigniorage(π_high)], 'ko', ms=6)
ax.axvline(peak, color='gray', lw=0.8, ls=':')
ax.set_xlabel(r"gross inflation $\pi$")
ax.set_ylabel("seigniorage")
ax.set_ylim(0, 1.5)
ax.legend(frameon=False)
plt.show()

print(f"peak of the Laffer curve at π = {peak:.4f}, raising {seigniorage(peak):.4f}
↪")
print(f"the two steady states raise {seigniorage(π_low):.4f} and {seigniorage(π_
↪high):.4f}")
```




```
peak of the Laffer curve at π = 2.4497, raising 1.0505
the two steady states raise 1.0000 and 1.0000
```

The two stationary equilibria are the two inflation rates at which the inflation-tax Laffer curve crosses the required revenue G, one on each side of the peak.

The low one sits on the rising branch, so financing a larger deficit there requires *more* inflation: the classical comparative static.

The high one sits on the falling branch, where more inflation raises *less* revenue, so a larger deficit is financed by a *lower* stationary inflation rate, the anti-classical comparative static that makes the high equilibrium so awkward.

As G rises the horizontal line moves up and the two intersections converge.

Past the peak of the curve there is no stationary equilibrium at all: the deficit exceeds the maximum revenue the inflation tax can raise.

The table tracks the two roots as G rises.

```
rows = []
for G_try in (0.8, 1.0, 1.04, 1.05, 1.0505, 1.06):
    A1_try = (w1_d + w2_d - 2*G_try)/w2_d
    disc = A1_try**2 - 4*A2
    if disc < 0:
        rows.append([G_try, np.nan, np.nan, "none"])
    else:
        r = np.sort(np.roots([1, -A1_try, A2]))
        rows.append([G_try, r[0], r[1], "two"])

pd.DataFrame(rows, columns=["$G$", "low $\\pi$", "high $\\pi$",
                            "stationary equilibria"]).round(4)
```

	G	low π	high π	stationary equilibria
0	0.8000	1.5642	3.8358	two
1	1.0000	2.0000	3.0000	two
2	1.0400	2.2328	2.6872	two
3	1.0500	2.4000	2.5000	two
4	1.0505	2.4424	2.4566	two
5	1.0600	NaN	NaN	none

The two roots move toward each other as G rises and collide at $G \approx 1.0505$, after which the model has no stationary equilibrium at all.

That number is not a coincidence.

Setting the discriminant of (91.8)'s fixed-point equation to zero gives a double root at $\pi = \sqrt{A_2} = \sqrt{w_1/w_2}$, and the deficit at which it occurs is exactly the peak of the Laffer curve computed above.

```
G_max = (w1_d + w2_d - 2*np.sqrt(A2)*w2_d) / 2
print(f"roots collide at G       = {G_max:.6f},  where π = √(w1/w2) = {np.
 ↪sqrt(A2):.6f}")
print(f"peak of the Laffer curve = {seigniorage(np.sqrt(A2)):.6f}  at π = {np.
 ↪sqrt(A2):.6f}")
```

```
roots collide at G       = 1.050510,  where π = √(w1/w2) = 2.449490
peak of the Laffer curve = 1.050510  at π = 2.449490
```

The largest deficit the government can finance is the maximum of the inflation tax, and at that deficit the two stationary equilibria have merged into one.

Exercise 91.7.5

This exercise brings the bubbles of *The bubble* into Part 2's model.

Set $G = 0$, keeping $w_1 = 6$ and $w_2 = 1$.

1. Show that the stationary inflation rates are $\pi = 1$ and $\pi = w_1/w_2 = 6$, and identify them with equilibria of *A Peculiar Definition of Bounded Rationality*'s money model with a constant money supply.
2. Show that the T-map (91.9) is constant, and say what that implies for least squares learning.
3. Run `ls_dynamics` with `G=0` from several initial beliefs below 6, and compare with the rational expectations paths from the same starting points. Which bubble does a society of least squares learners select?

Solution

With $G = 0$, $A_1 = 1 + w_1/w_2$ and $A_2 = w_1/w_2$, so the fixed-point equation $\pi^2 - A_1\pi + A_2 = 0$ factors as $(\pi - 1)(\pi - w_1/w_2) = 0$.

With no deficit the money supply is constant, so this is *A Peculiar Definition of Bounded Rationality*'s money model with $\mu = 1$.

There, $\pi = 1$ is the quantity-theory equilibrium $c = 0$, in which the price level is proportional to the (constant) money supply, and $\pi = w_1/w_2$ is the inflation rate that every bubble path with $c > 0$ approaches as real balances shrink to zero.

At $\pi = w_1/w_2$ saving is $(w_1 - w_2\pi)/2 = 0$, so it is the limit of a demonetizing economy, not a monetary steady state.

With $G = 0$ the T-map is $T(\beta) = s(\beta)/s(\beta) = 1$ for every belief that leaves saving positive: if agents expected *any* constant inflation rate, they would save a constant amount, and a constant stock of money chasing constant real balances means a constant price level.

So the differential equation for beliefs is $d\beta/d\tau = 1 - \beta$, and least squares beliefs should converge to 1 from anywhere below 6.

```
print("T-map at G = 0:", [T_map(β, G=0.0) for β in (0.5, 2.0, 4.0, 5.9)])

def re_path_G0(π0, T=25):
    "Rational expectations path π_{t+1} = A1 - A2/π_t with G = 0."
    A1_0, A2_0 = (w1_d + w2_d) / w2_d, w1_d / w2_d
    path = [π0]
    for _ in range(T):
        nxt = A1_0 - A2_0/path[-1]
        path.append(nxt if path[-1] > 0 and nxt > 0 else np.nan)
    return np.array(path)

dates = (10, 100, 1_000, 20_000)
rows = []
for π0 in (0.5, 1.5, 3.0, 5.0, 5.9):
    b, infl, ok = ls_dynamics(π0, T=dates[-1], G=0.0)
    rows.append([π0, "yes" if ok else "no", *b[[d - 1 for d in dates]],
                 re_path_G0(π0)[-1]])

pd.DataFrame(rows, columns=["initial belief $\\pi_0$", "survived",
                            *[f"belief at t = {d:,}" for d in dates],
                            "RE inflation at t = 25"]).set_index("initial belief $\
↪\pi_0$").round(4)
```

```
T-map at G = 0: [1.0, 1.0, 1.0, 1.0]
```

```
                        survived  belief at t = 10  belief at t = 100  \
initial belief $\pi_0$
0.5                          yes            0.8586             0.9843
1.5                          yes            1.0155             1.0015
3.0                          yes            0.9628             0.9954
5.0                          yes            0.7225             0.8876
5.9                          yes            0.5871             0.5932

                        belief at t = 1,000  belief at t = 20,000  \
initial belief $\pi_0$
0.5                                  0.9984                0.9999
1.5                                  1.0002                1.0000
3.0                                  0.9995                1.0000
5.0                                  0.9846                0.9992
5.9                                  0.6466                0.9117

                        RE inflation at t = 25
initial belief $\pi_0$
0.5                                        NaN
1.5                                        6.0
3.0                                        6.0
5.0                                        6.0
5.9                                        6.0
```

From every starting point below 6, least squares beliefs head to $\pi = 1$: the learners select $c = 0$, and learning kills the bubble.

Convergence is slow, though, and slowest from high initial beliefs.

Without inflation the price level stops growing, so least squares weights observations roughly equally and its gain falls like $1/t$, unlike the large effective gain we found in Part 2.

And from a belief near 6 the first realized inflation rate is exactly 1, which pulls the belief down at once, and the second is a deep deflation: saving jumps up from almost nothing, the price level collapses, and the early observations, taken at high price levels, dominate the regression for thousands of periods.

The rational expectations paths from the same starting points above 1 do the opposite: they are bubble paths, and inflation rises toward $w_1/w_2 = 6$.

(From a starting point below 1 the rational expectations path has no positive price level.)

The sign result of Exercise 91.7.6 does not apply at $\pi = 6$, where saving is zero and the T-map is undefined.

Exercise 91.7.6

This exercise proves that the stability reversal of *The stability reversal* is a theorem about the model, not a property of Economy 7C's parameters.

Let $F(x) = x^2 - A_1 x + A_2$, so that the stationary inflation rates are the roots of F, and let $s(\beta) = (w_1 - w_2\beta)/2$.

1. Show that $g(\pi) - \pi = -F(\pi)/\pi$ and $T(\beta) - \beta = \dfrac{w_2}{2}\,\dfrac{F(\beta)}{s(\beta) - G}$.
2. Deduce that at every steady state with positive saving, $\operatorname{sign}(T' - 1) = -\operatorname{sign}(g' - 1)$. Conclude that whenever there are two such steady states, the lower one is E-stable and unstable under the rational expectations dynamics, and the higher one is the reverse.
3. Check the result numerically for randomly drawn (w_1, w_2, G) with two steady states.

Solution

For the first identity,

$$g(\pi) - \pi = A_1 - \frac{A_2}{\pi} - \pi = -\frac{\pi^2 - A_1\pi + A_2}{\pi} = -\frac{F(\pi)}{\pi}.$$

For the second,

$$T(\beta) - \beta = \frac{s - \beta(s-G)}{s-G} = \frac{s(1-\beta) + G\beta}{s-G},$$

and

$$2\big[s(1-\beta) + G\beta\big] = (w_1 - w_2\beta)(1-\beta) + 2G\beta = w_2\beta^2 - (w_1 + w_2 - 2G)\beta + w_1 = w_2 F(\beta).$$

At a steady state π^* we have $F(\pi^*) = 0$, so differentiating the two identities gives

$$g'(\pi^*) - 1 = -\frac{F'(\pi^*)}{\pi^*}, \qquad T'(\pi^*) - 1 = \frac{w_2}{2}\,\frac{F'(\pi^*)}{s(\pi^*) - G}.$$

If saving is positive at the steady state, then $\pi^* = s/(s-G) > 0$ forces $s - G > 0$, so both denominators are positive and

$$\operatorname{sign}(T' - 1) = \operatorname{sign} F'(\pi^*) = -\operatorname{sign}(g' - 1).$$

F is an upward-opening parabola, so $F' < 0$ at its smaller root and $F' > 0$ at its larger one.

At π_{low}, therefore, $T' < 1$ (E-stable) and $g' > 1$ (unstable under rational expectations, since $g' = A_2/\pi^2 > 0$); at π_{high}, $T' > 1$ and $g' < 1$.

When the two roots merge, at the largest financeable deficit of Exercise 91.7.4, $F' = 0$ and both slopes equal one.

Here is a numerical check, using $T'(\beta) = (w_2/2)\,G/(s(\beta) - G)^2$ and $g'(\pi) = A_2/\pi^2$.

```
rng = np.random.default_rng(0)
checked = 0
while checked < 1_000:
    w2_r = rng.uniform(0.5, 5)
    w1_r = w2_r * rng.uniform(1.1, 10)                # w1 > w2
    G_r = rng.uniform(0, (np.sqrt(w1_r) - np.sqrt(w2_r))**2 / 2)    # below the␣
 ↪largest deficit
    A1_r, A2_r = (w1_r + w2_r - 2*G_r) / w2_r, w1_r / w2_r
    for π_star in np.roots([1, -A1_r, A2_r]).real:
        s = (w1_r - w2_r*π_star) / 2
        T_slope = (w2_r/2) * G_r / (s - G_r)**2
        g_slope = A2_r / π_star**2
        assert s > G_r and np.sign(T_slope - 1) == -np.sign(g_slope - 1)
    checked += 1
print(f"sign(T' - 1) = -sign(g' - 1) at both steady states in all {checked}␣
 ↪economies")
```

```
sign(T' - 1) = -sign(g' - 1) at both steady states in all 1000 economies
```

Exercise 91.7.7

Chung [1990] estimated the government's learning rule itself: he treated the parameters of the government's Kalman filter as free parameters of a likelihood function for post-war U.S. data on inflation and unemployment.

Here is a stripped-down version of his problem.

Simulate 16,000 periods of (π_t, U_t) from `constant_gain` with gain 0.02.

Suppose an econometrician knows everything about the model except the gain: θ, U^*, σ_η, σ_u, and the government's initial beliefs and moment matrix.

1. Explain why, for any candidate gain γ, the econometrician can reconstruct the government's beliefs and target inflation $g_{t-1}(\gamma)$ from the data.
2. Using the first T observations, compute the Gaussian log likelihood of inflation, $\ell_T(\gamma) = -\sum_{t \le T} (\pi_t - g_{t-1}(\gamma))^2 / (2\sigma_\eta^2)$ up to a constant, on a grid of gains around 0.02, for $T = 250$, 1,000, 4,000 and 16,000. How does the shape of ℓ_T change as T grows?
3. Why does information about the gain keep accumulating?

Solution

The government's beliefs are a deterministic function of the data: given the initial beliefs and moment matrix, the recursion that updates (α_t, M_t) uses only the observed (π_t, U_t) and the gain.

So for each candidate gain the econometrician can rerun the government's recursion on the data and read off the targets $g_{t-1}(\gamma)$ from (91.13).

The function below does this for a whole grid of gains at once.

```
def reconstruct_targets(π, U, gains):
    """
    Target inflation g_{t-1}(γ) that a government with gain γ would have chosen,
    reconstructed from data on (π_t, U_t), for each candidate γ in `gains`.
    """
    gains = np.asarray(gains)
    K = len(gains)
    α = np.tile(α_consistent, (K, 1))
    M = np.tile(np.eye(2), (K, 1, 1))
    targets = np.empty((len(π), K))
    for t in range(len(π)):
        targets[t] = -α[:, 0] * α[:, 1] / (1 + α[:, 1]**2)
        x = np.array([1.0, π[t]])
        M = M + gains[:, None, None] * (np.outer(x, x) - M)
        err = U[t] - α @ x
        step = np.linalg.solve(M, (err[:, None] * x)[:, :, None])[:, :, 0]
        α = α + gains[:, None] * step
    return targets

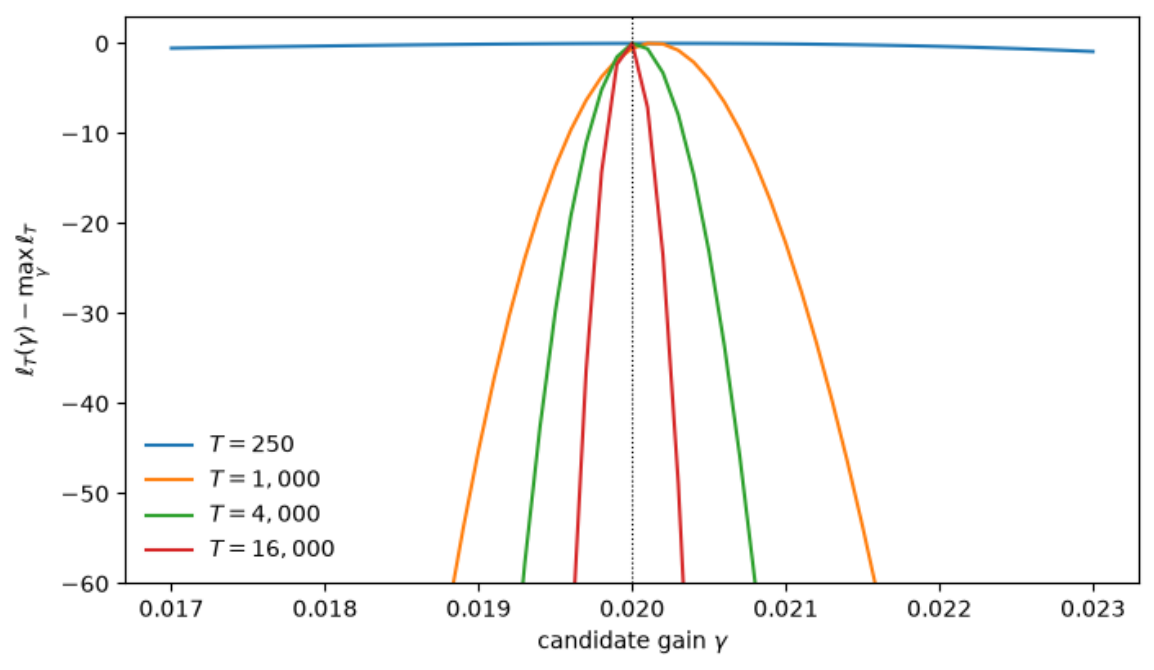
true_gain = 0.02
g_true, π_data, U_data = constant_gain(T=16_000, gain=true_gain, seed=1)
gains = np.round(np.linspace(0.017, 0.023, 61), 6)    # includes 0.02 exactly
targets = reconstruct_targets(π_data, U_data, gains)
print(f"largest error in reconstructed targets at the true gain: "
      f"{np.max(np.abs(targets[:, 30] - g_true)):.1e}")
print(f"first period with target inflation below 1: {np.argmax(g_true < 1)}")
```

```
largest error in reconstructed targets at the true gain: 8.9e-16
first period with target inflation below 1: 307
```

Now the log likelihood, normalized to zero at its maximum, for four sample sizes.

```
fig, ax = plt.subplots(figsize=(8, 4.5))
rows = []
for T in (250, 1_000, 4_000, 16_000):
    ssr = np.sum((π_data[:T, None] - targets[:T])**2, axis=0)
    loglik = -ssr / (2 * σ_η**2)
    loglik -= loglik.max()
    ax.plot(gains, loglik, label=f"$T = {T:,}$")
    rows.append([T, gains[np.argmax(loglik)], loglik[20], loglik[40]])
ax.axvline(true_gain, color='k', lw=0.8, ls=':')
ax.set_ylim(-60, 3)
ax.set_xlabel(r"candidate gain $\gamma$")
ax.set_ylabel(r"$\ell_T(\gamma) - \max_\gamma \ell_T$")
ax.legend(frameon=False)
plt.show()

pd.DataFrame(rows, columns=["T", "maximizing gain",
                            f"drop at γ = {gains[20]:.3f}",
                            f"drop at γ = {gains[40]:.3f}"]).set_index("T").
↪round(4)
```



T	maximizing gain	drop at γ = 0.019	drop at γ = 0.021
250	0.0202	-0.0980	-0.0601
1000	0.0201	-45.1487	-22.2241
4000	0.0200	-118.9857	-93.1558
16000	0.0200	-488.2670	-462.3829

At the true gain the reconstructed targets equal the actual ones, so the residuals are exactly the surprises η_t.

Any other gain implies a different history of beliefs and targets, and the log likelihood penalizes it.

With 250 observations the likelihood is nearly flat.

The sample ends before the first escape, and while beliefs sit near the self-confirming equilibrium they hardly move

whatever the gain, so the data say little about it.

Once escapes begin, the likelihood becomes sharply peaked at the true gain, and the drop at a fixed wrong gain keeps growing with T, roughly in proportion to T in the later samples.

Information about the gain keeps accumulating because a constant gain never stops shaping beliefs: every period the government responds to its newest observations with the same weight, so every period carries fresh evidence about that weight, and most of it arrives during escapes, when beliefs move a lot.

With a decreasing gain the influence of the gain's parameters on beliefs would die away, and so would the information about them; *1993 Prospects for Bounded Rationality in Macroeconomics* takes up this contrast.

Chung's problem was harder than this one: he also estimated θ, U^*, the shock variances and the government's prior, from a few decades of quarterly data rather than thousands of periods.

CHAPTER

NINETYTWO

LEARNING, APPROXIMATION, AND EQUILIBRIUM COMPUTATION

Contents

- *Learning, Approximation, and Equilibrium Computation*
  - *Overview*
  - *A continuous-state functional equation*
  - *A benchmark solution*
  - *Parameterized expectations*
  - *Learning is equilibrium computation*
  - *Learning without a functional form*
  - *Concluding remarks*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

92.1 Overview

In *Adaptive Agents in an Overlapping Generations Monetary Economy* a system of adaptive agents found a rational expectations equilibrium it was never told about.

The agents there learned a *separate* saving rate for each level of the government deficit — a **non-parametric** rule, one number per state.

That works when there are two states.

It runs into two walls otherwise, both noted in Sargent [1993]:

- a state that occurs rarely is learned about slowly, because its observations arrive slowly;
- when the number of states is large — and especially when the state is *continuous* — one parameter per state is hopeless.

The econometrician's response is to impose a **parametric** form on the decision rule, $s = f(G, \theta)$, with θ of low dimension, and use every observation to estimate θ.

This lecture works through what that buys and what it costs.

The costs and benefits are two sides of one idea.

A parametric family can be learned quickly because every observation informs every state.

But a learning scheme confined to a family of functions can converge to a rational expectations equilibrium *only if some member of the family supports one*.

Otherwise the best it can reach is an **approximate equilibrium**.

Working this out surfaces a theme that runs through the whole bounded-rationality program: algorithms that model learning and algorithms that compute equilibria resemble each other.

Marcet's method of parameterized expectations [den Haan and Marcet, 1990, Marcet and Marshall, 1992] is a standard tool for *computing* rational expectations equilibria.

Written recursively, it becomes a model of adaptive agents learning in real time.

An equilibrium computation is a centralized learning algorithm run by the modeller; a learning economy is a decentralized equilibrium computation run by the agents.

The two are different recursions, but they share their fixed points and, through a differential equation that both of them track, their stability condition.

The plan:

1. Put a deficit that takes many values — a fine discretization of a continuous process — into the overlapping generations model of *Adaptive Agents in an Overlapping Generations Monetary Economy*, turning the equilibrium condition into a **functional equation** in the saving rule $f(G)$.
2. Compute a benchmark solution, so we have ground truth to measure against.
3. Solve it by **parameterized expectations** with families of increasing richness, and watch approximate equilibria give way to the exact one.
4. Show that the recursive, real-time version of the same algorithm **is** a learning economy, and that one eigenvalue of the T map sets how fast it learns.
5. Learn the saving rule **non-parametrically** with a kernel estimator, and weigh the trade-off against the parametric route.

Let's start with some imports.

```
import warnings
import numpy as np
import pandas as pd
import quantecon as qe
import matplotlib.pyplot as plt
```

92.2 A continuous-state functional equation

We keep the overlapping generations monetary economy of *Adaptive Agents in an Overlapping Generations Monetary Economy* — two-period lived agents, log utility, endowments w_1 when young and w_2 when old, currency the only store of value, and a government financing a deficit by printing money.

We also keep the calibration of Part 1 of that lecture: $w_1 = 20$, $w_2 = 10$ and $N = 1$.

The one change: the deficit G_t now follows a **continuous** Markov process, with transition kernel $F(G', G) = \text{Prob}\{G_{t+1} \leq G' \mid G_t = G\}$.

We look for a stationary equilibrium in which saving is a *function* of the current deficit, $s_t = f(G_t)$.

The household's first-order condition, evaluated at $s_t = f(G_t)$, is

$$u'(w_1 - f(G_t)) = \mathbb{E}_t[u'(w_2 + f(G_t)R_t)\, R_t], \qquad \mathbb{E}_t(\cdot) = \mathbb{E}(\cdot \mid G_t),$$

and the government budget constraint with market clearing gives the return on currency in terms of the saving rule, exactly as before (with $N = 1$),

$$R_t = \frac{f(G_{t+1}) - G_{t+1}}{f(G_t)}.$$

Substituting R_t, and using $f(G_t)R_t = f(G_{t+1}) - G_{t+1}$ for old-age consumption, turns the first-order condition into a **functional equation** in f:

$$u'(w_1 - f(G_t)) = \mathbb{E}_t\left[u'(w_2 + f(G_{t+1}) - G_{t+1}) \cdot \frac{f(G_{t+1}) - G_{t+1}}{f(G_t)}\right]. \tag{92.1}$$

Because $f(G_t)$ is known at t, it pulls out of the expectation, whatever the utility function.

Writing

$$\psi(G_t) \equiv \mathbb{E}_t[u'(w_2 + f(G_{t+1}) - G_{t+1})\,(f(G_{t+1}) - G_{t+1})], \tag{92.2}$$

equation (92.1) becomes $u'(w_1 - f(G_t)) = \psi(G_t)/f(G_t)$.

Under log utility, $u'(c) = 1/c$, and this can be solved for saving in closed form:

$$\frac{1}{w_1 - f(G_t)} = \frac{\psi(G_t)}{f(G_t)}, \qquad \text{i.e.} \qquad f(G_t) = \frac{w_1\,\psi(G_t)}{1 + \psi(G_t)}.$$

So the whole problem reduces to finding the **conditional expectation** $\psi(G)$: once we have it, the saving rule follows in closed form.

That observation is the hinge of the entire lecture.

Every method below — parametric or not — is a way of estimating the one object $\psi(G)$.

Under log utility the random variable inside (92.2) is $y = (s' - G')/(w_2 + s' - G')$, where s' and G' are next period's saving and deficit.

We will call it the **regressand**, for reasons that will soon be clear.

```
w1, w2 = 20.0, 10.0

def s_from_psi(ψ):
    "Recover the saving rule from the conditional expectation ψ."
    return w1 * ψ / (1 + ψ)

def regressand(s, G):
    "u'(c_2) (s - G) under log utility, at saving s and deficit G."
    return (s - G) / (w2 + s - G)
```

92.2.1 The deficit process

We take $\log G_t$ to follow an AR(1), $\log G_{t+1} = (1-\rho)\log\bar{G} + \rho \log G_t + \sigma\varepsilon_{t+1}$.

To compute with it we replace it by a fine Markov chain, discretizing it on 201 grid points with the method of Tauchen [1986].

From here on that chain *is* the economy: the benchmark below is computed on it, and every learning economy is simulated from it.

Using one chain for both keeps the comparison honest.

If the learners lived in the AR(1) while the benchmark was computed on the chain, even a learner that solved its own economy perfectly would miss the benchmark, and the miss would be blamed on the learner.

With 201 states the chain is a close stand-in for the continuous process.

It is also exactly the case that defeats a state-by-state learner like those of *Adaptive Agents in an Overlapping Generations Monetary Economy*: such a learner would need 201 separate saving rates, and the rarest states are visited only about once every two thousand periods.

```
ρ, σ, G_bar = 0.7, 0.30, 0.6
n_states = 201

mc = qe.tauchen(n_states, ρ, σ, mu=np.log(G_bar) * (1 - ρ), n_std=2.5)
G, P = np.exp(mc.state_values), mc.P
ergodic = mc.stationary_distributions[0]      # where the deficit spends its time

def simulate_deficit(T, seed=0, P=P):
    "State indices of a simulated path of the deficit chain, started from its ergodic␣
 ↪distribution."
    rng = np.random.default_rng(seed)
    chain = qe.MarkovChain(P)
    init = rng.choice(len(P), p=chain.stationary_distributions[0])
    return chain.simulate_indices(T, init=init, random_state=rng)

G_perm = (np.sqrt(w1) - np.sqrt(w2))**2 / 2   # largest deficit that can be financed␣
 ↪forever

print(f"deficit grid: {n_states} points from {G.min():.3f} to {G.max():.3f}")
print(f"mean deficit E[G] under the ergodic distribution: {ergodic @ G:.3f}")
print(f"probability of the least-visited state: {ergodic.min():.1e}")
print(f"largest constant deficit with a monetary steady state: {G_perm:.3f}")
print(f"top of the grid relative to it: {G.max() / G_perm:.2f}")
print(f"share of time the deficit is above it: {ergodic[G > G_perm].sum():.2f}")
```

```
deficit grid: 201 points from 0.210 to 1.715
mean deficit E[G] under the ergodic distribution: 0.653
probability of the least-visited state: 4.4e-04
largest constant deficit with a monetary steady state: 0.858
top of the grid relative to it: 2.00
share of time the deficit is above it: 0.19
```

How large a deficit can this economy finance?

With a deficit held constant at G, a monetary steady state solves a quadratic (Exercise 92.8.2) that has real roots only if $G \leq (\sqrt{w_1} - \sqrt{w_2})^2/2 \approx 0.858$.

The top of our grid is twice that, and the deficit spends about a fifth of its time above it.

A monetary equilibrium nevertheless exists at every state, because high deficits are transitory.

With $\rho = 0.7$, a young agent in a high-deficit state expects the deficit, and with it the inflation tax on his saving, to fall back toward its mean.

92.3 A benchmark solution

To measure any learning scheme, we first need the true $f(G)$.

On the grid, (92.1) is a system of 201 equations, which we solve by **iterating the perceived-to-actual map**, precisely the T map of *A Peculiar Definition of Bounded Rationality*.

Guess a saving rule; use it on the right-hand side to compute the actual saving each state would call forth; repeat until the two agree.

On the chain the right-hand side is a matrix–vector product, $\psi = Py$ with $y_j = (s_j - G_j)/(w_2 + s_j - G_j)$, and the closed form then delivers the new saving rule.

This is a form of the time iteration of *Optimal Savings IV: Time Iteration*, and the damping weight on the new guess plays the role of the relaxation parameter λ of *A Peculiar Definition of Bounded Rationality*.

```
def T_map(s):
    "Actual saving rule called forth by a perceived rule s (one entry per state)."
    return s_from_psi(P @ regressand(s, G))

def solve_benchmark(damp=0.5, tol=1e-12, max_iter=5000):
    "Fixed point of the perceived-to-actual saving map on the deficit grid."
    s = np.full(n_states, (w1 - w2) / 2)
    for it in range(1, max_iter + 1):
        s_new = T_map(s)
        if np.max(np.abs(s_new - s)) < tol:
            return s_new, it
        s = damp * s_new + (1 - damp) * s
    warnings.warn("the benchmark iteration did not converge")
    return s, it


s_benchmark, iters = solve_benchmark()
print(f"converged in {iters} iterations")


def foc_residual(s_rule):
    "Max |first-order-condition error| of a saving rule under the true kernel P."
    return np.max(np.abs(1 / (w1 - s_rule) - (P @ regressand(s_rule, G)) / s_rule))


print(f"benchmark FOC residual: {foc_residual(s_benchmark):.2e}")
print(f"saving rule f(G) runs from {s_benchmark.max():.3f} (low deficit) "
      f"to {s_benchmark.min():.3f} (high deficit)")
```

```
converged in 223 iterations
benchmark FOC residual: 1.70e-14
saving rule f(G) runs from 4.471 (low deficit) to 2.840 (high deficit)
```

Here is the rule.

```
fig, ax = plt.subplots(figsize=(6.5, 4))
ax.plot(G, s_benchmark, 'k-', lw=2)
ax.set_xlabel("deficit $G$")
ax.set_ylabel("saving $f(G)$")
plt.show()
```

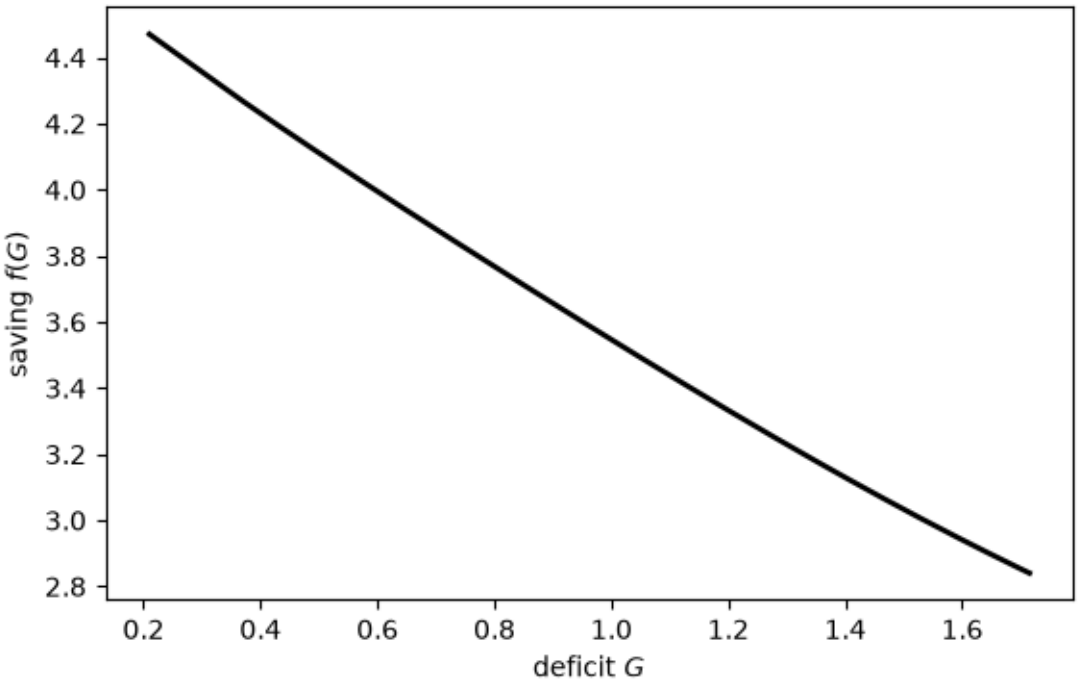


Fig. 92.1: Benchmark equilibrium saving rule

Saving falls as the deficit rises: a larger deficit means faster money creation, a poorer return on currency, and less saving.

The rule is gently curved, which will matter in a moment.

This $f(G)$ is the object no adaptive agent gets to see.

Everything below tries to recover it.

92.3.1 Which equilibrium?

Sargent [1993] notes that systems like (92.1) generally have more than one solution.

Our iteration started from $s = 5$, the saving rate of an economy with no deficit, and found a high-saving, low-inflation solution.

It would find the same solution from any nearby start, because the solution is **stable under the T map**: every eigenvalue of the map's Jacobian there is less than one.

```
def jacobian(f, x, eps=1e-6):
    "Numerical Jacobian of the map f at the point x."
    return np.column_stack([(f(x + eps * e) - f(x - eps * e)) / (2 * eps)
                            for e in np.eye(len(x))])

eig_benchmark = np.linalg.eigvals(jacobian(T_map, s_benchmark))
print(f"largest |imaginary part|: {np.abs(eig_benchmark.imag).max():.1e}")
print("largest eigenvalues of DT at the benchmark:",
      np.sort(eig_benchmark.real)[::-1][:4].round(3))
```

```
largest |imaginary part|: 3.0e-09
largest eigenvalues of DT at the benchmark: [0.773 0.542 0.353 0.208]
```

That is the E-stability condition of *One recursion, many learners*: the condition under which learners whose beliefs are pulled toward the T map's output settle on an equilibrium.

The economy with a constant deficit makes the point sharply.

It has two steady states, and at the high-inflation one the slope of the T map exceeds one (Exercise 92.8.2), so neither the iteration nor a learner can settle there.

That is the computational face of the stability reversal in *The stability reversal*, where least squares learners select the low-inflation steady state.

The largest eigenvalue printed above, about 0.77, will come back: it sets the speed of the algorithms below.

92.4 Parameterized expectations

Marcet's **method of parameterized expectations** [den Haan and Marcet, 1990, Marcet and Marshall, 1992] attacks (92.1) by putting a parametric form on the conditional expectation $\psi(G)$ from (92.2),

$$\psi(G) \approx \psi(G, \theta) = \phi(G)^\top \theta,$$

where $\phi(G)$ is a vector of basis functions and θ a short vector of coefficients.

The algorithm is a fixed-point iteration on θ:

1. Given θ, the saving rule is $f(G) = w_1 \psi(G, \theta)/(1 + \psi(G, \theta))$.
2. Simulate a long path of the deficit and compute the implied saving each period.
3. Form the *realized* value of the object inside (92.2), $y_t = u'(w_2 + s_{t+1} - G_{t+1})(s_{t+1} - G_{t+1})$, and **regress** it on $\phi(G_t)$ to get a new θ.
4. Iterate to convergence.

The regression in step 3 is doing the work of the conditional expectation: least squares projects the realized y_t onto functions of the current state, which is exactly what $\mathbb{E}_t[\cdot \mid G_t]$ is.

We use monomials in a rescaled $\log G$ as the basis, and vary the degree.

Every algorithm in this lecture starts from the same initial belief, $\psi = 1/3$, under which agents save $s = 5$, as they would if there were no deficit.

The functions take the deficit levels, the transition matrix, and the basis evaluated at each state as arguments, so that they can be reused on other economies.

Because the regressors do not change from one pass to the next, we evaluate the basis on the grid, and invert the regressor matrix, only once.

```
def basis(G_vals, degree):
    "Monomials in log-deficit rescaled to [-1, 1]."
    x = 2 * (np.log(G_vals) - np.log(G.min())) / (np.log(G.max()) - np.log(G.min())) -
↪ 1
    return np.vstack([x**k for k in range(degree + 1)]).T

ψ_init = 1 / 3                    # the belief under which agents save s = 5

def θ_init(k):
    "Initial coefficients: a constant ψ = ψ_init."
    θ = np.zeros(k)
    θ[0] = ψ_init
    return θ
```

(continues on next page)

(continued from previous page)

```
def rule(Φ, θ, G=G):
    "Saving rule at each state implied by coefficients θ, kept feasible."
    ψ = np.clip(Φ @ θ, 0.05, 50)
    return np.clip(s_from_psi(ψ), G + 0.05, w1 - 0.05)

def errors(s_rule):
    "Sup-norm and ergodic-weighted RMS distance from the benchmark."
    d = s_rule - s_benchmark
    return np.max(np.abs(d)), np.sqrt(ergodic @ d**2)

def parameterized_expectations(Φ, G=G, P=P, T=150_000, seed=0,
                               damp=0.6, tol=1e-11, max_iter=1000):
    """
    Batch parameterized expectations.  Φ holds the basis functions evaluated
    at each deficit state, one row per state.  Returns the implied saving rule
    at each state, the coefficients and the number of passes.
    """
    idx = simulate_deficit(T + 1, seed, P)
    X = Φ[idx[:-1]]                          # regressors ψ(G_t) along the history
    X_pinv = np.linalg.pinv(X)               # they never change, so invert once
    θ = θ_init(Φ.shape[1])
    for it in range(1, max_iter + 1):
        y = regressand(rule(Φ, θ, G), G)[idx[1:]]     # realized y_t
        θ_new = X_pinv @ y                            # regress y_t on ψ(G_t)
        if np.max(np.abs(θ_new - θ)) < tol:
            return rule(Φ, θ_new, G), θ_new, it
        θ = (1 - damp) * θ + damp * θ_new
    warnings.warn("parameterized expectations did not converge")
    return rule(Φ, θ, G), θ, it
```

92.4.1 Approximate versus exact equilibria

Now run it for families of increasing richness: a constant ψ (saving independent of the deficit), a linear one, and a quadratic one.

```
pea = {deg: parameterized_expectations(basis(G, deg)) for deg in (0, 1, 2)}
rules = {deg: pea[deg][0] for deg in pea}

fig, ax = plt.subplots(figsize=(7, 4.5))
ax.plot(G, s_benchmark, 'k-', lw=2.2, label="benchmark $f(G)$")
labels = {0: "degree 0 (constant)", 1: "degree 1 (linear)", 2: "degree 2 (quadratic)"}
for deg, colour in [(0, 'C0'), (1, 'C1'), (2, 'C2')]:
    ax.plot(G, rules[deg], '--', color=colour, lw=1.4, label=labels[deg])
ax.set_xlabel("deficit $G$")
ax.set_ylabel("saving $s$")
ax.legend(frameon=False)
plt.show()
```

The **constant** family cannot represent a saving rule that depends on the deficit at all, so the best it can do is a flat line through the middle of the data.

That flat line *is* an approximate equilibrium — a fixed point of the learning scheme — but it satisfies the first-order condition only on average, not state by state.

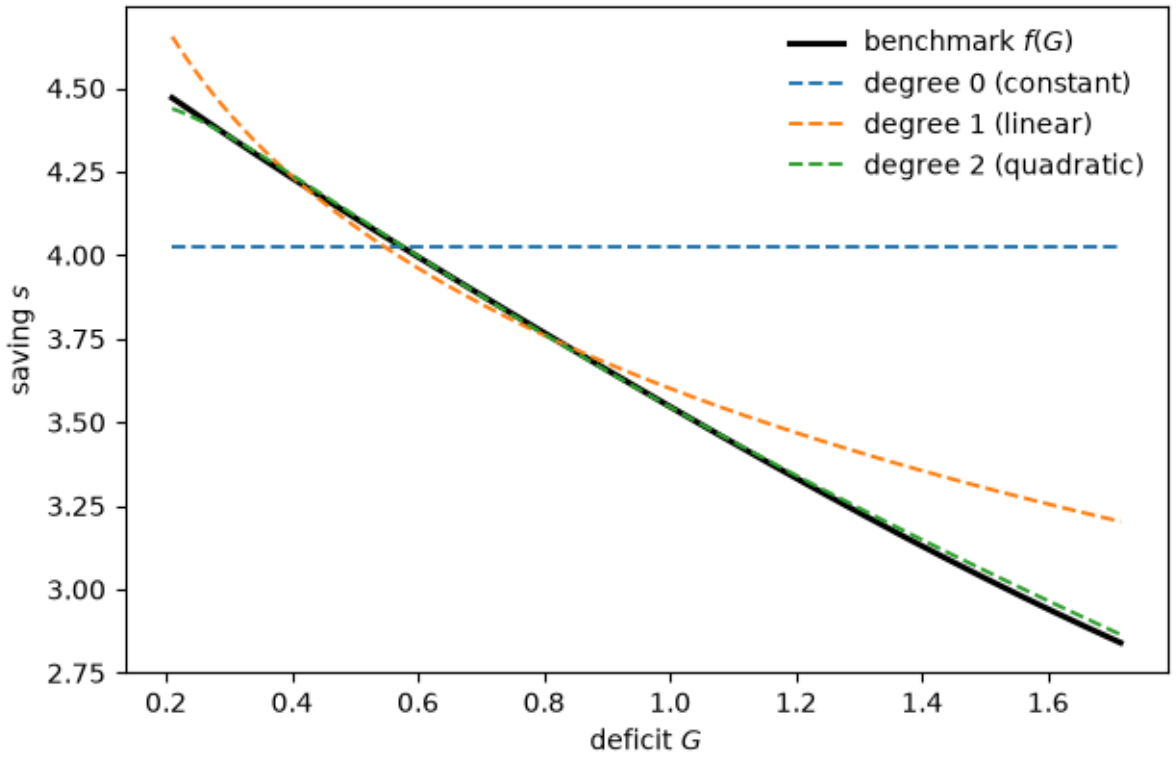


Fig. 92.2: Parameterized expectations rules of increasing degree

In the language of *Optimal Misspecified Beliefs* it is an equilibrium under forecast misspecification: the functional form is wrong, but the coefficient within it is the best that the data the agents themselves generate allow.

The **linear** family does much better, but it misses at the extremes: its rule lies above the benchmark at both ends of the deficit range and below it in the middle.

The **quadratic** family is almost indistinguishable from the benchmark: ψ is very nearly a quadratic function of $\log G$, the variable our basis is built on, so a three-parameter rule essentially nails it.

We can quantify "how close" two ways: by the distance to the benchmark, and by the residual in the first-order condition under the true kernel, the honest measure of how far an approximate equilibrium is from an exact one.

```
rows = []
for deg in (0, 1, 2):
    s_hat, _, passes = pea[deg]
    rows.append([labels[deg], *errors(s_hat), foc_residual(s_hat), passes])

pd.DataFrame(rows, columns=["family", "sup $|s - f|$", "ergodic RMS",
                            "FOC residual", "passes"]).set_index("family").round(4)
```

family	sup $\|s - f\|$	ergodic RMS	FOC residual	passes
degree 0 (constant)	1.1825	0.3212	0.0090	115
degree 1 (linear)	0.3626	0.0683	0.0039	135
degree 2 (quadratic)	0.0322	0.0073	0.0005	142

Every error column falls as the family grows richer.

The first-order-condition residual — the quantity that is exactly zero at a rational expectations equilibrium and positive

at an approximate one — drops by more than an order of magnitude from the constant family to the quadratic.

Marcet and Marshall [1992] give conditions under which this is no accident: as the family grows, the approximate equilibria computed by parameterized expectations converge to the rational expectations equilibrium.

The last column counts the passes the algorithm needed to converge: well over a hundred, for a reason we will come to.

Note

Richer is not *always* better with finite data.

The approximation error keeps falling as the degree rises, but the sampling error in the regression does not, and beyond the quadratic it takes over.

The extra terms then chase sampling noise in the tails of the deficit distribution, where observations are scarce: the ergodic-weighted error levels off, and the sup-norm error can tick back up, by an amount that depends on the simulated history.

That is the same phenomenon as the slowly-learned rare state in *Adaptive Agents in an Overlapping Generations Monetary Economy*, seen from the approximation side: a scheme fits well where the data is, and the cost of misfitting a rarely-visited region is small in terms of expected utility.

Exercise 92.8.1 separates the two sources of error.

92.5 Learning is equilibrium computation

So far parameterized expectations is an algorithm *we* run to compute an equilibrium.

Now comes the point of the whole lecture.

Write the same algorithm **recursively** — updating θ once per period as a single new observation arrives, rather than re-regressing a whole simulated history — and it becomes a model of *adaptive agents learning in real time*.

The recursive form is ordinary recursive least squares:

$$\begin{aligned} M_{t+1} &= M_t + \gamma_t(\phi(G_t)\phi(G_t)^\top - M_t), \\ \theta_{t+1} &= \theta_t + \gamma_t M_{t+1}^{-1}\phi(G_t)(y_t - \phi(G_t)^\top\theta_t), \end{aligned} \tag{92.3}$$

where y_t is the same realized regressand as before, M_t estimates the second-moment matrix of the regressors, and the gain γ_t declines like $1/t$.

In the code below, $\gamma_t = 1/(t+1)$, and M starts at a small multiple of the identity so that it can be inverted from the first period on.

This is a stochastic-approximation recursion with a decreasing gain, the same family as the Robbins–Monro learners of *Adaptive Agents in an Overlapping Generations Monetary Economy*, with M in the role of their curvature estimate.

It differs from their scheme in two ways, not one.

- **What is learned.** The agents of *Adaptive Agents in an Overlapping Generations Monetary Economy* learn a saving rate directly, taking Newton steps on *realized lifetime utility*. An agent's lifetime utility is known only when he is old, so the population there is split into odd and even classes that take turns. The agents here instead learn a *conditional expectation*, ψ, and solve their own Euler equation given it — what Sargent [1993] calls learning via model formation, as opposed to learning the hard way. Each period's realized regressand is known one period later, as soon as the next generation has saved, so no odd/even split is needed.
- **How far an observation reaches.** θ indexes a *parametric* rule, so a single observation updates the saving rule at every state at once, not just at the state that was visited.

Here is the learning economy.

```
def online_pea(Φ, G=G, P=P, T=500_000, seed=0, ridge=1e-3, a=1.0):
    """
    Recursive-least-squares parameterized expectations: the learning economy.

    The young at t save according to the current θ.  Their saving reveals the
    date t-1 realized regressand, and θ and the moment matrix M are updated
    once, with gain a / (t + a).  Returns the path of θ.
    """
    idx = simulate_deficit(T, seed, P)
    k = Φ.shape[1]
    θ, M = θ_init(k), ridge * np.eye(k)
    θ_path = np.empty((T, k))
    θ_path[0] = θ
    for t in range(1, T):
        g, x = G[idx[t]], Φ[idx[t - 1]]
        ψ = min(max(Φ[idx[t]] @ θ, 0.05), 50.0)
        s = min(max(s_from_psi(ψ), g + 0.05), w1 - 0.05)     # date-t saving
        y = regressand(s, g)                                  # date t-1's y, now known
        γ = a / (t + a)
        M += γ * (np.outer(x, x) - M)
        θ = θ + γ * np.linalg.solve(M, x * (y - x @ θ))
        θ_path[t] = θ
    return θ_path


Φ2 = basis(G, 2)
θ_path = online_pea(Φ2)
s_online = rule(Φ2, θ_path[-1])
print("online PEA after {:,} periods: sup error = {:.4f},  ergodic RMS = {:.4f}"
      .format(len(θ_path), *errors(s_online)))
```

```
online PEA after 500,000 periods: sup error = 0.0673,  ergodic RMS = 0.0224
```

How should we judge that?

The learner's target is not the benchmark but the approximate equilibrium of the quadratic family, the fixed point that the batch algorithm computes.

We can compute that target exactly, as a batch pass with infinitely much data would.

Let $T(\theta)$ be the coefficients of the population least squares projection of y_t on $\phi(G_t)$ when agents' beliefs are θ.

This is the T map of *A Peculiar Definition of Bounded Rationality* for beliefs confined to the family, and on the chain we can evaluate it exactly, using P for the conditional expectation and the ergodic distribution for the projection.

```
def T_pea(θ, Φ, G=G, P=P, π=ergodic):
    "Population version of one batch pass: project E[y_t | G_t] on the basis."
    Ey = P @ regressand(rule(Φ, θ, G), G)
    return np.linalg.solve(Φ.T @ (π[:, None] * Φ), Φ.T @ (π * Ey))

def approximate_equilibrium(Φ, G=G, P=P, π=ergodic, damp=0.6, tol=1e-13, max_iter=10_
↪000):
    "Fixed point θ* of the T map of the family with basis Φ."
    θ = θ_init(Φ.shape[1])
    for _ in range(max_iter):
        θ_new = T_pea(θ, Φ, G, P, π)
        if np.max(np.abs(θ_new - θ)) < tol:
```

(continues on next page)

(continued from previous page)

```
            return θ_new
        θ = (1 - damp) * θ + damp * θ_new
    warnings.warn("approximate_equilibrium did not converge")
    return θ


θ_star = approximate_equilibrium(Φ2)
s_star = rule(Φ2, θ_star)
eig_pea = np.sort(np.linalg.eigvals(jacobian(lambda θ: T_pea(θ, Φ2), θ_star)).
 ↪real)[::-1]
λ_max = eig_pea[0]

print("θ* (approximate equilibrium):", θ_star.round(4))
print("θ  (online, final)          :", θ_path[-1].round(4))
print("θ  (batch, T = 150,000)     :", pea[2][1].round(4))
print("approximate equilibrium: sup error = {:.4f},  ergodic RMS = {:.4f}"
      .format(*errors(s_star)))
print("eigenvalues of DT(θ*):", eig_pea.round(3))
```

```
θ* (approximate equilibrium): [ 0.2495 -0.059  -0.0232]
θ  (online, final)          : [ 0.251  -0.0579 -0.023 ]
θ  (batch, T = 150,000)     : [ 0.2498 -0.059  -0.0236]
approximate equilibrium: sup error = 0.0307,  ergodic RMS = 0.0070
eigenvalues of DT(θ*): [0.772 0.538 0.34 ]
```

The eigenvalues of $DT(\theta^*)$ are all below one, and they are close to the three largest eigenvalues of the benchmark's T map computed earlier: the quadratic family inherits the benchmark's slow directions.

Here is the learning economy at work.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))

t_grid = np.arange(1, len(θ_path) + 1)
for j in range(3):
    axes[0].plot(t_grid, θ_path[:, j], color=f'C{j}', lw=2, label=rf"$\theta_{j}$")
    axes[0].axhline(θ_star[j], color=f'C{j}', lw=0.8, ls=':')
axes[0].set_xscale('log')
axes[0].set_xlabel("$t$  (log scale)")
axes[0].set_ylabel(r"$\theta_t$")
axes[0].set_title(r"coefficients learned in real time ($\theta^*$ dotted)")
axes[0].legend(frameon=False, fontsize=9)

axes[1].plot(G, s_benchmark, 'k-', lw=2.2, label="benchmark $f(G)$")
axes[1].plot(G, s_star, 'C2:', lw=2, label="approximate equilibrium")
axes[1].plot(G, s_online, 'C3--', lw=1.5, label="online PEA")
axes[1].set_xlabel("deficit $G$")
axes[1].set_ylabel("saving $s$")
axes[1].set_title("saving rule after learning")
axes[1].legend(frameon=False)
plt.tight_layout()
plt.show()
```

The coefficients head for θ^*, but slowly: after half a million periods the online rule still lies visibly above the approximate equilibrium at high deficits, and it misses the benchmark by more than the batch rule does.

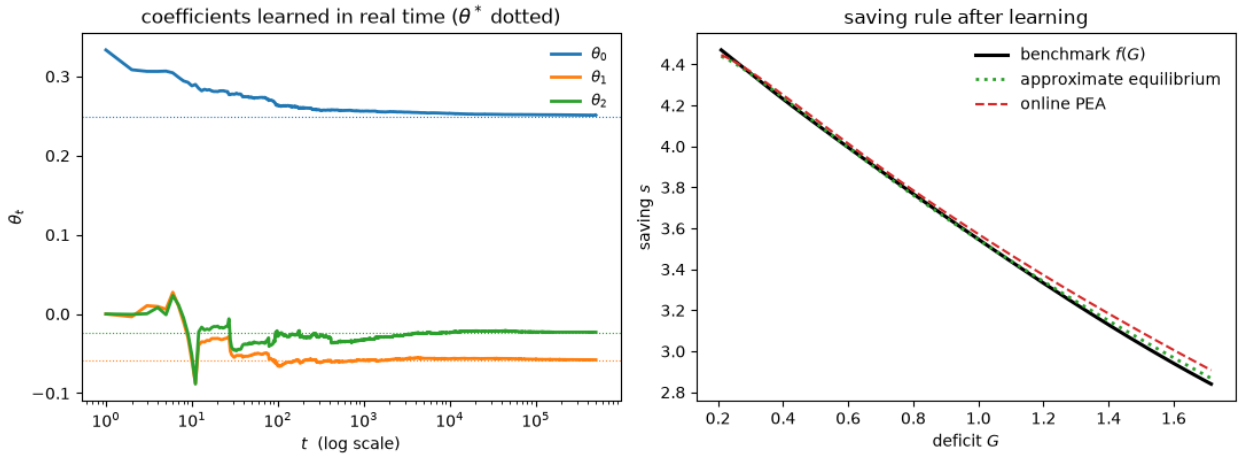


Fig. 92.3: Online learning of the saving rule

92.5.1 Same limit, different speeds

Why do the two algorithms share a limit, and why is the online one so slow?

A batch pass with damping $\delta = 0.6$ is $\theta_k = \theta_{k-1} + \delta\,[\hat{T}(\theta_{k-1}) - \theta_{k-1}]$, where $\hat{T}$ is the sample version of T.

That is the relaxation scheme of *A Peculiar Definition of Bounded Rationality*, with δ in the role of its relaxation parameter λ.

The online recursion (92.3) is a stochastic approximation algorithm.

The theory of such algorithms [Ljung, 1977, Marcet and Sargent, 1989, Woodford, 1990], later systematized and generalized by Evans and Honkapohja [2001], says that with a gain like $1/t$ the beliefs θ_t follow, on average, the ordinary differential equation

$$\frac{d\theta}{d\tau} = T(\theta) - \theta \tag{92.4}$$

in the rescaled time $\tau = \log t$, while M_t settles down to the second-moment matrix of the regressors.

One recursion, many learners sets out the general argument, and *Least Squares Learning in Self-Referential Models* works it through for linear models; Sargent [1993] illustrates it with the very economy of this lecture.

Three consequences follow.

- **Same fixed points.** The rest points of (92.4) are the fixed points of T, which the batch algorithm computes up to sampling error. Both schemes aim at the same approximate equilibrium θ^*.
- **Same stability condition.** Both converge to θ^* only if every eigenvalue of $DT(\theta^*)$ has real part below one, the E-stability condition. For the online algorithm this is the stability of (92.4). For the batch algorithm, the eigenvalues of its iteration matrix $I + \delta\,(DT - I)$ lie inside the unit circle only if this condition holds, and, for small enough δ, whenever it does.
- **Different speeds.** Near θ^*, a batch pass shrinks the error by the factor $1 - \delta(1 - \lambda_{\max})$, where $\lambda_{\max}$ is the largest eigenvalue of $DT(\theta^*)$. The differential equation shrinks it like $e^{-(1-\lambda_{\max})\tau}$, which is $t^{-(1-\lambda_{\max})}$ in calendar time.

```
δ = 0.6
print(f"largest eigenvalue of DT(θ*):          {λ_max:.3f}")
print(f"batch: error shrinks per pass by       {1 - δ * (1 - λ_max):.3f}")
print(f"online: error shrinks like             t^(-{1 - λ_max:.2f})")
print(f"500 times more data shrinks it by      {500**(1 - λ_max):.1f}")
```

```
largest eigenvalue of DT(θ*):          0.772
batch: error shrinks per pass by       0.863
online: error shrinks like             t^(-0.23)
500 times more data shrinks it by      4.1
```

So a batch pass cuts the error by only about 14 percent, which is why the batch algorithm needed well over a hundred passes.

The online error falls only like $t^{-0.23}$.

That is slower than the $t^{-1/2}$ rate at which sampling noise averages out, so it is the drift toward θ^*, not the noise, that dominates the error.

Sargent [1993], drawing on Benveniste *et al.* [1990], states the condition for the usual $\sqrt{t}$ rate: every eigenvalue of $DT - I$ must lie below $-1/2$, that is, $\lambda_{\max} < 1/2$, which fails here.

At that rate, multiplying the agents' experience by 500, from 1,000 to 500,000 periods, would cut their error only by a factor of about four.

Exercise 92.8.5 checks the rate.

The gap in speed measures the difference between an agent constrained to learn from experience as it arrives and a modeller free to re-use a whole simulated history at once.

But the limit is the same object.

Sargent's summary:

> Learning algorithms and equilibrium computation algorithms look like each other. Equilibrium computation algorithms often have interpretations as centralized learning algorithms whereby the model builder, acting in a role of 'social planner,' gropes for a set of pricing functions for markets and decision rules for agents that will satisfy all of the individual optimum conditions and market-clearing conditions. We have also seen that learning systems with boundedly rational agents sometimes have interpretations as decentralized equilibrium computation algorithms.

This is why, as Sargent [1993] remarks, it is understandable that Marcet arrived at parameterized expectations as a *computational* method by way of earlier work on the *dynamics of least squares learning* [Marcet and Sargent, 1989].

The two are not one recursion read in two directions.

They are two recursions that share their fixed points and, through (92.4), their stability condition; what separates them is speed, and what each asks of whoever runs it.

92.6 Learning without a functional form

The parametric route is fast but stakes everything on the family.

If no member of $\{f(\cdot, \theta)\}$ supports an equilibrium, the scheme converges to an approximate equilibrium and stops.

The **non-parametric** alternative imposes no functional form.

Following the recursive kernel estimators of Chen and White [1998], an agent estimates the conditional expectation $\psi(G)$ directly from a kernel-weighted average of past realized values, letting the data choose the shape.

The building block is the recursive kernel density estimator, itself a stochastic-approximation recursion,

$$\widehat{p}_t(x) = \widehat{p}_{t-1}(x) + \tfrac{1}{t}\left[h_t^{-d}\, K\!\left(\tfrac{x - x_t}{h_t}\right) - \widehat{p}_{t-1}(x)\right],$$

where $x \in \mathbb{R}^d$, K is a probability density such as the standard normal, and the bandwidth $h_t \searrow 0$.

The factor h_t^{-d} makes each term a density; without it, as the formula is printed in Sargent [1993], the estimate would shrink to zero along with the bandwidth.

Our agents need a conditional expectation rather than a density.

The ratio of two recursions of this kind gives the recursive **Nadaraya–Watson** estimator

$$\hat{\psi}_t(G) = \hat{\psi}_{t-1}(G) + \gamma_t(G)[y_t - \hat{\psi}_{t-1}(G)], \qquad \gamma_t(G) = \frac{K_{h_t}(\log G - \log G_t)}{\sum_{\tau \le t} K_{h_\tau}(\log G - \log G_\tau)}, \tag{92.5}$$

where $K_h(u) = h^{-1}K(u/h)$ and distance is measured in $\log G$.

It has the same shape as every other recursion in this lecture, now with a separate gain at each state: large when the deficit lands near G, and essentially zero when it lands far away.

Hold the bandwidth fixed at h and shrink it on a grid, and the gain becomes familiar.

As $h \to 0$, $K_h(\log G - \log G_t)/K_h(0)$ tends to one if $G_t = G$ and to zero otherwise.

The denominator then counts the visits to state G, and $\gamma_t(G) = 1/\tau_G$ when the deficit is at G, where τ_G is the number of visits so far, and zero otherwise.

That is the state-by-state $1/\tau$ clock of the Robbins–Monro learners of *Adaptive Agents in an Overlapping Generations Monetary Economy*, the book's "different clock for each event".

A positive bandwidth lets a visit to one state also advance the clocks of its neighbours: h measures how far the learner generalizes from the states it has seen to the states it has not.

What we implement below is the **batch** counterpart of (92.5): a Nadaraya–Watson estimate of $\psi(G)$ with a fixed bandwidth, formed as a kernel-weighted average of the realized regressand over a whole simulated history, and iterated to a fixed point.

With a fixed bandwidth and fixed data, the recursion reproduces the batch estimate exactly.

In a learning economy the data are not fixed, because the y_t depend on the rule being learned, so the recursive and batch versions are related as (92.3) is to the batch algorithm of *Parameterized expectations*: different recursions with the same fixed points.

Because the simulated deficit lives on the grid, the kernel-weighted sum over the history can be computed state by state, from how often each state was visited and the sum of the y_t that followed those visits.

The fixed-point iteration converges slowly, at a rate set by essentially the same eigenvalue as before, so we allow it many passes and warn if it has not converged.

```
def kernel_smoother(T=150_000, seed=0, h=0.06, damp=0.5, tol=1e-8, max_iter=2000):
    """
    Non-parametric parameterized expectations.  ψ(G) is estimated by a
    Nadaraya-Watson smoother of the realized regressand -- no functional form.
    The rule is iterated to a self-consistent fixed point.  Returns the saving
    rule at each state and the number of passes.
    """
    idx = simulate_deficit(T, seed)
    log_G = np.log(G)
    K = np.exp(-0.5 * ((log_G[:, None] - log_G[None, :]) / h)**2)    # weight of state␣
↪j at i
    visits = np.bincount(idx[:-1], minlength=n_states)

    s_grid = np.full(n_states, (w1 - w2) / 2)
    for it in range(1, max_iter + 1):
        y = regressand(s_grid, G)[idx[1:]]                            # realized y_t
        y_sums = np.bincount(idx[:-1], weights=y, minlength=n_states)
        ψ = (K @ y_sums) / (K @ visits)           # kernel-weighted average over the␣
```

(continues on next page)

(continued from previous page)

```
↪history
        s_new = np.clip(s_from_psi(ψ), G + 0.05, w1 - 0.05)
        if np.max(np.abs(s_new - s_grid)) < tol:
            return s_new, it
        s_grid = damp * s_new + (1 - damp) * s_grid
    warnings.warn(f"kernel_smoother did not converge in {max_iter} passes")
    return s_grid, it


s_kernel, passes = kernel_smoother()
print(f"converged in {passes} passes")
print("kernel smoother: sup error = {:.4f},  ergodic RMS = {:.4f}".format(*errors(s_
↪kernel)))
```

```
converged in 140 passes
kernel smoother: sup error = 0.0828,  ergodic RMS = 0.0203
```

The smoother needs well over a hundred passes, for the same reason as the batch parameterized expectations algorithm.

```
fig, ax = plt.subplots(figsize=(6.5, 4.2))
ax.plot(G, s_benchmark, 'k-', lw=2.2, label="benchmark $f(G)$")
ax.plot(G, s_kernel, 'C3o', ms=4, markevery=5, label="kernel smoother")
ax.set_xlabel("deficit $G$")
ax.set_ylabel("saving $s$")
ax.legend(frameon=False)
plt.show()
```

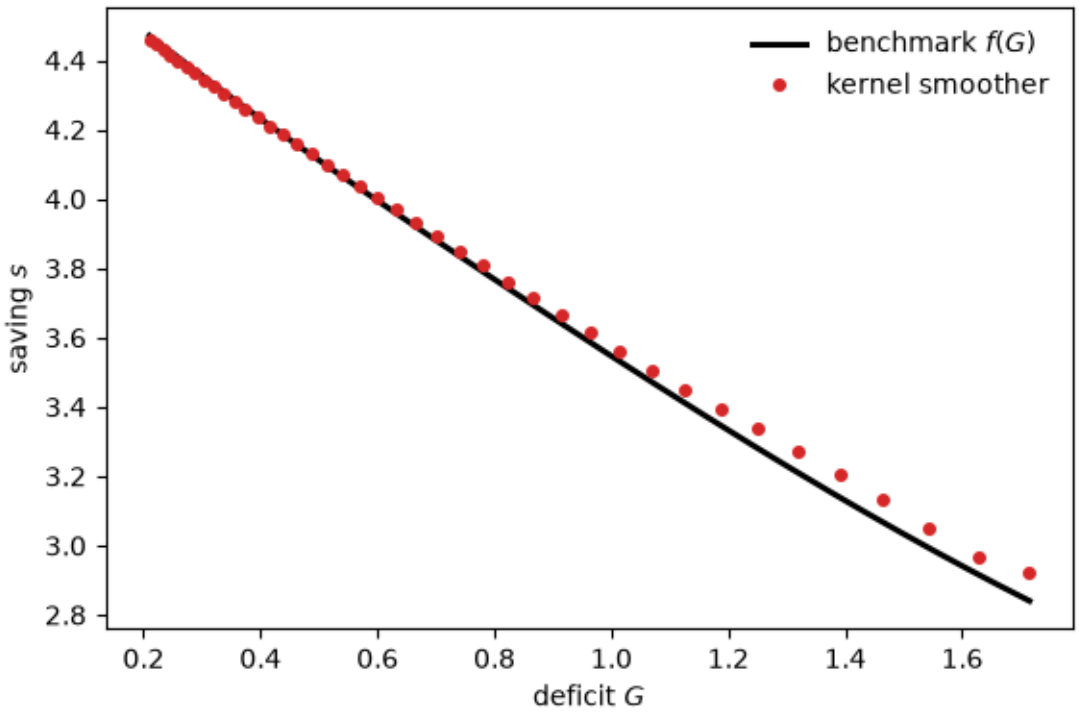


Fig. 92.4: Kernel smoother approximation to the saving rule

With no functional form imposed, the kernel learner recovers the benchmark well.

On this history it does less well than the quadratic rule fitted to the same data: both of its error measures are between two and three times as large.

The size of the gap varies from one simulated history to another, because a kernel estimate is noisier than a three-parameter fit, but much of the kernel's error here is systematic.

It grows toward the top of the deficit range, where the kernel overstates saving.

There the rule is steep, and most of the observations inside the kernel's window come from lower deficits, which are visited more often and call forth more saving; at the very top of the grid all of them do.

The quadratic family does so well because the true ψ *happens* to be almost exactly quadratic in $\log G$, so the parametric family we chose was very nearly correct.

Exercise 92.8.3 shows that a better choice of bandwidth closes most of the gap — but choosing it took knowledge of the benchmark, which no learner has.

When the family is not nearly correct, the two part ways, and the trade-off is the one familiar from econometrics:

- the **parametric** scheme learns fast and extrapolates gracefully, but is only as good as its family; a poorly chosen family delivers an approximate equilibrium and no warning;
- the **non-parametric** scheme cannot be led astray by a bad functional form, but pays for its flexibility with slower learning and with bias where the data is thin, at the edges of the deficit's range, where the kernel has few neighbours to average.

Both are recursions of the same shape.

Which one an adaptive agent should use is, once again, one of the modelling choices that the bounded-rationality program forces into the open and that rational expectations quietly made for us.

92.7 Concluding remarks

The overlapping generations model gave us a clean laboratory for a general lesson.

An adaptive agent in an environment with a large or continuous state cannot learn a separate response for every contingency.

It must **generalize**, imposing some structure that lets a finite amount of experience inform behavior across the whole state space.

Whether that structure is a low-order polynomial, a kernel bandwidth, or something else is a modelling choice with real consequences: it fixes whether the achievable limit is a rational expectations equilibrium or merely an approximation to one.

That same act of generalization is what makes learning and equilibrium computation two faces of one object.

Marcet's algorithm computes an equilibrium by parameterizing an expectation and regressing; an adaptive agent learns by parameterizing an expectation and regressing.

The two recursions are not identical, but they share their fixed points and their stability condition, and one eigenvalue of the T map sets how fast each of them converges.

The next lecture, *Exchange Rate Indeterminacy, Learning, and Experiments*, puts the two-currency economy of *A Peculiar Definition of Bounded Rationality* into this overlapping generations setting, where learners must cope with an exchange rate that rational expectations leaves undetermined.

The generalization problem in its sharpest form waits for *Networks, Genetic Algorithms, and Classifier Systems* and *Money as a Medium of Exchange among Artificially Intelligent Agents*.

There the state and action spaces are large enough that enumerating rules is out of the question, and agents must **discover** a compact set of good rules from scratch, using John Holland's genetic algorithm and classifier systems in place of the least squares and kernel machinery of this lecture.

92.8 Exercises

Exercise 92.8.1

The note above claimed that pushing the polynomial degree too high stops helping, because high-order terms chase sampling noise in the sparsely visited tails of the deficit distribution.

Check it, separating the two sources of error.

1. For degrees 0 through 5, compute the *population* approximate equilibrium, the fixed point of the family's T map with infinitely much data, as `approximate_equilibrium` does. Its distance from the benchmark is pure approximation error.
2. For the same degrees, run the batch parameterized-expectations algorithm on three simulated histories of length 150,000.

Report two error measures against the benchmark: the sup norm over the whole grid, and the RMS error weighted by the ergodic distribution of the deficit (i.e. weighted by where agents actually operate).

How does each measure behave as the degree rises, with and without sampling error?

Interpret.

Solution

```
seeds = (0, 1, 2)
sup_rows, rms_rows = [], []
for deg in range(6):
    Φ = basis(G, deg)
    sup_pop, rms_pop = errors(rule(Φ, approximate_equilibrium(Φ)))
    sups, rmss = [sup_pop], [rms_pop]
    for seed in seeds:
        sup, rms = errors(parameterized_expectations(Φ, seed=seed)[0])
        sups.append(sup)
        rmss.append(rms)
    sup_rows.append(sups)
    rms_rows.append(rmss)

cols = ["population"] + [f"history {seed}" for seed in seeds]
table = pd.concat({"sup error (whole grid)": pd.DataFrame(sup_rows, columns=cols),
                   "ergodic RMS (where agents live)": pd.DataFrame(rms_rows,␣
 ↪columns=cols)},
                  axis=1)
table.index.name = "degree"
table.round(4)
```

```
       sup error (whole grid)                                    \
                   population history 0 history 1 history 2
degree
0                      1.1793    1.1825    1.1795    1.1766
1                      0.3595    0.3626    0.3518    0.3414
2                      0.0307    0.0322    0.0333    0.0431
3                      0.0176    0.0170    0.0513    0.0882
4                      0.0095    0.0060    0.0517    0.0815
5                      0.0020    0.0038    0.0507    0.0705

       ergodic RMS (where agents live)
                          population history 0 history 1 history 2
degree
0                             0.3204    0.3212    0.3205    0.3198
1                             0.0677    0.0683    0.0668    0.0657
2                             0.0070    0.0073    0.0075    0.0116
3                             0.0020    0.0030    0.0062    0.0128
4                             0.0014    0.0027    0.0063    0.0126
5                             0.0004    0.0025    0.0062    0.0124
```

With infinitely much data, both measures fall steadily as the degree rises: adding terms to a polynomial in $\log G$ shrinks the approximation error, which by degree 5 is tiny.

With 150,000 observations, the errors fall steeply up to the quadratic, as in the population.

Beyond it the sample errors stop tracking the population errors: the ergodic-weighted errors stay above the population values on every history, and from degree 4 on the gap is about as large as the remaining approximation error, or larger.

The ergodic-weighted error levels off, at a level that differs from one history to another.

The sup-norm error behaves less regularly: on one history it keeps falling, while on the other two it rises at degree 3 and stays above its degree-2 value.

The high-order polynomials fit the center well but are free to move at the ends of the grid, where the ergodic distribution puts little mass and the regression has few data to discipline them.

This is the approximation-side image of the rare-state problem from *Adaptive Agents in an Overlapping Generations Monetary Economy*.

A learning scheme allocates its accuracy to where the observations are.

Misfitting a rarely-visited region shows up starkly in the sup norm but costs almost nothing in expected utility, which is why the ergodic-weighted measure is the economically relevant one.

Exercise 92.8.2

The constant family (degree 0) delivers an *approximate* equilibrium: a single saving rate, independent of the deficit, that satisfies the first-order condition only on average.

There is a natural benchmark for it.

Compute the single saving rate $\bar{s}$ that a degree-0 parameterized-expectations agent converges to, and compare it with the saving rate of the economy in which the deficit is fixed forever at its mean $\mathbb{E}[G]$.

That fixed-deficit economy has two steady states.

Find both, compute the slope of its T map at each, and say which one belongs in the comparison.

Are $\bar{s}$ and the fixed-deficit saving rate the same?

Should they be?

Which is larger, and why?

Solution

With the deficit fixed at G, the T map is $s \mapsto w_1 x/(w_2 + 2x)$ with $x = s - G$, so its slope is $w_1 w_2/(w_2 + 2x)^2$, and the steady states solve $2x^2 - (w_1 - w_2 - 2G)x + w_2 G = 0$.

```
# the degree-0 approximate equilibrium
s_const = rules[0][0]                          # constant across states
G_mean = ergodic @ G

# the deterministic economy with G fixed at its mean
x_roots = np.roots([2, -(w1 - w2 - 2 * G_mean), w2 * G_mean])
s_steady = np.sort(x_roots + G_mean)[::-1]     # high saving (low inflation) first
slopes = w1 * w2 / (w2 + 2 * (s_steady - G_mean))**2

s_iter = (w1 - w2) / 2                         # iterate the T map from s = 5
for _ in range(1000):
    s_iter = 0.5 * s_iter + 0.5 * s_from_psi(regressand(s_iter, G_mean))

print(f"mean deficit  E[G]                              : {G_mean:.4f}")
print(f"steady states with G fixed at E[G]              : {s_steady.round(4)}")
print(f"slope of the T map at each                      : {slopes.round(3)}")
print(f"iterating the T map from s = 5 lands on         : {s_iter:.4f}")
print(f"degree-0 approximate-equilibrium saving rate    : {s_const:.4f}")
```

```
mean deficit  E[G]                              : 0.6531
steady states with G fixed at E[G]              : [4.0342 1.619 ]
slope of the T map at each                      : [0.712 1.405]
iterating the T map from s = 5 lands on         : 4.0342
degree-0 approximate-equilibrium saving rate    : 4.0226
```

The fixed-deficit economy has two steady states, as in Part 2 of *Adaptive Agents in an Overlapping Generations Monetary Economy*.

At the high-saving, low-inflation one the slope of the T map is about 0.71, below one; at the low-saving one it is about 1.4, above one.

Iterating the saving map from $s = 5$ lands on the first, and nothing that iterates on beliefs can settle at the second.

That is the computational face of the stability reversal of *The stability reversal*, and it is why the low-inflation steady state is the right comparison.

The two saving rates are close but not equal, and the sign of the gap is predictable.

A regression on a constant returns the sample mean of y_t, so $\bar{s}$ is the fixed point of the map that sends s to $w_1\psi/(1+\psi)$ with $\psi = \mathbb{E}[y(s, G)]$, the expectation taken over the ergodic distribution of the deficit.

The fixed-deficit economy uses $\psi = y(s, \mathbb{E}[G])$ instead.

For given s, $y(s, G) = 1 - w_2/(w_2 + s - G)$ is concave in G, so by Jensen's inequality $\mathbb{E}[y(s, G)] < y(s, \mathbb{E}[G])$: spreading the deficit out lowers the expected payoff to saving.

The degree-0 map therefore lies below the fixed-deficit map at every s.

Both maps cross the 45-degree line from above at their stable fixed points, so $\bar{s}$ is below the fixed-deficit saving rate, as the numbers confirm.

They would agree only if the deficit were deterministic to begin with.

The lesson is that an approximate equilibrium is not a naive "certainty-equivalent" object.

Even the crudest one-parameter learner is responding to the *whole* distribution of outcomes it experiences, not to a point forecast.

Exercise 92.8.3

The non-parametric kernel learner has one free tuning constant: the bandwidth h.

Explore its effect on a shorter history, $T = 40{,}000$, where sampling noise matters.

For bandwidths from 0.003 to 0.2, run the kernel learner on ten simulated histories.

Split its ergodic-weighted mean squared error into a squared bias (the error of the rule averaged across histories) and a variance (the dispersion of the rules across histories).

Explain the shape of the trade-off.

What does the smallest bandwidth correspond to?

How does the kernel learner at its best bandwidth compare with the quadratic parametric rule fitted to the same histories?

Solution

```
T_short, seeds = 40_000, range(10)
bandwidths = (0.003, 0.01, 0.02, 0.03, 0.05, 0.08, 0.12, 0.20)

rows = []
for h in bandwidths:
    S = np.array([kernel_smoother(T=T_short, seed=seed, h=h)[0] for seed in seeds])
    bias2 = ergodic @ (S.mean(axis=0) - s_benchmark)**2
    var = ergodic @ S.var(axis=0)
    rows.append([h, np.sqrt(bias2), np.sqrt(var), np.sqrt(bias2 + var)])

table = pd.DataFrame(rows, columns=["bandwidth $h$", "bias", "std. deviation",
                                    "RMS error"]).set_index("bandwidth $h$")

S_quad = np.array([parameterized_expectations(Φ2, T=T_short, seed=seed)[0] for␣
 ↪seed in seeds])
rms_quad = np.sqrt(ergodic @ ((S_quad - s_benchmark)**2).mean(axis=0))
step = np.diff(np.log(G))[0]

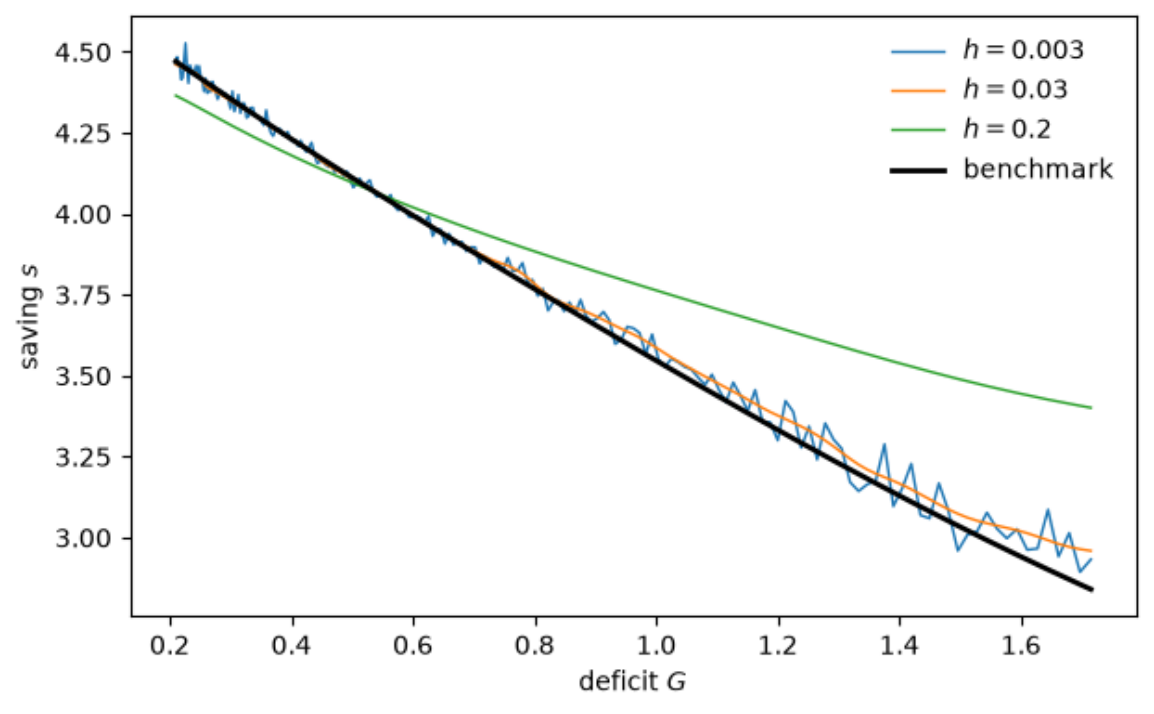
print(f"quadratic family on the same histories: RMS error {rms_quad:.4f}")
print(f"at h = 0.003 the weight on a neighbouring state is {np.exp(-0.5 * (step /␣
 ↪0.003)**2):.4f}")
table.round(4)
```

```
quadratic family on the same histories: RMS error 0.0133
at h = 0.003 the weight on a neighbouring state is 0.0022
```

```
                  bias  std. deviation  RMS error
bandwidth $h$
0.003           0.0088          0.0272     0.0286
0.010           0.0046          0.0178     0.0184
0.020           0.0027          0.0151     0.0153
0.030           0.0030          0.0138     0.0141
0.050           0.0110          0.0124     0.0166
0.080           0.0307          0.0110     0.0326
0.120           0.0636          0.0093     0.0643
0.200           0.1294          0.0075     0.1296
```

The rules themselves show the two failure modes.

```
fig, ax = plt.subplots(figsize=(7, 4.2))
for h, colour in [(0.003, 'C0'), (0.03, 'C1'), (0.20, 'C2')]:
    ax.plot(G, kernel_smoother(T=T_short, h=h)[0], '-', lw=1, color=colour, label=f
↪"$h = {h}$")
ax.plot(G, s_benchmark, 'k-', lw=2, label="benchmark")
ax.set_xlabel("deficit $G$")
ax.set_ylabel("saving $s$")
ax.legend(frameon=False)
plt.show()
```



The variance falls steadily as the bandwidth widens.

The bias is small up to about $h = 0.03$ and then rises steeply.

(At the smallest bandwidths the bias column mostly measures the noise in an average of ten noisy rules, about the standard deviation divided by $\sqrt{10}$.)

So the RMS error has an interior minimum, here near $h = 0.03$.

A **small** bandwidth averages over few neighbours, so the estimate is noisy: it wiggles around the benchmark, most of all toward the high-deficit end of the range.

This is the **variance** end of the trade-off.

At $h = 0.003$ the weight on even the nearest neighbouring state is negligible, so the learner is effectively the state-by-state learner of *Adaptive Agents in an Overlapping Generations Monetary Economy*, with one estimate per state: the $h \to 0$ limit of (92.5).

A **large** bandwidth averages over a wide window, smoothing the estimate but flattening it: the rule is pulled toward the saving of the middle of the range, below the benchmark at low deficits and above it at high ones.

This is the **bias** end.

At its best bandwidth the kernel learner is about as accurate as the quadratic family fitted to the same histories, but finding that bandwidth took the benchmark, which no learner has.

The bandwidth is the non-parametric counterpart of the polynomial degree in Exercise 92.8.1: both control how much structure the learner imposes, and both have a sweet spot that trades misfit-from-too-little-flexibility against noise-from-too-much.

The bounded-rationality program does not tell us where that spot is; it is one more choice that replaces the single discipline of rational expectations with a menu of plausible alternatives.

Exercise 92.8.4

Return to Economy 2 of *Adaptive Agents in an Overlapping Generations Monetary Economy*.

There the deficit is 0.8 in state 1 and 0 in state 2, the transition matrix is

$$\begin{bmatrix} 0.75 & 0.25 \\ 0.50 & 0.50 \end{bmatrix},$$

and the rational expectations saving rates are $(4.211, 4.364)$, which the Robbins–Monro learners of that lecture found state by state.

1. Run the batch parameterized-expectations algorithm on this economy with the basis $\phi(G) = (1, G)$. (The log-deficit basis of this lecture won't do, because $G = 0$ in state 2.) Show that it reproduces the rational expectations saving rates, and explain why it must.
2. Now use a constant basis, $\phi(G) = 1$. Run both the batch algorithm and the online learner (92.3), and compare their limits with each other and with the rational expectations equilibrium.

The functions `parameterized_expectations`, `online_pea` and `approximate_equilibrium` accept the deficit levels, the transition matrix and the basis evaluated at each state as arguments.

Solution

Here is part 1.

```
G_2 = np.array([0.8, 0.0])
P_2 = np.array([[0.75, 0.25],
                [0.50, 0.50]])
π_2 = qe.MarkovChain(P_2).stationary_distributions[0]
economy_2 = dict(G=G_2, P=P_2)

Φ_lin = np.column_stack([np.ones(2), G_2])
s_lin_batch = parameterized_expectations(Φ_lin, **economy_2)[0]
s_lin_pop = rule(Φ_lin, approximate_equilibrium(Φ_lin, π=π_2, **economy_2), G_2)
```

```
print(f"batch, ψ = (1, G), T = 150,000 : {s_lin_batch.round(4)}")
print(f"population fixed point         : {s_lin_pop.round(4)}")
print( "rational expectations          : [4.211 4.364]")
```

```
batch, ψ = (1, G), T = 150,000 : [4.2088 4.3629]
population fixed point         : [4.2111 4.3637]
rational expectations          : [4.211 4.364]
```

On a two-point state space every function of G is a linear function of G: two coefficients can match any two values.

So the linear family contains the rational expectations equilibria exactly, and every fixed point of the algorithm is one of them.

This economy also has a high-inflation equilibrium, with saving of about $(1.19, 1.63)$, but it is not E-stable, so the algorithm settles on the one shown.

The population fixed point matches it to every digit shown, and the batch version differs from it only by sampling error.

On two points the "parametric" family is non-parametric in disguise: it reaches the same equilibrium that the learners of *Adaptive Agents in an Overlapping Generations Monetary Economy* found state by state, but by model formation rather than the hard way.

Here is part 2, with the online learner run for 200,000 periods under both bases.

```
Φ_const = np.ones((2, 1))
s_const_batch = parameterized_expectations(Φ_const, **economy_2)[0]
s_const_pop = rule(Φ_const, approximate_equilibrium(Φ_const, π=π_2, **economy_2),␣
↪G_2)

T_2 = 200_000
paths = {"constant": (Φ_const, online_pea(Φ_const, T=T_2, **economy_2)),
         "linear":   (Φ_lin,   online_pea(Φ_lin,   T=T_2, **economy_2))}

print(f"constant basis, batch           : {s_const_batch.round(4)}")
print(f"constant basis, population      : {s_const_pop.round(4)}")
for name, (Φ, path) in paths.items():
    for t in (1_000, 10_000, 100_000, T_2 - 1):
        print(f"{name:8s} basis, online, t = {t:>7,}: {rule(Φ, path[t], G_2).
↪round(4)}")
```

```
constant basis, batch           : [4.2691 4.2691]
constant basis, population      : [4.271 4.271]
constant basis, online, t =   1,000: [4.3487 4.3487]
constant basis, online, t =  10,000: [4.3019 4.3019]
constant basis, online, t = 100,000: [4.2847 4.2847]
constant basis, online, t = 199,999: [4.2809 4.2809]
linear   basis, online, t =   1,000: [4.3041 4.4333]
linear   basis, online, t =  10,000: [4.251 4.392]
linear   basis, online, t = 100,000: [4.228  4.3789]
linear   basis, online, t = 199,999: [4.2235 4.3751]
```

The constant family cannot let saving depend on the state, so its fixed point is a single saving rate, between the two rational expectations rates: an approximate equilibrium.

The batch algorithm and the online learner aim at that same rate, not at the rational expectations equilibrium.

The online learner approaches it slowly, from above, because it starts from $s = 5$.

With the linear basis the online learner instead heads for the rational expectations saving rates.

What a learner can reach is set by its family, not by whether it learns in batch or in real time.

Exercise 92.8.5

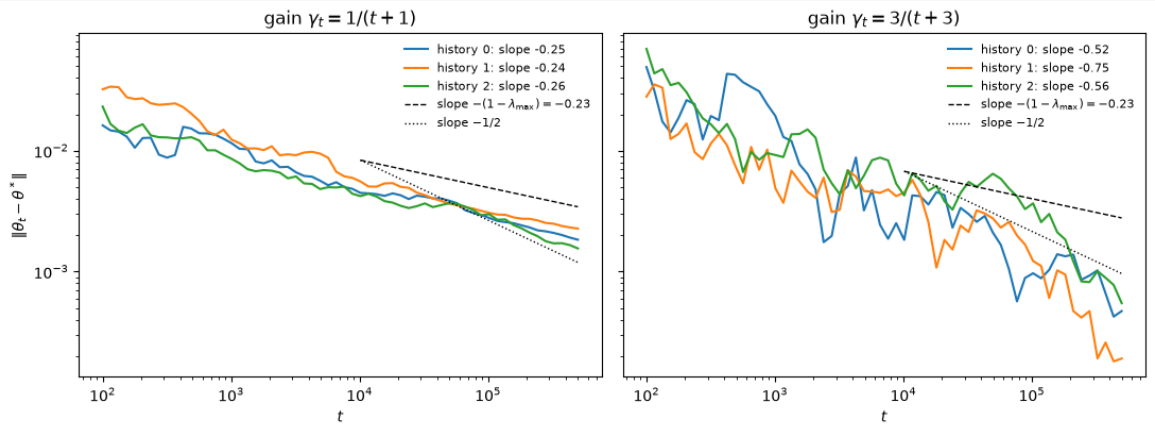
The text predicted that the online learner's distance from the approximate equilibrium θ^* shrinks like $t^{-(1-\lambda_{\max})}$, where $\lambda_{\max}$ is the largest eigenvalue of $DT(\theta^*)$.

1. Plot $\|\theta_t - \theta^*\|$ against t on log–log axes, for the path computed in the lecture and for two more simulated histories. Fit the slope over $t \geq 10{,}000$, and compare it with $-(1-\lambda_{\max})$ and with $-1/2$, the rate at which sampling noise averages out.
2. Repeat with the gain $\gamma_t = 3/(t+3)$ in place of $1/(t+1)$, which `online_pea` delivers with `a=3`. What changes, and why?

Solution

```
T_online = len(θ_path)
ts = np.unique(np.logspace(2, np.log10(T_online - 1), 60).astype(int))
late = ts >= 10_000

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2), sharey=True)
for ax, a in zip(axes, (1, 3)):
    starts = []
    for seed in (0, 1, 2):
        path = θ_path if (a == 1 and seed == 0) else online_pea(Φ2, seed=seed, a=a)
        dist = np.linalg.norm(path[ts] - θ_star, axis=1)
        slope = np.polyfit(np.log(ts[late]), np.log(dist[late]), 1)[0]
        ax.loglog(ts, dist, lw=1.5, label=f"history {seed}: slope {slope:.2f}")
        starts.append(dist[late][0])
    t_rel = ts[late] / ts[late][0]                # reference lines start at t =␣
↪10,000
    anchor = 1.5 * max(starts)
    ax.loglog(ts[late], anchor * t_rel**(-(1 - λ_max)), 'k--', lw=1,
              label=rf"slope $-(1-\lambda_{{\max}}) = {-(1 - λ_max):.2f}$")
    ax.loglog(ts[late], anchor * t_rel**(-0.5), 'k:', lw=1, label="slope $-1/2$")
    ax.set_xlabel("$t$")
    ax.set_title(f"gain $\\gamma_t = {a}/(t + {a})$")
    ax.legend(frameon=False, fontsize=8)
axes[0].set_ylabel(r"$\|\theta_t - \theta^*\|$")
plt.tight_layout()
plt.show()
```



With the gain $1/(t+1)$ the fitted slopes are close to $-(1-\lambda_{\max}) \approx -0.23$ and far from $-1/2$.

In the rescaled time $\tau = \log t$, the mean dynamics (92.4) contract at the rates $1-\lambda_i$, the slowest of which is $1-\lambda_{\max}$, so the drift toward θ^* decays like $t^{-(1-\lambda_{\max})}$, while sampling noise averages out like $t^{-1/2}$.

The slower of the two dominates, and here it is the drift.

A gain of a/t runs the differential equation a times faster, so the drift decays like $t^{-a(1-\lambda_{\max})}$.

With $a = 3$ that is faster than $t^{-1/2}$, so in the long run sampling noise takes over and the error falls at the usual $\sqrt{t}$ rate.

The fitted slopes are now near $-1/2$ or steeper, and after half a million periods the error is several times smaller than with the gain $1/(t+1)$.

This is the sense in which the book's condition for $\sqrt{t}$ convergence is a condition on the gain as well as on DT.

Larger gains are not free: they make the early phase, far from θ^*, more violent, and a gain that never declines would leave beliefs fluctuating around θ^* forever, the constant-gain behaviour described in *One recursion, many learners*.

CHAPTER

NINETYTHREE

EXCHANGE RATE INDETERMINACY, LEARNING, AND EXPERIMENTS

Contents

- *Exchange Rate Indeterminacy, Learning, and Experiments*
  - *Overview*
  - *The two-currency economy*
  - *Newton–Raphson learning*
  - *Evidence from the laboratory*
  - *A genetic algorithm economy*
  - *Concluding remarks*
  - *Exercises*

93.1 Overview

In *A Peculiar Definition of Bounded Rationality* we met a model in which rational expectations cannot pin down the exchange rate at all.

Two fiat currencies, perfect substitutes as long as their rates of return are equal, leave the exchange rate e *completely unrestricted*: if the equilibrium conditions have a solution for one e, they have a solution for every other, and the real allocation is identical across all of them.

This lecture, following Sargent [1993], expels the rational agents from that economy and puts adaptive ones in their place.

It is the hinge of the series.

Its first half uses the last of the series' econometric learners, a Newton–Raphson rule like the one in *Adaptive Agents in an Overlapping Generations Monetary Economy*.

Its second half uses the series' first genetic algorithm, a device that *Networks, Genetic Algorithms, and Classifier Systems* then sets out in general.

Two questions organize what follows.

First, does learning pin the exchange rate down?

An adaptive agent has beliefs, a rule for revising them, and initial conditions, and those initial conditions are exactly what the equilibrium conditions failed to determine.

We shall see that Newton–Raphson learners do converge to a constant exchange rate, but that learning does not *select* one.

Where they end up depends on where their portfolios and saving rates started, and also on how the learning algorithm is tuned.

Learning makes the exchange rate *history-dependent*: what Sargent calls the "dead hand of history" does the pinning that fundamentals refused to do.

Second, does a ghost of the indeterminacy survive?

It does.

The rest points of the learning algorithm reproduce the indeterminacy exactly: the condition that stops the algorithm is the same arbitrage condition that made the exchange rate free in the first place.

Sargent judges that an adaptive mechanism that selects the exchange rate path in this way "may seem a weak reed on which to base exchange rate determination."

Then we turn to evidence.

Arifovic [1996] ran this economy as a laboratory experiment with paying human subjects.

Their exchange rate did not settle down.

When she replaced the Newton–Raphson learners with a **genetic algorithm** — a population of binary-string agents bred by selection, crossover, and mutation — she got an economy whose exchange rate does not settle down either.

The genetic population learns to save at close to the rational expectations rate, but nothing pins down its portfolio share, which wanders along the continuum of equilibria.

The ghost of indeterminacy returns as volatility.

That contrast — a learning rule that converges versus one that generates permanent volatility — is where this lecture lands.

Let's start with some imports.

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from typing import NamedTuple
```

93.2 The two-currency economy

We use the overlapping generations incarnation of the Kareken and Wallace [1981] indeterminacy.

At each date t, N two-period-lived agents are born, endowed with w_1 when young and w_2 when old, with $w_1 > w_2$.

(As in *Adaptive Agents in an Overlapping Generations Monetary Economy*, the endowments are w_1 and w_2; *A Peculiar Definition of Bounded Rationality* writes them as $2w_1$ and $2w_2$.)

There are two fiat currencies in fixed supplies H_1 and H_2.

A young agent makes two decisions: how much to save, s_t, and what fraction λ_t of that saving to hold in currency 1 (the rest going to currency 2).

The realized lifetime utility of an agent who chooses (s_t, λ_t) is

$$U(s_t, \lambda_t) = u(w_1 - s_t) + u\left(w_2 + \lambda_t s_t \frac{p_{1t}}{p_{1,t+1}} + (1-\lambda_t) s_t \frac{p_{2t}}{p_{2,t+1}}\right), \tag{93.1}$$

where p_{it} is the price level in currency i and $p_{it}/p_{i,t+1}$ is the gross return on holding currency i.

We take $u(c) = \ln c$ throughout.

Equating each currency's supply to the demand for it gives the price levels

$$p_{1t} = \frac{H_1}{\sum_i \lambda_{it} s_{it}}, \qquad p_{2t} = \frac{H_2}{\sum_i (1-\lambda_{it}) s_{it}}, \tag{93.2}$$

and the exchange rate is $e_t = p_{1t}/p_{2t}$.

93.2.1 The indeterminacy, recalled

Consider a stationary equilibrium with constant prices.

Then each currency returns $p_{it}/p_{i,t+1} = 1$, so both returns are equal, and the portfolio return in (93.1) is 1 regardless of λ.

The saving decision then solves $\max_s \ln(w_1 - s) + \ln(w_2 + s)$, giving

$$s^\star = \frac{w_1 - w_2}{2},$$

but the portfolio share λ is *completely undetermined*: utility does not depend on it when the two returns are equal.

And λ is what fixes the exchange rate.

From (93.2), with a common (s, λ),

$$e = \frac{p_1}{p_2} = \frac{H_1}{H_2} \cdot \frac{1-\lambda}{\lambda}.$$

Every $\lambda \in (0, 1)$ is an equilibrium, so every $e \in (0, \infty)$ is an equilibrium.

This is the indeterminacy of *A Peculiar Definition of Bounded Rationality*, now expressed through an undetermined portfolio choice.

This lecture runs two economies with different parameters.

The Newton–Raphson section uses Sargent's $(w_1, w_2, H_1, H_2) = (20, 15, 100, 120)$, whose currency supplies are those of the two-currency example in *A Peculiar Definition of Bounded Rationality*.

The laboratory and genetic-algorithm sections use Arifovic's $(11, 1, 10, 10)$.

So we carry the parameters in a container and pass them explicitly rather than leaving them lying about as globals.

```
class Params(NamedTuple):
    w1: float          # endowment when young
    w2: float          # endowment when old
    H1: float          # fixed supply of currency 1
    H2: float          # fixed supply of currency 2

    @property
    def s_star(self):
        "Rational expectations saving rate, when both currencies return 1."
        return (self.w1 - self.w2) / 2


kw = Params(w1=20.0, w2=15.0, H1=100.0, H2=120.0)

print(f"rational expectations saving rate s* = {kw.s_star}")
print(f"any portfolio share λ is an equilibrium, and e = (H1/H2)(1-λ)/λ is free")
```

```
rational expectations saving rate s* = 2.5
any portfolio share λ is an equilibrium, and e = (H1/H2)(1-λ)/λ is free
```

93.3 Newton–Raphson learning

Now expel the rational agents.

As in *Adaptive Agents in an Overlapping Generations Monetary Economy*, a cohort's lifetime utility cannot be evaluated until it is old, so the population is split into an "even" and an "odd" class.

Each class carries its own rule (s, λ), updated only from the experience of earlier agents of the same class.

An agent revises (s, λ) by a **Newton–Raphson** step against realized utility: move in the direction that a second-order expansion of (93.1) says would have raised utility, given the returns the agent actually experienced.

Writing ∇U_τ and $\nabla^2 U_\tau$ for the gradient and Hessian of realized utility and M for a running estimate of the (negative-definite) Hessian, the recursion is

$$\begin{aligned} M_{\tau+1} &= M_\tau + \gamma_\tau(\nabla^2 U_\tau - M_\tau), \\ \begin{bmatrix} s \\ \lambda \end{bmatrix}_{\tau+1} &= \begin{bmatrix} s \\ \lambda \end{bmatrix}_\tau - \gamma_\tau M_{\tau+1}^{-1} \nabla U_\tau, \end{aligned} \tag{93.3}$$

with γ_τ a gain.

This is the two-variable version of the Robbins–Monro rule of *Adaptive Agents in an Overlapping Generations Monetary Economy*.

We use a **constant gain**, $\gamma_\tau = \gamma = 0.3$, a variant that Sargent mentions in a footnote.

Because this economy has no shocks, a constant gain can still converge to a point.

Here are the gradient and Hessian of (93.1) in the returns $R_1 = p_{1t}/p_{1,t+1}$ and $R_2 = p_{2t}/p_{2,t+1}$.

```
def grad_hess(p, s, lam, R1, R2):
    "Gradient and Hessian of realized utility U(s, λ)."
    A = lam*R1 + (1 - lam)*R2               # portfolio gross return
    c2 = p.w2 + s*A                         # old-age consumption
    dR = R1 - R2
    g = np.array([-1/(p.w1 - s) + A/c2,
                  s*dR/c2])
    U_ss = -1/(p.w1 - s)**2 - A**2/c2**2
    U_ll = -s**2 * dR**2 / c2**2
    U_sl = dR/c2 - s*A*dR/c2**2
    return g, np.array([[U_ss, U_sl], [U_sl, U_ll]])
```

93.3.1 The indeterminacy shows up as a singular Hessian

Look at the λ block of the Hessian.

When the two returns are equal, $R_1 = R_2$, every λ-derivative vanishes: the gradient component $s(R_1 - R_2)/c_2$ is zero, and so is the curvature $-s^2(R_1 - R_2)^2/c_2^2$.

Utility is **flat in** λ at equal returns.

That is the indeterminacy, seen locally: there is no force pushing λ in any direction.

Near equal returns the curvature is of second order in $R_1 - R_2$ while the gradient is of first order, so an undamped Newton step in λ, which divides the gradient by the curvature, is enormous.

We therefore add a constant **ridge** $\kappa > 0$ to the curvature estimate in the λ direction: each period we feed (93.3) the matrix $\nabla^2 U_\tau - \mathrm{diag}(0, \kappa)$ in place of $\nabla^2 U_\tau$.

This is a one-direction version of the Levenberg–Marquardt device of numerical optimization, and it is this lecture's addition, not the book's.

It is a *stabilizer*, not an enabler.

Without it the algorithm does not freeze; it explodes, as we check below.

With it, a λ step is roughly $\gamma\,|\partial U/\partial\lambda|/\kappa$, so a large κ makes portfolios adjust sluggishly.

```
def newton_learning(p, s0, lam0, T=400, gain=0.3, κ=0.5, w2_path=None):
    """
    Two-currency OLG with Newton-Raphson learning of (s, λ), one rule per class
    (even/odd).  κ is a ridge added to the curvature estimate M in the
    λ direction.  w2_path optionally gives the old-age endowment at each date.
    """
    s = np.array(s0, float)
    lam = np.array(lam0, float)
    M = [np.array([[-1.0, 0.0], [0.0, -κ]]) for _ in range(2)]   # initial curvature
    p1 = np.empty(T)
    p2 = np.empty(T)
    e = np.empty(T)
    s_hist = np.empty((T, 2))
    lam_hist = np.empty((T, 2))

    for t in range(T):
        j = t % 2                                        # class that is young at t
        p1[t] = p.H1 / (lam[j]*s[j])
        p2[t] = p.H2 / ((1 - lam[j])*s[j])
        e[t] = p1[t] / p2[t]
        if t >= 1:                                       # class young at t-1 is now old
            jp = 1 - j
            R1, R2 = p1[t-1]/p1[t], p2[t-1]/p2[t]
            q = p if w2_path is None else p._replace(w2=w2_path[t])
            g, h = grad_hess(q, s[jp], lam[jp], R1, R2)
            h[1, 1] -= κ                                 # the ridge
            M[jp] = M[jp] + gain*(h - M[jp])
            step = gain * np.linalg.solve(M[jp], g)
            s[jp] = np.clip(s[jp] - step[0], 0.1, p.w1 - 0.1)
            lam[jp] = np.clip(lam[jp] - step[1], 0.02, 0.98)
        s_hist[t] = s
        lam_hist[t] = lam

    return e, s_hist, lam_hist
```

93.3.2 Convergence to a history-dependent exchange rate

Run two experiments that are identical except for the initial saving rates and portfolio shares.

```
e1, s1, l1 = newton_learning(kw, [3.5, 2.0], [0.35, 0.40])
e2, s2, l2 = newton_learning(kw, [2.0, 3.5], [0.62, 0.58])

fig, axes = plt.subplots(1, 3, figsize=(14, 4))

axes[0].plot(np.log(e1), 'C0', lw=1.3, label="experiment 1")
axes[0].plot(np.log(e2), 'C1', lw=1.3, label="experiment 2")
axes[0].set_title("log exchange rate")
axes[0].set_xlabel("$t$")
axes[0].legend(frameon=False)

for s_h, c in [(s1, 'C0'), (s2, 'C1')]:
    axes[1].plot(s_h[:, 0], color=c, lw=1.0)
    axes[1].plot(s_h[:, 1], color=c, lw=1.0, ls=':')
axes[1].axhline(kw.s_star, color='k', lw=0.8, ls='--')
axes[1].set_title("saving  (solid = even, dotted = odd)")
axes[1].set_xlabel("$t$")

for l_h, c in [(l1, 'C0'), (l2, 'C1')]:
    axes[2].plot(l_h[:, 0], color=c, lw=1.0)
    axes[2].plot(l_h[:, 1], color=c, lw=1.0, ls=':')
axes[2].set_title(r"portfolio share $\lambda$")
axes[2].set_xlabel("$t$")
plt.tight_layout()
plt.show()
```

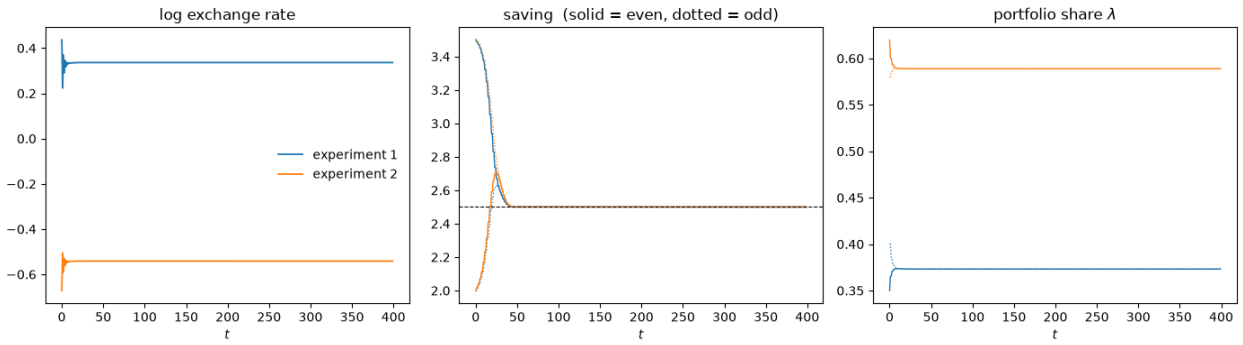


Fig. 93.1: Newton-Raphson learning from two initial conditions

Both economies converge.

Saving climbs or falls to the rational expectations rate $s^\star = 2.5$ from either side, and the two classes' portfolio shares merge to a common value.

But the two experiments converge to *different exchange rates*.

```
for name, e, s_h, l_h in [("experiment 1", e1, s1, l1), ("experiment 2", e2, s2, l2)]:
    still = np.abs(np.diff(l_h, axis=0)).max(axis=1) > 1e-6
    t_freeze = np.flatnonzero(still)[-1] + 2          # λ moves by < 1e-6 from here on
    print(f"{name}: e → {e[-1]:.4f} (log e → {np.log(e[-1]):.3f}),  s → {s_h[-1][0]:.
```

(continues on next page)

(continued from previous page)

```
↪3f},  "
          f"λ → {l_h[-1][0]:.4f},  λ frozen from t = {t_freeze}")
```

```
experiment 1: e → 1.3998 (log e → 0.336),  s → 2.500,  λ → 0.3732,  λ frozen from␣
↪t = 28
experiment 2: e → 0.5819 (log e → -0.541),  s → 2.500,  λ → 0.5888,  λ frozen from␣
↪t = 35
```

The saving rate is pinned down by fundamentals; the exchange rate is pinned down by history.

Experiment 1's limit is close to that of Sargent's experiment 1, whose λ settles near 0.37 and whose log exchange rate settles near 0.3.

Sargent's paths are still drifting after a few hundred periods, while with our constant gain λ stops moving within about 35 periods.

Before looking at how the limit depends on history, let's check what the ridge κ does.

Here is experiment 1 again with $\kappa = 0$.

```
e0, s0_h, l0_h = newton_learning(kw, [3.5, 2.0], [0.35, 0.40], κ=0.0)
print("last four exchange rates:", ", ".join(f"{x:.3f}" for x in e0[-4:]))
print(f"final λ (even, odd): {l0_h[-1]},  final s (even, odd): {s0_h[-1]}")
```

```
last four exchange rates: 40.833, 0.017, 40.833, 0.017
final λ (even, odd): [0.02 0.98],  final s (even, odd): [0.1 0.1]
```

Without the ridge the algorithm does not stand still.

The curvature estimate stays invertible, but its λ entry is tiny, so the λ steps are huge.

The portfolio shares crash into the bounds 0.02 and 0.98 that the code imposes, saving collapses to its lower bound, and from then on the exchange rate flips between two extreme values.

93.3.3 The dead hand of history

To see how history governs the outcome, sweep the initial portfolio shares and record the exchange rate each economy settles on.

The two classes must start with *different* portfolio shares.

If they start with the same λ_0, the exchange rate, which depends only on the young class's λ, is constant from the first period on.

Then $R_1 = R_2$ from the start, the λ-gradient is identically zero, and λ never moves, so the limit would be $(H_1/H_2)(1-\lambda_0)/\lambda_0$ by construction.

So we start the even and odd classes at $\bar\lambda_0 - 0.05$ and $\bar\lambda_0 + 0.05$, with experiment 1's saving rates, and plot each limit against the mean initial share $\bar\lambda_0$.

We compare the limits with the curve $(H_1/H_2)(1-\lambda)/\lambda$ evaluated at two points: the limiting share λ_∞, and the mean initial share $\bar\lambda_0$.

```
def e_of_λ(p, lam):
    "Exchange rate when every agent holds portfolio share λ."
    return (p.H1/p.H2) * (1 - lam) / lam
```

(continues on next page)

(continued from previous page)

```
λ_bar0 = np.linspace(0.15, 0.85, 15)
e_lim = np.empty_like(λ_bar0)
λ_lim = np.empty_like(λ_bar0)
for k, lb in enumerate(λ_bar0):
    e, s_h, l_h = newton_learning(kw, [3.5, 2.0], [lb - 0.05, lb + 0.05])
    e_lim[k], λ_lim[k] = e[-1], l_h[-1].mean()

fig, ax = plt.subplots(figsize=(7, 4.5))
ax.plot(λ_bar0, e_lim, 'C0o', ms=6, label="limiting $e$ from learning")
ax.plot(λ_bar0, e_of_λ(kw, λ_lim), 'C0-', lw=1,
        label=r"$(H_1/H_2)(1-\lambda_\infty)/\lambda_\infty$")
ax.plot(λ_bar0, e_of_λ(kw, λ_bar0), 'k--', lw=1,
        label=r"$(H_1/H_2)(1-\bar\lambda_0)/\bar\lambda_0$")
ax.set_yscale('log')
ax.set_xlabel(r"mean initial portfolio share $\bar\lambda_0$")
ax.set_ylabel("limiting exchange rate $e$")
ax.legend(frameon=False)
plt.show()
```

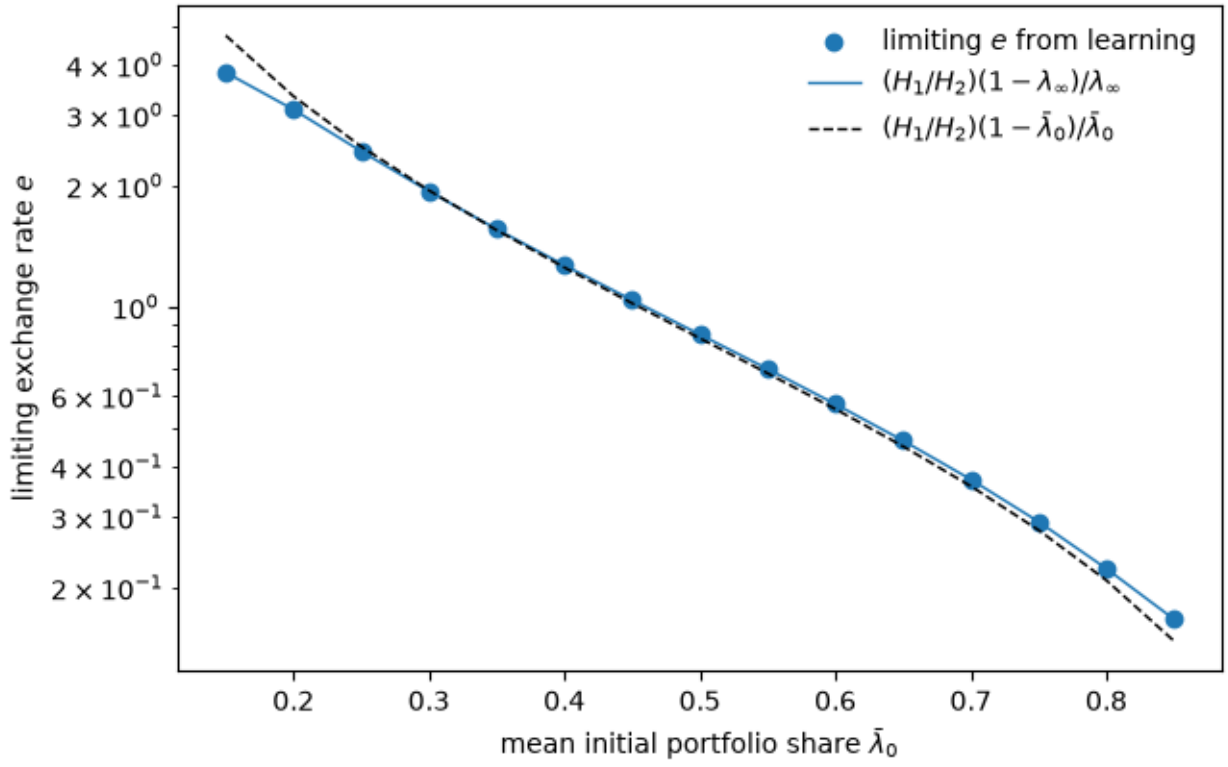


Fig. 93.2: Limiting exchange rate against the mean initial portfolio share

Here are the gaps between the limits and the two curves.

```
gap_lim = np.abs(e_lim / e_of_λ(kw, λ_lim) - 1)
gap_init = np.abs(e_lim / e_of_λ(kw, λ_bar0) - 1)
print(f"largest relative gap from the curve at λ∞:  {gap_lim.max():.1e}")
```

(continues on next page)

(continued from previous page)

```
print(f"relative gap from the curve at mean λ0:     "
      f"{gap_init.min():.3f} to {gap_init.max():.3f}")
```

```
largest relative gap from the curve at λ∞:  2.2e-16
relative gap from the curve at mean λ0:     0.005 to 0.191
```

Every limit lies exactly on the rational expectations continuum, evaluated at its own limiting λ: each one is an equilibrium.

The zero gap is close to automatic: in every period, transients included, the exchange rate is $(H_1/H_2)(1-\lambda)/\lambda$ at the young class's share, so all the gap confirms is that the two classes' shares have merged.

The informative comparison is with the dashed curve.

As $\bar{\lambda}_0$ sweeps from 0.15 to 0.85, the limits trace out a wide stretch of the continuum.

But the limit is not simply the equilibrium that the mean initial share implies.

While the two classes still differ, the two currencies' returns differ, and those transient return gaps push λ around before the arbitrage condition shuts the motion off.

The gaps from the dashed curve range from half a percent to almost twenty percent.

They are largest at the edges of the grid, where the classes' initial shares imply very different exchange rates.

This is what Sargent means by the dead hand of history:

> Our learning system renders the exchange rate path determinate by having the 'dead hand of history' put enough sluggishness into decisions. The exchange rate path can be said to be 'history-dependent' because the initial conditions assigned to (λ, s) assume an importance that does not vanish as time passes.

The relevant history is more than the initial portfolio shares.

The limit also depends on the initial saving rates and on the tuning of the algorithm, the ridge κ and the gain γ; Exercise 93.7.1 measures both.

Learning converts an *indeterminacy of levels* into a *dependence on history*.

It makes the exchange rate history-dependent; it does not select one.

93.3.4 The ghost of indeterminacy

Why does the exchange rate freeze rather than move to some particular value?

Because the algorithm comes to rest wherever the gradient vanishes, and the portfolio gradient $s(R_1 - R_2)/c_2$ vanishes precisely when $R_1 = R_2$, the **arbitrage condition** that made the exchange rate free under rational expectations.

Let's compute the returns and the gradient at the end of experiment 1, and the λ derivatives at other points of the continuum.

```
s_lim, lam_lim = s1[-1], l1[-1]                    # (even, odd) rules at the end
p1_lim = kw.H1 / (lam_lim*s_lim)                   # price levels when each class is␣
 ↪young
p2_lim = kw.H2 / ((1 - lam_lim)*s_lim)
R1, R2 = p1_lim[1]/p1_lim[0], p2_lim[1]/p2_lim[0]  # returns to an odd agent
g, h = grad_hess(kw, s_lim[1], lam_lim[1], R1, R2)
print(f"end of experiment 1: R1 = {R1:.6f}, R2 = {R2:.6f}, gradient = {g.round(8) + 0.
 ↪0}")

print("\nλ derivatives at other rest points (s = s*, R1 = R2 = 1):")
for lam in (0.2, 0.5, 0.8):
```

(continues on next page)

(continued from previous page)

```
    g, h = grad_hess(kw, kw.s_star, lam, 1.0, 1.0)
    print(f"  λ = {lam}:  ∂U/∂λ = {g[1] + 0.0:.1f},  ∂²U/∂λ² = {h[1, 1] + 0.0:.1f},  "
          f"e = {e_of_λ(kw, lam):.3f}")
```

```
end of experiment 1: R1 = 1.000000, R2 = 1.000000, gradient = [0. 0.]

λ derivatives at other rest points (s = s*, R1 = R2 = 1):
  λ = 0.2:  ∂U/∂λ = 0.0,  ∂²U/∂λ² = 0.0,  e = 3.333
  λ = 0.5:  ∂U/∂λ = 0.0,  ∂²U/∂λ² = 0.0,  e = 0.833
  λ = 0.8:  ∂U/∂λ = 0.0,  ∂²U/∂λ² = 0.0,  e = 0.208
```

The rest points of the learning dynamics *are* the rational expectations equilibria, exchange rate and all: a whole continuum of them.

One recursion, many learners explains how the limit of a learning recursion is governed by its linearization near a rest point: a rest point attracts when the linearized algorithm contracts in every direction.

A continuum of rest points rules that out in one direction, and shows up as an eigenvalue of exactly one.

Shift both classes' portfolio shares by the same amount and nothing pulls them back, because the shifted point is another rest point.

To see this, compute the Jacobian of two periods of the algorithm (one update of each class) at experiment 1's limit, by finite differences.

We hold the curvature estimate M at its limit, which loses nothing to first order, because the Newton step multiplies M^{-1} by a gradient that vanishes at the rest point.

```
def two_period_map(p, z, M, gain=0.3):
    "One update of each class's (s, λ), with the curvature estimate held at M."
    s, lam = z[:2].copy(), z[2:].copy()
    for jp in (1, 0):                                   # the odd class updates, then␣
↪the even
        j = 1 - jp
        R1 = lam[j]*s[j] / (lam[jp]*s[jp])              # p1 when jp was young / p1 now
        R2 = (1 - lam[j])*s[j] / ((1 - lam[jp])*s[jp])
        g, _ = grad_hess(p, s[jp], lam[jp], R1, R2)
        step = gain * np.linalg.solve(M, g)
        s[jp], lam[jp] = s[jp] - step[0], lam[jp] - step[1]
    return np.concatenate([s, lam])


def rest_point_jacobian(p, lam, κ=0.5, gain=0.3, ε=1e-6):
    "Jacobian of the two-period map at the rest point (s*, s*, λ, λ)."
    z = np.array([p.s_star, p.s_star, lam, lam])
    _, M = grad_hess(p, p.s_star, lam, 1.0, 1.0)      # Hessian at the rest point,
    M[1, 1] -= κ                                      # plus the ridge: the limit of M
    return np.column_stack([(two_period_map(p, z + ε*d, M, gain)
                             - two_period_map(p, z - ε*d, M, gain)) / (2*ε)
                            for d in np.eye(4)])


J = rest_point_jacobian(kw, lam_lim[0])
eigvals, eigvecs = np.linalg.eig(J)
order = np.argsort(-np.abs(eigvals))
labels = ["s even", "s odd", "λ even", "λ odd"]
print(pd.DataFrame(eigvecs[:, order].real, index=labels,
                   columns=[f"{v:.4f}" for v in eigvals[order].real]).round(3))
```

```
         1.0000  0.4014  0.2500  0.1600
s even    0.000   0.000   0.447   0.371
s odd     0.000   0.000   0.894   0.928
λ even    0.707  -0.535   0.000   0.000
λ odd     0.707   0.845   0.000   0.000
```

The columns are the eigenvectors, headed by their eigenvalues.

The eigenvalue of one belongs to the direction that moves both classes' λ together: the continuum.

The next eigenvalue belongs to a direction that moves the two classes' portfolio shares in opposite directions, and the two smallest belong to saving.

All three are below one, so saving returns to $s^\star$ and the classes' shares merge.

But the place on the continuum where they merge is whatever the stable directions left behind.

That is the dead hand of history, written as an eigenvalue.

The indeterminacy is not resolved; it is displaced into the initial conditions and the tuning of the algorithm.

Sargent's verdict is hedged but clear.

By tying down the expectations that the rational expectations equilibrium left free, he writes, the adaptive mechanism selects an exchange rate path, but

> This may seem a weak reed on which to base exchange rate determination.

A footnote adds:

> Put differently, a regime that allows the exchange rate to be history-dependent seems to be an ill-formed mechanism.

The ghost stays quiet here only because the Newton–Raphson learner stops experimenting once the returns are equal.

A learning rule that keeps experimenting can wake it, as the evidence below shows.

93.4 Evidence from the laboratory

Arifovic [1996] implemented this two-currency economy as an experiment with paying human subjects, using $w_1 = 11$, $w_2 = 1$, $H_1 = H_2 = 10$, so that $s^\star = 5$.

These are not the parameters of the Newton–Raphson section above; Exercise 93.7.4 runs the Newton–Raphson learner at Arifovic's parameters.

Each young subject chose a saving rate and a fraction to allocate between the two currencies; the experimenters cleared the two money markets against those choices, exactly as in (93.2).

Arifovic split her subjects into two groups of equal size, one young in odd periods and the other in even periods.

These alternating groups are the laboratory counterpart of the odd and even classes of our learners.

The experiment ended in a period that the subjects did not know in advance.

The money held by that last young generation was redeemed at the average of the previous two periods' price levels.

The result did not look like the Newton–Raphson simulations at all.

The exchange rate fluctuated persistently, roughly within a band from 0.5 to 2, and showed no sign of settling down.

If anything the amplitude grew between the second and third sessions.

The simple Newton–Raphson model, which at Sargent's parameters converges to a constant, does a poor job of matching that.

So Arifovic built a different model of the same economy, one in which the agents are a **population** bred by a genetic algorithm.

93.5 A genetic algorithm economy

Networks, Genetic Algorithms, and Classifier Systems sets out Holland's genetic algorithm in general; here we need only the version Arifovic used.

Now each class is a population of $N = 30$ agents, and each agent is a **binary string** of length 30: the first 20 bits encode its saving rate, the last 10 its portfolio share.

Every period the young population's strings are decoded into (s_i, λ_i) pairs, the two money markets clear against the aggregates via (93.2), and the exchange rate is read off.

A generation later, when the cohort is old, each string's realized utility (93.1) is its **fitness**, and the genetic algorithm breeds the next generation.

```
N, L_s, L_lam = 30, 20, 10          # population size; bits for saving and portfolio

arifovic = Params(w1=11.0, w2=1.0, H1=10.0, H2=10.0)

def decode(p, pop):
    "Binary strings → (s, λ), with s in (0, w1) and λ in [0.02, 0.98]."
    ints_s = pop[:, :L_s] @ (1 << np.arange(L_s)[::-1])
    ints_l = pop[:, L_s:] @ (1 << np.arange(L_lam)[::-1])
    s = 0.05 + (p.w1 - 0.10) * ints_s / (2**L_s - 1)
    lam = 0.02 + 0.96 * ints_l / (2**L_lam - 1)
    return s, lam

def fitness(p, pop, R1, R2):
    "Realized lifetime utility of each string given the returns it experienced."
    s, lam = decode(p, pop)
    c2 = p.w2 + s*(lam*R1 + (1 - lam)*R2)
    return np.log(p.w1 - s) + np.log(np.maximum(c2, 1e-9))
```

The genetic algorithm has the classic three operators — roulette-wheel **selection** of parents, single-point **crossover**, and bit-flip **mutation** — plus one that Arifovic introduced, the **election operator**: a child is admitted to the next generation only if it would have been fitter than its parent on the most recent returns, otherwise the parent survives.

Some details of our implementation are this lecture's own choices, not necessarily Arifovic's:

- a mutation probability of 0.033 per bit;
- crossover applied to every pair of parents;
- a roulette wheel whose weights are fitness minus the population's minimum fitness (a "windowed" wheel, needed because log utility can be negative);
- the decoding ranges for s and λ.

```
def genetic_step(p, pop, R1, R2, rng, p_mut=0.033, election=True):
    "Breed one new generation from the current one."
    fit = fitness(p, pop, R1, R2)
    weight = fit - fit.min() + 1e-6                  # windowed roulette wheel
    new = np.empty_like(pop)
```

(continues on next page)

(continued from previous page)

```
    for k in range(0, N, 2):
        i, j = np.searchsorted(np.cumsum(weight), rng.random(2) * weight.sum())
        parents = np.array([pop[i], pop[j]])
        cut = rng.integers(1, L_s + L_lam)           # single-point crossover
        kids = parents.copy()
        kids[0, cut:], kids[1, cut:] = parents[1, cut:], parents[0, cut:]
        for c in kids:                               # mutation
            c[rng.random(L_s + L_lam) < p_mut] ^= 1
        if election:                                 # Arifovic's election operator
            f_kids = fitness(p, kids, R1, R2)
            f_par = fitness(p, parents, R1, R2)
            for m in range(2):
                new[k + m] = kids[m] if f_kids[m] > f_par[m] else parents[m]
        else:
            new[k], new[k + 1] = kids
    return new


def genetic_economy(p, T=3000, seed=0, election=True):
    """
    Run the genetic-algorithm economy.  Returns the exchange rate e_t and, for
    the population that is young at t, the mean and standard deviation of s
    and of λ across its strings.
    """
    rng = np.random.default_rng(seed)
    pops = [rng.integers(0, 2, (N, L_s + L_lam)) for _ in range(2)]   # even, odd
    p1 = np.empty(T)
    p2 = np.empty(T)
    stats = np.empty((T, 4))                         # mean s, sd s, mean λ, sd λ
    for t in range(T):
        j = t % 2
        s, lam = decode(p, pops[j])
        p1[t], p2[t] = p.H1 / np.sum(lam*s), p.H2 / np.sum((1 - lam)*s)
        stats[t] = s.mean(), s.std(), lam.mean(), lam.std()
        if t >= 1:
            jp = 1 - j
            R1, R2 = p1[t-1]/p1[t], p2[t-1]/p2[t]
            pops[jp] = genetic_step(p, pops[jp], R1, R2, rng, election=election)
    return p1 / p2, stats
```

93.5.1 Volatility that does not die out

Let's run the economy for 4000 periods and drop the first 500 as a burn-in.

Because every string's λ lies in $[0.02, 0.98]$ and $H_1 = H_2$, the exchange rate is confined to $[1/49, 49]$, so $|\log e_t| \leq \ln 49 \approx 3.89$.

The dotted lines in the figure mark those bounds, and the shaded band marks the experimental range $[0.5, 2]$.

```
burn = 500
e_ga, stats_ga = genetic_economy(arifovic, T=4000, seed=1)
t_ga = np.arange(burn, len(e_ga))
log_e = np.log(e_ga[burn:])
bound = np.log(0.98 / 0.02)

fig, ax = plt.subplots(figsize=(10, 4))
```

(continues on next page)

(continued from previous page)

```
ax.plot(t_ga, log_e, lw=0.5)
ax.axhspan(np.log(0.5), np.log(2), color='C1', alpha=0.2, label="experimental range")
for b in (bound, -bound):
    ax.axhline(b, color='k', ls=':', lw=0.8)
ax.set_xlabel("$t$")
ax.set_ylabel(r"$\log e_t$")
ax.legend(frameon=False, loc='upper right')
plt.show()
```

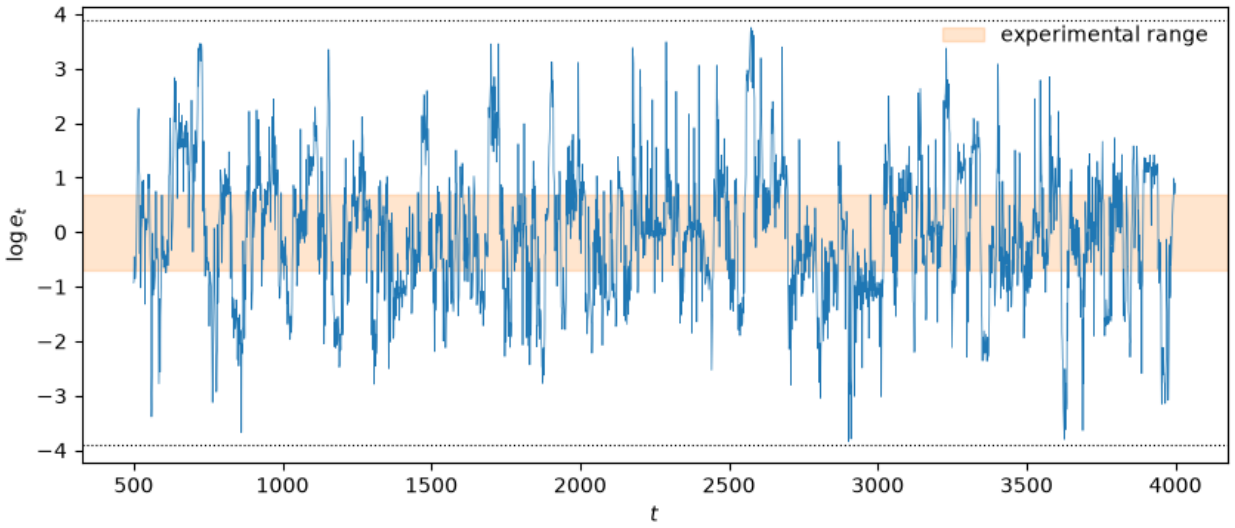


Fig. 93.3: Exchange rate in the genetic-algorithm economy

The exchange rate wanders and keeps wandering, and it does so far outside the experimental range.

```
in_band = np.mean((e_ga[burn:] > 0.5) & (e_ga[burn:] < 2))
print(f"e ranges from {e_ga[burn:].min():.3f} to {e_ga[burn:].max():.1f}; "
      f"share of periods in [0.5, 2]: {in_band:.2f}")
print(f"largest |log e| = {np.abs(log_e).max():.2f}, against the bound {bound:.2f}")
print(f"standard deviation of log e, t = 500-1499: {np.log(e_ga[500:1500]).std():.3f}
↪")
print(f"standard deviation of log e, t = 3000-3999: {np.log(e_ga[3000:]).std():.3f}")
```

```
e ranges from 0.021 to 42.8; share of periods in [0.5, 2]: 0.43
largest |log e| = 3.85, against the bound 3.89
standard deviation of log e, t = 500-1499: 1.279
standard deviation of log e, t = 3000-3999: 1.194
```

The volatility is a permanent feature, not a transient: it is about as large late in the run as early.

It is also bigger than in the laboratory.

Fewer than half the periods fall inside the experimental band, and the series comes close to the bounds that the encoding imposes.

Sargent noted the same thing about Arifovic's genetic economy, whose exchange rates were "more variable than those found in Arifovic's experiments."

93.5.2 What the population learns

Unlike the Newton–Raphson learner, the genetic population never settles down completely, since mutation keeps injecting new strings.

But it does learn, and what it learns and what it cannot learn line up with the stable and the neutral directions of *The ghost of indeterminacy*.

To see this, look at the young population's mean and spread of s and of λ over a stretch of the run.

```
window = np.arange(2000, 2300)
fig, axes = plt.subplots(2, 1, figsize=(10, 6), sharex=True)
for ax, (m, sd), name in zip(axes, [(0, 1), (2, 3)], ["saving $s$", r"portfolio share
↪$\lambda$"]):
    mean, spread = stats_ga[window, m], stats_ga[window, sd]
    ax.plot(window, mean, lw=1, label="population mean")
    ax.fill_between(window, mean - spread, mean + spread, alpha=0.3,
                    label="± within-population s.d.")
    ax.set_ylabel(name)
axes[0].axhline(arifovic.s_star, color='k', ls='--', lw=0.8, label="$s^\\star$")
axes[0].legend(frameon=False, ncol=3, loc='upper right')
axes[1].set_ylim(0, 1)
axes[1].set_xlabel("$t$")
plt.tight_layout()
plt.show()
```

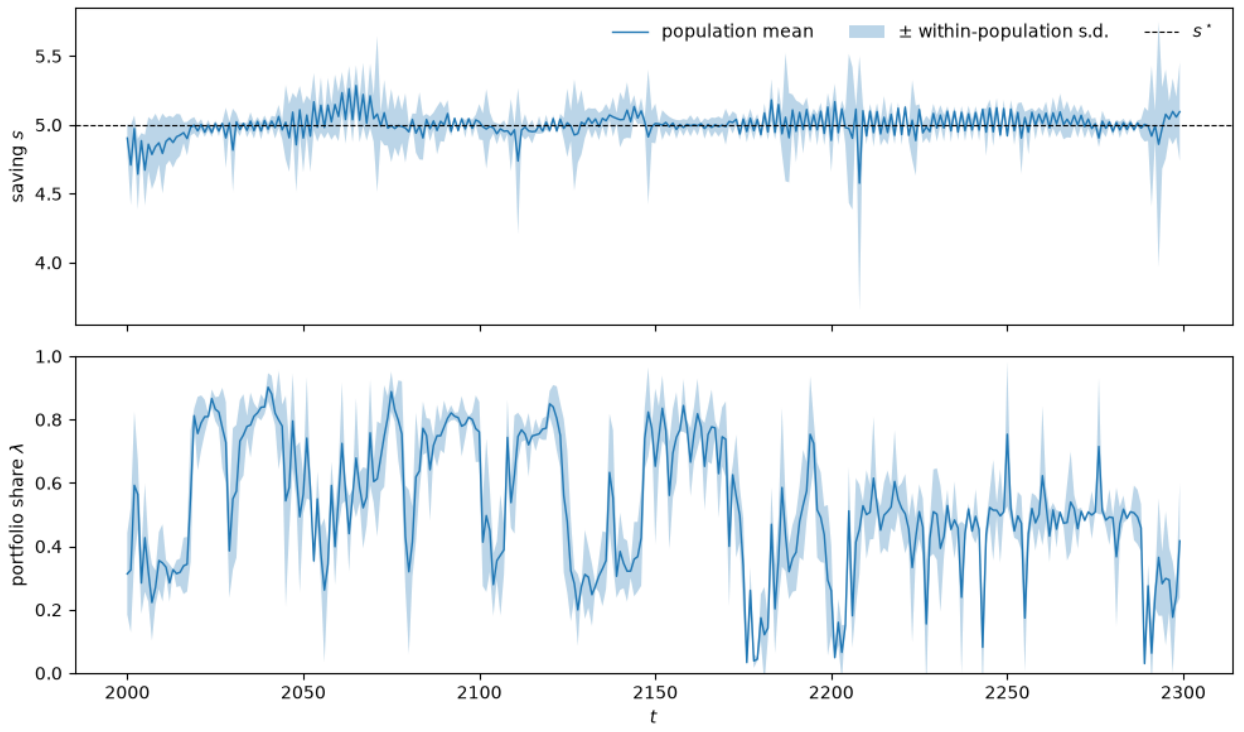


Fig. 93.4: Saving and portfolio shares inside the genetic population

Here are the same statistics averaged over the whole run after the burn-in.

```
S = stats_ga[burn:]
print(f"mean saving: {S[:, 0].mean():.3f}   (s* = {arifovic.s_star})")
print(f"within-population s.d. of s: {S[:, 1].mean():.3f}")
print(f"within-population s.d. of λ: {S[:, 3].mean():.3f}")
print(f"s.d. over time of the population's mean λ: {S[:, 2].std():.3f}  "
      f"(a uniform draw on [0.02, 0.98] has s.d. {0.96/np.sqrt(12):.3f})")
print(f"5th and 95th percentiles of the population's mean λ: "
      f"{np.percentile(S[:, 2], 5):.2f}, {np.percentile(S[:, 2], 95):.2f}")
```

```
mean saving: 4.990   (s* = 5.0)
within-population s.d. of s: 0.110
within-population s.d. of λ: 0.092
s.d. over time of the population's mean λ: 0.237  (a uniform draw on [0.02, 0.98]␣
↪has s.d. 0.277)
5th and 95th percentiles of the population's mean λ: 0.11, 0.88
```

The genetic economy settles in one direction and drifts in the other.

In the determinate direction, saving, the population learns: mean saving sits at $s^\star$ and the strings agree closely with one another.

The strings also agree fairly closely about λ at any moment.

But the population's common λ wanders over most of the range the encoding allows, and the exchange rate wanders with it.

The reason is the ghost of indeterminacy.

Nothing in the fundamentals favors one portfolio share over another: at equal returns every share is equally fit.

A share becomes fitter than another only when the two currencies' returns differ, and they differ only because one cohort's portfolio differs from the previous cohort's.

So fitness is neutral in λ apart from return gaps that the population creates itself, and nothing anchors the population's portfolio.

Mutation, crossover, and those self-generated return gaps push it along the continuum of equilibria, and nothing holds it in place.

The indeterminacy that the Newton–Raphson learner froze into a history-dependent constant returns here as volatility.

The election operator does much of the population's learning; Exercise 93.7.2 shows what happens without it.

93.5.3 The shape of the volatility

Sargent describes the exchange rate in Arifovic's genetic economy as close to a random walk — but with **mean reversion**, showing up as a dip at zero frequency in the spectrum of its first difference.

We plot the spectrum of the winsorized $\Delta \log e_t$ against frequency in radians on a log scale, as Sargent does, together with the autocorrelation of $\log e_t$.

```
def averaged_spectrum(x, n_seg=16):
    "Bartlett-averaged periodogram; frequencies in radians per period."
    x = x - x.mean()
    seg = len(x) // n_seg
    acc = np.zeros(seg // 2 + 1)
    for k in range(n_seg):
        acc += np.abs(np.fft.rfft(x[k*seg:(k+1)*seg]))**2 / seg
    return 2 * np.pi * np.fft.rfftfreq(seg), acc / n_seg
```

(continues on next page)

(continued from previous page)

```
d_log_e = np.diff(log_e)
d_log_e = np.clip(d_log_e, *np.percentile(d_log_e, [1, 99]))     # winsorize tails
freq, spec = averaged_spectrum(d_log_e)

def acf(x, K):
    x = x - x.mean()
    return np.array([np.sum(x[k:]*x[:len(x)-k]) / np.sum(x*x) for k in range(K)])


fig, axes = plt.subplots(1, 2, figsize=(12, 4))
axes[0].semilogy(freq[1:], spec[1:], lw=1.2)
axes[0].set_title(r"spectrum of $\Delta \log e$")
axes[0].set_xlabel("frequency (radians)")

axes[1].bar(range(25), acf(log_e, 25))
axes[1].set_title(r"autocorrelation of $\log e$")
axes[1].set_xlabel("lag")
plt.tight_layout()
plt.show()
```

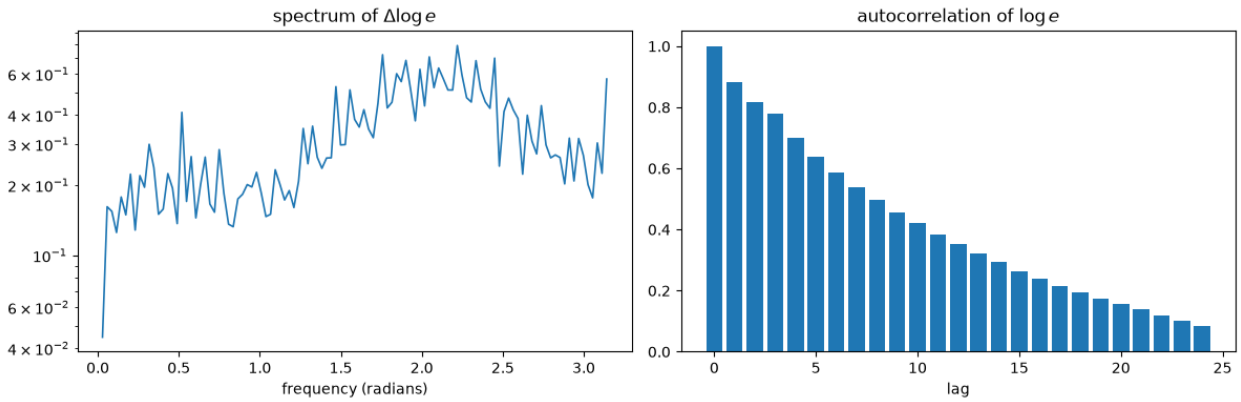


Fig. 93.5: Spectrum and autocorrelation of the exchange rate

The spectrum is hump-shaped.

It is low near zero frequency, rises to a peak at intermediate frequencies, and falls off toward π, much like the spectrum in Sargent's figure, which peaks near 1.5 radians.

The autocorrelation of the level falls to one half by about lag 8 and is close to zero by lag 25.

That is much faster than a near unit root would allow.

But the sample autocorrelation is only an informal guide, because a random walk's sample autocorrelation also decays in a finite sample.

A direct test is to regress $\log e_t$ on a constant and $\log e_{t-1}$ and compare the coefficient $\hat{\phi}$ with one.

The Dickey–Fuller statistic $(\hat{\phi} - 1)/\mathrm{se}(\hat{\phi})$ rejects a unit root at the 5% level when it is below about -2.86.

```
def ar1(x):
    "OLS of x_t on (1, x_{t-1}): returns ψ̂ and the Dickey-Fuller statistic."
    X = np.column_stack([np.ones(len(x) - 1), x[:-1]])
    b = np.linalg.lstsq(X, x[1:], rcond=None)[0]
    u = x[1:] - X @ b
    se = np.sqrt(u @ u / (len(u) - 2) * np.linalg.inv(X.T @ X)[1, 1])
    return b[1], (b[1] - 1) / se


band = slice(1, len(spec)//2)
ψ, df_stat = ar1(log_e)
print(f"spectrum of Δlog e at the lowest frequency / mean over (0, π/2): "
      f"{spec[1] / spec[band].mean():.2f}")
print(f"spectral peak at frequency {freq[1:][np.argmax(spec[1:])]:.2f} radians")
print(f"AR(1) coefficient ψ = {ψ:.3f},  half-life = {np.log(0.5)/np.log(ψ):.1f}␣
↪periods")
print(f"Dickey-Fuller statistic = {df_stat:.1f}   (5% critical value ≈ -2.86)")
```

```
spectrum of Δlog e at the lowest frequency / mean over (0, π/2): 0.21
spectral peak at frequency 2.22 radians
AR(1) coefficient ψ = 0.882,  half-life = 5.5 periods
Dickey-Fuller statistic = -14.8   (5% critical value ≈ -2.86)
```

The dip at zero frequency is pronounced.

The spectrum peaks near 2.2 radians, a cycle of about three periods.

The level is strongly mean-reverting, with a half-life of about five and a half periods, and the Dickey–Fuller test rejects a unit root decisively; Exercise 93.7.3 checks other seeds.

So in our implementation the genetic exchange rate is a stationary series, not a near random walk.

Part of the mean reversion is mechanical: the encoding's bounds on λ cap $|\log e_t|$, so the series cannot wander off for good.

Sargent's assessment is that real floating exchange rates for hard-currency pairs have spectra of log differences that look much like Arifovic's, except *without* the dip at zero frequency.

Actual exchange rates are even closer to pure random walks: it is hard to find a model that forecasts them better than "no change," a benchmark that the neural-network forecasts of Kuan and Liu [1995] set out to beat.

So the genetic economy matches the higher-frequency texture of floating rates but not their lowest frequencies.

That it gets that far with no fundamental shocks at all, from a population of adapting agents repricing two intrinsically identical currencies, is the striking part.

93.6 Concluding remarks

Two models of the same indeterminate economy gave two very different verdicts.

The Newton–Raphson learner **converges** — and in converging, exposes the indeterminacy rather than curing it.

Its rest points satisfy the arbitrage condition that leaves the exchange rate free.

They form a continuum, along which the algorithm has a unit eigenvalue.

So the limit is wherever the transient leaves the economy, and it depends on the initial portfolios and saving rates and on the tuning of the algorithm.

Learning makes the exchange rate history-dependent; it does not select one on any economic ground.

Sargent's hedged verdict is that such a mechanism "may seem a weak reed."

The genetic-algorithm economy **does not converge**, but not because it fails to learn.

With Arifovic's election operator, the population learns a saving rate close to $s^\star$ and becomes nearly homogeneous.

What it cannot learn is a portfolio share, because at equal returns no share is fitter than any other.

So the population's common λ wanders along the continuum of equilibria, and the exchange rate wanders with it.

The ghost of indeterminacy returns as volatility, in an economy with no fundamental disturbances whatsoever.

That volatility is an instance of one of the payoffs that *What the program is good for* lists for the bounded-rationality program: a new source of dynamics.

It resembles the higher-frequency texture of floating exchange rates but not their low frequencies, where the genetic economy's strong mean reversion has no counterpart in the data.

The two halves of the lecture also use different calibrations.

Exercise 93.7.4 runs the Newton–Raphson learner at Arifovic's, and finds that it converges there only if it is damped enough.

With this lecture's damping it sits on the edge of stability, and it generates recurrent bursts of volatility of its own.

Which learning model an economist reaches for, and how she tunes it, is one more of the choices that the bounded-rationality program forces into the open.

Networks, Genetic Algorithms, and Classifier Systems sets out Holland's genetic algorithm in general, together with its richer cousin, the classifier system.

Money as a Medium of Exchange among Artificially Intelligent Agents then puts populations of classifier systems to work: they must not merely tune a portfolio but *discover from scratch* which good will serve as money.

93.7 Exercises

Exercise 93.7.1

The main text varied the initial portfolio shares.

This exercise asks whether anything else matters for the limit.

1. Hold the initial portfolio shares fixed at $\lambda_0 = (0.30, 0.50)$ for the (even, odd) classes, and vary the initial saving rates s_0 over a few values. Confirm that the classes' shares merge and saving converges to $s^\star$. Does the limiting exchange rate change with s_0?
2. Start from experiment 1's initial condition and vary the ridge κ over $(0.1, 0.2, 0.5, 1, 2)$, holding $\gamma = 0.3$. Then vary the gain γ over $(0.05, 0.1, 0.3)$, holding $\kappa = 0.5$. Record the limiting exchange rate each time.

Is the limiting exchange rate a function of the initial portfolio shares alone?

Solution

Here is part 1.

```
rows = []
for s0 in ([3.0, 2.2], [2.2, 3.0], [4.0, 4.0], [1.0, 1.0], [6.0, 2.0]):
    e, s_h, l_h = newton_learning(kw, s0, [0.30, 0.50], T=800)
    rows.append([str(s0), l_h[-1][0], l_h[-1][1], s_h[-1].mean(), e[-1]])
```

```
pd.DataFrame(rows, columns=["initial s", "λ even (limit)", "λ odd (limit)",
                            "s (limit)", "e (limit)"]).round(4)
```

	initial s	λ even (limit)	λ odd (limit)	s (limit)	e (limit)
0	[3.0, 2.2]	0.4117	0.4117	2.5	1.1908
1	[2.2, 3.0]	0.4471	0.4471	2.5	1.0304
2	[4.0, 4.0]	0.4444	0.4444	2.5	1.0418
3	[1.0, 1.0]	0.4145	0.4145	2.5	1.1773
4	[6.0, 2.0]	0.3772	0.3772	2.5	1.3759

The two classes' shares merge (columns 2 and 3 agree), and saving converges to $s^\star = 2.5$.

But the limiting exchange rate changes with the initial saving rates, even though the initial portfolio shares are the same in every row.

The λ-gradient $s(R_1 - R_2)/c_2$ is proportional to saving, and the returns during the transient depend on both classes' saving.

So saving's path to $s^\star$ shapes how far λ moves before the arbitrage condition stops it.

Here is part 2.

```
rows = []
for κ in (0.1, 0.2, 0.5, 1.0, 2.0):
    e, s_h, l_h = newton_learning(kw, [3.5, 2.0], [0.35, 0.40], T=2000, κ=κ)
    rows.append([κ, 0.3, l_h[-1][0], e[-1]])
for gain in (0.05, 0.1, 0.3):
    e, s_h, l_h = newton_learning(kw, [3.5, 2.0], [0.35, 0.40], T=4000, gain=gain)
    rows.append([0.5, gain, l_h[-1][0], e[-1]])

pd.DataFrame(rows, columns=["κ", "gain γ", "λ (limit)", "e (limit)"]).round(4)
```

	κ	gain γ	λ (limit)	e (limit)
0	0.1	0.30	0.5445	0.6970
1	0.2	0.30	0.3859	1.3263
2	0.5	0.30	0.3732	1.3998
3	1.0	0.30	0.3709	1.4136
4	2.0	0.30	0.3707	1.4145
5	0.5	0.05	0.3694	1.4225
6	0.5	0.10	0.3701	1.4185
7	0.5	0.30	0.3732	1.3998

From the same initial condition, the limit moves with both κ and γ.

The ridge and the gain set how far λ travels in response to a given transient return gap, and so where on the continuum the motion stops.

The effect is modest for moderate values but large for small κ, where the λ steps are big and the transient is violent.

So the limiting exchange rate is not a function of the initial portfolio shares alone.

It depends on the whole initial condition and on the tuning of the algorithm, which makes the weak reed weaker still: two economists with identical data on initial beliefs, but different learning algorithms, would predict different exchange rates.

Exercise 93.7.2

Arifovic's **election operator** admits a child to the next generation only if it beats its parent on the latest returns.

Investigate what the operator does.

Run the genetic economy with and without it for several seeds, and report for each case:

- the standard deviation of $\log e_t$;
- the population's mean saving, compared with $s^\star$;
- the within-population standard deviations of s and of λ;
- the standard deviation over time of the population's mean λ.

Which way does the operator push exchange rate volatility, and why?

Solution

Here are the six seeds with and without the operator.

```
rows = []
for election in (True, False):
    for sd in range(6):
        e, st = genetic_economy(arifovic, T=2500, seed=sd, election=election)
        st = st[500:]
        rows.append([election, sd, np.log(e[500:]).std(), st[:, 0].mean(),
                     st[:, 1].mean(), st[:, 3].mean(), st[:, 2].std()])

table = pd.DataFrame(rows, columns=["election", "seed", "s.d. log e", "mean s",
                                    "within s.d. s", "within s.d. λ", "s.d. of␣
↪mean λ"])
table.drop(columns="seed").groupby("election").agg(["min", "max"]).T.round(3)
```

```
election             False  True
s.d. log e     min   0.473  1.081
               max   0.553  1.350
mean s         min   4.535  4.973
               max   5.399  5.286
within s.d. s  min   1.327  0.088
               max   1.358  0.110
within s.d. λ  min   0.229  0.083
               max   0.238  0.095
s.d. of mean λ min   0.111  0.214
               max   0.129  0.255
```

The table reports, for each statistic, its smallest and largest value across the six seeds.

The election operator **raises** exchange rate volatility, and the ranges across seeds do not overlap.

The other rows show why.

Without the operator the genetic algorithm learns much less.

Every child is admitted however unfit, so mutation and crossover keep the population a broad cloud.

Saving is spread widely within the population, and depending on the seed its mean lies anywhere between about 4.5 and 5.4, against $s^\star = 5$.

Portfolio shares are almost as dispersed as a uniform draw on $[0.02, 0.98]$, whose standard deviation is 0.277.

The aggregates in (93.2) average over that diverse population, so the aggregate portfolio share, and with it the exchange rate, moves comparatively little.

The lower volatility comes mainly from averaging over that cloud.

With the operator, selection works.

Mean saving stays within about six percent of $s^\star$ in every seed, and the population becomes nearly homogeneous in both s and λ.

But a homogeneous population whose fitness is flat in λ moves its portfolio share together, and nothing anchors where it goes.

So the population's mean λ wanders much more, and the exchange rate with it.

The operator that makes the population learn is the same one that makes the indeterminacy visible: learning the determinate part of the equilibrium exposes the indeterminate part.

Exercise 93.7.3

The lecture reported that the genetic economy's exchange rate is strongly mean-reverting.

Check that the finding does not depend on the seed, and check the diagnostic.

1. For six seeds, run the genetic economy for 4000 periods, drop 500 as a burn-in, and estimate ϕ in $\log e_t = \mu + \phi \log e_{t-1} + \varepsilon_t$. Report ϕ, the half-life $\ln(1/2)/\ln\phi$, and the Dickey–Fuller statistic. Is the series stationary?
2. Simulate a pure random walk of the same length and compute its sample autocorrelations at lags 1, 5, 10, and 25. Can a decaying sample autocorrelation distinguish a mean-reverting series from a random walk?

Solution

Here is part 1.

```
rows = []
for sd in range(6):
    le = np.log(genetic_economy(arifovic, T=4000, seed=sd)[0][500:])
    ψ_sd, df_sd = ar1(le)
    rows.append([sd, ψ_sd, np.log(0.5)/np.log(ψ_sd), df_sd])

pd.DataFrame(rows, columns=["seed", "ψ", "half-life", "Dickey-Fuller"]).round(3)
```

	seed	ψ	half-life	Dickey-Fuller
0	0	0.903	6.819	-13.316
1	1	0.882	5.530	-14.794
2	2	0.893	6.118	-14.071
3	3	0.876	5.244	-15.184
4	4	0.884	5.632	-14.658
5	5	0.896	6.337	-13.814

Across seeds ϕ is well below one, the half-life is a handful of periods, and the Dickey–Fuller statistic is far below the 5% critical value of about -2.86.

The series is stationary, and strongly mean-reverting.

Here is part 2.

```
rw = np.cumsum(np.random.default_rng(0).standard_normal(3500))
print("sample autocorrelation of a random walk at lags 1, 5, 10, 25:",
      acf(rw, 26)[[1, 5, 10, 25]].round(3))
print("sample autocorrelation of the GA log e at lags 1, 5, 10, 25:   ",
      acf(log_e, 26)[[1, 5, 10, 25]].round(3))
```

```
sample autocorrelation of a random walk at lags 1, 5, 10, 25: [0.999 0.995 0.
↪991 0.978]
sample autocorrelation of the GA log e at lags 1, 5, 10, 25:    [0.882 0.639 0.
↪42  0.067]
```

The random walk's sample autocorrelation decays too, because in a finite sample the autocorrelation is computed around the sample mean.

So "it decays" is not by itself evidence against a unit root.

What distinguishes the two series is *how fast* the autocorrelation decays, and the Dickey–Fuller statistic turns that into a formal test.

The genetic economy's autocorrelation falls nearly to zero within 25 lags, while the random walk's barely moves.

Exercise 93.7.4

The Newton–Raphson section used Sargent's parameters $(w_1, w_2, H_1, H_2) = (20, 15, 100, 120)$, while the genetic-algorithm section used Arifovic's $(11, 1, 10, 10)$.

To compare like with like, run the Newton–Raphson learner at Arifovic's parameters.

1. With the lecture's tuning ($\kappa = 0.5$, $\gamma = 0.3$), run it for 3000 periods from $s_0 = (6, 4)$, $\lambda_0 = (0.45, 0.55)$. Does it converge?
2. Linearize around a rest point $(s^\star, s^\star, \lambda, \lambda)$. Show that each update multiplies the difference between the two classes' portfolio shares by $1 - a$, where

$$a = \frac{\gamma\, s^\star}{\kappa\, c_2^\star\, \lambda(1-\lambda)}, \qquad c_2^\star = w_2 + s^\star.$$

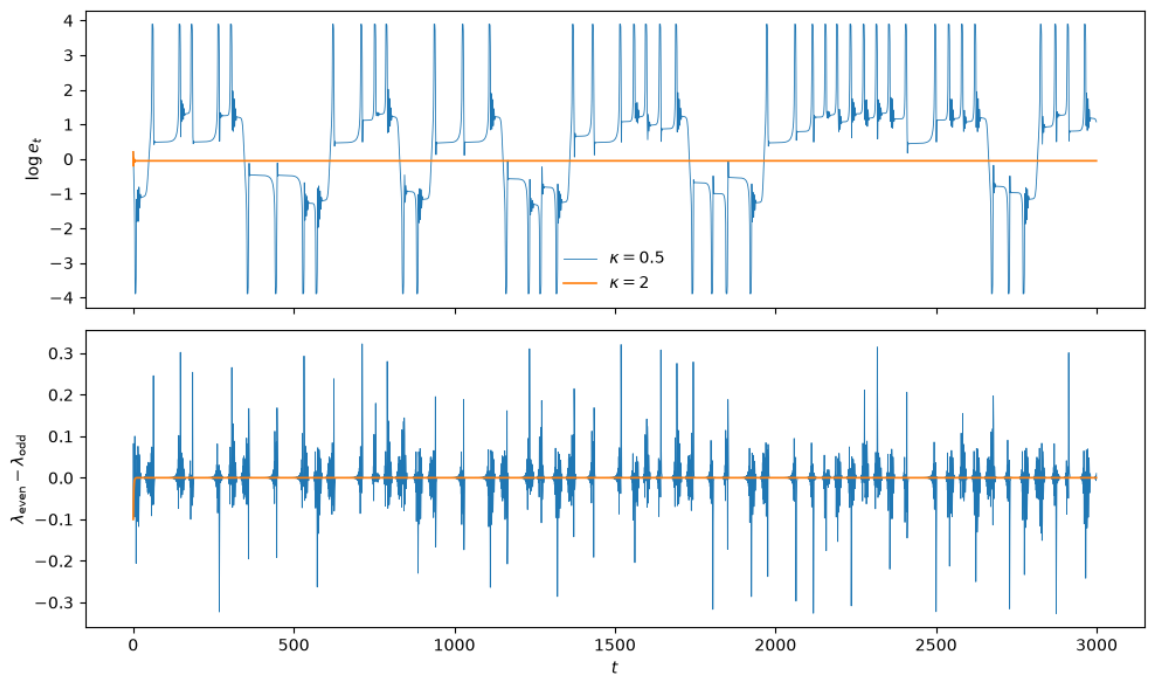
   Hint: the portfolio return $A = \lambda R_1 + (1-\lambda) R_2$ does not depend on the portfolio shares at all, and near a rest point $R_1 - R_2 \approx (\lambda_{\text{young}} - \lambda_{\text{old}})/(\lambda(1-\lambda))$. Compute a at both calibrations, and check the result against `rest_point_jacobian`.
3. How large must κ be for some rest points to be stable at Arifovic's parameters? Rerun part 1 with $\kappa = 2$.

Solution

Here is part 1.

```
ea, sa, la = newton_learning(arifovic, [6.0, 4.0], [0.45, 0.55], T=3000)
eb, sb, lb = newton_learning(arifovic, [6.0, 4.0], [0.45, 0.55], T=3000, κ=2.0)

fig, axes = plt.subplots(2, 1, figsize=(10, 6), sharex=True)
axes[0].plot(np.log(ea), lw=0.6, label=r"$\kappa = 0.5$")
axes[0].plot(np.log(eb), lw=1.2, label=r"$\kappa = 2$")
axes[0].set_ylabel(r"$\log e_t$")
axes[0].legend(frameon=False)
axes[1].plot(la[:, 0] - la[:, 1], lw=0.6)
axes[1].plot(lb[:, 0] - lb[:, 1], lw=1.2)
axes[1].set_ylabel(r"$\lambda_{\mathrm{even}} - \lambda_{\mathrm{odd}}$")
axes[1].set_xlabel("$t$")
plt.tight_layout()
plt.show()
```



With the lecture's tuning the learner does not converge.

The two classes' portfolio shares nearly merge, saving reaches $s^\star = 5$, and the exchange rate rests for a while.

Then the gap between the classes' shares grows in an alternating pattern until it explodes.

The exchange rate shoots to the bounds and comes back to rest somewhere else, again and again.

For part 2, let δ be the young class's share minus the old class's.

The portfolio return is $A = \lambda_{\text{old}} R_1 + (1 - \lambda_{\text{old}}) R_2 = s_{\text{young}} / s_{\text{old}}$, which does not involve the shares.

So the saving equations do not involve λ at all.

Near a rest point, the cross term $\partial^2 U / \partial s \partial \lambda$ is proportional to $R_1 - R_2$ and vanishes.

The curvature estimate's λ entry is then $-\kappa$, and the old class's update is

$$\lambda_{\text{old}} \leftarrow \lambda_{\text{old}} + \frac{\gamma}{\kappa} \frac{s^\star}{c_2^\star} (R_1 - R_2) \approx \lambda_{\text{old}} + a\,\delta.$$

Next period the classes swap roles, so the new young-minus-old difference is $-(1 - a)\delta$.

Measured as even minus odd, the gap is multiplied by $1 - a$ at each update, and by $(1 - a)^2$ over two periods.

That is the eigenvalue we met in *The ghost of indeterminacy*.

The gap dies out if $0 < a < 2$ and explodes if $a > 2$.

```
def a_coef(p, lam, κ=0.5, gain=0.3):
    return gain * p.s_star / (κ * (p.w2 + p.s_star) * lam * (1 - lam))


for p, name in [(kw, "Sargent"), (arifovic, "Arifovic")]:
    for lam in (0.3, 0.5, 0.7):
        a = a_coef(p, lam)
        ev = np.sort(np.abs(np.linalg.eigvals(rest_point_jacobian(p, lam))))[::-1]
        print(f"{name:8s} λ = {lam}:  a = {a:.3f},  (1-a)² = {(1 - a)**2:.3f},  "
              f"|eigenvalues| = {ev.round(3)}")
```

```
Sargent  λ = 0.3:  a = 0.408,  (1-a)² = 0.350,  |eigenvalues| = [1.   0.35 0.25␣
↪0.16]
Sargent  λ = 0.5:  a = 0.343,  (1-a)² = 0.432,  |eigenvalues| = [1.    0.432 0.
↪25  0.16 ]
Sargent  λ = 0.7:  a = 0.408,  (1-a)² = 0.350,  |eigenvalues| = [1.   0.35 0.25␣
↪0.16]
Arifovic λ = 0.3:  a = 2.381,  (1-a)² = 1.907,  |eigenvalues| = [1.907 1.    0.
↪695 0.646]
Arifovic λ = 0.5:  a = 2.000,  (1-a)² = 1.000,  |eigenvalues| = [1.    1.    0.
↪695 0.646]
Arifovic λ = 0.7:  a = 2.381,  (1-a)² = 1.907,  |eigenvalues| = [1.907 1.    0.
↪695 0.646]
```

The formula matches the numerical Jacobian.

At Sargent's parameters a is far below 2 unless λ is below about 0.05 or above about 0.95, so the rest points are stable apart from the unit eigenvalue along the continuum.

At Arifovic's parameters the λ-gradient is much steeper, because the ratio $s^\star/c_2^\star$ is $5/6$ instead of $1/7$.

With $\kappa = 0.5$ and $\gamma = 0.3$, then, $a = 0.5/(\lambda(1-\lambda)) \geq 2$ at *every* rest point.

The rest point at $\lambda = 1/2$ is exactly on the edge, and all the others are unstable.

That is why the learner keeps escaping.

For part 3, $\lambda(1-\lambda) \leq 1/4$, so some rest points are stable if and only if $\kappa > 2\gamma s^\star/c_2^\star = 0.5$ at Arifovic's parameters.

The lecture's κ sits exactly on that boundary.

With $\kappa = 2$ the rest points with $\lambda(1-\lambda) > 1/16$, that is, with λ between about 0.07 and 0.93, are stable, and the figure shows the learner converging to a constant.

```
print(f"κ = 2: e → {eb[-1]:.4f},  s → {sb[-1].round(3)},  λ → {lb[-1].round(4)}")
print(f"largest |Δ log e| over the last 500 periods: "
      f"{np.abs(np.diff(np.log(eb[-500:]))).max():.1e}")
```

```
κ = 2: e → 0.9475,  s → [5. 5.],  λ → [0.5135 0.5135]
largest |Δ log e| over the last 500 periods: 0.0e+00
```

So the gap between the two halves of the lecture is about the adaptive machinery, not the calibration.

A Newton–Raphson learner that is damped enough converges at Arifovic's parameters too, to a history-dependent constant.

But how much damping is "enough" depends on the economy: with the same tuning, one calibration converges and the other bursts.

Exercise 93.7.5

In a footnote, Sargent distinguishes *history* dependence from *path* dependence.

In a stochastic version of the model, realizations that start from identical initial conditions can converge to different exchange rates, because different shocks hit the transient dynamics.

Add IID shocks to the old-age endowment, $w_{2t} = w_2 + \sigma\varepsilon_t$ with $\varepsilon_t \sim N(0,1)$ and $\sigma = 1$, at Sargent's parameters.

(The `w2_path` argument of `newton_learning` does this.)

1. Start 200 economies at the same initial condition, $s_0 = (3.0, 2.2)$ and $\lambda_0 = (0.30, 0.50)$, with different shock sequences, and run each for 600 periods. Does the exchange rate converge in each run? Does saving?

2. Plot a histogram of the limiting exchange rates, and mark the limit without shocks.
3. Why do the shocks stop affecting the exchange rate after a while, even though they never stop hitting the economy?

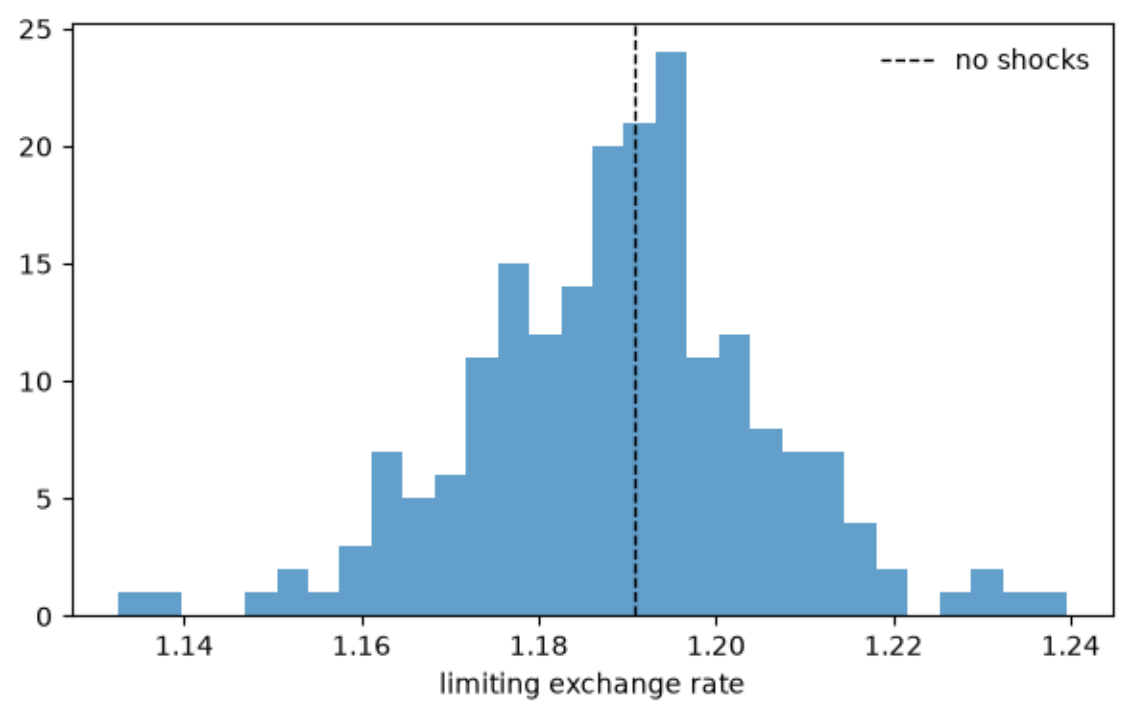
Solution

Here are parts 1 and 2.

```
T, σ = 600, 1.0
e_det = newton_learning(kw, [3.0, 2.2], [0.30, 0.50], T=T)[0][-1]

limits, last_move, s_spread = [], [], []
for seed in range(200):
    shocks = σ * np.random.default_rng(seed).standard_normal(T)
    e, s_h, l_h = newton_learning(kw, [3.0, 2.2], [0.30, 0.50], T=T,
                                  w2_path=kw.w2 + shocks)
    limits.append(e[-1])
    last_move.append(np.abs(np.diff(np.log(e[-300:]))).max())
    s_spread.append(s_h[-300:, 0].std())

fig, ax = plt.subplots(figsize=(7, 4))
ax.hist(limits, bins=30, alpha=0.7)
ax.axvline(e_det, color='k', ls='--', lw=1, label="no shocks")
ax.set_xlabel("limiting exchange rate")
ax.legend(frameon=False)
plt.show()
```



Here is a summary of the runs.

```
limits = np.array(limits)
print(f"limit without shocks: {e_det:.4f}")
print(f"limits with shocks: {limits.min():.4f} to {limits.max():.4f}, s.d. {limits.
 ↪std():.4f}")
print(f"largest |Δ log e| over the last 300 periods, across all runs: {max(last_
 ↪move):.1e}")
print(f"s.d. of even-class saving over the last 300 periods: "
      f"{min(s_spread):.3f} to {max(s_spread):.3f}")
```

```
limit without shocks: 1.1908
limits with shocks: 1.1326 to 1.2395, s.d. 0.0169
largest |Δ log e| over the last 300 periods, across all runs: 7.8e-16
s.d. of even-class saving over the last 300 periods: 0.193 to 0.274
```

Every exchange rate path converges: over the last 300 periods it does not move at all, up to rounding.

Saving does not converge; with a constant gain and shocks that never stop, it keeps fluctuating around $s^\star$.

Yet the economies, which differ only in their shocks, settle on different exchange rates.

That is path dependence in Sargent's sense.

Sargent [1993] quotes Evans and Honkapohja's description of how adaptive learning resolves indeterminacy:

> models with multiple solutions are converted into models with path dependence in which the trajectory of the economy, and the [rational expectations equilibrium] attained in the limit, are determined through a learning rule by initial forecasts and by the sequence of exogenous shocks.

The shocks matter only during the transient, because the exchange rate depends only on the young class's λ.

Once the two classes' shares merge, $R_1 = R_2$ whatever saving does, so the λ-gradient $s(R_1 - R_2)/c_2$ is zero whatever the endowment.

From then on, shocks to w_2 move saving but cannot move the portfolio share, and the exchange rate stays wherever the transient left it.

Because the shocks have only the brief transient in which to act, the spread of limits around the no-shock value is narrow.

CHAPTER

NINETYFOUR

NETWORKS, GENETIC ALGORITHMS, AND CLASSIFIER SYSTEMS

Contents

- *Networks, Genetic Algorithms, and Classifier Systems*
  - *Overview*
  - *The perceptron: an agent as a discriminant function*
  - *Associative memory: the Hopfield network*
  - *The genetic algorithm*
  - *Classifier systems*
  - *Evolutionary programming*
  - *Concluding remarks*
  - *Exercises*

94.1 Overview

The lectures so far gave adaptive agents a fairly econometric brain.

In *Adaptive Agents in an Overlapping Generations Monetary Economy* they ran stochastic approximation and recursive least squares, in *Learning, Approximation, and Equilibrium Computation* they fitted conditional expectations, and in *Exchange Rate Indeterminacy, Learning, and Experiments* they took Newton steps against realized utility.

Each of those learners was told what to estimate (a saving rate, a conditional expectation, a portfolio) and updated its estimate by a recursion of the kind described in *One recursion, many learners*.

Exchange Rate Indeterminacy, Learning, and Experiments also gave a first look at a learner of another kind: Arifovic's genetic-algorithm economy (*A genetic algorithm economy*), whose exchange rate never settles down.

This lecture sets out that algorithm and its relatives in general.

They make up the toolbox that Sargent [1993] draws from John Holland's work on adaptive systems [Holland, 1975, Holland *et al.*, 1986] and from the connectionist literature on neural networks.

Sargent's framing is that the bounded-rationality program asks us to populate models with agents who "behave like econometricians," and that the artificial-intelligence literature is a catalogue of candidate brains for them:

> It is from the store of methods built up in these literatures that we shall select the 'brains' to give our boundedly rational agents.

We tour four of them:

1. the **perceptron**, the simplest neural network, which turns out to be a linear discriminant function, a direct link to the "agents as econometricians" reading;
2. the **Hopfield network**, an associative memory that recovers stored patterns from corrupted inputs by rolling downhill on an energy function;
3. the **genetic algorithm**, Holland's population-based search, which we set loose on Axelrod's iterated prisoner's dilemma;
4. the **classifier system**, Holland's "brain as a competitive economy" of if-then rules, whose simplest instance, a two-armed bandit, already exposes a subtle limitation, and whose **bucket brigade** passes credit back along chains of decisions.

The first two are bridges from what came before.

The perceptron is a parametric rule that an econometrician recognizes on sight, and the Hopfield network is a memory rather than a rule.

The last two do something new: instead of tuning the coefficients of a fixed rule, they **discover rules**, searching a large space of possibilities with experience as the only guide.

We close with **evolutionary programming**: the idea that a population of adaptive agents, plodding toward an equilibrium, can be used to *compute* equilibria that we cannot solve for directly.

That is the idea *Money as a Medium of Exchange among Artificially Intelligent Agents* carries out in full, and this lecture builds the machinery it uses.

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
```

94.2 The perceptron: an agent as a discriminant function

The simplest neural network is a single **perceptron**: k inputs x_i, weights w_i, and one output

$$y = S\left(\sum_{i=1}^{k} w_i x_i\right) = S(w^\top x),$$

where S is a "squasher" mapping $\mathbb{R}$ onto $[0, 1]$: a step function, or the sigmoid $S(z) = 1/(1 + e^{-z})$, or any cumulative distribution function.

For fixed weights the perceptron is a **classifier**: it fires ($y = 1$) when $w^\top x > 0$ and stays quiet otherwise, so the boundary $w^\top x = 0$ is a hyperplane separating two classes.

Training the perceptron, that is, choosing w to minimize $\sum_t (y_t - S(w^\top x_t))^2$, is a nonlinear least squares problem, solved by the stochastic-gradient recursion

$$w_t = w_{t-1} + \gamma_t \nabla S(w_{t-1}, x_t)\,(y_t - S(w_{t-1}^\top x_t)),$$

where $\nabla S(w, x)$ is the gradient of $S(w^\top x)$ with respect to w, equal to $S(w^\top x)(1 - S(w^\top x))\,x$ for the sigmoid.

This has the form $\theta_t = \theta_{t-1} + \gamma_t Q(\theta_{t-1}, x_t)$ shared by every learner in this series (*One recursion, many learners*).

The book suggests a simple scheme: set γ_t to a small constant and run the sample through the recursion repeatedly until the weights settle.

Take the book's example: separate football players from economists on two features, weight and salary (synthetic data, in arbitrary units).

```
rng = np.random.default_rng(0)
n = 100
economists = rng.multivariate_normal([-1.0, -0.3], [[0.5, 0.1], [0.1, 0.5]], n)
players    = rng.multivariate_normal([1.2,  0.8], [[0.6, 0.0], [0.0, 0.6]], n)
X = np.vstack([economists, players])
y = np.r_[np.zeros(n), np.ones(n)]
X_aug = np.column_stack([np.ones(len(X)), X])          # prepend an intercept

def sigmoid(z):
    return 1 / (1 + np.exp(-z))

def train_perceptron(X_aug, y, y=0.1, passes=200, nls=True, seed=1):
    "Constant-gain stochastic gradient, run repeatedly through the sample."
    rng = np.random.default_rng(seed)
    w = np.zeros(X_aug.shape[1])
    for _ in range(passes):
        for t in rng.permutation(len(y)):
            S = sigmoid(X_aug[t] @ w)
            grad = S * (1 - S) * X_aug[t] if nls else X_aug[t]
            w += y * grad * (y[t] - S)
    return w

w = train_perceptron(X_aug, y)                         # the recursion displayed above
print(f"training accuracy: {np.mean((sigmoid(X_aug @ w) > 0.5) == y):.3f}")
```

```
training accuracy: 0.945
```

The `nls=False` branch drops the factor $S(1-S)$; we use it in a moment.

The point Sargent [1993] stresses is that this is nothing exotic to an econometrician: the perceptron's decision boundary is essentially a **linear discriminant function**.

Compare the perceptron's boundary normal to Fisher's linear discriminant direction.

```
def unit(v):
    return v / np.linalg.norm(v)

def angle(u, v):
    "Angle in degrees between two directions, ignoring sign."
    return np.degrees(np.arccos(min(1.0, abs(unit(u) @ unit(v)))))

m0, m1 = X[y == 0].mean(0), X[y == 1].mean(0)
S_within = np.cov(X[y == 0].T) * n + np.cov(X[y == 1].T) * n
lda_direction = unit(np.linalg.solve(S_within, m1 - m0))
perceptron_direction = unit(w[1:])

print(f"perceptron boundary normal   : {perceptron_direction.round(3)}")
print(f"Fisher discriminant direction: {lda_direction.round(3)}")
print(f"cosine similarity: {abs(perceptron_direction @ lda_direction):.4f}, "
      f"angle {angle(w[1:], lda_direction):.1f} degrees")
```

```
perceptron boundary normal   : [0.982 0.188]
Fisher discriminant direction: [0.925 0.38 ]
cosine similarity: 0.9798, angle 11.5 degrees
```

Dropping the factor $S(1-S)$ from the recursion turns it into stochastic-gradient ascent on the logistic log-likelihood, that is, into **logistic regression**.

```
w_logit = train_perceptron(X_aug, y, nls=False)
print(f"logistic regression: cosine {abs(unit(w_logit[1:]) @ lda_direction):.4f}, "
      f"angle {angle(w_logit[1:], lda_direction):.1f} degrees")
```

```
logistic regression: cosine 0.9978, angle 3.8 degrees
```

Here is the perceptron's boundary.

```
fig, ax = plt.subplots(figsize=(6.5, 5))
ax.scatter(*economists.T, s=18, c='C0', label="economists ($y=0$)")
ax.scatter(*players.T, s=18, c='C1', label="football players ($y=1$)")
xs = np.linspace(X[:, 0].min(), X[:, 0].max(), 2)
ax.plot(xs, -(w[0] + w[1]*xs) / w[2], 'k-', lw=1.5, label="perceptron boundary")
ax.set_ylim(X[:, 1].min() - 0.5, X[:, 1].max() + 0.5)
ax.set_xlabel("weight")
ax.set_ylabel("salary")
ax.legend(frameon=False)
plt.show()
```

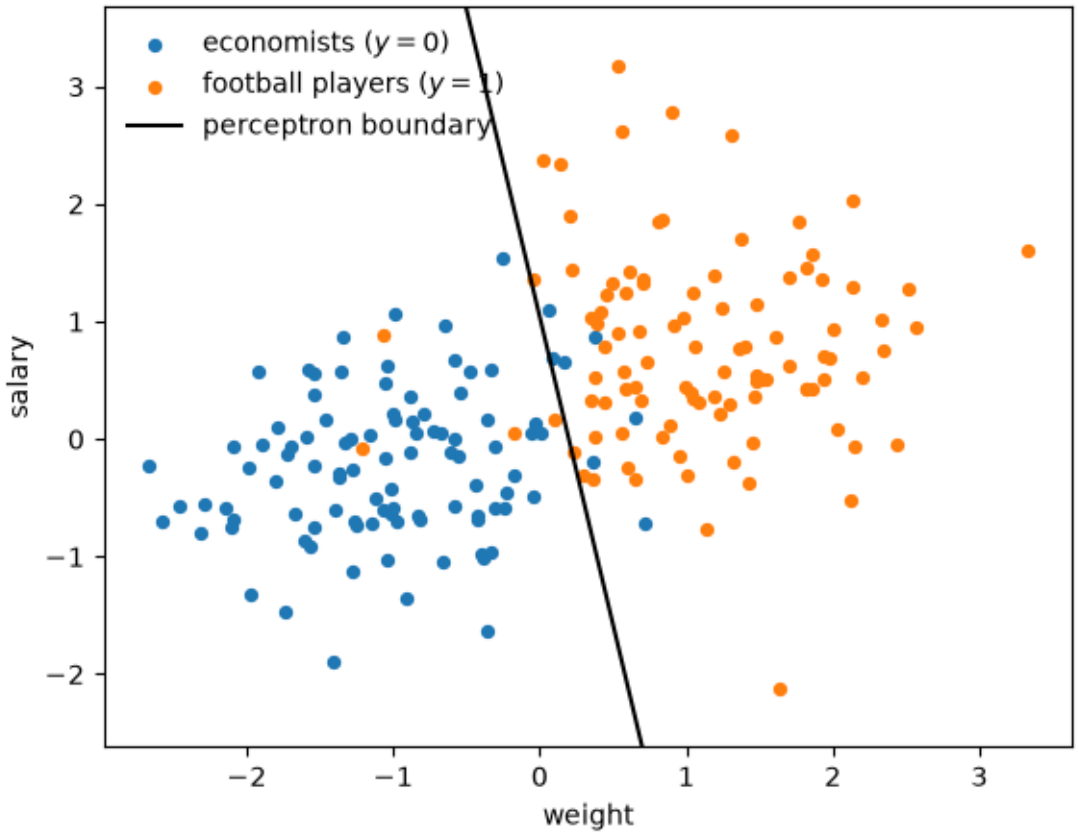


Fig. 94.1: Perceptron boundary separating the two groups

Both boundaries point close to Fisher's direction: the perceptron's is about 11.5 degrees away and logistic regression's about 4.

The closeness is no accident.

If the two classes are Gaussian with a common covariance matrix Σ, then $\Pr(y = 1 \mid x)$ is exactly a sigmoid in x whose slope is $\Sigma^{-1}(\mu_1 - \mu_0)$, Fisher's direction, so both recursions aim at it.

Here the two covariance matrices differ a little and the sample is finite, and the least squares criterion weights the observations differently from the likelihood, so the directions differ.

Minsky and Papert's famous critique [Minsky and Papert, 1969] was precisely that a single perceptron can represent *only* a linear discriminant, and so cannot separate classes that a line cannot.

The revival of the field came with the recognition that **layering** perceptrons, feeding their outputs into further perceptrons, approximates any nonlinear discriminant, the subject of *Introduction to Artificial Neural Networks*.

For our purposes the perceptron is the entry in the catalogue closest to what the previous lectures already did: a parametric rule, trained by a stochastic-gradient recursion, that an econometrician would recognize on sight.

The Hopfield network is a bridge of another kind, and the last two brains are stranger still.

94.3 Associative memory: the Hopfield network

The second device stores patterns and recalls them from corrupted fragments.

Represent a pattern as a vector of ± 1 of length N.

We want to store p patterns $\sigma^1, \dots, \sigma^p$, the columns of an $N \times p$ matrix σ, so that each is a fixed point of a dynamical system, and so that feeding in a corrupted version converges quickly to the nearest stored pattern.

The **Hopfield network** [Hopfield, 1982] does this with the dynamics

$$s(t+1) = \operatorname{sgn}(w\, s(t)),$$

and a weight matrix built from the patterns themselves.

When the patterns are orthogonal, Hebb's rule $w = \frac{1}{N}\sigma\sigma^\top$ suffices; for merely linearly independent (correlated) patterns one uses the projection rule $w = \frac{1}{N}\sigma V^{-1}\sigma^\top$ with $V = \frac{1}{N}\sigma^\top\sigma$.

Both make each stored pattern an exact fixed point.

The projection rule has a simple geometry: $w = \sigma(\sigma^\top\sigma)^{-1}\sigma^\top$ is the orthogonal projection onto the span of the stored patterns.

So the energy $E(s) = -\frac{1}{2}s^\top w s$ equals $-\frac{1}{2}\|ws\|^2$, which is at least $-\frac{1}{2}\|s\|^2 = -N/2$, with equality exactly when s lies in that span.

The stored patterns are therefore *global* minima of the energy.

We store five letters on a 5×5 pixel grid, and because letters share many pixels, they are correlated, so the projection rule is the right one.

(The code stores the patterns as the rows of σ, the transpose of the matrix in the formulas.)

```
PATTERNS = {
    'A': ["01110", "10001", "11111", "10001", "10001"],
    'E': ["11111", "10000", "11110", "10000", "11111"],
    'I': ["11111", "00100", "00100", "00100", "11111"],
    'O': ["01110", "10001", "10001", "10001", "01110"],
    'T': ["11111", "00100", "00100", "00100", "00100"],
}

def to_vector(rows):
    return np.array([1 if c == '1' else -1 for r in rows for c in r])
```

(continues on next page)

(continued from previous page)

```
letters = list(PATTERNS)
σ = np.array([to_vector(PATTERNS[c]) for c in letters])      # (p, N): one pattern␣
↪per row

def projection_rule(patterns):
    "Weight matrix making each (correlated) pattern an exact fixed point."
    N = patterns.shape[1]
    Σ = patterns.T                                          # (N, p)
    V = Σ.T @ Σ / N
    return Σ @ np.linalg.inv(V) @ Σ.T / N

def energy(w, s):
    "Hopfield energy; stored patterns are global minima."
    return -0.5 * s @ w @ s

def recall(w, s0, max_iter=30):
    "Iterate sgn(w s) to a fixed point."
    s = s0.copy()
    for _ in range(max_iter):
        s_new = np.sign(w @ s)
        s_new[s_new == 0] = 1
        if np.array_equal(s_new, s):
            break
        s = s_new
    return s

w_hop = projection_rule(σ)
fixed = all(np.array_equal(np.sign(w_hop @ σ[i]), σ[i]) for i in range(len(letters)))
print(f"all stored letters are fixed points: {fixed}")
print(f"energy of each stored letter: {[round(energy(w_hop, σ[i]), 1) for i in␣
↪range(len(letters))]}")
```

```
all stored letters are fixed points: True
energy of each stored letter: [np.float64(-12.5), np.float64(-12.5), np.float64(-
↪12.5), np.float64(-12.5), np.float64(-12.5)]
```

Every stored letter sits at the same energy $-N/2$ and is a fixed point.

Now corrupt a few pixels and let the network fall back to the nearest memory.

```
rng = np.random.default_rng(3)
show = ['E', 'O', 'T']

fig, axes = plt.subplots(len(show), 3, figsize=(6, 6))
for row, letter in enumerate(show):
    i = letters.index(letter)
    corrupt = σ[i].copy()
    corrupt[rng.choice(25, 4, replace=False)] *= -1             # flip 4 pixels
    recovered = recall(w_hop, corrupt)
    for col, (img, title) in enumerate([(σ[i], "stored"),
                                        (corrupt, "corrupted"),
                                        (recovered, "recovered")]):
        axes[row, col].imshow(img.reshape(5, 5), cmap='binary', vmin=-1, vmax=1)
        axes[row, col].set_xticks([])
        axes[row, col].set_yticks([])
        if row == 0:
```

(continues on next page)

(continued from previous page)

```
            axes[row, col].set_title(title)
plt.tight_layout()
plt.show()
```

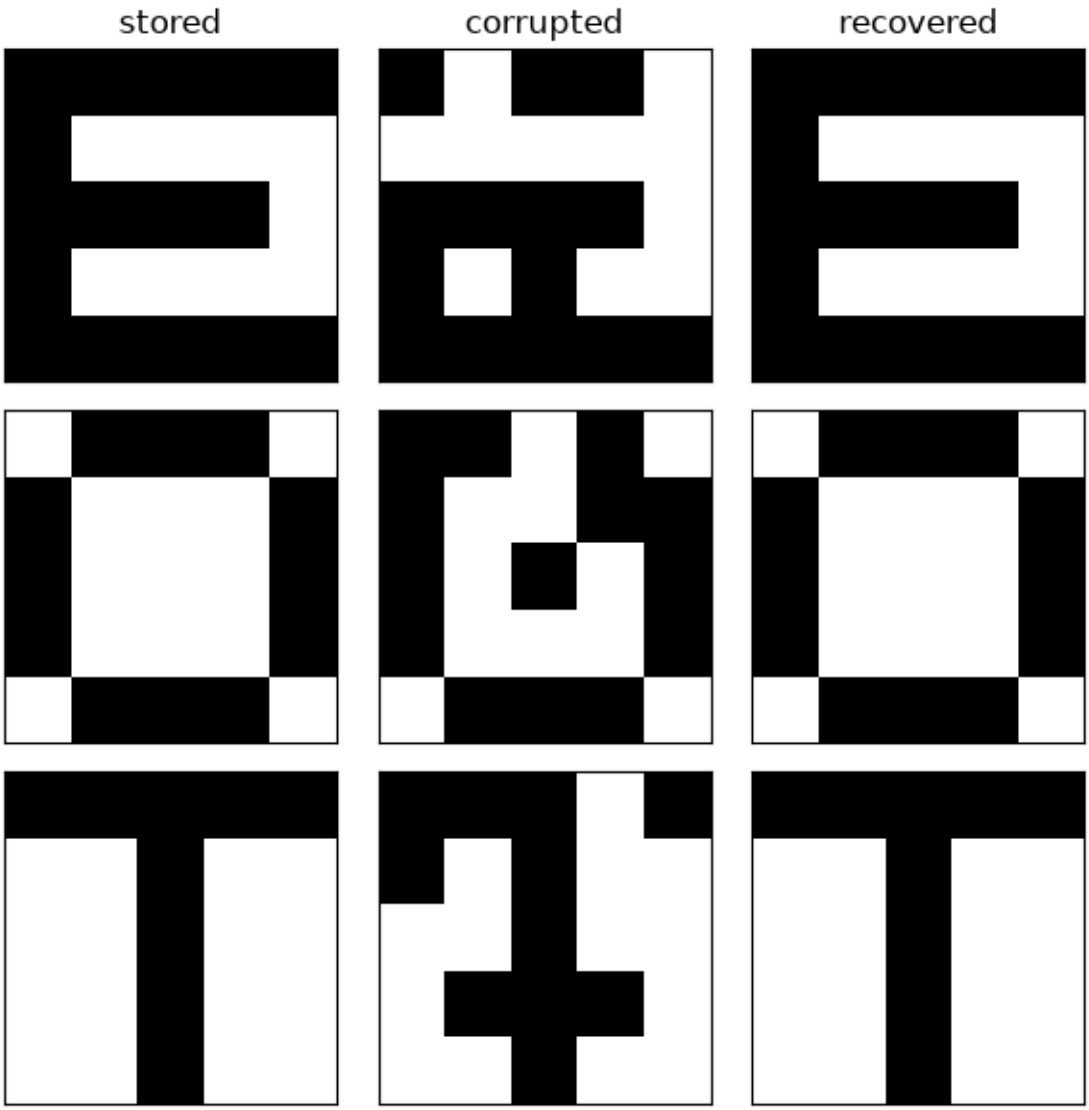


Fig. 94.2: Hopfield recall of letters from corrupted inputs

The corrupted letters snap back to their originals.

How reliably depends on how much is corrupted and on how correlated the patterns are.

A corrupted letter can occasionally land in a *different* letter's basin, or in a spurious mixture the designer never stored.

```
for n_flip in (3, 5, 7):
    correct = 0
    for i in range(len(letters)):
```

(continues on next page)

(continued from previous page)

```
        for _ in range(40):
            corrupt = σ[i].copy()
            corrupt[rng.choice(25, n_flip, replace=False)] *= -1
            if np.array_equal(recall(w_hop, corrupt), σ[i]):
                correct += 1
    print(f"{n_flip} pixels flipped: recovered exactly {correct}/{len(letters)*40}")
```

```
3 pixels flipped: recovered exactly 179/200
5 pixels flipped: recovered exactly 138/200
7 pixels flipped: recovered exactly 77/200
```

Five patterns on 25 pixels is a heavy load ($p/N = 0.2$), and the letters overlap a lot: I and T differ in only 4 of the 25 pixels.

That is why even 3-pixel corruptions fail about one time in ten.

The book's analysis of recall assumes many fewer patterns than pixels, $p \ll N$.

The mechanism is worth naming because it recurs throughout this program.

Recall is **descent on an energy function** whose minima include the stored patterns.

Feeding in a corrupted pattern places the system on the energy surface near the right minimum, and the dynamics roll it down.

This is exactly the geometry behind **simulated annealing**, the method the book invokes for escaping *unwanted* local minima: add a controlled amount of random "shaking," large at first and declining over time, so the system can hop out of shallow spurious basins before settling into a deep one.

Annealing and stochastic Newton methods modify Newton's method to improve its global behavior without giving up its good local behavior.

The same principle, to experiment a lot early and less later, governs the classifier systems of *Money as a Medium of Exchange among Artificially Intelligent Agents*, whose genetic algorithm fires with a probability that declines like $1/\sqrt{t}$, and the softmax auction of Exercise 94.8.3, whose "temperature" $1/\beta$ plays the role of the shaking.

94.4 The genetic algorithm

The third brain departs from Newton's method more fundamentally: it does not descend a smooth surface at all.

Holland's **genetic algorithm** [Goldberg, 1989, Holland, 1975] searches "rugged landscapes" that lack the smoothness Newton's method needs, by evolving a *population* of candidate solutions encoded as binary strings.

The two approaches can be combined.

Ellen McGrattan applied genetic algorithms to estimate nonlinear rational expectations models: she first used a genetic algorithm to search for the vicinity of a maximum, then switched to a Newton method [Sargent, 1993].

The algorithm is handed two things: a nonnegative fitness function f to maximize, and an *encoding* that maps each binary string of length S to a candidate solution x.

Given a population of N strings, one generation applies four operators:

1. **Evaluation**: compute each string's fitness $f(x_i)$.
2. **Reproduction**: copy strings into the next generation, string i with probability $f(x_i)/\sum_j f(x_j)$, proportional to fitness (a "biased roulette wheel").

3. **Crossover**: pair strings up and swap their tails at a random cut point, forming children that recombine their parents' segments.
4. **Mutation**: flip each bit independently with a small probability p_m.

The parameters of the algorithm are N, S and p_m.

Our code below uses a common variant of reproduction called *windowing*: it subtracts the generation's lowest fitness before spinning the wheel.

When all fitnesses lie close together, as they will below, plain proportional selection barely discriminates among strings; windowing sharpens it, and gives the worst string essentially no chance to reproduce.

Reproduction concentrates the population on what already works; crossover and mutation inject new candidates.

Holland's account of why crossover helps is in terms of **schemata**.

A schema H is a template in $\{0, 1, *\}^S$, where $*$ matches either bit: `1**0*` is the set of strings of length 5 with a 1 in the first position and a 0 in the fourth.

Its *order* $o(H)$ is the number of fixed positions, and its *defining length* $\delta(H)$ is the distance between the first and the last of them.

Holland's **schema theorem** says that with fitness-proportional reproduction, single-point crossover applied to every pair, and mutation at rate p_m, the expected number of strings in H satisfies

$$\mathbb{E}\, m(H, t+1) \;\geq\; m(H, t)\, \frac{f(H)}{\bar{f}} \left[1 - \frac{\delta(H)}{S-1} - o(H)\, p_m \right],$$

where $f(H)$ is the average fitness of the strings in H and $\bar{f}$ that of the whole population (with windowing, read f as the windowed fitness).

Short, low-order schemata with above-average fitness therefore tend to multiply.

The bound covers one generation only.

The "building-block hypothesis," that the algorithm works by assembling short, fit schemata into better strings, is a heuristic rather than a theorem.

Whether it applies depends, in the book's words, "precisely on how a given problem has been encoded": on whether bits that work well together sit near one another.

The $*$ of a schema reappears as the $\#$ of a classifier's condition below.

Crucially, the individual strings do not learn: each lives one generation and dies.

Only the *society*, the sequence of populations, learns.

(Arifovic's election operator in *Exchange Rate Indeterminacy, Learning, and Experiments* softens this: a parent survives unless its child would have done better, so a good string can live for many generations.)

That is what makes the genetic algorithm awkward as a model of one brain, and a better fit as a model of a population or a market, a point that matters when we reach *Money as a Medium of Exchange among Artificially Intelligent Agents*.

94.4.1 Axelrod's iterated prisoner's dilemma

Axelrod [1987] used the genetic algorithm to evolve strategies for the **iterated prisoner's dilemma**, the game whose round-robin computer tournament, famously, was won by tit-for-tat.

He bred strategies against eight strategies representative of those submitted to his second tournament; we use a smaller panel.

Two players each choose to Cooperate or Defect; the payoffs reward defection against a cooperator but punish mutual defection.

```
# 1 = Cooperate, 0 = Defect;  T > R > P > S is the prisoner's dilemma ranking
T_pay, R_pay, P_pay, S_pay = 5, 3, 1, 0
PAYOFF = np.array([[P_pay, T_pay],       # row: my move (D, C); column: opponent's (D,
 ↪ C)
                   [S_pay, R_pay]])
```

Following Axelrod, a strategy is a **70-bit string**: a strategy conditions on the outcomes of the last three rounds.

Each round is one of four outcomes (my move, opponent's move), so three rounds give $4^3 = 64$ possible histories, and 64 bits specify the action in each.

The remaining 6 bits encode a presumed pre-game history to get the first moves going.

Our `play` function plays a whole population of strings at once, one row per string.

```
POW = 2 ** np.arange(5, -1, -1)          # reads 6 history bits as an integer 0..63

def play(pop, opponent, n_rounds=150):
    "Play every 70-bit string in pop against opponent; return average payoffs."
    pop = np.atleast_2d(pop)
    N = len(pop)
    action = pop[:, :64]
    idx = pop[:, 64:] @ POW                      # the presumed pre-game history
    mine = np.zeros((N, n_rounds), int)
    theirs = np.zeros((N, n_rounds), int)
    my_total, opp_total = np.zeros(N), np.zeros(N)
    for t in range(n_rounds):
        a = action[np.arange(N), idx]            # my move from the lookup table
        b = opponent(mine[:, :t], theirs[:, :t], t)
        my_total += PAYOFF[a, b]
        opp_total += PAYOFF[b, a]
        mine[:, t], theirs[:, t] = a, b
        idx = 32 * a + 16 * b + (idx >> 2)       # newest round in, oldest round out
    return my_total / n_rounds, opp_total / n_rounds
```

The strategy is bred to do well against a fixed panel of opponents.

Each panel member is a function of the histories *as that member sees them*: its first argument is what its opponent has played and its second is what it has played itself.

The Random member flips a fair coin each round.

To make comparisons fair, each generation draws one sequence of coin flips, and every string in the population faces that same sequence.

This device, **common random numbers**, makes differences in fitness within a generation reflect differences in strategy rather than luck of the draw.

```
def all_cooperate(opp_moves, own_moves, t): return np.ones(len(opp_moves), int)
def all_defect(opp_moves, own_moves, t):    return np.zeros(len(opp_moves), int)

def tit_for_tat(opp_moves, own_moves, t):
    return opp_moves[:, -1] if t > 0 else np.ones(len(opp_moves), int)

def grudger(opp_moves, own_moves, t):
    return opp_moves.all(axis=1).astype(int)     # cooperate until the first defection

def random_opponent(flips):
    "An opponent who plays coin flips drawn in advance: row k of flips in game k."
    return lambda opp_moves, own_moves, t: flips[:, t]

PANEL = {'AllC': all_cooperate, 'AllD': all_defect,
         'TFT': tit_for_tat, 'Grudger': grudger}

def fitness(pop, flips):
    "Average payoff of each string against the panel, Random included."
    opponents = list(PANEL.values()) + [random_opponent(flips)]
    return np.mean([play(pop, opp)[0] for opp in opponents], axis=0)
```

Tit-for-tat copies what its opponent last did, and the grudger defects forever once its opponent has defected once.

Both are themselves 70-bit strings.

Tit-for-tat's table copies the opponent's most recent move.

The grudger's table cooperates only after three rounds of mutual cooperation, which, starting from a cooperative pre-game history, is the same as never having seen a defection.

We use these two strings as benchmarks, and score any string by its payoff against each panel member, averaging over 400 fresh Random games (the same 400 for every string we compare).

```
TFT_GENE = np.r_[(np.arange(64) >> 4) & 1, np.ones(6, int)]
GRUDGER_GENE = np.r_[(np.arange(64) == 63).astype(int), np.ones(6, int)]

def panel_scores(gene, reps=400, seed=1):
    "Payoff against each panel member; Random is averaged over reps fresh games."
    flips = np.random.default_rng(seed).integers(0, 2, (reps, 150))
    scores = {name: play(gene, opp)[0][0] for name, opp in PANEL.items()}
    scores['Random'] = play(np.tile(gene, (reps, 1)), random_opponent(flips))[0].
 ↪mean()
    return scores

def panel_fitness(gene):
    return np.mean(list(panel_scores(gene).values()))

def first_move(gene):
    return 'C' if gene[:64][gene[64:] @ POW] else 'D'
```

Now evolve.

We display one run, with seed 16, chosen because its champion's fitness is close to the median of the twenty runs examined below.

```
def evolve(N=60, generations=80, p_mut=0.01, seed=0):
    rng = np.random.default_rng(seed)
    pop = rng.integers(0, 2, (N, 70))
    best_hist, mean_hist = [], []
```

(continues on next page)

(continued from previous page)

```
    for _ in range(generations):
        flips = rng.integers(0, 2, (1, 150))           # one coin sequence for␣
↪everyone
        fit = fitness(pop, flips)
        best_hist.append(fit.max())
        mean_hist.append(fit.mean())
        weight = fit - fit.min() + 1e-6                # windowing
        cum = np.cumsum(weight)
        parents = pop[np.searchsorted(cum, rng.random(N) * cum[-1])]
        mom, dad = parents[0::2], parents[1::2]
        left = np.arange(70) < rng.integers(1, 70, (N // 2, 1))   # single-point cuts
        pop = np.vstack([np.where(left, mom, dad), np.where(left, dad, mom)])
        pop[rng.random(pop.shape) < p_mut] ^= 1        # mutation
    fit = fitness(pop, rng.integers(0, 2, (1, 150)))
    return pop[fit.argmax()], np.array(best_hist), np.array(mean_hist)


champion, best_hist, mean_hist = evolve(seed=16)
print(f"generation 0:  best {best_hist[0]:.3f}, mean {mean_hist[0]:.3f}")
print(f"generation {len(best_hist)-1}: best {best_hist[-1]:.3f}, mean {mean_hist[-1]:.
↪3f}")
```

```
generation 0:  best 2.429, mean 1.941
generation 79: best 2.608, mean 2.397
```

The figure tracks the best and the mean fitness within each generation.

```
fig, ax = plt.subplots(figsize=(7.5, 4))
ax.plot(best_hist, label="best in population", lw=2)
ax.plot(mean_hist, label="population mean", lw=2)
ax.set_xlabel("generation")
ax.set_ylabel("panel fitness")
ax.legend(frameon=False)
plt.show()
```

The population mean climbs quickly for about a dozen generations and then fluctuates around 2.4, held down by the crossover and mutation that keep creating weaker strings.

The best string's fitness jumps around and drifts up slowly, to about 2.6 by the end.

It jumps around because each generation faces a new sequence of coin flips from Random, and because nothing guarantees that the best string survives reproduction, crossover, and mutation.

What kind of strategy did evolution find?

```
strategies = {'champion': champion, 'TFT': TFT_GENE, 'Grudger': GRUDGER_GENE}
table = {name: panel_scores(gene) for name, gene in strategies.items()}

print(f"{'payoff vs':>10}" + "".join(f"{name:>10}" for name in strategies))
for opp in table['champion']:
    print(f"{opp:>10}" + "".join(f"{table[name][opp]:>10.3f}" for name in strategies))
print(f"{'panel avg':>10}" + "".join(f"{np.mean(list(table[name].values())):>10.3f}"
                                     for name in strategies))
print(f"champion's first move: {first_move(champion)}")
```

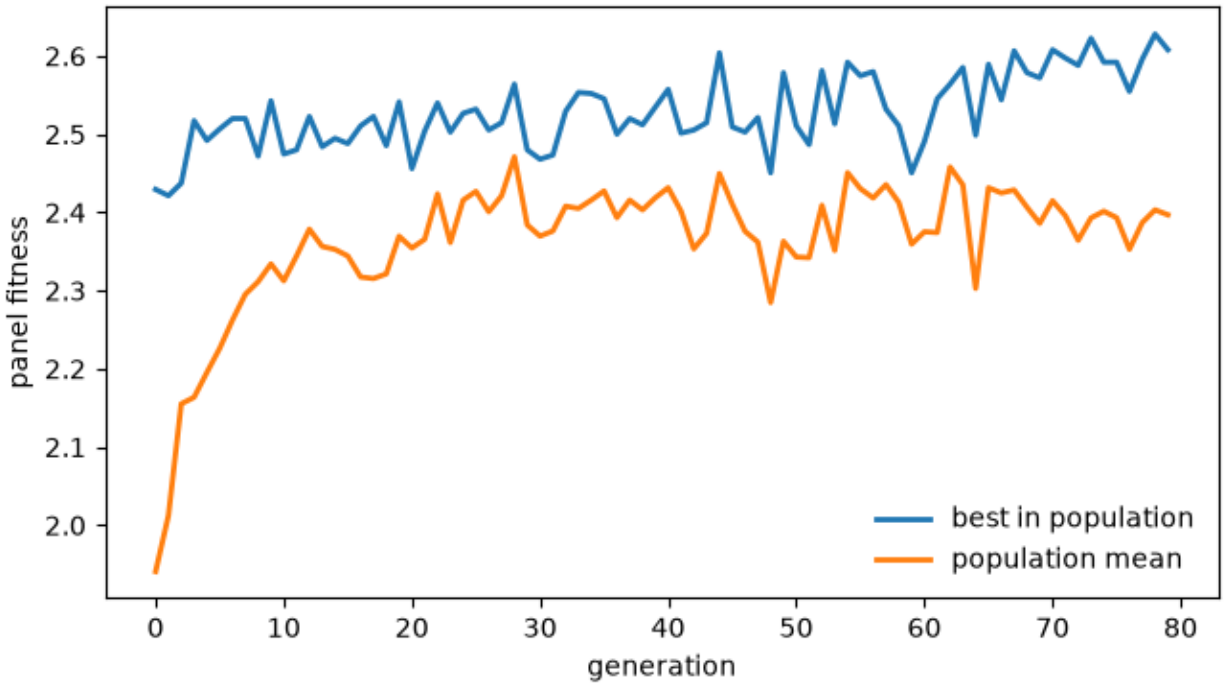


Fig. 94.3: Population fitness across generations

```
 payoff vs  champion       TFT   Grudger
      AllC     3.000     3.000     3.000
      AllD     0.987     0.993     0.993
       TFT     3.000     3.000     3.000
   Grudger     3.000     3.000     3.000
    Random     2.881     2.239     2.971
 panel avg     2.573     2.446     2.593
champion's first move: C
```

The evolved strategy is *nice*, in Axelrod's language: it is never the first to defect against AllC, TFT, or the grudger, and so earns the mutual-cooperation payoff of 3 against all three.

It refuses to be suckered by AllD after its first two rounds, and it defects against Random often enough to earn much more than tit-for-tat does.

Against this panel it beats tit-for-tat, and the margin comes almost entirely from Random.

It does not beat the grudger, a member of the panel, which does slightly better still.

Is this run typical?

The next cell repeats the experiment with twenty seeds.

```
results = []
for seed in range(20):
    sc = panel_scores(evolve(seed=seed)[0])
    nice = all(sc[k] == 3.0 for k in ('AllC', 'TFT', 'Grudger'))
    results.append((np.mean(list(sc.values())), nice))
fits = np.array([r[0] for r in results])
nice = np.array([r[1] for r in results])
tft_fit, grudger_fit = panel_fitness(TFT_GENE), panel_fitness(GRUDGER_GENE)
```

(continues on next page)

(continued from previous page)

```
print(f"nice champions: {nice.sum()} of 20")
print(f"champion fitness: median {np.median(fits):.3f}, "
      f"range [{fits.min():.3f}, {fits.max():.3f}]")
print(f"  nice ones: mean {fits[nice].mean():.3f};  others: mean {fits[~nice].mean():.
 ↪3f}")
print(f"beat TFT ({tft_fit:.3f}): {np.sum(fits > tft_fit)} of 20")
print(f"within 0.01 of Grudger ({grudger_fit:.3f}): {np.sum(fits > grudger_fit - 0.
 ↪01)} of 20;"
      f"  above it: {np.sum(fits > grudger_fit)}")
```

```
nice champions: 15 of 20
champion fitness: median 2.576, range [2.481, 2.594]
  nice ones: mean 2.550;  others: mean 2.556
beat TFT (2.446): 20 of 20
within 0.01 of Grudger (2.593): 7 of 20;  above it: 1
```

Most champions are nice, but not all.

The rest are *probe-and-exploit* strategies: at some point they defect unprovoked, which wins 5 a round forever against AllC and plunges them into mutual defection with the grudger.

The panel makes the two families almost exactly equal.

AllC and the grudger behave identically until someone defects, so a strategy cannot tell them apart without probing.

A probe gains 2 a round against AllC and loses 2 a round against the grudger, so the two families tie on that pair of opponents.

Which family the algorithm lands on is a matter of the seed: a small instance of a rugged landscape with two peaks of almost equal height.

The panel also has a ceiling.

No strategy can average more than 1 against AllD, more than 3 against Random (against a coin, defection is better whatever the coin shows), more than about 3 against tit-for-tat, or more than about 6 against AllC and the grudger together.

So panel fitness cannot exceed about $13/5 = 2.6$, and the grudger is already within 0.01 of that ceiling.

Every champion beats tit-for-tat, but none beats the grudger by more than a rounding error.

Axelrod's own genetic algorithm, as Sargent [1993] summarizes it, "produced a strategy that would have won the tournament, and in particular would have outplayed the 'tit-for-tat' strategy that won the tournament."

Our toy version reproduces the second half of that sentence, and shows where the edge comes from.

It is not extra memory: the grudger needs only one bit of memory, and does as well as any champion.

It is refusing to forgive: against Random, tit-for-tat keeps returning to cooperation and gets suckered, while the grudger and the best champions settle into defection.

And, as Sargent notes, such comparisons flatter the genetic algorithm, because it "got to practice against" its opponents, while the strategies submitted to a tournament could only guess what they would face.

94.5 Classifier systems

The genetic algorithm evolves a population in which no individual learns.

Holland's **classifier system** puts the same evolutionary machinery *inside a single agent*, as a model of one brain [Holland *et al.*, 1986].

Sargent describes it as Holland's vision of the mind as a **competitive economy**:

> The statements compete with one another for the opportunity to decide. The classifier system incorporates elements of the genetic algorithm with other aspects in a way that represents a brain in terms that Holland describes as a competitive economy.

A classifier system has five parts:

- **Classifiers**, if-then rules encoded as strings over the trinary alphabet $\{0, 1, \#\}$, with a condition part that is matched against the state and an action part that prescribes a move.
- **A decoder** that, given the current state, finds which classifiers' conditions match.
- **An auction** that selects one matching classifier to act.
- **An accounting system** that updates each classifier's **strength** from the payoffs that follow its decisions and, in sequential problems, passes payoffs *backward* along the chain of rules that produced them.
- **A device for creating new classifiers**, so that the system's very vocabulary evolves:
  - when no classifier matches the state, one is created whose condition matches it and whose action is random;
  - now and then new rules replace weak ones, so that untried actions get tried;
  - rules are *generalized* (0's and 1's replaced by #'s) or *specialized* (#'s replaced by 0's and 1's), often by genetic operators.

We take the decoder and the auction first, then the accounting system, through two examples: the bandit and the bucket brigade.

94.5.1 Matching, specificity, and bids

A state matches a condition if they agree in every position where the condition does not have a #, the "I don't care" symbol.

The condition `1#0` matches the states `100` and `110` and nothing else.

A rule with many #'s is *general*: its condition is met often.

A rule with few is *specific*.

Marimon *et al.* [1990] measure the specificity of a classifier e by $\sigma_e = 1/(1 + \text{number of } \#\text{'s in } e)$.

```
def matches(condition, state):
    "A state matches if it agrees with every position that is not '#'."
    return all(c in ('#', s) for c, s in zip(condition, state))

def specificity(condition):
    return 1 / (1 + condition.count('#'))

state = '110'
for condition in ('110', '1#0', '##0', '###', '0##'):
    print(f"condition {condition}: matches {state}? {str(matches(condition, state)):5}
 ↪  "
          f"specificity {specificity(condition):.2f}")
```

```
condition 110: matches 110? True   specificity 1.00
condition 1#0: matches 110? True   specificity 0.50
condition ##0: matches 110? True   specificity 0.33
condition ###: matches 110? True   specificity 0.25
condition 0##: matches 110? False  specificity 0.33
```

Several rules usually match at once, and they may disagree about what to do.

Each matched rule e bids a fraction $b(e)$ of its strength S_e, a fraction that rises with its specificity:

$$B_e = b(e)\, S_e, \qquad b(e) = b_{11} + b_{12}\, \sigma_e,$$

where, in the notation of Marimon *et al.* [1990], b_{11} and b_{12} are positive constants that sum to less than one.

When a general rule and a specific one are equally strong, the specific one outbids it, so an auction on bids favors it.

This is how a classifier system handles the trade-off that the book calls *generality versus discrimination*.

General rules are tested often, because their conditions are often met, but they prescribe the same action across many states.

Specific rules can fit the action to the state, but they need long histories before their strengths mean much.

The trade-off resembles the one between parametric and non-parametric estimation (*Learning without a functional form*).

With a specificity premium, a specific rule can carve out an exception to a general one, and the balance between generality and specificity can itself evolve as rules are generalized and specialized.

94.5.2 The auction

The auction turns bids into a decision.

The book lists two principles: the highest bidder wins, or each matched rule wins with probability equal to its relative strength.

Both belong to one family, the **softmax**:

$$\Pr(e \text{ wins}) = \frac{\exp(\beta B_e)}{\sum_{e' \in M} \exp(\beta B_{e'})},$$

where M is the set of matched rules.

As $\beta \to \infty$ the softmax puts all its weight on the highest bid, so the **argmax auction** is its limit.

Applied to the logarithms of strengths with $\beta = 1$, it gives the relative-strength rule, since $\exp(\ln S_e) / \sum_{e'} \exp(\ln S_{e'}) = S_e / \sum_{e'} S_{e'}$.

The parameter β sets how much the system experiments.

The argmax auction never lets a rule that looks worse act, so it can never discover that the rule is in fact better; with a finite β, the system always might.

Money as a Medium of Exchange among Artificially Intelligent Agents uses an argmax auction applied to strengths rather than bids: there the strongest matched rule acts, bids only set the payments of the bucket brigade, and all experimentation is left to the operators for creating rules.

94.5.3 A two-armed bandit

The simplest classifier system is an example of Brian Arthur and Carl Simon, who presented it independently at the Santa Fe Institute in 1989 (see also Arthur [1991]).

Arm i pays a random reward with mean $\mu_i > 0$, with $\mu_1 > \mu_2$, and the player knows neither distribution.

The classifier system holds two rules, `#1` ("whatever the history, pull arm 1") and `#0` ("pull arm 2"), whose conditions are always met.

Arthur's accounting sets the strength of rule i to the average payoff its arm has delivered,

$$S_{i,\tau_i} = S_{i,\tau_i - 1} + \frac{1}{\tau_i}\left(x_{i,\tau_i} - S_{i,\tau_i - 1}\right),$$

where τ_i counts the times arm i has been pulled, and the auction picks rule i with probability $S_i/(S_1 + S_2)$.

The two rules start with **equal** strengths: the system is told nothing about either arm and must build its estimates out of the rewards it experiences.

```
def two_armed_bandit(μ, noise_sd=0.5, T=20_000, seed=0):
    "Strength = running average of an arm's payoff; pull ∝ strength."
    rng = np.random.default_rng(seed)
    S = np.ones(2)                          # equal strengths: no prior knowledge
    τ = np.array([1, 1])                    # pull counters
    pulls = np.empty(T, int)
    for t in range(T):
        w = np.maximum(S, 1e-9)             # guard against a dip below zero
        i = 0 if rng.random() < w[0] / w.sum() else 1
        reward = μ[i] + noise_sd * rng.standard_normal()
        τ[i] += 1
        S[i] += (reward - S[i]) / τ[i]      # running average
        pulls[t] = i
    return pulls

pulls = two_armed_bandit([3.0, 1.0])
frac_best = np.cumsum(pulls == 0) / np.arange(1, len(pulls) + 1)

fig, ax = plt.subplots(figsize=(7.5, 4))
ax.plot(frac_best, lw=1)
ax.axhline(0.75, color='k', ls='--', label="probability match  $\\mu_1/(\\mu_1+\\mu_2)
↪$")
ax.axhline(1.0, color='C3', ls=':', label="optimal (always best arm)")
ax.set_xlabel("$t$")
ax.set_ylabel("fraction of pulls on the best arm")
ax.set_ylim(0.5, 1.05)
ax.legend(frameon=False)
plt.show()
```

The fraction of pulls on the better arm converges, but **not to one**.

It converges to $\mu_1/(\mu_1 + \mu_2)$, the arm's share of total expected reward.

```
for μ in ([1.0, 0.5], [2.0, 1.0], [3.0, 1.0]):
    frac = np.mean(two_armed_bandit(μ)[-5000:] == 0)
    print(f"μ = {μ}: fraction on best arm = {frac:.3f}, "
          f"probability match = {μ[0]/(μ[0]+μ[1]):.3f}  (optimal = 1.0)")
```

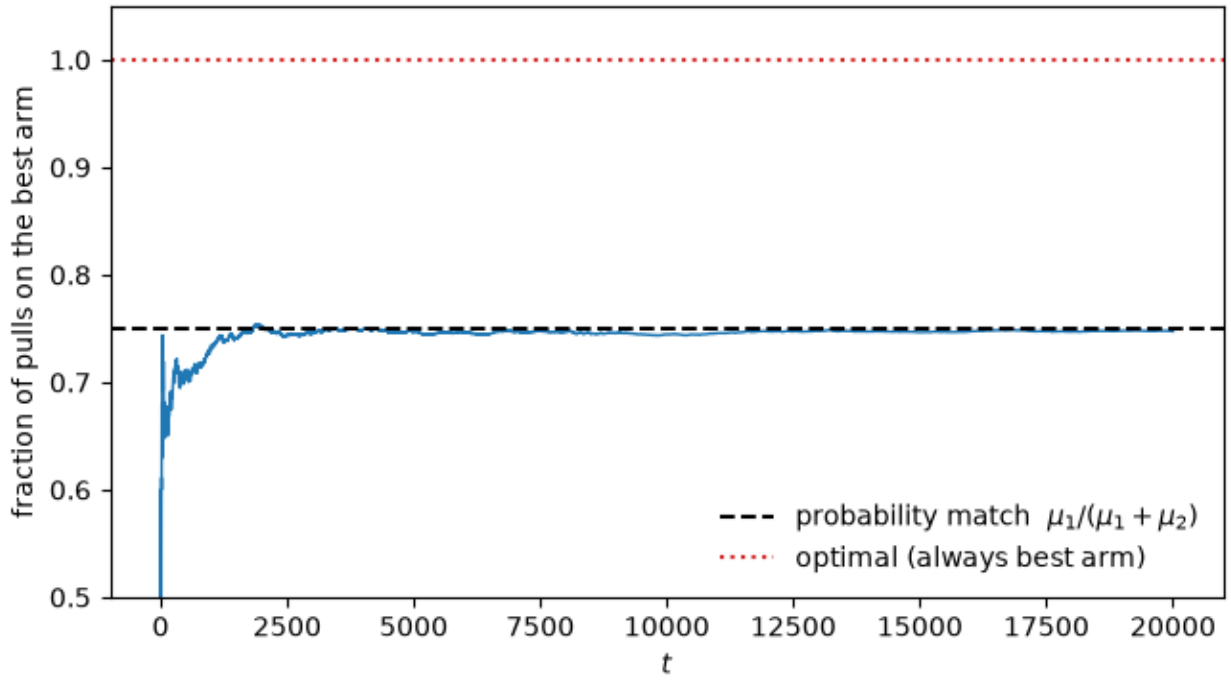


Fig. 94.4: The classifier bandit converges to probability matching

```
μ = [1.0, 0.5]: fraction on best arm = 0.655, probability match = 0.667  (optimal␣
↪= 1.0)
μ = [2.0, 1.0]: fraction on best arm = 0.659, probability match = 0.667  (optimal␣
↪= 1.0)
μ = [3.0, 1.0]: fraction on best arm = 0.750, probability match = 0.750  (optimal␣
↪= 1.0)
```

Arthur showed, by the methods of stochastic approximation, that the strength-proportional classifier **probability-matches**.

The argument is short.

As long as the strengths stay positive, which needs $\mu_i > 0$, each arm keeps a positive probability of being pulled, and so is pulled infinitely often.

Each strength is the sample mean of its arm's payoffs, so it converges to μ_i by the strong law of large numbers.

The auction's probabilities therefore converge to $\mu_1/(\mu_1 + \mu_2)$, and so does the fraction of pulls on arm 1.

The classifier plays the arms in proportion to their expected rewards rather than concentrating on the best one, so it leaves reward on the table forever.

With $\mu = (3, 1)$ it earns $0.75 \times 3 + 0.25 \times 1 = 2.5$ a pull, against 3 for always pulling arm 1.

The proportional rule also cares only about the ratio μ_1/μ_2: $\mu = (1, 0.5)$ and $\mu = (2, 1)$ produce the same shares.

Notice where the trouble lies.

The accounting is fine: each strength converges to the right number, its arm's mean payoff.

The trouble is the **auction**, which turns correct strengths into choice probabilities proportional to them.

Exercise 94.8.3 cures probability matching by changing the auction alone.

That separates two jobs.

The auction governs the balance between exploring and exploiting.

The accounting governs *credit assignment*, which is trivial in the bandit, where each payoff belongs to the arm that produced it, but hard in sequential problems.

94.5.4 The bucket brigade

In a sequential problem, a rule's decision may pay off only later, through a chain of further decisions.

The accounting must then reward a rule for merely *setting up* a profitable future decision.

Holland's device for this is the **bucket brigade**: each rule that wins an auction pays its bid to the rule that acted just before it, the one that moved the system into the state where the current rule could act.

A rule's receipts are the payoff it collects from the environment plus the bid paid to it by its successor.

Take the shortest chain.

A *setup* rule s acts first and moves the system into a state where a *payoff* rule p acts and collects R.

Rule p pays its bid bS_p back to s; to keep things simple, every rule bids the same fraction b of its strength.

Holland kept strengths as **cumulative balances**: a rule that acts pays out its bid and takes in its receipts, like a bank account, so that $S \leftarrow S - bS + \text{receipts}$.

Arthur, in the bandit, and Marimon *et al.* [1990], in the economy of the next lecture, **average** instead:

$$S_\tau = S_{\tau-1} + \frac{1}{\tau}\left[\text{receipts}_\tau - (1+b)\, S_{\tau-1}\right],$$

where τ counts the times the rule has acted, so that strength is a running average of receipts net of the bid.

With $b = 0$ this is the bandit's running average.

The code runs both accountings on the chain, starting from zero strengths, with $R = 1$ and $b = 0.1$.

```
def bucket_brigade(R=1.0, b=0.1, passes=10_000, accounting='average'):
    "Two-link chain: the setup rule acts, then the payoff rule acts and earns R."
    S = np.zeros(2)                                  # strengths of [setup, payoff]
    path = []
    for τ in range(1, passes + 1):
        receipts = np.array([b * S[1], R])           # the payoff rule pays its bid back
        if accounting == 'average':                  # Arthur; Marimon-McGrattan-Sargent
            S = S + (receipts - (1 + b) * S) / τ
        else:                                        # Holland's cumulative balance
            S = S - b * S + receipts
        path.append(S.copy())
    return np.array(path)

for accounting in ('balance', 'average'):
    path = bucket_brigade(accounting=accounting)
    print(f"{accounting}: [setup, payoff] after passes 1, 2, 3: "
          f"{path[0].round(3)}, {path[1].round(3)}, {path[2].round(3)};  "
          f"limit {path[-1].round(3)}")
```

```
balance: [setup, payoff] after passes 1, 2, 3: [0. 1.], [0.1 1.9], [0.28 2.71]; ␣
↪limit [10. 10.]
average: [setup, payoff] after passes 1, 2, 3: [0. 1.], [0.05 0.95], [0.063 0.935];
↪  limit [0.083 0.909]
```

Two things stand out.

First, credit moves back **one link per pass**.

In the first pass the payoff rule collects R but the setup rule receives nothing, because the payoff rule's strength, and hence its bid, was still zero when it paid.

In a chain of k links, the payoff takes $k-1$ passes to reach the first rule.

Second, the two accountings disagree about what the setup rule is worth.

Cumulative balances come to rest where receipts equal the bid, $bS_p = R$ and $bS_s = bS_p$, so $S_s = S_p = R/b$: the setup rule ends up exactly as strong as the payoff rule.

Averages come to rest where $(1+b)S_p = R$ and $(1+b)S_s = bS_p$, so

$$S_p = \frac{R}{1+b}, \qquad S_s = \frac{b}{1+b}\, S_p.$$

Each link back along the chain scales strength by $b/(1+b)$, about 0.09 when $b = 0.1$.

The averaging accounting values a setup rule as if it discounted the future very heavily, a fact with large consequences in *Money as a Medium of Exchange among Artificially Intelligent Agents* (*The bucket brigade's discount factor*).

Exercise 94.8.5 asks when a setup rule can beat a rule that pays off at once.

Why average, then?

Holland's balance with bids proportional to strength can be written $S \leftarrow S + b\,(\text{receipts}/b - S)$, a *constant-gain* average of receipts$/b$, which keeps fluctuating as long as the receipts are random.

The $1/\tau$ gain turns it into a decreasing-gain stochastic approximation of the kind studied in *One recursion, many learners*.

Its limit is a rest point of the mean dynamics $dS/d\tau = \mathbb{E}[\text{receipts}] - (1+b)S$, which is what makes strengths converge and lets the next lecture characterize their limits.

The bucket brigade is also an ancestor of **temporal-difference learning** [Sutton and Barto, 2018].

A rule's strength is moved toward what it earns now plus a fraction of the strength of the rule it hands the system to.

The Q-learning of *Job Search X: Search with Q-Learning* and *Inventory Management via Q-Learning* moves the value of a state-action pair toward the current reward plus the discounted value of the next state in the same way, with the successor's bid playing the role of that discounted value.

Designing that backward flow is the central craft of building a classifier system, and it is exactly what *Money as a Medium of Exchange among Artificially Intelligent Agents* must get right to make agents learn to accept money today for the sake of a trade tomorrow.

94.6 Evolutionary programming

We have surveyed four brains.

The last idea is about what to *do* with them.

Many of the systems of adaptive agents in this series, however plodding, find their way to an equilibrium.

The least squares learners of *Adaptive Agents in an Overlapping Generations Monetary Economy* converge to a stationary equilibrium, and the learners of *Learning, Approximation, and Equilibrium Computation* to the best approximate equilibrium their functional form allows.

At Sargent's calibration, the Newton learners of *Exchange Rate Indeterminacy, Learning, and Experiments* converge to one of a continuum of equilibria, picked out by history.

Not all of them do.

The same Newton learners fail to converge at Arifovic's calibration.

And the genetic-algorithm economy of *Exchange Rate Indeterminacy, Learning, and Experiments* never settles down (*A genetic algorithm economy*): its saving rate settles near the equilibrium value, but mutation keeps injecting new strings, and its exchange rate keeps fluctuating with a volatility that does not die out.

Evolutionary programming turns the tendency to converge into a tool.

If a population of adaptive agents reliably converges to an equilibrium, we can run the population as a *method for computing* the equilibrium, especially in models too complicated to solve by hand.

The genetic-algorithm economy is a reminder that the reliability has to be designed in: the experimentation must die away, as annealing's shaking does, if the population is to come to rest.

Sargent is careful about what is and isn't going on:

> The adaptive agents are 'teaching' the economist in the same sense that any numerical algorithm for solving nonlinear equations 'teaches' a mathematician. When these agents can 'teach' us something, it is because we designed them to do so.

This is the same duality that ran through *Learning, Approximation, and Equilibrium Computation*: a learning economy is a decentralized equilibrium computation, and an equilibrium computation is a centralized learning algorithm.

The genetic algorithm and the classifier system are simply richer computational engines than recursive least squares, able to search rugged landscapes and to discover the structure of a good rule, not just tune the coefficients of a fixed one.

The showcase is Kiyotaki and Wright [1989]'s search-theoretic model of money, in which the medium of exchange is not assumed but must **emerge** from how agents choose to trade.

The equilibrium is a set of trading strategies and matching probabilities, and for enriched versions of the model it is hard to characterize analytically.

Marimon *et al.* [1990] put populations of Holland classifier systems into that environment and watched them converge to an equilibrium, then built a five-good version, with no known analytical solution, and let the classifier systems suggest what the equilibrium looked like.

94.7 Concluding remarks

The four brains in this catalogue differ in what they take as given.

The perceptron is handed a functional form and asked only for its coefficients, which is why an econometrician recognizes it immediately.

The Hopfield network is handed the patterns themselves and asked only to recall them.

The genetic algorithm is handed an encoding and a fitness function and must search a space in which it cannot compute a gradient.

The classifier system is handed a vocabulary in which rules can be written and must discover which rules are worth holding.

Moving down that list, we hand the agent less and ask it to find more, which is precisely the direction the bounded-rationality program pushes us.

Three cautions emerged along the way, and all three return in the next lecture.

The genetic algorithm's individuals do not learn: only the population does, which makes it a model of a society rather than of a mind, and a population that never stops experimenting, as in *Exchange Rate Indeterminacy, Learning, and Experiments*, need not settle down.

The bandit showed that a classifier system's behavior is decided by its auction and its accounting rather than by the fact of having classifiers at all: the auction decides how much it experiments, and the accounting how it assigns credit.

And the bucket brigade showed that the accounting builds in a valuation of the future: with averaged strengths, a rule that sets up a payoff one link away is worth only a fraction $b/(1+b)$ of the rule that collects it.

Money as a Medium of Exchange among Artificially Intelligent Agents assembles this machinery into agents who must learn, from scratch, which good to use as money.

Their rules have conditions over $\{0, 1, \#\}$ and bid in proportion to strength with a specificity premium; an argmax auction on strength picks among them; a bucket brigade links each trade to the consumption it makes possible; and a genetic algorithm whose activity dies away over time refreshes them.

All three cautions bear directly on what those agents turn out to be able to learn.

94.8 Exercises

Exercise 94.8.1

The Hopfield network's recall is descent on the energy $E(s) = -\frac{1}{2} s^\top w s$.

Verify the descent directly.

Take a stored letter, corrupt several pixels, and record the energy at each step of the recall dynamics.

Confirm that energy never increases, and that a recall that ends at a stored letter ends at energy $-N/2$, the lowest energy any state can have.

Then prove that energy cannot increase under the synchronous update $s' = \operatorname{sgn}(ws)$ when w is the projection-rule matrix.

Finally, corrupt patterns more heavily and classify the failures: does recall land on a *different* stored letter, or on a *spurious* state that was never stored?

Compare the energies in each case, and use them to explain why the error happens.

Solution

Here is a version of `recall` that records the energy at every step.

```
def recall_with_energy(w, s0, max_iter=30):
    s = s0.copy()
    trace = [energy(w, s)]
    for _ in range(max_iter):
        s_new = np.sign(w @ s)
        s_new[s_new == 0] = 1
        trace.append(energy(w, s_new))
        if np.array_equal(s_new, s):
            break
        s = s_new
    return s, trace

rng = np.random.default_rng(1)
i = letters.index('E')
corrupt = σ[i].copy()
corrupt[rng.choice(25, 5, replace=False)] *= -1
final, trace = recall_with_energy(w_hop, corrupt)
```

```
print(f"energy along the recall path: {[round(e, 2) for e in trace]}")
print(f"monotonically non-increasing: {all(np.diff(trace) <= 1e-9)}")
print(f"recovered the intended letter 'E': {np.array_equal(final, σ[i])}")
```

```
energy along the recall path: [np.float64(-6.38), np.float64(-12.5), np.
↪float64(-12.5)]
monotonically non-increasing: True
recovered the intended letter 'E': True
```

The energy falls from its corrupted value to $-N/2 = -12.5$ and stays there.

Here is why it cannot rise.

Let $s' = \text{sgn}(ws)$ and $d = s' - s$.

Expanding the quadratic form, and using the symmetry of w,

$$E(s') - E(s) = -d^\top ws - \tfrac{1}{2}\, d^\top w d.$$

Each nonzero entry of d is $d_i = 2s'_i = 2\,\text{sgn}((ws)_i)$, so the first term is $-2\sum_i |(ws)_i|$, summed over the flipped pixels, and is never positive.

The second term is never positive either, because the projection matrix w is positive semidefinite.

So the energy never rises, and it falls strictly whenever a pixel flips against a nonzero field.

Since there are finitely many states, recall cannot cycle and must halt at a fixed point.

(Ties are broken toward $+1$, as in the code, so a pixel can only flip from $+1$ to -1 against a strictly negative field, and that strict fall rules out cycles.)

The argument uses positive semidefiniteness, which the projection rule has.

The usual Hebb rule with its diagonal set to zero does not have it, and with such weights, updating all pixels at once can cycle between two states.

Hopfield's own energy argument [Hopfield, 1982] is for updating one pixel at a time.

Now corrupt more heavily and sort the outcomes into three kinds: the intended letter, a *different* stored letter, and a spurious state that is a fixed point but was never taught.

The code also checks monotonicity on every one of the 1000 recall paths, and looks at the spurious states that sit at the lowest possible energy.

```
def classify_outcome(final, i):
    if np.array_equal(final, σ[i]):
        return "intended"
    if any(np.array_equal(final, σ[k]) for k in range(len(letters))):
        return "wrong letter"
    return "spurious"

counts = {"intended": 0, "wrong letter": 0, "spurious": 0}
wrong_energy, spurious_energy, deepest = [], [], []
always_downhill = True
for i in range(len(letters)):
    for seed in range(200):
        r = np.random.default_rng(1000*i + seed)
        corrupt = σ[i].copy()
        corrupt[r.choice(25, 7, replace=False)] *= -1
        final, trace = recall_with_energy(w_hop, corrupt)
        always_downhill &= all(np.diff(trace) <= 1e-9)
        kind = classify_outcome(final, i)
```

```
        counts[kind] += 1
        if kind == "wrong letter":
            wrong_energy.append(energy(w_hop, final))
        elif kind == "spurious":
            spurious_energy.append(energy(w_hop, final))
            if np.isclose(energy(w_hop, final), -12.5):
                deepest.append(final)

negatives = sum(any(np.array_equal(s, -σ[k]) for k in range(len(letters))) for s␣
↪in deepest)
print(f"7-pixel corruptions ({sum(counts.values())} trials): {counts}")
print(f"energy non-increasing on every path: {always_downhill}")
print(f"stored patterns all have energy {energy(w_hop, σ[0]):.1f}")
print(f"wrong-letter results: energy in [{min(wrong_energy):.1f}, {max(wrong_
↪energy):.1f}]")
print(f"spurious results:     energy in [{min(spurious_energy):.1f}, {max(spurious_
↪energy):.1f}]")
print(f"spurious results at energy -N/2: {len(deepest)}, "
      f"of which sign-reversed letters: {negatives}")
```

```
7-pixel corruptions (1000 trials): {'intended': 404, 'wrong letter': 62,
↪'spurious': 534}
energy non-increasing on every path: True
stored patterns all have energy -12.5
wrong-letter results: energy in [-12.5, -12.5]
spurious results:     energy in [-12.5, -10.4]
spurious results at energy -N/2: 41, of which sign-reversed letters: 3
```

Both failure modes occur, and at this level of corruption spurious states outnumber wrong letters by almost nine to one.

A **wrong-letter** result sits at exactly the stored energy $-N/2$: the corruption pushed the starting point clear across a basin boundary, into the domain of attraction of a *different* stored memory that is just as deep.

Energy descent converges to *a* minimum but cannot guarantee the *nearest* one when correlated patterns have interlocking basins.

A **spurious** result is a fixed point the designer never intended, created as an artifact of the storage rule.

Most spurious results are shallower than the stored patterns: they are fixed points outside the span of the letters, most of them local minima.

A minority sit at exactly $-N/2$, which requires them to lie in that span.

Some of those are photographic negatives of letters: because $\operatorname{sgn}(w(-s)) = -\operatorname{sgn}(ws)$, the vector $-\sigma$ is a fixed point whenever σ is.

Most are signed mixtures of letters that happen to be ± 1 vectors, such as $E - I + T$ and $A - E + I$.

Depth alone therefore does not identify a memory as one we asked for.

Because all of these are fixed points of the update, the dynamics cannot escape any of them.

That is why **simulated annealing** matters: adding declining random shaking lets the system hop out of a shallow spurious basin, or across a basin boundary, before settling, something pure energy descent can never do.

Exercise 94.8.2

The genetic algorithm's exploration comes from two operators, crossover and mutation.

The book argues that crossover "lies at the heart of the algorithm," while mutation alone "is a poor mechanism for injecting diversity."

Test the claim on Axelrod's game.

Write a mutation-only variant (each child is a mutated copy of one selected parent, with no crossover) and compare the fitness it reaches against the full algorithm.

Score each run by its champion's `panel_fitness`, so that every strategy faces the same Random games, and use several seeds.

The book is careful to say *when* mutation alone is weak: "when the mutation rate is set to a very low value, mutation alone is a poor mechanism for injecting diversity."

So run the comparison at a **low** mutation rate, $p_m = 0.002$, where crossover has to carry the exploration, and then at a higher one, $p_m = 0.02$.

Solution

The mutation-only variant keeps the windowed roulette wheel and drops crossover.

```
def evolve_mutation_only(N=60, generations=60, p_mut=0.002, seed=0):
    "Each child is a mutated copy of one roulette-selected parent."
    rng = np.random.default_rng(seed)
    pop = rng.integers(0, 2, (N, 70))
    for _ in range(generations):
        fit = fitness(pop, rng.integers(0, 2, (1, 150)))
        weight = fit - fit.min() + 1e-6
        cum = np.cumsum(weight)
        pop = pop[np.searchsorted(cum, rng.random(N) * cum[-1])]
        pop[rng.random(pop.shape) < p_mut] ^= 1
    fit = fitness(pop, rng.integers(0, 2, (1, 150)))
    return pop[fit.argmax()]

n_seeds = 12
for p_mut in (0.002, 0.02):
    gaps = []
    for seed in range(n_seeds):
        with_x = panel_fitness(evolve(generations=60, p_mut=p_mut, seed=seed)[0])
        without_x = panel_fitness(evolve_mutation_only(p_mut=p_mut, seed=seed))
        gaps.append(with_x - without_x)
    gaps = np.array(gaps)
    print(f"p_mut = {p_mut}: crossover minus mutation-only = {gaps.mean():.3f} "
          f"(standard error {gaps.std(ddof=1) / np.sqrt(n_seeds):.3f}); "
          f"crossover ahead in {np.sum(gaps > 0)} of {n_seeds} seeds")
```

```
p_mut = 0.002: crossover minus mutation-only = 0.070 (standard error 0.017);␣
↪crossover ahead in 11 of 12 seeds
p_mut = 0.02: crossover minus mutation-only = -0.010 (standard error 0.020);␣
↪crossover ahead in 6 of 12 seeds
```

At the low mutation rate crossover clearly wins: its champion is fitter in almost every seed.

At the higher rate the advantage shrinks to near zero: in these seeds the average gap is within one standard error of zero, and crossover is ahead in only about half of them.

The reason is the one the book gives.

With little mutation, a mutation-only population can only inch forward one rare bit-flip at a time, and it quickly loses

diversity as selection copies its few best strings.

Crossover instead recombines segments of strings that have already proved useful.

In this encoding a block of 16 consecutive table entries is a set of histories that share the most recent round's outcome, so crossover splices one string's responses to some recent outcomes onto another string's responses to others.

That is the schema story, though the schema theorem guarantees only that short, fit blocks tend to survive, not that splicing them improves the string.

When mutation generates plenty of diversity on its own, crossover's contribution is no longer pivotal, which is the book's qualification: it is at a very low mutation rate that mutation alone is weak.

Exercise 94.8.3

The two-armed bandit classifier probability-matches, which is suboptimal: it keeps pulling the worse arm a fixed fraction of the time forever.

A natural fix is to make the auction **greedier**.

Replace the strength-proportional choice rule with a softmax,

$$\pi_1 = \frac{e^{\beta S_1}}{e^{\beta S_1} + e^{\beta S_2}},$$

where β controls greediness ($\beta \to \infty$ always picks the stronger arm).

1. Implement it and show how the long-run fraction on the best arm depends on β, for $\mu = (3, 1)$.
2. How large must β be for the softmax to beat probability matching?
3. Now make the arms hard to tell apart: $\mu = (1.2, 1.0)$ with reward noise of standard deviation 2. Compare $\beta = 2$ with a very greedy $\beta = 32$ across 40 seeds. What does greed cost?

Solution

The softmax bandit changes only the line that chooses the arm.

```
def bandit_softmax(μ, β, noise_sd=0.5, T=20_000, seed=0):
    "Share of the last 5000 pulls on arm 1 under a softmax auction."
    rng = np.random.default_rng(seed)
    S = np.ones(2)                              # equal strengths, as before
    τ = np.array([1, 1])
    pulls = np.empty(T, int)
    for t in range(T):
        p1 = 1 / (1 + np.exp(-β * (S[0] - S[1])))
        i = 0 if rng.random() < p1 else 1
        reward = μ[i] + noise_sd * rng.standard_normal()
        τ[i] += 1
        S[i] += (reward - S[i]) / τ[i]
        pulls[t] = i
    return np.mean(pulls[-5000:] == 0)

μ = [3.0, 1.0]
β_star = np.log(μ[0] / μ[1]) / (μ[0] - μ[1])
print(f"probability match target = {μ[0]/(μ[0]+μ[1]):.3f},  optimal = 1.000,  "
      f"break-even β = {β_star:.3f}\n")
```

```
for β in (0.5, 1.0, 2.0, 4.0, 8.0):
    theory = 1 / (1 + np.exp(-β * (μ[0] - μ[1])))
    print(f"β = {β:>4}: fraction on best arm = {bandit_softmax(μ, β):.3f}  "
          f"(limit {theory:.3f})")
```

```
probability match target = 0.750,  optimal = 1.000,  break-even β = 0.549

β =  0.5: fraction on best arm = 0.732  (limit 0.731)
β =  1.0: fraction on best arm = 0.888  (limit 0.881)
β =  2.0: fraction on best arm = 0.980  (limit 0.982)
β =  4.0: fraction on best arm = 1.000  (limit 1.000)
β =  8.0: fraction on best arm = 1.000  (limit 1.000)
```

The softmax keeps every arm's probability bounded away from zero, so each arm is pulled infinitely often and its strength converges to μ_i, as in the proportional case.

The long-run share on arm 1 is therefore $1/(1+e^{-\beta(\mu_1-\mu_2)})$, the "limit" column, which the simulations match.

It beats the probability-matching share $\mu_1/(\mu_1+\mu_2)$ exactly when $e^{-\beta(\mu_1-\mu_2)} < \mu_2/\mu_1$, that is, when

$$\beta > \frac{\ln(\mu_1/\mu_2)}{\mu_1-\mu_2} = \frac{\ln 3}{2} \approx 0.55.$$

So $\beta = 0.5$ does slightly *worse* than probability matching, and from $\beta = 1$ on the softmax does better, approaching the optimal policy as β rises.

Unlike the proportional rule, the softmax depends on the difference $\mu_1 - \mu_2$, not the ratio, so the threshold depends on the units of the payoffs.

Greed has a cost when the arms are hard to tell apart.

```
for β in (2.0, 32.0):
    shares = np.array([bandit_softmax([1.2, 1.0], β, noise_sd=2.0, seed=s)
                       for s in range(40)])
    print(f"β = {β:>4}: mean share on the better arm {shares.mean():.3f}; "
          f"runs ending mostly on the worse arm: {np.sum(shares < 0.5)} of 40")
```

```
β =  2.0: mean share on the better arm 0.597; runs ending mostly on the worse␣
↪arm: 0 of 40
β = 32.0: mean share on the better arm 0.675; runs ending mostly on the worse␣
↪arm: 13 of 40
```

With $\beta = 32$ the auction concentrates quickly, and in these 40 runs that raises the average share on the better arm above that of $\beta = 2$.

The gain is not robust, though: over 400 seeds the two averages are 0.63 and 0.60, about one standard error apart.

And in about a third of the runs the auction concentrates on the *worse* arm.

A few unlucky early payoffs make the better arm look worse, after which it is almost never pulled, so its strength is almost never corrected: the system **locks in**.

With $\beta = 2$ the system never locks onto the worse arm, but it never concentrates either.

The softmax is the Boltzmann distribution of statistical mechanics with temperature $1/\beta$.

Getting both benefits requires starting hot and cooling slowly, letting $\beta_t \to \infty$ over time, which is simulated annealing again.

The argmax auction of *Money as a Medium of Exchange among Artificially Intelligent Agents* is the $\beta = \infty$ limit, which locks in whenever it is wrong; that is why its classifier systems need separate operators for experimentation, whose activity is made to decline over time.

So the exercise makes concrete why the *auction* and the *accounting*, and not just the fact of having classifiers, determine how well a classifier system performs.

Exercise 94.8.4

Return to Bray's market (*From relaxation to learning: Bray's market*, *Least Squares Learning in Self-Referential Models*), in which the price is

$$p_t = a + b\,\bar{\beta}_t + u_t,$$

where $\bar{\beta}_t$ is the average forecast of p_t and u_t is IID $N(0, \sigma_u^2)$.

Least squares learners converge to the rational expectations price $a/(1-b)$ when $b < 1$ and do not when $b > 1$ [Bray, 1982].

Replace Bray's representative least squares learner by a population of $N = 60$ forecasters, each a 16-bit string decoded into a forecast in $[0, 20]$.

Each period:

- the price is set by the population's average forecast;
- each forecaster's fitness is minus its squared forecast error;
- the population breeds next period's forecasters by this lecture's genetic algorithm (windowed roulette, single-point crossover, mutation at rate 0.01), followed by the election operator of *Exchange Rate Indeterminacy, Learning, and Experiments*: a child replaces its parent only if it would have forecast this period's price better.

Set $a = (1-b)p^*$, so that the rational expectations price is $p^* = 10$ whatever b is, and take $\sigma_u = 0.1$.

1. With $b = 0.7$, show that across 20 seeds the average forecast settles at the rational expectations price within 300 periods.
2. Try $b = 1.05$, slightly above Bray's bound. What happens, and why?
3. Sargent's book reports that, in her 1991 thesis, Jasmina Arifovic found that a population of this kind "can sometimes learn its way to a rational expectations equilibrium even when Bray's necessary condition $b < 1$ is violated" [Sargent, 1993]. In what sense does your experiment agree? (Arifovic [1994] studies genetic-algorithm learning in the cobweb model.)

Solution

Here is the population of forecasters.

```
def decode(pop, lo=0.0, hi=20.0):
    "Read each row of bits as a binary number and map it into [lo, hi]."
    L = pop.shape[1]
    return lo + (hi - lo) * (pop @ 2.0 ** np.arange(L - 1, -1, -1)) / (2**L - 1)

def ga_forecasters(b, p_star=10.0, σ_u=0.1, N=60, L=16, T=300, p_mut=0.01, seed=0):
    "Path of the average forecast in Bray's market with a GA population."
    rng = np.random.default_rng(seed)
    a = (1 - b) * p_star
    pop = rng.integers(0, 2, (N, L))
    path = np.empty(T)
    for t in range(T):
        β = decode(pop)
```

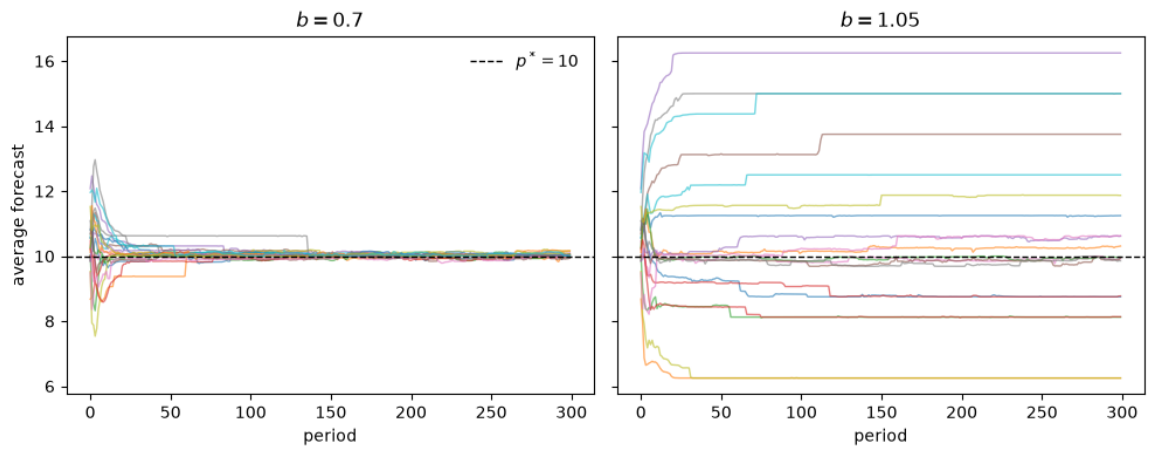
```
        path[t] = β.mean()
        p = a + b * β.mean() + σ_u * rng.standard_normal()
        fit = -(p - β)**2
        weight = fit - fit.min() + 1e-6
        cum = np.cumsum(weight)
        parents = pop[np.searchsorted(cum, rng.random(N) * cum[-1])]
        mom, dad = parents[0::2], parents[1::2]
        left = np.arange(L) < rng.integers(1, L, (N // 2, 1))
        kids = np.vstack([np.where(left, mom, dad), np.where(left, dad, mom)])
        kids[rng.random(kids.shape) < p_mut] ^= 1
        parents = np.vstack([mom, dad])
        better = (p - decode(kids))**2 < (p - decode(parents))**2    # election
        pop = np.where(better[:, None], kids, parents)
    return path

paths = {b: np.array([ga_forecasters(b, seed=s) for s in range(20)]) for b in (0.7,
 ↪ 1.05)}
for b, P in paths.items():
    gap = np.abs(P - 10)
    print(f"b = {b}: median |average forecast - p*| at t = 10, 100, 299: "
          f"{np.median(gap[:, 10]):.3f}, {np.median(gap[:, 100]):.3f}, "
          f"{np.median(gap[:, -1]):.3f};  runs within 0.25 of p* at the end: "
          f"{np.sum(gap[:, -1] < 0.25)} of 20")
```

```
b = 0.7: median |average forecast - p*| at t = 10, 100, 299: 0.347, 0.055, 0.
 ↪046;  runs within 0.25 of p* at the end: 20 of 20
b = 1.05: median |average forecast - p*| at t = 10, 100, 299: 1.313, 1.398, 1.
 ↪557;  runs within 0.25 of p* at the end: 4 of 20
```

The figure shows all forty paths.

```
fig, axes = plt.subplots(1, 2, figsize=(10, 4), sharey=True)
for ax, (b, P) in zip(axes, paths.items()):
    ax.plot(P.T, lw=1, alpha=0.6)
    ax.axhline(10, color='k', ls='--', lw=1, label="$p^* = 10$")
    ax.set_title(f"$b = {b}$")
    ax.set_xlabel("period")
axes[0].set_ylabel("average forecast")
axes[0].legend(frameon=False)
plt.tight_layout()
plt.show()
```



At $b = 0.7$ every one of these runs settles at the rational expectations price.

At $b = 1.05$ only a few runs end near it.

The others are pushed away in the first few periods, while the population is still diverse.

They then stall, apart from occasional jumps, at a forecast that is wrong in the same direction every period.

The T-map of *One recursion, many learners* explains the direction.

A population whose average forecast is $\bar{\beta}$ induces the price $T(\bar{\beta}) = a + b\bar{\beta}$ on average, and selection and election favor forecasts near the realized price.

Since $T(\bar{\beta}) - \bar{\beta} = (b - 1)(\bar{\beta} - p^*)$, the favored forecasts lie on the side of p^* when $b < 1$ and on the far side when $b > 1$.

The genetic algorithm lies outside the theory of stochastic approximation, but here it moves in the direction that the theory's mean dynamics would predict.

Why does it stall rather than run off to the edge of $[0, 20]$?

Because b is close to 1, the push is weak: a population 1.5 away from p^* makes systematic forecast errors of only $0.05 \times 1.5 = 0.075$ a period, less than the noise.

And once the strings are nearly identical, crossover creates nothing new, and many small moves of the forecast require several bits to flip at once (a "Hamming cliff" of binary encoding), so the population can freeze.

So the experiment agrees with the book in a weak sense.

When $b > 1$ the population ends near the rational expectations price only in runs that happened to stall near it.

The equilibrium does not attract it, and it does not attract least squares either, which diverges when $b > 1$, though very slowly when b is this close to 1 (*Least Squares Learning in Self-Referential Models*).

The election operator matters even when $b = 0.7$.

Remove it, and the population keeps its children however badly they forecast, and it settles much less tightly on p^*: as in *Exchange Rate Indeterminacy, Learning, and Experiments*, election is what lets a genetic population learn, by making it homogeneous.

Exercise 94.8.5

Extend the two-link chain of *The bucket brigade* by giving the system a choice at the first stage.

A *myopic* rule m collects $r_m = 1$ at once and ends the episode.

The setup rule s collects nothing but moves the system to a second stage, where the payoff rule p collects R.

To separate the accounting from the auction, choose m or s at random with equal probability in every episode, and use the averaging accounting with $b = 0.1$.

1. Suppose there are no backward payments: the payoff rule's bid is paid to no one. What happens to the setup rule's strength?
2. With the payoff rule paying its bid to the setup rule, find the stationary strengths of m and s, and the smallest ratio R/r_m at which an argmax auction would choose the setup rule. Check your answer by simulation.
3. Repeat part 2 with Holland's cumulative balances. Which accounting is more patient?

Solution

Here is the chain with a choice at its first stage.

```
def two_stage(R, r_m=1.0, b=0.1, T=20_000, accounting='average', backward=True,␣
↪seed=0):
    "Strengths of [myopic, setup, payoff] when the first-stage rule is chosen at␣
↪random."
    rng = np.random.default_rng(seed)
    S, τ = np.zeros(3), np.zeros(3)
    for _ in range(T):
        if rng.random() < 0.5:
            acts = [(0, r_m)]                                   # myopic: paid at␣
↪once
        else:
            acts = [(1, b * S[2] if backward else 0.0),         # setup: paid the␣
↪next bid
                    (2, R)]                                     # payoff: paid by␣
↪the world
        for k, receipts in acts:
            τ[k] += 1
            if accounting == 'average':
                S[k] += (receipts - (1 + b) * S[k]) / τ[k]
            else:
                S[k] += receipts - b * S[k]
    return S

b = 0.1
print(f"predicted threshold R / r_m = (1 + b) / b = {(1 + b) / b:.1f}\n")
for R in (5.0, 10.0, 12.0, 20.0):
    rows = {'no backward payment': two_stage(R, backward=False),
            'average': two_stage(R),
            'balance': two_stage(R, accounting='balance')}
    for name, S in rows.items():
        winner = 'setup' if S[1] > S[0] else 'myopic'
        print(f"R = {R:>4}, {name:>19}: S_m = {S[0]:7.3f}, S_s = {S[1]:7.3f}  ->
↪{winner}")
    print()
```

```
predicted threshold R / r_m = (1 + b) / b = 11.0

R =  5.0, no backward payment: S_m =   0.909, S_s =   0.000  -> myopic
R =  5.0,             average: S_m =   0.909, S_s =   0.413  -> myopic
R =  5.0,             balance: S_m =  10.000, S_s =  50.000  -> setup
R = 10.0, no backward payment: S_m =   0.909, S_s =   0.000  -> myopic
R = 10.0,             average: S_m =   0.909, S_s =   0.826  -> myopic
R = 10.0,             balance: S_m =  10.000, S_s = 100.000  -> setup

R = 12.0, no backward payment: S_m =   0.909, S_s =   0.000  -> myopic
R = 12.0,             average: S_m =   0.909, S_s =   0.992  -> setup
R = 12.0,             balance: S_m =  10.000, S_s = 120.000  -> setup
R = 20.0, no backward payment: S_m =   0.909, S_s =   0.000  -> myopic
R = 20.0,             average: S_m =   0.909, S_s =   1.653  -> setup
R = 20.0,             balance: S_m =  10.000, S_s = 200.000  -> setup
```

Part 1. Without backward payments the setup rule's receipts are always zero, so its strength never rises above zero, however large R is.

No auction based on strength would ever choose it.

Part 2. With averaging, the stationary strengths solve $(1+b)S_m = r_m$, $(1+b)S_p = R$ and $(1+b)S_s = bS_p$, so

$$S_m = \frac{r_m}{1+b}, \qquad S_s = \frac{b}{(1+b)^2}\,R.$$

The setup rule wins exactly when $R/r_m > (1+b)/b = 11$.

The simulation confirms it: the setup rule loses at $R = 10$ and wins at $R = 12$.

Part 3. With cumulative balances, $S_m = r_m/b$ and $S_s = S_p = R/b$, so the setup rule wins whenever $R > r_m$.

Holland's accounting is patient; the averaging accounting is very impatient.

A classifier system that averages its strengths with $b = 0.1$ prefers 1 now to 10 one step later.

The bucket brigade's discount factor shows how the same arithmetic, applied to the chain that links trading to consuming, makes the agents of *Money as a Medium of Exchange among Artificially Intelligent Agents* effectively very impatient.

CHAPTER

NINETYFIVE

MONEY AS A MEDIUM OF EXCHANGE AMONG ARTIFICIALLY INTELLIGENT AGENTS

Contents

- *Money as a Medium of Exchange among Artificially Intelligent Agents*
  - *Overview*
  - *The Kiyotaki-Wright environment*
  - *Classifier systems for the Kiyotaki-Wright environment*
  - *Implementation*
  - *Economy A1.1: does a medium of exchange emerge?*
  - *Economy A2.1: when theory predicts speculation*
  - *The bucket brigade's discount factor*
  - *Economy B.1: selecting between two equilibria*
  - *The genetic algorithm*
  - *Economy A1.2: learning from random rules*
  - *Economy C: fiat money*
  - *Economy D: five goods, five types*
  - *Concluding remarks*
  - *Exercises*

95.1 Overview

Kiyotaki and Wright [Kiyotaki and Wright, 1989] studied an economy in which there is no double coincidence of wants.

Trade can occur only if some good is accepted not because it is wanted, but because it can be passed along later.

A good that plays that role is a **medium of exchange**.

Kiyotaki and Wright characterized *stationary Nash equilibria* of such an economy under the assumption that agents are fully rational: agents know the distribution of goods across their trading partners and best-respond to it.

Marimon, McGrattan and Sargent [Marimon *et al.*, 1990] asked a different question.

Suppose we drop rationality altogether and put in its place a collection of **artificially intelligent agents** who begin with arbitrary, even random, rules of thumb, and who adapt those rules only by keeping score of what has paid off in the past.

Will such agents *learn* to use a medium of exchange?

And when the economy has two equilibria, or when its only equilibrium asks agents to be patient, what will they settle on?

These are instances of the question that *A Peculiar Definition of Bounded Rationality* asked of rational expectations (*The knowledge that rational expectations imputes*): can adaptive agents find an equilibrium, and which one?

The learning device that Marimon, McGrattan and Sargent used is John Holland's **classifier system** [Holland, 1975, Holland *et al.*, 1986].

Networks, Genetic Algorithms, and Classifier Systems introduced classifier systems and the bucket brigade in the abstract; here they are assembled into agents.

Each type of agent gets a population of if-then rules, an auction that decides which rule acts, and an accounting system that credits rules that lead to good outcomes and debits rules that lead to bad ones.

When the rules are generated at random, a **genetic algorithm** breeds new rules and retires old ones.

This lecture rebuilds their computational experiments and runs each of them for many random seeds, so that we can say how often each outcome occurs.

The main findings are:

1. With a complete list of rules, agents in Economy A1 learn its unique stationary equilibrium in every run: good 1, the cheapest good to store, circulates as money. Starting from random rules they get most of the way, as in the paper.
2. In Economy A2, where patient agents would *speculate* by accepting an expensive good, the artificial agents never do. Most runs settle on the fundamental pattern of trade, which is not an equilibrium of that economy, some only after thousands of periods stuck in a pattern with too little trade, and a few stay stuck for good. We show why: the bucket brigade makes the agents behave as if they discounted the future at a factor of about 0.016.
3. In Economy B, which has a fundamental and a speculative equilibrium, the agents select the fundamental one in most runs and the speculative one in none.
4. In nearly every run, an intrinsically worthless, costlessly stored object – **fiat money** – is accepted in trade by agents who have to discover its usefulness for themselves.
5. In an economy with five goods and five types, for which the authors had no analytical characterization of equilibrium, the algorithm proposes a pattern of trade, with qualifications that we describe.

Let's start with some imports.

```
import numpy as np
import pandas as pd
import matplotlib.pyplot as plt
from dataclasses import dataclass, replace
from typing import NamedTuple
from numba import njit
```

95.2 The Kiyotaki-Wright environment

There are three types of agents, indexed by $i = 1, 2, 3$, and three goods, indexed by $k = 1, 2, 3$.

A type i agent derives utility only from consuming good i.

He has a technology for producing good i^*, where $i^* \neq i$.

In Kiyotaki and Wright's *model A*, the production pattern is the "Wicksell triangle"

type i	produces i^*	consumes
1	2	1
2	3	2
3	1	3

so there is no double coincidence of wants: the agent who has what you want never wants what you have.

All goods are indivisible and each agent can store exactly one unit of exactly one good from one period to the next.

Storing good k for one period costs s_k, with

$$s_3 > s_2 > s_1 > 0.$$

There are equal numbers A_i of agents of each type, so $A = 3A_i$ in total.

Each period every agent is randomly matched with exactly one other agent, without regard to type.

Write x_{at} for the good that agent a carries into date t and $\rho_t(a)$ for the agent with whom a is matched.

The **pre-trade state** of agent a is the pair

$$z_{at} = (x_{at},\ x_{\rho_t(a)t}).$$

Each period each agent makes two decisions in sequence.

First, having seen z_{at}, he decides whether to propose a trade,

$$\lambda_{at} = \begin{cases} 1 & \text{propose to trade } x_{at} \text{ for } x_{\rho_t(a)t} \\ 0 & \text{refuse} \end{cases}$$

Trade occurs if and only if $\lambda_{at}\lambda_{\rho_t(a)t} = 1$, so post-trade holdings are

$$x^+_{at} = (1 - \lambda_{at}\lambda_{\rho_t(a)t})x_{at} + \lambda_{at}\lambda_{\rho_t(a)t}x_{\rho_t(a)t}. \tag{95.1}$$

Second, he decides whether to consume what he is left holding,

$$\gamma_{at} = \begin{cases} 1 & \text{consume } x^+_{at} \\ 0 & \text{carry } x^+_{at} \text{ into } t+1 \end{cases}$$

If he consumes he immediately produces good $f(a) = i^*$ and carries that into $t + 1$.

Hence

$$x_{a,t+1} = \gamma_{at} f(a) + (1 - \gamma_{at})x^+_{at}. \tag{95.2}$$

The one-period net payoff is

$$U_a(\gamma_{at}) = \gamma_{at}[u_i(x^+_{at}) - s(f(a))] - (1 - \gamma_{at})\, s(x^+_{at}), \tag{95.3}$$

where $u_i(k) = u_i > 0$ if $k = i$ and $u_i(k) = 0$ otherwise.

Note that an agent is *permitted* to consume a good he does not want; he simply gets zero utility from it while still producing and paying to store $f(a)$.

This "consume anything" feature has a consequence worth stating once.

An agent who swaps his own production good for some other good and then eats it ends the period exactly where refusing to trade would have left him: holding his production good and paying its storage cost.

Such trades are payoff-neutral, so nothing in the accounting system described below pushes them out, and they will turn up in several of our economies.

Note

Kiyotaki and Wright ranked payoff streams by expected discounted utility.

Marimon, McGrattan and Sargent instead assume that each agent cares about his **long-run average** utility, and they build the accounting system inside a classifier system out of running averages for that reason.

We shall see in *The bucket brigade's discount factor* that the agents this accounting system produces nevertheless behave as if they discounted the future very heavily.

95.2.1 Two equilibria

Since agents' payoffs depend on what other agents do, the model can have more than one stationary equilibrium.

Kiyotaki and Wright characterize equilibria by a set of probabilities, of which the most useful for us is

$$\pi_{it}^h(k) = \text{probability that a type } i \text{ agent holds good } k \text{ at } t.$$

In the **fundamental equilibrium**, good 1 – the cheapest to store – is the general medium of exchange:

- type 1 agents always hold good 2 (which they produce) and trade it for good 1;
- type 3 agents always hold good 1 and trade it for good 3;
- type 2 agents hold good 1 half the time and good 3 half the time.

So the equilibrium holding probabilities are

	$k=1$	$k=2$	$k=3$
$i=1$	0	1	0
$i=2$	0.5	0	0.5
$i=3$	1	0	0

Type 2 agents accept good 1 even though they never consume it: they use it as money.

In the **speculative equilibrium**, type 1 agents additionally accept good 3 – the *most* expensive good to store – because they expect to trade it away quickly for good 1.

The reason is that a type 1 agent holding good 2 can get good 1 only from a type 2 agent, whereas a type 1 agent holding good 3 can get it from a type 3 agent, who wants good 3 and holds good 1.

Each partner type is met with probability $1/3$, so holding good 3 instead of good 2 raises the chance of acquiring good 1 each period by $(\pi_3^h(1) - \pi_2^h(1))/3$, at an extra storage cost of $s_3 - s_2$.

Kiyotaki and Wright show that, in the limit as the discount rate goes to zero, the fundamental equilibrium is the unique stationary equilibrium if

$$s_3 - s_2 > (\pi_3^h(1) - \pi_2^h(1))\tfrac{1}{3}u_1, \tag{95.4}$$

evaluated at fundamental holdings, and that the speculative equilibrium is the unique one when the inequality is reversed at speculative holdings.

Footnote 7 of the paper prints $\pi_1^h(3) - \pi_1^h(2)$, which transposes the indices: with type 1's own holdings the right side would be $-u_1/3$ and the inequality could never fail.

At the fundamental equilibrium $\pi_3^h(1) = 1$ and $\pi_2^h(1) = 1/2$, and at the speculative equilibrium (tabulated below) $\pi_2^h(1) = 0.586$.

The following function solves (95.4) for u_1; the argument β will be used in *The bucket brigade's discount factor*.

```
s_A = np.array([0.1, 1.0, 20.0])           # storage costs in Economy A


def speculation_threshold(s, π31, π21, β=1.0):
    "The u_1 above which a type 1 agent with discount factor β prefers good 3."
    return 3 * (s[2] - s[1]) / (β * (π31 - π21))


u_fund = speculation_threshold(s_A, π31=1.0, π21=0.5)     # at fundamental holdings
u_spec = speculation_threshold(s_A, π31=1.0, π21=0.586)   # at speculative holdings
print(f"fundamental equilibrium requires u_1 < {u_fund:.1f}")
print(f"speculative equilibrium requires u_1 > {u_spec:.1f}")
```

```
fundamental equilibrium requires u_1 < 114.0
speculative equilibrium requires u_1 > 137.7
```

Economy A1 has $u_1 = 100$, which satisfies the fundamental condition, though the cutoff is only 14 percent above it.

Economy A2 has $u_1 = 500$, which violates it, so for patient agents the unique stationary equilibrium of A2 is the speculative one.

Each of these economies thus has one stationary equilibrium for patient agents, and the question it poses is whether that equilibrium is learnable.

Economy B below has two, and poses the question of selection.

95.2.2 The economies

The following table, modeled on the paper's Table 4, lists the economies we simulate.

All have fifty agents of each type and equal utilities $u_i = u$ for every type.

economy	type i produces	storage costs	u	bids $b_{11}, b_{12}, b_{21}, b_{22}$	initial rules	equilibrium for patient agents
A1.1	2, 3, 1	0.1, 1, 20	100	0.025, 0.025, 0.25, 0.25	all 72 + 12	fundamental
A2.1	2, 3, 1	0.1, 1, 20	500	0.025, 0.025, 0.25, 0.25	all 72 + 12	speculative
B.1	3, 1, 2	1, 4, 9	100	0.25, 0.25, 0.25, 0.25	all 72 + 12	fundamental and speculative
A1.2	2, 3, 1	0.1, 1, 20	100	0.025, 0.025, 0.25, 0.25	72 + 12 random	fundamental
A2.2	2, 3, 1	0.1, 1, 20	500	0.025, 0.025, 0.25, 0.25	72 + 12 random	speculative
B.2	3, 1, 2	1, 4, 9	100	0.025, 0.025, 0.25, 0.25	72 + 12 random	fundamental and speculative
C	2, 3, 1	9, 14, 29; fiat 0	100	0.025, 0.025, 0.025, 0.25	150 + 20 random	fundamental
D	3, 4, 5, 1, 2	1, 4, 9, 20, 30	200	0.025, 0.025, 0.25, 0.25	180 + 20 random	not characterized

The bids are explained below; "all 72 + 12" means a complete enumeration of exchange and consumption rules.

A2.2 and B.2 appear in the exercises.

The paper's Table 4 lists Economy C's storage costs as 0.1, 20 and 70, but its Table 15, which reports the simulation, lists 9, 14 and 29, and we use those.

95.3 Classifier systems for the Kiyotaki-Wright environment

Networks, Genetic Algorithms, and Classifier Systems set out the general machinery in *Classifier systems*: the alphabet $\{0, 1, \#\}$, matching, specificity, and auctions among rules with strengths.

Its *The bucket brigade* showed how a rule that sets up a later reward gets paid for it.

Here we record only what is specific to Marimon, McGrattan and Sargent: how goods are encoded, the two decisions each agent makes every period, and the two strength equations that link them.

95.3.1 Encoding

Each type has two classifier systems, one for exchange decisions and one for consumption decisions.

The paper's Table 1 encodes goods with three positions each: a one-hot code for "good k" and a code with a 0 in position k for "not good k".

```
WILD = -1                                   # the "don't care" symbol #


def table1(n_goods):
    "The paper's Table 1: one-hot codes for 'good k', then codes for 'not good k'."
    good = np.eye(n_goods, dtype=np.int8)
    not_good = np.full((n_goods, n_goods), WILD, dtype=np.int8)
    np.fill_diagonal(not_good, 0)
    return np.vstack([good, not_good])
```

(continues on next page)

(continued from previous page)

```
def code_string(c):
    return ''.join('#' if b == WILD else str(int(b)) for b in c)


def rule_string(cond, action, n_goods):
    "Print a classifier the way the paper does, for example '010 100 -> 1'."
    s = code_string(cond)
    if len(cond) == 2 * n_goods:
        s = s[:n_goods] + ' ' + s[n_goods:]
    return f"{s} -> {int(action)}"


pd.DataFrame({'code': [code_string(c) for c in table1(3)],
              'meaning': [f'good {k}' for k in (1, 2, 3)]
                         + [f'not good {k}' for k in (1, 2, 3)]})
```

	code	meaning
0	100	good 1
1	010	good 2
2	001	good 3
3	0##	not good 1
4	#0#	not good 2
5	##0	not good 3

An **exchange classifier** is a string of length 7: three positions for the agent's own holding, three for the partner's holding, and a binary action (1 = propose trade, 0 = refuse).

The paper's two examples are

```
100 001 -> 1     "if I hold good 1 and my partner holds good 3, propose to trade"
100 ##0 -> 0     "if I hold good 1 and my partner does not hold good 3, refuse"
```

A **consumption classifier** is a string of length 4: three positions for the post-trade holding x_{at}^{+} and one action digit (1 = consume).

A *complete enumeration* pairs every condition in Table 1 with both actions.

That gives $6 \times 6 \times 2 = 72$ exchange classifiers and $6 \times 2 = 12$ consumption classifiers, enough to express every Markov strategy.

Table 1 has no condition meaning "any good", so the enumeration contains no rule such as "eat whatever you hold".

When rules are drawn at random, each position is 0, 1 or # with equal probability.

There are then $3^6 \times 2 = 1458$ possible exchange strings and $3^3 \times 2 = 54$ consumption strings, and many of them match no state at all: `110`, for instance, is no good's code.

With four objects (Economy C) or five goods (Economy D) the codes are one-hot strings of length four or five.

For C that gives the paper's count of $3^8 \times 2 = 13{,}122$ possible exchange strings.

The paper does not say how it encoded the goods of Economy D, so the five-position code is our choice.

A rule's **specificity** is $\sigma = 1/(1 + \text{number of } \#\text{'s})$.

An exchange rule about two particular goods has $\sigma = 1$, a rule like `100 ##0` has $\sigma = 1/3$, and a rule about two "not" conditions has $\sigma = 1/5$.

95.3.2 The auction

Attached to classifier e at date t is a **strength** $S_e^a(t)$.

Given the state z_{at}, let $M_e(z_{at})$ be the set of exchange classifiers whose conditions z_{at} satisfies.

The classifier that acts is the strongest matched one,

$$e_t(z_{at}) = \arg\max \{S_e^a(t) : e \in M_e(z_{at})\}, \tag{95.5}$$

and the consumption classifier is chosen the same way from the classifiers that match x_{at}^+.

This is the paper's equation (6); we break ties at random, a detail the paper does not specify.

Each matched classifier also has a **bid**, a fraction of its strength,

$$b_1(e) = b_{11} + b_{12}\sigma_e \qquad \text{and} \qquad b_2(c) = b_{21} + b_{22}\sigma_c$$

for exchange and consumption classifiers, respectively.

Under (95.5) the bid plays no role in deciding who wins: it only determines the payments described below.

The paper's verbal description calls the winner "the highest bidder" and says that bids that rise with specificity "favor specific rules over more general rules".

That would be true of an auction on $b \cdot S$ among rules of positive strength, but under (95.5) a larger bid fraction pulls a rule's strength at a rest point *toward zero*, as the note below shows.

So among rules whose downstream payments are positive, higher bids lower stationary strength and favor general rules; only among rules whose downstream payments are negative do they favor specific rules, as the paper says.

The argmax auction is the greedy limit of the softmax auction in *Networks, Genetic Algorithms, and Classifier Systems*'s two-armed bandit exercise (Exercise **??**), in which matched rule r wins with probability proportional to $e^{\eta S_r}$ and $\eta \to \infty$ gives the argmax; that lecture calls the intensity β, a letter we reserve here for discount factors.

The paper's footnote 5 mentions such a stochastic auction as an alternative.

A greedy auction avoids probability matching, but it leaves all experimentation to the initial strengths and to the operators that create new rules, a choice that will matter in Economy A2.

95.3.3 The bucket brigade

Strengths are updated by a system of internal payments that Holland calls a *bucket brigade*.

Only the classifier that wins an auction changes its strength, so each classifier carries a counter $\tau_e^a(t)$ recording the number of auctions it has won, initialized at 1.

Payments flow as follows.

- The external payoff $U_a(\gamma_t)$ goes to the winning **consumption** classifier at t.
- The winning consumption classifier at t pays its bid to the winning **exchange** classifier at t, which created the state that gave it a chance to act.
- The winning exchange classifier at t pays its bid to the winning **consumption** classifier at $t-1$, which set up the state z_{at}.

The resulting laws of motion are the paper's equations (12) and (13),

$$S^a_{c,\tau_c(t)} = S^a_{c,\tau_c(t)-1} - \frac{1}{\tau_c(t)-1}\Big[(1+b_2(c))S^a_{c,\tau_c(t)-1} - \sum_e I_e^a(t) b_1(e) S^a_{e,\tau_e(t)} - U_a(\gamma_{ct})\Big] \tag{95.6}$$

$$S^a_{e,\tau_e(t)+1} = S^a_{e,\tau_e(t)} - \frac{1}{\tau_e(t)}\Big[(1+b_1(e))S^a_{e,\tau_e(t)} - \sum_c I^a_c(t) b_2(c) S^a_{c,\tau_c(t)}\Big] \tag{95.7}$$

where $I^a_e(t)$ and $I^a_c(t)$ are indicators for winning the auction at t.

So the consumption classifier updated at date t is the one that won at $t-1$, and each bid travels one step back along the chain

$$\cdots \to c_{t-1} \to e_t \to c_t \to e_{t+1} \to \cdots$$

An exchange classifier that proposes a trade which is *not* reciprocated is not charged, its counter does not advance, and it pays nothing backward: the agent learns nothing from an offer that was refused.

All agents of a given type share one classifier system, as in the paper.

The paper runs (95.6)-(95.7) in "synthetic time", counting matches rather than periods, so read literally a consumption classifier is paid by the next exchange classifier of its *type* to act, which usually belongs to a different agent.

We instead link each agent's own consecutive decisions, which is what the equations describe for a single agent, and the reading under which a decision is credited with its own consequences.

Note

Equations (95.6)-(95.7) make strength a running **average** of past net receipts rather than a cumulative total, which is what Holland's original specification used; the paper's footnote 8 says the averaging was chosen so that strengths could converge.

With $1/\tau$ gains these are stochastic approximation recursions, the same decreasing-gain recursion that drives the learners of *Adaptive Agents in an Overlapping Generations Monetary Economy* and *Learning, Approximation, and Equilibrium Computation* (*One recursion, many learners*).

So any limit point must satisfy the paper's equation (17),

$$\mathbb{E}\Big[(1+b_2(c))S_c - \sum_e I_e b_1(e) S_e - U(\gamma_c)\Big] = 0, \qquad \mathbb{E}\Big[(1+b_1(e))S_e - \sum_c I_c b_2(c) S_c\Big] = 0.$$

Marimon, McGrattan and Sargent call strengths solving these equations *stationary*, and say a stationary Nash equilibrium is *supported* when the rules that win auctions at stationary strengths are exactly the rules that implement equilibrium behavior.

The second equation gives $S_e = \mathbb{E}[b_2(c)S_c]/(1+b_1(e))$, so a larger bid fraction $b_1(e)$ shrinks a rule's stationary strength toward zero: it lowers S_e when $\mathbb{E}[b_2(c)S_c] > 0$ and raises it when $\mathbb{E}[b_2(c)S_c] < 0$.

One more feature of the system drives much of what follows.

All strengths start at zero, while nearly every net payoff an agent experiences is negative, because storage is costly.

So an untried rule beats every rule whose experience so far has been net costly, and under the argmax auction an agent sweeps through its untried rules early on.

This *optimistic initialization* is the main way an agent experiments when it starts from a complete enumeration.

It has a failure mode.

A rule that has never acted, or whose offers were all refused so that it learned nothing, keeps its strength of zero, and it loses every auction to a rival whose strength happens to sit slightly above zero, for as long as that rival's average stays there.

95.4 Implementation

The state spaces are tiny, so for every rule we precompute the set of states it matches and store it as the bits of an integer: state (x, y) of an exchange rule is bit $xG + y$, where G is the number of goods.

The simulation loops run under Numba's `njit`.

Classifier systems are stored as parallel arrays indexed by type and rule.

```
class Rules(NamedTuple):
    "Classifier systems for all types, stored as arrays indexed [type, rule]."
    cond: np.ndarray        # condition part, entries in {0, 1, WILD}
    action: np.ndarray      # action part: 1 = trade (or consume), 0 = refuse
    strength: np.ndarray    # S, a running average of net receipts
    used: np.ndarray        # the counter tau, initialized at 1
    match: np.ndarray       # bit s is set if the rule matches state s
    sigma: np.ndarray       # specificity, 1 / (1 + number of #'s)
    born: np.ndarray        # a stamp that changes whenever the rule is replaced


@njit
def goods_matched(field):
    "Bitmask of the goods whose one-hot code satisfies a condition field."
    n_goods = field.shape[0]
    mask = 0
    for g in range(n_goods):
        ok = True
        for j in range(n_goods):
            if field[j] != WILD and field[j] != (1 if j == g else 0):
                ok = False
        if ok:
            mask |= 1 << g
    return mask


@njit
def refresh(R, i, r, n_goods):
    "Record which states rule r of type i matches, and its specificity."
    c = R.cond[i, r]
    if c.shape[0] == n_goods:                    # consumption rule: state = holding
        m = goods_matched(c)
    else:                                        # exchange rule: state = (own,␣
↪partner)
        own, partner = goods_matched(c[:n_goods]), goods_matched(c[n_goods:])
        m = 0
        for x in range(n_goods):
            for y in range(n_goods):
                if (own >> x) & 1 and (partner >> y) & 1:
                    m |= 1 << (x * n_goods + y)
    R.match[i, r] = m
    R.sigma[i, r] = 1.0 / (1.0 + np.sum(c == WILD))


@njit
def set_rule(R, i, r, cond, action, strength, n_goods, stamp):
    "Overwrite rule r of type i with a new classifier, whose counter starts at 1."
    R.cond[i, r] = cond
    R.action[i, r] = action
```

(continues on next page)

(continued from previous page)

```
    R.strength[i, r] = strength
    R.used[i, r] = 1
    stamp[0] += 1
    R.born[i, r] = stamp[0]
    refresh(R, i, r, n_goods)
```

A complete enumeration pairs every condition in Table 1 with both actions, while a random population draws every position independently.

Following the paper, all initial strengths are zero.

```
def make_rules(cond, action, n_goods):
    n_types, n = action.shape
    R = Rules(cond.astype(np.int8), action.astype(np.int8),
              np.zeros((n_types, n)), np.ones((n_types, n), np.int64),
              np.zeros((n_types, n), np.int64), np.zeros((n_types, n)),
              np.zeros((n_types, n), np.int64))
    for i in range(n_types):
        for r in range(n):
            refresh(R, i, r, n_goods)
    return R


def complete_rules(n_types, n_goods, exchange, specific_only=False):
    "Every Table 1 condition (or pair of conditions) with both actions."
    conds = table1(n_goods)[:n_goods] if specific_only else table1(n_goods)
    if exchange:
        conds = np.array([np.concatenate([a, b]) for a in conds for b in conds])
    cond = np.tile(np.repeat(conds, 2, axis=0), (n_types, 1, 1))
    action = np.tile([0, 1], (n_types, len(conds)))
    return make_rules(cond, action, n_goods)


def random_rules(n_types, n_rules, length, n_goods, rng):
    "Rules whose positions are drawn independently and uniformly from {0, 1, #}."
    return make_rules(rng.integers(-1, 2, size=(n_types, n_rules, length)),
                      rng.integers(0, 2, size=(n_types, n_rules)), n_goods)
```

The auction (95.5) follows.

For the exercises it also offers the softmax auction with intensity η.

Several operators below must delete a "weak" rule, and the paper measures weakness "in terms of a combination of strength and cumulated number of times the classifier has won the bidding in the past".

Our combination adds the two, each rescaled to $[0, 1]$ across the candidates.

```
@njit
def auction(R, i, s, eta, rng):
    """
    The rule of type i that decides in state s, or -1 if no rule matches.

    eta = inf is the paper's auction: the strongest matched rule wins, with
    ties broken at random.  A finite eta gives a softmax auction instead.
    """
    S, match = R.strength[i], R.match[i]
    best, top, n_ties = -1, -np.inf, 0
    for r in range(S.shape[0]):
```

(continues on next page)

(continued from previous page)

```
        if (match[r] >> s) & 1:
            if S[r] > top:
                best, top, n_ties = r, S[r], 1
            elif S[r] == top:
                n_ties += 1
                if rng.random() * n_ties < 1:
                    best = r
    if best < 0 or eta == np.inf:
        return best
    total = 0.0
    for r in range(S.shape[0]):
        if (match[r] >> s) & 1:
            total += np.exp(eta * (S[r] - top))
    x = rng.random() * total
    for r in range(S.shape[0]):
        if (match[r] >> s) & 1:
            x -= np.exp(eta * (S[r] - top))
            if x <= 0:
                return r
    return best


@njit
def weakest(R, i, s, exclude):
    """
    The weakest rule of type i other than `exclude`, scoring each candidate by
    its strength plus its number of wins, each rescaled to [0, 1].  Candidates
    are the rules matching state s, or all rules if none does or if s < 0.
    """
    S, used, match = R.strength[i], R.used[i], R.match[i]
    n = S.shape[0]
    cand = np.zeros(n, np.bool_)
    for r in range(n):
        cand[r] = r != exclude and (s < 0 or ((match[r] >> s) & 1) == 1)
    if not cand.any():
        for r in range(n):
            cand[r] = r != exclude
    lo, hi = S[cand].min(), S[cand].max()
    u_lo, u_hi = used[cand].min(), used[cand].max()
    best, best_score = -1, np.inf
    for r in range(n):
        if cand[r]:
            score = ((S[r] - lo) / (hi - lo + 1e-12)
                     + (used[r] - u_lo) / (u_hi - u_lo + 1e-12))
            if score < best_score:
                best, best_score = r, score
    return best
```

Three of the four operators that the paper adds for random rules (its Section 6) act inside an auction, so we implement them here.

Creation fires when no rule matches the state: the weakest rule is replaced by one whose condition is exactly the state, with a random action.

Diversification fires when every matched rule calls for the same action: the weakest *other* matched rule is replaced by a rule for exactly this state with the opposite action and the winner's strength, which then ties with the winner.

Specialization is switched on in period t with probability $f_s(t) = 1/(2\sqrt{t})$.

When it is on, each # in the winner's condition is replaced, with probability $p_s = 0.01$, by the bit that matches the current state, and the sharpened copy replaces a weak matched rule.

Under a complete enumeration every state is matched and every condition appears with both actions, so creation and diversification never fire, as the paper notes; specialization is used only with random rules.

```
@njit
def decide(R, i, s, code, eta, specialize, ps, n_goods, rng, stamp):
    """
    Run the auction of type i in state s, whose encoding is `code`, applying
    the paper's creation, diversification and specialization operators.
    """
    w = auction(R, i, s, eta, rng)
    if w < 0:                                          # creation
        j = weakest(R, i, -1, -1)
        set_rule(R, i, j, code, rng.integers(0, 2), 0.0, n_goods, stamp)
        return j

    one_action = True                                  # diversification
    for r in range(R.action.shape[1]):
        if (R.match[i, r] >> s) & 1 and R.action[i, r] != R.action[i, w]:
            one_action = False
            break
    if one_action:
        j = weakest(R, i, s, w)
        set_rule(R, i, j, code, 1 - R.action[i, w], R.strength[i, w],
                 n_goods, stamp)
        w = auction(R, i, s, eta, rng)

    if specialize:                                     # specialization
        new = R.cond[i, w].copy()
        changed = False
        for k in range(new.shape[0]):
            if new[k] == WILD and rng.random() < ps:
                new[k] = code[k]
                changed = True
        if changed:
            j = weakest(R, i, s, w)
            set_rule(R, i, j, new, R.action[i, w], R.strength[i, w],
                     n_goods, stamp)
    return w
```

An `Economy` collects the parameters of the physical environment together with the settings of the learning algorithm.

Goods are numbered from 0 in the code and from 1 in the text, and fiat money, when present, is the last good.

The last three fields are used only in the exercises.

```
@dataclass
class Economy:
    name: str
    produces: np.ndarray               # produces[i] = good produced by type i
    storage_costs: np.ndarray          # one per good; fiat money, if any, is last
    u: float = 100.0                   # utility from consuming one's own good
    n_agents_per_type: int = 50
    random_rules: bool = False         # False: complete enumeration, no GA
    n_trade_rules: int = 72            # population sizes when rules are random
    n_consume_rules: int = 12
    b_trade: tuple = (0.025, 0.025)    # (b11, b12)
```

(continues on next page)

(continued from previous page)

```
    b_consume: tuple = (0.25, 0.25)     # (b21, b22)
    n_fiat: int = 0                     # units of fiat money handed out at t = 0
    ps: float = 0.01                    # specialization: chance a # is filled in
    eta: float = np.inf                 # inf: argmax auction; finite: softmax
    no_eat: bool = False                # True: only one's own good can be eaten
    specific_only: bool = False         # True: enumerate fully specific rules only

    @property
    def n_types(self):
        return len(self.produces)

    @property
    def n_goods(self):
        return len(self.storage_costs)

    @property
    def fiat(self):
        "Index of fiat money, or -1 if there is none."
        return self.n_goods - 1 if self.n_goods > self.n_types else -1
```

The bucket brigade updates two strengths for each agent each period: the exchange rule that won today, as in (95.7), and the consumption rule that won for the same agent yesterday, as in (95.6), whose second receipt is only now known.

The `born` stamps guard against paying a rule that has been overwritten in the meantime.

```
@njit
def bucket_brigade(E, C, i, a, e, active, c, payoff, b, pending):
    """
    Update strengths for agent a of type i: first the exchange rule e that
    won today, then the consumption rule that won for this agent yesterday.
    """
    b11, b12, b21, b22 = b
    last_c, last_payoff, last_born = pending
    S_e = E.strength[i, e]                          # before its own update
    b1 = b11 + b12 * E.sigma[i, e]
    if active:                                      # equation (13)
        tau = E.used[i, e]
        E.used[i, e] = tau + 1
        b2 = b21 + b22 * C.sigma[i, c]
        E.strength[i, e] = S_e - ((1 + b1) * S_e - b2 * C.strength[i, c]) / tau
    p = last_c[a]
    if p >= 0 and C.born[i, p] == last_born[a]:    # equation (12)
        inflow = b1 * S_e if active else 0.0
        b2p = b21 + b22 * C.sigma[i, p]
        tau = C.used[i, p]
        C.used[i, p] = tau + 1
        C.strength[i, p] -= ((1 + b2p) * C.strength[i, p]
                             - inflow - last_payoff[a]) / tau
    last_c[a], last_payoff[a], last_born[a] = c, payoff, C.born[i, c]
```

Each period all A agents are shuffled into $A/2$ pairs, and each pair makes its exchange decisions, trades if both propose, makes its consumption decisions, and settles accounts.

```
@njit
def run_period(E, C, hold, types, produces, s, u, fiat, b, pending, eta,
               specialize, ps, no_eat, rng, stamp, exch, cons):
    "One period: random matching, two auctions per agent, and the accounting."
```

(continues on next page)

(continued from previous page)

```
    n_goods = s.shape[0]
    order = rng.permutation(hold.shape[0])
    code = np.zeros(2 * n_goods, np.int8)
    e = np.zeros(2, np.int64)
    e_born = np.zeros(2, np.int64)
    offer = np.zeros(2, np.int64)
    for k in range(hold.shape[0] // 2):
        pair = (order[2 * k], order[2 * k + 1])
        x = (hold[pair[0]], hold[pair[1]])

        for m in range(2):                          # exchange decisions
            i = types[pair[m]]
            code[:] = 0
            code[x[m]] = 1
            code[n_goods + x[1 - m]] = 1
            e[m] = decide(E, i, x[m] * n_goods + x[1 - m], code, eta,
                          specialize, ps, n_goods, rng, stamp)
            e_born[m], offer[m] = E.born[i, e[m]], E.action[i, e[m]]
        trade = offer[0] == 1 and offer[1] == 1
        if trade:
            hold[pair[0]], hold[pair[1]] = x[1], x[0]
            exch[types[pair[0]], x[0], x[1]] += 1
            exch[types[pair[1]], x[1], x[0]] += 1

        for m in range(2):                          # consumption decisions
            a = pair[m]
            i, g = types[a], hold[a]
            code[:] = 0
            code[g] = 1
            c = decide(C, i, g, code[:n_goods], eta, specialize, ps,
                       n_goods, rng, stamp)
            cons[i, g, 0] += 1
            if C.action[i, c] == 1 and g != fiat and (not no_eat or g == i):
                cons[i, g, 1] += 1
                hold[a] = produces[i]
                payoff = (u if g == i else 0.0) - s[produces[i]]
            else:
                payoff = -s[g]
            active = (trade or offer[m] == 0) and E.born[i, e[m]] == e_born[m]
            bucket_brigade(E, C, i, a, e[m], active, c, payoff, b, pending)
```

A `Simulation` holds the state of an economy and runs it period by period, recording holdings, trades and consumption decisions.

Initial holdings are random, as in the paper, and fiat money is handed to randomly chosen agents.

With random rules, the genetic algorithm `generalize` runs after each period; it is defined in *The genetic algorithm*, before the first economy that needs it.

```
class Simulation:

    def __init__(self, econ, seed=0):
        self.econ = econ
        self.rng = rng = np.random.default_rng(seed)
        K, G = econ.n_types, econ.n_goods
        if econ.random_rules:
            self.E = random_rules(K, econ.n_trade_rules, 2 * G, G, rng)
```

(continues on next page)

(continued from previous page)

```
            self.C = random_rules(K, econ.n_consume_rules, G, G, rng)
        else:
            self.E = complete_rules(K, G, True, econ.specific_only)
            self.C = complete_rules(K, G, False, econ.specific_only)
        self.types = np.repeat(np.arange(K), econ.n_agents_per_type)
        A = len(self.types)
        self.hold = rng.integers(0, K, size=A)       # random commodity holdings
        if econ.n_fiat:
            self.hold[rng.choice(A, econ.n_fiat, replace=False)] = econ.fiat
        self.pending = (np.full(A, -1), np.zeros(A), np.zeros(A, np.int64))
        self.stamp = np.zeros(1, np.int64)
        self.t = 0
        self.hold_hist, self.exch_hist, self.cons_hist = [], [], []

    def run(self, T):
        econ, rng = self.econ, self.rng
        K, G = econ.n_types, econ.n_goods
        b = tuple(float(v) for v in econ.b_trade + econ.b_consume)
        for _ in range(T):
            self.t += 1
            f = 1 / (2 * np.sqrt(self.t))   # rate of specialization and the GA
            specialize = econ.random_rules and rng.random() < f
            exch, cons = np.zeros((K, G, G)), np.zeros((K, G, 2))
            run_period(self.E, self.C, self.hold, self.types, econ.produces,
                       econ.storage_costs, econ.u, econ.fiat, b, self.pending,
                       econ.eta, specialize, econ.ps, econ.no_eat, rng,
                       self.stamp, exch, cons)
            if econ.random_rules:
                for i in range(K):
                    if rng.random() < f:
                        generalize(self.E, i, G, rng, self.stamp, n_crowd=8)
                    if rng.random() < f:
                        generalize(self.C, i, G, rng, self.stamp, n_crowd=4)
            self.hold_hist.append([np.bincount(self.hold[self.types == i],
                                               minlength=G) for i in range(K)])
            self.exch_hist.append(exch)
            self.cons_hist.append(cons)
```

The next cell defines helpers that put the output into the format of the paper's tables (which report ten-period moving averages) and figures.

`winners` lists, for each state, the rule that would win the auction there, and `contenders` lists every rule competing in a state.

`across_seeds` reruns an economy for twenty seeds.

95.5 Economy A1.1: does a medium of exchange emerge?

Our first economy is the Wicksell triangle with

$$s_1 = 0.1, \quad s_2 = 1, \quad s_3 = 20, \quad u_i = 100,$$

fifty agents of each type, and a complete enumeration of the 72 exchange and 12 consumption classifiers, all with strength zero.

Every economy we display is run with seed 0, and its typicality is then checked across 20 seeds (10 for Economy D).

```
economy_a11 = Economy(
    name='A1.1',
    produces=np.array([1, 2, 0]),                # type 1 -> good 2, etc. (0-indexed)
    storage_costs=s_A,
    u=100.0,
)

sim_a11 = Simulation(economy_a11, seed=0)
sim_a11.run(1000)

fundamental = np.array([[0.0, 1.0, 0.0],
                        [0.5, 0.0, 0.5],
                        [1.0, 0.0, 0.0]])
paper_a11 = np.array([[0.0, 1.0, 0.0],           # the paper's Table 6b, t = 1000
                      [0.506, 0.0, 0.494],
                      [1.0, 0.0, 0.0]])

print("holdings at t = 500\n", holdings(sim_a11, t=500), "\n")
print("holdings at t = 1000\n", holdings(sim_a11), "\n")
print(f"max |simulated - fundamental equilibrium| = "
      f"{distance(holdings(sim_a11).to_numpy(), fundamental):.3f}")
print(f"max |simulated - paper's table|           = "
      f"{distance(holdings(sim_a11).to_numpy(), paper_a11):.3f}")
```

```
holdings at t = 500
         good 1  good 2  good 3
type 1   0.000     1.0   0.000
type 2   0.428     0.0   0.572
type 3   1.000     0.0   0.000

holdings at t = 1000
         good 1  good 2  good 3
type 1   0.000     1.0   0.000
type 2   0.494     0.0   0.506
type 3   1.000     0.0   0.000

max |simulated - fundamental equilibrium| = 0.006
max |simulated - paper's table|           = 0.012
```

By $t = 500$ type 1 holds good 2, type 3 holds good 1, and type 2 splits its time between goods 1 and 3, as in the fundamental equilibrium tabulated above.

The time paths show how quickly types 1 and 3 settle down, and a pause in type 2's learning.

```
plot_holdings(sim_a11)
```

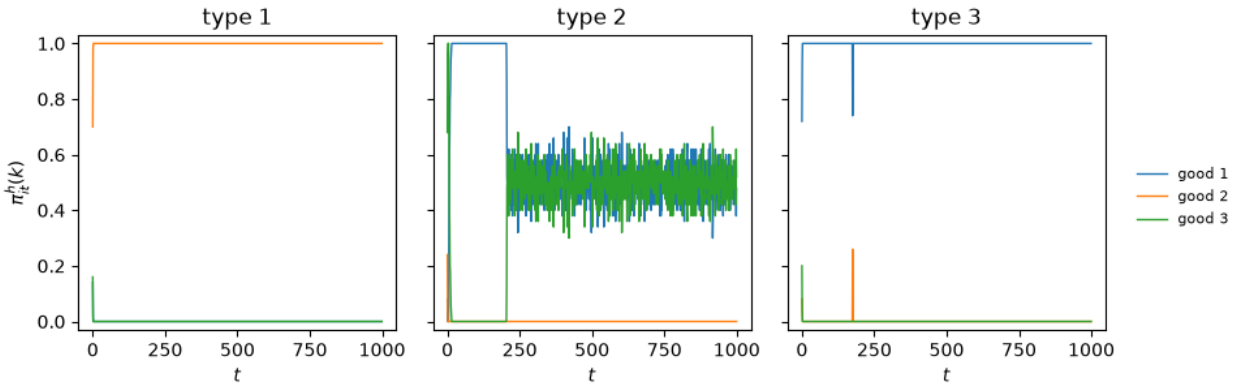


Fig. 95.1: Holdings in Economy A1.1

For roughly the first two hundred periods type 2 agents hold good 1 and nothing else: they have acquired money but do not yet spend it.

The rule that buys good 2 with good 1 has not yet been given a chance, for a reason we explain in *Economy A2.1: when theory predicts speculation*, and once it wins its first auctions type 2 enters the cycle "acquire money, spend money".

From then on type 2's holdings fluctuate, because which good a type 2 agent holds depends on where it is in that cycle.

The next figure shows the trades themselves, averaged over the last hundred periods.

```
plot_flows(sim_a11)
```

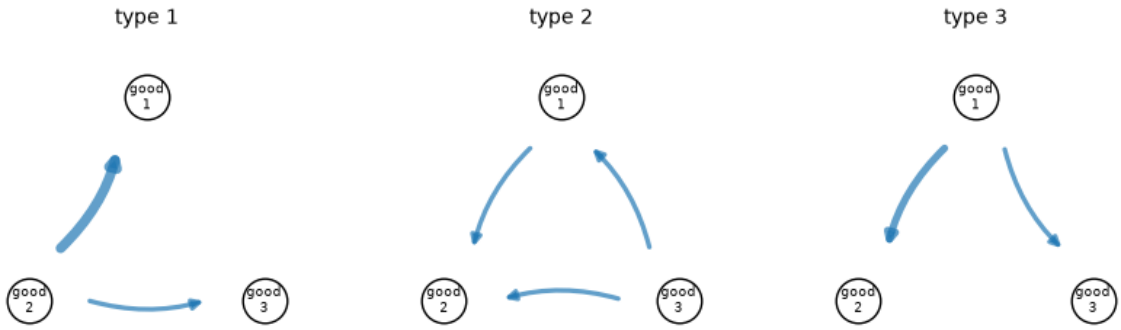


Fig. 95.2: Exchange pattern discovered in Economy A1.1

The fundamental equilibrium triangle is there: type 1 gives up good 2 for good 1, type 3 gives up good 1 for good 3, and type 2 runs both legs, giving up good 3 for good 1 and later good 1 for good 2.

Good 1 circulates as money.

The picture also has arrows that the fundamental equilibrium does not, and they are not thin: type 3 gives up good 1 for good 2 about twice as often as it gives up good 1 for good 3, and type 1 gives up good 2 for good 3 about as often as type 2 runs either of its equilibrium legs.

Each is a payoff-neutral trade of the kind described above, in which the agent eats what it receives and produces its own

good again; the partner gains, because type 1 now gets good 1 from type 3 as well as from type 2, and type 2 gets good 2 for its good 3 directly.

The paper's run had none of these trades, but nothing in the accounting rules them out.

We can also look inside type 1's classifier systems and read off the rule that wins in each state, first for exchange and then for consumption.

```
pd.concat([winners(sim_a11, 0), winners(sim_a11, 0, exchange=False)])
```

```
                winning rule  strength  times used
state
good 1, good 1  ##0 100 -> 1     32.09       22404
good 1, good 2  #0# 0## -> 1      7.93           5
good 1, good 3  #0# 0## -> 1      7.93           5
good 2, good 1  ##0 100 -> 1     32.09       22404
good 2, good 2  010 0## -> 1     -0.14        9320
good 2, good 3  010 001 -> 1     -0.10        6370
good 3, good 1  001 #0# -> 1     40.69           2
good 3, good 2  #0# 0## -> 1      7.93           5
good 3, good 3  001 #0# -> 1     40.69           2
good 1              100 -> 1     66.35       22389
good 2              010 -> 1     -0.33       18730
good 3              #0# -> 1     -0.34        6556
```

The rule that does the work, `##0 100 -> 1`, says "if I hold anything but good 3 and my partner holds good 1, trade": a general rule.

In the state (good 2, good 3) the winning rule *accepts* good 3, but a type 1 agent who acquires good 3 eats it at once and produces good 2 again, which leaves him exactly where refusing would have.

That is the payoff-neutral trade described above, and it is why we read the outcome from holdings, which never show type 1 holding good 3, rather than from the winning actions.

The paper's run (its Table 6) had type 1 refusing good 3 instead, which leads to the same holdings.

Is seed 0 typical?

```
sims_a11 = across_seeds(economy_a11, 1000)
H_a11 = final_holdings(sims_a11)
F_a11 = np.array([trade_rates(sim) for sim in sims_a11])
d, n = distance(H_a11, fundamental), len(sims_a11)
print(f"runs within 0.1 of the fundamental equilibrium: {np.sum(d < 0.1)} of {n}")
print(f"largest distance in any run: {d.max():.3f}")
print(f"runs in which type 3 gives good 1 for good 2:   {np.sum(F_a11[:, 2, 0, 1] > 0.
↪02)} of {n}")
print(f"runs in which type 1 gives good 2 for good 3:   {np.sum(F_a11[:, 0, 1, 2] > 0.
↪02)} of {n}")
```

```
runs within 0.1 of the fundamental equilibrium: 20 of 20
largest distance in any run: 0.060
runs in which type 3 gives good 1 for good 2:   17 of 20
runs in which type 1 gives good 2 for good 3:   15 of 20
```

In every run the holdings settle on the unique stationary equilibrium, and in most runs the payoff-neutral trades appear alongside it.

The medium of exchange is learnable.

95.6 Economy A2.1: when theory predicts speculation

Economy A2 differs from A1 only in that $u_i = 500$ instead of 100.

That is enough to violate the Kiyotaki-Wright inequality (95.4), so for patient agents the unique stationary equilibrium is now the **speculative** one, in which type 1 agents accept good 3 in the expectation of trading it for good 1.

Do our adaptive agents find it?

```
economy_a21 = replace(economy_a11, name='A2.1', u=500.0)

sim_a21 = Simulation(economy_a21, seed=0)
sim_a21.run(1000)

speculative = pd.DataFrame([[0, 0.707, 0.293],        # the paper's Table 8b
                            [0.586, 0, 0.414],
                            [1, 0, 0]],
                           index=type_names(economy_a21),
                           columns=good_names(economy_a21))
print("simulated holdings at t = 1000\n", holdings(sim_a21), "\n")
print("speculative equilibrium\n", speculative)
```

```
simulated holdings at t = 1000
         good 1  good 2  good 3
type 1   0.000     1.0   0.000
type 2   0.524     0.0   0.476
type 3   1.000     0.0   0.000

speculative equilibrium
         good 1  good 2  good 3
type 1   0.000   0.707   0.293
type 2   0.586   0.000   0.414
type 3   1.000   0.000   0.000
```

The answer is no.

The simulated holdings are those of the *fundamental* pattern: type 1 agents hold good 2 and never good 3, whereas speculation requires them to hold good 3 nearly thirty percent of the time.

Here is what happens across seeds.

```
sims_a21 = across_seeds(economy_a21, 1000)
H_a21 = final_holdings(sims_a21)
n = len(sims_a21)
print(f"type 1 holds good 3 more than 5% of the time:  {np.sum(H_a21[:, 0, 2] > 0.05)}
 ↪ of {n}")
print(f"within 0.1 of the fundamental pattern:         "
      f"{np.sum(distance(H_a21, fundamental) < 0.1)} of {n}")
print(f"type 2 holds good 1 less than 10% of the time: {np.sum(H_a21[:, 1, 0] < 0.1)}␣
 ↪of {n}")
print(f"type 2 holds good 1 more than 90% of the time: {np.sum(H_a21[:, 1, 0] > 0.9)}␣
 ↪of {n}")
```

```
type 1 holds good 3 more than 5% of the time:  0 of 20
within 0.1 of the fundamental pattern:         10 of 20
type 2 holds good 1 less than 10% of the time: 9 of 20
type 2 holds good 1 more than 90% of the time: 1 of 20
```

Type 1 never speculates.

Half the runs are near the fundamental pattern at $t = 1000$.

In nearly all the others type 2 almost never holds good 1: types 2 and 3 have stopped trading with each other, so type 2 never acquires the money it needs to buy good 2.

In one run type 2 holds good 1 all the time: it has money but never spends it, as in the first two hundred periods of our A1.1 run.

Here is the first run of the more common kind, seen through the exchange rules of type 3 agents who hold good 1 and meet a type 2 agent holding good 3.

```
locked = int(np.flatnonzero(H_a21[:, 1, 0] < 0.1)[0])
sim_locked = sims_a21[locked]
print(f"seed {locked}\n", holdings(sim_locked), "\n")
c = contenders(sim_locked, 2, 0, 2)          # type 3 holds good 1, meets good 3
pd.concat([c.head(2), c[c['rule'] == '100 001 -> 1']])
```

```
seed 2
         good 1  good 2  good 3
type 1      0.0     1.0     0.0
type 2      0.0     0.0     1.0
type 3      1.0     0.0     0.0
```

```
           rule  strength  times used
0  ##0 001 -> 0       0.1       16714
1  100 001 -> 0       0.0           1
2  100 001 -> 1       0.0           0
```

The type 3 agent wants good 3, which is worth 500 to it, yet the winning rule `##0 001 -> 0` refuses to trade.

Its strength is only 0.1, but it is positive, while the specific rule that would trade, `100 001 -> 1`, has never been used and still has its initial strength of zero.

The refusal rule earned a large payment in the first periods, when a general consumption rule that also covered good 3 followed it.

With $1/\tau$ averaging, the average remembers that windfall for thousands of uses.

This is optimistic initialization's failure mode, and the greedy auction never gives the trade rule a chance until the refusal rule's average sinks below zero.

That does happen eventually, so runs locked in this way are slow rather than stuck for good.

We continue the runs that were not near the fundamental pattern for another 5,000 periods.

```
slow = [sim for sim, h in zip(sims_a21, H_a21) if distance(h, fundamental) >= 0.1]
for sim in slow:
    sim.run(5000)
H_slow = final_holdings(slow)
print(f"runs not near the fundamental pattern at t = 1000: {len(slow)}")
print(f"of these, within 0.1 of it at t = 6000:            "
      f"{np.sum(distance(H_slow, fundamental) < 0.1)}")
print(f"of these, type 1 holds good 3 more than 5%:        {np.sum(H_slow[:, 0, 2] >␣
↪0.05)}")
```

```
runs not near the fundamental pattern at t = 1000: 10
of these, within 0.1 of it at t = 6000:            10
of these, type 1 holds good 3 more than 5%:        0
```

All of these runs reach the fundamental pattern by $t = 6000$, and type 1 still never speculates.

Lock-in is more frequent here than in Economy A1.1, plausibly because u_1 is five times larger, and so are the early windfalls that the averages must forget.

Not every lock-in wears off, though.

Other seeds turn up a second failure mode, and seed 24 shows it.

```
sim_24 = Simulation(economy_a21, seed=24)
sim_24.run(6000)
print(holdings(sim_24), "\n")
winners(sim_24, 1, exchange=False)
```

```
        good 1  good 2  good 3
type 1     0.0     1.0     0.0
type 2     0.0     0.0     1.0
type 3     1.0     0.0     0.0
```

```
       winning rule  strength  times used
state
good 1     ##0 -> 1    146.86      123378
good 2     ##0 -> 1    146.86      123378
good 3     001 -> 0    -12.61      175852
```

This time type 3 does accept good 3; the trouble is with type 2.

It trades good 3 for good 1, as in the fundamental pattern, but then eats the good 1 instead of keeping it to buy good 2 from type 1, so it ends every period holding good 3 again.

The rule responsible is the general consumption rule `##0 -> 1`, which says "eat" whether type 2 holds good 1 or good 2.

Eating good 2 is right and earns type 2 its utility of consumption, which keeps the rule's strength high; eating good 1 is wrong, but the same rule decides that state too.

The specific rule that would keep good 1, `100 -> 0`, was tried only a few times and sits far below the general rule's strength, an average that includes every meal of good 2, so under the argmax auction it never wins again.

So this run is stuck for good: seed 24, and seed 26 in the same way, are still in this pattern at $t = 24{,}000$.

Optimistic initialization lets a rule that was once lucky shut out an untried rival; this second lock-in is the work of an overgeneral consumption rule, one that earns well in one state and shuts out the specific rule that would act correctly in another.

Type 1 still never speculates: across seeds 20 to 59 it holds good 3 no more than 5 percent of the time in every run.

So the agents do not find A2's equilibrium.

For patient agents the speculative equilibrium is the *unique* stationary equilibrium of this economy, so this is not learning selecting among equilibria.

It is convergence to an outcome that is not an equilibrium of the economy the agents live in.

The paper reached the same verdict from a single run, and its diagnosis is worth quoting in full:

> *Patience requires experience.* The transfer system inside the classifier system is designed to converge to a set of long-run average strengths. In the limit, the artificially intelligent agents should behave as long-run average payoff maximizers, since the steady-state strengths weight payoffs by their relative frequencies. It takes time, however, for optimal rules to achieve the desired strengths. The behavior of our artificially intelligent agents can be very myopic at the beginning. In economies, such as Economy A2, in which the nature of the equilibrium changes with the discount rate, this early myopia might have a perverse effect in diverting

> the economy towards a low discount-factor stationary equilibrium, such as the fundamental equilibrium. The underlying algorithm has to provide enough experimentation to avoid an early 'lock in'. The present algorithm seems defective in that it has too little experimentation to support the speculative equilibrium even in the long simulations we have run.

The paper's footnote 14 adds that a fundamental equilibrium exists in A2 only if the discount factor is sufficiently low.

Its footnote 15 reports a telling contrast: in experiments by Marimon and Miller with a genetic algorithm that breeds whole trading strategies (the approach the paper's footnotes 1 and 3 describe), 25 of 30 runs of Economy A2 converged to the speculative equilibrium.

Their agents had no bucket brigade to credit: each strategy was judged as a whole.

The paper also points to a second culprit.

In its run, type 1's winning exchange rule *accepted* good 3 (its Table 9d), but type 1's winning consumption rules were "too general (have too many # signs) to distinguish among goods stored", so type 1 ate good 3 rather than carrying it to a type 3 agent.

Accepting good 3 and then eating it is payoff-equivalent to refusing it, which is why the holdings looked fundamental.

Our seed 0 shows the same thing.

```
pd.concat([winners(sim_a21, 0).loc[['good 2, good 1', 'good 2, good 3']],
           winners(sim_a21, 0, exchange=False).loc[['good 1', 'good 3']]])
```

state	winning rule	strength	times used
good 2, good 1	0## ##0 -> 1	96.82	29260
good 2, good 3	010 0## -> 1	0.30	10707
good 1	100 -> 1	334.29	17421
good 3	#0# -> 1	1.08	5983

Type 1's winning exchange rule in the state (good 2, good 3), `010 0## -> 1`, accepts good 3, and the consumption rule that wins when type 1 holds good 3 is the general rule `#0# -> 1`, "if I do not hold good 2, eat".

So far this is the paper's Table 9d.

But generality is not what makes type 1 eat good 3 here.

Here are all of type 1's consumption rules that match good 3.

```
c = sim_a21.C
rows = [r for r in range(c.cond.shape[1]) if (c.match[0, r] >> 2) & 1]
pd.DataFrame({'rule': [rule_string(c.cond[0, r], c.action[0, r], 3) for r in rows],
              'strength': c.strength[0, rows].round(2),
              'times used': c.used[0, rows] - 1})
```

	rule	strength	times used
0	001 -> 0	-11.59	1
1	001 -> 1	-0.48	13
2	0## -> 0	-9.32	6
3	0## -> 1	-0.68	141
4	#0# -> 0	-13.09	5
5	#0# -> 1	1.08	5983

The specific rule that eats good 3 has a strength near zero, while the specific rule that carries it, tried only once, has a strength near -12.

Carrying good 3 costs 20 this period, against 1 for eating it, and the reward for carrying it, a chance to trade it for good 1 next period, reaches the carrying rule only through the bucket brigade, heavily discounted.

Type 1 therefore never carries good 3 into the next period, and the holdings are fundamental.

Three mechanisms are therefore at work in A2:

1. **Impatience.** Even at rest, the bucket brigade values the future at a tiny discount factor, which the next section derives.
2. **Greedy lock-in.** The argmax auction stops experimenting once untried rules are used up, as in the trapped runs above.
3. **Overconsumption by general consumption rules**, which makes accepting good 3 pointless, as in the paper's Table 9, and in seed 24 keeps type 2 from holding money.

Exercise 95.14.3, Exercise 95.14.4 and Exercise 95.14.5 take them apart one at a time.

95.7 The bucket brigade's discount factor

The paper hoped that strengths built from long-run averages would make the agents behave like long-run average payoff maximizers once the strengths settled down.

The stationarity conditions say otherwise.

Take an agent whose exchange rule e' at $t+1$ follows a consumption rule c at t, and whose consumption rule c' at $t+1$ follows e'.

From the second stationarity condition, $(1+b_1(e'))S_{e'} = \mathbb{E}[b_2(c')S_{c'}]$.

Substituting it into the first gives

$$(1+b_2(c))S_c = \mathbb{E}[U(\gamma_c)] + \mathbb{E}\left[\beta_{\text{eff}}\,(1+b_2(c'))S_{c'}\right], \qquad \beta_{\text{eff}} = \frac{b_1(e')\,b_2(c')}{(1+b_1(e'))(1+b_2(c'))}. \tag{95.8}$$

So the scaled strengths $V_c = (1+b_2(c))S_c$ satisfy a Bellman equation for the agent's current policy, $V_c = \mathbb{E}[U] + \mathbb{E}[\beta_{\text{eff}}V_{c'}]$, with discount factor β_{eff}.

The exchange strengths inherit the ranking: among exchange rules of equal specificity, S_e is proportional to the expected value of the consumption rule that follows e.

When an offer is refused, e' pays nothing backward, which lowers the effective discount factor further.

```
def beta_eff(b1, b2):
    "Effective discount factor of the bucket brigade for bid fractions b1, b2."
    return b1 * b2 / ((1 + b1) * (1 + b2))


b11, b12, b21, b22 = 0.025, 0.025, 0.25, 0.25            # Economy A's bids
β_A = beta_eff(b11 + b12, b21 + b22)                     # fully specific rules
print(f"β_eff with A's bids, specific rules:   {β_A:.4f}")
print(f"β_eff with A's bids, general rules:    {beta_eff(b11 + b12 / 5, b21 + b22 /␣
↪3):.4f}")
print(f"β_eff with B.1's bids, specific rules: {beta_eff(0.5, 0.5):.4f}")
print(f"β_eff as both bid fractions approach 1: {beta_eff(1, 1):.4f}")
```

```
β_eff with A's bids, specific rules:   0.0159
β_eff with A's bids, general rules:    0.0073
β_eff with B.1's bids, specific rules: 0.1111
β_eff as both bid fractions approach 1: 0.2500
```

With the paper's bids the agents behave, at rest, as if they discounted the future at a factor of about 0.016 per period, and the most general rules, which bid less, at about 0.007.

Because $b/(1+b) < 1/2$ for any bid fraction below one, β_{eff} can never reach $1/4$.

The source is the averaging itself: a rule's own bid enters the average it keeps, so each link of the chain passes back only the fraction $b/(1+b)$ of the value it receives.

What does an agent with discount factor β do in Economy A?

Redoing the argument behind (95.4) for an agent who discounts each period's payoffs by β, a type 1 agent facing given holdings of the other types prefers to hold good 3 if and only if

$$s_3 - s_2 < \beta(\pi_3^h(1) - \pi_2^h(1))\tfrac{1}{3}u_1,$$

so lowering β scales up the threshold for u_1 by $1/\beta$.

```
for label, β in [("patient agents, β = 1", 1.0), ("bucket brigade, β = β_eff", β_A)]:
    print(f"{label}: fundamental for u_1 < "
          f"{speculation_threshold(s_A, 1.0, 0.5, β):,.0f}, speculative for u_1 > "
          f"{speculation_threshold(s_A, 1.0, 0.586, β):,.0f}")
print(f"A2's u_1 = 500 supports speculation only if β > "
      f"{speculation_threshold(s_A, 1.0, 0.586) / 500:.3f}")
```

```
patient agents, β = 1: fundamental for u_1 < 114, speculative for u_1 > 138
bucket brigade, β = β_eff: fundamental for u_1 < 7,182, speculative for u_1 > 8,674
A2's u_1 = 500 supports speculation only if β > 0.275
```

For an agent with $\beta = \beta_{\text{eff}}$, A2's $u_1 = 500$ is far inside the region where the fundamental pattern is the unique equilibrium.

The figure plots the two thresholds against β.

```
βs = np.logspace(-2.5, 0, 200)
fig, ax = plt.subplots(figsize=(7, 4))
ax.loglog(βs, speculation_threshold(s_A, 1.0, 0.5, βs),
          label="fundamental equilibrium exists below")
ax.loglog(βs, speculation_threshold(s_A, 1.0, 0.586, βs), '--',
          label="speculative equilibrium exists above")
for u, name in [(100, "A1"), (500, "A2")]:
    ax.axhline(u, color='k', lw=0.6)
    ax.annotate(f"{name}: $u_1 = {u}$", (βs[0], u * 1.1), fontsize=9)
ax.axvline(β_A, color='C3', lw=1)
ax.axvline(0.25, color='C3', lw=1, ls=':')
ax.annotate(r"$\beta_{\rm eff}$", (β_A * 1.08, 150), color='C3')
ax.annotate("1/4", (0.25 * 1.08, 150), color='C3')
ax.set_xlabel(r"discount factor $\beta$")
ax.set_ylabel(r"$u_1$")
ax.legend(frameon=False, loc='upper right')
plt.show()
```

Economy A2 supports speculation only for β above 0.275, beyond the reach of any bid fractions, so at rest the bucket brigade cannot support A2's equilibrium with fully specific rules, whatever its bids.

This sharpens the paper's diagnosis.

Early myopia is real, but it is not only early: the stationary strengths themselves are impatient.

The fundamental pattern – accept a good only if you eat it or it is cheaper to store – is exactly what an impatient agent does, which is the paper's "low discount-factor stationary equilibrium".

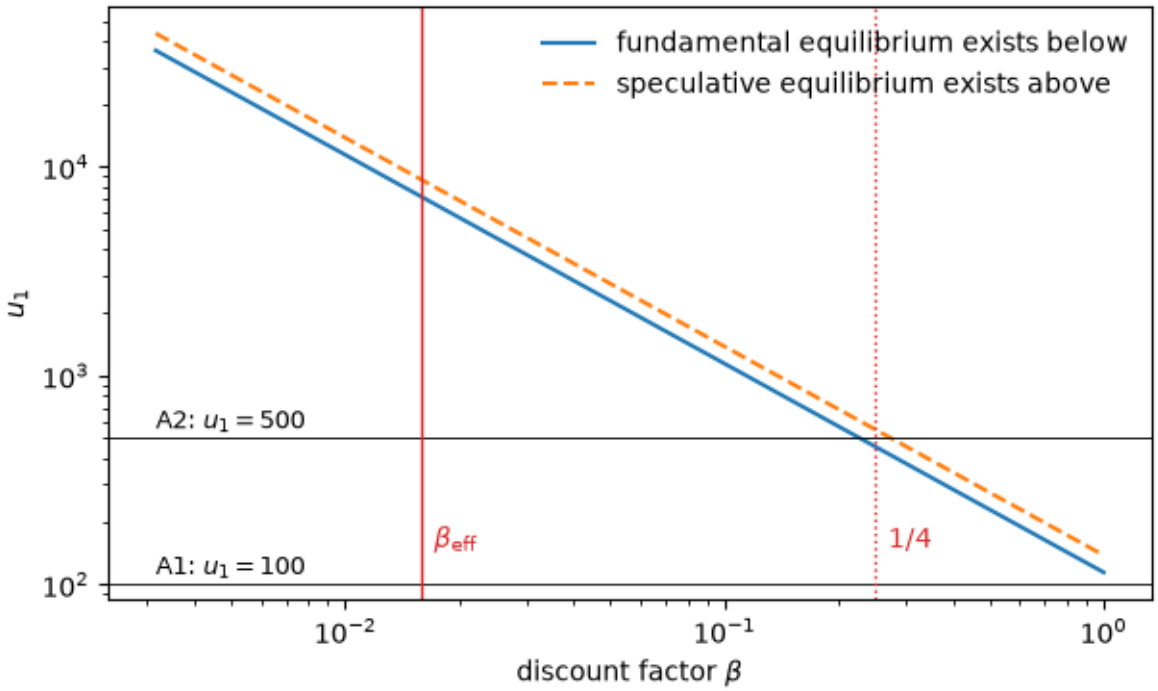


Fig. 95.3: Speculation thresholds for u_1 against the discount factor

The same is true of the patterns the classifier systems find in every other economy of this lecture: the fundamental equilibrium of Economy B, the acceptance of fiat money in Economy C (fiat costs nothing to store), and the "lower storage cost" pattern the paper reports for Economy D.

Equation (95.8) is a familiar object.

The scaled strengths are evaluated like the temporal-difference value estimates of reinforcement learning [Sutton and Barto, 2018], which *Job Search X: Search with Q-Learning* and *Inventory Management via Q-Learning* apply to a job seeker and to a firm.

The difference is that here the discount factor is not chosen to match the agent's preferences: it is a by-product of the bid fractions.

The bucket brigade finds the same factor $b/(1+b)$ per link in a two-link chain, and shows that Holland's cumulative balances, which do not average, pass a payoff back undiminished but need not converge.

One recursion, many learners explains why, with $1/\tau$ averaging, the rest points of (95.6)-(95.7) are the objects to study.

Exercise 95.14.5 checks the threshold by simulation.

95.8 Economy B.1: selecting between two equilibria

Economy B changes the production pattern – type 1 produces good 3, type 2 produces good 1, type 3 produces good 2 – and compresses storage costs to $s = (1, 4, 9)$, with $u_i = 100$.

Now there are two stationary equilibria for patient agents: in the fundamental one goods 1 and 2 serve as media of exchange, and in the speculative one goods 2 and 3 do.

B.1 also uses exchange bids ten times A's, $b_{11} = b_{12} = 0.25$, as in the paper's Table 12, which raises β_{eff} to about 0.11.

```
economy_b1 = Economy(
    name='B.1',
    produces=np.array([2, 0, 1]),               # type 1 -> good 3, type 2 -> good 1,␣
 ↪...
    storage_costs=np.array([1.0, 4.0, 9.0]),
    u=100.0,
    b_trade=(0.25, 0.25),
)

sim_b1 = Simulation(economy_b1, seed=0)
sim_b1.run(1000)

fundamental_b = np.array([[0, 0.293, 0.707],     # the paper's Table 10b
                          [1, 0, 0],
                          [0.586, 0.414, 0]])
speculative_b = np.array([[0, 0.586, 0.414],     # the paper's Table 11b
                          [0.707, 0, 0.293],
                          [0, 1, 0]])
paper_b1 = np.array([[0.0, 0.28, 0.72],          # the paper's Table 12b, t = 1000
                     [0.994, 0.0, 0.006],
                     [0.526, 0.474, 0.0]])

print("holdings at t = 500\n", holdings(sim_b1, t=500), "\n")
print("holdings at t = 1000\n", holdings(sim_b1), "\n")
h = holdings(sim_b1).to_numpy()
for name, target in [("fundamental equilibrium", fundamental_b),
                     ("speculative equilibrium", speculative_b),
                     ("paper's table", paper_b1)]:
    print(f"max |simulated - {name}| = {distance(h, target):.3f}")
print(f"max |paper's table - fundamental equilibrium| = "
      f"{distance(paper_b1, fundamental_b):.3f}")
```

```
holdings at t = 500
         good 1  good 2  good 3
type 1    0.00   0.292   0.708
type 2    1.00   0.000   0.000
type 3    0.57   0.430   0.000

holdings at t = 1000
         good 1  good 2  good 3
type 1   0.000   0.266   0.734
type 2   1.000   0.000   0.000
type 3   0.556   0.444   0.000

max |simulated - fundamental equilibrium| = 0.030
max |simulated - speculative equilibrium| = 0.556
max |simulated - paper's table| = 0.030
max |paper's table - fundamental equilibrium| = 0.060
```

The run ends near the fundamental equilibrium of Economy B, closer to it than the paper's own run.

```
plot_holdings(sim_b1)
```

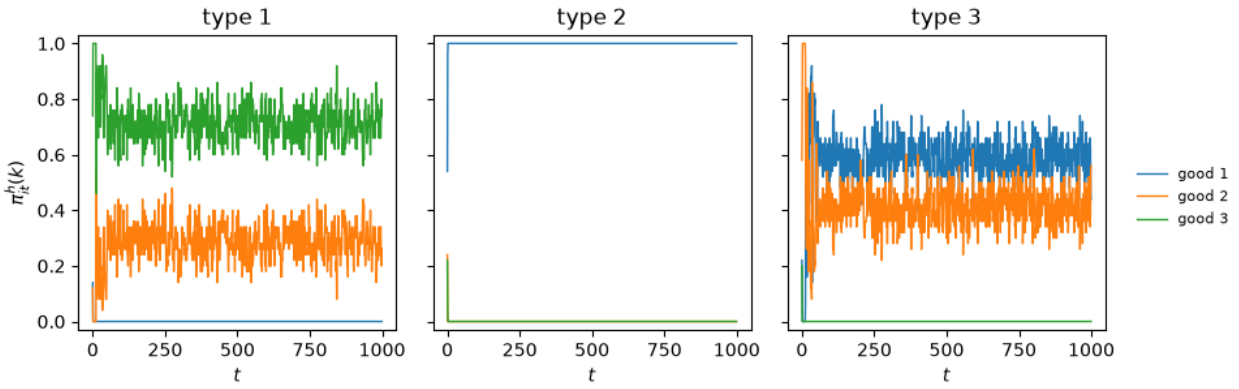


Fig. 95.4: Holdings in Economy B.1

Across seeds:

```
H_b1 = final_holdings(across_seeds(economy_b1, 1000))
d_f, d_s = distance(H_b1, fundamental_b), distance(H_b1, speculative_b)
print(f"within 0.1 of the fundamental equilibrium: {np.sum(d_f < 0.1)} of {len(H_b1)}
 ↪")
print(f"within 0.1 of the speculative equilibrium: {np.sum(d_s < 0.1)} of {len(H_b1)}
 ↪")
print(f"smallest distance from the speculative equilibrium: {d_s.min():.2f}")
print("holdings in the runs farther than 0.1 from the fundamental equilibrium:")
print(H_b1[d_f >= 0.1].round(2))
```

```
within 0.1 of the fundamental equilibrium: 18 of 20
within 0.1 of the speculative equilibrium: 0 of 20
smallest distance from the speculative equilibrium: 0.29
holdings in the runs farther than 0.1 from the fundamental equilibrium:
[[[0.   0.49 0.51]
  [1.   0.   0.  ]
  [0.   1.   0.  ]]

 [[0.   0.27 0.73]
  [1.   0.   0.  ]
  [0.47 0.53 0.  ]]]
```

In 18 of the 20 runs the holdings end within 0.1 of the fundamental equilibrium, and no run comes within 0.29 of the speculative one.

Of the other two runs, one is just outside the 0.1 band, and in the other type 3 agents never trade their good 2 for good 1.

Here learning does select between two equilibria, and it selects the fundamental one.

The paper reports one more feature that our run does not reproduce: at $t = 500$ its B.1 looked speculative, and only by $t = 1000$ had it "practically converged to the fundamental equilibrium".

The paper read that movement as evidence that the selection was not "the result of myopic behavior".

In our run there is no speculative phase beyond the first few periods, which are dominated by the random initial holdings, so the movement is not something we can confirm.

And *The bucket brigade's discount factor* suggests a different reading: the fundamental equilibrium of Economy B is also the one an impatient agent plays, and the bucket brigade is impatient at rest, not just early on.

95.9 The genetic algorithm

Complete enumeration is feasible only because the Kiyotaki-Wright state space is tiny.

In larger problems an agent must work with a limited population of rules that is continually revised.

The paper starts such agents from randomly generated rules and adds four operators: the three we have already met inside the auction, and **generalization**, a genetic algorithm.

Generalization is called with probability $f_g(t) = 1/(2\sqrt{t})$ in period t.

When it is called, it works as follows (the paper's Section 6 and Table 7):

- **Potential exterminants** are rules that are both weak and seldom used: strength below $\bar{S} = 0$ *and* a counter below $p_3 = 0.2$ times the largest counter.
- **Potential parents** are a fraction $p_2 = 0.7$ of the rules, chosen randomly "in accordance with strength and frequency"; we draw them with probabilities proportional to their counters and then select pairs within them by a roulette wheel on strength.
- Pairs mate until a fraction $p_1 = 0.2$ of the population has been replaced.
- **Generalizing crossover** (the paper's Fig. 4): the pair draws two cut points and focuses on the positions inside them, or, with probability one half, outside them; wherever the parents carry different specific bits in the focus, both children get a $\#$. Children therefore match everything their parents matched, and more.
- Each child gets the average of its parents' strengths and replaces the potential exterminant it most resembles, by **crowding**: $N = 8$ times (4 for consumption rules) a sample of half the potential exterminants is drawn, the weakest member of each sample is scored by the number of positions it shares with the child plus one if its action differs, and the highest score loses its place.

The paper's footnote 12 says that more standard Holland algorithms "were first tried without much success, prompting us to produce the modified algorithm described in section 6".

We therefore use the generalizing crossover in every economy with random rules, and, since the paper's Table 7 lists no mutation rate, no mutation.

Where the paper is silent we have made the following choices:

- ties in the auction are broken at random;
- generalization is drawn independently for each type's exchange and consumption systems;
- new rules, including children, start with a counter of 1, so after their first experience the strength they were born with survives only through the bid they paid out of it;
- creation gives the new rule strength zero;
- "weakness" outside the genetic algorithm is the score in `weakest`;
- specialization, when switched on in a period, applies to every auction in that period;
- each agent's own decisions form a bucket-brigade chain, as discussed above;
- the goods of Economy D get a five-position one-hot code.

```
@njit
def roulette(w, rng):
    "Draw an index with probability proportional to the nonnegative weights w."
    x = rng.random() * w.sum()
    for k in range(w.shape[0]):
        x -= w[k]
        if x <= 0:
            return k
    return w.shape[0] - 1


@njit
def generalize(R, i, n_goods, rng, stamp, n_crowd,
               p1=0.2, p2=0.7, p3=0.2, p4=0.5, S_bar=0.0):
    "One call of the paper's genetic algorithm on the classifier system of type i."
    cond, action, S, used = R.cond[i], R.action[i], R.strength[i], R.used[i]
    n, L = cond.shape

    # potential exterminants: weak *and* seldom used
    doomed = [r for r in range(n) if S[r] < S_bar and used[r] < p3 * used.max()]
    if len(doomed) == 0:
        return

    # potential parents: p2 * n rules drawn without replacement, weighted by use
    w = used.astype(np.float64)
    parents = np.zeros(int(p2 * n), np.int64)
    for k in range(parents.shape[0]):
        parents[k] = roulette(w, rng)
        w[parents[k]] = 0.0
    fitness = S[parents] - S[parents].min() + 1e-6

    for _ in range(max(1, int(p1 * n) // 2)):
        mum = parents[roulette(fitness, rng)]
        dad = parents[roulette(fitness, rng)]
        kids = (cond[mum].copy(), cond[dad].copy())

        # generalizing crossover
        cuts = np.sort(rng.integers(0, L + 1, size=2))
        inside = rng.random() < 0.5
        for j in range(L):
            if ((cuts[0] <= j < cuts[1]) == inside and kids[0][j] != WILD
                    and kids[1][j] != WILD and kids[0][j] != kids[1][j]):
                kids[0][j] = WILD
                kids[1][j] = WILD
        strength = 0.5 * (S[mum] + S[dad])

        # crowding: each child replaces the potential exterminant it most resembles
        for kid, act in ((kids[0], action[mum]), (kids[1], action[dad])):
            if len(doomed) == 0:
                return
            size = max(1, int(p4 * len(doomed)))
            victim, best = doomed[0], -1
            for _ in range(n_crowd):
                sample = rng.permutation(len(doomed))[:size]
                cand = doomed[sample[0]]
                for q in sample:
                    if S[doomed[q]] < S[cand]:
```

(continues on next page)

(continued from previous page)

```
                    cand = doomed[q]
                likeness = np.sum(kid == cond[cand]) + (act != action[cand])
                if likeness > best:
                    victim, best = cand, likeness
            set_rule(R, i, victim, kid, act, strength, n_goods, stamp)
            doomed.remove(victim)
```

95.10 Economy A1.2: learning from random rules

Economy A1.2 has the same parameters as A1.1, but agents begin with 72 exchange rules and 12 consumption rules drawn *at random* from the 1,458 and 54 possible strings.

Most of these rules are nonsense, and there is no guarantee that the population contains the rules needed to support an equilibrium.

The operators must manufacture them.

```
economy_a12 = replace(economy_a11, name='A1.2', random_rules=True)

sim_a12 = Simulation(economy_a12, seed=0)
sim_a12.run(2000)

paper_a12 = np.array([[0.0, 0.98, 0.02],        # the paper's Table 7b, t = 2000
                      [0.318, 0.0, 0.682],
                      [1.0, 0.0, 0.0]])
print("holdings at t = 1000\n", holdings(sim_a12, t=1000), "\n")
print("holdings at t = 2000\n", holdings(sim_a12), "\n")
print("paper's table at t = 2000\n", paper_a12, "\n")
print(f"max |paper's table - fundamental equilibrium| = "
      f"{distance(paper_a12, fundamental):.3f}")
```

```
holdings at t = 1000
         good 1  good 2  good 3
type 1   0.000   0.942   0.058
type 2   0.424   0.000   0.576
type 3   0.842   0.158   0.000

holdings at t = 2000
         good 1  good 2  good 3
type 1   0.000   0.946   0.054
type 2   0.422   0.000   0.578
type 3   0.788   0.212   0.000

paper's table at t = 2000
 [[0.    0.98  0.02 ]
 [0.318 0.    0.682]
 [1.    0.    0.   ]]

max |paper's table - fundamental equilibrium| = 0.182
```

Here are the time paths.

```
plot_holdings(sim_a12)
```

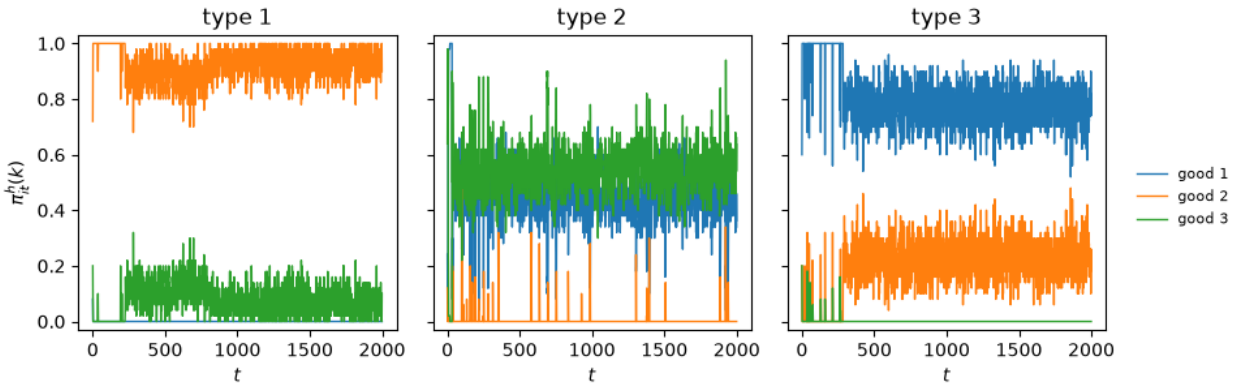


Fig. 95.5: Holdings in Economy A1.2

The fundamental pattern emerges again, though less cleanly than under complete enumeration.

Type 1 mostly holds good 2, type 2 holds only goods 1 and 3, and type 3 mostly holds good 1.

But types 1 and 3 also hold goods that the fundamental equilibrium excludes: in this run type 1 holds good 3 about 5 percent of the time and type 3 holds good 2 about a fifth of the time.

The paper's run was cleaner for types 1 and 3 but not for type 2: at $t = 2000$ its type 2 agents held good 1 only 32 percent of the time, and its Table 7e shows type 2 eating good 1 – its money – with probability 0.45 at $t = 1000$ and 0.23 at $t = 2000$.

The paper calls this overeating "a minor problem in several of our economies", associated with states the classifier system visits too rarely to learn about.

Here are the eating frequencies over the last 200 periods of our run.

```
eating_rates(sim_a12)
```

```
         good 1  good 2  good 3
type 1     1.00    0.50    0.53
type 2     0.02    0.98    0.97
type 3     0.56    0.46    1.00
```

In this run type 2 hardly ever eats its money, which is why its holdings are close to the equilibrium's even split.

Types 1 and 3 eat "wrong" goods about half the time, which costs them nothing when the good is their own production good (good 2 for type 1, good 1 for type 3).

But type 1 also carries good 3 into the next period about half the times it gets it, at a cost of 20 rather than 1, and type 3 carries good 2 about half the time: these are the rarely visited states the paper blames.

Across seeds:

```
sims_a12 = across_seeds(economy_a12, 2000)
H_a12 = final_holdings(sims_a12)
eat_money = np.array([eating_rates(sim).iloc[1, 0] for sim in sims_a12])
n = len(H_a12)
print(f"type 1 holds good 2 at least 90% of the time: {np.sum(H_a12[:, 0, 1] >= 0.9)}␣
↪of {n}")
print(f"type 3 holds good 1 at least 90% of the time: {np.sum(H_a12[:, 2, 0] >= 0.9)}␣
↪of {n}")
print(f"type 2 holds only goods 1 and 3:              {np.sum(H_a12[:, 1, 1] < 0.05)}␣
```

(continues on next page)

(continued from previous page)

```
 ↪of {n}")
print(f"type 2's share of good 1: median {np.median(H_a12[:, 1, 0]):.2f}, "
      f"range {H_a12[:, 1, 0].min():.2f} to {H_a12[:, 1, 0].max():.2f}")
print(f"within 0.1 of the fundamental equilibrium:    "
      f"{np.sum(distance(H_a12, fundamental) < 0.1)} of {n}")
print(f"type 2 eats good 1 at least 40% of the time:  {np.sum(eat_money >= 0.4)} of
 ↪{n}")
```

```
type 1 holds good 2 at least 90% of the time: 16 of 20
type 3 holds good 1 at least 90% of the time: 13 of 20
type 2 holds only goods 1 and 3:              20 of 20
type 2's share of good 1: median 0.19, range 0.09 to 0.50
within 0.1 of the fundamental equilibrium:    5 of 20
type 2 eats good 1 at least 40% of the time:  11 of 20
```

Seed 0 is one of the better runs.

In every run type 2 holds only goods 1 and 3, and in most runs types 1 and 3 hold their equilibrium goods at least 90 percent of the time.

But type 2's median share of good 1 is only 0.19, and in about half the runs type 2 eats good 1 at least 40 percent of the time, the overeating of the paper's run.

Only a quarter of the runs end within 0.1 of the fundamental equilibrium; the paper's own run was 0.18 away.

General rules survive the genetic algorithm.

```
share = lambda R: np.mean(R.cond == WILD)
print(f"share of # in exchange rules:    {share(sim_a12.E):.2f}")
print(f"share of # in consumption rules: {share(sim_a12.C):.2f}")
```

```
share of # in exchange rules:    0.32
share of # in consumption rules: 0.16
```

A random string has a # in a third of its positions, and generalizing crossover keeps adding them while specialization slowly removes them.

95.11 Economy C: fiat money

Now add a fourth object, fiat money, which the paper calls good 0.

It

- costs nothing to store, $s_0 = 0$;
- yields no utility to anybody; and
- cannot be consumed.

It is intrinsically worthless.

It is introduced by handing 48 units of it to 48 randomly chosen agents at $t = 0$, and commodity storage costs are raised to $s = (9, 14, 29)$ so that no commodity is nearly as cheap to store as money.

The paper's footnote 17 reports that a version with Economy A's storage costs did not converge, "probably because the low cost of storing good 1 (0.1) made this good too good a substitute for fiat money".

Agents start from 150 random exchange rules and 20 random consumption rules each, out of 13,122 and 162 possible strings.

```
economy_c = Economy(
    name='C',
    produces=np.array([1, 2, 0]),
    storage_costs=np.array([9.0, 14.0, 29.0, 0.0]),    # fiat money is the last good
    u=100.0,
    random_rules=True,
    n_trade_rules=150,
    n_consume_rules=20,
    b_consume=(0.025, 0.25),
    n_fiat=48,
)

sim_c = Simulation(economy_c, seed=0)
sim_c.run(2000)

fundamental_c = np.array([[0, 0.74, 0, 0.26],       # the paper's Table 14b
                          [0.26, 0, 0.42, 0.32],
                          [0.62, 0, 0, 0.38]])
paper_c = np.array([[0, 0.54, 0, 0.46],             # the paper's Table 15b, t = 1250
                    [0.18, 0, 0.53, 0.28],
                    [0.77, 0.01, 0, 0.21]])
print("holdings at t = 750\n", holdings(sim_c, t=750), "\n")
print("holdings at t = 1250\n", holdings(sim_c, t=1250), "\n")
h = holdings(sim_c, t=1250).to_numpy()
print(f"max |simulated - fundamental equilibrium| = {distance(h, fundamental_c):.3f}")
print(f"max |simulated - paper's table|           = {distance(h, paper_c):.3f}")
```

```
holdings at t = 750
         good 1  good 2  good 3   fiat
type 1   0.000   0.676   0.042  0.282
type 2   0.108   0.000   0.588  0.304
type 3   0.626   0.000   0.000  0.374

holdings at t = 1250
         good 1  good 2  good 3   fiat
type 1   0.000   0.628   0.070  0.302
type 2   0.096   0.000   0.538  0.366
type 3   0.708   0.000   0.000  0.292

max |simulated - fundamental equilibrium| = 0.164
max |simulated - paper's table|           = 0.158
```

Every type holds fiat money roughly 30 percent of the time, close to the theoretical fundamental equilibrium (the paper's Table 14b); the paper's own run (its Table 15b) has type 1 holding fiat money more often, 46 percent of the time.

The largest discrepancy from the theory is type 2's holding of good 1, about 0.10 in our run against 0.26 in theory and 0.18 in the paper's run: our type 2 agents rely on fiat money rather than good 1 as their medium of exchange more than the theory says.

```
plot_holdings(sim_c)
```

The exchange pattern shows money changing hands.

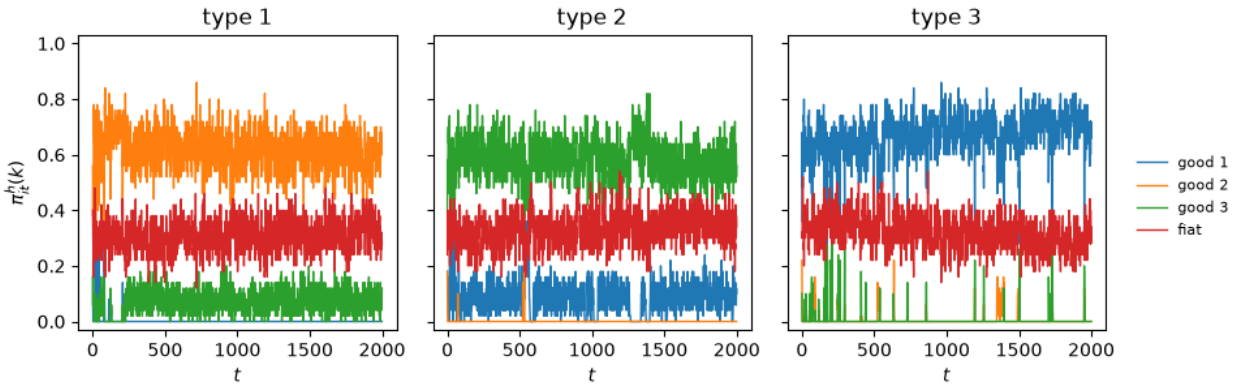


Fig. 95.6: Holdings in Economy C

```
plot_flows(sim_c)
```

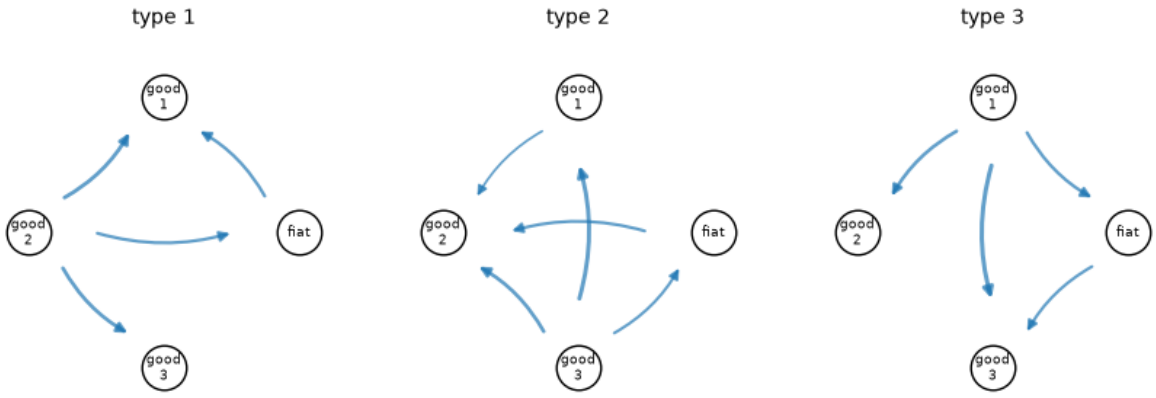


Fig. 95.7: Exchange pattern discovered in Economy C

The arrows into and out of the fiat node show money changing hands in both directions for every type: agents give up their production good to acquire it, and give it up to acquire the good they consume.

Good 1 still serves as a second money for type 2, which gives up good 3 for good 1 and good 1 for good 2.

The payoff-neutral trades of Economy A1.1 appear again: type 1 gives up good 2 for good 3, and type 3 gives up good 1 for good 2.

Across seeds:

```
H_c = final_holdings(across_seeds(economy_c, 1250))
n = len(H_c)
print(f"every type holds fiat at least 10% of the time: "
      f"{np.sum(H_c[:, :, 3].min(axis=1) >= 0.1)} of {n}")
print("holdings averaged over runs:")
print(pd.DataFrame(H_c.mean(axis=0), index=type_names(economy_c),
                   columns=good_names(economy_c)).round(2))
print(f"within 0.1 of the fundamental equilibrium:      "
```

(continues on next page)

(continued from previous page)

```
      f"{np.sum(distance(H_c, fundamental_c) < 0.1)} of {n}")
print(f"within 0.2 of the fundamental equilibrium:      "
      f"{np.sum(distance(H_c, fundamental_c) < 0.2)} of {n}")
print("holdings in runs where some type holds fiat less than 10% of the time:")
print(H_c[H_c[:, :, 3].min(axis=1) < 0.1].round(2))
```

```
every type holds fiat at least 10% of the time: 18 of 20
holdings averaged over runs:
        good 1  good 2  good 3  fiat
type 1    0.00    0.60    0.05  0.35
type 2    0.16    0.00    0.54  0.30
type 3    0.63    0.06    0.00  0.31
within 0.1 of the fundamental equilibrium:      2 of 20
within 0.2 of the fundamental equilibrium:      14 of 20
holdings in runs where some type holds fiat less than 10% of the time:
[[[0.   0.04 0.   0.96]
  [0.96 0.   0.04 0.  ]
  [0.98 0.01 0.01 0.  ]]

 [[0.   0.49 0.   0.51]
  [0.17 0.06 0.67 0.1 ]
  [0.52 0.13 0.   0.36]]]
```

In 18 of the 20 runs every type holds fiat money at least 10 percent of the time.

In one of the other two, type 1 agents hoard fiat money while types 2 and 3 hold almost none; in the other, type 2's share of fiat money is just under 10 percent.

On average across runs the holdings are within about 0.15 of the theoretical equilibrium: type 1 holds more fiat money and less good 2 than in theory, and type 2 less good 1 and more good 3, as in seed 0.

Those gaps are why only 2 runs come within 0.1 of the equilibrium, while 14 come within 0.2.

Nothing in the environment told the agents to do this.

Each agent discovered only that rules which end up holding the costless object earn more than rules that do not, and, because everyone was discovering this at once, the belief became self-confirming.

The discovery asks little of the agents: fiat money is the cheapest object to store, so even an impatient agent accepts it, provided others will take it in turn.

95.12 Economy D: five goods, five types

The last economy has five types and five goods, with production

type i	produces	consumes
1	good 3	good 1
2	good 4	good 2
3	good 5	good 3
4	good 1	good 4
5	good 2	good 5

storage costs $s = (1, 4, 9, 20, 30)$ from the paper's Table 4, and $u_i = 200$.

Each agent carries 180 exchange rules and 20 consumption rules.

The authors emphasize that they had *no analytical characterization of equilibrium* for this economy, and used the simulation to suggest what an equilibrium might look like.

```
economy_d = Economy(
    name='D',
    produces=np.array([2, 3, 4, 0, 1]),          # type 1 -> good 3, type 2 -> good 4,␣
 ↪...
    storage_costs=np.array([1.0, 4.0, 9.0, 20.0, 30.0]),
    u=200.0,
    random_rules=True,
    n_trade_rules=180,
    n_consume_rules=20,
)

sim_d = Simulation(economy_d, seed=0)
sim_d.run(2000)
holdings(sim_d)
```

	good 1	good 2	good 3	good 4	good 5
type 1	0.000	0.114	0.886	0.000	0.000
type 2	0.032	0.000	0.112	0.814	0.042
type 3	0.046	0.052	0.000	0.028	0.874
type 4	0.832	0.018	0.094	0.000	0.056
type 5	0.208	0.546	0.246	0.000	0.000

The time paths show how the holdings got there.

```
plot_holdings(sim_d)
```

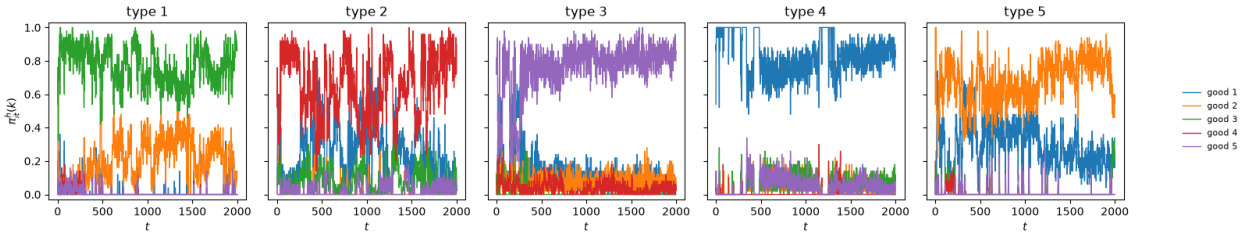


Fig. 95.8: Holdings in Economy D

Two features stand out.

First, the diagonal is zero: no type ends a period holding the good it consumes, because it eats it.

Second, every type holds mostly the good it produces; at the end of this run only type 5 holds other goods much of the time, and type 1 holds the cheaper good 2 about a tenth of the time.

The time paths have not fully settled; type 2's holdings in particular swing between goods 4 and 1 for much of the run.

The paper's run looks different in its details: at $t = 1750$ its type 1 agents held the cheaper good 2 83 percent of the time, and its type 3 agents split their holdings between good 2 and good 5.

Let us classify the trades that actually take place by what the agent acquires.

```
def trade_composition(sim, window=200):
    """
```

(continues on next page)

(continued from previous page)

```
    Shares of realized trades in which the agent acquires its own consumption
    good, a good cheaper to store than the one given up, or neither; and the
    share of 'neither' trades that give up the agent's own production good
    and are then eaten (payoff-neutral).
    """
    econ = sim.econ
    f = np.array(sim.exch_hist)[-window:].sum(axis=0)
    c = np.array(sim.cons_hist[-window:]).sum(axis=0)
    eat = c[:, :, 1] / np.maximum(c[:, :, 0], 1)
    s = econ.storage_costs
    out = {'own consumption good': 0.0, 'a cheaper good': 0.0, 'neither': 0.0}
    neutral = 0.0
    for i in range(econ.n_types):
        for j in range(econ.n_goods):
            for k in range(econ.n_goods):
                if j == k:
                    continue
                key = ('own consumption good' if k == i else
                       'a cheaper good' if s[k] < s[j] else 'neither')
                out[key] += f[i, j, k]
                if key == 'neither' and j == econ.produces[i]:
                    neutral += f[i, j, k] * eat[i, k]
    total = sum(out.values())
    out = {k: v / total for k, v in out.items()}
    out['  of which payoff-neutral'] = neutral / total
    return pd.Series(out, name='share of realized trades').round(3)


trade_composition(sim_d)
```

```
own consumption good         0.339
a cheaper good               0.307
neither                      0.354
  of which payoff-neutral    0.183
Name: share of realized trades, dtype: float64
```

About two thirds of realized trades acquire either the agent's own consumption good or a good that is cheaper to store, which is the pattern the paper describes.

About half of the remaining third are payoff-neutral: the agent gives up its own production good for a dearer good, eats it, and produces its own good again.

The rest, about a sixth of all trades, fit neither description.

Here are the same statistics for ten seeds, together with the number of times per period an agent of each type eats its own good.

```
sims_d = across_seeds(economy_d, 2000, seeds=range(10))
comp = pd.DataFrame([trade_composition(sim) for sim in sims_d])
print(comp.describe().loc[['mean', 'min', 'max']].round(2), "\n")
own = np.array([[np.array(sim.cons_hist[-200:]).sum(axis=0)[i, i, 1] / (200 * 50)
                 for i in range(5)] for sim in sims_d])
print("own-good consumption per agent per period, smallest over runs:",
      own.min(axis=0).round(3))
```

```
      own consumption good  a cheaper good  neither    of which payoff-neutral
mean                  0.33            0.31     0.37                       0.18
```

(continues on next page)

(continued from previous page)

```
min                   0.29            0.28     0.33                          0.15
max                   0.38            0.34     0.40                          0.23

own-good consumption per agent per period, smallest over runs: [0.04  0.1   0.041␣
↪0.059 0.07 ]
```

The shares change little from run to run, and every type eats its own good in every run.

The paper summarizes its own run this way:

> From the simulation results, we can see that the trading patterns nearly seem to describe a fundamental equilibrium in which agents are only willing to trade for commodities of lower cost than the one currently in storage, except that they always accept the commodity of their type. Some speculative moves can be detected.

Our runs share the first part of that description, with two qualifications: a third of the trades fit neither pattern, half of them because the "consume anything" rule makes them payoff-neutral, and in our displayed run the holdings have not settled down.

The paper's example of a speculative move – type 2 agents accepting good 3 from type 1 in exchange for good 1, in order to trade it to type 3 for good 2 – comes from its own run.

What the simulation delivers is a candidate pattern of trade, which would still have to be checked against the equilibrium conditions, as the authors note.

95.13 Concluding remarks

Marimon, McGrattan and Sargent's agents know very little.

They do not know their own utility functions, they do not know storage costs, they do not know the distribution of goods across the population, and they certainly do not solve dynamic programs.

They recognize utility when they experience it and costs when they bear them, and they keep running averages.

Out of this the following emerges.

- **A medium of exchange is learnable.** With a complete list of rules, Economy A1 reaches its unique stationary equilibrium in every run; from random rules the agents get most of the way, overeating money as the paper's agents did.
- **Selection, and its limits.** In Economy B, with two equilibria, the classifier systems select the fundamental one in most runs and the speculative one in none. In Economy A2, whose only equilibrium for patient agents is speculative, they never speculate: most runs converge, sometimes very slowly, to an outcome that is not an equilibrium of that economy, and a few lock into a pattern with even less trade.
- **Why: the agents are impatient.** At stationary strengths the bucket brigade values the future with a discount factor $b_1 b_2/((1+b_1)(1+b_2))$, about 0.016 with the paper's bids and below $1/4$ for any bids, and the patterns the agents find are those of impatient agents. The greedy auction and overgeneral consumption rules add to the effect.
- **Institutions can be discovered.** In nearly every run, Economy C's agents built a fiat monetary system out of nothing but their own experience of storage costs.
- **The method proposes candidates.** Economy D produced a pattern of trade in a model its authors had not solved, though not one that has been verified to be an equilibrium.

The paper is candid about what is missing.

There are no convergence theorems, only a sketch of how stochastic approximation arguments might supply them, and the authors judge their algorithm to provide "too little experimentation".

Experimentation is part of the story, but the stationary strengths show that more of it would not be enough: an agent built this way values a payoff one period ahead at less than a quarter of its worth, whatever its bids.

1993 Prospects for Bounded Rationality in Macroeconomics takes stock of what the series has delivered, and enters Economy A2 on the debit side of the ledger.

95.14 Exercises

Exercise 95.14.1

Economy A2.2 is Economy A2 – with $u_i = 500$ – but started from randomly generated rules and evolved with the genetic algorithm, exactly as A1.2 relates to A1.1.

The paper's summary table lists its equilibrium type as speculative, and the authors report that after 1000 iterations the economy had not converged, with trading patterns closer to the fundamental equilibrium than to the speculative one in period 1000 – "(The reverse was true in period 500.)"

Simulate it for 2000 periods and report holdings at $t = 500$, $t = 1000$ and $t = 2000$.

Then repeat for 20 seeds.

Does the extra experimentation supplied by the genetic algorithm produce speculation at any of these dates?

Solution

```
economy_a22 = replace(economy_a21, name='A2.2', random_rules=True)

sim_a22 = Simulation(economy_a22, seed=0)
sim_a22.run(2000)
for t in (500, 1000, 2000):
    print(f"t = {t}\n", holdings(sim_a22, t=t), "\n")
```

```
t = 500
          good 1  good 2  good 3
type 1   0.000   1.000   0.000
type 2   0.438   0.000   0.562
type 3   0.788   0.212   0.000

t = 1000
          good 1  good 2  good 3
type 1   0.000   1.000   0.000
type 2   0.428   0.000   0.572
type 3   0.722   0.278   0.000

t = 2000
          good 1  good 2  good 3
type 1   0.000   1.000   0.000
type 2   0.404   0.000   0.596
type 3   0.732   0.268   0.000
```

Now the same experiment for 20 seeds, recording holdings at each date.

```
sims_a22 = across_seeds(economy_a22, 2000)
for t in (500, 1000, 2000):
    H = final_holdings(sims_a22, t)
    closer = distance(H, speculative.to_numpy()) < distance(H, fundamental)
    print(f"t = {t}: type 1 holds good 3 > 5% in {np.sum(H[:, 0, 2] > 0.05)} runs "
          f"(at most {H[:, 0, 2].max():.2f} of the time); closer to speculative "
          f"than to fundamental in {np.sum(closer)} runs")
```

```
t = 500: type 1 holds good 3 > 5% in 2 runs (at most 0.16 of the time); closer␣
↪to speculative than to fundamental in 0 runs
t = 1000: type 1 holds good 3 > 5% in 1 runs (at most 0.12 of the time); closer␣
↪to speculative than to fundamental in 0 runs
t = 2000: type 1 holds good 3 > 5% in 3 runs (at most 0.15 of the time); closer␣
↪to speculative than to fundamental in 1 runs
```

No.

In this run type 1 never holds good 3 at any of the three dates.

Across 20 runs, at each date type 1 holds good 3 more than 5 percent of the time in at most three runs, and never as often as the speculative equilibrium's 29 percent.

Only one run, at $t = 2000$, is even marginally closer to the speculative holdings than to the fundamental ones, and our runs do not reproduce the paper's report that its run looked more speculative at $t = 500$.

Randomizing the initial rules and adding the genetic algorithm can manufacture rules that accept good 3, but it does not change the discount factor at which the bucket brigade values holding it (*The bucket brigade's discount factor*).

Exercise 95.14.2

Economy B.2 is Economy B started from random rules, with the genetic algorithm and A's bids.

The paper reports that it "had not converged after 2000 periods" but was "moving towards the fundamental equilibrium", with holdings at $t = 2000$ of $(0, 0.354, 0.646)$ for type 1, $(0.996, 0, 0.004)$ for type 2 and $(0.268, 0.732, 0)$ for type 3.

Simulate it, compare its holdings with the two equilibria of Economy B and with the paper's run, and compare its distance from the fundamental equilibrium with B.1's across 20 seeds.

Solution

```
economy_b2 = replace(economy_b1, name='B.2', random_rules=True,
                     b_trade=(0.025, 0.025))
sim_b2 = Simulation(economy_b2, seed=0)
sim_b2.run(2000)

paper_b2 = np.array([[0, 0.354, 0.646], [0.996, 0, 0.004], [0.268, 0.732, 0]])
rows = {'B.2 simulated, t = 2000': holdings(sim_b2).to_numpy(),
        'fundamental equilibrium': fundamental_b,
        'speculative equilibrium': speculative_b,
        "paper's B.2, t = 2000": paper_b2}
pd.concat({k: pd.DataFrame(v, index=type_names(economy_b2),
                           columns=good_names(economy_b2)).round(3)
           for k, v in rows.items()})
```

```
                                  good 1  good 2  good 3
B.2 simulated, t = 2000 type 1     0.000   0.100   0.900
                        type 2     0.856   0.000   0.144
                        type 3     0.112   0.888   0.000
fundamental equilibrium type 1     0.000   0.293   0.707
                        type 2     1.000   0.000   0.000
                        type 3     0.586   0.414   0.000
speculative equilibrium type 1     0.000   0.586   0.414
                        type 2     0.707   0.000   0.293
                        type 3     0.000   1.000   0.000
paper's B.2, t = 2000   type 1     0.000   0.354   0.646
                        type 2     0.996   0.000   0.004
                        type 3     0.268   0.732   0.000
```

Now compare B.2 with B.1 across seeds.

```
H_b2 = final_holdings(across_seeds(economy_b2, 2000))
for name, H in [("B.1", H_b1), ("B.2", H_b2)]:
    print(f"{name}: median distance to fundamental "
          f"{np.median(distance(H, fundamental_b)):.2f}, to speculative "
          f"{np.median(distance(H, speculative_b)):.2f}; "
          f"median share of good 1 for type 3 {np.median(H[:, 2, 0]):.2f}")
```

```
B.1: median distance to fundamental 0.06, to speculative 0.56; median share of␣
↪good 1 for type 3 0.56
B.2: median distance to fundamental 0.25, to speculative 0.52; median share of␣
↪good 1 for type 3 0.35
```

This run of B.2 is not near either equilibrium: type 3 holds good 1 only 11 percent of the time (59 percent in the fundamental equilibrium, 27 percent in the paper's run), and type 1 holds good 2 only 10 percent of the time (29 percent in the fundamental equilibrium, 59 in the speculative one).

Across seeds, B.2's median distance from the fundamental equilibrium is 0.25, against 0.06 for B.1, while both stay about 0.5 from the speculative one.

The typical type 3 agent in B.2 holds good 1 about a third of the time, between the paper's 27 percent and the equilibrium's 59.

Like the paper's run, B.2 ends much closer to the fundamental equilibrium than to the speculative one without reaching it in 2000 periods: building the needed rules from random material takes time that the enumerated economy never had to spend.

Exercise 95.14.3

Kiyotaki and Wright's agents cannot consume goods they do not value.

Impose that restriction on Economy A2.1 by setting `no_eat=True`, which removes the overconsumption mechanism of *Economy A2.1: when theory predicts speculation*.

Across 20 seeds, how often does type 1 hold good 3?

In the runs where it does, which rule of type 1's wins in the state (good 2, good 3), and what else does that rule cover?

Solution

```
economy_noeat = replace(economy_a21, name='A2.1, no eating', no_eat=True)
```

```
rows = []
for seed in range(20):
    sim = Simulation(economy_noeat, seed)
    sim.run(1000)
    h = holdings(sim, window=100).to_numpy()
    w = strongest(sim.E, 0, 1 * 3 + 2)
    f = trade_rates(sim)[0]
    rows.append({'seed': seed, 'type 1 holds good 3': h[0, 2],
                 'type 1 trades 2 for 3': f[1, 2], 'type 1 trades 3 for 1': f[2,␣
 ↪0],
                 'winning rule in (good 2, good 3)': rule_string(sim.E.cond[0, w],
                                                                 sim.E.action[0,␣
 ↪w], 3)})
noeat = pd.DataFrame(rows).set_index('seed').round(3)
noeat[noeat['type 1 holds good 3'] > 0.05]
```

```
      type 1 holds good 3  type 1 trades 2 for 3  type 1 trades 3 for 1  \
seed
3                   0.239                 0.179                 0.119
5                   0.360                 0.000                 0.000
6                   0.380                 0.000                 0.000
8                   0.278                 0.165                 0.101
10                  0.273                 0.164                 0.095
11                  0.363                 0.199                 0.089
12                  0.480                 0.000                 0.000
13                  0.600                 0.000                 0.000
18                  0.480                 0.000                 0.000

     winning rule in (good 2, good 3)
seed
3                          0## #0# -> 1
5                          010 #0# -> 0
6                          ##0 #0# -> 0
8                          010 #0# -> 1
10                         ##0 #0# -> 1
11                         ##0 #0# -> 1
12                         ##0 001 -> 0
13                         ##0 #0# -> 0
18                         ##0 #0# -> 0
```

Type 1 holds good 3 more than 5 percent of the time in 9 of the 20 runs, but for two different reasons.

In four runs (seeds 3, 8, 10 and 11) type 1 trades good 2 for good 3 and good 3 for good 1, and its holdings look like the speculative equilibrium's.

But the rule that accepts good 3 is general, a version of "trade if my partner does not hold good 2", which also covers the state in which the partner holds good 1.

Its strength is earned in that state, where trading leads to eating good 1, and carried over to good 3: speculation by generality, not by the bucket brigade crediting the patience it requires.

In the other five runs type 1 holds good 3 but never trades it.

These agents acquired good 3 early on, cannot eat it, and meet type 3 agents who have locked into refusing it.

So removing overconsumption does not make the agents speculate for the right reason, and generality cuts both ways: it can carry rewards from one state to another, which here mimics speculation, and it produces the overconsumption of *Economy A2.1: when theory predicts speculation*.

This is the tradeoff between generality and discrimination that Sargent [1993] emphasizes in his discussion of classifier systems.

Exercise 95.14.4

Replace the argmax auction of Economy A2.1 by a softmax auction with intensity $\eta = 0.5$ (set `eta=0.5`), in the spirit of Exercise **??** and the paper's footnote 5, leaving everything else unchanged.

Across 20 seeds, how often does type 1's winning exchange rule in the state (good 2, good 3) accept good 3, and how often does type 1 actually hold good 3?

Compare your answer with the paper's Table 9d.

Solution

```
economy_soft = replace(economy_a21, name='A2.1, softmax', eta=0.5)

rows = []
for seed in range(20):
    sim = Simulation(economy_soft, seed)
    sim.run(1000)
    w, c = strongest(sim.E, 0, 1 * 3 + 2), strongest(sim.C, 0, 2)
    rows.append({'accepts good 3': sim.E.action[0, w] == 1,
                 'type 1 holds good 3': holdings(sim, window=100).iloc[0, 2],
                 'consumption rule for good 3': rule_string(sim.C.cond[0, c],
                                                            sim.C.action[0, c], 3)}
 ↪)
soft = pd.DataFrame(rows)
print(f"winning rule accepts good 3: {soft['accepts good 3'].sum()} of 20")
print(f"largest share of time type 1 holds good 3: {soft['type 1 holds good 3'].
 ↪max():.3f}")
soft['consumption rule for good 3'][soft['accepts good 3']].value_counts()
```

```
winning rule accepts good 3: 14 of 20
largest share of time type 1 holds good 3: 0.008
```

```
consumption rule for good 3
#0# -> 1    11
0## -> 1     3
Name: count, dtype: int64
```

With exploration, the rule that accepts good 3 wins in 14 of the 20 runs, yet type 1 almost never holds good 3.

In every run in which it accepts good 3, the winning consumption rule for good 3 is a general eating rule: either "eat if I do not hold good 2", which also covers good 1, or "eat if I do not hold good 1", which also covers good 2, where eating does no harm.

This is the paper's Table 9d pattern: the speculative exchange rule wins, and the general consumption rule eats good 3 at once.

A softmax auction keeps experimenting, so no rule is shut out for good by an early windfall, but it leaves the other two mechanisms in place.

Exercise 95.14.5

The bucket brigade's discount factor predicts that a type 1 agent whose rules are all fully specific will not speculate in Economy A2 unless u_1 exceeds roughly $6(s_3 - s_2)/\beta_{\text{eff}}$.

Test the prediction by simulation.

Give the agents only fully specific rules (`specific_only=True`), let them explore with a softmax auction ($\eta = 0.5$), forbid eating unwanted goods, and vary u_1.

For six seeds at each u_1, record whether type 1's winning rule in the state (good 2, good 3) accepts good 3.

At what u_1 does speculation start to win?

Solution

```
print(f"predicted thresholds: {speculation_threshold(s_A, 1, 0.5, β_A):,.0f} "
      f"(fundamental holdings), {speculation_threshold(s_A, 1, 0.586, β_A):,.0f} "
      f"(speculative holdings)\n")

for u in (500, 4_000, 8_000, 12_000, 20_000, 40_000):
    econ = replace(economy_a21, u=float(u), specific_only=True, eta=0.5,
                   no_eat=True)
    accept, share = 0, []
    for seed in range(6):
        sim = Simulation(econ, seed)
        sim.run(2000)
        accept += sim.E.action[0, strongest(sim.E, 0, 1 * 3 + 2)] == 1
        share.append(holdings(sim, window=200).iloc[0, 2])
    print(f"u_1 = {u:>6,}: accepts good 3 in {accept} of 6 runs; "
          f"type 1 holds good 3 {np.mean(share):.2f} of the time")
```

```
predicted thresholds: 7,182 (fundamental holdings), 8,674 (speculative holdings)
u_1 =    500: accepts good 3 in 0 of 6 runs; type 1 holds good 3 0.03 of the␣
 ↪time
u_1 =  4,000: accepts good 3 in 0 of 6 runs; type 1 holds good 3 0.03 of the␣
 ↪time
u_1 =  8,000: accepts good 3 in 0 of 6 runs; type 1 holds good 3 0.03 of the␣
 ↪time
u_1 = 12,000: accepts good 3 in 0 of 6 runs; type 1 holds good 3 0.05 of the␣
 ↪time
u_1 = 20,000: accepts good 3 in 2 of 6 runs; type 1 holds good 3 0.11 of the␣
 ↪time
u_1 = 40,000: accepts good 3 in 4 of 6 runs; type 1 holds good 3 0.32 of the␣
 ↪time
```

In these six runs, at A2's $u_1 = 500$ and at every u_1 up to 12,000, the rule that accepts good 3 never wins.

It starts to win between 12,000 and 20,000, and at 40,000 type 1 holds good 3 about as often as in the speculative equilibrium.

The prediction gets the order of magnitude right: speculation needs a utility tens of times larger than A2's, against a threshold of 138 for patient agents.

These six runs need somewhat more than the prediction, plausibly because the auction compares noisy strengths in runs of finite length; on other seeds the accepting rule already wins occasionally at 12,000.

CHAPTER

NINETYSIX

1993 PROSPECTS FOR BOUNDED RATIONALITY IN MACROECONOMICS

Contents

- *1993 Prospects for Bounded Rationality in Macroeconomics*
  - *Overview*
  - *The quest, restated*
  - *The debit side: how much we must hard-wire*
  - *Why the econometricians have not returned the compliment*
  - *The credit side: selection, computation, and a returned gift*
  - *Reading the ledger against the quest*
  - *Postscript: what became of the program*
  - *Exercises*

96.1 Overview

This is the closing lecture of the series, and it is different in kind from the others.

It carries no new model.

Instead it gathers the judgments that Sargent [1993] recorded in 1993 — the opinions, reservations, and hopes with which he ended *Bounded Rationality in Macroeconomics* — and loops them back to the quest that opened the book, and that opened *A Peculiar Definition of Bounded Rationality*, the first lecture here.

That quest had a destination: a theory of **transition dynamics**, the out-of-equilibrium adjustment that the Eastern European reformers of 1989 had to manage with no map.

The route was to expel the rational agents from our models and replace them with "artificially intelligent" agents who behave like **econometricians**, who gather data, form theories, estimate, and adapt.

We can now ask how far that route carried us, by 1993, toward the destination.

Sargent's own answer is a ledger, with entries on both sides.

On the credit side: adaptive dynamics as a device for **selecting** among equilibria, as a tool — evolutionary programming — for **computing** them, and as a source of new tools for econometricians.

On the debit side: the original prize, a theory of transition dynamics, is largely unclaimed; the agents must be heavily prompted; and a striking asymmetry has emerged.

The program set out to make the agents in our models behave more like econometricians.

The econometricians, Sargent observes, have not returned the compliment.

This lecture explains that asymmetry, gives it a concrete numerical face, and asks when it fails.

It then settles, one by one, the four payoffs that *A Peculiar Definition of Bounded Rationality* promised, including one, new sources of dynamics, that the series delivers but Sargent's closing ledger leaves out.

The lecture closes with a postscript on what became of the program after 1993.

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
```

96.2 The quest, restated

The knowledge that rational expectations imputes set out the argument against which everything below is measured.

Rational expectations imposes two requirements, individual rationality and mutual consistency of perceptions, and when a rational expectations model is taken to data it imputes to the agents inside it far more knowledge than the econometrician studying them has.

The bounded rationality program proposes to close that gap from the other side, by making the agents learn the distributions from data, as the econometrician must.

The hope was a description of the system *while it is still adjusting*, before beliefs and outcomes have settled into mutual consistency: a theory of transition dynamics.

What the program is good for listed four payoffs the program might deliver: equilibrium selection, computation, transition dynamics and the analysis of regime changes, and new sources of dynamics.

We take stock following Sargent [1993]'s own accounting: first the debits, then the central asymmetry, then the credits; and then we settle each of the four payoffs.

96.3 The debit side: how much we must hard-wire

The first reservation is about **arbitrariness**.

Bounded rationality is most easily defined by what it is *not* — rational expectations — and that very malleability is a liability.

Once we stop insisting that agents know the equilibrium, we must decide, case by case, exactly what they *do* know and how they learn it.

Do they know their own utility and profit functions, or must they learn those too?

Do they know calculus and dynamic programming, or only trial and error?

Do they learn from their own experience alone, or from others'?

To what class of approximating functions do we confine what they learn about?

Every model in this series answered these questions by **hard-wiring**, by prompting the agents heavily, with an eye on the outcome we hoped they would reach.

The series descended a gradient of prompting, telling its agents a little less at each step.

Bray's least squares learners (*From relaxation to learning: Bray's market*; see also *Least Squares Learning in Self-Referential Models*), who reappear below, know the correct supply curve and need only estimate one conditional expectation to plug into it.

The Marcet–Sargent agents of the book's investment example (see *Least Squares Learning in Self-Referential Models*) know dynamic programming, know their return function, and know the parametric form of the law of motion — they lack only its coefficients, which they update by vector autoregression.

The Robbins–Monro agents of *Adaptive Agents in an Overlapping Generations Monetary Economy* know their utility function but not the distributions they face; the agents of *Learning, Approximation, and Equilibrium Computation* must in addition choose a functional form; the Newton–Raphson learners of *Exchange Rate Indeterminacy, Learning, and Experiments* know the gradient of their own payoff.

Even the least prompted agents — the genetic-algorithm strings of *Exchange Rate Indeterminacy, Learning, and Experiments* and the classifier systems of *Networks, Genetic Algorithms, and Classifier Systems* and *Money as a Medium of Exchange among Artificially Intelligent Agents*, which see only a payoff signal and recognize utility only when they experience it — are still told *when* to choose and *what* information to condition on, and the entire apparatus of their accounting system and genetic operators is designed by hand, with the Kiyotaki–Wright equilibrium in view.

Sargent named the consequence: "a concomitant sensitivity of outcomes to details of adaptive algorithms."

Economy A2 of *Money as a Medium of Exchange among Artificially Intelligent Agents* is the sharpest case in the series (*Economy A2.1: when theory predicts speculation*).

For patient agents, the speculative equilibrium is the *unique* stationary equilibrium of that economy, yet the classifier agents settle on the fundamental pattern of trade.

That is not learning selecting among equilibria, since there is nothing to select among.

It is convergence to an outcome that is not an equilibrium of the economy the agents live in.

Marimon *et al.* [1990] suspected early myopia of "diverting the economy towards a low discount-factor stationary equilibrium, such as the fundamental equilibrium", and *The bucket brigade's discount factor* shows why: the bucket brigade makes the agents effectively very impatient.

The outcome is a property of how the agents were built, and the paper itself reports (its footnote 15) that a genetic algorithm studied by Marimon and Miller reached the speculative equilibrium in 25 of 30 runs.

The second reservation is about **simplicity**.

The learning tasks we set our agents are trivial next to those in a first econometrics course, let alone those that real firms and households are implicitly solving.

We ask an agent to learn a single time-invariant decision rule, or a fixed collection of conditional expectations.

We do not ask it to learn the parameters of a simultaneous-equations system, or to infer a mapping from policy regimes to distributions, the tasks that make econometrics hard.

Even the choice of what to learn is a prompt.

Agents in *Learning, Approximation, and Equilibrium Computation* who are confined to a low-order polynomial converge to an approximate equilibrium, a fixed point of the restricted problem rather than the rational expectations equilibrium, and the modeller's choice of family decides how close they come.

Put together, these reservations bear directly on the original prize.

The environments into which we have cast our adaptive agents are far more stable and hospitable than the transitions we actually care about.

Convergence-rate results are scarce, and tractability forces us to restrict the distribution of agents' beliefs severely.

So the literature on adaptive processes, Sargent judged in 1993, falls well short of a secure foundation for a theory of real-time transition dynamics, the very thing the quest set out to find.

He declines to end on that failure, though, and the baseball metaphor he reaches for is deliberately modest:

> It would not be wise or fair to end this essay by dwelling on the failure of adaptive methods so far to have 'hit a home run' by giving us a good theory of transition dynamics. The problem of transition dynamics is difficult and long-standing. So maybe it should count as a single, or at least a sacrifice fly, that these methods have sharpened our appreciation of the problem.

96.4 Why the econometricians have not returned the compliment

Here is the asymmetry that gives this lecture its theme.

The bounded rationality program is, at bottom, a movement to make the agents in our models behave more like the econometricians who build and estimate those models.

Imitation is the sincerest form of flattery.

We might therefore have expected macroeconometricians to rush to fit these models to data.

There was no rush.

Chung [1990] estimated the Sims policy-maker model (*A government learning the Phillips curve*), in which a government runs a Kalman filter (*A First Look at the Kalman Filter*) on a Phillips curve whose coefficients it believes drift as random walks.

Along with the structural parameters, Chung estimated the variances of the drift and the prior mean that started the government's filter.

Sargent called this "the only econometrically serious macroeconomic implementation of bounded rationality of which I know."

Why the reluctance?

The reasons are worth spelling out, because they are not about taste.

The governing dictum among applied econometricians is Lucas's: *beware of theorists bearing free parameters*.

Replacing a rational agent with a boundedly rational one **adds** parameters: parameters describing beliefs and how beliefs move.

Take the simplest case, Bray's model.

Relative to its rational expectations version, the adaptive version adds at least the initial belief and a parameter setting the gain sequence, and one might want more parameters to describe the gain's shape.

That would already give Lucas's warning something to bite on.

But there is a deeper problem, and it is the reason the added parameters are not just unwelcome but genuinely hard to estimate.

Because Bray's adaptive system **converges** to the rational expectations equilibrium, which in her model is unique, the extra parameters influence only the *transient*.

The asymptotic distribution of the data contains no information about them.

Let us see this directly.

96.4.1 The nuisance parameter, made concrete

Bray's cobweb economy [Bray, 1982] sets the market price by

$$p_t = a + b\,\beta_t + u_t,$$

where β_t is the price agents expect, formed by averaging past prices,

$$\beta_t = \beta_{t-1} + \gamma_t\,(p_{t-1} - \beta_{t-1}), \qquad \gamma_t = \frac{1}{t + t_0},$$

and u_t is an IID standard normal shock.

The constant t_0 is the weight the agents attach to the belief they start with, measured in observations: with $t_0 = 50$ they treat β_0 as though it summarized fifty prior prices, so that after t periods β_0 still carries weight $t_0/(t + t_0)$.

Some such weight is needed for the exercise to have any content.

With $t_0 = 0$ the first gain is $\gamma_1 = 1$, so $\beta_1 = p_0$ exactly and the initial belief is erased after a single period — there would be no transient left for the econometrician to try to estimate.

When $b < 1$ the belief converges to the rational expectations value $\beta^\star = a/(1 - b)$, whatever value it started from.

That is Bray's result, and it is the E-stability condition of *One recursion, many learners*: the associated differential equation $d\beta/d\tau = T(\beta) - \beta = a + (b - 1)\beta$ is stable exactly when $b < 1$.

```
a, b = 5.0, 0.7
t0 = 50                                          # weight on the initial belief
β_star = a / (1 - b)                             # rational expectations belief

def simulate(β0, u):
    "Bray's cobweb under least squares learning, given a shock path u."
    T = len(u)
    β = np.empty(T)
    p = np.empty(T)
    β[0] = β0
    for t in range(T):
        p[t] = a + b * β[t] + u[t]
        if t + 1 < T:
            β[t + 1] = β[t] + (1 / (t + 1 + t0)) * (p[t] - β[t])
    return p, β

print(f"rational expectations belief β* = {β_star:.3f}")
```

```
rational expectations belief β* = 16.667
```

The initial belief β_0 is the extra "bounded rationality" parameter.

Watch three economies with very different initial beliefs head toward $\beta^\star$.

```
rng = np.random.default_rng(0)
u = rng.standard_normal(400)

fig, ax = plt.subplots(figsize=(7.5, 4))
for β0, colour in [(2.0, 'C0'), (16.667, 'C1'), (40.0, 'C2')]:
    _, β = simulate(β0, u)
    ax.plot(β, color=colour, lw=1.3, label=fr"$\beta_0 = {β0}$")
ax.axhline(β_star, color='k', ls='--', lw=0.8, label=r"$\beta^\star$")
ax.set_xlabel("$t$")
```

(continues on next page)

(continued from previous page)

```
ax.set_ylabel(r"belief $\beta_t$")
ax.legend(frameon=False)
plt.show()
```

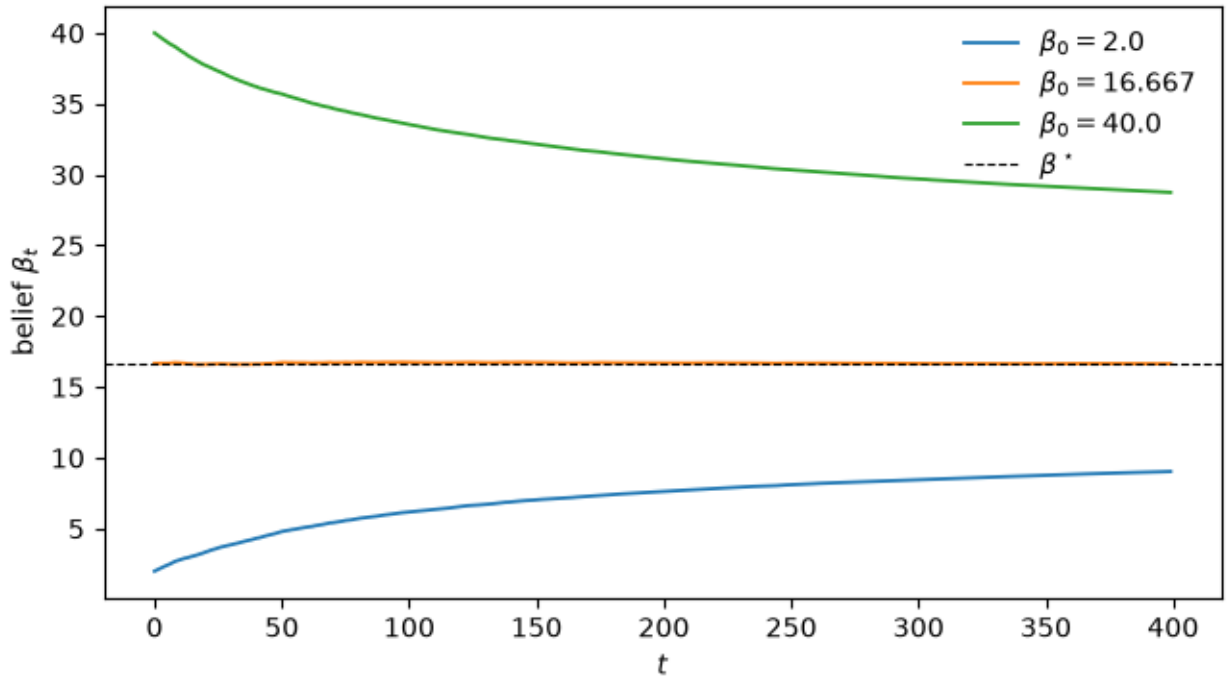


Fig. 96.1: Beliefs drift toward the rational expectations value, slowly

All three head toward $\beta^\star$, but slowly.

Each period the learning recursion closes the gap $\beta_t - \beta^\star$ by the factor $1-(1-b)/(t+t_0)$ on average, so after t periods roughly $(t_0/(t+t_0))^{1-b}$ of it remains.

```
for β0 in (2.0, 40.0):
    _, β = simulate(β0, u)
    share = (β[-1] - β_star) / (β0 - β_star)
    print(f"β0 = {β0:4.1f}: share of the initial gap left after {len(u)} periods =
 ↪{share:.2f}")
print(f"approximation (t0 / (t + t0))^(1 - b) = {(t0 / (len(u) - 1 + t0)) ** (1 - b):.
 ↪2f}")
```

```
β0 =  2.0: share of the initial gap left after 400 periods = 0.52
β0 = 40.0: share of the initial gap left after 400 periods = 0.52
approximation (t0 / (t + t0))^(1 - b) = 0.52
```

About half the initial gap is still there after 400 periods.

That is slower than the $1/\sqrt{t}$ rate at which an ordinary estimator settles down, as the book's chapter 5 warns: least squares learning converges at the $\sqrt{T}$ rate only when the slope of the associated differential equation, here $b-1$, is below $-1/2$.

Slow forgetting by the *economy*, however, is not the same thing as information for the *econometrician*.

Now put on the econometrician's hat.

Given a sample of prices and the structural parameters (a, b), the model implies a shock $\hat{u}_t = p_t - a - b\,\beta_t(\beta_0)$ for any candidate initial belief β_0, because the belief path is pinned down by β_0 and the observed prices.

The sum of squared implied shocks measures how well a given β_0 fits the data.

```
def belief_path(β0, p):
    "Belief sequence implied by an initial belief and an observed price path."
    T = len(p)
    β = np.empty(T)
    β[0] = β0
    for t in range(1, T):
        β[t] = β[t - 1] + (1 / (t + t0)) * (p[t - 1] - β[t - 1])
    return β

def ssr(β0, p):
    "Sum of squared implied shocks, as a function of the belief parameter β0."
    β = belief_path(β0, p)
    return np.sum((p - a - b * β) ** 2)
```

Whether the data can pin β_0 down is a question of how fast *information* about it accumulates as the sample grows.

The natural way to read that off is the **excess** sum of squares — how much worse a wrong β_0 fits than the best one.

For an ordinary, well-identified parameter every new observation adds to the penalty for being wrong, so the excess grows in proportion to T and the confidence interval shrinks like $1/\sqrt{T}$.

Because the shocks have unit variance, the excess is also a likelihood-ratio statistic, so a penalty above 3.84 rejects the wrong value at the 5 percent level.

Watch what happens here.

```
grid = np.linspace(2, 40, 80)

def excess_curve(T, seed=1):
    "SSR(β₀) − min SSR across the β₀ grid, for a sample of length T."
    u = np.random.default_rng(seed).standard_normal(T)
    p, _ = simulate(β_star, u)
    curve = np.array([ssr(b0, p) for b0 in grid])
    return curve - curve.min()

fig, ax = plt.subplots(figsize=(7.5, 4))
for T, colour in zip((50, 500, 5000), ('C0', 'C1', 'C2')):
    ax.plot(grid, excess_curve(T), color=colour, lw=1.5, label=f"$T = {T}$")
ax.axvline(β_star, color='k', ls='--', lw=0.8)
ax.set_xlabel(r"belief parameter $\beta_0$")
ax.set_ylabel("excess sum of squares")
ax.legend(frameon=False)
plt.show()
```

The curve for $T = 50$ is lower, but beyond a few hundred observations the curves lie nearly on top of one another instead of getting steeper.

To see that this is not an artifact of the range plotted, compare the penalty for a wrong β_0 with the penalty for a wrong *slope* b, estimated from the same data.

We compute the slope penalty twice: once holding the intercept a at its true value, and once re-estimating a along with b, and we average each penalty over 20 samples.

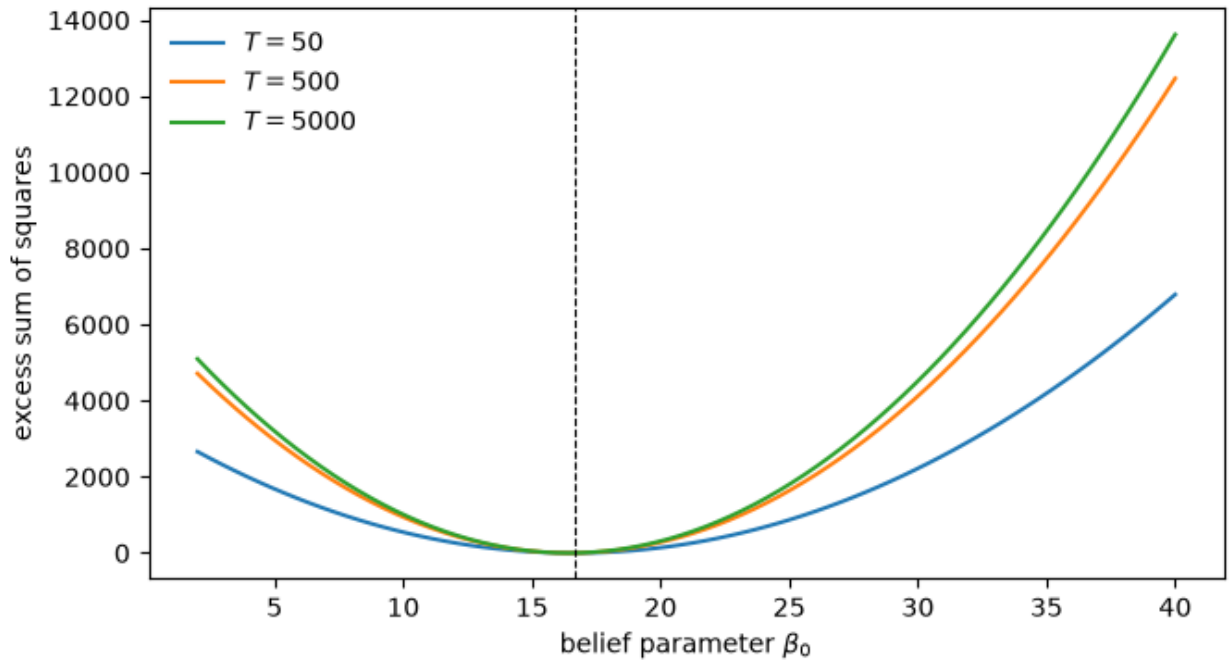


Fig. 96.2: Information about the belief parameter stops accumulating

```
def ssr_a_free(r):
    "Sum of squared residuals after re-estimating the intercept (removing the mean)."
    return np.sum((r - r.mean()) ** 2)

print(f"{'':>6}   {'β₀ = 2':>8}   {'b = 0.75':>10}   {'b = 0.75':>14}   {'b = 0.75, a␣
↪re-est.':>19}")
print(f"{'T':>6}   {'':>8}   {'a known':>10}   {'a re-estimated':>14}   {'largest of␣
↪20':>19}")
for T in (50, 500, 5_000, 50_000):
    pen_β0, pen_b_known, pen_b_free = [], [], []
    for seed in range(20):
        u = np.random.default_rng(seed).standard_normal(T)
        p, _ = simulate(β_star, u)
        β = belief_path(β_star, p)              # belief path at the true β0
        pen_β0.append(ssr(2.0, p) - ssr(β_star, p))
        pen_b_known.append(np.sum((p - a - 0.75 * β) ** 2)
                           - np.sum((p - a - b * β) ** 2))
        pen_b_free.append(ssr_a_free(p - 0.75 * β) - ssr_a_free(p - b * β))
    print(f"{T:>6}   {np.mean(pen_β0):8.1f}   {np.mean(pen_b_known):10.1f}   "
          f"{np.mean(pen_b_free):14.3f}   {np.max(pen_b_free):19.3f}")
```

```
           β₀ = 2     b = 0.75         b = 0.75   b = 0.75, a re-est.
     T                 a known   a re-estimated         largest of 20
    50      2655.4        37.0            0.032                 0.061
   500      4807.9       354.8            0.130                 0.301
  5000      5229.7      3447.5            0.204                 0.290
 50000      5278.7     34702.3            0.274                 0.671
```

The columns behave completely differently.

The penalty for getting the initial belief wrong stops growing.

Past a few thousand observations it is pinned near a constant, and every further observation is uninformative about β_0.

The confidence interval for β_0 never shrinks; the parameter is not consistently estimable at all.

The reason is visible in the belief recursion.

Given the observed prices, the econometrician's reconstructed belief at t depends on β_0 only through the weight $t_0/(t+t_0)$, and the squares of those weights have a finite sum.

The rest of β_0's long shadow on the economy is carried by the observed prices themselves.

The penalty for a wrong slope with a held at its true value grows in proportion to the sample: a hundredfold more data makes a hundredfold stronger case against the wrong value, which is what identification looks like.

But this identification is borrowed.

With a held fixed, a wrong b shifts the predicted mean price, and the mean price $a/(1-b)$ is something the rational expectations version of the model identifies too.

What a learning model adds is the possibility of separating a from b, and for that the data have to see beliefs move.

The third column shows how little they see.

With a re-estimated, the penalty is a fraction of one and creeps up by a tenth or less with each tenfold increase in the sample; in none of the 20 samples does it come near 3.84, even with 50,000 observations.

Under a $1/t$ gain belief motion dies out, so the learning transient is the only thing that separates a from b: the weakness that makes β_0 unidentifiable afflicts the slope too, in milder form.

This is the technical heart of the matter.

The parameters that bounded rationality adds live in the transient, and under a decreasing gain so does the information that would separate the old parameters from one another.

A transient contributes only a limited amount of information no matter how long we watch the economy afterwards, so those parameters become a **nuisance to estimate** — they enter the likelihood, and the data have little to say about them.

Sargent adds one more reason for the econometricians' cool response, with a characteristic "maybe".

Many applied macroeconometricians are in the market for methods that *reduce* the number of parameters needed to explain the data.

A reduction is precisely what bounded rationality does not offer.

It offers more parameters, most of them weakly identified.

So the flattery ran one way.

The theorists remade their agents in the econometrician's image; the econometricians, offered models full of extra, poorly-identified parameters, and warned by Lucas against exactly that, declined the gift.

96.4.2 When the equilibrium is not unique

The argument assumed that learning converges to a unique, E-stable equilibrium, as it does in Bray's model with $b < 1$.

Then initial beliefs, and the shape of the gain, can only affect the transient.

Under multiplicity that premise can fail.

In *Adaptive Agents in an Overlapping Generations Monetary Economy* the initial belief decides whether the economy settles at the low-inflation steady state or collapses (*The stability reversal*).

In *Exchange Rate Indeterminacy, Learning, and Experiments*, at the calibration where the Newton–Raphson learner converges, the initial saving rates and portfolio shares fix the limiting exchange rate forever (*The dead hand of history*).

Where initial beliefs select among a continuum of limits, as in the exchange-rate economy, their effect is permanent, and a permanent effect is one that every new observation helps to measure.

In a deterministic economy like that one, observed with error, the data measure the effect of the initial beliefs at the usual rate, as they would a structural parameter.

Exercise 96.8.2 shows this, and also shows that permanence is not always enough.

When beliefs are driven by observed prices, as in Bray's model, the econometrician can reconstruct them from the data, and the initial belief's weight in that reconstruction dies away whatever the equilibrium structure.

Shocks blur the picture in the exchange-rate economy too: once they enter, even identical starting points converge to different exchange rates (Exercise **??**), so the limit no longer reflects the initial beliefs alone.

96.5 The credit side: selection, computation, and a returned gift

The ledger is not one-sided.

Set against those debits, Sargent lists three successes, and the last of them quietly undoes some of the asymmetry just described.

In his opening chapter he also counts as a credit the models' "success in inspiring useful tests of equilibrium selection", meaning laboratory experiments.

Equilibrium selection.

Where a rational expectations model has many equilibria, a system of adaptive agents often converges to a *particular* one, turning learning into a device for selecting among them.

The clearest case in the series is the economy of *Adaptive Agents in an Overlapping Generations Monetary Economy*, where least squares learners select the low-inflation steady state that the rational expectations dynamics reject (*The stability reversal*), and where the laboratory subjects of Marimon and Sunder [1993] went the same way.

With no deficit, the same learners kill the bubbles of *A Peculiar Definition of Bounded Rationality* and select its bubbleless equilibrium (Exercise **??**).

In *Money as a Medium of Exchange among Artificially Intelligent Agents*, Economy B has both a fundamental and a speculative equilibrium, and the classifier systems select the fundamental one; in Economy A1, where the fundamental equilibrium is the only one, they find it.

The exchange rate of *Exchange Rate Indeterminacy, Learning, and Experiments* earns credit only in a weaker sense: learning determines the exchange rate only by remembering where it started, and the limit also moves with the tuning of the algorithm, its stabilizer κ and its gain (*The dead hand of history*).

Even that much holds only where the Newton–Raphson learner converges; at Arifovic's calibration, with the same tuning, it does not converge at all.

The book granted that this "may seem a weak reed on which to base exchange rate determination".

Economy A2 belongs on the debit side, as we saw above.

Nor is selection a welfare theorem.

In the book's example in the spirit of Brock, which *Adaptive Agents in an Overlapping Generations Monetary Economy* recalls as a warning, least squares selects a stationary equilibrium that every nonstationary equilibrium Pareto-dominates.

Sargent is candid that his affection for this use sits in some tension with his doubts about the dynamics that perform the selection:

> I know that it is inconsistent to doubt the real-time dynamics but keep the equilibria selected by them. I confess that my affection for the selection performed in the monetary models described in Chapter 6 is partly driven by my prior conviction that the selected equilibria seem sensible to me.

Chapter 6 of the book is its chapter on laboratory experiments: Marimon and Sunder's with the one-currency economy of *Adaptive Agents in an Overlapping Generations Monetary Economy*, and Arifovic's with the two-currency economy of *Exchange Rate Indeterminacy, Learning, and Experiments*.

The selection Sargent has in mind is the low-inflation one, on which least squares learning and Marimon and Sunder's subjects agreed.

Evolutionary programming.

If a population of adaptive agents reliably converges to an equilibrium, we can run the population as a *method of computing* the equilibrium, especially in models too complicated to solve by hand (*Evolutionary programming*).

The procedure is guess-and-verify: the adaptive system supplies the guess, and the economist checks that it satisfies the equilibrium conditions.

That is what Marimon *et al.* [1990] proposed for a five-good Kiyotaki–Wright economy, for which no analytical characterization was in hand (*Money as a Medium of Exchange among Artificially Intelligent Agents*): their classifier economy settled down to trading strategies and holding probabilities that suggested an equilibrium ("we could then presumably verify whether it is an equilibrium"), although the paper itself reports only that the patterns "nearly seem to describe" a fundamental equilibrium.

The cleanest instance in the series is *Learning, Approximation, and Equilibrium Computation*.

There Marcet's method of parameterized expectations, a tool for computing equilibria, has a learning counterpart in which agents update their forecasting rule as data arrive, and the two share their fixed points and, through the associated differential equation, their stability condition (*Parameterized expectations*).

The book had already remarked that Marcet's method "has many aspects of a learning or adaptation algorithm".

Used this way, the dynamics need not be credible as dynamics.

The book's chapter 7 describes how Randall Wright, computing a Kiyotaki–Wright equilibrium with endogenous shares of types, told his evolving population the theoretical payoffs it could have learned only from very long histories: hopeless as a model of real time, harmless as a computation.

That division of labor is how one can doubt the dynamics and still compute with them.

Sargent expected this use to be applied often.

New tools for the econometrician.

Here the asymmetry bends back on itself.

The literatures on parallel and genetic algorithms have handed econometricians new computational gadgets — genetic algorithms and stochastic Gauss–Newton procedures among them — for solving their *own* estimation and optimization problems.

Ellen McGrattan, for instance, applied genetic algorithms to estimate nonlinear rational expectations models, using them to search for the vicinity of a maximum before switching to a Newton method, a use noted in *Networks, Genetic Algorithms, and Classifier Systems*.

So the econometricians took up the adaptive algorithms after all, not as *models* of the agents they study, but as *tools* in their own hands.

The compliment was returned, but through a side door: the algorithms crossed over, the models did not.

New sources of dynamics.

The book's second chapter listed one more payoff: adaptive agents might generate behavior that rational expectations models find hard to produce, such as the trading volume that the no-trade theorem rules out.

Sargent's closing ledger does not post it, but the series delivers two instances.

Arifovic's genetic-algorithm economy in *Exchange Rate Indeterminacy, Learning, and Experiments* produces an exchange rate that never settles down, even though saving settles near its equilibrium value (*A genetic algorithm economy*), and even though every equilibrium of that economy in which both currencies are valued has a constant exchange rate.

And the constant-gain government of *Adaptive Agents in an Overlapping Generations Monetary Economy* produces a recurrent sawtooth of escapes from the self-confirming inflation rate toward zero inflation (*A government learning the Phillips curve*), dynamics that the model's unique self-confirming equilibrium does not have.

96.6 Reading the ledger against the quest

Line the two columns up against the destination we set out for, and against the four payoffs of *What the program is good for*.

payoff	verdict	where the series shows it
equilibrium selection	credit, qualified	*Adaptive Agents in an Overlapping Generations Monetary Economy* (low inflation); *Money as a Medium of Exchange among Artificially Intelligent Agents* (Economy B); only in a weaker sense in *Exchange Rate Indeterminacy, Learning, and Experiments*; Economy A2 is a debit
computation	credit	*Learning, Approximation, and Equilibrium Computation*; *Networks, Genetic Algorithms, and Classifier Systems*; *Money as a Medium of Exchange among Artificially Intelligent Agents* (the five-good economy); tools for econometricians
transition dynamics and regime changes	debit in 1993, partly paid since	escape dynamics in *Adaptive Agents in an Overlapping Generations Monetary Economy*, characterized in *Escaping Nash Inflation*
new sources of dynamics	credit, absent from the 1993 ledger	*Exchange Rate Indeterminacy, Learning, and Experiments* (genetic-algorithm volatility); *Adaptive Agents in an Overlapping Generations Monetary Economy* (the escape sawtooth)

The **debit** column holds the prize itself.

Sargent had argued in his second chapter that comparing rational expectations equilibria across policy regimes is itself inconsistent with rational expectations, and a theory of how adaptive agents move from one regime to another was supposed to repair that.

A theory of real-time transition dynamics — the map the Eastern European reformers lacked — remained, in 1993, largely unclaimed, blocked by arbitrariness, prompting, over-simple learning tasks, and a shortage of empirical traction.

The **credit** column holds what the journey delivered along the way: a way to select among multiple equilibria where the selection does not hinge on the details of the algorithm, a practical way to compute equilibria that resist analysis, new sources of dynamics, and a transfer of computational technique into econometrics itself.

The organizing image of the whole book is the gap between the econometrician and the agents inside the model.

Rational expectations closes that gap by lifting the agents up to a knowledge the econometrician lacks.

Bounded rationality proposed to close it from the other side, by bringing the agents down to the econometrician's level and making them learn.

By 1993 the program had not, on Sargent's own accounting, delivered the transition dynamics that motivated it, and the econometricians it sought to imitate had kept the models at arm's length — for the sound reason that those models add parameters the data identify only weakly, and not consistently, when what the econometrician wants is fewer.

But it had sharpened the questions, selected equilibria, computed them, and lent its tools to the very econometricians who declined its models.

That was the state of the prospects for bounded rationality in macroeconomics, as they stood in 1993.

96.7 Postscript: what became of the program

Three decades is long enough to see which of the 1993 worries were permanent and which were about a field that had not yet found its footing.

96.7.1 The arbitrariness was disciplined

The first reservation was that once we stop insisting agents know the equilibrium, nothing tells us what to put in its place.

The discipline that matured is **expectational stability**.

Its core was already in the 1993 book.

A footnote to the book's chapter 5 writes the limiting dynamics of Bray's belief as the differential equation $d\beta/d\tau = T(\beta) - \beta = a + (b-1)\beta$, where T maps the perceived forecast into the optimal one.

The same footnote credits DeCanio [1979] and Evans [1985], Evans [1989] with using T to define expectational stability, and credits Marcet and Sargent [1989] with showing, by the differential-equation methods of Ljung [1977], that $T(\beta)-\beta$ governs the convergence of least squares learning in a class of models.

Woodford [1990] used the same methods to study learning of sunspot equilibria, and the book's chapter 5 closes with a section on the associated differential equation.

One recursion, many learners sets out this thread for the whole series.

What came after 1993 was systematization.

Evans and Honkapohja [2001] generalized these results and organized them into a theory of learnability that is now the standard test of whether a rational expectations equilibrium is a plausible outcome of adaptive behavior.

For least squares and related decreasing-gain estimators near a rational expectations equilibrium, E-stability is the right condition, and one can check which equilibria satisfy it without simulating anything.

That narrows the arbitrariness: the selection that least squares performs is not an artifact of the particular recursion the modeller happened to write down.

The condition has limits, and those were mapped too.

It is not a theorem about every algorithm.

Stochastic-gradient learners can fail to learn equilibria that are E-stable [Giannitsarou, 2005].

Population algorithms lie outside the theory altogether: drawing on Arifovic's 1991 dissertation, the book reports that a genetic-algorithm population can sometimes learn its way to the equilibrium of Bray's model even when $b > 1$, where least squares cannot (Sargent [1993], chapter 5), and the classifier agents of Economy A2 converge to something that is not an equilibrium at all.

Within those limits the condition does its job, and the stability reversal of *Adaptive Agents in an Overlapping Generations Monetary Economy* is a case in point.

Marcet and Sargent's differential equation already accounted for it in 1993.

In the language of E-stability, it says that $T'(2) = 1/2 < 1$ at the low-inflation steady state and $T'(3) = 2 > 1$ at the high one (*The stability reversal*), and the ranking is a theorem about the model rather than about these parameters (Exercise **??**).

That is why Bruno and Fischer [1990], using adaptive expectations instead of least squares, found the same reversal: rules that move beliefs toward realized inflation inherit the same stability condition.

The verdict does not even need small gains: with the timing of *Adaptive Agents in an Overlapping Generations Monetary Economy*, the low steady state is locally stable for every constant gain below 2 and the high one for none (*The stability reversal*, Exercise 96.8.1).

Bullard and Duffy [1993] supplied the warning against generalizing too far: in economies whose agents live four or more periods, least squares learning can fail to converge locally to a rational expectations equilibrium.

96.7.2 The transition dynamics arrived, in a narrower form than hoped

The prize was a theory of out-of-equilibrium adjustment.

What the program delivered instead was a theory of *departures from* equilibrium: escape dynamics.

In 1993, the escapes of Sims's government were known only from simulation, which the book presented as formalizations of a story and a counter-story about postwar inflation.

The recurrent sawtooth simulated in *A government learning the Phillips curve* comes from the constant-gain version of the model that Sargent [1999] studied after 1993.

Cho *et al.* [2002] characterized its escapes analytically with large-deviations theory, computing the most likely escape path and the rate at which escapes occur, and Williams [2019] extended the characterization considerably.

Sargent and Williams [2005] then explained why the simulated economies of Sims and Chung escape and *stay* near the Ramsey outcome while Sargent's are pulled back again and again: the difference lies in how much variance the government attributes to its regression error.

The QuantEcon lectures *Escaping Nash Inflation* and *Priors, Escapes, and Learning Cycles* work through both.

This is less than the original quest asked for.

It describes recurrent excursions away from a self-confirming equilibrium, not the arrival of a market economy in a country that never had one.

But it is a genuine theory of a system that does not settle down, derived rather than simulated, and in 1993 there was none.

96.7.3 The econometricians did return the compliment

Here the 1993 assessment was largely overtaken.

The obstacle, we saw above, was that the added parameters live in a vanishing transient.

But that argument applies to a learning scheme whose gain goes to zero, so that beliefs converge and then stop moving.

With a gain that does not vanish, beliefs never settle.

Their motion is part of the stationary distribution of the data, and the parameters that govern it are identified the way ordinary structural parameters are, from the whole sample.

The 1993 book already had the exception in view.

It noted that its transients came from algorithms that eventually settle down, adding that "the exceptions are the constant-gain algorithms".

And Chung's government, the one serious empirical case, ran a Kalman filter on coefficients it believed drift: for given drift variances its gain does not go to zero, which is why those variances could be estimated from the whole sample.

Let us check the claim on Bray's market, giving the econometrician the harder problem of re-estimating a and b along with the gain.

We hold the initial belief at its true value; under a constant gain its influence dies out geometrically.

```
def simulate_constant_gain(β0, gain, u):
    "Bray's cobweb when agents discount old prices at a fixed rate."
    T = len(u)
    β, p = np.empty(T), np.empty(T)
    β[0] = β0
    for t in range(T):
        p[t] = a + b * β[t] + u[t]
        if t + 1 < T:
            β[t + 1] = β[t] + gain * (p[t] - β[t])
    return p, β

def belief_path_cg(β0, gain, p):
    "Belief sequence implied by a candidate gain and an observed price path."
    T = len(p)
    β = np.empty(T)
    β[0] = β0
    for t in range(1, T):
        β[t] = β[t - 1] + gain * (p[t - 1] - β[t - 1])
    return β

def ssr_ab_free(p, β):
    "Sum of squared residuals after re-estimating a and b by least squares."
    X = np.column_stack([np.ones_like(β), β])
    coef, *_ = np.linalg.lstsq(X, p, rcond=None)
    resid = p - X @ coef
    return resid @ resid

gain_true, gain_wrong = 0.05, 0.08
print(f"{'T':>6}   {'gain 0.08 penalty':>17}   {'share > 3.84':>12}   "
      f"{'b = 0.75 penalty':>16}")
for T in (500, 5_000, 50_000, 200_000):
    pen_gain, pen_b = [], []
    for seed in range(20):                      # average over 20 samples
        u = np.random.default_rng(seed).standard_normal(T)
        p, β = simulate_constant_gain(β_star, gain_true, u)
        pen_gain.append(ssr_ab_free(p, belief_path_cg(β_star, gain_wrong, p))
                        - ssr_ab_free(p, belief_path_cg(β_star, gain_true, p)))
        pen_b.append(ssr_a_free(p - 0.75 * β) - ssr_a_free(p - b * β))
    pen_gain = np.array(pen_gain)
    print(f"{T:>6}   {pen_gain.mean():17.2f}   {np.mean(pen_gain > 3.84):12.0%}   "
          f"{np.mean(pen_b):16.2f}")
```

```
     T   gain 0.08 penalty   share > 3.84   b = 0.75 penalty
   500               -0.09             0%              0.33
  5000                4.78            60%              1.47
 50000               83.60           100%             10.26
200000              349.95           100%             41.00
```

Both penalties now keep growing with the sample, and in the largest samples they grow in proportion to it: each roughly quadruples between 50,000 and 200,000 observations.

Contrast the slope penalty with a re-estimated under the $1/t$ gain, which never got past a fraction of one.

Persistent belief motion supplies the identifying variation that a vanishing gain withholds.

The information per observation is small, though.

With 500 observations a gain of 0.08 cannot be told from 0.05: the average penalty is about zero and no sample rejects

the wrong gain at the 5 percent level.

With 5,000 observations a majority of samples reject it, and with 50,000 all do.

The gain is identified at the usual rate, but weakly per observation; Exercise 96.8.4 shows how the information shrinks with the size of the gain, and Exercise **??** poses the same question for the Phillips-curve government of *Adaptive Agents in an Overlapping Generations Monetary Economy*.

Later work built on Chung's foundation.

Sargent [1999] made a constant-gain government the center of an account of postwar U.S. inflation, and Sargent and Williams [2005] showed how the government's prior about drifting coefficients, encoded in a drift covariance V, shapes its learning dynamics (*Priors, Escapes, and Learning Cycles*).

Sargent *et al.* [2006] estimated a model in which the government runs a Kalman filter on drifting coefficients, the Sims–Chung specification, using postwar U.S. data to estimate the drift covariance V that sets its non-vanishing gain.

Primiceri [2006] estimated a model in which policy makers learn about the natural rate of unemployment and the persistence of inflation.

The estimates revived a version of the 1993 worry.

To fit the data, SWZ needed a large V, and critics, including Sargent [2008] himself, objected that the implied beliefs were implausibly volatile (*Priors, Escapes, and Learning Cycles*).

The response was to bring in direct evidence on beliefs.

Survey data. Carboni and Ellison [2009] required the model's beliefs to reproduce the forecasts in the Federal Reserve's Greenbook, which shrank the estimated drift variance of the constant by a factor of about fifty while keeping the account of the rise and fall of inflation.

Cohorts. Malmendier and Nagel [2016] found that people's inflation expectations in survey data weight the inflation they have lived through, with gains that decline with age.

In an economy of overlapping cohorts someone is always young, so the aggregate never finishes its transient, and parameters that would be transient for a single agent become parameters of the stationary distribution (Exercise 96.8.3).

Experiments. Laboratory experiments, which the book counted as a credit, offer a third route: there the experimenter knows the environment and observes choices, and sometimes forecasts, directly, as in Marimon and Sunder [1993] and Arifovic [1996].

The econometricians also adopted the agents' apparatus outright.

Cogley and Sargent [2005], whom *Drifts and Volatilities* follows, fit a vector autoregression whose coefficients drift as random walks, the very specification Sims and Chung had handed their government, and asked whether the Great Inflation was bad policy or bad luck.

So the flattery did run both ways in the end.

The learning technology that made the models estimable was already present in Chung's 1990 model.

What changed after 1993 was that a gain that never vanishes moved from exception to standard specification.

It was adopted largely for reasons of economics, since an agent who suspects that its world is changing should keep discounting old data, and it turned out to pay an econometric dividend.

96.7.4 A second retreat, made differently

The program in this book keeps individual rationality and gives up mutual consistency: agents optimize, but against beliefs they are still estimating.

A parallel literature retreats along a different axis.

In the **robustness** work of Hansen and Sargent [Hansen and Sargent, 2008], agents share an approximating model with the econometrician, a model that can be estimated, but they distrust it.

They optimize against the worst case among nearby models that the data could not easily distinguish from it, and the size of that set can be calibrated with detection-error probabilities (*Robust Permanent Income and Pricing*).

Where bounded rationality changes the agents' beliefs, robustness keeps a model common to agents and econometrician and changes the agents' preferences, toward caution.

The two are complements rather than rivals, and both are answers to the question that opens *A Peculiar Definition of Bounded Rationality*: what do we do about the knowledge that rational expectations imputes?

One answer is that agents should learn what the econometrician is learning.

The other is that they should act well while doubting what they have learned.

A third answer, which the book set aside for lack of space, keeps rationality whole: Bayesian agents learning inside a rational expectations equilibrium, as in *Rational Learning and Rational Expectations*.

96.7.5 The algorithms kept crossing over

The 1993 ledger noted that econometricians had adopted the adaptive algorithms as computational tools even while declining the models.

That traffic increased, and reversed direction again.

Holland's bucket brigade (*The bucket brigade*) — pay part of your reward backward to whatever set you up — is a precursor of *temporal-difference learning*, which became an organizing idea of modern reinforcement learning [Sutton and Barto, 2018]; *Job Search X: Search with Q-Learning* applies temporal-difference learning to a McCall worker.

The resemblance runs deep enough to be useful: *The bucket brigade's discount factor* reads the strengths of the bucket brigade in *Money as a Medium of Exchange among Artificially Intelligent Agents* as values computed with a very low discount factor, which helps explain Economy A2.

The classifier systems of *Money as a Medium of Exchange among Artificially Intelligent Agents* are recognizable, in retrospect, as reinforcement learners with a hand-built function approximator; and the perceptrons of *Networks, Genetic Algorithms, and Classifier Systems* are ancestors of the deep networks of *Introduction to Artificial Neural Networks*, which now serve as the function approximators.

Sargent said in 1993 that he would select the "brains" of his agents from the store of methods built up in statistics, econometrics, and artificial intelligence.

Those methods turned out to be early versions of what that last field went on to build.

96.7.6 Reading the ledger again

The 1993 debits have not all been paid.

The choices remain many, the learning tasks we set our agents remain simple next to the ones real firms solve, and no one has produced the theory of transition dynamics that the Eastern European reforms called for.

But the entries have moved.

Selection's theory, sketched in 1993, was systematized and its limits mapped; non-convergence acquired an analytical characterization instead of a simulation; and the econometricians, offered models whose gains never vanish, together with survey, cohort, and laboratory evidence on beliefs themselves, took them up.

The gap between the econometrician and the agents inside the model is still there.

It is narrower, and we now know a good deal about its width.

The next lecture, *The Rise and Fall of U.S. Inflation*, opens a suite built on *The Conquest of American Inflation* [Sargent, 1999], in which a government that keeps re-estimating its Phillips curve — the policy maker of Sims and Chung, grown up — is asked to explain the rise and fall of U.S. inflation.

96.8 Exercises

Exercise 96.8.1

This exercise checks the E-stability account of the stability reversal in *Adaptive Agents in an Overlapping Generations Monetary Economy* numerically, and extends that lecture's linearization (*The stability reversal*), which covers gains between 0 and 1, to find exactly how large the gain can be.

In the constant-deficit economy, a young agent who expects gross inflation β saves $s(\beta) = (w_1 - w_2\beta)/2$, and the government finances a real deficit G per young person by printing money.

If everyone expects the same constant β, realized inflation is constant at

$$T(\beta) = \frac{s(\beta)}{s(\beta) - G}.$$

Use Marimon and Sunder's Economy 7C, $w_1 = 6$, $w_2 = 1$, $G = 1$.

1. Find the fixed points of T and compute T' at each, analytically and numerically. Which steady states are E-stable? Compare with the slopes of the perfect-foresight map $g(\pi) = A_1 - A_2/\pi$, with $A_1 = (w_1 + w_2 - 2G)/w_2$ and $A_2 = w_1/w_2$, at the same points.
2. Replace least squares with Cagan's adaptive expectations, $\beta_{t+1} = \beta_t + \lambda(\pi_t - \beta_t)$, where realized inflation is $\pi_t = s(\beta_{t-1})/(s(\beta_t) - G)$ and $\beta_{-1} = \beta_0$. Simulate for $\lambda \in \{0.1, 0.5, 1, 1.5\}$ from initial beliefs $1.5, 2.5, 2.9, 2.99$ and 3.01, stopping if $s(\beta_t) \leq G$ (no positive price level). Where do the paths go?
3. Linearize the system in part 2 around each steady state and find, analytically, all the values of $\lambda > 0$ for which it is locally stable. Compare with the condition $|1 + \lambda(T' - 1)| < 1$ that a version of the recursion without the lag in realized inflation would give.

Solution

Here $s(\beta) = (6 - \beta)/2$, so $T(\beta) = (6 - \beta)/(4 - \beta)$.

Its fixed points solve $\beta^2 - 5\beta + 6 = 0$, so they are $\beta = 2$ and $\beta = 3$, the two steady states, and $T'(\beta) = 2/(4-\beta)^2$ gives $T'(2) = 1/2$ and $T'(3) = 2$.

```
w1, w2, G = 6.0, 1.0, 1.0

def s(β):
    return (w1 - w2 * β) / 2

def T_map(β):
    return s(β) / (s(β) - G)

A1, A2 = (w1 + w2 - 2 * G) / w2, w1 / w2
steady_states = np.sort(np.roots([1, -A1, A2]).real)

h = 1e-6
for β in steady_states:
    T_prime = (T_map(β + h) - T_map(β - h)) / (2 * h)
    g_prime = A2 / β**2
    print(f"β = {β:.0f}:  T(β) = {T_map(β):.4f}   T'(β) = {T_prime:.4f}   "
          f"g'(β) = {g_prime:.4f}   E-stable: {T_prime < 1}")
```

```
β = 2:  T(β) = 2.0000   T'(β) = 0.5000   g'(β) = 1.5000   E-stable: True
β = 3:  T(β) = 3.0000   T'(β) = 2.0000   g'(β) = 0.6667   E-stable: False
```

Only the low-inflation steady state is E-stable.

Under perfect foresight the ranking is reversed: $g'(2) = 3/2 > 1$ and $g'(3) = 2/3 < 1$.

At both steady states $T' - 1$ and $g' - 1$ have opposite signs, which is the reversal; Exercise **??** proves that this holds for all parameters.

Now the adaptive-expectations version.

```
def adaptive_path(β0, lam, n=2000):
    "Cagan adaptive expectations with the timing of olg_adaptive_money."
    β_prev, β = β0, β0
    for t in range(n):
        if s(β) <= G:
            return β, False                          # no positive price level
        π = s(β_prev) / (s(β) - G)                   # realized inflation
        β_prev, β = β, β + lam * (π - β)
    return β, True

starts = (1.5, 2.5, 2.9, 2.99, 3.01)
print("λ      " + "".join(f"β0 = {β0:<9}" for β0 in starts))
for lam in (0.1, 0.5, 1.0, 1.5):
    cells = []
    for β0 in starts:
        β_end, alive = adaptive_path(β0, lam)
        cells.append(f"{β_end:.4f}" if alive else "collapse")
    print(f"{lam:<6} " + "".join(f"{c:<14}" for c in cells))
```

```
λ      β0 = 1.5      β0 = 2.5      β0 = 2.9      β0 = 2.99     β0 = 3.01
0.1    2.0000        2.0000        2.0000        2.0000        collapse
0.5    2.0000        2.0000        2.0000        2.0000        collapse
1.0    2.0000        2.0000        2.0000        2.0000        collapse
1.5    2.0000        2.0000        2.0000        2.0000        collapse
```

Every path that starts below 3 converges to the low-inflation steady state, and the path that starts just above 3 collapses: the reversal survives Cagan's adaptive expectations, as Bruno and Fischer [1990] found.

For the linearization, write $\delta_t = \beta_t - \beta^\star$.

Realized inflation responds to the current and lagged belief with coefficients $c_0 = -s^\star s'/(s^\star - G)^2$ and $c_1 = s'/(s^\star - G)$, where $s' = -w_2/2$ and $s^\star = s(\beta^\star)$, and $c_0 + c_1 = T'(\beta^\star)$.

So $\delta_{t+1} = (1 + \lambda(c_0 - 1))\,\delta_t + \lambda c_1\,\delta_{t-1}$.

```
def spectral_radius(β_ss, lam):
    "Largest root modulus of the linearized adaptive-expectations system."
    s_ss, s_prime = s(β_ss), -w2 / 2
    c0 = -s_ss * s_prime / (s_ss - G)**2
    c1 = s_prime / (s_ss - G)
    A = np.array([[1 + lam * (c0 - 1), lam * c1],
                  [1.0, 0.0]])
    return np.max(np.abs(np.linalg.eigvals(A)))

for β_ss in steady_states:
    radii = [spectral_radius(β_ss, lam) for lam in (0.5, 1.0, 1.9, 2.0, 2.1)]
    print(f"β* = {β_ss:.0f}: spectral radius at λ = 0.5, 1, 1.9, 2, 2.1:",
          np.round(radii, 3))
```

```
β* = 2: spectral radius at λ = 0.5, 1, 1.9, 2, 2.1: [0.5   0.707 0.975 1.    1.
↪025]
β* = 3: spectral radius at λ = 0.5, 1, 1.9, 2, 2.1: [1.707 2.618 4.365 4.562 4.
↪759]
```

At $\beta^\star = 2$, $c_0 = 1$ and $c_1 = -1/2$, so the characteristic polynomial is $z^2 - z + \lambda/2$; the product of its roots is $\lambda/2$, and the system is stable exactly when $0 < \lambda < 2$.

At $\beta^\star = 3$, $c_0 = 3$ and $c_1 = -1$, so the polynomial is $z^2 - (1 + 2\lambda)z + \lambda$, which is negative at $z = 1$; there is always a root above one, and the high steady state is unstable for every $\lambda > 0$.

Without the lag the multiplier would be $1 + \lambda(T' - 1)$: $1 - \lambda/2$ at the low steady state, stable for $\lambda < 4$, and $1 + \lambda$ at the high one, never stable.

The lag in realized inflation halves the admissible range of λ but does not touch the ranking, which is fixed by the sign of $T' - 1$.

The admissible range includes every gain in $(0, 1]$, and so the effective gain of about 0.75 at which least squares operates near the low steady state (*The stability reversal*).

For small λ any rule that moves beliefs gradually toward realized inflation inherits the E-stability condition $T' < 1$, which is the logic of *One recursion, many learners*; here the condition survives gains far from small.

Exercise 96.8.2

This exercise asks when initial beliefs can be identified under multiplicity.

Set $a = 0$ and $b = 1$ in Bray's market.

Then every constant belief is self-fulfilling: $T(\beta) = \beta$, so $T' = 1$, the unit eigenvalue of the indeterminate direction in *Exchange Rate Indeterminacy, Learning, and Experiments*.

Keep the gain $1/(t + t_0)$ with $t_0 = 50$.

1. Simulate the learning economy for 20,000 periods from $\beta_0 \in \{5, 10, 15\}$, with five shock sequences each, and report the final beliefs. Do the same for $a = 5$, $b = 0.7$.
2. Make the economy deterministic, like the economies of *Adaptive Agents in an Overlapping Generations Monetary Economy* and *Exchange Rate Indeterminacy, Learning, and Experiments*: agents see $p_t = a + b\beta_t$, and the econometrician, who knows a, b and t_0, observes $y_t = p_t + e_t$ with e_t IID standard normal. Compute the expected penalty for $\beta_0 = 12$ when the truth is $\beta_0 = 10$, for $T = 10^2, 10^3, 10^4, 10^5$, and for $b = 0.3$ and

$b = 0.7$ (with $a = 5$) and $b = 1$ (with $a = 0$). Hint: the expected penalty is $\sum_t d_t^2$, where d_t is the difference between the two implied price paths. Explain the three growth rates.

3. Return to Bray's setting, where the shock u_t enters the market and moves beliefs, and the econometrician observes p_t. Compute the average penalty for $\beta_0 = 12$ against the truth $\beta_0 = 10$ for $b = 0.7$ (with $a = 5$) and $b = 1$ (with $a = 0$), for $T = 10^2, 10^3, 10^4$. Why is it bounded even when $b = 1$, where β_0 shifts the limit permanently?

Solution

We need a version of the economy and of the belief reconstruction that takes (a, b) as arguments.

```
t0_ex = 50

def bray(a_, b_, β0, u):
    "Bray's market with gain 1/(t + t0) and shocks u that move beliefs."
    T = len(u)
    β, p = np.empty(T), np.empty(T)
    β[0] = β0
    for t in range(T):
        p[t] = a_ + b_ * β[t] + u[t]
        if t + 1 < T:
            β[t + 1] = β[t] + (p[t] - β[t]) / (t + 1 + t0_ex)
    return p, β

print("final beliefs after 20,000 periods, five shock sequences each")
for a_, b_ in ((0.0, 1.0), (5.0, 0.7)):
    for β0 in (5.0, 10.0, 15.0):
        finals = [bray(a_, b_, β0, np.random.default_rng(seed).standard_normal(20_
↪000))[1][-1]
                  for seed in range(5)]
        print(f"a = {a_}, b = {b_}, β0 = {β0:4.1f}:", np.round(finals, 2))
```

```
final beliefs after 20,000 periods, five shock sequences each
a = 0.0, b = 1.0, β0 =  5.0: [4.97 4.82 4.93 5.08 5.05]
a = 0.0, b = 1.0, β0 = 10.0: [ 9.97  9.82  9.93 10.08 10.05]
a = 0.0, b = 1.0, β0 = 15.0: [14.97 14.82 14.93 15.08 15.05]
a = 5.0, b = 0.7, β0 =  5.0: [14.72 14.67 14.72 14.76 14.76]
a = 5.0, b = 0.7, β0 = 10.0: [15.55 15.5  15.55 15.59 15.59]
a = 5.0, b = 0.7, β0 = 15.0: [16.38 16.33 16.38 16.42 16.42]
```

With $b = 1$ the limit stays close to where the economy started, and it also differs across shock sequences: the outcome is path dependent, in the sense of the book's footnote on the exchange-rate model, which Exercise **??** demonstrates in that economy.

With $b = 0.7$ every economy heads toward $\beta^\star = 16.67$, slowly, whatever β_0.

For part 2 the economy is deterministic, so the expected penalty can be computed without simulating the measurement errors.

```
def deterministic_prices(a_, b_, β0, T):
    "Prices in Bray's market with no shocks in the market itself."
    β, p = β0, np.empty(T)
    for t in range(T):
        p[t] = a_ + b_ * β
        β = β + (p[t] - β) / (t + 1 + t0_ex)
    return p

Ts = (100, 1_000, 10_000, 100_000)
for a_, b_ in ((5.0, 0.3), (5.0, 0.7), (0.0, 1.0)):
    d = deterministic_prices(a_, b_, 12.0, Ts[-1]) - deterministic_prices(a_, b_,␣
↪10.0, Ts[-1])
    expected = [np.sum(d[:T] ** 2) for T in Ts]
    print(f"b = {b_}: expected penalty at T = 1e2, 1e3, 1e4, 1e5:", np.
↪round(expected, 1))
```

```
b = 0.3: expected penalty at T = 1e2, 1e3, 1e4, 1e5: [16.2 31.9 39.9 43.2]
b = 0.7: expected penalty at T = 1e2, 1e3, 1e4, 1e5: [ 135.9  585.7 1806.7 4900.
↪5]
b = 1.0: expected penalty at T = 1e2, 1e3, 1e4, 1e5: [   400.   4000.  40000.␣
↪400000.]
```

In the deterministic economy, the effect of β_0 on the price at t decays like $(t_0/(t+t_0))^{1-b}$, so the expected penalty is roughly a sum of terms of order $t^{-2(1-b)}$.

With $b = 0.3$ that sum converges, and the penalty levels off: the effect dies too fast to be measured well.

With $b = 0.7$ the effect dies slowly, and the penalty keeps growing, but far more slowly than T (eventually like $T^{2b-1} = T^{0.4}$).

With $b = 1$ the effect never dies: the price stays at β_0 forever, and the penalty is exactly $(12-10)^2 T = 4T$.

Here the multiplicity makes the initial belief permanent, and the data identify it at the usual rate, as they would a structural parameter.

Now part 3, with shocks that move beliefs.

```
def reconstructed_beliefs(β0, p):
    "Beliefs the econometrician reconstructs from observed prices and a candidate␣
↪β0."
    T = len(p)
    β = np.empty(T)
    β[0] = β0
    for t in range(1, T):
        β[t] = β[t - 1] + (p[t - 1] - β[t - 1]) / (t + t0_ex)
    return β

for a_, b_ in ((5.0, 0.7), (0.0, 1.0)):
    rows = []
    for T in (100, 1_000, 10_000):
        pens = []
        for seed in range(20):
            p, _ = bray(a_, b_, 10.0, np.random.default_rng(seed).standard_
↪normal(T))
            resid_wrong = p - a_ - b_ * reconstructed_beliefs(12.0, p)
            resid_true = p - a_ - b_ * reconstructed_beliefs(10.0, p)
            pens.append(np.sum(resid_wrong**2) - np.sum(resid_true**2))
        rows.append(np.mean(pens))
    print(f"b = {b_}: average penalty at T = 1e2, 1e3, 1e4:", np.round(rows, 1))
```

```
b = 0.7: average penalty at T = 1e2, 1e3, 1e4: [ 70.2 101.1 103.8]
b = 1.0: average penalty at T = 1e2, 1e3, 1e4: [140.8 202.1 208.6]
```

The penalty is bounded in both cases.

With gain $1/(t+t_0)$ the reconstructed belief is $\beta_t = (t_0\beta_0 + \sum_{s<t} p_s)/(t+t_0)$, so a wrong β_0 shifts it by $(\beta_0' - \beta_0)\, t_0/(t+t_0)$, whatever b is.

The expected penalty is therefore about $b^2(\beta_0' - \beta_0)^2 t_0$, which is close to $0.49 \times 4 \times 50 = 98$ for $b = 0.7$ and $4 \times 50 = 200$ for $b = 1$.

When $b = 1$, β_0 does shift the limit permanently, but so do the early shocks, and the observed prices tell the econometrician where the economy is without telling him how much of its position is due to the initial belief.

Permanence is necessary for identification at the usual rate, but not sufficient: the initial belief must leave a mark that the econometrician cannot attribute to anything he observes.

Exercise 96.8.3

Malmendier and Nagel [2016] find that people form inflation expectations from the inflation they have lived through, with gains that decline with age.

Put such cohorts into Bray's market with $a = 5$ and $b = 0.7$.

At each date $L = 40$ cohorts are alive; a cohort of age $k = 0, 1, \ldots, L-1$ updates its belief with gain $1/(k+1+t_0)$ with $t_0 = 5$; each newborn cohort starts from the prior $\beta_0 = 10$; and the market expectation is the average belief across living cohorts.

1. Simulate 50,000 periods and compare the average price after the first 1,000 periods with $\beta^\star = a/(1-b)$. Why does the economy never reach the rational expectations equilibrium?
2. An econometrician knows b, β_0 and the cohort structure, and re-estimates a. Compute the penalty for $t_0 = 10$ when the truth is $t_0 = 5$, for $T = 500, 5000, 50000$, averaged over 8 samples. Do the same for a single, infinitely lived learner with gain $1/(t+1+t_0)$, the same β_0 and the same t_0.
3. Explain the difference, and relate it to the postscript.

Solution

We carry the cohorts' beliefs in an array ordered by age; each period the oldest cohort dies and a newborn enters at the prior.

```
L = 40

def cohort_beliefs(β0, t0_c, p):
    "Average belief of the living cohorts implied by an observed price path."
    gains = 1 / (np.arange(L) + 1 + t0_c)
    B = np.full(L, β0)
    β_bar = np.empty(len(p))
    for t in range(len(p)):
        β_bar[t] = B.mean()
        B = B + gains * (p[t] - B)
        B[1:] = B[:-1].copy()                    # everyone ages; the oldest dies
        B[0] = β0                                # a newborn enters at the prior
    return β_bar
```

```
def cohort_economy(β0, t0_c, u):
    "Bray's market when the expectation is the average over overlapping cohorts."
    gains = 1 / (np.arange(L) + 1 + t0_c)
    B = np.full(L, β0)
    p = np.empty(len(u))
    for t in range(len(u)):
        p[t] = a + b * B.mean() + u[t]
        B = B + gains * (p[t] - B)
        B[1:] = B[:-1].copy()
        B[0] = β0
    return p

def single_beliefs(β0, t0_c, p):
    "A single infinitely lived learner with gain 1/(t + 1 + t0)."
    β = np.empty(len(p))
    β[0] = β0
    for t in range(1, len(p)):
        β[t] = β[t - 1] + (p[t - 1] - β[t - 1]) / (t + t0_c)
    return β

def single_economy(β0, t0_c, u):
    p = np.empty(len(u))
    β = β0
    for t in range(len(u)):
        p[t] = a + b * β + u[t]
        β = β + (p[t] - β) / (t + 1 + t0_c)
    return p

def penalty_a_free(p, β_wrong, β_true, Ts):
    "Excess SSR (intercept re-estimated) on the first T observations, for each T."
    r_w, r_t = p - b * β_wrong, p - b * β_true
    return np.array([ssr_a_free(r_w[:T]) - ssr_a_free(r_t[:T]) for T in Ts])

Ts = (500, 5_000, 20_000, 50_000)
pen_cohort, pen_single, mean_price = [], [], []
for seed in range(8):
    u = np.random.default_rng(seed).standard_normal(Ts[-1])
    p = cohort_economy(10.0, 5.0, u)
    mean_price.append(p[1_000:].mean())
    pen_cohort.append(penalty_a_free(p, cohort_beliefs(10.0, 10.0, p),
                                     cohort_beliefs(10.0, 5.0, p), Ts))
    p = single_economy(10.0, 5.0, u)
    pen_single.append(penalty_a_free(p, single_beliefs(10.0, 10.0, p),
                                     single_beliefs(10.0, 5.0, p), Ts))

print(f"cohort economy: average price {np.mean(mean_price):.2f}, β* = {β_star:.2f}
 ↪")
print("penalty for t0 = 10, at T = 500, 5000, 20000, 50000")
print("   cohorts:      ", np.round(np.mean(pen_cohort, axis=0), 2))
print("   single agent: ", np.round(np.mean(pen_single, axis=0), 2))
```

```
cohort economy: average price 14.00, β* = 16.67
penalty for t0 = 10, at T = 500, 5000, 20000, 50000
   cohorts:       [ 1.99  6.11 19.42 41.23]
   single agent:  [1.22 1.57 2.76 3.52]
```

The cohort economy's average price stays well below $\beta^\star$.

Every period a newborn cohort arrives believing $\beta_0 = 10$, and the young weight their own short experience heavily, so the average belief is permanently pulled toward the prior: the transient never ends in the aggregate.

The penalty for a wrong t_0 stays bounded for the single learner, whose t_0 matters only while its belief is young: the reconstructed belief is $(t_0\beta_0 + \sum_{s<t} p_s)/(t + t_0)$, so t_0's influence on it dies out like $1/t$.

In the cohort economy the penalty keeps growing, roughly in proportion to T beyond a few thousand observations, even though re-estimating a absorbs the permanent shift in the mean.

The information comes from the dynamics: t_0 governs how strongly the young respond to recent prices, and there are always young cohorts responding.

This is the identification route that Malmendier and Nagel [2016] took with survey data on the expectations of different cohorts: in an economy of overlapping generations, parameters that would be transient for one agent become parameters of the stationary distribution.

Exercise 96.8.4

The constant-gain demonstration in the postscript showed that a gain is identified at the usual rate, but weakly per observation.

This exercise asks how the information depends on the size of the gain.

Use Bray's market with $a = 5$, $b = 0.7$, initial belief $\beta^\star$, and constant gain g.

1. For $g \in \{0.005, 0.01, 0.02, 0.05, 0.1, 0.2\}$, with a and b re-estimated, compute the penalty for the wrong gain $1.6g$ at $T = 20{,}000$, averaged over 20 samples, and divide it by T. How does the information per observation depend on g?
2. At $T = 200$, roughly the length of the postwar record in quarterly data, in what share of the samples does the penalty exceed 3.84?
3. What do the answers imply for the $1/t$ learner, and for the debate over the large drift covariance estimated by Sargent *et al.* [2006]?

Solution

We reuse `simulate_constant_gain`, `belief_path_cg` and `ssr_ab_free` from the postscript.

```
print(f"{'g':>6}   {'penalty / T, T = 20000':>22}   {'share > 3.84, T = 20000':>23}
↪    "
      f"{'share > 3.84, T = 200':>21}")
for g in (0.005, 0.01, 0.02, 0.05, 0.1, 0.2):
    results = {}
    for T in (200, 20_000):
        pens = []
        for seed in range(20):
            u = np.random.default_rng(seed).standard_normal(T)
            p, β = simulate_constant_gain(β_star, g, u)
            pens.append(ssr_ab_free(p, belief_path_cg(β_star, 1.6 * g, p))
                        - ssr_ab_free(p, β))
        results[T] = np.array(pens)
    big = results[20_000]
    print(f"{g:>6}   {big.mean() / 20_000:22.5f}   {np.mean(big > 3.84):23.0%}   "
          f"{np.mean(results[200] > 3.84):21.0%}")
```

```
    g   penalty / T, T = 20000   share > 3.84, T = 20000   share > 3.84, T =␣
↪200
0.005                  0.00010                       25%                       0
↪%
 0.01                  0.00025                       55%                       0
↪%
 0.02                  0.00052                       70%                       0
↪%
 0.05                  0.00152                       95%                       0
↪%
  0.1                  0.00357                      100%                       0
↪%
  0.2                  0.00839                      100%                      10
↪%
```

The information per observation rises with the gain, roughly in proportion to it or a little faster.

The identifying variation is the motion of beliefs, and for small gains the stationary variance of beliefs around $\beta^\star$ is roughly proportional to the gain.

With 20,000 observations the wrong value of the largest gains is rejected in every sample, while that of the smallest is rejected in only a quarter of them; with 200 observations the data say almost nothing about any of the gains.

As $g \to 0$ the constant-gain learner approaches the $1/t$ learner, whose belief motion dies out and whose learning parameters are not consistently estimable: the two cases are the ends of a continuum.

In a sample as short as the postwar record, the data are informative about learning only when beliefs move a lot.

That is one way to read the large drift covariance that Sargent *et al.* [2006] estimated, and the objection that the implied beliefs were implausibly volatile.

Direct evidence on beliefs, such as the Greenbook forecasts that Carboni and Ellison [2009] used, supplies information that outcome data of that length cannot.

Part XIII

Phillips Curve Tradeoffs

CHAPTER

NINETYSEVEN

THE RISE AND FALL OF U.S. INFLATION

Contents

- *The Rise and Fall of U.S. Inflation*
  - *Overview*
  - *Facts*
  - *The Phillips curve in the data*
  - *Two interpretations*
  - *Ignoring the Lucas Critique*
  - *A premature summary: triumph or vindication?*
  - *Data patterns after 1999*
  - *A note on notation*
  - *Exercises*

97.1 Overview

1993 Prospects for Bounded Rationality in Macroeconomics closed the Bounded Rationality suite by handing on a government that keeps re-estimating its Phillips curve: the policy maker of Sims and Chung, grown up.

This lecture opens a suite that asks that government to explain the history of U.S. inflation.

The suite is based on Thomas Sargent's *The Conquest of American Inflation* [Sargent, 1999], and this first lecture combines chapters 1 and 2 of that book, with a preview of its concluding chapter.

Readers arriving from the Bounded Rationality suite have already met a condensed version of the model.

A government learning the Phillips curve studies a constant-gain government with the calibration of the static model in *Escaping Nash Inflation*, *One recursion, many learners* sets out the stochastic-approximation and E-stability tools that we use to analyze such governments, and *The econometricians did return the compliment* tells how econometricians came to estimate them.

The suite asks a question about post-war U.S. macroeconomic history:

> If we take for granted that inflation is under the control of the Federal Reserve, how can we explain the rise of U.S. inflation into the 1970s and its abrupt fall under Paul Volcker in the early 1980s?

The essay evaluates two interpretations, both based on policy makers' *beliefs* about the Phillips curve.

In both stories, the Federal Reserve learns the natural-rate-of-unemployment theory from a combination of experience and *a priori* reasoning.

The stories differ in how that theory is cast:

- **The triumph of natural-rate theory.** Academic economists discovered the natural-rate hypothesis, taught that any inflation-unemployment tradeoff is temporary, and eventually persuaded policy makers to pursue low inflation.
- **The vindication of econometric policy evaluation.** Policy makers never abandoned the methods that Robert Lucas criticized in his famous Critique.
  - Recurrently re-estimating a Phillips curve and using it to choose a target, they were led toward lower inflation by the *data itself*, an adversely shifting empirical Phillips curve.

This lecture presents the facts that motivate both stories, sketches the two interpretations, and reviews the Lucas Critique that chapter 2 both invokes and modifies.

It then previews, in *The road ahead*, how the remaining twelve lectures build the models, and closes by carrying the book's data through to the present.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from statsmodels.tsa.filters.bk_filter import bkfilter
```

Note

The figures in the next two sections reproduce the ones in chapter 1 of Sargent [1999], which were drawn from data available in the late 1990s.

We download the underlying series directly from FRED and restrict attention to the same historical window.

The section *Data patterns after 1999* then carries the most enlightening of these figures through to the present and asks what the additional quarter-century of data means for the two stories.

97.2 Facts

We begin with the single fact that the whole essay seeks to explain: the hump-shaped path of U.S. inflation since World War II.

For a much longer view of price levels, reaching back centuries and across countries, see Price Level Histories.

We measure inflation as the book does, by the annualized monthly change in the consumer price index (all items, not seasonally adjusted), smoothed with a 13-month centered moving average to remove seasonal and high-frequency noise.

The function below reads a series straight from FRED; we use it for the price index and for the unemployment rate that we need in the next section.

```
def fred_series(series_id):
    "Download a monthly series from FRED as a pandas Series."
    url = f'https://fred.stlouisfed.org/graph/fredgraph.csv?id={series_id}'
    frame = pd.read_csv(url, index_col=0, parse_dates=True, na_values='.')
    return frame.iloc[:, 0].rename(series_id)
```

(continues on next page)

(continued from previous page)

```
cpi_all = fred_series('CPIAUCNS')       # CPI, all items, not seasonally adjusted
u_all = fred_series('LNS14000028')      # unemployment rate, white men 20 and over

book_window = slice('1948', '1999-01')
cpi = cpi_all[book_window]

# annualized monthly inflation, then a 13-month centered moving average
inflation = 1200 * np.log(cpi).diff()
inflation_ma = inflation.rolling(13, center=True).mean()
```

Figure 1.1 plots the smoothed series over the book's window.

```
fig, ax = plt.subplots(figsize=(9, 5))
ax.plot(inflation_ma, lw=1.2)
ax.axhline(0, color='k', lw=0.5)
ax.set_xlabel('year')
ax.set_ylabel('inflation (percent, annualized)')
ax.set_title('Figure 1.1: Monthly inflation, CPI all items, '
             '13-month centered moving average')
plt.show()
```

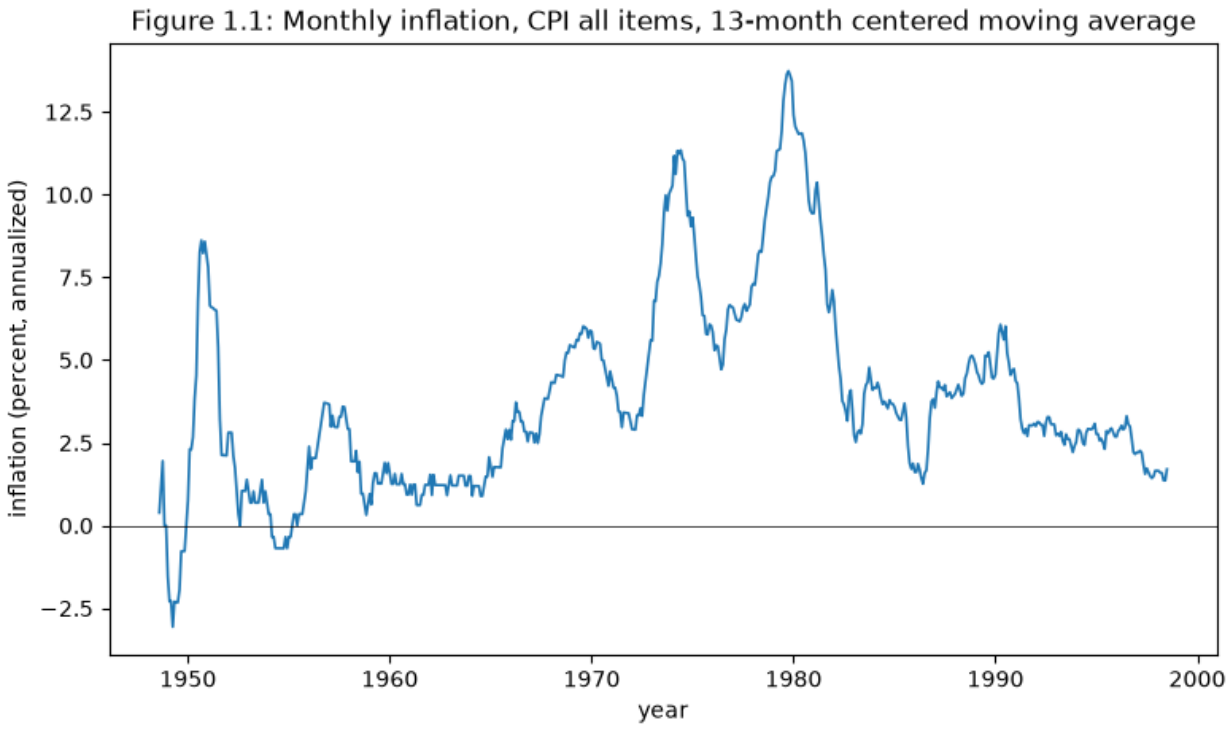


Fig. 97.1: Monthly CPI inflation, 13-month centered moving average, 1948-1999

Inflation was low during the late 1950s and early 1960s, swept upward into the 1970s, and then fell abruptly with Volcker's stabilization in the early 1980s.

Any explanation that treats inflation as under the Federal Reserve's control must account for this rise and fall.

97.3 The Phillips curve in the data

Despite its disrepute in some academic and policy-making circles, the Phillips curve persists in U.S. data, and simple procedures detect it.

To coax the Phillips curve from the data, we follow the book in two ways.

First, we use the unemployment rate for a single demographic group — white men 20 years and over — rather than the aggregate rate, which is contaminated by slow-moving shifts in the demographic mix.

Second, we look at *business-cycle* frequencies, filtering out slowly moving components so that the eye can spot the inverse relationship between inflation and unemployment.

The unemployment series begins in January 1954, so the joint sample does too.

```
u = u_all[book_window]

data = pd.concat([inflation.rename('inflation'),
                  u.rename('unemployment')], axis=1).dropna()
data.head()
```

```
                  inflation  unemployment
observation_date
1954-01-01         0.000000           3.6
1954-02-01         0.000000           3.8
1954-03-01         0.000000           4.1
1954-04-01        -4.469279           4.7
1954-05-01         4.469279           4.6
```

Figure 1.2 plots the two series together, with inflation smoothed as in Figure 1.1.

```
fig, ax = plt.subplots(figsize=(9, 5))
ax.plot(inflation_ma[data.index[0]:], 'C0', lw=1,
        label='inflation (CPI, 13-month centered moving average)')
ax.plot(u, 'C1:', lw=1.2, label='unemployment (white men 20+)')
ax.axhline(0, color='k', lw=0.5)
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.set_title('Figure 1.2: Monthly unemployment and inflation rates')
ax.legend()
plt.show()
```

To isolate the business-cycle relationship, we apply the finite-lag bandpass filter of Baxter and King [Baxter and King, 1999] to the monthly series.

Following the book, we keep fluctuations with periods between 24 and 84 months and use a lead-lag truncation of 84 months.

Because the filter needs 84 months of data on each side of a date, it drops the first and last seven years of the sample, so the filtered series begin in 1961.

```
# Baxter-King bandpass: periods between 24 and 84 months, truncation 84
bk = bkfilter(data, low=24, high=84, K=84)
bk.columns = ['inflation_cycle', 'unemployment_cycle']
print(f"filtered sample: {bk.index[0]:%Y-%m} to {bk.index[-1]:%Y-%m}")
```

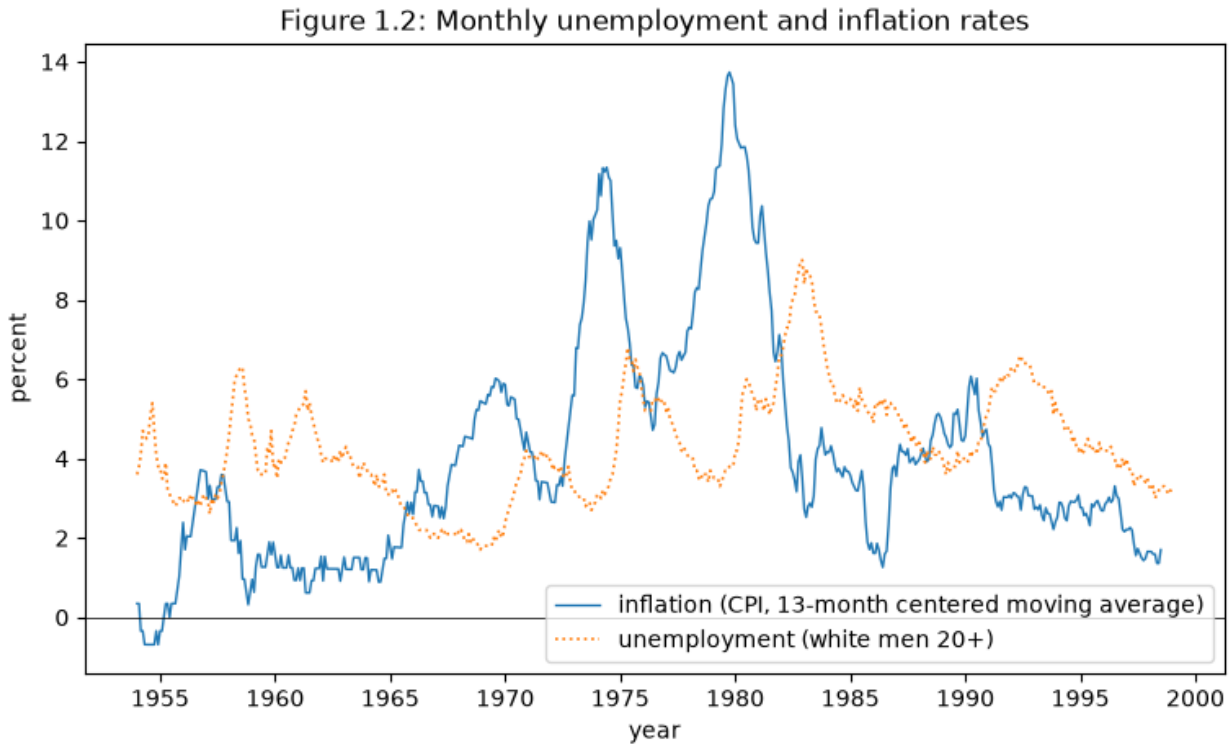


Fig. 97.2: Monthly unemployment (white men 20+) and inflation (13-month centered moving average)

```
filtered sample: 1961-01 to 1992-01
```

Figure 1.3 plots the filtered components.

```
fig, ax = plt.subplots(figsize=(9, 5))
ax.plot(bk.index, bk['inflation_cycle'], 'C0', lw=1, label='inflation')
ax.plot(bk.index, bk['unemployment_cycle'], 'C1:', lw=1.2,
        label='unemployment')
ax.axhline(0, color='k', lw=0.5)
ax.set_xlabel('year')
ax.set_ylabel('deviation from trend (percent)')
ax.set_title('Figure 1.3: Business-cycle components '
             '(Baxter-King bandpass filter)')
ax.legend()
plt.show()
```

The filtered components tend to move in opposite directions: a business-cycle Phillips curve.

We can see the tradeoff more directly in scatter plots for the subperiod that interests us most, 1960-1982.

As in the book, Figure 1.4 plots smoothed inflation against unemployment and Figure 1.5 plots the business-cycle components, with unemployment on the horizontal axis in both.

In Figure 1.5 a faint line joins successive months.

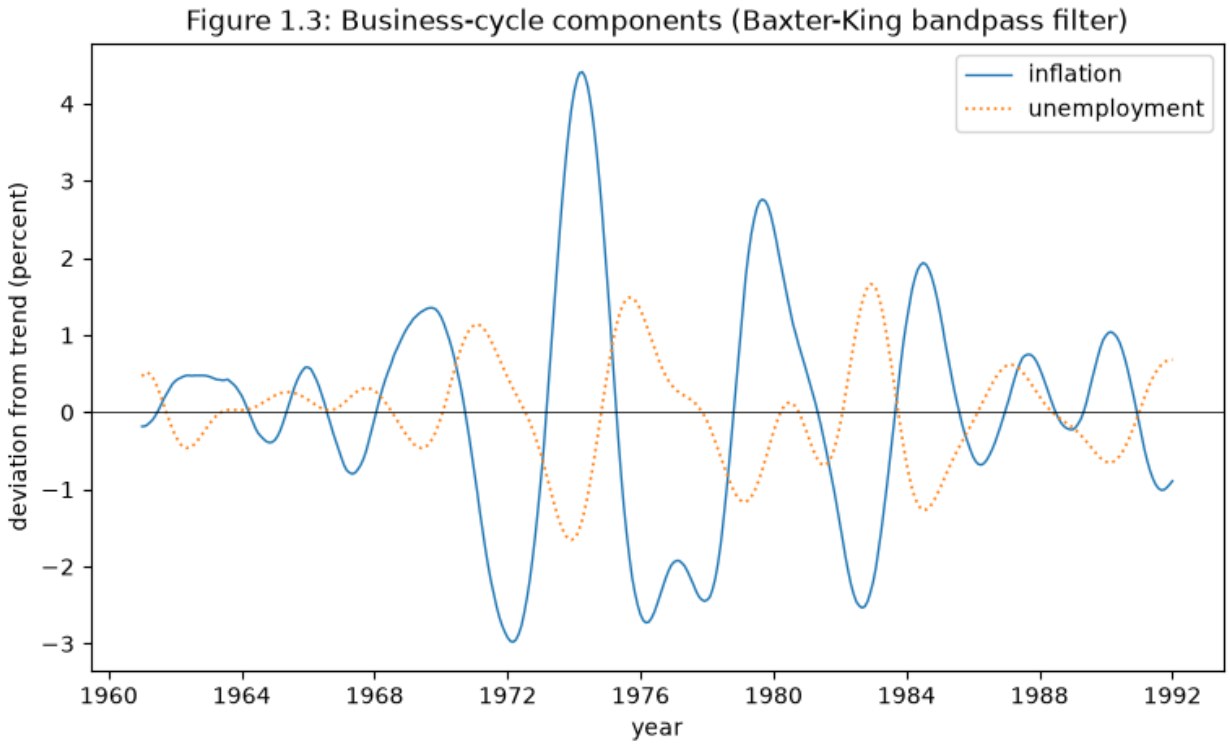


Fig. 97.3: Business-cycle components of inflation and unemployment, Baxter-King bandpass filter

```
sub = slice('1960', '1982')
cycle_sub = bk.loc[sub]

fig, axes = plt.subplots(1, 2, figsize=(12, 5))

axes[0].scatter(u[sub], inflation_ma[sub], s=8, alpha=0.6)
axes[0].set_xlabel('unemployment (white men 20+)')
axes[0].set_ylabel('inflation (13-month centered moving average)')
axes[0].set_title('Figure 1.4: Smoothed series, 1960-1982')

axes[1].plot(cycle_sub['unemployment_cycle'], cycle_sub['inflation_cycle'],
             lw=0.5, alpha=0.4)
axes[1].scatter(cycle_sub['unemployment_cycle'],
                cycle_sub['inflation_cycle'], s=8, alpha=0.6)
axes[1].set_xlabel('unemployment (business-cycle component)')
axes[1].set_ylabel('inflation (business-cycle component)')
axes[1].set_title('Figure 1.5: Business-cycle components, '
                  f'{cycle_sub.index[0].year}-{cycle_sub.index[-1].year}')

plt.tight_layout()
plt.show()
```

Focusing on the business-cycle components sharpens the apparent Phillips curve.

Figure 1.5 reveals **Phillips loops**: followed month by month, the business-cycle components trace out *clockwise* loops rather than a single stable curve.

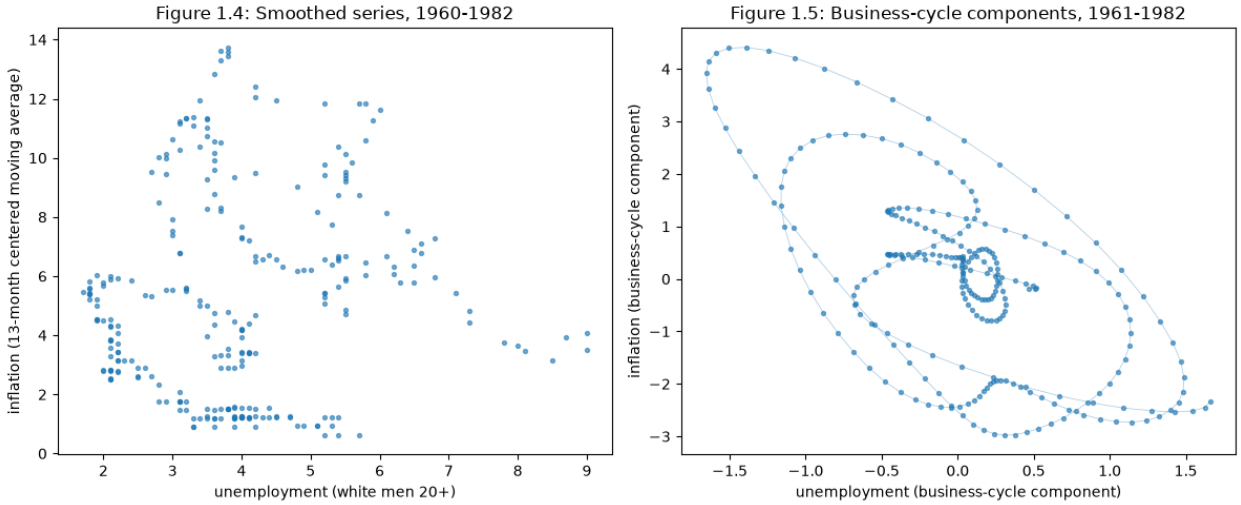


Fig. 97.4: Inflation against unemployment, 1960-1982: smoothed series and business-cycle components

Clockwise loops mean that inflation *lags* unemployment.

As unemployment falls, inflation is slow to rise, so the path runs leftward along the bottom of a loop; as unemployment rises again, inflation is slow to fall, so the path runs rightward along the top.

At these frequencies the negative correlation between inflation and unemployment is strongest when inflation is paired with unemployment about six months earlier (Exercise 97.9.1 asks you to verify this).

Slowly adjusting expectations of inflation, which the natural-rate theory places at the center of the story, are one possible source of that lag.

Note

The book adjusts for demographic change by choosing a single unemployment series, and in chapter 9 it uses these same data to estimate an adaptive model of policy making (*Econometric Policy Evaluation*).

A broader definition of unemployment would inject additional low-frequency demographic components, which one might model with a unit-root process.

The essay instead puts a unit root into the inflation-unemployment process from a different source: the *drifting beliefs* of a monetary authority cut loose from the discipline of Bretton Woods.

97.4 Two interpretations

Both stories start from the same initial conditions — the history of inflation and unemployment, and the state of expectations, inherited by policy makers around 1960 — and both assume that the data conform to the natural-rate hypothesis, whether or not policy makers realized it.

97.4.1 The triumph of natural-rate theory

Adherence to the gold standard and then to Bretton Woods gave the U.S. low inflation and low expected inflation.

In 1960, Paul Samuelson and Robert Solow [Samuelson and Solow, 1960] found a Phillips curve in U.S. data and taught that it was *exploitable* — that policy could raise inflation to reduce unemployment.

That is the lesson of their page 192; on page 193 they themselves warned that "it would be wrong, though, to think that our Figure 2 menu that relates obtainable price and unemployment behavior will maintain its same shape in the longer run."

Within a decade the recommendation was widely endorsed and implemented.

To everyone's dismay, the Phillips curve then shifted adversely: inflation rose, but unemployment on average did not fall.

Meanwhile Edmund Phelps [Phelps, 1967], Milton Friedman [Friedman, 1968], and Robert Lucas [Lucas, 1972] created and refined the concept of the natural rate of unemployment, which assigns a central role to expectations of inflation in locating the Phillips curve.

The natural-rate theory allowed only a *temporary* tradeoff and explained the adverse shifts; its rational expectations version implied that policy makers should ignore the temporary tradeoff and strive only for low inflation.

In this story, these ideas diffused from academics to policy makers and ultimately produced the lower inflation of the 1980s and 1990s.

Events were shaped by policy makers' beliefs — some false, others true — and the actions those beliefs inspired.

Lucas [Lucas, 1976] then used Samuelson and Solow's method of deducing policy from a statistical Phillips curve as an example of erroneous methodology.

His reasoning soon banished the method from academic research centers, but, as the book puts it, "the method survived and prospered within the Federal Reserve System" — a fact that inspires the second story.

97.4.2 The vindication of econometric policy evaluation

The alternative story ascribes Volcker's conquest partly to the *success* of the very econometric and policy-making procedures that Lucas challenged.

It had been told before, by Christopher Sims [Sims, 1988] and Heetaik Chung [Chung, 1990].

Policy makers accepted the Samuelson-Solow Phillips curve as an exploitable tradeoff and adopted their methods for learning from data and deducing policy.

Recurrently, they re-estimated a distributed-lag Phillips curve and used it to reset a target inflation-unemployment pair.

Interpreted mechanically — without identifying expectations as the hidden state variable — the adversely shifting empirical Phillips curve eventually led policy makers to pursue lower inflation.

This story is told with an *adaptive* theory of policy that departs minimally from rational expectations.

It connects a sequence of ideas that the rest of the suite develops: drifting coefficients, self-confirming equilibria, least squares and other recursive learning algorithms, convergence of least squares learners to self-confirming equilibria, and recurrent dynamics along *escape routes* from those equilibria.

The key idea, taken from Sims [1988], is that an adaptive model lets a government *learn* from its past attempts to exploit the Phillips curve, and eventually discover a version of the natural-rate hypothesis that instructs it to reduce inflation.

To Sims's insight the book adds an account of how the *escape-route dynamics* of the adaptive system can spontaneously generate regime shifts.

In this story inflation policy is not made in a vacuum or delivered as a natural experiment; it emerges gradually from an adaptive process.

Although it backs away slightly from rational expectations, the vindication story imposes *more* restrictions on government policy than the triumph story, which occasionally shifts free parameters describing government behavior to fit the facts.

97.5 Ignoring the Lucas Critique

Chapter 2 of the book confronts the obvious objection to the vindication story: doesn't the Lucas Critique forbid the mechanical, exploitable Phillips curve on which it rests?

The essay resurrects the econometric and policy-evaluation procedures that Lucas decisively criticized [Lucas, 1976], and emphasizes a neglected aspect of his Critique: **drifting coefficients**.

97.5.1 The Critique

An econometric model is a collection of stochastic difference equations, some of which describe private agents' decision rules.

Econometric policy evaluation in the Tinbergen-Theil tradition holds those private decision rules *fixed* while the government optimizes its own rule against an objective function.

Lucas noted that if private agents solve intertemporal optimization problems, their decision rules *depend on* the government's rule.

By missing this dependence, the Tinbergen-Theil method mistranslates the government's preferences over outcomes into an ordering over decision rules, and so gives unreliable policy advice.

A rational expectations equilibrium builds the dependence in, by requiring the law of motion that private agents perceive to coincide with the actual law of motion that their own decisions help generate (*Rational Expectations Equilibrium*).

Least Squares Learning in Self-Referential Models shows how agents who estimate their perceived law of motion by least squares can find such a fixed point, and the self-confirming equilibria below are fixed points of the same kind.

97.5.2 The appeal to drifting coefficients

Lucas conceded the impressive forecasting record of Keynesian models but argued that good forecasting is no evidence for the *invariance under intervention* that Tinbergen-Theil assumes.

He stressed that forecasters routinely adjusted the constant terms of key equations, and interpreted these adjustments as an adaptive-coefficients model in the spirit of Cooley and Prescott [Cooley and Prescott, 1973].

The intertemporal instability of estimated relationships — coefficient drift, presumably present even under stable operating rules for policy — undermined treating them as invariant to systematic changes in policy rules.

But Lucas did not read the drift as *confirming* his theoretical point that private decision rules vary with the government's rule.

His three examples compare stationary environments indexed by different government rules, and they are silent about drifting coefficients; the book recalls that many early readers misunderstood his reference to drifting coefficients as supporting his basic point.

Lucas left the drift *unexplained*, and neither the macroeconomic theory nor the rational expectations econometrics built after the Critique accounts for it: both focus on environments with time-invariant transition functions.

97.5.3 Parameter drift as a point of departure

The essay starts from parameter drift, treating it as a *smoking gun* — the key evidence that the government's beliefs about the economy, and hence its policy toward inflation, have evolved over time.

It builds a model from two components:

1. a Tinbergen-Theil theory of government decision making — the Phelps problem of *Adaptive Expectations and the Phelps Problem*; and
2. a drifting-coefficients econometric procedure for the government, featuring the constant adjustments that Lucas wrote about.

It disciplines the exercise by assuming that a rational expectations version of the natural-rate hypothesis is true, and it sets other explanations aside: "I do not examine unusual supply shocks and other plausible explanations of post-World War II U.S. inflation."

If the fundamentals of the environment are time-invariant, the forces behind consistency proofs in econometrics make the government's estimates and its decision rule converge [Marcet and Sargent, 1989].

In the limit there can emerge a **self-confirming equilibrium** (*Self-Confirming Equilibria*), in which the government's estimates are reinforced jointly by its own behavior and the private sector's reactions to it.

Within a self-confirming equilibrium, some of the force of the Lucas Critique vanishes.

Although the government's invariance assumption is wrong, it is not disappointed in outcomes, because those outcomes are statistically consistent with its beliefs.

A self-confirming equilibrium also restricts the government's model and its behavior: because the government's rule is derived from its fitted model, the free parameters that described government policy in Lucas's models disappear.

So a self-confirming equilibrium is a rational expectations equilibrium with *fewer* free parameters than the models of Lucas [1972] and Lucas [1976], and the lost parameters are exactly the ones we would shift to represent a regime change.

Sims [1982] implicitly appealed to a self-confirming equilibrium when he argued that regime changes are rare because existing policies are approximately optimal.

To admit regime changes and drifting coefficients, convergence to a self-confirming equilibrium must be *resisted*.

The essay arrests convergence by replacing the government's least squares estimator with a constant-gain, adaptive-coefficients algorithm like Cooley and Prescott's, which overweights recent data — endowing the government with a suspicion that the environment is unstable.

This weakens the pull toward a self-confirming equilibrium and sustains dynamics along an escape route, along which regime changes occur.

Ironically, as we shall see in *Adaptive Learning and Escape Dynamics*, the procedures that *violate* the Lucas Critique can yield better outcomes than ones that respect it.

97.6 A premature summary: triumph or vindication?

The rest of this suite builds models.

Before we start, it is worth previewing where they lead and the tension they leave unresolved — a premature summary of the journey, adapted from the concluding chapter of Sargent [1999].

Everything can be organized around two *benchmark* models.

In the first, due to Phelps [1967], the public forms expectations *adaptively* while the government chooses policy *optimally*, taking the public's rule as given.

In the second, the rational-expectations natural-rate model, the public is *rational* while the government's policy is treated as *exogenous and arbitrary*.

Lucas recommended replacing the first benchmark with the second.

Coming to grips with the two stories drives us to propose models that make various *compromises* between these poles — and the remaining lectures are those compromises.

97.6.1 The road ahead

We begin by imposing rationality on *both* sides.

- *The Credibility Problem* (chapter 3) studies the one-period model of Kydland and Prescott [1977], whose time-consistent (Nash) outcome is a pessimistic prediction: high inflation, with nothing gained for it.
- *Credible Government Policies* (chapter 4) repeats that economy forever. Computing the whole set of sustainable values with the recursive methods of Abreu *et al.* [1990], it finds so many outcomes sustainable that the theory replaces pessimism with *agnosticism*.

That agnosticism is the first reason to hesitate before declaring the triumph of natural-rate theory.

Following Sims and Chung, we then turn back from the Lucas Critique and start again from the Phelps benchmark.

- *Adaptive Expectations and the Phelps Problem* (chapter 5) solves the Phelps problem against an arbitrary adaptive-expectations Phillips curve and shows why the *induction hypothesis* matters.

Next we make one change: the government's model of the private sector is no longer arbitrary — it is *fit to historical data*.

Varying the details of that fitting problem generates the next five lectures.

- *Optimal Misspecified Beliefs* (chapter 6) introduces equilibria with optimal but misspecified forecasting functions, in Bray's market.
- *Self-Confirming Equilibria* (chapter 7) requires the government's Phillips curve to be self-confirming, and finds that self-confirming equilibria reproduce the Nash outcome or something worse, before closing with an equilibrium in which a subtly misspecified public forecast does better than Nash.

- *Adaptive Learning and Escape Dynamics* (chapter 8) makes the government a real-time econometrician, an "anticipated-utility" learner in the sense of Kreps [1998], and finds spontaneous, Volcker-like stabilizations when the gain is constant.
- *Escaping Nash Inflation* characterizes those escapes analytically, following Cho *et al.* [2002].
- *Priors, Escapes, and Learning Cycles* follows Sargent and Williams [2005] in showing how the government's prior about drifting coefficients shapes convergence, cycles, and escapes.

The next three lectures take the models to data.

- *Econometric Policy Evaluation* (chapter 9) estimates the adaptive model on U.S. data by maximum likelihood, following Chung, and asks whether the estimates vindicate econometric policy evaluation.
- *Drifts and Volatilities* follows Cogley and Sargent [2005] in fitting a vector autoregression with drifting coefficients and stochastic volatility to data through 2000, and asks whether the Great Inflation was bad policy or bad luck.
- *Drifts and Volatilities after 2000* brings that model up to date with the data since 2000.

Finally, *The Lost Conquest: Fed Policy in the 2020s* follows Sargent and Williams [2025] in turning the suite's tools, including an anticipated-utility Federal Reserve, on the inflation of the 2020s and the Federal Reserve's slow response to it, and so returns to the question with which this lecture's survey of the data ends: was the conquest lost?

Symbols change from lecture to lecture, because each lecture follows its source; *A note on notation* records the translations and the gain each lecture uses.

The adaptive models are a *disciplined* retreat from rational expectations, not an abandonment of it.

Apart from a single gain parameter, they carry no free parameters governing expectations, and period by period they impose the same cross-equation restrictions as a rational expectations model.

And because a self-confirming equilibrium is the attractor of their *mean dynamics*, under tranquil conditions, with a gain that eventually falls like $1/t$, they converge back to it, and hence, on the equilibrium path, to rational expectations, satisfying a desideratum of Kreps [1998].

That convergence holds for least squares; *Priors, Escapes, and Learning Cycles* exhibits other priors under which the self-confirming equilibrium is unstable.

But to match the data, our main interest, like Sims's, is in the dynamics that adaptation adds.

Those dynamics come in two kinds.

Sims [1988] and Chung [1990] gave their governments *Kalman filters* on drifting coefficients, whose gains do not vanish; Sims's government fit a static Phillips curve and Chung's a distributed lag.

In their simulations the government escapes from the Nash outcome once and then stays near the Ramsey outcome, apparently indefinitely.

The governments of Sargent [1999] and Cho *et al.* [2002] instead use constant-gain recursive least squares, and their escapes *recur*: each is followed by a return toward Nash, producing the sawtooth of *A government learning the Phillips curve*.

Why the two kinds of history differ was a puzzle until Sargent and Williams [2005] traced it to the government's prior: the drift covariance and the noise variance that its filter assumes.

Priors, Escapes, and Learning Cycles works through their account.

The learning governments of *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation* are Sargent's: suspecting that the Phillips curve is prone to wander, they use a constant-gain algorithm, the sensible choice when coefficients drift.

The payoff is a striking one: the adaptive models produce abrupt *stabilizations* of inflation that defy the inferior self-confirming outcome toward which the mean dynamics point.

These regime shifts arise not from any change in the government's procedures, nor from large shocks, but from *changes in beliefs created by the government's own econometrics* — the mathematics of escape routes in the space of approximating models.

97.6.2 The induction hypothesis: villain and hero

At the center of the story sits the **induction hypothesis** — the restriction that the weights on lagged inflation in an expectations equation sum to one, so that a permanently higher inflation rate is eventually fully expected.

It was built almost without comment into the adaptive expectations hypothesis of Friedman [1957] and Cagan [1956], and it was the basis of Solow's and Tobin's early tests of the natural-rate hypothesis [Solow, 1968, Tobin, 1968].

Cast as a *villain* in Lucas's Critique — a naive restriction that rational expectations does not imply — the induction hypothesis re-emerges as the *hero* of the adaptive models: activating it is exactly what makes the government's Phelps problem call for near-Ramsey, low-inflation policy.

In the distributed-lag model of *Adaptive Learning and Escape Dynamics*, the escape routes are precisely the paths along which the government's estimated model, by using a unit root to approximate a constant, stumbles into believing the induction hypothesis.

Wrestling with such approximation problems, with several models simultaneously in play, is what led Sims [1980] to call bounded rationality a *wilderness*, set apart from the tidy one-model world of rational expectations.

97.6.3 The reservation

Which story, then — triumph or vindication?

The contest is not rational expectations versus an alternative, because *both* stories selectively apply and withdraw from rational expectations.

And the vindication story, however well it fits, is an exercise in *positive* economics, not *normative* economics.

It is tempting to read its long stretches of near-Ramsey inflation as an endorsement of the adaptive policy-making procedures that produce them — but that temptation should be resisted.

For the same adaptive dynamics that permit a stabilization also make it temporary: once beliefs come close to the induction hypothesis, the mean dynamics point *away* from it, back toward the region where the Phelps problem recommends resuscitating inflation.

It can take a long time for the mean dynamics to push the system back to the self-confirming equilibrium, but it is bound to happen.

The simulations contain long episodes that look like Paul Volcker — and others that look like Arthur Burns.

Theoretical work after Kydland and Prescott [1977], such as Rogoff [1985], insists that durable low inflation must rest on *commitment mechanisms* that keep a monetary authority from choosing sequentially — not on the hope that an adaptive government, armed with an approximate model, will by chance eventually learn to do approximately the right thing.

So the book ends on a hope rather than a verdict: we *hope* that the triumph story is the right one — that policy makers have learned a correct rational expectations version of the natural-rate hypothesis and have found devices to commit themselves to low inflation.

Because if instead the vindication story is closer to the truth, then the mean dynamics that pull beliefs back toward the self-confirming equilibrium are always waiting, eventually, to carry inflation back up.

The quarter-century of data since 1999 — a long, quiet *Great Moderation* interrupted by a sudden surge in 2021-2022 — is a running test of exactly this hope, and we turn to it next.

97.7 Data patterns after 1999

Sargent [1999] was written at the end of the long disinflation that began with Volcker.

We now have another quarter-century of data.

This section carries the book's most enlightening figures through to the present and asks what the new observations mean for the two stories.

We use the full FRED series, through the latest month available, and first check each for months that are missing inside the sample.

```
def missing_months(series):
    "Months between the first and last observation that have no value."
    months = pd.date_range(series.first_valid_index(),
                           series.last_valid_index(), freq='MS')
    return months.difference(series.dropna().index)

for series in (cpi_all, u_all):
    gaps = [f'{m:%Y-%m}' for m in missing_months(series)]
    print(f"{series.name:12s} last observation "
          f"{series.last_valid_index():%Y-%m}, missing months: {gaps}")
```

```
CPIAUCNS     last observation 2026-09, missing months: ['2025-10']
LNS14000028  last observation 2026-10, missing months: ['2025-10']
```

The federal government shutdown of late 2025 left the October 2025 observations of both series uncollected.

One missing month would turn every 13-month moving average and every 12-month change that touches it into a missing value, silently cutting the smoothed series short.

We therefore fill isolated missing months by interpolation, geometric for the price index and linear for unemployment, and leave longer gaps alone.

```
def fill_isolated(series, log=False):
    "Interpolate single missing months; leave longer gaps missing."
    s = np.log(series) if log else series
    isolated = s.isna() & s.shift(1).notna() & s.shift(-1).notna()
    s = s.where(~isolated, s.interpolate(limit_area='inside'))
    return np.exp(s) if log else s

cpi_full = fill_isolated(cpi_all['1948':], log=True)
u_full = fill_isolated(u_all)
print("still missing:",
      [f'{m:%Y-%m}' for m in missing_months(cpi_full).union(missing_months(u_full))])

# the book's measure (annualized monthly change, 13-month centered MA)
inflation_full = 1200 * np.log(cpi_full).diff()
inflation_ma_full = inflation_full.rolling(13, center=True).mean()

# year-over-year inflation: smoother, and the measure usually quoted today
inflation_yoy = 100 * (cpi_full / cpi_full.shift(12) - 1)
print(f"13-month centered average available through "
      f"{inflation_ma_full.last_valid_index():%Y-%m}")
```

```
still missing: []
13-month centered average available through 2026-03
```

A centered 13-month average needs six months of data on each side of a date, so the smoothed series ends six months before the latest observation.

97.7.1 The rise and fall, extended

Figure 1.1 showed inflation rising into the 1970s and falling under Volcker.

Extending it to the present adds three chapters the book could not see: the *Great Moderation* of low, stable inflation from the mid-1980s; a decade of low, below-target inflation after the 2008 financial crisis, dipping below zero in 2008-09 and briefly again in 2015; and a sudden surge in 2021-2022 to the highest rate since 1981, followed by a rapid decline.

```
fig, ax = plt.subplots(figsize=(11, 5))
ax.plot(inflation_ma_full, lw=1)
ax.axhline(0, color='k', lw=0.5)
ax.axvspan(pd.Timestamp('1948-01-01'), pd.Timestamp('1999-01-01'),
           color='C0', alpha=0.06, label="the book's window")
for date, y, txt in [('1951-01-01', 9.3, 'Korean\nWar'),
                     ('1970-01-01', 12.0, '1970s\nacceleration'),
                     ('1984-06-01', 0.5, 'Volcker'),
                     ('2009-07-01', -1.5, '2009\ndeflation scare'),
                     ('2022-06-01', 8.9, '2021-22\nsurge')]:
    ax.annotate(txt, (pd.Timestamp(date), y), ha='center', fontsize=9,
                color='C3')
ax.set_xlabel('year')
ax.set_ylabel('inflation (percent, annualized)')
ax.set_title('Inflation, CPI all items, 13-month centered moving average, '
             f'1948-{inflation_ma_full.last_valid_index().year}')
ax.legend(loc='upper right')
plt.show()
```

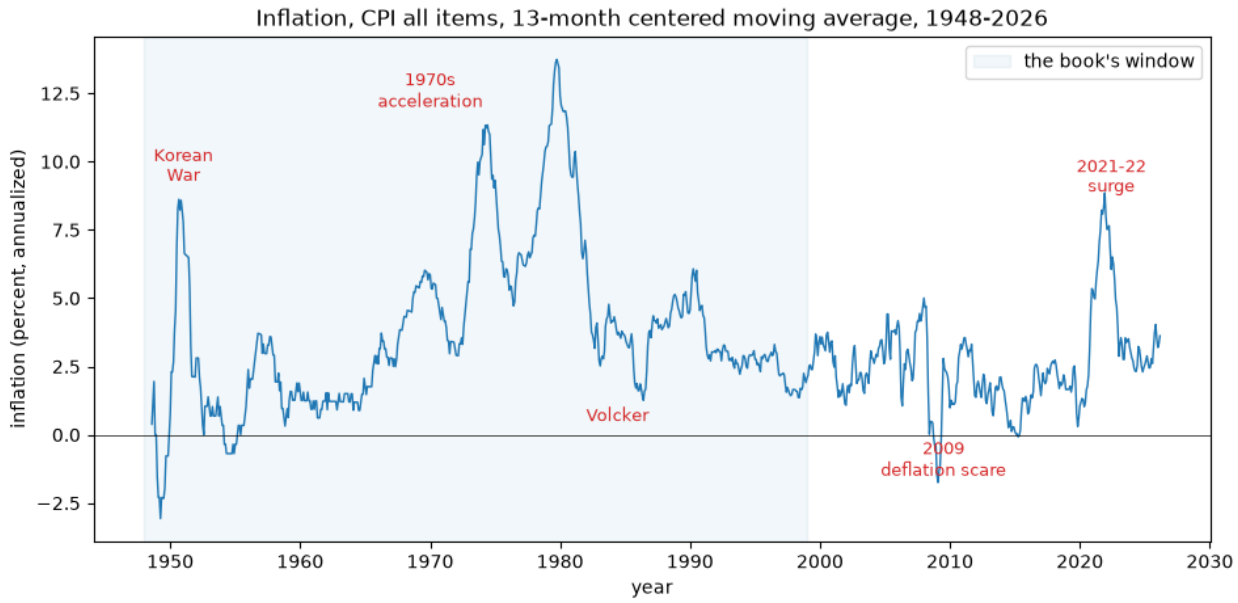


Fig. 97.5: Inflation extended to the present, with the book's window shaded

The hump that the book set out to explain is now one of *two*.

The second compressed a fast acceleration and an almost-as-fast disinflation into about three years.

Its closest precedent is not the 1970s but the Korean War spike of 1950-1951, inside the book's own window: inflation jumped by about eight percentage points within a year and was back near 2 percent within another, with no drawn-out acceleration of the kind that preceded Volcker.

97.7.2 Unemployment and inflation to the present

Figure 1.2 plotted the two series together for the post-war period.

Extending it shows the two most dramatic macroeconomic events of the new data: the COVID unemployment spike of 2020 — for white men 20 and over, the highest rate since the series began in 1954 — and the inflation surge that followed.

```
recent = slice('1990', None)

fig, ax = plt.subplots(figsize=(11, 5))
ax.plot(inflation_yoy[recent], 'C0', lw=1, label='inflation (CPI, year-over-year)')
ax.plot(u_full[recent], 'C1:', lw=1.2, label='unemployment (white men 20+)')
ax.axhline(0, color='k', lw=0.5)
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()
```

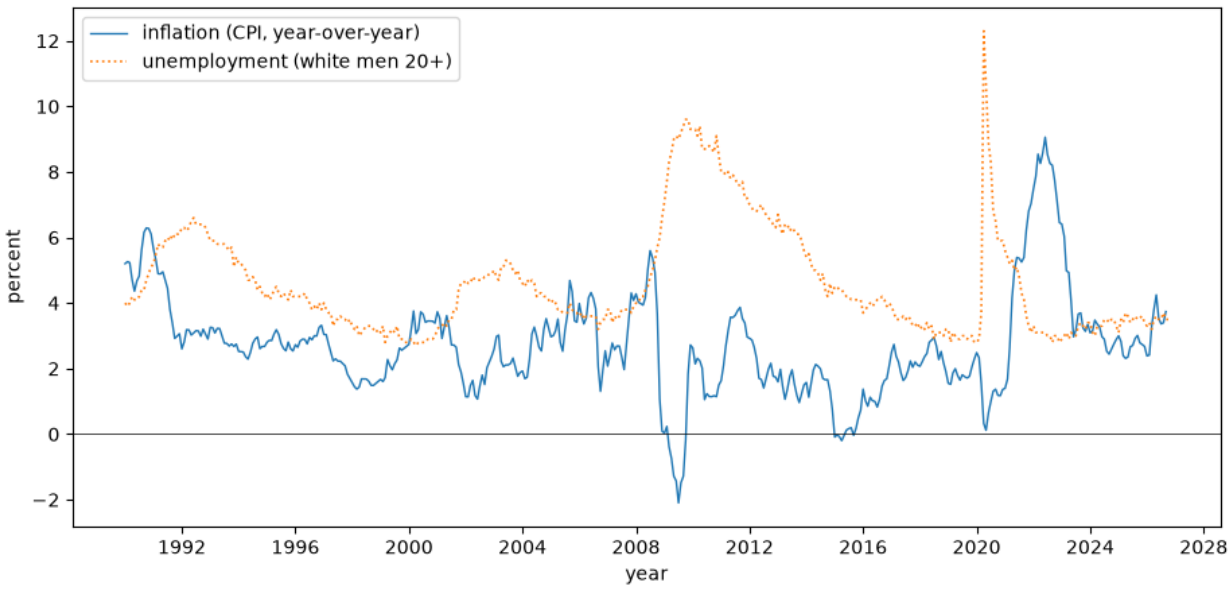


Fig. 97.6: Unemployment and inflation since 1990

Two features stand out.

From the mid-1990s to 2020, inflation stayed remarkably quiet even as unemployment swung widely: it fell in 2000 and again in 2019 to levels last seen in the early 1970s, and roughly tripled in the 2008 recession.

Then, after the COVID spike, unemployment fell back quickly while inflation surged.

Falling unemployment with rising inflation is a move *along* a downward-sloping Phillips curve, the signature of a demand disturbance.

But by itself the pattern cannot separate a demand-driven move along a steep Phillips curve from supply disturbances, and the mix of the two is still debated.

How far has inflation come back down?

```
since_peak = inflation_yoy['2022-07':]
at_or_below_2 = inflation_yoy[inflation_yoy <= 2]
print(f"peak year-over-year inflation: {inflation_yoy['2021':'2023'].max():.1f}% "
      f"({inflation_yoy['2021':'2023'].idxmax():%Y-%m})")
print(f"lowest reading since then:     {since_peak.min():.1f}% "
      f"({since_peak.idxmin():%Y-%m})")
print(f"latest reading:                {inflation_yoy.iloc[-1]:.1f}% "
      f"({inflation_yoy.index[-1]:%Y-%m})")
print(f"last month at or below 2%:     {at_or_below_2.index[-1]:%Y-%m}")
```

```
peak year-over-year inflation: 9.1% (2022-06)
lowest reading since then:     2.3% (2025-04)
latest reading:                3.7% (2026-09)
last month at or below 2%:     2021-02
```

Year-over-year CPI inflation came down quickly from its 2022 peak, but as of this writing it has not been back to 2 percent since early 2021.

97.7.3 The Phillips curve across three eras

The book coaxed a Phillips curve from 1960-1982 data.

The most striking post-1999 pattern is how *unstable* the inflation-unemployment scatter has been across eras.

We split the sample into the book's acceleration era, the Great Moderation, and the post-2008 period, and plot inflation against unemployment in each.

```
scatter_data = pd.concat([inflation_yoy.rename('inflation'),
                          u_full.rename('unemployment')], axis=1).dropna()

last_year = scatter_data.index[-1].year
eras = [('1960', '1983', '1960–1983 (acceleration)'),
        ('1984', '2007', '1984–2007 (Great Moderation)'),
        ('2008', None, f'2008–{last_year} (crisis, COVID, surge)')]

fig, axes = plt.subplots(1, 3, figsize=(14, 4.5), sharex=True, sharey=True)
for ax, (lo, hi, title) in zip(axes, eras):
    era_data = scatter_data.loc[lo:hi]
    ax.scatter(era_data['unemployment'], era_data['inflation'], s=8, alpha=0.5)
    ax.axhline(0, color='k', lw=0.5)
    ax.set_xlabel('unemployment')
    ax.set_title(title, fontsize=10)
axes[0].set_ylabel('inflation (year-over-year)')
plt.tight_layout()
plt.show()
```

Here are the correlation and the least squares slope of inflation on unemployment in each era.

```
for lo, hi, title in eras:
    era_data = scatter_data.loc[lo:hi]
```

(continues on next page)

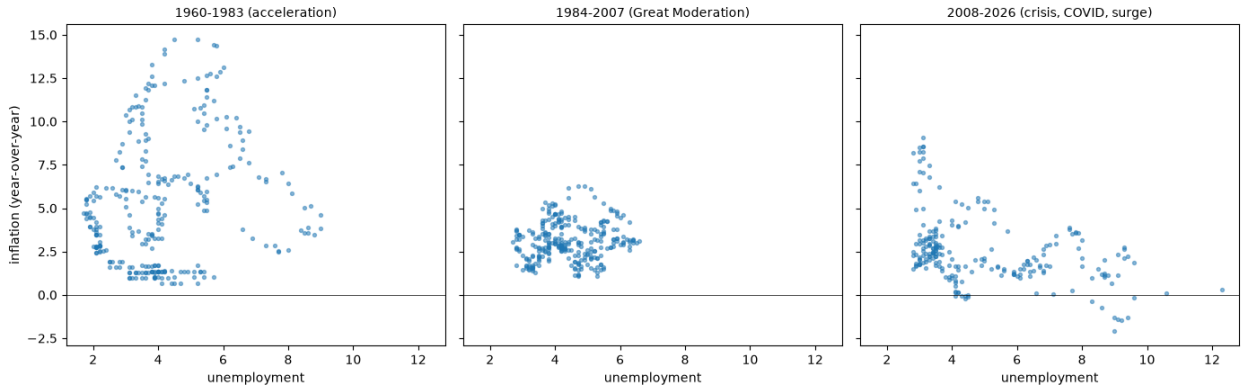


Fig. 97.7: The inflation-unemployment scatter across three eras

(continued from previous page)

```
    corr = era_data.corr().iloc[0, 1]
    slope = np.polyfit(era_data['unemployment'], era_data['inflation'], 1)[0]
    print(f"{title:38s} correlation {corr:+.2f}   slope {slope:+.2f}")
```

```
1960-1983 (acceleration)               correlation +0.19   slope +0.43
1984-2007 (Great Moderation)           correlation +0.11   slope +0.13
2008-2026 (crisis, COVID, surge)       correlation -0.42   slope -0.40
```

In 1960-1983 the points sprawl across a wide range of inflation rates — the era of shifting expectations and Phillips *loops*; in levels the correlation is weakly *positive*, because the adverse shifts swamp the tradeoff.

In 1984-2007 they collapse into a tight, low, nearly flat cloud — the Great Moderation, in which inflation barely responded to unemployment at all.

Since 2008 the scatter is L-shaped: flat at moderate and high unemployment, through the slow recovery from the financial crisis and the COVID recession, and nearly vertical at low unemployment, where the 2021-2022 surge sits.

Paradoxically, it is the only era whose correlation in levels is clearly negative.

An L shape is consistent with a convex Phillips curve, steep only when labor markets are tight, or with supply shocks that happened to arrive when unemployment was low.

Whatever the Phillips curve is, it is not a stable *reduced-form* relationship.

That is exactly what a natural-rate economy with drifting government beliefs predicts: the structural curve, which links unemployment to *surprise* inflation, can stay put while the empirical curve, which links it to inflation itself, drifts with expectations and policy.

Exercise 97.9.3 repeats the comparison at business-cycle frequencies, where the answer is different.

97.7.4 What the new data mean for the two stories

The additional quarter-century neither refutes nor confirms either story cleanly, but it sharpens both.

For the triumph of natural-rate theory. The Great Moderation reads as the triumph completed: with the natural-rate consensus entrenched and central-bank independence established, inflation stayed low and expectations stayed *anchored* for a generation.

In the book's language, policy makers had learned a correct rational expectations version of the natural-rate hypothesis and found devices to commit themselves to low inflation — the outcome the book hoped for.

Even the 2021-2022 surge, on this reading, supports the triumph: once the Federal Reserve moved decisively, inflation came down quickly from its peak and long-run expectations never came unmoored, although inflation has not yet returned to its pre-2021 level.

For the vindication of econometric policy evaluation. The surge is also a reminder that the monetary authority's *model* can still mislead it: the widely held 2021 view that inflation would be "transitory" was a model that the data falsified, and policy adjusted only after beliefs did.

The decade of low, below-target inflation before 2020 — the apparently *flat* Phillips curve, with neither the "missing disinflation" of 2009-2013 nor the "missing inflation" of 2015-2019 fitting a stable curve — is precisely the kind of drifting empirical relationship whose changing slope and intercept the book's adaptive government tracks in real time.

Exercise 97.9.2 asks you to track them as such a government would.

Note

A caveat the book itself insists on, for its own sample as much as for ours: its mechanisms assume that the *fundamentals* — the true data-generating process — are stable, so that all the action comes from the government's evolving beliefs.

As we saw above, the book declines to "examine unusual supply shocks", and that leaves out the oil shocks of 1973 and 1979 just as it leaves out the pandemic disruptions and energy prices of 2021-2022.

Disentangling shifting beliefs from shifting fundamentals is exactly the identification problem that makes this history so hard, and so interesting; *Drifts and Volatilities* takes it on directly.

The tools built in the rest of this suite — self-confirming equilibria, drifting coefficients, and escape dynamics — remain a natural language for asking the question the new data pose: will low inflation keep re-anchoring after each shock, or can a sequence of surprises still set beliefs drifting, as they did after 1965?

Put in the terms of the book's title: has the conquest of American inflation been lost?

Drifts and Volatilities and *Drifts and Volatilities after 2000* ask the data directly whether the coefficients of post-war U.S. macroeconomic dynamics really drifted, before and after 2000.

The closing lecture, *The Lost Conquest: Fed Policy in the 2020s*, turns the suite's tools on the 2021-2022 surge and asks why the Federal Reserve was so slow to respond.

We begin building the models in *The Credibility Problem*.

97.8 A note on notation

The lectures follow the notation of the source each one is based on, and those sources do not agree with one another.

Rather than impose a single scheme and diverge from the papers a reader may want to consult, we keep each lecture faithful to its source and record the translations here.

Readers arriving from the Bounded Rationality suite should note that there γ is the learning gain and the government's beliefs are α (*Notation across the series*, and the dictionary in *A government learning the Phillips curve*); from *Adaptive Expectations and the Phelps Problem* to *Priors, Escapes, and Learning Cycles*, γ is the government's beliefs.

object	symbol	where
inflation	y	*The Credibility Problem* – *Adaptive Learning and Escape Dynamics* (the book's notation)
	π	*Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*, *Drifts and Volatilities*, *Drifts and Volatilities after 2000*, *The Lost Conquest: Fed Policy in the 2020s*; *A government learning the Phillips curve*
public's expected inflation	x	*The Credibility Problem* – *Adaptive Learning and Escape Dynamics*
	$\hat{x}$	*Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*
government's systematic inflation	$h(\gamma)X$	*Self-Confirming Equilibria*, *Adaptive Learning and Escape Dynamics*
	x	*Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*
	g_{t-1}	*A government learning the Phillips curve*
output gap	x	*The Lost Conquest: Fed Policy in the 2020s*
natural rate of unemployment	U^*	*The Credibility Problem* – *Adaptive Learning and Escape Dynamics*; *A government learning the Phillips curve*
	u	*Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*
slope of the true Phillips curve	θ	*The Credibility Problem* – *Escaping Nash Inflation*; *A government learning the Phillips curve*
slope of the perceived Phillips curve	κ	*The Lost Conquest: Fed Policy in the 2020s*
government's beliefs	γ (classical fit), β (Keynesian fit)	*Adaptive Expectations and the Phelps Problem* – *Adaptive Learning and Escape Dynamics*
	$\gamma = (\gamma_1, \gamma_{-1})$, slope first	*Escaping Nash Inflation*
	$\gamma = (a, b)$, intercept first	*Priors, Escapes, and Learning Cycles*; (α_0, α_1) in *A government learning the Phillips curve*
	$\theta_t = (\alpha_{0,t}, \rho_t, \kappa_t)$	*The Lost Conquest: Fed Policy in the 2020s*
econometrician's VAR coefficients	θ_t	*Drifts and Volatilities*, *Drifts and Volatilities after 2000*
learning gain	$1/t$, a_n	*The Credibility Problem* (least squares and its appendix)
	$1 - \lambda$	*Adaptive Expectations and the Phelps Problem* (Cagan–Friedman), *Adaptive Learning and Escape Dynamics* (constant gain $g_t = 1 - \lambda$)
	C	*Optimal Misspecified Beliefs*, *Self-Confirming Equilibria* (a gain chosen to fit the data)
	ε	the primer in *Adaptive Learning and Escape Dynamics*; *Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*
	γ	*The Lost Conquest: Fed Policy in the 2020s*; the Bounded Rationality suite
discount factor	δ	*Credible Government Policies* – *Priors, Escapes, and Learning Cycles*
	β	*The Lost Conquest: Fed Policy in the 2020s*

Econometric Policy Evaluation estimates the model of chapter 9 of the book in the book's notation, that of *Adaptive Learning and Escape Dynamics*: inflation y, beliefs γ and β, forgetting factor λ, and discount factor δ.

Several letters change jobs, and a reader carrying symbols forward will otherwise be caught by them.

- The letter x changes sides: up to *Adaptive Learning and Escape Dynamics* it is what the *public* expects, while in *Escaping Nash Inflation* and *Priors, Escapes, and Learning Cycles* it is what the *government* sets. The two coincide in equilibrium, which is exactly what makes the switch easy to miss. In *The Lost Conquest: Fed Policy in the 2020s*, x is the output gap.
- The letter θ is the slope of the Phillips curve in the early lectures, the government's whole belief vector in *The Lost Conquest: Fed Policy in the 2020s*, and an *econometrician's* coefficients, not a government's beliefs, in *Drifts and Volatilities* and *Drifts and Volatilities after 2000*.
- The letter γ is the government's beliefs from *Adaptive Expectations and the Phelps Problem* to *Priors, Escapes, and Learning Cycles*, but the learning gain in *The Lost Conquest: Fed Policy in the 2020s*, as in the Bounded Rationality suite.
- The letter β is the government's Keynesian Phillips-curve coefficients in *Adaptive Expectations and the Phelps Problem*, *Self-Confirming Equilibria*, and *Adaptive Learning and Escape Dynamics*, the Federal Reserve's discount factor in *The Lost Conquest: Fed Policy in the 2020s* (where β_n is the discount factor of its New Keynesian model of the economy), and a block of covariance parameters in *Drifts and Volatilities*.
- The letter λ works several shifts: the Cagan–Friedman adaptation parameter in *Adaptive Expectations and the Phelps Problem*, a forgetting factor in *Adaptive Learning and Escape Dynamics*, a prior-tightening parameter in *Priors, Escapes, and Learning Cycles*, a block of parameters in *Drifts and Volatilities*, and a measured persistence root in *The Lost Conquest: Fed Policy in the 2020s*, where λ_x is also the weight on the output gap in the Federal Reserve's loss function.

The static model of *Escaping Nash Inflation* and *Priors, Escapes, and Learning Cycles* comes in two orderings, because each lecture follows its source paper.

Escaping Nash Inflation follows Cho *et al.* [2002] and lists the slope of the government's regression first, while *Priors, Escapes, and Learning Cycles* and *A government learning the Phillips curve* follow Sargent and Williams [2005] and list the intercept first.

The constant gains differ across lectures too, again following the sources.

gain	where	source
0.025	*Adaptive Learning and Escape Dynamics*	the book's forgetting factor $\lambda = .975$ (chapter 8)
0.02	*A government learning the Phillips curve*	the condensed model of the Bounded Rationality suite
a range	*Escaping Nash Inflation*	escape times as the gain varies
0.05 and 0.02	*Priors, Escapes, and Learning Cycles*	0.05, the largest gain in the tables of Sargent and Williams [2005], for the analysis; 0.02 for the constant-gain comparison with *A government learning the Phillips curve*
0.015	chapter 9 of the book and *Econometric Policy Evaluation*, quarterly data	the book fixes $\lambda = .985$ and $\delta = .999$ and estimates the rest
0.03	*The Lost Conquest: Fed Policy in the 2020s*	Sargent and Williams [2025], quarterly data

Because the length of a period differs across these settings, the numbers are not directly comparable: a gain g gives the data an effective memory of about $1/g$ periods.

Smaller gains make escapes from the self-confirming equilibrium rarer, but at gains like those of *Adaptive Learning and Escape Dynamics* and *A government learning the Phillips curve* escapes are frequent; *Escaping Nash Inflation* measures how escape times grow as the gain shrinks.

97.9 Exercises

Exercise 97.9.1

The claim that "focusing on the business-cycle components sharpens the apparent Phillips curve" can be made quantitative, and the loops in Figure 1.5 can be explained.

Work with the monthly series in `data` over 1961-1982, the part of the book's window that the filtered series cover.

1. Compute the correlation between inflation and unemployment for the raw monthly series and for the 13-month moving average of inflation plotted in Figure 1.4.
2. Split each raw series into three frequency bands — periods longer than 84 months, periods between 24 and 84 months (the Baxter-King components of Figure 1.5), and periods shorter than 24 months — and compute the correlation within each band. Hint: `bkfilter(data, low=2, high=84, K=84)` keeps all periods shorter than 84 months; subtracting it from the data leaves the low-frequency band, and subtracting the business-cycle band from it leaves the high-frequency band.
3. Which band carries the Phillips curve, and which carries the adverse shifts? Why is the raw correlation close to zero?
4. For the business-cycle components, compute $\mathrm{corr}(\pi_t, U_{t-k})$ for $k = 0, 1, \ldots, 24$ months, where π_t is inflation and U_t unemployment. At what lag is the correlation most negative, and how does that lag explain the direction of the loops in Figure 1.5?

Solution

We split the data into the three bands and compute the correlations.

```
window = slice('1961', '1982')

below_84 = bkfilter(data, low=2, high=84, K=84)    # periods under 84 months
below_84.columns = data.columns
cycle = bk.set_axis(data.columns, axis=1)           # periods of 24-84 months
bands = {'over 84 months': data.loc[below_84.index] - below_84,
         '24 to 84 months': cycle,
         'under 24 months': below_84 - cycle}

smoothed = pd.concat([inflation_ma, u], axis=1).dropna()
raw_var = data.loc[window, 'inflation'].var()

print(f"{'raw monthly series':28s} correlation "
      f"{data.loc[window].corr().iloc[0, 1]:+.2f}")
print(f"{'13-month moving average':28s} correlation "
      f"{smoothed.loc[window].corr().iloc[0, 1]:+.2f}")
for name, band in bands.items():
    b = band.loc[window]
    share = b['inflation'].var() / raw_var
    print(f"{'periods ' + name:28s} correlation {b.corr().iloc[0, 1]:+.2f}"
          f"   share of inflation variance {share:.2f}")
```

```
raw monthly series            correlation +0.04
13-month moving average       correlation +0.11
periods over 84 months        correlation +0.46   share of inflation variance 0.
 ↪45
periods 24 to 84 months       correlation -0.64   share of inflation variance 0.
 ↪14
periods under 24 months       correlation -0.15   share of inflation variance 0.
 ↪40
```

The raw correlation is close to zero because it mixes three bands with correlations of opposite sign.

At periods longer than seven years the correlation is *positive*: these are the adverse shifts of the Phillips curve, the slowly moving component that the natural-rate theory attributes to expected inflation, and the vindication story to the government's drifting beliefs.

At business-cycle periods the correlation is strongly negative: this band carries the Phillips curve.

At periods shorter than two years, which hold a large share of the variance of raw monthly inflation, much of it rounding and seasonal noise in the unadjusted index, the correlation is weak.

Smoothing alone does not help: the 13-month moving average removes most of the high-frequency band but leaves the positive low-frequency shifts in place, so its correlation is, if anything, slightly more positive than the raw one.

The bandpass filter works because it removes both.

Now the lead-lag pattern.

```
lags = np.arange(25)
cross = np.array([cycle['inflation'].corr(cycle['unemployment'].shift(k).
 ↪loc[window])
                  for k in lags])
k_star = lags[np.argmin(cross)]
print(f"corr(π_t, U_t)        = {cross[0]:+.2f}")
print(f"most negative at k = {k_star} months: corr(π_t, U_t-k) = {cross.min():+.2f}
 ↪")
```

```
corr(π_t, U_t)        = -0.64
most negative at k = 6 months: corr(π_t, U_t-k) = -0.81
```

The correlation is most negative when inflation is paired with unemployment about six months earlier, so inflation lags unemployment.

That lag is what makes the loops in Figure 1.5, with unemployment on the horizontal axis, run clockwise.

Start at a peak of unemployment, on the right of a loop.

As unemployment falls, inflation has not yet responded, so the path moves left along the bottom; after unemployment bottoms out, inflation reaches its peak about half a year later, so the path climbs the left side; as unemployment rises, inflation is still high, so the path moves right along the top; and so on around.

Exercise 97.9.2

Lucas observed that forecasters kept adjusting the constant terms of their equations.

Let's watch an econometrician do this to the empirical Phillips curve, regressing unemployment on inflation as the governments of *A government learning the Phillips curve*, *Escaping Nash Inflation*, and *Priors, Escapes, and Learning Cycles* do.

Let π_t be year-over-year CPI inflation (`inflation_yoy`), which uses only past data, and U_t unemployment

(`u_full`).

Estimate

$$U_t = a_t + b_t \pi_t + \text{error}$$

by recursive least squares.

Start from ordinary least squares estimates on 1954-1959 and update monthly from January 1960 to the latest month, using

$$R_t = R_{t-1} + g_t \left(z_t z_t^\top - R_{t-1} \right),$$
$$\begin{bmatrix} a_t \\ b_t \end{bmatrix} = \begin{bmatrix} a_{t-1} \\ b_{t-1} \end{bmatrix} + g_t R_t^{-1} z_t \left(U_t - a_{t-1} - b_{t-1} \pi_t \right),$$

where $z_t = (1, \pi_t)^\top$.

1. Use the constant gain $g_t = 0.02$ of *A government learning the Phillips curve*, which with monthly data gives an effective memory of about four years, and plot a_t and b_t.
2. Repeat with the decreasing gain $g_t = 1/t$ of ordinary least squares, counting t from the start of the 1954-1959 pre-sample, and compare.
3. How did the estimated slope b_t behave in the late 1960s, in the 1970s, and after 2008? What happened to the intercept a_t?
4. The static government of *A government learning the Phillips curve* would set target inflation $-a_t b_t / (1 + b_t^2)$. Compare that target with actual inflation.

Solution

Here is a recursive least squares routine that handles both gains.

```
def rls_phillips(π, U, coef0, R0, gain=None, t0=0):
    """
    Recursive least squares for U_t = a + b π_t, starting from estimates
    coef0 = (a, b) and moment matrix R0 based on t0 observations.  Uses a
    constant gain if `gain` is given and the gain 1/t of ordinary least
    squares otherwise.
    """
    coef, R = coef0.copy(), R0.copy()
    path = np.empty((len(π), 2))
    for t in range(len(π)):
        g = gain if gain is not None else 1 / (t0 + t + 1)
        z = np.array([1.0, π[t]])
        R = R + g * (np.outer(z, z) - R)
        coef = coef + g * np.linalg.solve(R, z * (U[t] - z @ coef))
        path[t] = coef
    return path

rls_data = pd.concat([inflation_yoy.rename('π'), u_full.rename('U')],
                     axis=1).dropna()
pre, est = rls_data.loc[:'1959'], rls_data.loc['1960':]

Z = np.column_stack([np.ones(len(pre)), pre['π']])
coef0 = np.linalg.lstsq(Z, pre['U'], rcond=None)[0]    # OLS on the pre-sample
R0 = Z.T @ Z / len(pre)
```

```
paths = {}
for label, gain in [('constant gain 0.02', 0.02), ('least squares, gain 1/t',␣
 ↪None)]:
    path = rls_phillips(est['π'].values, est['U'].values, coef0, R0,
                        gain=gain, t0=len(pre))
    paths[label] = pd.DataFrame(path, index=est.index, columns=['a', 'b'])
    paths[label]['target'] = (-paths[label]['a'] * paths[label]['b']
                              / (1 + paths[label]['b']**2))
```

We plot the drifting coefficients and the static rule's target.

```
fig, axes = plt.subplots(3, 1, figsize=(10, 9), sharex=True)
for label, p in paths.items():
    axes[0].plot(p['a'], lw=1.2, label=label)
    axes[1].plot(p['b'], lw=1.2, label=label)
    axes[2].plot(p['target'], lw=1.2, label=f'static target, {label}')
axes[2].plot(est['π'], 'k', lw=0.6, alpha=0.6, label='actual inflation')
axes[1].axhline(0, color='k', lw=0.5)
axes[0].set_ylabel('intercept $a_t$')
axes[1].set_ylabel('slope $b_t$')
axes[2].set_ylabel('percent')
axes[2].set_xlabel('year')
for ax in axes:
    ax.legend(fontsize=8, loc='upper left')
plt.tight_layout()
plt.show()
```



A few numbers make the comparison precise.

```
cg, ls = paths['constant gain 0.02'], paths['least squares, gain 1/t']
for lo, hi in [('1960', '1971'), ('1972', '2007'), ('2008', '2015'), ('2016',␣
 ↪None)]:
    b = cg.loc[lo:hi, 'b']
    print(f"steepest constant-gain slope, {lo}-{b.index[-1].year}: "
          f"{b.min():+.2f} ({b.idxmin():%Y-%m})")
print(f"constant-gain slope, 1974-1981: {cg.loc['1974':'1981', 'b'].min():+.2f} "
      f"to {cg.loc['1974':'1981', 'b'].max():+.2f}")
print(f"constant-gain intercept, range: {cg['a'].min():.1f} to {cg['a'].max():.1f}
 ↪")
print(f"least squares, after 1975:  intercept {ls.loc['1975':, 'a'].min():.1f} "
      f"to {ls.loc['1975':, 'a'].max():.1f},  slope {ls.loc['1975':, 'b'].min():+.
 ↪2f} "
      f"to {ls.loc['1975':, 'b'].max():+.2f}")
print(f"constant-gain static target, 1960-1985: "
      f"{cg.loc[:'1985', 'target'].min():.1f} to {cg.loc[:'1985', 'target'].max():.
 ↪1f};"
      f"  actual inflation peak {est.loc[:'1985', 'π'].max():.1f}")
```

```
steepest constant-gain slope, 1960-1971: -0.62 (1967-12)
steepest constant-gain slope, 1972-2007: -0.24 (1983-11)
steepest constant-gain slope, 2008-2015: -0.92 (2011-02)
steepest constant-gain slope, 2016-2026: -0.98 (2021-02)
constant-gain slope, 1974-1981: -0.11 to +0.17
constant-gain intercept, range: 2.7 to 8.1
least squares, after 1975:  intercept 3.4 to 4.6,  slope -0.09 to +0.10
constant-gain static target, 1960-1985: -0.5 to 1.9;  actual inflation peak 14.8
```

With the constant gain, the estimated slope is close to zero in the early 1960s and then turns strongly negative in the late 1960s: as the data then showed it, the tradeoff of Samuelson and Solow looked real and exploitable.

During the 1970s it flattens back to zero and at times turns positive, as inflation rose without bringing unemployment down.

That is the adverse shift, and in the static model of *Escaping Nash Inflation* a slope near zero is exactly the belief that leads the government to aim for zero inflation.

After 2008 the slope turns sharply negative twice, when the unemployment spikes of 2009 and 2020 arrived with falling inflation; these are the steepest slopes of the whole sample.

The intercept drifts too, jumping in the early 1980s, when the Volcker disinflation coincided with high unemployment, and again after 2008 and in 2020: these are Lucas's adjusted constants.

With the decreasing gain the estimates move only slowly after the mid-1970s, because least squares weights an observation from 1955 as heavily as one from last month and so stops tracking drift.

Before 1985 the static government's target never rises much above 2 percent, nowhere near the inflation of the 1970s.

A static, classical government fit to these data would never have recommended the Great Inflation; chapter 9 of the book finds that its own classical model "fails even to match the acceleration in inflation leading up to 1970" and prefers a distributed-lag Keynesian specification (*Econometric Policy Evaluation*).

Constant-gain least squares approximates a Kalman filter for coefficients that drift as random walks, under a particular prior about the drift (*Priors, Escapes, and Learning Cycles*); *Adaptive Learning and Escape Dynamics* puts such a learner inside the model, and *Drifts and Volatilities* lets an econometrician's VAR coefficients drift in the same way.

Exercise 97.9.3

The three-eras figure compares the eras in levels.

Repeat the comparison at business-cycle frequencies.

1. Apply the Baxter-King filter of Figure 1.3 to raw monthly inflation (`inflation_full`) and unemployment (`u_full`) over the whole sample. Where do the filtered series end, and why?
2. To reach the end of the sample, use the Christiano-Fitzgerald bandpass filter with the same band, `cffilter` in `statsmodels.tsa.filters.cf_filter`, which uses asymmetric weights near the ends of the sample instead of dropping them. Check that the two filters agree where both are available.
3. For 1960-1983, 1984-2007, and 2008 to the end of the sample, compute the correlation between the business-cycle components and the least squares slope of the inflation component on the unemployment component. Was the Great Moderation's Phillips curve flat at business-cycle frequencies? What happened after 2008?

Solution

We filter the full sample both ways.

```
from statsmodels.tsa.filters.cf_filter import cffilter

full = pd.concat([inflation_full.rename('inflation'),
                  u_full.rename('unemployment')], axis=1).dropna()

bk_full = bkfilter(full, low=24, high=84, K=84).set_axis(full.columns, axis=1)
cf_full = cffilter(full, low=24, high=84, drift=True)[0].set_axis(full.columns,␣
 ↪axis=1)

print(f"data end in {full.index[-1]:%Y-%m}; "
      f"Baxter-King components end in {bk_full.index[-1]:%Y-%m}")
overlap = bk_full.index
for col in full.columns:
    print(f"{col}: correlation of the two filters' components "
          f"{bk_full[col].corr(cf_full.loc[overlap, col]):.2f}")
```

```
data end in 2026-09; Baxter-King components end in 2019-09
inflation: correlation of the two filters' components 0.98
unemployment: correlation of the two filters' components 0.97
```

The Baxter-King filter is a two-sided moving average with 84 months on each side, so it loses the last seven years of data, including the whole 2021-2022 surge.

The Christiano-Fitzgerald components track the Baxter-King ones closely where both exist; near the end of the sample they rest on one-sided information and will be revised as new data arrive.

For comparison with Exercise 97.9.1, we also compute the correlation in the band of periods longer than 84 months, using the same filter.

```
below_84_cf = cffilter(full, low=2, high=84, drift=True)[0].set_axis(full.columns,␣
 ↪axis=1)
low_cf = full - below_84_cf                          # periods over 84 months

eras_bc = [('1960', '1983'), ('1984', '2007'), ('2008', None)]
for lo, hi in eras_bc:
    e = cf_full.loc[lo:hi]
    corr = e.corr().iloc[0, 1]
    slope = np.polyfit(e['unemployment'], e['inflation'], 1)[0]
    corr_low = low_cf.loc[lo:hi].corr().iloc[0, 1]
    print(f"{lo}-{e.index[-1].year}: business cycle: correlation {corr:+.2f}, "
          f"slope {slope:+.2f}, sd of inflation {e['inflation'].std():.2f};"
          f"  over 84 months: correlation {corr_low:+.2f}")
b = bk_full.loc['2008':'2019']
print(f"2008-{b.index[-1].year}, Baxter-King business cycle: correlation "
      f"{b.corr().iloc[0, 1]:+.2f}")
```

```
1960-1983: business cycle: correlation -0.67, slope -1.62, sd of inflation 1.72;
 ↪  over 84 months: correlation +0.31
1984-2007: business cycle: correlation -0.68, slope -1.10, sd of inflation 0.81;
 ↪  over 84 months: correlation +0.00
2008-2026: business cycle: correlation -0.25, slope -0.39, sd of inflation 1.45;
 ↪  over 84 months: correlation -0.38
2008-2019, Baxter-King business cycle: correlation -0.07
```

At business-cycle frequencies the Great Moderation's Phillips curve was about as tight as that of 1960-1983, with a similar negative correlation, but flatter, and its swings were smaller.

What disappeared after 1984 was the low-frequency drift that swamps the tradeoff in levels, whose correlation falls from positive to about zero, not the business-cycle tradeoff itself.

After 2008 the business-cycle correlation weakens sharply, and with the Baxter-King filter over 2008-2019, before the pandemic and the surge, it nearly vanishes.

So the "flat Phillips curve" of the 2010s, unlike the flatness of the Great Moderation, shows up at business-cycle frequencies too.

The ranking is the reverse of the one in levels, where the post-2008 era was the only one with a clearly negative correlation; that correlation comes mainly from movements slower than the business cycle.

Near the end of the sample these low-frequency components rest on one-sided information, so the last of these numbers is the least reliable.

CHAPTER

NINETYEIGHT

THE CREDIBILITY PROBLEM

Contents

- *The Credibility Problem*
  - *Overview*
  - *A one-period economy*
  - *Nash and Ramsey outcomes*
  - *Best response dynamics*
  - *Least squares learning converges to Nash*
  - *More foresight*
  - *Appendix: stochastic approximation*
  - *Exercises*

98.1 Overview

This lecture describes a basic expectational Phillips curve model of the sort studied by Kydland and Prescott [1977] and Barro and Gordon [1983].

It is the first *modeling* lecture in a suite based on chapters of Sargent [1999], following *The Rise and Fall of U.S. Inflation*, which lays out the two stories and reviews the Lucas Critique.

This lecture follows chapter 3 of the book, and *Credible Government Policies* follows chapter 4.

Together the two lectures formalize

- the temptation to inflate that is unleashed by the discovery of a Phillips curve,
- the value of a commitment technology for resisting that temptation, and
- the fragility of reputational mechanisms as substitutes for commitment.

Throughout, rational expectations is the only equilibrium concept that we use.

Alterations in the *timing* of decisions by a government and a private sector induce different economies with distinct outcomes.

A government faces a **credibility problem** whenever it wishes to make decisions sooner than it must.

We shall compare outcomes under two timing protocols:

- In one, the government chooses inflation *before* the private sector sets its expectations, so the government takes into account how its choice shapes those expectations.
- In the other, the government chooses inflation *after* the private sector has set its expectations.

The deterioration in outcomes under the second protocol measures the loss from an inability to commit.

Time Inconsistency of Ramsey Plans studies the same time-inconsistency problem in Calvo's dynamic monetary model.

We also study two *out-of-equilibrium* dynamics that converge to the no-commitment (Nash) outcome:

- best response dynamics, and
- least squares learning.

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

98.2 A one-period economy

Although credibility problems are intrinsically dynamic, it is possible to describe them in a one-period model under different patterns of within-period timing.

This prepares the way for the multi-period analyses in subsequent lectures.

We describe a version of the one-period model of Kydland and Prescott [1977] in the terms used by Nancy Stokey [Stokey, 1989].

Let (U, y, x) be the unemployment rate, the inflation rate, and the public's expectation of the inflation rate, respectively.

The government's one-period payoff is

$$-\frac{1}{2}\left(U^2 + y^2\right). \tag{98.1}$$

Unemployment is determined by an expectations-augmented Phillips curve

$$U = U^* - \theta(y - x), \qquad \theta > 0. \tag{98.2}$$

The equation asserts that unemployment deviates from U^*, the natural rate, only when there is surprise inflation or deflation.

Substituting (98.2) into (98.1) expresses the government's payoff as a function $r(x, y)$:

$$r(x, y) = -\frac{1}{2}\left[\left(U^* - \theta(y - x)\right)^2 + y^2\right]. \tag{98.3}$$

We work with the following objects.

Rational expectations equilibrium: a triple (U, x, y) satisfying (98.2) and $y = x$.

Government (one-period) best response: given the public's expectation x, a decision rule $B(x) = \text{argmax}_y\, r(x, y)$ for setting y.

Nash equilibrium: a pair (x, y) satisfying (i) $x = y$, and (ii) $y = B(x)$.

Ramsey problem: $\max_y r(y, y)$.

The *Ramsey outcome* is the value of y that attains the maximum.

Best response dynamics: the dynamical system $y_t = B(y_{t-1})$, y_0 given.

A rational expectations equilibrium is a (U, x, y) triple that lies on the Phillips curve and for which private agents are not fooled, given x.

Substituting $x = y$ into the Phillips curve (98.2) shows that $U = U^*$ in any rational expectations equilibrium.

This identifies U^* as the natural rate of unemployment.

98.3 Nash and Ramsey outcomes

A Nash equilibrium builds in a best response by the government, taking the state of expectations x as given, together with a response $x = y$ by the market, i.e., rational expectations for a given y.

Maximizing (98.3) with respect to y for a fixed x gives the government's best response function

$$y = B(x) = \frac{\theta}{\theta^2 + 1} U^* + \frac{\theta^2}{\theta^2 + 1} x. \tag{98.4}$$

The Nash equilibrium sets $x = y = B(x)$, which gives

$$y^N = x^N = \theta U^*, \qquad U = U^*.$$

The Ramsey problem instead imposes $x = y$ *before* maximizing, so it maximizes $r(y, y) = -\frac{1}{2}(U^{*2} + y^2)$, which yields the Ramsey outcome

$$y^R = x^R = 0, \qquad U = U^*.$$

Thus $r(x^R, y^R) = -\frac{1}{2} U^{*2}$ while $r(x^N, y^N) = -\frac{1}{2}(1 + \theta^2) U^{*2}$.

Both outcomes deliver the natural rate U^*, but the Nash equilibrium delivers it with positive inflation and hence a strictly lower payoff.

Let's collect these formulas in a class.

```
class CredibilityModel:
    """
    A one-period expectational Phillips curve economy.
    """

    def __init__(self, θ=1.0, U_star=5.0):
        self.θ, self.U_star = θ, U_star

    def phillips(self, y, x):
        "Unemployment implied by inflation y and expected inflation x."
        return self.U_star - self.θ * (y - x)

    def r(self, x, y):
        "Government one-period payoff."
        U = self.phillips(y, x)
        return -0.5 * (U**2 + y**2)

    def B(self, x):
        "Government best response to expected inflation x."
        θ = self.θ
        return θ / (θ**2 + 1) * self.U_star + θ**2 / (θ**2 + 1) * x
```

(continues on next page)

(continued from previous page)

```
    def nash(self):
        "Nash equilibrium inflation (= expected inflation)."
        return self.θ * self.U_star

    def ramsey(self):
        "Ramsey inflation (= expected inflation)."
        return 0.0
```

The default calibration is $\theta = 1$ and $U^* = 5$.

```
cm = CredibilityModel()

y_N, y_R = cm.nash(), cm.ramsey()
print(f"Nash inflation      y^N = {y_N:.2f}")
print(f"Ramsey inflation    y^R = {y_R:.2f}")
print(f"Nash payoff         r   = {cm.r(y_N, y_N):.3f}")
print(f"Ramsey payoff       r   = {cm.r(y_R, y_R):.3f}")
```

```
Nash inflation      y^N = 5.00
Ramsey inflation    y^R = 0.00
Nash payoff         r   = -25.000
Ramsey payoff       r   = -12.500
```

The Nash payoff is worse than the Ramsey payoff.

The government would prefer the Ramsey outcome, but it cannot attain it without a technology for committing to $y = 0$ before the public forms its expectations.

The Nash equilibrium is supported by a timing protocol in which the government decides *after* the private sector sets its expectations.

The Ramsey outcome is associated with a timing protocol in which the government chooses *first*, knowing that in a rational expectations equilibrium the public's expectations will move with its choice because $y = x$.

98.3.1 A picture of the two outcomes

The government's indifference curves are circles centered at the origin in (U, y) space, because the payoff (98.1) depends only on $U^2 + y^2$.

For a given expectation x, the Phillips curve (98.2) is a downward-sloping line in (U, y) space.

The government's best response for y, given x, occurs where an indifference curve is tangent to the Phillips curve indexed by x.

```
fig, ax = plt.subplots(figsize=(7, 6))

U_grid = np.linspace(0, 12, 200)

# a family of Phillips curves indexed by expected inflation x
for x in [0.0, y_N / 2, y_N]:
    # U = U_star - θ (y - x)  =>  y = x + (U_star - U) / θ
    y_line = x + (cm.U_star - U_grid) / cm.θ
    ax.plot(U_grid, y_line, 'C0', lw=1)

# government indifference curves (circles U^2 + y^2 = const)
ξ = np.linspace(0, 2 * np.pi, 200)
```

(continues on next page)

(continued from previous page)

```
for R in [np.hypot(cm.U_star, y_R), np.hypot(cm.U_star, y_N)]:
    if R > 0:
        ax.plot(R * np.cos(ξ), R * np.sin(ξ), 'C1--', lw=1)

ax.plot(cm.U_star, y_N, 'ko')
ax.annotate('Nash', (cm.U_star, y_N), (cm.U_star + 0.4, y_N + 0.4))
ax.plot(cm.U_star, y_R, 'ko')
ax.annotate('Ramsey', (cm.U_star, y_R), (cm.U_star + 0.4, y_R + 0.4))

ax.set_xlim(0, 12)
ax.set_ylim(0, 10)
ax.set_xlabel('unemployment $U$')
ax.set_ylabel('inflation $y$')
plt.show()
```

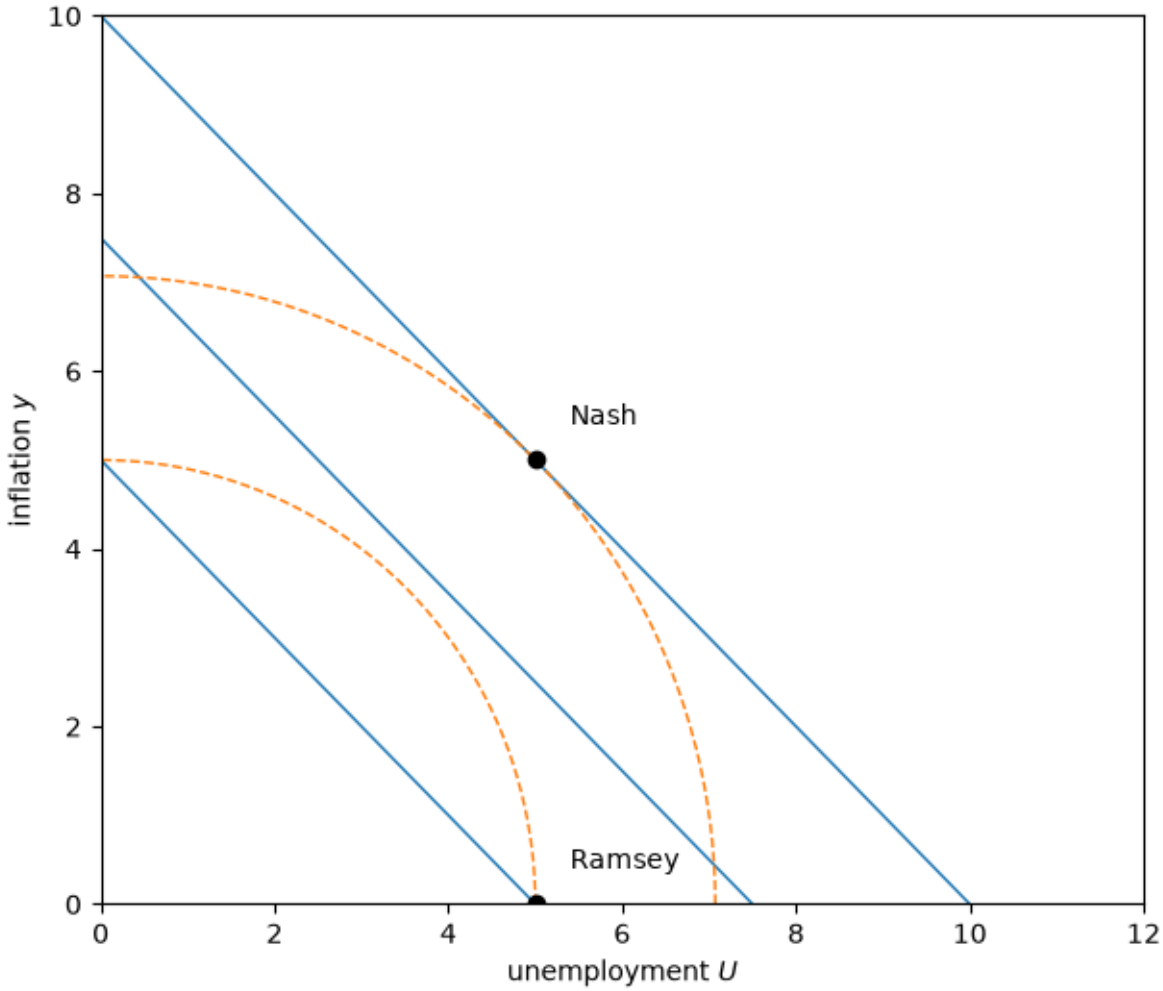


Fig. 98.1: Phillips curves indexed by expected inflation, government indifference curves, and the Nash and Ramsey outcomes

Solid lines are Phillips curves for expected inflation $x \in \{0, y^N/2, y^N\}$; dashed circles are government indifference

curves.

The Nash outcome (U^*, y^N) lies on a larger circle (lower payoff) than the Ramsey outcome $(U^*, 0)$.

98.4 Best response dynamics

Best response dynamics convert the one-period model into a dynamic one by positing an adaptive mechanism in which expected inflation equals last period's inflation, $x_t = y_{t-1}$.

This leads to the dynamics

$$y_t = B(y_{t-1}), \qquad y_0 \text{ given.}$$

Because B is an affine map with slope $\theta^2/(\theta^2+1) \in (0,1)$, iterating on it converges to the fixed point $y^N = \theta U^*$ from any starting point.

Let's plot the best response function against the 45-degree line and simulate the dynamics.

```
def best_response_path(cm, y0, T=20):
    "Iterate y_{t+1} = B(y_t)."
    y = np.empty(T + 1)
    y[0] = y0
    for t in range(T):
        y[t + 1] = cm.B(y[t])
    return y

y_path = best_response_path(cm, y0=0.0, T=20)
```

The figure traces this path as a cobweb between the best response function and the 45-degree line.

```
fig, ax = plt.subplots(figsize=(6, 6))

x_grid = np.linspace(0, y_N + 1, 100)
ax.plot(x_grid, cm.B(x_grid), 'C0', label='$B(x)$')
ax.plot(x_grid, x_grid, 'k--', lw=1, label='45 degrees')

# cobweb of the best response dynamics
for t in range(len(y_path) - 1):
    ax.plot([y_path[t], y_path[t]], [y_path[t], y_path[t + 1]], 'C1', lw=0.8)
    ax.plot([y_path[t], y_path[t + 1]], [y_path[t + 1], y_path[t + 1]],
            'C1', lw=0.8)

ax.plot(y_N, y_N, 'ko')
ax.annotate('Nash', (y_N, y_N), (y_N - 1.5, y_N + 0.3))
ax.set_xlabel('$x$')
ax.set_ylabel('$B(x)$')
ax.legend()
plt.show()
```

Starting from $x = 0$ (the Ramsey inflation rate), the government sets $y = B(0) > 0$.

This provokes the public to raise its expectation, which leads the government to raise inflation further.

The limit of this process is the Nash outcome $y = x = y^N$, a self-confirming situation.

Thus best response dynamics converge to the Nash equilibrium and reinforce (U^*, y^N) as the prediction of the model without a commitment technology.

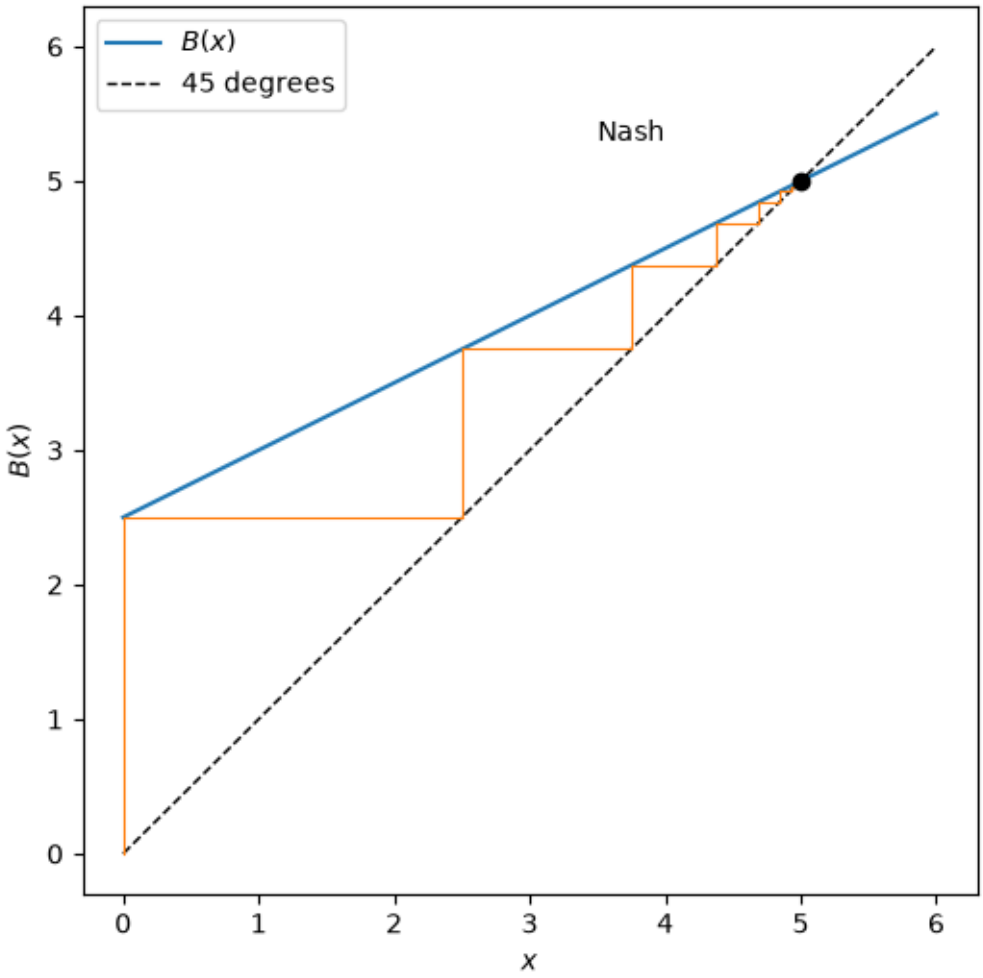


Fig. 98.2: Best response dynamics cobwebbing up to the Nash inflation rate

98.5 Least squares learning converges to Nash

A version of the best response dynamics also emerges from least squares learning.

Least squares learning plays a key role throughout this suite of lectures, and this simple example introduces analytical elements that reappear later in more complex settings.

Bray [1982] studied a competitive market in which traders formed the expected price by averaging past prices.

In the same spirit, assume that for $t \geq 2$ expected inflation x_t is the average of past inflation rates,

$$x_t = \frac{1}{t-1} \sum_{s=1}^{t-1} y_s,$$

which can be represented recursively as

$$x_t = x_{t-1} + \frac{1}{t-1}(y_{t-1} - x_{t-1}), \qquad x_1 = 0. \tag{98.5}$$

Actual inflation is a disturbed version of the best response mapping evaluated at x_t,

$$y_t = B(x_t) + \eta_t, \tag{98.6}$$

where η_t is an IID mean-zero term that represents the government's imperfect control of inflation.

Substituting (98.6), dated $t-1$, into (98.5) gives the stochastic recursion

$$x_t = x_{t-1} + \frac{1}{t-1}\left[B(x_{t-1}) - x_{t-1} + \eta_{t-1}\right]. \tag{98.7}$$

The forecast x_t depends only on shocks dated $t-1$ and earlier, as it must, because the public forms it before it sees y_t.

(The corresponding equation (8) of Sargent [1999] dates this shock t, a slip.)

This is Bray's market of *From relaxation to learning: Bray's market*, written in the notation of this lecture.

- The government's best response B plays the role of the firms' best response h.
- The forecast x_t plays the role of the firms' belief X^*_{t-1}, formed from data through $t-1$.
- Inflation y_t plays the role of realized average output X_t.

The feedback from beliefs to outcomes is

$$h' = B'(x) = \frac{\theta^2}{1+\theta^2} \in (0,1),$$

so Bray's condition $h' < 1$ for least squares learners to converge holds for every θ.

Least Squares Learning in Self-Referential Models analyzes Bray's model (its Example 2) in the general framework of Marcet and Sargent [1989].

By the theory of stochastic approximation, the limiting behavior of x_t is described by the associated ordinary differential equation (ODE)

$$\frac{dx}{dt} = B(x) - x. \tag{98.8}$$

The rest point of this ODE satisfies $x = B(x)$, which is the Nash equilibrium inflation rate $x = \theta U^*$.

Because the map B is affine, the ODE is linear with slope

$$\mathcal{M} = \frac{d}{dx}\left(B(x) - x\right) = B'(x) - 1 = -\frac{1}{\theta^2 + 1}.$$

Since $\mathcal{M} < 0$, the ODE is stable about its rest point.

In the language of *One recursion, many learners*, the Nash equilibrium is **E-stable** (expectationally stable): B maps perceived into actual inflation, and the derivative of that map, $h' = B'$, is less than one.

$\mathcal{M} = h' - 1$ is the E-stability eigenvalue of this one-dimensional system.

E-stability is what brings to bear the theorems of Marcet and Sargent [1989], which give conditions under which x_t converges to y^N globally.

The same number recurs throughout this suite.

When the government itself estimates a Phillips curve by least squares, the mean dynamics of its beliefs have eigenvalues -1 and $-1/(1+\theta^2)$ (*Uniqueness and E-stability*, *Mean dynamics near the self-confirming equilibrium*).

Let's simulate the recursion (98.7) and confirm convergence to the Nash outcome.

```
def ls_learning(cm, T=2000, σ_η=1.0, seed=0, ε=None):
    """
    Simulate learning of expected inflation.

    x[t] is the forecast formed after observing y_1, ..., y_t, which is the
    x_{t+1} of pc_expect1; x[0] = x_1 = 0, and y[0] is unused.

    With ε=None the gain is 1/t (least squares: the 1/(t-1) gain of
    pc_expect1, reindexed); otherwise it is the constant ε.
    """
    rng = np.random.default_rng(seed)
    x = np.empty(T + 1)
    y = np.full(T + 1, np.nan)
    x[0] = 0.0
    for t in range(1, T + 1):
        y[t] = cm.B(x[t - 1]) + σ_η * rng.standard_normal()
        gain = 1.0 / t if ε is None else ε
        x[t] = x[t - 1] + gain * (y[t] - x[t - 1])
    return x, y
```

We simulate 2000 periods at the default calibration.

```
x, y = ls_learning(cm, T=2000, σ_η=1.0)

fig, ax = plt.subplots(figsize=(9, 5))
ax.plot(x, 'C0', lw=1, label='expected inflation $x_t$')
ax.axhline(y_N, color='k', ls='--', lw=1, label='Nash $y^N$')
ax.axhline(y_R, color='C2', ls=':', lw=1, label='Ramsey $y^R$')
ax.set_xlabel('$t$')
ax.set_ylabel('inflation')
ax.legend()
plt.show()
```

The least squares dynamics confirm the pessimism of the best response dynamics.

As Sargent [1999] puts it, this pessimism undermines the *triumph of natural-rate theory* story of *The Rise and Fall of U.S. Inflation*, especially its triumph aspect.

Given an initial condition in the form of a low, gold-standard or Bretton Woods value of x, the best response or least squares dynamics can explain an *acceleration* of inflation.

But they cannot explain a Volcker-style *stabilization* that brings inflation back down.

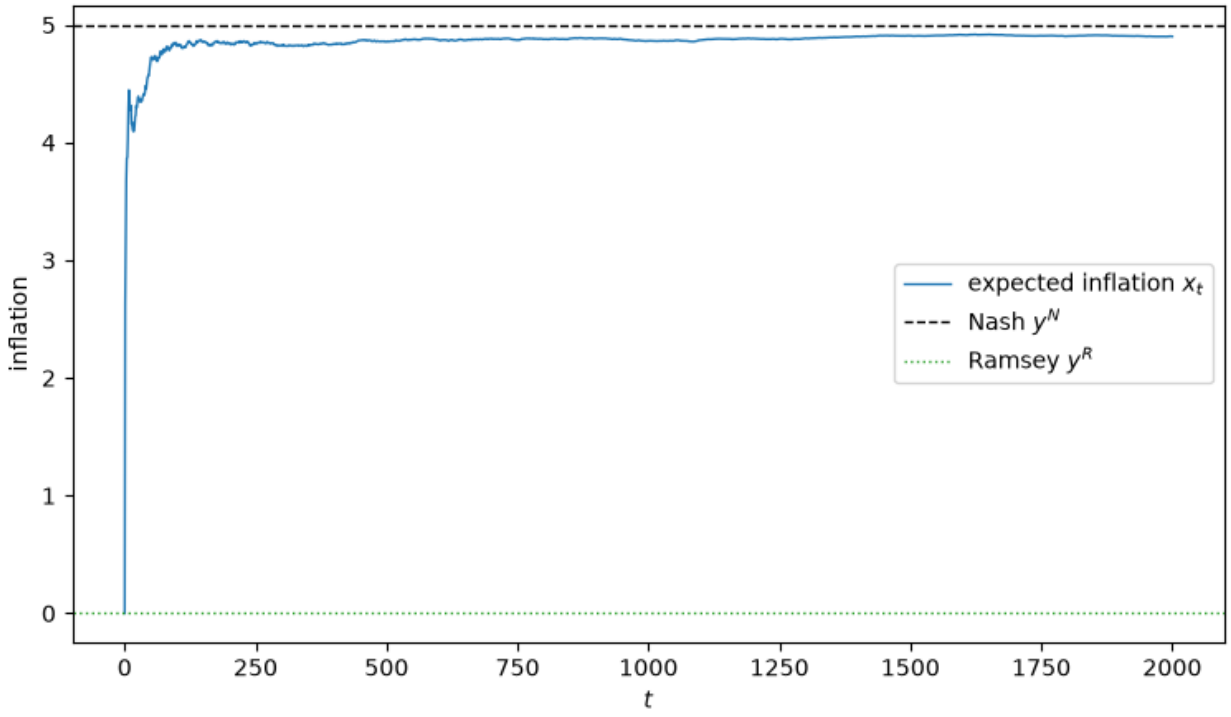


Fig. 98.3: Least squares learning of expected inflation converges to the Nash outcome

Later lectures reformulate versions of least squares learning in ways designed to moderate this pessimism, and in doing so they move us away from the triumph story and toward the vindication story.

98.5.1 How fast?

$\mathcal{M}$ also governs the *speed* of convergence (footnote 3 of chapter 3 of Sargent [1999]; Evans and Honkapohja [2001]).

Taking expectations of (98.7) shows that the mean error $Ex_t - y^N$ obeys the noise-free version of the recursion, so it shrinks like $t^{\mathcal{M}}$.

A necessary condition for convergence at the usual rate, with $\sqrt{t}\,(x_t - y^N)$ settling into a normal distribution, is $\mathcal{M} < -1/2$, which here requires $\theta < 1$.

When $-1/2 \leq \mathcal{M} < 0$, x_t still converges to y^N, but more slowly, at a rate governed by $|\mathcal{M}|$.

At $\theta = 1$, the value used here and in most later lectures, $\mathcal{M} = -1/2$ sits exactly on the boundary: the standard deviation of x_t shrinks like $\sqrt{\log t / t}$ instead of $1/\sqrt{t}$.

In the market of *From relaxation to learning: Bray's market*, by contrast, $h' = -1.5$ gives $\mathcal{M} = -2.5$, and errors shrink at the usual rate $t^{-1/2}$.

Exercise Exercise 98.8.1 measures these rates.

The $-1/2$ boundary concerns the convergence rate of a learner whose gain *decreases* like $1/t$.

It is not what governs how often a *constant*-gain learner escapes from the Nash outcome in later lectures; a large-deviation rate does that (*A primer on recursive algorithms*).

98.6 More foresight

Best response and least squares learning are out-of-equilibrium dynamics tacked onto a one-period economy.

They force all movement through expectation formation: in choosing inflation, the government forgets that the economy lasts more than one period.

Better outcomes can occur if the government plans for the future.

Subsequent lectures describe three ways of modeling foresight, which impute varying amounts of rationality and predict different qualities of outcomes:

1. A reputational approach that attributes rational expectations to both the government and the public, the subject of *Credible Government Policies*, which comes next.
   - Many outcomes are sustainable, ranging from repetition of the Ramsey outcome to paths worse than repetition of the Nash outcome, a multiplicity that turns out to be the theory's chief lesson.
2. An approach that keeps the government rational but gives the public *adaptive* expectations in the original Cagan-Friedman sense, the subject of *Adaptive Expectations and the Phelps Problem*.
   - Depending on a comparison between a discount factor and an adaptation parameter, this setup can improve outcomes and possibly sustain repetition of the Ramsey outcome.
3. Approaches that make the government boundedly rational as well.
   - In *Self-Confirming Equilibria*, prepared by *Optimal Misspecified Beliefs*, the government holds a possibly misspecified model of the Phillips curve that the data it generates confirm.
   - From *Adaptive Learning and Escape Dynamics* on, it re-estimates that model in real time.
   - The outcomes depend on details of the beliefs assigned to the government, and a vindication of econometric policy evaluation can emerge.

98.7 Appendix: stochastic approximation

Here we sketch why the ODE (98.8) governs the tail behavior of the stochastic difference equation (98.7).

The argument, due to Kushner and Clark [1978], has two key components: a shift in time scale and a liberal application of averaging.

The subsection "Stochastic approximation and E-stability" of *One recursion, many learners* presents the same tools from the standpoint of the Bounded Rationality suite, and *Least Squares Learning in Self-Referential Models* states the Marcet–Sargent convergence theorems for vectors of beliefs.

Let $\{a_n\}_{n\geq 0}$ be a positive sequence of gains satisfying

$$\lim_{n\to\infty} a_n = 0, \qquad \sum_n a_n = +\infty, \qquad \sum_n a_n^2 < +\infty.$$

The choice $a_n = 1/(n+1)$ satisfies these assumptions.

Relabel (98.7) so that its x_{n+1} and η_{n+1} become x_n and η_n.

Then the gain is $a_n = 1/(n+1)$ and the recursion reads

$$x_{n+1} = x_n + a_n \left[B(x_n) - x_n + \eta_n \right], \tag{98.9}$$

where η_n is IID with mean zero and finite variance.

Introduce the transformed time scale $t_0 = 0$, $t_n = \sum_{i=0}^{n-1} a_i$, and interpolate the discrete sequence $\{x_n\}$ into a continuous-time process $x^0(t)$.

The map $m(t) = \max\{n : t_n \leq t\}$ takes transformed time back to the original index.

Kushner and Clark show that on this transformed time scale the interpolated process is well approximated by the integral equation

$$x^0(t) = x^0(0) + \int_0^t [B(x^0(s)) - x^0(s)]\, ds + R(t),$$

where the approximation error $R(t)$ has two components — one from approximating a distributed lag in $B(x) - x$ by an integral, and one from a distributed lag in η_s.

Studying a sequence of left-shifted versions of the process, they show that both error components can be driven to zero as $n \to \infty$.

For the noise component, the key step is to note that over the stretch of transformed time from t_n to $t_n + t$ the noise enters through the martingale $W_n(t) = \sum_{i=n}^{m(t_n+t)} a_i \eta_i$.

Its variance is $\mathbb{V}[\eta] \sum_{i=n}^{m(t_n+t)} a_i^2$, a *tail* sum of the convergent series $\sum_i a_i^2$, and so it goes to zero as $n \to \infty$.

The remaining error is sent to zero because $a_i \to 0$ shrinks the mesh of the Riemann sum used to approximate the integral.

In the limit, the stochastic difference equation (98.9) shares the behavior of the non-stochastic ODE

$$\frac{d}{dt}\tilde{x}(t) = B(\tilde{x}(t)) - \tilde{x}(t),$$

which is said to describe the *mean dynamics* of the original system.

Later lectures study systems like (98.9) in which a_i does *not* approach zero as i grows — so-called constant gain algorithms.

Exercise Exercise 98.8.2 gives a first look at one.

The mean dynamics (98.8) and these constant-gain algorithms become the central tools of *Adaptive Learning and Escape Dynamics*, *Escaping Nash Inflation*, and *Priors, Escapes, and Learning Cycles*, where the scalar expectation x studied here grows into a whole vector of drifting Phillips-curve coefficients.

98.8 Exercises

Exercise 98.8.1

This exercise measures how the speed of least squares learning depends on $\mathcal{M} = -1/(1+\theta^2)$.

Let $e_t = x_t - \theta U^*$ be the forecast error relative to Nash.

1. Show that the mean error Ee_t obeys the noise-free version of (98.7), and that it shrinks like $t^{\mathcal{M}}$.
2. For $\theta \in \{0.5, 1, 2\}$ and $U^* = 5$, simulate 200 paths of length $T = 3000$ (seeds $0, \ldots, 199$). On log–log axes, plot $|\bar{e}_t|/(\theta U^*)$, where $\bar{e}_t$ is the average error across paths, and compare it with the exact mean error. Dividing by θU^*, the initial error, makes the three cases comparable.
3. Plot $\sqrt{t}$ times the standard deviation of x_t across paths, and explain why it levels off only for $\theta < 1$.
4. What fraction of the $\theta = 2$ paths is still below the Nash rate at $t = 3000$? Does the $\theta = 2$ learner wander around its limit?

Solution

Write $B(x) - x = \mathcal{M}(x - y^N)$.

Shifted forward one period, (98.7) says $e_{t+1} = (1 + \mathcal{M}/t)\, e_t + \eta_t/t$, with $e_1 = -\theta U^*$, and taking expectations removes the shock.

So the mean error is exactly the path that `ls_learning` produces with $\sigma_\eta = 0$:

$$Ee_{t+1} = e_1 \prod_{s=1}^{t} \left(1 + \frac{\mathcal{M}}{s}\right) \approx \frac{e_1\, t^{\mathcal{M}}}{\Gamma(1 + \mathcal{M})},$$

because $\log \prod_s (1 + \mathcal{M}/s) \approx \mathcal{M} \sum_{s \le t} 1/s \approx \mathcal{M} \log t$.

The variance v_t of e_t obeys $v_{t+1} = (1 + \mathcal{M}/t)^2 v_t + \sigma_\eta^2/t^2$.

Its solution behaves like $\sigma_\eta^2/[(-2\mathcal{M} - 1)\, t]$ when $\mathcal{M} < -1/2$, like $\sigma_\eta^2 \log t/t$ when $\mathcal{M} = -1/2$, and like a multiple of $t^{2\mathcal{M}}$ when $\mathcal{M} > -1/2$.

The simulation below computes the mean error both ways and fits slopes on $100 \le t \le 3000$.

```
T, n_paths = 3000, 200
t = np.arange(1, T + 1)
fit = t >= 100

fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))
for k, θ in enumerate([0.5, 1.0, 2.0]):
    cm_θ = CredibilityModel(θ=θ, U_star=5.0)
    y_N_θ, M = cm_θ.nash(), -1 / (1 + θ**2)

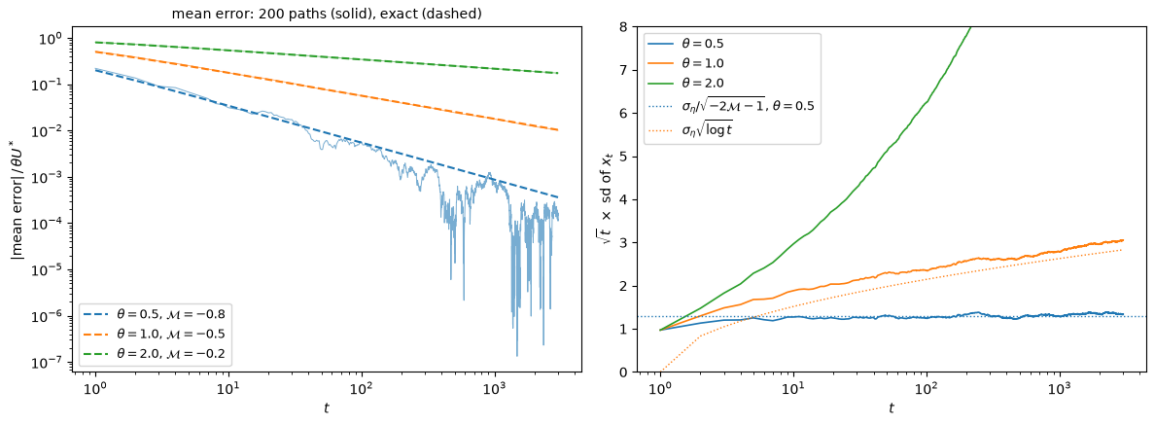
    # errors after t = 1, ..., T observations, one row per path
    e = np.array([ls_learning(cm_θ, T=T, seed=s)[0][1:]
                  for s in range(n_paths)]) - y_N_θ
    exact = ls_learning(cm_θ, T=T, σ_η=0.0)[0][1:] - y_N_θ     # E e_t
    mean, sd = e.mean(axis=0), e.std(axis=0)

    slope_exact = np.polyfit(np.log(t[fit]), np.log(np.abs(exact[fit])), 1)[0]
    slope_mc = np.polyfit(np.log(t[fit]), np.log(np.abs(mean[fit])), 1)[0]
    noisy = np.abs(exact) < 2 * sd / np.sqrt(n_paths)
    print(f"θ = {θ}:  M = {M:+.2f},  slope of exact mean {slope_exact:+.3f},  "
          f"slope of 200-path mean {slope_mc:+.3f}")
    if noisy.any():
        print(f"         exact mean below 2 Monte Carlo standard errors "
              f"from t = {t[noisy][0]}")
    print(f"         at t = {T}: mean error {mean[-1]:+.3f}, sd {sd[-1]:.3f}, "
          f"share of paths below Nash {np.mean(e[:, -1] < 0):.2f}")

    axes[0].loglog(t, np.abs(mean) / y_N_θ, color=f'C{k}', lw=0.7, alpha=0.6)
    axes[0].loglog(t, np.abs(exact) / y_N_θ, color=f'C{k}', ls='--', lw=1.5,
                   label=rf'$\theta = {θ}$, $\mathcal{{M}} = {M:.1f}$')
    axes[1].semilogx(t, np.sqrt(t) * sd, color=f'C{k}', lw=1.2,
                     label=rf'$\theta = {θ}$')

M_half = -1 / (1 + 0.5**2)                          # M at θ = 0.5
axes[1].axhline(1 / np.sqrt(-2 * M_half - 1), color='C0', ls=':', lw=1,
                label=r'$\sigma_\eta / \sqrt{-2\mathcal{M} - 1}$, $\theta = 0.5$')
axes[1].semilogx(t, np.sqrt(np.log(t)), color='C1', ls=':', lw=1,
                 label=r'$\sigma_\eta \sqrt{\log t}$')
axes[0].set_xlabel('$t$')
axes[0].set_ylabel(r'$|$mean error$|$ \,/\, \theta U^*$')
axes[0].set_title('mean error: 200 paths (solid), exact (dashed)', fontsize=10)
axes[0].legend(fontsize=9)
axes[1].set_xlabel('$t$')
axes[1].set_ylabel(r'$\sqrt{t}\ \times$ sd of $x_t$')
axes[1].set_ylim(0, 8)
axes[1].legend(fontsize=9)
plt.tight_layout()
```

```
θ = 0.5:  M = -0.80,  slope of exact mean -0.800,  slope of 200-path mean -1.057
          exact mean below 2 Monte Carlo standard errors from t = 43
          at t = 3000: mean error -0.000, sd 0.024, share of paths below Nash 0.
↪43
θ = 1.0:  M = -0.50,  slope of exact mean -0.500,  slope of 200-path mean -0.505
          at t = 3000: mean error -0.050, sd 0.056, share of paths below Nash 0.
↪82
θ = 2.0:  M = -0.20,  slope of exact mean -0.200,  slope of 200-path mean -0.200
          at t = 3000: mean error -1.735, sd 0.322, share of paths below Nash 1.
↪00
```



The exact mean errors fall with slopes equal to $\mathcal{M}$: -0.8, -0.5, and -0.2.

For $\theta = 1$ and $\theta = 2$ the 200-path average lies on top of the exact mean.

For $\theta = 0.5$ it does so only for a few dozen periods.

After that the mean error is smaller than the Monte Carlo standard error, and the average is noise.

That is the first lesson: when $\mathcal{M} < -1/2$, the initial condition is forgotten (like $t^{-0.8}$) faster than sampling noise dies out (like $t^{-1/2}$).

So convergence is limited by the noise, and $\sqrt{t}$ times the standard deviation levels off at $\sigma_\eta / \sqrt{-2\mathcal{M} - 1} \approx 1.29$.

When $\mathcal{M} > -1/2$ the ranking reverses.

For $\theta = 2$ both the mean error and the standard deviation shrink only like $t^{-0.2}$, so $\sqrt{t}$ times the standard deviation keeps growing, and the mean error dominates.

After 3000 periods the average $\theta = 2$ learner is still about 1.7 points below Nash, several standard deviations away, and every one of the 200 paths is below Nash.

The $\theta = 2$ learner does not wander around its limit: it climbs slowly, and nearly monotonically, from below.

The default $\theta = 1$ is the boundary between the two cases.

There the mean error and the standard deviation both shrink like $t^{-1/2}$, up to the factor $\sqrt{\log t}$ that makes $\sqrt{t}$ times the standard deviation creep upward.

Exercise 98.8.2

Replace the decreasing gain $1/(t-1)$ in (98.5) by a constant gain $\varepsilon \in (0, 1)$, so that $x_{t+1} = x_t + \varepsilon(y_t - x_t)$.

This is the adaptive-expectations rule of *Adaptive Expectations and the Phelps Problem* with $\varepsilon = 1 - \lambda$, but here the

government keeps best-responding one period at a time.

1. Show that $e_t = x_t - y^N$ is a first-order autoregression with coefficient $1 + \varepsilon\mathcal{M}$. Show that, when it is stationary, its variance is

$$V(\varepsilon) = \frac{\varepsilon\sigma_\eta^2}{(1-h')(2-\varepsilon(1-h'))}, \qquad h' = \frac{\theta^2}{1+\theta^2}.$$

   This is the formula of Exercise **??** for Bray's market, with $h' = B'$. For which gains is the recursion stable?

2. With $\theta = 1$ and $\sigma_\eta = 1$, use the option `ε` of `ls_learning` to simulate 100,000 periods for $\varepsilon \in \{0.01, 0.025, 0.05, 0.1, 0.2\}$. Drop the first 5,000 periods, and compare the sample variance of e_t with $V(\varepsilon)$.

3. Plot 3000 periods of x_t under $\varepsilon = 0.025$ and under the least squares gain, using the same shocks. What does a constant gain buy, and what does it cost?

Solution

Because $B(x) - x = \mathcal{M}(x - y^N)$, the update gives $e_{t+1} = (1 + \varepsilon\mathcal{M})\, e_t + \varepsilon\eta_t$.

With $\rho = 1+\varepsilon\mathcal{M} = 1-\varepsilon(1-h')$, the autoregression is stable when $|\rho| < 1$, that is, when $\varepsilon < 2/(1-h') = 2(1+\theta^2)$, which every gain below one satisfies.

Its variance is $\varepsilon^2\sigma_\eta^2/(1-\rho^2)$, and $1 - \rho^2 = \varepsilon(1-h')(2-\varepsilon(1-h'))$ gives the formula.

For small gains $V(\varepsilon) \approx \varepsilon\sigma_\eta^2/(2|\mathcal{M}|)$, which equals ε when $\theta = \sigma_\eta = 1$.

```
cm_1 = CredibilityModel(θ=1.0, U_star=5.0)
h_prime = cm_1.θ**2 / (1 + cm_1.θ**2)

print("   ε      sample variance    V(ε)    variance / ε")
for ε in (0.01, 0.025, 0.05, 0.1, 0.2):
    x_ε, _ = ls_learning(cm_1, T=100_000, seed=3, ε=ε)
    e_ε = x_ε[5_000:] - cm_1.nash()
    V = ε / ((1 - h_prime) * (2 - ε * (1 - h_prime)))
    print(f"{ε:6.3f}    {e_ε.var():12.5f}    {V:10.5f}    {e_ε.var() / ε:9.3f}")
```

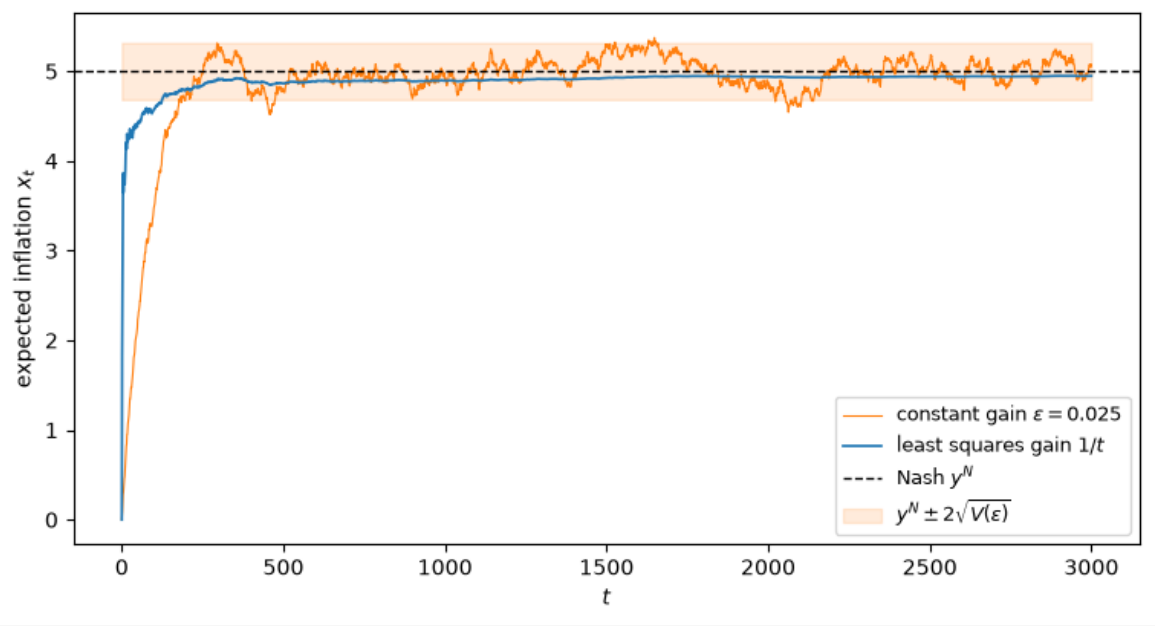
```
   ε      sample variance    V(ε)    variance / ε
 0.010         0.00968      0.01003       0.968
 0.025         0.02513      0.02516       1.005
 0.050         0.05054      0.05063       1.011
 0.100         0.10223      0.10256       1.022
 0.200         0.20959      0.21053       1.048
```

The sample variances match the formula, and the variance is nearly proportional to the gain.

The figure compares a constant gain with the least squares gain on the same shocks.

```
x_c, _ = ls_learning(cm_1, T=3000, seed=4, ε=0.025)
x_ls, _ = ls_learning(cm_1, T=3000, seed=4)
band = 2 * np.sqrt(0.025 / ((1 - h_prime) * (2 - 0.025 * (1 - h_prime))))

fig, ax = plt.subplots(figsize=(9, 4.5))
ax.plot(x_c, 'C1', lw=0.8, label=r'constant gain $\varepsilon = 0.025$')
ax.plot(x_ls, 'C0', lw=1.2, label='least squares gain $1/t$')
ax.axhline(cm_1.nash(), color='k', ls='--', lw=1, label='Nash $y^N$')
ax.fill_between(np.arange(3001), cm_1.nash() - band, cm_1.nash() + band,
                color='C1', alpha=0.15, label=r'$y^N \pm 2\sqrt{V(\varepsilon)}$')
ax.set_xlabel('$t$')
ax.set_ylabel('expected inflation $x_t$')
ax.legend(fontsize=9)
plt.show()
```



A constant gain buys a faster rate of forgetting.

The constant-gain learner forgets its initial condition geometrically, at rate $1 - \varepsilon(1 - h')$ per period.

The least squares learner moves faster at first, because its gain $1/t$ exceeds 0.025 for the first 40 periods.

But its mean error then shrinks only like $t^{-1/2}$, as Exercise 98.8.1 showed, and after a few hundred periods it is larger than the constant-gain learner's.

The cost is that the constant-gain learner never settles.

It converges only in distribution, to a cloud around y^N whose variance is proportional to the gain.

In this scalar model the cloud is all that happens, because the mean dynamics pull x straight back toward y^N and there is nowhere else to go.

In *Adaptive Learning and Escape Dynamics* the government's beliefs are a whole vector of Phillips-curve coefficients, learned with the same gain 0.025.

There the noise can push beliefs along a path that leads toward the Ramsey outcome, and the mean dynamics then pull them back toward Nash.

These are the recurrent escapes of Sargent [1999] and Cho *et al.* [2002].

Exercise 98.8.3

This exercise studies the *inflation bias* $y^N - y^R$ and two remedies discussed in *Remedies*.

Replace the government's payoff (98.1) by $-\frac{1}{2}\left[(U - kU^*)^2 + \alpha y^2\right]$, with $k \in [0, 1]$ and $\alpha > 0$.

The lecture's case is $k = 0$, $\alpha = 1$.

1. Show that the best response is $B(x) = [\theta(1 - k)U^* + \theta^2 x]/(\theta^2 + \alpha)$ and that Nash inflation is $y^N = \theta(1 - k)U^*/\alpha$. Check the formula by iterating the best response.
2. Alan Blinder's remedy is an authority who knows the natural rate and never wants unemployment to differ from it ($k = 1$). Rogoff [1985]'s remedy is a conservative central banker who weights inflation more heavily than

society does ($\alpha > 1$). What does each do to the inflation bias?

3. Now let the Phillips curve be $U = U^* - \theta(y - x) + \zeta$. Here ζ is an IID shock with mean zero and variance σ_ζ^2, which the policymaker sees before setting y but the public does not see when it forms x. Show that in the Nash equilibrium $x = \theta(1-k)U^*/\alpha$, $y = x + \theta\zeta/(\theta^2 + \alpha)$, and $U = U^* + \alpha\zeta/(\theta^2 + \alpha)$.
4. Society keeps the payoff (98.1). Set $\theta = 1$ and $U^* = 5$. For a Rogoff banker ($k = 0$) and $\sigma_\zeta \in \{1, 2.5, 5\}$, plot society's expected loss $\frac{1}{2}E(U^2 + y^2)$ against α and find the best α. Compare with Blinder's authority ($k = 1$, $\alpha = 1$).

Solution

1. The first-order condition is $\theta\left[(1-k)U^* - \theta(y - x)\right] = \alpha y$, which gives $B(x)$.

Setting $x = B(x)$ gives $\alpha y^N = \theta(1-k)U^*$.

Because the slope $\theta^2/(\theta^2 + \alpha)$ of B is less than one, best response dynamics converge to y^N, which lets us check the formula.

```
def B_general(x, θ, U_star, k, α):
    "Best response of a policymaker with payoff -[(U - k U*)^2 + α y^2]/2."
    return (θ * (1 - k) * U_star + θ**2 * x) / (θ**2 + α)


θ_b, U_b = 1.0, 5.0
for k, α in [(0, 1), (0, 2), (0, 5), (0.5, 1), (1, 1)]:
    y = 0.0
    for _ in range(2000):
        y = B_general(y, θ_b, U_b, k, α)
    print(f"k = {k:3}, α = {α}:  iterated best response {y:.4f},  "
          f"θ(1-k)U*/α = {θ_b * (1 - k) * U_b / α:.4f}")
```

```
k =   0, α = 1:  iterated best response 5.0000,  θ(1-k)U*/α = 5.0000
k =   0, α = 2:  iterated best response 2.5000,  θ(1-k)U*/α = 2.5000
k =   0, α = 5:  iterated best response 1.0000,  θ(1-k)U*/α = 1.0000
k = 0.5, α = 1:  iterated best response 2.5000,  θ(1-k)U*/α = 2.5000
k =   1, α = 1:  iterated best response 0.0000,  θ(1-k)U*/α = 0.0000
```

2. Both remedies leave unemployment at the natural rate, which is all any policymaker can achieve on average, so what they change is inflation.

Blinder's authority ($k = 1$) eliminates the bias altogether: it has no reason to push unemployment below U^*, and $y^N = 0$ is the Ramsey outcome.

Rogoff's banker shrinks the bias in proportion to $1/\alpha$, and eliminates it only as $\alpha \to \infty$.

3. With the shock, the first-order condition becomes $\theta\left[(1-k)U^* + \zeta - \theta(y - x)\right] = \alpha y$.

Taking expectations, and using $x = Ey$, gives $x = \theta(1-k)U^*/\alpha$; subtracting gives $y - x = \theta\zeta/(\theta^2 + \alpha)$.

Then $U = U^* + \zeta - \theta(y - x) = U^* + \alpha\zeta/(\theta^2 + \alpha)$.

A more conservative banker offsets less of the shock, so a larger α makes unemployment *more* variable and inflation less variable.

4. Society's expected loss is

$$\frac{1}{2}\left[U^{*2} + \left(\frac{\theta(1-k)U^*}{\alpha}\right)^2 + \sigma_\zeta^2 \frac{\alpha^2 + \theta^2}{(\theta^2 + \alpha)^2}\right].$$

The middle term is the inflation bias and the last term is the cost of the shocks.

The bias term falls with α.

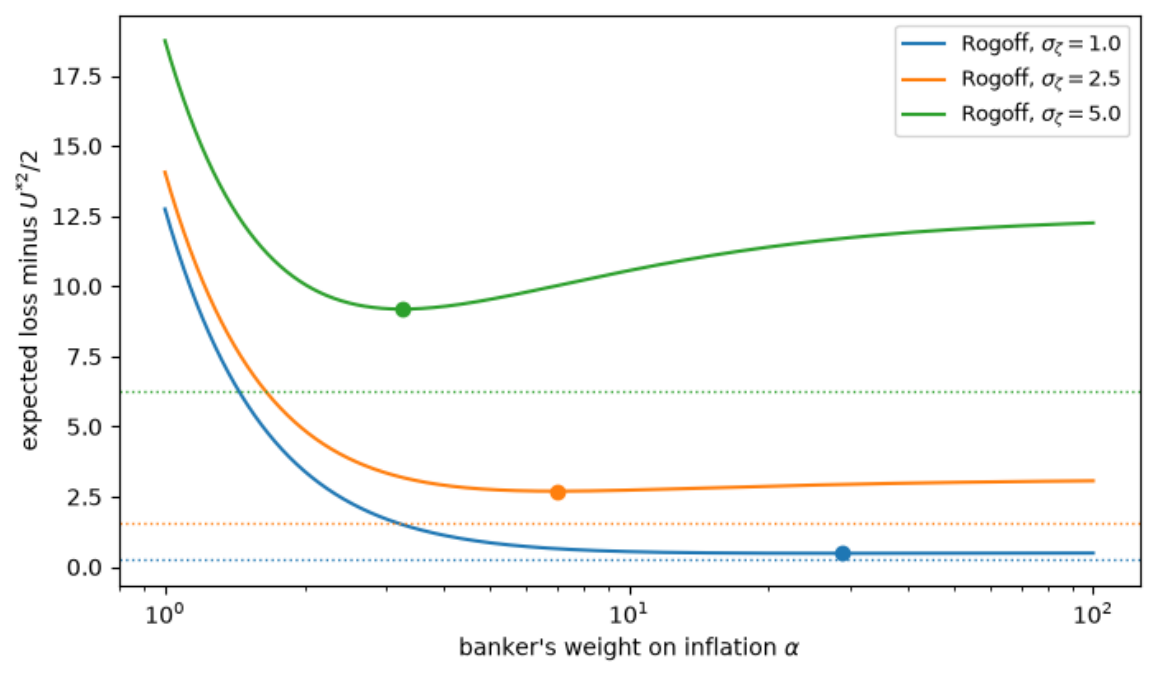
The shock term is smallest at $\alpha = 1$, where it equals $\sigma_\zeta^2/(1+\theta^2)$, and it rises toward σ_ζ^2 as $\alpha \to \infty$.

That is Rogoff's tradeoff.

```
def society_loss(α, k, σ_ζ, θ=θ_b, U_star=U_b):
    "Society's expected loss (1/2) E(U^2 + y^2) when an authority (k, α) sets y."
    bias = θ * (1 - k) * U_star / α
    return 0.5 * (U_star**2 + bias**2
                  + σ_ζ**2 * (α**2 + θ**2) / (θ**2 + α)**2)


α_grid = np.geomspace(1, 100, 4001)
fig, ax = plt.subplots(figsize=(8, 4.5))
for j, σ_ζ in enumerate([1.0, 2.5, 5.0]):
    loss = society_loss(α_grid, 0, σ_ζ) - 0.5 * U_b**2
    i = loss.argmin()
    blinder = society_loss(1.0, 1, σ_ζ) - 0.5 * U_b**2
    ax.plot(α_grid, loss, color=f'C{j}', label=rf'Rogoff, $\sigma_\zeta = {σ_ζ}$')
    ax.plot(α_grid[i], loss[i], 'o', color=f'C{j}')
    ax.axhline(blinder, color=f'C{j}', ls=':', lw=1)
    print(f"σ_ζ = {σ_ζ}:  best α = {α_grid[i]:.2f},  excess loss there {loss[i]:.
↪3f};"
          f"  at α = 1 {loss[0]:.3f};  Blinder {blinder:.3f}")
ax.set_xscale('log')
ax.set_xlabel(r"banker's weight on inflation $\alpha$")
ax.set_ylabel(r'expected loss minus $U^{*2}/2$')
ax.legend(fontsize=9)
plt.show()
```

```
σ_ζ = 1.0:  best α = 28.71,  excess loss there 0.483;  at α = 1 12.750; ␣
↪Blinder 0.250
σ_ζ = 2.5:  best α = 6.97,  excess loss there 2.697;  at α = 1 14.062;  Blinder␣
↪1.562
σ_ζ = 5.0:  best α = 3.24,  excess loss there 9.185;  at α = 1 18.750;  Blinder␣
↪6.250
```



The dotted lines are Blinder's authority.

For each σ_ζ the best Rogoff banker is more conservative than society ($\alpha > 1$) but not infinitely so, and the noisier the economy, the less conservative the best banker.

Blinder's authority beats every Rogoff banker.

It removes the bias without distorting stabilization: its loss, $\frac{1}{2}\left[U^{*2} + \sigma_\zeta^2/(1+\theta^2)\right]$, is what society would get by committing to the best state-contingent rule $y = \theta\zeta/(1+\theta^2)$.

The reason is that $-\frac{1}{2}(U - U^*)^2$ differs from $-\frac{1}{2}U^2$ by a term that is linear in U, and that linear term is the whole source of the temptation to inflate.

CHAPTER

NINETYNINE

CREDIBLE GOVERNMENT POLICIES

Contents

- *Credible Government Policies*
  - *Overview*
  - *The repeated economy*
  - *Recursive strategies and promised values*
  - *The Abreu–Pearce–Stacchetti method*
  - *The best and the worst equilibrium*
  - *Examples of recursive equilibria*
  - *Multiplicity*
  - *Numerical example*
  - *Interpretations*
  - *Exercises*

99.1 Overview

The Credibility Problem left the government in a trap.

Without a technology for committing itself, a government that re-optimizes every period ends at the Nash outcome, inflating at rate θU^* and getting nothing for it.

That was a one-period story, and it invited an obvious objection: a government that expects to face the same public again tomorrow has a *reputation* to protect.

This lecture, following chapter 4 of Sargent [1999], takes the objection seriously.

We repeat the Kydland–Prescott economy forever, let the government's current action depend on the whole history of outcomes, and ask which outcome paths can be supported as **subgame perfect equilibria**.

The answer is not the one the objection anticipated.

Reputation does not single out the Ramsey outcome, and it does not confirm the Nash outcome either.

It delivers a *continuum* of equilibrium values, some better than Nash and some considerably worse, with no principle inside the model to select among them.

That is the conclusion this lecture exists to establish, and Sargent states it in one sentence:

> My conclusion is that the multiplicity of credible plans replaces pessimism with agnosticism.

It matters for the argument of the whole suite.

The Rise and Fall of U.S. Inflation rests part of its case against the *triumph of natural-rate theory* on exactly this weakness: a theory with so many equilibria makes few predictions, so a story that depends on policy makers having learned the right one is doing work the theory cannot do for it.

99.1.1 What we build

The machinery is the recursive method of Abreu, Pearce, and Stacchetti [Abreu *et al.*, 1990], which computes not an optimal *value* but the whole *set* of equilibrium values.

Ordinary dynamic programming iterates an operator that maps continuation *values* into values.

APS iterate an operator that maps *sets* of continuation values into sets of values, and the set of subgame perfect equilibrium values is its largest fixed point.

Note

Making a promised value into a state variable, and then doing dynamic programming in that state, is the technique sometimes called **dynamic programming squared**.

It is the organizing idea of the QuantEcon lecture on Stackelberg plans and of the recursive treatments of Ramsey problems in Ljungqvist and Sargent [2018].

This lecture is a compact instance: the state is a promised value, and the object we compute is a set.

A *scalar* promised value suffices here because the private sector's problem is static: the public's forecast x_t concerns only this period's inflation, so it is pinned down by the government's current prescription.

In the monetary model of Chang [1998], households' money demand looks forward, so the government must also promise them a continuation marginal value of money, and the state becomes a pair; Competitive Equilibria of a Model of Chang and Credible Government Policies in a Model of Chang compute the resulting two-dimensional sets with the same set-iteration idea.

Sustainable Plans for a Calvo Model builds Abreu's stick-and-carrot plans, which appear below, in a linear-quadratic version of Calvo's model.

We then use the machinery three ways.

We construct particular equilibria by guess-and-verify — infinite repetition of Nash, infinite repetition of something better, and Abreu's stick-and-carrot, which can be *worse* than Nash.

We compute the best and worst equilibrium values from two small programming problems, and check them against a direct iteration of the APS operator.

And we exhibit three quite different equilibria that all attain the same worst value, which is the sharpest illustration of how little the equilibrium concept pins down.

Note

Sargent's own advice to the reader is worth passing on: this chapter is difficult if you have not seen the theory before, and a reader willing to accept its verdict — that the theory of credible policy yields agnosticism rather than a prediction — can go straight to *Adaptive Expectations and the Phelps Problem* without losing the thread of the argument.

Let's start with our imports:

```
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from typing import NamedTuple
```

99.2 The repeated economy

The one-period economy is the one from *The Credibility Problem*.

Let (U, y, x) be unemployment, inflation, and the public's expectation of inflation.

Unemployment obeys the expectations-augmented Phillips curve $U = U^* - \theta(y - x)$, and the government's one-period payoff, written as a function of (x, y), is

$$r(x, y) = -\frac{1}{2}\left[(U^* - \theta(y - x))^2 + y^2\right]. \tag{99.1}$$

The government's one-period best response to an expectation x is

$$B(x) = \frac{\theta\,(U^* + \theta x)}{\theta^2 + 1}, \tag{99.2}$$

the Nash outcome is $y^N = \theta U^*$, and the Ramsey outcome is $y^R = 0$.

Two things are new.

First, the economy repeats for $t = 1, 2, \ldots$, and the government ranks outcome paths $(x, y) = \{x_t, y_t\}_{t=1}^{\infty}$ by

$$V^g(x, y) = (1 - \delta)\sum_{t=1}^{\infty} \delta^{t-1} r(x_t, y_t), \qquad \delta \in (0, 1). \tag{99.3}$$

The factor $(1 - \delta)$ puts V^g in the same units as a one-period payoff, so values and one-period returns can be compared directly.

Second, inflation is confined to a bounded interval $Y = [0, y^{\#}]$.

The lower bound is the Ramsey outcome; the upper bound $y^{\#}$ is what makes the government's problem non-trivial, as we shall see when we look for the *worst* equilibrium.

We use the parameter values of the numerical example in chapter 4 of Sargent [1999]: $\theta = 1.25$, $U^* = 5.5$, $y^{\#} = 10$, and $\delta = 0.95$.

They differ from the $\theta = 1$, $U^* = 5$ of *The Credibility Problem*, so that our numbers can be checked against the book's.

```
class Model(NamedTuple):
    θ: float = 1.25          # slope of the Phillips curve
    U_star: float = 5.5      # natural rate of unemployment
    y_max: float = 10.0      # y^#, the highest admissible inflation rate
    δ: float = 0.95          # discount factor


def r(m, x, y):
    "One-period government payoff when the public expects x and inflation is y."
    U = m.U_star - m.θ * (y - x)
    return -0.5 * (U**2 + y**2)
```

(continues on next page)

(continued from previous page)

```
def B(m, x):
    "The government's one-period best response to expected inflation x."
    return m.θ * (m.U_star + m.θ * x) / (m.θ**2 + 1)


def y_nash(m):
    return m.θ * m.U_star
```

Two payoff schedules do all the work below, both read as functions of a single inflation rate y that the public has come to expect.

The first is the **rational-expectations payoff** $r(y, y)$: what the government gets if it delivers the inflation the public expects.

The second is the **deviation payoff** $r(y, B(y))$: what it gets if the public expects y but the government yields to temptation and best-responds.

Both have closed forms worth recording, because they explain everything that follows:

$$r(y, y) = -\frac{1}{2}\left(U^{*2} + y^2\right), \qquad r(y, B(y)) = -\frac{\left(U^* + \theta y\right)^2}{2\left(1 + \theta^2\right)}. \tag{99.4}$$

```
def r_keep(m, y):
    "Payoff from delivering the expected inflation rate: r(y, y)."
    return -0.5 * (m.U_star**2 + y**2)


def r_cheat(m, y):
    "Payoff from best-responding to an expectation of y: r(y, B(y))."
    return -(m.U_star + m.θ * y)**2 / (2 * (1 + m.θ**2))


m = Model()
ys = np.linspace(0, m.y_max, 7)
print("closed forms agree with the definitions:",
      np.allclose(r_keep(m, ys), r(m, ys, ys)),
      np.allclose(r_cheat(m, ys), r(m, ys, B(m, ys))))
```

```
closed forms agree with the definitions: True True
```

The gap between them is a perfect square,

$$r(y, B(y)) - r(y, y) = \frac{\left(y - y^N\right)^2}{2\left(1 + \theta^2\right)} \;\geq\; 0. \tag{99.5}$$

So the deviation schedule lies above the other everywhere except at the Nash rate, where the two curves are *tangent*: there the temptation to deviate vanishes because $B(y^N) = y^N$.

```
grid = np.linspace(0, m.y_max, 400)
yN = y_nash(m)
v_R, v_N = r_keep(m, 0.0), r_keep(m, yN)
v_cand = r_cheat(m, m.y_max)                    # candidate worst value

fig, ax = plt.subplots(figsize=(7.5, 5.5))
ax.plot(grid, r_keep(m, grid), 'C0', lw=1.6, label='$r(y, y)$: deliver what is␣
```

(continues on next page)

(continued from previous page)

```
↪expected')
ax.plot(grid, r_cheat(m, grid), 'C1', lw=1.6, label='$r(y, B(y))$: deviate')
ax.plot(0.0, v_R, 'ko', ms=6)
ax.annotate('$v^R$', (0.0, v_R), textcoords='offset points', xytext=(8, 4))
ax.plot(yN, v_N, 'ko', ms=6)
ax.annotate('$v^N$', (yN, v_N), textcoords='offset points', xytext=(8, 4))
ax.plot(m.y_max, v_cand, 'ko', ms=6)
ax.annotate(r'$r(y^\#, B(y^\#))$', (m.y_max, v_cand), textcoords='offset points',
            xytext=(-70, 6))
ax.axvline(yN, color='k', lw=0.6, ls=':')
ax.set_xlabel('inflation rate $y$ expected by the public')
ax.set_ylabel('one-period payoff')
ax.legend(loc='lower left')
plt.show()
```

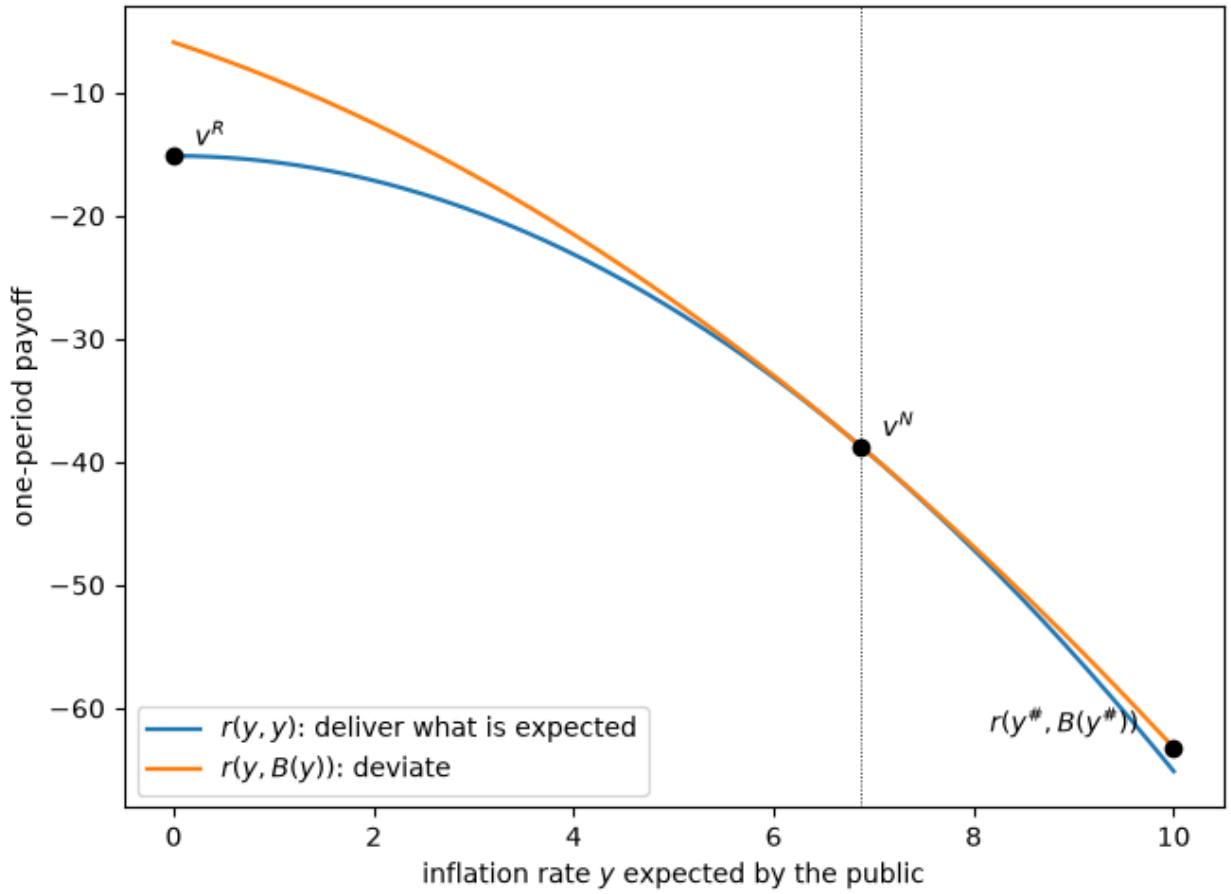


Fig. 99.1: The two payoff schedules: delivering expected inflation against deviating from it

The vertical gap between the curves is the government's one-period gain from deviating.

Below y^N, which includes the Ramsey outcome, $B(y) > y$, so the temptation is to inflate *more* than expected.

Above y^N, $B(y) < y$, and the temptation is to inflate *less* than expected.

By (99.5) the gap is symmetric about y^N, so the gap alone does not favor the top of Y.

What will put the *worst* equilibrium at the top of Y is the level of the deviation schedule.

$r(y, B(y))$ falls monotonically in y, so the least a government can secure for itself by deviating is smallest when the public expects $y^{\#}$.

The point $r(y^{\#}, B(y^{\#}))$ marked in the figure is our first candidate for the worst equilibrium value.

99.3 Recursive strategies and promised values

A government strategy must be able to condition today's action on the entire past.

Carrying whole histories around is unmanageable, so we follow Sargent [1999] in restricting attention to strategies with a recursive representation — a restriction that costs nothing, since it excludes no equilibrium *values*.

Definition 99.3.1 (Recursive government strategy)

A recursive government strategy is a pair of functions $\sigma = (\sigma_1, \sigma_2)$ together with an initial condition v_1, with the structure

$$v_1 \in \mathbb{R} \text{ given}, \qquad y_t = \sigma_1(v_t), \qquad v_{t+1} = \sigma_2(v_t, x_t, y_t),$$

where v_t is a state variable that summarizes the history of outcomes before t.

Each member of the private sector knows v_t and knows σ, and so forecasts

$$x_t = \sigma_1(v_t). \tag{99.6}$$

Equation (99.6) builds rational expectations into the private sector: the public is never surprised along the equilibrium path.

A strategy (σ, v_1) generates an entire outcome path, and hence a value $V^g(\sigma, v_1)$ through (99.3).

The theory of credible policy now does something that at first looks like a trick.

It ties the past to the future by requiring the state variable to *be* the value it generates,

$$v = V^g(\sigma, v). \tag{99.7}$$

So v works two shifts at once.

In $v_{t+1} = \sigma_2(v_t, x_t, y_t)$ it is a bookkeeping device that records what has happened.

In (99.7) it is a **promised value** — the discounted future the government is owed on entering the period.

Reading it the second way is what makes the machinery go: the strategy hands the government current and future outcomes that make it *want* to do what is expected of it.

Note

Two readings of σ are available and, within an equilibrium, impossible to tell apart.

It can be a decision rule that the government chooses, or a description of a system of public expectations to which the government conforms.

We return to this ambiguity in the interpretation section, because it is where the theory's agnosticism comes from.

99.3.1 Historical antecedents

Formalizing credibility was an achievement of the 1980s, but a sophisticated grasp of the idea is much older.

In 1784 the finance minister Jacques Necker explained to Louis XVI why an absolute monarch, free to renege on debt payments, found it hard to borrow [Sargent and Velde, 1995]:

> Therefore one can rekindle or sustain public trust only by giving reassurances on the sovereign's intentions, and by proving that no motive can incite him to fail his obligations.

Every clause of that sentence is in the modern definition.

A king proves he will never fail his obligations by never *wanting* to fail them — by conforming to a system of expectations that carries its own incentive constraints, so that at every date and in every contingency his current payoff plus continuation value is higher if he confirms expectations than if he disappoints them.

Definition 99.3.2 (Subgame perfect equilibrium)

A recursive strategy with promised value v is a **subgame perfect equilibrium** (SPE) if and only if

(a) $\sigma_2(v, \sigma_1(v), \eta)$ is itself a value attainable by a subgame perfect equilibrium, for every $\eta \in Y$; and

(b) writing $y = \sigma_1(v)$,

$$v = (1-\delta)\, r(y,y) + \delta\, \sigma_2(v,y,y) \;\geq\; (1-\delta)\, r(y,\eta) + \delta\, \sigma_2(v,y,\eta), \qquad \forall\, \eta \in Y.$$

The definition attaches four objects to an equilibrium: a promised value v; a first-period outcome (y, y); a continuation value v' if the prescribed outcome is observed; and another continuation value $\tilde{v}$ if it is not.

In those terms condition (b) reads

$$v = (1-\delta) r(y,y) + \delta v' \;\geq\; (1-\delta) r(y,\eta) + \delta \tilde{v}, \qquad \forall \eta \in Y,$$

which simply says the government does better adhering than departing.

Condition (a) says the continuation values must themselves be equilibrium values.

The definition is circular — equilibrium values appear on both sides — and that circularity is exactly what recursivity buys, and what APS learned to exploit.

99.4 The Abreu–Pearce–Stacchetti method

Dynamic programming computes an optimal value function by iterating the Bellman equation, a map that turns tomorrow's value function into today's.

APS adapted the idea to equilibrium *sets*.

Start with a candidate set $W \subset \mathbb{R}$ of continuation values.

Pick a first-period rational-expectations outcome (y, y) and two continuation values from W: a w_1 to reward adherence and a w_2 to punish deviation.

If w_1 is high enough and w_2 low enough, the pair supports y, and delivers the value

$$w = (1-\delta) r(y,y) + \delta w_1 \;\geq\; (1-\delta) r(y,\eta) + \delta w_2, \qquad \forall \eta \in Y.$$

Definition 99.4.1 (Admissibility)

A pair (y, w) is **admissible** with respect to a set of continuation values W if there exist $w_1, w_2 \in W$ with $w = (1-\delta)r(y,y) + \delta w_1 \geq (1-\delta)r(y,\eta) + \delta w_2$ for all $\eta \in Y$.

Let $\mathcal{T}(W)$ collect the w components of all admissible pairs.

APS and Sargent [1999] call this operator B; we write $\mathcal{T}$ to keep B for the government's best response.

This construction builds in condition (b) of *Definition 99.3.2* but ignores condition (a), because the continuation values were drawn from an arbitrary set.

Definition 99.4.2 (Self-generation)

A set W of prospective continuation values is **self-generating** if $W \subseteq \mathcal{T}(W)$.

Every value in a self-generating set is supported by continuation values drawn from that same set — which is condition (a).

APS proved four things:

1. the set V of SPE values is the largest self-generating set;
2. $\mathcal{T}$ maps compact sets into compact sets;
3. $\mathcal{T}$ is monotone: $W_2 \subseteq W_1$ implies $\mathcal{T}(W_2) \subseteq \mathcal{T}(W_1)$;
4. starting from any compact W_0 with $V \subseteq \mathcal{T}(W_0) \subseteq W_0$, the iteration $W_j = \mathcal{T}(W_{j-1})$ converges monotonically to $V = \mathcal{T}(V)$ (Abreu *et al.* [1990], Theorem 5).

The requirement $V \subseteq \mathcal{T}(W_0)$ in the last item cannot be dropped, although the book's statement of the theorem omits it.

The single point $W_0 = \{v^N\}$ satisfies $\mathcal{T}(W_0) \subseteq W_0$, yet it is a fixed point of $\mathcal{T}$ far smaller than V (exercise Exercise 99.10.6).

A safe choice is the set of all feasible values, $W_0 = [r(y^\#, y^\#),\ 0]$.

It contains V, and $\mathcal{T}$ maps it into itself, so monotonicity gives $V = \mathcal{T}(V) \subseteq \mathcal{T}(W_0) \subseteq W_0$.

99.4.1 Computing the operator

Two simplifications make $\mathcal{T}$ easy to evaluate here.

Because the government's temptation is worst when it best-responds, the constraint in *Definition 99.4.1* binds hardest at $\eta = B(y)$, so we may replace "for all η" by that single deviation.

And because a lower punishment relaxes the constraint, we may always set w_2 to the smallest element of W.

Writing $W = [\underline{w}, \overline{w}]$, the outcome y is then admissible provided some $w_1 \in W$ satisfies

$$w_1 \;\geq\; \frac{(1-\delta)\,[r(y, B(y)) - r(y,y)]}{\delta} + \underline{w} \;\equiv\; \ell(y), \tag{99.8}$$

and the values it generates sweep out the interval

$$\left[(1-\delta)r(y,y) + \delta \max(\underline{w}, \ell(y)),\ (1-\delta)r(y,y) + \delta\overline{w}\right].$$

```
def T_operator(m, W, n_grid=4001):
    """
    One application of the APS operator 𝒯 to an interval W = [w_lo, w_hi].

    Returns the new interval, or None if no first-period outcome is admissible.
    """
    w_lo, w_hi = W
    y = np.linspace(0.0, m.y_max, n_grid)
    ℓ = (1 - m.δ) * (r_cheat(m, y) - r_keep(m, y)) / m.δ + w_lo    # the bound ℓ(y) of␣
↪cp_bound
    ok = ℓ <= w_hi                                                  # admissible␣
↪outcomes
    if not ok.any():
        return None
    keep = (1 - m.δ) * r_keep(m, y)
    lows = keep[ok] + m.δ * np.maximum(w_lo, ℓ[ok])
    highs = keep[ok] + m.δ * w_hi
    return lows.min(), highs.max()


def solve_aps(m, W0=None, tol=1e-12, max_iter=10_000):
    """
    Iterate the APS operator from W0 until it settles.

    The default W0 is the set of all feasible values, [r(y^#, y^#), 0].
    """
    W = (r_keep(m, m.y_max), 0.0) if W0 is None else W0
    for it in range(max_iter):
        W_new = T_operator(m, W)
        if W_new is None:
            return None, it
        if max(abs(W_new[0] - W[0]), abs(W_new[1] - W[1])) < tol:
            return W_new, it
        W = W_new
    return W, max_iter
```

We start from $W_0 = [r(y^\#, y^\#),\, 0]$ and iterate.

```
W, iters = solve_aps(m)
print(f"converged in {iters} iterations")
print(f"set of SPE values V = [{W[0]:.4f}, {W[1]:.4f}]")
```

```
converged in 534 iterations
set of SPE values V = [-63.2195, -15.1250]
```

99.5 The best and the worst equilibrium

APS's iteration is general but tells us little about *why* the set ends where it does.

For this economy both endpoints can be found by hand, and the arguments are instructive.

99.5.1 The worst

The worst equilibrium value solves

$$\underline{v} = \min_{y \in Y,\, v_1 \in V} [(1-\delta) r(y,y) + \delta v_1] \quad \text{subject to} \quad (1-\delta) r(y,y) + \delta v_1 \geq (1-\delta) r(y, B(y)) + \delta \underline{v},$$

where the *worst* value is used as the continuation in the event of a deviation — the harshest punishment available.

For a given y, the objective is smallest when v_1 is as low as the constraint allows, so at the minimum the constraint binds, and the two sides collapse to $\underline{v} = r(y, B(y))$ for the minimizing y.

If every $y \in Y$ could be used, finding the worst equilibrium would reduce to a one-dimensional problem:

$$\underline{v} = \min_{y \in Y} r(y, B(y)) = -\frac{\left(U^* + \theta y^\#\right)^2}{2\left(1+\theta^2\right)}, \tag{99.9}$$

where the closed form follows from (99.4), and the minimizing action is $y^\#$ because the deviation payoff falls monotonically in y.

But (99.9) is only a *candidate*.

To deliver it with a first period at $y^\#$, the strategy must reward adherence with the continuation value v_1 that solves $\underline{v} = (1-\delta) r(y^\#, y^\#) + \delta v_1$,

$$v_1 = \frac{r(y^\#, B(y^\#)) - (1-\delta)\, r(y^\#, y^\#)}{\delta}, \tag{99.10}$$

and that reward must itself be an equilibrium value: $v_1 \leq \overline{v}$.

That check is the last step of the book's construction: compute the candidate, compute the best value $\overline{v}$ that it implies (below), and verify $v_1 \leq \overline{v}$.

The check can fail.

The reward (99.10) grows like $1/\delta$: an impatient government cannot be promised enough future reward to accept a period at $y^\#$.

The worst equilibrium then starts at a lower inflation rate, the two programming problems must be solved jointly, and we let the APS iteration do that.

Now the role of the upper bound $y^\#$ is clear.

The coefficient on $y^{\#2}$ in (99.10) is $\left[(1-\delta) - \theta^2/(1+\theta^2)\right]/(2\delta)$.

So when $\delta > 1/(1+\theta^2)$, the reward needed falls without limit as $y^\#$ rises, the check keeps passing, and $y^\#$ is what keeps the worst punishment finite.

For a less patient government the check eventually fails, and beyond some point raising $y^\#$ changes nothing.

Proposition 99.5.1 (The worst SPE is self-enforcing)

At the worst equilibrium the continuation value following a deviation equals the initial promised value, so a deviation simply restarts the equilibrium.

99.5.2 The best

Given $\underline{v}$ as the threat, the best value solves

$$\overline{v} = \max_{y \in Y} r(y, y) \quad \text{subject to} \quad r(y, y) \geq (1-\delta) r(y, B(y)) + \delta \underline{v}, \tag{99.11}$$

where we have used the fact that the best value must reward adherence with itself, so that $\overline{v} = (1-\delta) r(y, y) + \delta \overline{v} = r(y, y)$.

Proposition 99.5.2 (The best SPE is self-rewarding)

At the best equilibrium the continuation value following adherence equals the promised value.

The code below implements both programming problems, with the check.

```
def next_value(m, v):
    "Solve v = (1-δ) r(y^#, y^#) + δ v' for the continuation value v'."
    return (v - (1 - m.δ) * r_keep(m, m.y_max)) / m.δ


def best_value(m, v_lo, n_grid=200_001):
    "The best SPE value, given the worst value as the punishment threat."
    y = np.linspace(0.0, m.y_max, n_grid)
    feasible = r_keep(m, y) >= (1 - m.δ) * r_cheat(m, y) + m.δ * v_lo
    if not feasible.any():
        return None, None
    vals = np.where(feasible, r_keep(m, y), -np.inf)
    k = vals.argmax()
    return vals[k], y[k]


def worst_value(m):
    """
    The worst SPE value and the action that attains it.

    The candidate r(y^#, B(y^#)) of cp_worst is accepted if the reward v_1 of
    cp_reward is no larger than the best value the candidate implies.
    Otherwise the worst value comes from the APS iteration, and it is attained
    at the highest action whose required reward ℓ(y) of cp_bound equals the
    best value (solved for y using cp_gap).
    """
    v_cand = r_cheat(m, m.y_max)
    v_best, _ = best_value(m, v_cand)
    if next_value(m, v_cand) <= v_best:
        return v_cand, m.y_max
    (w_lo, w_hi), _ = solve_aps(m)
    y_hat = y_nash(m) + np.sqrt(2 * (1 + m.θ**2) * m.δ * (w_hi - w_lo) / (1 - m.δ))
    return w_lo, y_hat


v_lo, y_sharp = worst_value(m)
v_hi, y_best = best_value(m, v_lo)

print(f"worst value  v_lo = {v_lo:.4f}   attained with y = {y_sharp:.2f}")
print(f"reward required by the candidate  v_1 = {next_value(m, v_lo):.4f}"
      f"  <= v_hi? {next_value(m, v_lo) <= v_hi}")
```

(continues on next page)

(continued from previous page)

```
print(f"best  value  v_hi = {v_hi:.4f}   attained with y = {y_best:.2f}")
print(f"APS iteration gave [{W[0]:.4f}, {W[1]:.4f}]")
print(f"the two agree: {np.allclose(W, (v_lo, v_hi), atol=1e-6)}")
```

```
worst value  v_lo = -63.2195   attained with y = 10.00
reward required by the candidate  v_1 = -63.1192  <= v_hi? True
best  value  v_hi = -15.1250   attained with y = 0.00
APS iteration gave [-63.2195, -15.1250]
the two agree: True
```

The programming problems and the set iteration agree, and at this discount factor the best equilibrium is Ramsey itself.

The check passes easily at $\delta = 0.95$, but not for every discount factor.

```
print("    δ    candidate   reward v_1   best given candidate   check   worst value␣
↪at y")
for δ in (0.02, 0.04, 0.05, 0.2, 0.95):
    md = m._replace(δ=δ)
    v_cand = r_cheat(md, md.y_max)
    v_1, v_best = next_value(md, v_cand), best_value(md, v_cand)[0]
    v_worst, y_worst = worst_value(md)
    print(f"{δ:6.2f} {v_cand:11.4f} {v_1:12.4f} {v_best:16.4f} {str(v_1 <= v_best):>
↪11}"
          f" {v_worst:12.4f} {y_worst:6.2f}")
```

```
     δ    candidate   reward v_1   best given candidate   check   worst value  at y
  0.02    -63.2195      30.1494         -27.3125       False     -49.0384   8.28
  0.04    -63.2195     -17.4878         -22.8882       False     -60.5483   9.69
  0.05    -63.2195     -27.0152         -21.2924        True     -63.2195  10.00
  0.20    -63.2195     -55.5976         -15.1250        True     -63.2195  10.00
  0.95    -63.2195     -63.1192         -15.1250        True     -63.2195  10.00
```

For the two smallest discount factors the reward the candidate needs exceeds the best equilibrium value.

The worst equilibrium then starts below $y^{\#}$, and its value lies above $r(y^{\#}, B(y^{\#}))$; exercise Exercise 99.10.2 finds the discount factor at which this begins.

The next cell varies the upper bound $y^{\#}$ for a patient and an impatient government.

```
print(f"1/(1 + θ²) = {1 / (1 + m.θ**2):.3f}")
for δ in (0.2, 0.95):
    for y_max in (10.0, 30.0, 100.0):
        v_worst, y_worst = worst_value(m._replace(δ=δ, y_max=y_max))
        print(f"δ = {δ:4.2f}, y^# = {y_max:5.1f}:  worst value {v_worst:10.2f}"
              f"  attained with y = {y_worst:6.2f}")
```

```
1/(1 + θ²) = 0.390
δ = 0.20, y^# =  10.0:  worst value     -63.22  attained with y =  10.00
δ = 0.20, y^# =  30.0:  worst value    -249.29  attained with y =  24.20
δ = 0.20, y^# = 100.0:  worst value    -248.94  attained with y =  24.18
δ = 0.95, y^# =  10.0:  worst value     -63.22  attained with y =  10.00
δ = 0.95, y^# =  30.0:  worst value    -360.78  attained with y =  30.00
δ = 0.95, y^# = 100.0:  worst value   -3322.98  attained with y = 100.00
```

For the patient government the worst value keeps falling as $y^{\#}$ rises, and it is always attained at $y^{\#}$.

For the impatient government it stops falling: once $y^{\#}$ is large enough, the worst equilibrium starts at the same inflation rate below $y^{\#}$, whatever the bound (the small differences in the output come from the grid in `T_operator`).

Without a bound on inflation, almost anything could be supported, but only for a patient government.

99.6 Examples of recursive equilibria

We now construct particular equilibria by guess-and-verify, in the order in which the literature found them.

99.6.1 Infinite repetition of Nash

The easiest equilibrium repeats the one-period Nash outcome forever.

Take $v_1 = v^N = r(y^N, y^N)$, $\sigma_1(v) = y^N$ for every v, and $\sigma_2(v, x, y) = v^N$ for every (v, x, y).

Condition (a) holds by construction, and condition (b) collapses to $r(y^N, y^N) \geq r(y^N, B(y^N))$, which holds *with equality* because y^N is a fixed point of the best response map.

Nothing deters the government here, and nothing needs to: it is already doing what it most wants to do.

This is the Markov perfect equilibrium of the repeated game.

The economy has no payoff-relevant state variable, so a Markov strategy cannot condition on anything, and the government simply plays its one-period best response every period (*Markov Perfect Equilibrium* studies Markov perfect equilibria of dynamic games that do have states).

99.6.2 Infinite repetition of something better

Let $v^b = r(y^b, y^b) > v^N$ be a better-than-Nash value, and suppose

$$r(y^b, B(y^b)) - r(y^b, y^b) \;\leq\; \frac{\delta}{1-\delta}\,(v^b - v^N)\,. \tag{99.12}$$

The left side is the one-period gain from cheating; the right side is the discounted loss from reverting to Nash forever.

When (99.12) holds, the strategy that prescribes y^b while the promised value is v^b and reverts to Nash after any deviation is an SPE.

Barro and Gordon [1983] and Stokey [1989] studied the case $y^b = y^R$: anticipated reversion to Nash supports Ramsey forever.

Whether it does depends on patience, and for this economy the threshold has a remarkably clean form.

Setting $y^b = 0$ in (99.12) and using the closed forms (99.4), every U^* and every scale factor cancels, leaving

$$\delta \;\geq\; \delta^\star = \frac{1}{2+\theta^2}. \tag{99.13}$$

```
def supports_ramsey_by_nash(m):
    "Does reversion to Nash sustain Ramsey forever?"
    gain = r_cheat(m, 0.0) - r_keep(m, 0.0)
    loss = m.δ / (1 - m.δ) * (r_keep(m, 0.0) - r_keep(m, y_nash(m)))
    return gain <= loss


δ_star = 1 / (2 + m.θ**2)
print(f"closed-form cutoff  δ* = 1/(2 + θ²) = {δ_star:.6f}\n")
for δ in (0.20, 0.27, δ_star - 1e-6, δ_star + 1e-6, 0.50, 0.95):
    md = m._replace(δ=δ)
    print(f"  δ = {δ:.8f}: Nash reversion supports Ramsey? "
          f"{supports_ramsey_by_nash(md)}")
```

```
closed-form cutoff  δ* = 1/(2 + θ²) = 0.280702

  δ = 0.20000000: Nash reversion supports Ramsey? False
  δ = 0.27000000: Nash reversion supports Ramsey? False
  δ = 0.28070075: Nash reversion supports Ramsey? False
  δ = 0.28070275: Nash reversion supports Ramsey? True
  δ = 0.50000000: Nash reversion supports Ramsey? True
  δ = 0.95000000: Nash reversion supports Ramsey? True
```

Above $\delta^\star \approx 0.28$ the threat of reverting to Nash is enough to hold the government at zero inflation; below it, the threat is too weak.

99.6.3 Something worse: Abreu's stick and carrot

When reversion to Nash is not a strong enough threat, a *worse* equilibrium might serve as the punishment.

Abreu [1986], studying oligopoly, and Stokey [1991], studying credible public policy, showed how to construct one; Abreu [1988] developed the general theory of such optimal punishments.

The device is a stick-and-carrot strategy: the government is required to inflate at the maximum rate $y^\#$ for one period — the stick — after which it is rewarded with the Ramsey outcome forever.

The value is

$$\tilde{v} = (1-\delta)\, r(y^\#, y^\#) + \delta\, v^R, \tag{99.14}$$

and a deviation at any date, in either phase, *restarts* the whole strategy, so the continuation value after a deviation is always $\tilde{v}$ itself.

The strategy is an SPE only if neither phase tempts the government.

The stick. In the first period the incentive constraint is

$$\tilde{v} = (1-\delta)\, r(y^\#, y^\#) + \delta\, v^R \;\geq\; (1-\delta)\, r(y^\#, B(y^\#)) + \delta\, \tilde{v}.$$

Moving $\delta\tilde{v}$ to the left and dividing by $1-\delta$ leaves $\tilde{v} \geq r(y^\#, B(y^\#))$: the stick must be worth at least the candidate worst value (99.9).

The carrot. Once at Ramsey, the government must prefer staying there to grabbing the one-period gain from $B(0)$ and restarting the stick:

$$v^R \;\geq\; (1-\delta)\, r(0, B(0)) + \delta\, \tilde{v}. \tag{99.15}$$

Substituting (99.14), and using (99.5) and $v^R - r(y^\#, y^\#) = y^{\#2}/2$, the two constraints become

$$\delta \;\geq\; \frac{\left(y^\# - y^N\right)^2}{(1+\theta^2)\, y^{\#2}} \quad \text{(stick)}, \qquad \delta \;\geq\; \delta_A = \frac{\theta^2 U^{*2}}{(1+\theta^2)\, y^{\#2}} \quad \text{(carrot)}.$$

Here $y^N > y^\# - y^N$, so the carrot is the constraint that binds.

```
def abreu_value(m):
    "Value of the stick-and-carrot strategy: one period at y^#, then Ramsey."
    return (1 - m.δ) * r_keep(m, m.y_max) + m.δ * r_keep(m, 0.0)


def supports_ramsey_with_threat(m, threat):
    "Does a permanent threat value sustain Ramsey forever?"
```

(continues on next page)

(continued from previous page)

```
    return r_keep(m, 0.0) >= (1 - m.δ) * r_cheat(m, 0.0) + m.δ * threat


def is_abreu_spe(m):
    """
    Check both incentive constraints of the stick-and-carrot strategy, in which
    any deviation restarts the strategy.

    Returns (stick constraint holds, carrot constraint holds, both hold).
    """
    v_tilde = abreu_value(m)
    stick = v_tilde >= (1 - m.δ) * r_cheat(m, m.y_max) + m.δ * v_tilde
    carrot = supports_ramsey_with_threat(m, v_tilde)        # cp_abreu_carrot
    return stick, carrot, stick and carrot


δ_stick = (m.y_max - y_nash(m))**2 / ((1 + m.θ**2) * m.y_max**2)
δ_A = m.θ**2 * m.U_star**2 / ((1 + m.θ**2) * m.y_max**2)
print(f"stick holds for δ >= {δ_stick:.4f};  carrot holds for δ >= δ_A = {δ_A:.4f}\n")
for δ in (0.10, 0.20, 0.95):
    md = m._replace(δ=δ)
    va, vn = abreu_value(md), r_keep(md, y_nash(md))
    stick, carrot, spe = is_abreu_spe(md)
    print(f"δ = {δ:4.2f}:  v_abreu = {va:8.4f}   v^N = {vn:8.4f}   "
          f"worse than Nash? {str(va < vn):5}   stick? {str(stick):5}   "
          f"carrot? {str(carrot):5}   SPE? {spe}")
```

```
stick holds for δ >= 0.0381;  carrot holds for δ >= δ_A = 0.1845

δ = 0.10:  v_abreu = -60.1250   v^N = -38.7578   worse than Nash? True    stick?␣
 ↪True    carrot? False   SPE? False
δ = 0.20:  v_abreu = -55.1250   v^N = -38.7578   worse than Nash? True    stick?␣
 ↪True    carrot? True    SPE? True
δ = 0.95:  v_abreu = -17.6250   v^N = -38.7578   worse than Nash? False   stick?␣
 ↪True    carrot? True    SPE? True
```

At $\delta = 0.95$ the stick-and-carrot value is *better* than Nash, because a single bad period is heavily outweighed by an eternity of Ramsey.

At $\delta = 0.2$ it is far *worse* than Nash — and that is the point.

At $\delta = 0.1$ the government would still accept the stick, but it would not stay at Ramsey: the one-period gain from inflating outweighs the discounted cost of restarting the stick.

So at $\delta = 0.1$ the strategy is not an equilibrium, even though its first-period constraint holds.

The carrot constraint (99.15) asks whether the threat $\tilde{v}$ sustains Ramsey forever.

A punishment worse than Nash is a stronger deterrent than reversion to Nash, so it can support outcomes that Nash reversion cannot.

```
m2 = m._replace(δ=0.20)

print(f"at δ = 0.2:")
print(f"  threat = Nash     ({r_keep(m2, y_nash(m2)):8.4f}): "
      f"supports Ramsey? {supports_ramsey_with_threat(m2, r_keep(m2, y_nash(m2)))}")
print(f"  threat = Abreu    ({abreu_value(m2):8.4f}): "
      f"supports Ramsey? {supports_ramsey_with_threat(m2, abreu_value(m2))}")
```

```
at δ = 0.2:
  threat = Nash    (-38.7578): supports Ramsey? False
  threat = Abreu   (-55.1250): supports Ramsey? True
```

At $\delta = 0.2$, an impatient government cannot be held at Ramsey by the prospect of Nash, but *can* be held there by the prospect of Abreu's stick.

The worst equilibrium value $\underline{v}$ is a harsher punishment still, and it holds Ramsey in place down to an even lower discount factor (exercise Exercise 99.10.4).

Two definitions name the structures we have just met.

Definition 99.6.1 (Self-enforcing and self-rewarding)

A recursive SPE is **self-enforcing** if the continuation value following a deviation equals the initial promised value, and **self-rewarding** if the continuation value following adherence equals the promised value.

Abreu's stick-and-carrot is self-enforcing; infinite repetition of a better-than-Nash outcome is self-rewarding.

Proposition 99.5.1 and *Proposition 99.5.2* say that these two structures are not curiosities: they are exactly what the extreme equilibria look like.

99.7 Multiplicity

Two layers of multiplicity inhabit this theory.

There is a continuum of equilibrium *values*, and — the sharper point — many different outcome *paths* attain the same value.

To see the second, we construct three equilibria that all deliver the worst value $\underline{v}$.

All three begin the same way: the first-period promised value is $\underline{v}$, the prescribed action is $y^{\#}$, and the continuation value v_2 solves $\underline{v} = (1-\delta) r(y^{\#}, y^{\#}) + \delta v_2$.

They differ in what they do afterwards.

The code reuses `next_value` from the worst-value section.

```
def action_delivering(m, v):
    "The inflation rate ȳ with r(ȳ, ȳ) = v."
    return np.sqrt(max(-2.0 * v - m.U_star**2, 0.0))


def method_1(m, v_lo, v_hi, max_steps=500):
    "Climb at y^# until one more step would overshoot v_hi, then settle at v_hi."
    v, y = [v_lo], []
    for _ in range(max_steps):
        nxt = next_value(m, v[-1])
        if nxt > v_hi:
            break
        y.append(m.y_max)
        v.append(nxt)
    target = (v[-1] - m.δ * v_hi) / (1 - m.δ)      # r(ȳ, ȳ) for the switching period
    y.append(action_delivering(m, target))
    v.append(v_hi)
    y.append(action_delivering(m, v_hi))           # stay at v_hi forever after
```

(continues on next page)

(continued from previous page)

```
    return np.array(v), np.array(y)


def method_2(m, v_lo, max_steps=500):
    "Climb at y^# only until the promised value first exceeds v^N, then freeze."
    v_N = r_keep(m, y_nash(m))
    v, y = [v_lo], []
    for _ in range(max_steps):
        nxt = next_value(m, v[-1])
        y.append(m.y_max)
        v.append(nxt)
        if nxt > v_N:
            break
    y.append(action_delivering(m, v[-1]))          # freeze at v** forever
    return np.array(v), np.array(y)


def method_3(m, v_lo):
    "One period at y^#, then freeze immediately at v_2."
    v_2 = next_value(m, v_lo)
    return (np.array([v_lo, v_2, v_2]),
            np.array([m.y_max, action_delivering(m, v_2),
                      action_delivering(m, v_2)]))
```

The figure plots the promised values and inflation rates along the three paths.

```
paths = {'method 1': method_1(m, v_lo, v_hi),
         'method 2': method_2(m, v_lo),
         'method 3': method_3(m, v_lo)}

fig, axes = plt.subplots(2, 3, figsize=(13, 6.5), sharex='col')
for k, (name, (v, y)) in enumerate(paths.items()):
    axes[0, k].plot(range(1, len(v) + 1), v, 'C0o-', ms=3.5, lw=1)
    axes[0, k].axhline(v_lo, color='k', ls='--', lw=0.8)
    axes[0, k].axhline(v_hi, color='C2', ls=':', lw=0.8)
    axes[0, k].set_title(f'{name}: continuation values', fontsize=10)
    axes[1, k].plot(range(1, len(y) + 1), y, 'C1o-', ms=3.5, lw=1)
    axes[1, k].set_ylim(-0.4, m.y_max + 0.4)
    axes[1, k].set_xlabel('$t$')
    axes[1, k].set_title(f'{name}: inflation', fontsize=10)
axes[0, 0].set_ylabel('promised value $v_t$')
axes[1, 0].set_ylabel('inflation $y_t$')
plt.tight_layout()
plt.show()
```

Here are their lengths and terminal values.

```
print(f"Nash value {r_keep(m, y_nash(m)):.4f} at inflation {y_nash(m):.4f};  "
      f"worst value {v_lo:.4f}\n")
for name, (v, y) in paths.items():
    print(f"{name:9s}: {len(y):3d} periods before settling,  "
          f"first-period value {v[0]:.4f},  "
          f"terminal value {v[-1]:8.4f} at inflation {y[-1]:.4f}")
```

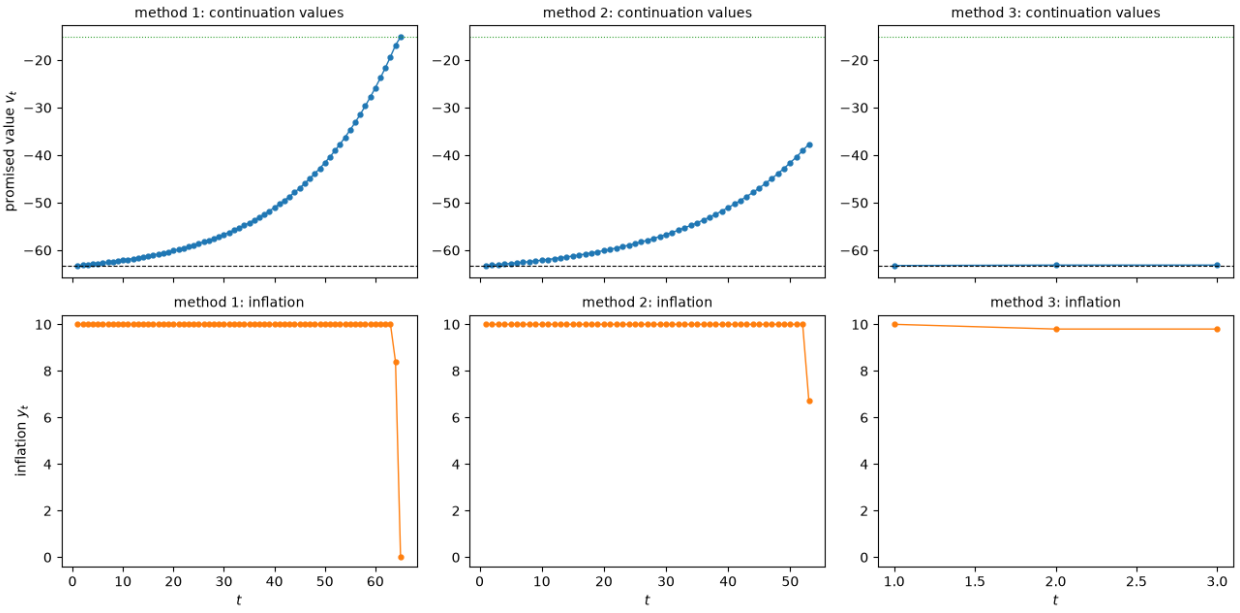


Fig. 99.2: Three subgame perfect equilibria that attain the same worst value (the book's Figures 4.2-4.4)

```
Nash value -38.7578 at inflation 6.8750;  worst value -63.2195

method 1 :  65 periods before settling,  first-period value -63.2195,  terminal␣
 ↪value -15.1250 at inflation 0.0000
method 2 :  53 periods before settling,  first-period value -63.2195,  terminal␣
 ↪value -37.6853 at inflation 6.7172
method 3 :   3 periods before settling,  first-period value -63.2195,  terminal␣
 ↪value -63.1192 at inflation 9.7974
```

The three paths could hardly be less alike.

The first climbs at maximum inflation for about sixty periods before easing down to Ramsey.

The second climbs for a shorter spell and then freezes forever at a value just *better* than Nash, sustained by an inflation rate just below the Nash rate.

The third inflates at the maximum for a single period and then settles immediately, at a value barely above the worst one.

All three are subgame perfect, and all three deliver exactly the same value to the government.

The equilibrium concept has nothing to say about which of them describes the world.

99.8 Numerical example

Here is the numerical example reported in chapter 4 of Sargent [1999], computed from scratch.

```
rows = [
    ("$\\theta$", m.θ, ""),
    ("$U^*$", m.U_star, ""),
    ("$y^\\#$", m.y_max, ""),
    ("$\\delta$", m.δ, ""),
    ("$y^N$ (Nash inflation)", y_nash(m), "6.8750"),
```

(continues on next page)

(continued from previous page)

```
    ("$y^R$ (Ramsey inflation)", 0.0, "0"),
    ("$v^R$", r_keep(m, 0.0), "-15.1250"),
    ("$v^N$", r_keep(m, y_nash(m)), "-38.7578"),
    ("$\\underline v$", v_lo, "-63.2195"),
    ("$v_{\\rm abreu}$", abreu_value(m), "-17.6250"),
]
pd.DataFrame(rows, columns=["object", "computed", "reported in the book"]).set_index(
 ↪"object")
```

object	computed	reported in the book
θ	1.250000	
U^*	5.500000	
$y^\#$	10.000000	
δ	0.950000	
y^N (Nash inflation)	6.875000	6.8750
y^R (Ramsey inflation)	0.000000	0
v^R	-15.125000	-15.1250
v^N	-38.757812	-38.7578
$\underline v$	-63.219512	-63.2195
$v_{\rm abreu}$	-17.625000	-17.6250

The book also reports the Nash-reversion cutoff and the value of the stick and carrot at $\delta = 0.2$.

```
print(f"cutoff discount factor δ*        : {1 / (2 + m.θ**2):.4f}  (book reports 0.
 ↪2807)")
print(f"Abreu stick-and-carrot at δ = 0.2: {abreu_value(m._replace(δ=0.2)):.4f}"
      f"  (book reports -55.125)")
```

```
cutoff discount factor δ*        : 0.2807  (book reports 0.2807)
Abreu stick-and-carrot at δ = 0.2: -55.1250  (book reports -55.125)
```

Every entry reproduces the book.

99.9 Interpretations

The literature on credible plans bears mixed news for the *triumph of natural-rate theory* story of *The Rise and Fall of U.S. Inflation*.

It certainly rescues the government from the pessimism of *The Credibility Problem*: better outcomes than Nash are available, and at plausible discount factors Ramsey itself is an equilibrium.

But it rescues too much.

Values worse than Nash are equilibria too, and between the extremes lies a continuum.

The multitude of outcomes mutes the model empirically, and it undoes the very thing early rational-expectations researchers wanted from the hypothesis — the elimination of free parameters describing expectations.

99.9.1 Whose expectations are they?

In his 1979 review of an OECD report edited by Paul McCracken, Lucas protested the report's recommendation that "governments should try to promote good expectations," as though expectations were an extra set of policy instruments.

In 1979 that protest was well aimed: rational-expectations models then took government policy as exogenous and made expectations a *function* of it, linked by cross-equation restrictions.

The theory of credible policy changes the picture in a way that vindicates neither side cleanly.

It turns systems of expectations into free parameters that influence outcomes — but into parameters that nobody inside the model gets to choose.

The government complies with equilibrium expectations about its own behavior.

Within an equilibrium, the government's strategy is simultaneously a decision rule and a description of the public's expectations, and the two cannot be disentangled.

As Sargent puts it: the authors of the McCracken report believed in multiplicity and manipulation, while Lucas doubted both; the literature on credible plans supports multiplicity but not manipulation.

99.9.2 Remedies

Reputation alone is a weak foundation for anti-inflation policy, and that weakness has inspired proposals to change the game rather than to hope for a good equilibrium within it.

Rogoff [1985] proposed delegating monetary policy to someone who cares less about unemployment than society does.

Assigning it to someone unaware even of a *temporary* tradeoff would be equally effective, and Alan Blinder later suggested a related device: an authority who knows the natural rate and never wants unemployment to differ from it.

Exercise Exercise **??** works out both remedies in the one-period economy of *The Credibility Problem*, including the price Rogoff's remedy pays when there are shocks to stabilize.

Maintaining a pool of potential central bankers with different inflation–unemployment preferences, some like Rogoff's conservative and others like the authority in Kydland and Prescott's model, can also improve outcomes [Ball, 1995, Barro and Gordon, 1983].

When the public does not know which kind of banker is in office, even a government of the inflationary kind can gain by acting like the conservative kind and building a reputation; Backus and Driffill [1985] and Barro [1986] model this as a game of incomplete information.

99.9.3 Where this leaves us

Every remedy just listed changes the *institutions*, not the theory.

Within the theory, the multiplicity is irreducible, and Sargent locates its source precisely: it stems from the rationality imputed to *everyone* in the system.

Perfection is what generates the continuum, because a system of expectations sophisticated enough to support one equilibrium is sophisticated enough to support many.

So the book retreats from perfection.

The remaining lectures replace fully rational participants with ones whose understanding of the economy is limited.

First the public forecasts adaptively, in *Adaptive Expectations and the Phelps Problem*, which comes next.

Then agents fit misspecified models to the data their own beliefs generate, in *Optimal Misspecified Beliefs* and *Self-Confirming Equilibria*; in the self-confirming equilibria of the latter, the misspecified agent is the government.

Finally, from *Adaptive Learning and Escape Dynamics* on, the government re-estimates its model in real time.

Those models are closer to what Lucas had in mind when he criticized the McCracken report, and, unlike the theory of this lecture, they make predictions.

99.10 Exercises

Exercise 99.10.1

The cutoff (99.13) claims that reversion to Nash sustains Ramsey forever if and only if $\delta \geq 1/(2+\theta^2)$ — a threshold that depends on θ alone, and not on the natural rate U^* or the upper bound $y^\#$.

Verify both claims numerically.

For a grid of θ values, find the smallest δ on a fine grid at which reversion to Nash supports Ramsey, and compare it with $1/(2+\theta^2)$.

Then check that changing U^* leaves the answer untouched.

Solution

```
δ_grid = np.linspace(0.01, 0.99, 9801)

rows = []
for θ in (0.5, 1.0, 1.25, 2.0, 3.0):
    for U_star in (5.5, 11.0):
        md = Model(θ=θ, U_star=U_star)
        ok = [δ for δ in δ_grid
              if supports_ramsey_by_nash(md._replace(δ=δ))]
        rows.append((θ, U_star, min(ok), 1 / (2 + θ**2)))

pd.DataFrame(rows, columns=["$\\theta$", "$U^*$", "smallest $\\delta$ found",
                            "$1/(2+\\theta^2)$"]).round(4)
```

	θ	U^*	smallest δ found	$1/(2+\theta^2)$
0	0.50	5.5	0.4445	0.4444
1	0.50	11.0	0.4445	0.4444
2	1.00	5.5	0.3334	0.3333
3	1.00	11.0	0.3334	0.3333
4	1.25	5.5	0.2808	0.2807
5	1.25	11.0	0.2808	0.2807
6	2.00	5.5	0.1667	0.1667
7	2.00	11.0	0.1667	0.1667
8	3.00	5.5	0.0910	0.0909
9	3.00	11.0	0.0910	0.0909

The grid search matches the closed form to within the grid spacing, and doubling U^* changes nothing.

The reason both U^* and $y^\#$ drop out is visible in (99.4): the one-period temptation at $y = 0$ and the Nash–Ramsey value gap are both proportional to U^{*2}, so the scale cancels from the inequality, and $y^\#$ never enters because neither side involves the upper bound.

A larger unemployment response θ to surprise inflation makes the Nash outcome worse relative to Ramsey, which strengthens the threat and lets a *less* patient government be held at zero inflation.

(In the (U, y) plane of *The Credibility Problem* the Phillips curve has slope $-1/\theta$, so a larger θ means a *flatter* curve.)

Exercise 99.10.2

The set of equilibrium values shrinks as the government becomes impatient.

Use `solve_aps` to compute $V = [\underline{v}, \overline{v}]$ for discount factors between 0.001 and 0.98, and plot both endpoints against δ on a logarithmic axis, marking the Nash and Ramsey values.

1. Below what discount factor δ_R does the best equilibrium value stop being Ramsey? Find δ_R exactly, by asking when the punishment $r(y^\#, B(y^\#))$ makes the government prefer Ramsey to the one-period gain from deviating, and confirm it with `solve_aps`.
2. At small δ the lower edge of V bends away from $r(y^\#, B(y^\#))$. Find where, and explain why, using the check of the worst-value section.
3. What happens to V as $\delta \to 0$, and why is that the answer you should expect?

Solution

Let $G = r(0, B(0)) - v^R$ be the one-period gain from deviating at Ramsey.

With $\underline{v}$ as a permanent punishment, Ramsey is sustained when $v^R \geq (1-\delta)(v^R + G) + \delta \underline{v}$, that is, when $\delta \geq \delta_R = G/(G + v^R - \underline{v})$.

For the bend, we find by bisection the discount factor δ_w at which the reward (99.10) that the candidate $r(y^\#, B(y^\#))$ requires just equals the best value it implies.

```
v_R, v_N = r_keep(m, 0.0), r_keep(m, y_nash(m))
v_cand = r_cheat(m, m.y_max)
G = r_cheat(m, 0.0) - v_R                    # one-period gain from deviating at␣
↪Ramsey
δ_R = G / (G + v_R - v_cand)


def check_gap(δ):
    "Reward required by the candidate worst value minus the best value it implies."
    md = m._replace(δ=δ)
    return next_value(md, v_cand) - best_value(md, v_cand)[0]


lo, hi = 0.01, 0.1                       # the check fails at lo and passes at hi
for _ in range(50):
    mid = (lo + hi) / 2
    lo, hi = (mid, hi) if check_gap(mid) > 0 else (lo, mid)
δ_w = hi

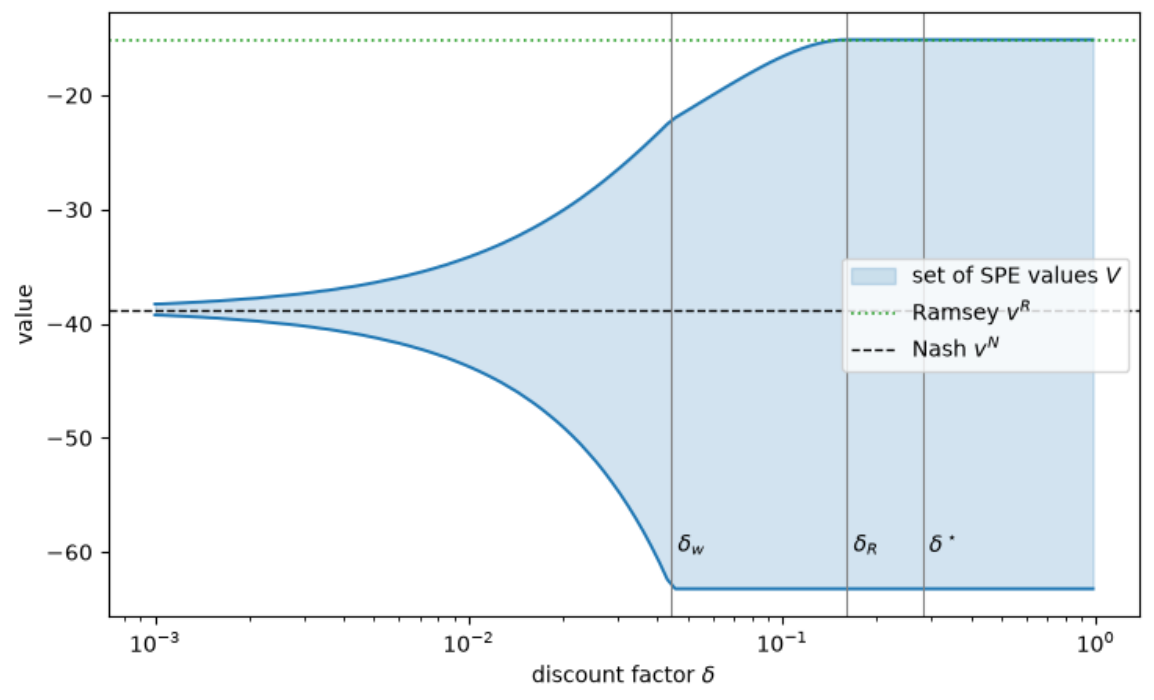
print(f"δ_R = {δ_R:.4f};  δ_w = {δ_w:.4f};  δ* = {1 / (2 + m.θ**2):.4f}")
for δ in (δ_R - 0.001, δ_R + 0.001):
    Wδ, _ = solve_aps(m._replace(δ=δ))
    print(f"  δ = {δ:.4f}: best value {Wδ[1]:.4f}, equals Ramsey? "
          f"{np.isclose(Wδ[1], v_R, atol=1e-6)}")
for δ in (0.001, 0.01, 0.02):
    Wδ, _ = solve_aps(m._replace(δ=δ))
    print(f"  δ = {δ:.3f}: V = [{Wδ[0]:.3f}, {Wδ[1]:.3f}], "
          f"width / δ = {(Wδ[1] - Wδ[0]) / δ:.0f}")
```

```
δ_R = 0.1609;  δ_w = 0.0443;  δ* = 0.2807
  δ = 0.1599: best value -15.1254, equals Ramsey? False
  δ = 0.1619: best value -15.1250, equals Ramsey? True
  δ = 0.001: V = [-39.241, -38.279], width / δ = 962
  δ = 0.010: V = [-43.738, -34.175], width / δ = 956
  δ = 0.020: V = [-49.038, -30.072], width / δ = 948
```

The figure marks the three discount factors.

```
δs = np.geomspace(0.001, 0.98, 120)
lows, highs = [], []
for δ in δs:
    Wδ, _ = solve_aps(m._replace(δ=δ))
    lows.append(Wδ[0])
    highs.append(Wδ[1])

fig, ax = plt.subplots(figsize=(8.5, 5))
ax.fill_between(δs, lows, highs, alpha=0.2, color='C0', label='set of SPE values $V
↪$')
ax.plot(δs, highs, 'C0', lw=1.4)
ax.plot(δs, lows, 'C0', lw=1.4)
ax.axhline(v_R, color='C2', ls=':', lw=1.2, label='Ramsey $v^R$')
ax.axhline(v_N, color='k', ls='--', lw=1, label='Nash $v^N$')
for δ_mark, name in [(δ_w, r'$\delta_w$'), (δ_R, r'$\delta_R$'),
                     (1 / (2 + m.θ**2), r'$\delta^\star$')]:
    ax.axvline(δ_mark, color='gray', lw=0.8)
    ax.annotate(name, (δ_mark, -60), textcoords='offset points', xytext=(3, 0))
ax.set_xscale('log')
ax.set_xlabel(r'discount factor $\delta$')
ax.set_ylabel('value')
ax.legend(loc='center right')
plt.show()
```



The best equilibrium value equals Ramsey for $\delta \geq \delta_R \approx 0.161$, well *below* the $\delta^\star \approx 0.28$ of (99.13).

The cutoff $\delta^\star$ uses reversion to *Nash* as the threat, while the APS set uses the worst equilibrium, a much harsher threat, so Ramsey survives at lower discount factors than Barro and Gordon's argument alone would suggest.

Below $\delta_w \approx 0.044$ the check of the worst-value section fails.

The reward (99.10) needed to make one period at $y^\#$ acceptable grows like $1/\delta$, and below δ_w it exceeds the best equilibrium value.

The worst equilibrium must then start at a lower inflation rate, and the lower edge of V rises above $r(y^\#, B(y^\#))$ toward v^N.

As $\delta \to 0$ the set collapses toward the single point v^N, with a width roughly proportional to δ.

That is what it must do: an entirely impatient government cares only about the current period, promises about the future carry no weight, and the only outcome that can be sustained is the one the government would choose myopically — the Nash outcome.

The sets are nested, growing with δ; Abreu *et al.* [1990] prove that, in their setting, the set of average equilibrium values is monotone in the discount factor.

Exercise 99.10.3

Method 3 in the multiplicity section is the boldest of the three: it inflates at $y^\#$ for a single period and then freezes forever at a constant inflation rate.

Because it settles immediately, its incentive constraint is the tightest of the three, so it is worth checking rather than assuming.

Verify directly that method 3 is a subgame perfect equilibrium: confirm that its first-period promise is honoured, that its frozen continuation value lies inside V, and that the government prefers adherence to deviation in *both* phases, using $\underline{v}$ as the punishment.

Solution

```
v_2 = next_value(m, v_lo)
y_tilde = action_delivering(m, v_2)

# phase 1: promised v_lo, prescribed action y^#, continuation v_2 on adherence
lhs_1 = (1 - m.δ) * r_keep(m, m.y_max) + m.δ * v_2
rhs_1 = (1 - m.δ) * r_cheat(m, m.y_max) + m.δ * v_lo

# phase 2: promised v_2, prescribed action ỹ, continuation v_2 on adherence
lhs_2 = (1 - m.δ) * r_keep(m, y_tilde) + m.δ * v_2
rhs_2 = (1 - m.δ) * r_cheat(m, y_tilde) + m.δ * v_lo

print(f"v_2 = {v_2:.4f},  ỹ = {y_tilde:.4f}")
print(f"v_2 lies inside V = [{v_lo:.4f}, {v_hi:.4f}]: "
      f"{v_lo <= v_2 <= v_hi}")
print(f"phase 1 delivers the promise: {np.isclose(lhs_1, v_lo)}")
print(f"phase 1 incentive: {lhs_1:.4f} vs {rhs_1:.4f}  ->  "
      f"holds with equality? {np.isclose(lhs_1, rhs_1)}")
print(f"phase 2 delivers the promise: {np.isclose(lhs_2, v_2)}")
print(f"phase 2 incentive: {lhs_2:.4f} >= {rhs_2:.4f}  ->  {lhs_2 >= rhs_2}")
print(f"slack in phase 2: {lhs_2 - rhs_2:.6f}")
```

```
v_2 = -63.1192,  ỹ = 9.7974
v_2 lies inside V = [-63.2195, -15.1250]: True
phase 1 delivers the promise: True
phase 1 incentive: -63.2195 vs -63.2195  ->  holds with equality? True
phase 2 delivers the promise: True
phase 2 incentive: -63.1192 >= -63.1312  ->  True
slack in phase 2: 0.011955
```

All the conditions hold, and the first-period promise is delivered exactly.

The phase-1 constraint holds with equality, and that is just the definition of $\underline{v}$: at $y^\#$ the government can always deviate to $B(y^\#)$ and accept the punishment $\underline{v}$, which is worth $(1-\delta)r(y^\#, B(y^\#)) + \delta\underline{v}$, and $\underline{v} = r(y^\#, B(y^\#))$ promises exactly that much.

The margin in the second phase is very thin, and that is not an accident either.

The frozen value v_2 is only barely above $\underline{v}$, so the punishment for deviating is only barely worse than the equilibrium itself — which is precisely what it means to be near the bottom of the set of equilibrium values.

Exercise 99.10.4

This exercise compares three ways of holding the government at Ramsey forever.

1. Suppose any deviation from Ramsey is followed by a continuation value v_p that does not depend on δ. Show that Ramsey is sustained if and only if $\delta \geq G/(G + v^R - v_p)$, where $G = r(0, B(0)) - v^R$.
2. Evaluate this threshold for Nash reversion ($v_p = v^N$) and for the worst equilibrium ($v_p = \underline{v} = r(y^\#, B(y^\#))$, the worst value for every δ above the δ_w of exercise Exercise 99.10.2).
3. Abreu's stick and carrot punishes with $\tilde{v}$, which depends on δ; its threshold is the δ_A of *Something worse: Abreu's stick and carrot*. Compare $\tilde{v}$ at δ_A with $\underline{v}$, and explain the difference.
4. Order the three thresholds, and check δ_A with `is_abreu_spe`.

Solution

1. Deviating from Ramsey earns $r(0, B(0)) = v^R + G$ for one period and v_p thereafter, so the constraint is $v^R \geq (1-\delta)(v^R + G) + \delta v_p$, which rearranges to $\delta(G + v^R - v_p) \geq G$.

2. By (99.5), $G = \theta^2 U^{*2}/(2(1+\theta^2))$, and $v^R - v^N = \theta^2 U^{*2}/2$, so Nash reversion gives $1/(2+\theta^2)$, the cutoff (99.13).

The worst equilibrium gives the δ_R of exercise Exercise 99.10.2.

```
def ramsey_threshold(m, v_p):
    "Smallest δ at which the permanent punishment v_p sustains Ramsey."
    G = r_cheat(m, 0.0) - r_keep(m, 0.0)
    return G / (G + r_keep(m, 0.0) - v_p)


δ_nash = ramsey_threshold(m, r_keep(m, y_nash(m)))
δ_opt = ramsey_threshold(m, r_cheat(m, m.y_max))
δ_A = m.θ**2 * m.U_star**2 / ((1 + m.θ**2) * m.y_max**2)
print(f"Nash reversion      δ* = {δ_nash:.4f}")
print(f"Abreu stick/carrot  δ_A = {δ_A:.4f}")
print(f"worst equilibrium   δ_R = {δ_opt:.4f}\n")
```

```
m_A = m._replace(δ=δ_A)
print(f"at δ_A: Abreu's punishment {abreu_value(m_A):.4f},  worst value "
      f"{r_cheat(m, m.y_max):.4f},  reward after the stick in the worst␣
 ↪equilibrium "
      f"{next_value(m_A, r_cheat(m, m.y_max)):.4f}\n")
for δ in (δ_A - 1e-4, δ_A + 1e-4):
    print(f"δ = {δ:.5f}: stick and carrot is an SPE? {is_abreu_spe(m._
 ↪replace(δ=δ))[2]}")
```

```
Nash reversion      δ* = 0.2807
Abreu stick/carrot  δ_A = 0.1845
worst equilibrium   δ_R = 0.1609

at δ_A: Abreu's punishment -55.9024,  worst value -63.2195,  reward after the␣
 ↪stick in the worst equilibrium -54.7944

δ = 0.18435: stick and carrot is an SPE? False
δ = 0.18455: stick and carrot is an SPE? True
```

3. Both punishments begin with a period at $y^{\#}$.

Abreu's then pays the full carrot, Ramsey forever.

The worst equilibrium pays a smaller reward after the stick, the v_1 of (99.10), which is itself worse than Nash, so its punishment is harsher.

4. The ordering is $\delta^\star > \delta_A > \delta_R$.

The harsher the punishment, the lower the discount factor at which it deters the one-period gain G.

Nash reversion is the mildest of the three punishments, the worst equilibrium the harshest, and Abreu's simple stick and carrot lies between them.

Exercise 99.10.5

For $\delta < \delta^\star$, reversion to Nash cannot sustain Ramsey.

The best value it can sustain is equation (25) of chapter 4 of Sargent [1999]:

$$v^{nr} = r(y^{nr}, y^{nr}) = (1-\delta)\, r(y^{nr}, B(y^{nr})) + \delta\, v^N,$$

where y^{nr} is the lowest inflation rate at which adhering is at least as good as deviating and reverting to Nash forever.

1. Using (99.5), show that for $\delta \leq \delta^\star$

$$y^{nr} = \theta U^* \frac{1 - \delta(2+\theta^2)}{1+\delta\theta^2}.$$

2. Plot v^{nr} against $\delta \in (0, \delta^\star]$, together with the best SPE value $\overline{v}$ from `solve_aps`, and compare the two at $\delta = 0.2$.

Solution

1. The constraint $r(y,y) \geq (1-\delta) r(y, B(y)) + \delta v^N$ is equivalent to $\delta\left[r(y,y) - v^N\right] \geq (1-$

$\delta)\,[r(y, B(y)) - r(y, y)]$.

Write $z = y^N - y \geq 0$ for the distance below Nash.

Then $r(y, y) - v^N = z(2y^N - z)/2$, and by (99.5) the right side is $(1-\delta)z^2/(2(1+\theta^2))$.

Dividing by $z > 0$, the constraint becomes $z \leq 2y^N\delta(1+\theta^2)/(1+\delta\theta^2)$.

The lowest admissible inflation rate is y^N minus that bound, which gives the formula.

It reaches zero exactly at $\delta = 1/(2+\theta^2) = \delta^\star$.

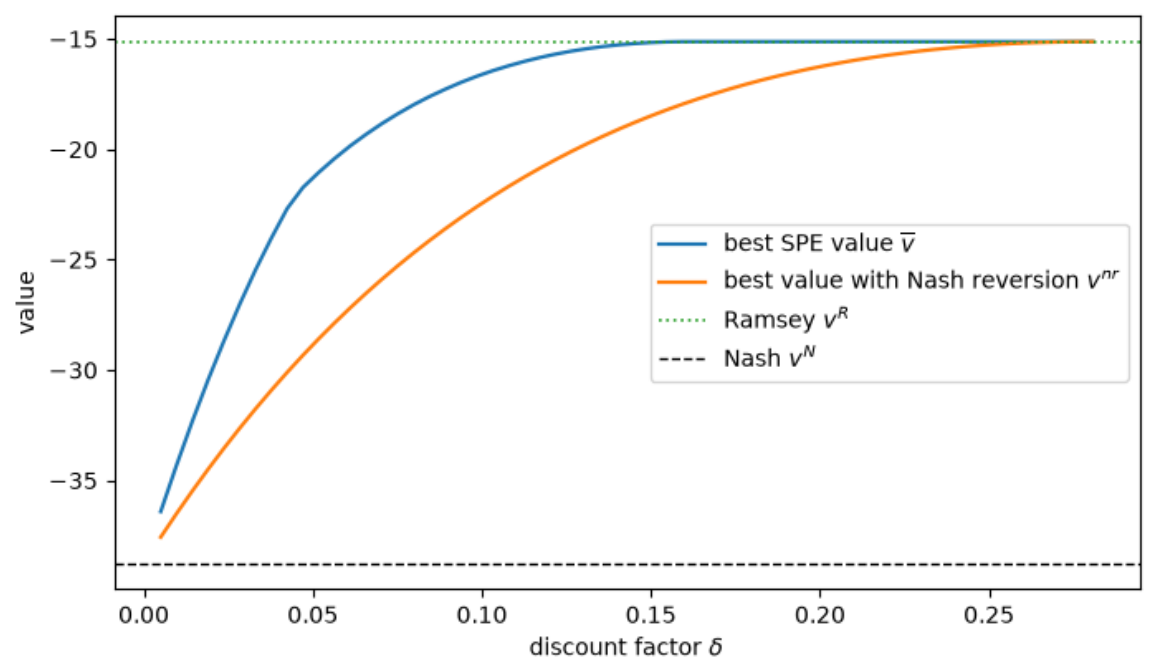
Every value between v^N and v^{nr} can also be sustained by Nash reversion, but no value above v^{nr}.

```
def nash_reversion_best(m):
    "The best value sustainable by reversion to Nash (book eq. (25)) and its␣
 ↪action."
    y_nr = max(y_nash(m) * (1 - m.δ * (2 + m.θ**2)) / (1 + m.δ * m.θ**2), 0.0)
    return r_keep(m, y_nr), y_nr


δ_star = 1 / (2 + m.θ**2)
δs_nr = np.linspace(0.005, δ_star, 60)
v_nr = [nash_reversion_best(m._replace(δ=δ))[0] for δ in δs_nr]
v_best = [solve_aps(m._replace(δ=δ))[0][1] for δ in δs_nr]

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(δs_nr, v_best, 'C0', lw=1.6, label=r'best SPE value $\overline{v}$')
ax.plot(δs_nr, v_nr, 'C1', lw=1.6, label=r'best value with Nash reversion $v^{nr}$
 ↪')
ax.axhline(r_keep(m, 0.0), color='C2', ls=':', lw=1.2, label='Ramsey $v^R$')
ax.axhline(r_keep(m, y_nash(m)), color='k', ls='--', lw=1, label='Nash $v^N$')
ax.set_xlabel(r'discount factor $\delta$')
ax.set_ylabel('value')
ax.legend()
plt.show()

m_02 = m._replace(δ=0.2)
v_nr_02, y_nr_02 = nash_reversion_best(m_02)
print(f"δ = 0.2: Nash reversion sustains at best {v_nr_02:.4f}, with y = {y_nr_02:.
 ↪4f};"
      f"  best SPE value {solve_aps(m_02)[0][1]:.4f}")
```




```
δ = 0.2: Nash reversion sustains at best -16.2589, with y = 1.5060;  best SPE␣
↪value -15.1250
```

At $\delta = 0.2$, Nash reversion can at best hold inflation down to about 1.5 percent, while the optimal punishment sustains Ramsey itself.

The two curves meet at $\delta^\star$, where Nash reversion first sustains Ramsey.

As $\delta \to 0$ both curves fall back to v^N, because the worst equilibrium value itself approaches v^N (exercise Exercise 99.10.2).

The gap closes only in levels: measured from v^N, the optimal punishment still sustains about twice the improvement that Nash reversion does.

```
v_N = r_keep(m, y_nash(m))
for δ in [0.02, 0.005]:
    m_δ = m._replace(δ=δ)
    ratio = (solve_aps(m_δ)[0][1] - v_N) / (nash_reversion_best(m_δ)[0] - v_N)
    print(f"δ = {δ}: (v̄ - v^N) / (v^nr - v^N) = {ratio:.2f}")
```

```
δ = 0.02: (v̄ - v^N) / (v^nr - v^N) = 1.95
δ = 0.005: (v̄ - v^N) / (v^nr - v^N) = 1.98
```

Exercise 99.10.6

The algorithm theorem of *The Abreu–Pearce–Stacchetti method* requires the starting set to satisfy $V \subseteq \mathcal{T}(W_0) \subseteq W_0$.

1. Show that $W_0 = \{v^N\}$ satisfies $\mathcal{T}(W_0) = W_0$. Confirm it with `T_operator` and `solve_aps`, and say which hypothesis of the theorem fails.
2. Start instead from $W_0 = \{v^R\}$. Where does the iteration go, and why?
3. What do the two experiments say about what a set of continuation values needs in order to support anything other than Nash?

Solution

1. With a single continuation value w, the reward and the punishment coincide, $w_1 = w_2 = w$, so admissibility requires $(1-\delta)r(y,y) \geq (1-\delta)r(y,B(y))$.

By (99.5) only $y = y^N$ qualifies, and it generates the single value $(1-\delta)v^N + \delta w$.

So $\mathcal{T}(\{w\}) = \{(1-\delta)v^N + \delta w\}$, and $\mathcal{T}(\{v^N\}) = \{v^N\}$.

The hypothesis that fails is $V \subseteq \mathcal{T}(W_0)$.

The Nash rate $y^N = 6.875$ is a point of the grid in `T_operator`, so the numerical operator finds it.

```
v_N, v_R = r_keep(m, y_nash(m)), r_keep(m, 0.0)
print(f"v^N = {v_N:.4f};  T({{v^N}}) = {T_operator(m, (v_N, v_N))}")
W_N, it_N = solve_aps(m, W0=(v_N, v_N))
print(f"iterating from {{v^N}}: stops after {it_N} iterations at {W_N}")
W_R, it_R = solve_aps(m, W0=(v_R, v_R))
print(f"iterating from {{v^R}}: stops after {it_R} iterations at "
      f"[{W_R[0]:.4f}, {W_R[1]:.4f}]")
```

```
v^N = -38.7578;  T({v^N}) = (np.float64(-38.7578125), np.float64(-38.7578125))
iterating from {v^N}: stops after 0 iterations at (np.float64(-38.7578125), np.
↪float64(-38.7578125))
iterating from {v^R}: stops after 542 iterations at [-38.7578, -38.7578]
```

2. Starting from $\{v^R\}$, the iterates are the single points $(1-\delta^j)\,v^N + \delta^j v^R$, which converge to $\{v^N\}$, not to V.

This start fails the other hypothesis, since $\mathcal{T}(\{v^R\}) \not\subseteq \{v^R\}$.

3. Incentives come from the *spread* between the reward for adhering and the punishment for deviating.

A single continuation value offers no spread, so the only outcome it can support is the one the government chooses without incentives, the static best response.

A starting set must be wide enough to contain every reward and every punishment that equilibria use, which is what $V \subseteq \mathcal{T}(W_0)$ guarantees, and what the set of all feasible values $[r(y^\#, y^\#), 0]$ provides.

CHAPTER

ADAPTIVE EXPECTATIONS AND THE PHELPS PROBLEM

Contents

- *Adaptive Expectations and the Phelps Problem*
  - *Overview*
  - *Adaptive expectations*
  - *The Phelps problem*
  - *Disinflation paths*
  - *The general Phelps problem*
  - *Testing the natural-rate hypothesis*
  - *Sacrifice ratios and the subversion of the Phelps model*
  - *Exercises*

In addition to what's in Anaconda, this lecture will use the following library:

```
!pip install quantecon
```

100.1 Overview

> They cannot look out far. They cannot look in deep. But when was that ever a bar To any watch they keep?
>
> – Robert Frost

This lecture continues the study of Phillips curve tradeoffs begun in *The Credibility Problem*.

It follows chapter 5 of Sargent [1999].

Credible Government Policies pursued better-than-Nash outcomes while keeping *everyone* rational, and found too many of them: a continuum of sustainable values, with nothing inside the theory to choose among them.

Here we retreat from that perfection in the smallest way we can, by keeping the government rational and giving the public a mechanical forecasting rule.

We describe

- the Cagan-Friedman adaptive expectations hypothesis,
- how Phelps [1967] used it to formulate a government control problem,
- the sacrifice ratio that the hypothesis implies, and

- the role of adaptive expectations in early econometric tests of the natural-rate hypothesis, and the rational expectations critique of those tests.

The key object is the **Phelps problem**: a government that is rational solves an optimal control problem while the public forecasts inflation with a fixed, mechanical adaptive rule.

Unlike the one-period model of *The Credibility Problem*, the government now takes into account that the economy lasts more than one period, and that today's inflation shapes tomorrow's expectations.

This intertemporal link can improve outcomes and, in a limiting case, even sustain the Ramsey outcome.

The Phelps problem is a linear-quadratic (LQ) control problem, so we solve it with the *LQ control* tools from QuantEcon.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
from scipy.signal import lfilter
```

100.2 Adaptive expectations

Phelps [1967] formulated a control problem for a natural-rate model.

He dropped rationality for the public but not for the government, and assigned the public a particular mechanical forecasting rule that is known to the government.

The public uses the adaptive expectations scheme of Milton Friedman and Cagan [1956] (see Monetarist Theory of Price Levels with Adaptive Expectations for Cagan's own use of it):

$$x_t - x_{t-1} = (1-\lambda)(y_{t-1} - x_{t-1}), \qquad \lambda \in (0,1), \tag{100.1}$$

where x_t is the public's expectation of inflation and y_t is inflation.

Rearranging, $x_t = \lambda x_{t-1} + (1-\lambda) y_{t-1}$, so expected inflation is a geometric distributed lag of past inflation,

$$x_t = (1-\lambda)\sum_{i=1}^{\infty} \lambda^{i-1} y_{t-i}. \tag{100.2}$$

Notice that (100.1) is a *constant gain* version of the least squares learning algorithm from *The Credibility Problem*, with the constant $(1-\lambda)$ playing the role that the decreasing gain t^{-1} played there.

It is also the relaxation algorithm of *From relaxation to learning: Bray's market*, run in calendar time with gain $1-\lambda$.

Note

Here λ is the weight on the old forecast, so the *gain* is $1-\lambda$.

The Bounded Rationality lectures write the gain itself as λ (*Notation across the series*), and *Optimal Misspecified Beliefs* and *Self-Confirming Equilibria* write it as $C = 1-\lambda$.

Equation (100.2) possesses what Cho and Matsui [1995] call an **induction property**: if the government keeps repeating a constant $y_t = \tilde{y}$ policy, eventually the public comes to set $x_t \approx \tilde{y}$.

The weights $(1-\lambda)\lambda^{i-1}$ sum to one, so a permanently maintained inflation rate is eventually expected.

Let's confirm the induction property numerically.

```
def adaptive_forecast(y, λ, x0=0.0):
    "Simulate x_t = λ x_{t-1} + (1-λ) y_{t-1}."
    T = len(y)
    x = np.empty(T)
    x[0] = x0
    for t in range(1, T):
        x[t] = λ * x[t - 1] + (1 - λ) * y[t - 1]
    return x

T = 60
y_const = np.full(T, 10.0)                  # a constant inflation policy

fig, ax = plt.subplots(figsize=(9, 4.5))
for λ in [0.7, 0.9]:
    x = adaptive_forecast(y_const, λ)
    ax.plot(x, lw=1.5, label=rf'$\lambda = {λ}$')
ax.axhline(10.0, color='k', ls='--', lw=1, label='policy $\\tilde y$')
ax.set_xlabel('$t$')
ax.set_ylabel('expected inflation $x_t$')
ax.legend()
plt.show()
```

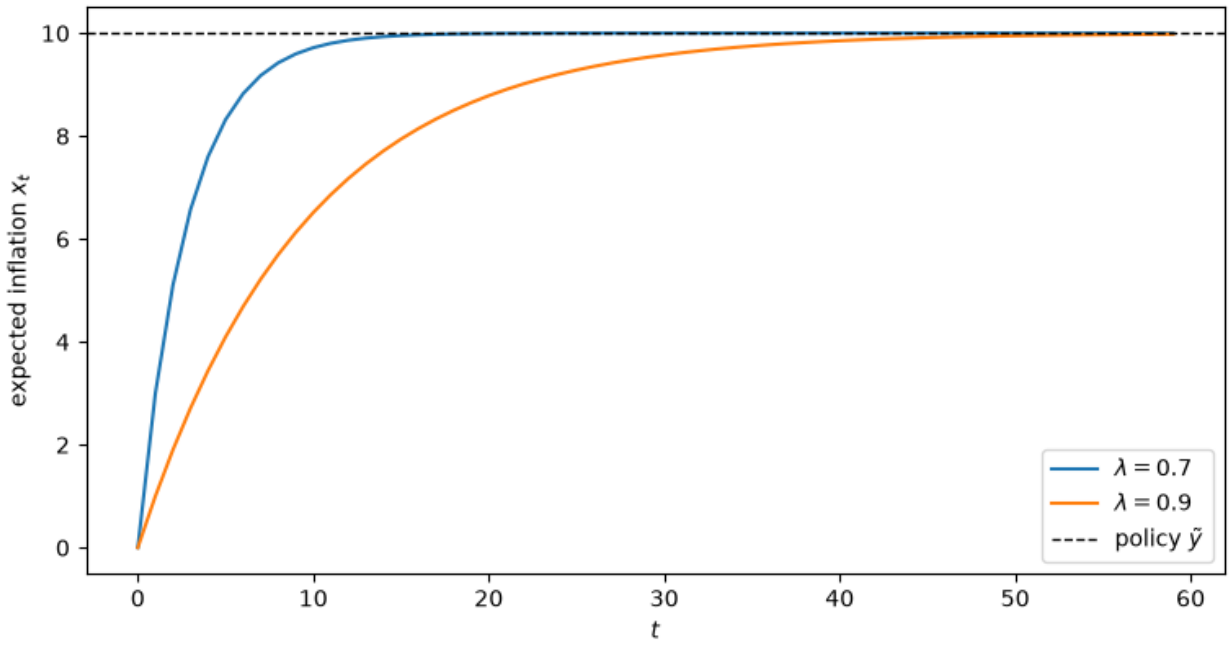


Fig. 100.1: Adaptive expectations converge to a constant inflation policy: the induction property

Under a constant policy the public's expectation converges to the policy, more slowly the larger is λ.

Solow and Tobin exploited this induction property when they tested the natural-rate hypothesis, as we discuss *below*.

100.3 The Phelps problem

The economy repeats forever, and the government evaluates outcome sequences with the discounted criterion

$$V^g = (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}p(U_t,y_t), \qquad p(U_t,y_t) = -\frac{1}{2}(U_t^2+y_t^2), \qquad \delta\in(0,1]. \tag{100.3}$$

When $\delta = 1$ we interpret (100.3) in the limit-of-means (Cesàro) sense.

Under that criterion only the limit of the outcome sequence matters, so every path that drives y_t to 0 is optimal and the transition path is not pinned down.

The $\delta = 1$ results below are therefore limits as $\delta \uparrow 1$: the code sets $\beta = 1 - 10^{-7}$, and the rule barely changes as that number shrinks further.

The government maximizes (100.3) by choice of a rule for inflation y_t, subject to the adaptive expectations scheme (100.1) and the expectations-augmented Phillips curve

$$U_t = U^* - \theta(y_t - x_t). \tag{100.4}$$

Because expected inflation x_t is predetermined at the start of period t, it is the only endogenous state variable.

The government's problem is therefore a discounted LQ control problem with

- state $s_t = \begin{bmatrix}1 & x_t\end{bmatrix}^\top$,
- control y_t, and
- transition $x_{t+1} = \lambda x_t + (1-\lambda)y_t$.

100.3.1 Casting the problem in LQ form

Write $U_t = a^\top s_t - \theta y_t$ with $a = \begin{bmatrix}U^* & \theta\end{bmatrix}^\top$.

The per-period loss $\frac{1}{2}(U_t^2+y_t^2)$ then equals

$$\frac{1}{2}\left[s_t^\top(aa^\top)s_t + (\theta^2+1)y_t^2 - 2\theta\, y_t\,(a^\top s_t)\right].$$

Matching this to the QuantEcon LQ loss $s_t^\top R s_t + y_t^\top Q y_t + 2y_t^\top N s_t$ (see *Eliminating Cross Products* for how such cross products are handled), and matching the transition to $s_{t+1} = As_t + By_t$, gives

$$R = \tfrac{1}{2}aa^\top, \quad Q = \tfrac{1}{2}(\theta^2+1), \quad N = -\tfrac{1}{2}\theta\, a^\top, \quad A = \begin{bmatrix}1 & 0\\ 0 & \lambda\end{bmatrix}, \quad B = \begin{bmatrix}0\\ 1-\lambda\end{bmatrix}.$$

The discount factor is $\beta = \delta$; the scaling $(1-\delta)$ in (100.3) does not affect the optimal policy.

```
class PhelpsProblem:
    """
    The Phelps optimal-control problem: a rational government facing a
    public that forecasts inflation adaptively with parameter λ.
    """

    def __init__(self, θ=1.0, U_star=5.0, λ=0.7, δ=0.96):
        self.θ, self.U_star, self.λ, self.δ = θ, U_star, λ, δ

        a = np.array([[U_star], [θ]])
```

(continues on next page)

(continued from previous page)

```
        R = 0.5 * (a @ a.T)
        Q = np.array([[0.5 * (θ**2 + 1)]])
        N = -0.5 * θ * a.T
        A = np.array([[1.0, 0.0], [0.0, λ]])
        B = np.array([[0.0], [1 - λ]])

        # δ = 1 (limit of means) is handled as the limit δ → 1
        β = min(δ, 1 - 1e-7)
        self.lq = qe.LQ(Q, R, A, B, N=N, beta=β)
        P, F, d = self.lq.stationary_values()

        # optimal rule y_t = f1 + f2 x_t
        self.f1, self.f2 = -F[0, 0], -F[0, 1]

    def simulate(self, x0=12.0, T=60):
        "Disinflation path (U_t, y_t) starting from expectation x0."
        θ, U_star, λ = self.θ, self.U_star, self.λ
        U, y, x = np.empty(T), np.empty(T), x0
        for t in range(T):
            y[t] = self.f1 + self.f2 * x
            U[t] = U_star - θ * (y[t] - x)
            x = λ * x + (1 - λ) * y[t]
        return U, y
```

The optimal rule takes the form $y_t = f_1 + f_2 x_t$ with $f_1 \neq 0$ and $f_2 \neq 1$.

These inequalities reflect that the public does not use an optimal forecasting rule; if instead $f_1 = 0$ and $f_2 = 1$, we would have $y_t = x_t$ for all histories.

```
pp = PhelpsProblem(θ=1.0, U_star=5.0, λ=0.7, δ=0.96)
print(f"optimal rule:  y_t = {pp.f1:.3f} + {pp.f2:.3f} x_t")
```

```
optimal rule:  y_t = 0.406 + 0.334 x_t
```

100.3.2 A proposition

The reason the Phelps problem is interesting is the following result.

Proposition 100.3.1 (δ = 1 eventually sustains Ramsey)

In the absence of discounting ($\delta = 1$), the government drives y_t to 0, the Ramsey outcome.

When $\delta = 1$, λ governs the speed of convergence to the Ramsey outcome.

When $\delta < 1$, the limit point of y_t depends on a comparison of λ with δ, which we can make exact.

Let q_t denote the marginal value to the government of expected inflation x_t.

The first-order condition for y_t and the envelope condition for x_t are

$$\theta U_t - y_t + \delta(1-\lambda) q_{t+1} = 0, \qquad q_t = -\theta U_t + \delta\lambda q_{t+1}.$$

In a steady state $x = y$, so $U = U^*$ and $q = -\theta U^*/(1-\delta\lambda)$, and the first-order condition gives

$$y_\infty = \theta U^* \frac{1-\delta}{1-\delta\lambda}. \tag{100.5}$$

The formula has a marginal interpretation.

A permanent one-point cut in inflation from the steady state lowers the inflation loss by y_∞ every period, a discounted benefit of $y_\infty/(1-\delta)$.

It also raises unemployment by $\theta\lambda^t$ in period t, while expectations catch up, and each point of unemployment costs about U^* at the margin, a discounted cost of $\theta U^*/(1-\delta\lambda)$.

Equating the two gives (100.5).

Because $1-\delta\lambda = (1-\delta)+\delta(1-\lambda)$, steady-state inflation is a fraction $(1-\delta)/[(1-\delta)+\delta(1-\lambda)]$ of the Nash rate θU^*.

It is close to the Ramsey value 0 only when the government's impatience $1-\delta$ is small relative to the public's gain $1-\lambda$.

That is what chapter 5 of Sargent [1999] means by $\lambda \ll \delta$ with δ close to one.

It is not enough that $\lambda < \delta$: with $\lambda = 0.98$ and $\delta = 0.99$, formula (100.5) gives $y_\infty \approx 0.34\,\theta U^*$.

The public's expectations are wrong along the transition path but are correct in the steady state, by virtue of the induction property.

The steady state of the rule $y_t = f_1 + f_2 x_t$ is $f_1/(1-f_2)$, so we can check (100.5) against the LQ solution.

```
for λ, δ in [(0.7, 0.96), (0.9, 0.96), (0.7, 0.99), (0.98, 0.99)]:
    pp = PhelpsProblem(λ=λ, δ=δ)
    y_lq = pp.f1 / (1 - pp.f2)
    y_cf = pp.θ * pp.U_star * (1 - δ) / (1 - δ * λ)
    print(f"λ = {λ}, δ = {δ}:  LQ steady state = {y_lq:.4f},  formula = {y_cf:.4f}")
```

```
λ = 0.7, δ = 0.96:  LQ steady state = 0.6098,  formula = 0.6098
λ = 0.9, δ = 0.96:  LQ steady state = 1.4706,  formula = 1.4706
λ = 0.7, δ = 0.99:  LQ steady state = 0.1629,  formula = 0.1629
λ = 0.98, δ = 0.99:  LQ steady state = 1.6779,  formula = 1.6779
```

100.4 Disinflation paths

We now reproduce the disinflation experiments of chapter 5 of Sargent [1999].

Set $\theta = 1$ and $U^* = 5$, and start the government's problem from late-1970s initial conditions $x_{-1} = y_{-1} = 12$, which imply $U = U^* = 5$.

The following tables record paths of unemployment U and inflation y at selected lags, for two discount factors $\delta \in \{0.96, 1\}$ and two adaptation parameters $\lambda \in \{0.7, 0.9\}$.

```
def disinflation_table(δ, lags=(1, 5, 20, 50)):
    rows = []
    for λ in [0.7, 0.9]:
        U, y = PhelpsProblem(λ=λ, δ=δ).simulate(x0=12.0, T=max(lags) + 1)
        for lag in lags:
            rows.append((λ, lag, U[lag - 1], y[lag - 1]))
    return rows

for δ in [0.96, 1.0]:
    print(f"\n δ = {δ}")
    print(f" {'λ':>4} {'lag':>4} {'U':>7} {'y':>7}")
    for λ, lag, U, y in disinflation_table(δ):
        print(f" {λ:>4} {lag:>4} {U:>7.2f} {y:>7.2f}")
```

```
δ = 0.96
   λ  lag       U       y
 0.7    1   12.59    4.41
 0.7    5    8.11    2.17
 0.7   20    5.11    0.66
 0.7   50    5.00    0.61
 0.9    1   11.87    5.13
 0.9    5   10.24    4.27
 0.9   20    6.91    2.49
 0.9   50    5.25    1.61

δ = 1.0
   λ  lag       U       y
 0.7    1   13.21    3.79
 0.7    5    8.28    1.51
 0.7   20    5.10    0.05
 0.7   50    5.00    0.00
 0.9    1   13.40    3.60
 0.9    5   11.28    2.70
 0.9   20    7.12    0.91
 0.9   50    5.24    0.10
```

For each parameter setting the government engineers a major recession and immediately brings inflation down more than halfway toward its eventual limiting value.

In the discounted case ($\delta = 0.96$), inflation settles at the positive level y_∞ of (100.5), and the government accepts a longer but milder recession when $\lambda = 0.9$ than when $\lambda = 0.7$.

In the undiscounted case ($\delta = 1$), inflation is driven all the way to the Ramsey value of zero, more slowly the larger is λ.

Let's plot the full paths, with the inherited inflation rate of 12 shown at $t = 0$.

```
T = 40
t = np.arange(T + 1)
fig, axes = plt.subplots(2, 2, figsize=(11, 7), sharex=True)

for j, δ in enumerate([0.96, 1.0]):
    for λ in [0.7, 0.9]:
        U, y = PhelpsProblem(λ=λ, δ=δ).simulate(x0=12.0, T=T)
        axes[0, j].plot(t, np.concatenate([[12.0], y]), lw=1.5,
                        label=rf'$\lambda = {λ}$')
        axes[1, j].plot(t, np.concatenate([[5.0], U]), lw=1.5,
                        label=rf'$\lambda = {λ}$')
    axes[0, j].set_title(rf'$\delta = {δ}$')
    axes[0, j].set_ylabel('inflation $y_t$')
    axes[1, j].axhline(5.0, color='k', ls=':', lw=1)
    axes[1, j].set_ylabel('unemployment $U_t$')
    axes[1, j].set_xlabel('period ($t = 0$: inherited $y = 12$)')
    axes[0, j].legend()

plt.tight_layout()
plt.show()
```

The undiscounted problem drives inflation to the Ramsey outcome; the discounted problem stops short of it.

In both, the government buys its immediate disinflation with a recession: unemployment jumps far above U^* in the first period and returns to U^* as expectations catch up.

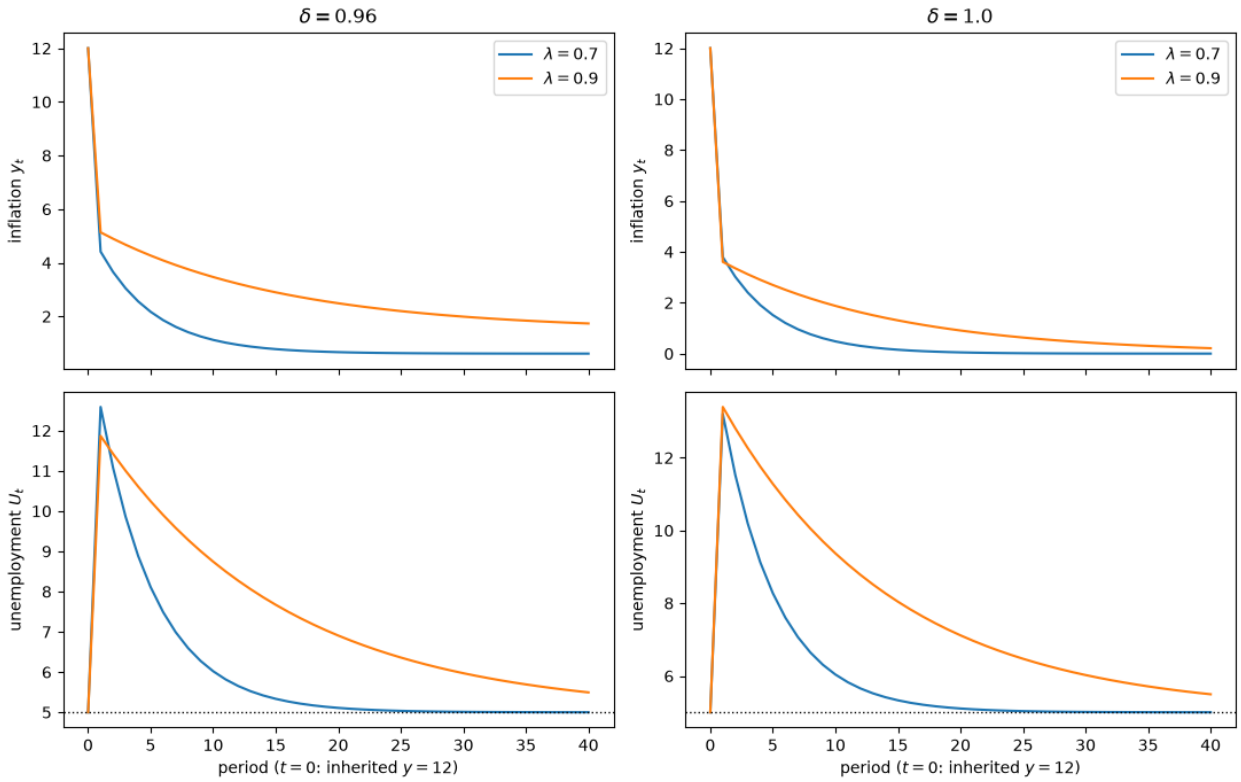


Fig. 100.2: Optimal disinflation paths from the Phelps problem, discounted and undiscounted

100.5 The general Phelps problem

For the government's control problem, what matters is the *reduced form* of the Phillips curve, not the underlying structure that identifies x_t.

It is useful to state a more general version of the Phelps problem in which the government's model is a reduced-form distributed lag Phillips curve.

Following Sargent [1999], define the vectors

$$X_{U,t-1} = \begin{bmatrix} U_{t-1} & \cdots & U_{t-m_U} \end{bmatrix}^\top, \qquad X_{y,t-1} = \begin{bmatrix} y_{t-1} & \cdots & y_{t-m_y} \end{bmatrix}^\top,$$

and $X_{t-1} = \begin{bmatrix} X_{U,t-1}^\top & X_{y,t-1}^\top & 1 \end{bmatrix}^\top$, which collects information dated $t-1$ and earlier.

X_{t-1} is the state vector of the general Phelps problem.

We can write two reduced-form Phillips curves that differ only in their *direction of fit*:

$$\text{Classical:} \quad U_t = \gamma^\top X_{C,t} + \varepsilon_{C,t}, \qquad X_{C,t} = \begin{bmatrix} y_t & X_{t-1}^\top \end{bmatrix}^\top,$$

$$\text{Keynesian:} \quad y_t = \beta^\top X_{K,t} + \varepsilon_{K,t}, \qquad X_{K,t} = \begin{bmatrix} U_t & X_{t-1}^\top \end{bmatrix}^\top.$$

The subscripts C and K stand for *Classical* (regress U on y) and *Keynesian* (regress y on U).

The general **Phelps problem** is to choose a control law $\hat{y}_t = hX_{t-1}$ to maximize the expected value of (100.3) subject to the government's believed Phillips curve and to $y_t = \hat{y}_t + v_{2t}$, where v_{2t} is a control error.

This induces a mapping from the government's beliefs γ to its decision rule h:

$$h = h(\gamma). \tag{100.6}$$

These objects — γ, β, $h(\gamma)$, and the two directions of fit — are exactly the ingredients we will need to define **self-confirming equilibria** in *Self-Confirming Equilibria*.

100.5.1 The adaptive case in general form

The specific problem solved above is the special case in which γ is restricted by substituting the adaptive expectations hypothesis (100.1) into the Phillips curve (100.4).

To find the restriction, use (100.4) at $t-1$ to write $x_{t-1} = y_{t-1} + (U_{t-1} - U^*)/\theta$, substitute it into (100.1), and substitute the result into (100.4):

$$U_t = (1-\lambda)U^* - \theta y_t + \lambda U_{t-1} + \theta y_{t-1}. \tag{100.7}$$

This is a classical Phillips curve with one lag of each variable ($m_U = m_y = 1$) and coefficients $\gamma = \begin{bmatrix} -\theta & \lambda & \theta & (1-\lambda)U^* \end{bmatrix}^\top$ on $X_{C,t} = \begin{bmatrix} y_t & U_{t-1} & y_{t-1} & 1 \end{bmatrix}^\top$.

Once that substitution is made, the state variable x_t disappears from view.

The induction hypothesis is visible in (100.7): the weights on current and lagged inflation sum to $-\theta + \theta = 0$.

Solved for y_t, (100.7) becomes the Keynesian curve $y_t = y_{t-1} + \theta^{-1}\left[(1-\lambda)U^* + \lambda U_{t-1} - U_t\right]$, with a weight of exactly one on lagged inflation.

Because $x_t = y_{t-1} + \lambda(U_{t-1} - U^*)/\theta$, the rule $y_t = f_1 + f_2 x_t$ of the original problem is, in terms of $X_{t-1} = \begin{bmatrix} U_{t-1} & y_{t-1} & 1 \end{bmatrix}^\top$,

$$h = \begin{bmatrix} f_2\lambda/\theta & f_2 & f_1 - f_2\lambda U^*/\theta \end{bmatrix}. \tag{100.8}$$

Let's confirm this by solving the general Phelps problem directly.

With state X_{t-1}, control y_t, and a believed curve $U_t = \gamma_y y_t + c^\top X_{t-1}$, where γ_y is the coefficient on y_t and c collects the rest, the problem is again an LQ problem.

Adaptive Learning and Escape Dynamics solves the same problem, with two lags of each variable, for a government that re-estimates γ every period.

```
def phelps_general(γ, δ):
    """
    Solve the general Phelps problem for the believed classical curve
    U_t = γ[0] y_t + γ[1] U_{t-1} + γ[2] y_{t-1} + γ[3].
    Returns h, with ŷ_t = h @ [U_{t-1}, y_{t-1}, 1].
    """
    γ_y = γ[0]
    c = np.asarray(γ[1:], dtype=float).reshape(-1, 1)
    R = 0.5 * (c @ c.T)
    Q = np.array([[0.5 * (γ_y**2 + 1)]])
    N = 0.5 * γ_y * c.T
    A = np.vstack([c.T, [0.0, 0.0, 0.0], [0.0, 0.0, 1.0]])  # next state (U_t, y_t, 1)
    B = np.array([[γ_y], [1.0], [0.0]])
    lq = qe.LQ(Q, R, A, B, N=N, beta=min(δ, 1 - 1e-7))
    P, F, d = lq.stationary_values()
    return -F[0]
```

(continues on next page)

(continued from previous page)

```
θ, U_star, λ, δ = 1.0, 5.0, 0.7, 0.96
y_adaptive = np.array([-θ, λ, θ, (1 - λ) * U_star])
pp = PhelpsProblem(θ=θ, U_star=U_star, λ=λ, δ=δ)

h_general = phelps_general(y_adaptive, δ)
h_from_x = np.array([pp.f2 * λ / θ, pp.f2, pp.f1 - pp.f2 * λ * U_star / θ])
print("h from the general problem:  ", h_general.round(4))
print("h from the x-state rule:     ", h_from_x.round(4))
```

```
h from the general problem:   [ 0.2338  0.334  -0.7628]
h from the x-state rule:      [ 0.2338  0.334  -0.7628]
```

The two agree: the original Phelps problem is the general one with γ restricted by (100.7).

Note

The **induction hypothesis** is the restriction that, in the Keynesian Phillips curve $y_t = \beta^\top X_{K,t} + \varepsilon_{K,t}$, the weights on lagged y's sum to unity (equivalently, in the classical form the weights on current and lagged y's sum to zero).

Under adaptive expectations this holds because the weights in (100.2) sum to one, as (100.7) shows.

100.6 Testing the natural-rate hypothesis

Robert Solow and James Tobin [Solow, 1968, Tobin, 1968] exploited the induction hypothesis to test the natural-rate hypothesis.

Substituting the geometric distributed lag (100.2) into an inverted Phillips curve gives

$$y_t = (1-\lambda)\sum_{i=1}^{\infty}\lambda^{i-1}y_{t-i} + \theta^{-1}(U^* - U_t). \tag{100.9}$$

They proposed to test the natural-rate hypothesis by running the regression

$$y_t = b_0 + b_1(1-\lambda)\sum_{i=1}^{\infty}\lambda^{i-1}y_{t-i} + b_2 U_t + \varepsilon_t, \tag{100.10}$$

and interpreting a finding of $b_1 < 1$ as evidence of a long-run tradeoff between inflation and unemployment.

In a steady state with constant inflation, (100.10) implies the long-run slope $dU/dy = (1-b_1)/b_2$, which equals $\theta(b_1-1)$ when b_2 takes its value $-1/\theta$ from (100.9).

Early implementations found $b_1 < 1$ and so rejected the natural-rate hypothesis in favor of a long-run tradeoff.

Later implementations of the same tests failed to reject it.

Lucas [1972] and Sargent [1971] argued that the rejections did not bear on the natural-rate hypothesis.

Solow and Tobin had deduced the restriction $b_1 = 1$ by using the induction hypothesis to identify expected inflation.

But the induction hypothesis describes one particular experiment, a permanent and fixed inflation rate, which may be very unlikely under the stochastic process that actually governs inflation in the sample.

Under rational expectations, the weights that link expected to past inflation are restricted by that process, and they need not sum to one.

For example, suppose that inflation is a stationary AR(1) process with root ρ, that $x_t = \mathbb{E}_{t-1} y_t$, and that the natural-rate Phillips curve holds.

Then the population regression (100.10) delivers

$$b_1 = \frac{\rho(1+\lambda)}{1+\lambda\rho} < 1, \tag{100.11}$$

a "rejection" of a hypothesis that is true (Exercise 100.8.2 derives the formula and checks it by simulation).

The estimate approaches one as $\rho \to 1$.

According to an argument of Sargent [1971] and King and Watson [1994], this is why the pattern of rejections and non-rejections is consistent with the tendency of inflation to have a unit root after the 1960s but not before: the unit-sum restriction $b_1 = 1$ is *compatible* with rational expectations when y_t has a unit root.

The self-confirming equilibria of *Self-Confirming Equilibria* give another illustration: rational expectations and the natural-rate hypothesis both hold there, and the induction hypothesis fails.

From the viewpoint of Phelps's control problem, it is incidental whether the natural-rate hypothesis holds: the Phelps problem imparts interesting dynamics to the inflation-unemployment choice whether or not $b_1 = 1$.

The public's expectations themselves disappear from the problem; what matters is the zero-frequency behavior of the distributed lag of y in (100.10).

100.7 Sacrifice ratios and the subversion of the Phelps model

Large *sacrifice ratios* — estimated amounts of foregone output required to bring inflation down one percentage point — circulated widely in the late 1970s.

The adaptive model delivers a sacrifice ratio in a simple exact form.

100.7.1 The sacrifice ratio

By (100.4) and (100.1),

$$U_t - U^* = -\theta(y_t - x_t) = -\frac{\theta}{1-\lambda}(x_{t+1} - x_t).$$

Summing over $t = 1, \dots, T$ telescopes the right side:

$$\sum_{t=1}^{T}(U_t - U^*) = \frac{\theta}{1-\lambda}\,(x_1 - x_{T+1}). \tag{100.12}$$

Cumulative excess unemployment per point of reduction in expected inflation is therefore $\theta/(1-\lambda)$, for *any* policy path and whatever the government's discount factor.

Because expected and actual inflation coincide in a steady state, this is also the cost of each point of permanent disinflation.

The government cannot change the ratio.

It can choose only how far to disinflate and how to spread the unemployment over time; with a quadratic loss it spreads it out.

With $\theta = 1$ the ratio is 3.33 when $\lambda = 0.7$ and 10 when $\lambda = 0.9$.

Let's check this along the optimal paths.

```
for λ in [0.7, 0.9]:
    for δ in [0.96, 1.0]:
        U, y = PhelpsProblem(λ=λ, δ=δ).simulate(x0=12.0, T=2000)
        ratio = np.sum(U - 5.0) / (12.0 - y[-1])
        print(f"λ = {λ}, δ = {δ}:  sacrifice ratio = {ratio:.3f},  "
              f"θ/(1-λ) = {1 / (1 - λ):.3f}")
```

```
λ = 0.7, δ = 0.96:  sacrifice ratio = 3.333,  θ/(1-λ) = 3.333
λ = 0.7, δ = 1.0:  sacrifice ratio = 3.333,  θ/(1-λ) = 3.333
λ = 0.9, δ = 0.96:  sacrifice ratio = 10.000,  θ/(1-λ) = 10.000
λ = 0.9, δ = 1.0:  sacrifice ratio = 10.000,  θ/(1-λ) = 10.000
```

Long expectations lags, λ near one, imply large sacrifice ratios.

100.7.2 The subversion of the Phelps model

Despite its encouraging implications for sustaining the Ramsey outcome under the induction hypothesis, Phelps's control problem carries a tattered past.

For a fixed $\delta < 1$, it is always possible to find a λ close enough to one that high inherited inflationary expectations make the government want to avoid the Ramsey outcome.

Formula (100.5) shows why: $y_\infty \to \theta U^*$, the Nash rate, as $\lambda \to 1$.

The next cell fixes $\delta = 0.96$ and lets λ approach one.

```
δ = 0.96
print(f"{'λ':>6} {'first-period y':>15} {'y_∞':>7} {'θ/(1-λ)':>9}")
for λ in [0.7, 0.9, 0.95, 0.99, 0.999]:
    pp = PhelpsProblem(λ=λ, δ=δ)
    y_first = pp.f1 + pp.f2 * 12.0
    y_inf = pp.θ * pp.U_star * (1 - δ) / (1 - δ * λ)
    print(f"{λ:>6} {y_first:>15.2f} {y_inf:>7.2f} {pp.θ / (1 - λ):>9.1f}")
```

```
     λ  first-period y     y_∞   θ/(1-λ)
   0.7            4.41    0.61       3.3
   0.9            5.13    1.47      10.0
  0.95            5.92    2.27      20.0
  0.99            7.64    4.03     100.0
 0.999            8.40    4.88    1000.0
```

The government still cuts inflation sharply in the first period, but it stops far short of zero, and each point of disinflation costs $\theta/(1-\lambda)$ points of unemployment.

In the late 1970s, models with long expectations-adjustment lags were used in this way to recommend *against* reducing inflation, fortifying those who advocated learning to live with high inflation.

By stressing the cross-equation restrictions of rational expectations, Lucas [1972] and Sargent [1971] argued that historical estimates of λ are silent about the optimal speed of disinflation: λ describes how expectations responded to past policy, not how they would respond to a new one.

The lesson to carry forward is that activating the induction hypothesis can eventually lead to better outcomes, but in the form that Phelps, Tobin, and Solow used it, the induction hypothesis retreats from rational expectations.

In *Optimal Misspecified Beliefs* and *Self-Confirming Equilibria* the adaptation parameter becomes an equilibrium outcome, $C = 1 - \lambda$, chosen by the public to fit the data that its own forecasts help generate.

From *Adaptive Learning and Escape Dynamics* on, the government applies updating schemes like (100.1) to functions, the coefficients of its Phillips curve, rather than to numbers.

The LQ Phelps problem solved here returns in *Adaptive Learning and Escape Dynamics*, *Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*, and *The Lost Conquest: Fed Policy in the 2020s*, where a learning government re-solves it every period.

In *Adaptive Learning and Escape Dynamics*, it is the drift of the government's estimated Phillips curve toward the induction hypothesis that triggers the Volcker-like stabilizations.

100.8 Exercises

Exercise 100.8.1

Formula (100.5) makes the comparison of λ with δ exact.

1. For $\lambda = 0.8$, compute the limiting inflation rate y_∞ (the value that y_t settles at after a long simulation) for a grid of discount factors $\delta \in \{0.90, 0.91, \ldots, 0.99\}$, plot it against δ, and overlay (100.5).
2. Plot $y_\infty / \theta U^*$ against $\lambda \in [0, 0.999]$ for $\delta \in \{0.9, 0.96, 0.99\}$.
3. For each of these δ, how small must λ be for y_∞ to be within 10% of the Ramsey value, that is, $y_\infty \leq 0.1\, \theta U^*$?

Solution

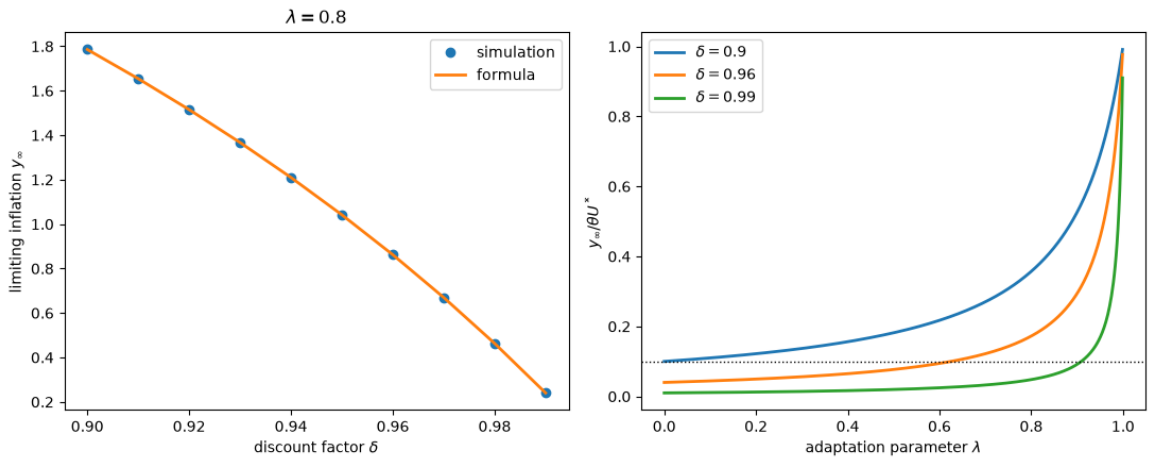
```
δ_grid = np.arange(0.90, 0.995, 0.01)
y_sim = [PhelpsProblem(λ=0.8, δ=δ).simulate(x0=12.0, T=400)[1][-1]
         for δ in δ_grid]
y_formula = 5.0 * (1 - δ_grid) / (1 - δ_grid * 0.8)

λ_grid = np.linspace(0, 0.999, 500)

fig, axes = plt.subplots(1, 2, figsize=(11, 4.5))
axes[0].plot(δ_grid, y_sim, 'o', label='simulation')
axes[0].plot(δ_grid, y_formula, lw=2, label='formula')
axes[0].set_xlabel(r'discount factor $\delta$')
axes[0].set_ylabel(r'limiting inflation $y_\infty$')
axes[0].set_title(r'$\lambda = 0.8$')
axes[0].legend()

for δ in [0.9, 0.96, 0.99]:
    axes[1].plot(λ_grid, (1 - δ) / (1 - δ * λ_grid), lw=2,
                 label=rf'$\delta = {δ}$')
    λ_max = (1 - 10 * (1 - δ)) / δ
    print(f"δ = {δ}: y_∞ ≤ 0.1 θU* requires λ ≤ {λ_max:.3f}")
axes[1].axhline(0.1, color='k', ls=':', lw=1)
axes[1].set_xlabel(r'adaptation parameter $\lambda$')
axes[1].set_ylabel(r'$y_\infty / \theta U^*$')
axes[1].legend()
plt.tight_layout()
plt.show()
```

```
δ = 0.9: y_∞ ≤ 0.1 θU* requires λ ≤ 0.000
δ = 0.96: y_∞ ≤ 0.1 θU* requires λ ≤ 0.625
δ = 0.99: y_∞ ≤ 0.1 θU* requires λ ≤ 0.909
```



The simulated limits lie on the formula.

As δ rises toward one, the government becomes willing to accept the transitional recession needed to reap the long-run benefit of low expected inflation, so the limiting inflation rate falls toward the Ramsey value of zero.

From (100.5), $y_\infty \leq 0.1\,\theta U^*$ requires $1 - \delta\lambda \geq 10(1-\delta)$, that is, $\lambda \leq [1 - 10(1-\delta)]/\delta$.

When $\delta = 0.9$ no $\lambda \in (0,1)$ qualifies, and even when $\delta = 0.99$ the public must adapt fast enough that λ is below about 0.91.

As $\lambda \to 1$ every curve rises to one: very slow adaptation subverts the Phelps model, whatever the discount factor.

Exercise 100.8.2

This exercise reproduces the Lucas–Sargent critique of the Solow–Tobin test, in the form of (100.11).

Consider an economy in which both the natural-rate hypothesis and rational expectations hold:

$$y_t = (1-\rho)\mu + \rho y_{t-1} + e_t, \qquad x_t = \mathbb{E}_{t-1} y_t, \qquad U_t = U^* - \theta(y_t - x_t) + v_{1t},$$

where e_t and v_{1t} are independent white noises with variances σ_e^2 and σ_1^2, and $|\rho| < 1$.

1. Show that the population regression (100.10) has $b_1 = \rho(1+\lambda)/(1+\lambda\rho)$ and $b_2 = -\theta\sigma_e^2/(\theta^2\sigma_e^2 + \sigma_1^2)$.
2. Verify both by simulation, with $\lambda = 0.7$, $\theta = 1$, $U^* = 5$, $\mu = 4$, $\sigma_e = 1$, $\sigma_1 = 0.5$, and $\rho \in \{0, 0.5, 0.9, 0.99\}$.
3. Interpret the results.

Solution

Write $z_t = (1-\lambda)\sum_{i\geq 1} \lambda^{i-1} y_{t-i}$ for the regressor in (100.10).

Under rational expectations $y_t = x_t + e_t$ with $x_t = (1-\rho)\mu + \rho y_{t-1}$, and $U_t - U^* = -\theta e_t + v_{1t}$.

Because U_t depends only on current shocks, it is uncorrelated with z_t and with x_t, so the two slopes can be computed one at a time.

The coefficient on U_t is $b_2 = \operatorname{cov}(e_t, U_t)/\mathbb{V}[U_t] = -\theta\sigma_e^2/(\theta^2\sigma_e^2 + \sigma_1^2)$.

The coefficient on z_t is $b_1 = \operatorname{cov}(x_t, z_t)/\mathbb{V}[z_t] = \rho \operatorname{cov}(y_{t-1}, z_t)/\mathbb{V}[z_t]$.

With $\gamma_0 = \mathbb{V}[y_t]$ and autocovariances $\gamma_0 \rho^j$,

$$\operatorname{cov}(y_{t-1}, z_t) = \frac{(1-\lambda)\gamma_0}{1-\lambda\rho}, \qquad \mathbb{V}[z_t] = \frac{(1-\lambda)^2\gamma_0(1+\lambda\rho)}{(1-\lambda^2)(1-\lambda\rho)},$$

and the ratio gives $b_1 = \rho(1+\lambda)/(1+\lambda\rho)$.

```
rng = np.random.default_rng(0)
λ, θ, U_star, μ, σ_e, σ_1, T = 0.7, 1.0, 5.0, 4.0, 1.0, 0.5, 200_000
burn = 2_000

for ρ in [0.0, 0.5, 0.9, 0.99]:
    e = σ_e * rng.standard_normal(T)
    v1 = σ_1 * rng.standard_normal(T)
    y = μ + lfilter([1.0], [1.0, -ρ], e)                    # AR(1) around μ
    U = U_star - θ * e + v1                                  # y_t - x_t = e_t
    z = lfilter([0.0, 1 - λ], [1.0, -λ], y)                 # (1-λ) Σ λ^{i-1} y_{t-i}
    X = np.column_stack([np.ones(T - burn), z[burn:], U[burn:]])
    b0, b1, b2 = np.linalg.lstsq(X, y[burn:], rcond=None)[0]
    print(f"ρ = {ρ:<4}:  b1 = {b1:.3f} (formula {ρ * (1 + λ) / (1 + λ * ρ):.3f}),"
          f"  b2 = {b2:.3f} (formula {-θ * σ_e**2 / (θ**2 * σ_e**2 + σ_1**2):.3f})
↪")
```

```
ρ = 0.0 :  b1 = -0.003 (formula 0.000),  b2 = -0.799 (formula -0.800)
ρ = 0.5 :  b1 = 0.633 (formula 0.630),  b2 = -0.801 (formula -0.800)
ρ = 0.9 :  b1 = 0.940 (formula 0.939),  b2 = -0.800 (formula -0.800)
ρ = 0.99:  b1 = 0.994 (formula 0.994),  b2 = -0.800 (formula -0.800)
```

The estimates match the formulas.

In every case $b_1 < 1$, so the Solow–Tobin test rejects the natural-rate hypothesis even though it is true here.

The induction hypothesis builds in weights that sum to one, but rational expectations of a stationary process do not.

As $\rho \to 1$ the bias vanishes, which is the King–Watson reading of why later tests, run on data in which inflation had a unit root, stopped rejecting.

The coefficient b_2 is the slope that a *Keynesian* regression of inflation on unemployment recovers.

It is the slope β_1 of the Keynesian self-confirming equilibrium in *Self-Confirming Equilibria*, with σ_e in the role of σ_2.

Exercise 100.8.3

The sacrifice ratio (100.12) does not depend on the policy.

Fix $\delta = 0.96$ and take $\lambda \in \{0.7, 0.9, 0.95, 0.99\}$, starting from $x = 12$.

1. Compute cumulative excess unemployment per point of disinflation along the optimal path.
2. Do the same for a "cold turkey" policy that sets $y_t = 0$ forever, and for a gradual policy that lowers inflation by one point per period, starting from 11, until it reaches zero.
3. Compare the three with $\theta/(1-\lambda)$, and explain why the optimal policy nonetheless differs from cold turkey.

Solution

```
def sacrifice_ratio(y_path, λ, θ=1.0, x0=12.0):
```

```
    "Cumulative excess unemployment per point of fall in expected inflation."
    x, excess = x0, 0.0
    for y in y_path:
        excess += -θ * (y - x)           # U_t - U*
        x = λ * x + (1 - λ) * y
    return excess / (x0 - x)

T = 5_000
δ = 0.96
for λ in [0.7, 0.9, 0.95, 0.99]:
    U, y_opt = PhelpsProblem(λ=λ, δ=δ).simulate(x0=12.0, T=T)
    y_cold = np.zeros(T)
    y_grad = np.maximum(11.0 - np.arange(T), 0.0)
    print(f"λ = {λ}:  optimal {sacrifice_ratio(y_opt, λ):6.2f},  "
          f"cold turkey {sacrifice_ratio(y_cold, λ):6.2f},  "
          f"gradual {sacrifice_ratio(y_grad, λ):6.2f},  "
          f"θ/(1-λ) = {1 / (1 - λ):6.2f}")
```

```
λ = 0.7:  optimal   3.33,  cold turkey   3.33,  gradual   3.33,  θ/(1-λ) =   3.
↪33
λ = 0.9:  optimal  10.00,  cold turkey  10.00,  gradual  10.00,  θ/(1-λ) =  10.
↪00
λ = 0.95:  optimal  20.00,  cold turkey  20.00,  gradual  20.00,  θ/(1-λ) =  20.
↪00
λ = 0.99:  optimal 100.00,  cold turkey 100.00,  gradual 100.00,  θ/(1-λ) = 100.
↪00
```

All three policies cost $\theta/(1-\lambda)$ points of unemployment per point of disinflation, as (100.12) says they must.

The policies differ in how far they disinflate and in how they spread the unemployment over time.

Cold turkey takes inflation all the way to zero, and front-loads the recession.

The optimal discounted policy stops at $y_\infty > 0$ for the reason given after (100.5): near the steady state, the discounted benefit of the last point of disinflation, $y_\infty/(1-\delta)$, just equals its discounted unemployment cost, $\theta U^*/(1-\delta\lambda)$.

As $\lambda \to 1$ that cost rises toward $\theta U^*/(1-\delta)$, so y_∞ rises toward θU^*.

With a quadratic loss, the government also prefers to spread the unemployment out rather than take it all at once.

CHAPTER

ONE

OPTIMAL MISSPECIFIED BELIEFS

Contents

- *Optimal Misspecified Beliefs*
  - *Overview*
  - *An experiment in Bray's lab*
  - *Optimal misspecification*
  - *Comparing the true and forecasting models*
  - *How the equilibrium depends on ρ*
  - *Lessons*
  - *Exercises*

101.1 Overview

This lecture continues the study of Phillips curve tradeoffs.

It follows chapter 6 of Sargent [1999].

In *Adaptive Expectations and the Phelps Problem* the public forecast inflation with a fixed adaptive rule whose parameter λ we simply chose.

That was unsatisfying in a way the *vindication* story of *The Rise and Fall of U.S. Inflation* cannot afford: a free parameter describing expectations is exactly what rational expectations was meant to eliminate.

Here we take the first step toward earning that parameter, by letting agents choose it to fit the data their own beliefs generate.

We describe three conceptual issues that recur throughout this suite of lectures:

1. how to formulate equilibria in which agents share a common *misspecified* least squares forecasting model,
2. how expectations can contribute independent dynamics within equilibria, and
3. how the classic adaptive expectations scheme can use second moments to approximate a first moment.

To expose these issues we temporarily set aside the Phillips curve and work with Bray [1982]'s simple model of the price of a single good, a workhorse for studying bounded rationality.

From relaxation to learning: Bray's market presents Bray's model as the relaxation algorithm with a gain of $1/t$, and *Least Squares Learning in Self-Referential Models* studies least squares learning in it.

We alter Bray's model to illustrate an equilibrium concept that merges aspects of rational and adaptive expectations in a new way, and that we apply to the Phillips curve in *Self-Confirming Equilibria*.

The focus is on **market equilibrium with optimal but misspecified forecasts**.

- *Optimal* means the free parameters of the forecasting scheme are chosen by (nonlinear) least squares.
- *Misspecified* means the forecasting model is wrong in functional form, or uses too small an information set.

A distinctive feature is that the true model *depends on* how the agents' model is misspecified: agents' beliefs affect their behavior, which shapes the data they then fit.

We will also see that the equilibrium depends on a detail of the agents' model, a root ρ that they place just below one, and that the forecasting gain it delivers vanishes as $\rho \to 1$.

We work in the frequency domain, so let's import what we need:

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.optimize import minimize_scalar
from scipy.signal import lfilter
```

101.2 An experiment in Bray's lab

Following Bray [1982], assume that

$$p_t = a + b\, p^e_{t+1} + u_t, \tag{101.1}$$

where u_t is IID with mean zero and variance σ_u^2, $a > 0$, $b \in (0, 1)$, p_t is the market price, and p^e_{t+1} is the market's expectation of next period's price.

In the notation of *From relaxation to learning: Bray's market*, Bray's b is the feedback coefficient h'.

The rational expectations equilibrium has $p^e_{t+1} = \frac{a}{1-b}$ and $p_t = \frac{a}{1-b} + u_t$.

Bray posited that p^e_{t+1} is the empirical average of past prices, and showed that when $0 < b < 1$ this decreasing-gain scheme converges almost surely to the rational expectation $\frac{a}{1-b}$.

During the transition, the state variable p^e_t contributes dynamics and makes the price serially correlated, but these dynamics are transitory: at the rational expectations equilibrium, p_t is a constant plus a serially uncorrelated shock.

101.2.1 Constant-gain adaptive expectations

To let expectations impart *persistent* serial correlation, we depart from Bray and assume that the market has **constant-gain** adaptive expectations,

$$p^e_{t+1} = C p_t + (1 - C) p^e_t, \qquad |C| < 1. \tag{101.2}$$

Bray's scheme replaces C by $\frac{1}{t}$, so that p^e_{t+1} becomes a sample average.

Fixing C instead *discounts* past observations, which arrests convergence to rational expectations and prevents p^e_{t+1} from converging to a constant.

Written as a distributed lag,

$$p^e_{t+1} = \frac{C}{1 - (1 - C)L} p_t, \tag{101.3}$$

where L is the lag operator.

Equation (101.2) would be the linear least squares forecast if the price followed the integrated moving-average process

$$p_t = p_{t-1} + \epsilon_t - (1-C)\epsilon_{t-1}, \tag{101.4}$$

so the market perceives the price as composed of purely permanent and transitory components.

This is the answer to Muth [1960]'s inverse optimal-prediction problem in *A Peculiar Definition of Bounded Rationality*, with Muth's moving-average coefficient equal to $1-C$.

Equivalently, the perceived model is a *local-level* model, a random walk observed with white noise.

For it, (101.2) is the steady-state Kalman filter of *Convergence*, and the gain C is tied to the ratio q/r of the variance of the random-walk innovation to that of the noise by $q/r = C^2/(1-C)$ (Exercise 101.7.4).

Muth asked for which process a given smoothing rule is optimal.

Here the process is generated by the rule, and agents choose the rule to fit the process.

101.2.2 The actual law of motion

Substituting the belief (101.3) into (101.1) shows that when the market forecasts this way, its actions make the *actual* law of motion for the price

$$p_t = \frac{a}{1-b} + \frac{1}{1-bC}\left[\frac{1-(1-C)L}{1-\frac{1-C}{1-bC}L}\right] u_t = \nu + f(L)u_t, \tag{101.5}$$

where $\nu = \frac{a}{1-b}$ and $f(L)$ is defined to match.

The price has mean ν and spectral density

$$F(\omega) = f(e^{i\omega})f(e^{-i\omega})\,\sigma_u^2, \qquad \omega \in [-\pi, \pi].$$

Notice that F depends on C through f.

Let's encode the true process.

```
class BrayModel:
    """
    Bray's price model with constant-gain adaptive expectations.

    ρ is the root of the agents' perceived ARMA(1,1) model, which stands in
    for the unit root of the integrated moving-average model (see below).
    """

    def __init__(self, a=1.0, b=0.5, σ_u=1.0, ρ=0.995, N=1024):
        self.a, self.b, self.σ_u, self.ρ, self.N = a, b, σ_u, ρ, N
        self.ν = a / (1 - b)
        ω = 2 * np.pi * np.arange(N) / N
        self.ω = ω
        self.z = np.exp(1j * ω)

    def true_spectrum(self, C):
        "Spectral density F(ω) of the actual price process, given belief C."
        b, z = self.b, self.z
        ψ = (1 - C) / (1 - b * C)
        scale = 1 / (1 - b * C)
```

(continues on next page)

(continued from previous page)

```
        f = scale * (1 - (1 - C) * z) / (1 - ψ * z)
        return np.abs(f)**2 * self.σ_u**2

    def approx_spectrum(self, c, σ_ε2=1.0):
        "Spectral density G(ω) of the agent's approximating model."
        g = (1 - (1 - c) * self.z) / (1 - self.ρ * self.z)
        return np.abs(g)**2 * σ_ε2
```

101.3 Optimal misspecification

Two facts about the actual law (101.5) motivate an equilibrium restriction on C:

1. Given that the price obeys (101.5), the true linear least squares one-step forecasting rule is *not* a geometric distributed lag like (101.2).
2. Even restricting the forecast to the form (101.2), the *best* such rule would make C solve a forecast-error-minimization problem, so C is an outcome, not a free parameter.

A rational expectations equilibrium would repair both features.

Chapter 6 of Sargent [1999] softens the equilibrium concept, proceeding "in the spirit (but not the letter)" of Lucas and Prescott [1971] and Brock [1972]: we leave feature 1 untouched (agents keep the wrong functional form) while fixing feature 2 (they choose the best parameter within that form).

Think of putting a single individual into a market where everyone else (the "representative agent") uses C, so the price obeys (101.5).

This "big C, little c" device is the "big Y, little y" trick of *Rational Expectations Equilibrium*, applied to a forecasting rule rather than to a decision rule.

The individual chooses c to fit the best model of the form

$$p_t = \frac{1-(1-c)L}{1-L}\epsilon_t, \tag{101.6}$$

by minimizing the one-step-ahead forecast error variance.

101.3.1 The perceived model and its root ρ

A model with the exact unit root of (101.6) cannot be fitted to data that have a mean, and its spectral density is infinite at frequency zero.

Following Sargent [1999] (chapter 6, footnotes 4 and 5), the agents' perceived model replaces the unit root by a root ρ slightly below one:

$$p_t = \frac{1-(1-c)L}{1-\rho L}\epsilon_t = g(L)\epsilon_t, \qquad 0<\rho<1. \tag{101.7}$$

This is a stationary, mean-zero ARMA(1,1) model.

It has no constant, so it must use the near-unit-root persistence of $1/(1-\rho L)$ to fit the nonzero mean ν of the price.

We treat ρ as a primitive of the perceived model, part of what agents believe, not as a numerical device: the equilibrium depends on it, as *How the equilibrium depends on \rho* shows.

Definition 101.3.1 (Best-estimate map)

Given C and the consequent price process (101.5), the individual's best forecast parameter $c = B(C)$ is the nonlinear least squares estimator of c in (101.7), where the data are generated by (101.5).

Inverting (101.7) gives the perceived model's one-step-ahead forecast error $\epsilon_t = g(L)^{-1} p_t$.

Applied to the actual price, this error has mean $\nu / g(1) = \nu(1-\rho)/c$ and a variance that we can compute from the spectral density F.

Following the frequency-domain method of Hansen and Sargent [Hansen and Sargent, 1993], the best approximating (c, σ_ϵ^2) minimizes

$$A(c, \sigma_\epsilon^2) = \frac{1}{N} \sum_{j=0}^{N-1} \left\{ \log G(\omega_j, c) + \frac{F(\omega_j)}{G(\omega_j, c)} \right\} + \frac{\nu^2}{G(0)}, \tag{101.8}$$

where $\omega_j = \frac{2\pi j}{N}$ and $G(\omega, c) = |g(e^{i\omega})|^2 \sigma_\epsilon^2$ is the spectral density of the perceived model.

The term $\frac{\nu^2}{G(0)} = \frac{\nu^2 (1-\rho)^2}{c^2 \sigma_\epsilon^2}$ charges the perceived model for the squared mean of its forecast errors; this is how the model is made to use its near-unit root to fit the mean.

Concentrating out σ_ϵ^2 sets it equal to the mean square forecast error

$$\sigma_\epsilon^2(c) = \frac{1}{N} \sum_{j=0}^{N-1} \frac{F(\omega_j)}{|g(e^{i\omega_j})|^2} + \frac{\nu^2 (1-\rho)^2}{c^2}. \tag{101.9}$$

The average of $\log |g(e^{i\omega})|^2$ over frequencies is zero, by Kolmogorov's formula for an invertible filter with $g(0) = 1$.

So minimizing (101.8) over c is the same as minimizing $\sigma_\epsilon^2(c)$: the agent chooses the c that minimizes the mean square one-step-ahead forecast error, computed under his model (101.7).

The mean term shows where ρ enters: a larger c shrinks the squared mean $\nu^2 (1-\rho)^2 / c^2$, so the closer ρ is to one, the weaker the pull toward a large c.

Definition 101.3.2 (Equilibrium under forecast misspecification)

An equilibrium under forecast misspecification is a fixed point $C = B(C)$.

At such a fixed point the representative agent is representative: the single individual's best parameter equals the one everyone uses.

```
def forecast_mse(F, ν, c, ρ, z):
    """
    Mean square one-step forecast error σ_ε^2(c) of the perceived model
    g(L) = (1 - (1-c)L)/(1 - ρL), when the true process has mean ν and
    spectral density F on the grid z = exp(iω).  ρ = 1 is allowed.
    """
    W = np.abs((1 - ρ * z) / (1 - (1 - c) * z))**2      # |1/g(e^{iω})|^2
    return np.mean(F * W) + ν**2 * W[0]                 # variance + squared mean

def best_estimate(model, C):
    "The best-estimate map c = B(C)."
    F = model.true_spectrum(C)
    res = minimize_scalar(lambda c: forecast_mse(F, model.ν, c, model.ρ, model.z),
                          bounds=(1e-4, 0.99), method='bounded',
```

(continues on next page)

(continued from previous page)

```
                           options={'xatol': 1e-10})
    return res.x

def solve_equilibrium(model, C0=0.3, tol=1e-10, maxit=500):
    "Iterate the best-estimate map to a fixed point."
    C = C0
    for _ in range(maxit):
        C_new = best_estimate(model, C)
        if abs(C_new - C) < tol:
            break
        C = C_new
    return C_new
```

Now we solve for the equilibrium of the baseline model.

```
bray = BrayModel(a=1.0, b=0.5, σ_u=1.0, ρ=0.995)
C_star = solve_equilibrium(bray)
print(f"equilibrium belief   C = {C_star:.4f}")
```

```
equilibrium belief   C = 0.0805
```

For these parameters the equilibrium belief is $C \approx 0.08$, reproducing the value reported in chapter 6 of Sargent [1999].

101.3.2 Forecast errors

Three standard deviations of one-step-ahead forecast errors describe the equilibrium:

- $\bar{\sigma}_\epsilon$, the root mean square forecast error of the perceived model (101.7) at $c = C$, including the mean of its errors: the square root of the minimized (101.9).
- the root mean square error of the rule (101.2) that agents actually use. Its weights sum to one, so its errors have mean zero; it is σ_ϵ computed with $\rho = 1$.
- $\sigma_u/(1 - bC)$, the standard deviation of the innovation in the true process (101.5), the smallest one-step error that any forecaster who knew the true law could attain.

```
F_star = bray.true_spectrum(C_star)
c_star = best_estimate(bray, C_star)

σ_bar = np.sqrt(forecast_mse(F_star, bray.ν, c_star, bray.ρ, bray.z))
σ_rule = np.sqrt(forecast_mse(F_star, bray.ν, C_star, 1.0, bray.z))
σ_true = bray.σ_u / (1 - bray.b * C_star)

print(f"perceived model (ρ = {bray.ρ}):    {σ_bar:.4f}")
print(f"rule actually used (ρ = 1):     {σ_rule:.4f}")
print(f"true innovation σ_u/(1 - bC):   {σ_true:.4f}")
```

```
perceived model (ρ = 0.995):    1.0579
rule actually used (ρ = 1):     1.0530
true innovation σ_u/(1 - bC):   1.0419
```

The first number is the $\bar{\sigma}_\epsilon$ that chapter 6 of Sargent [1999] reports.

The book gives 1.0208 for the last one, which is the square root of the value printed here.

The forecasts that agents actually make are only slightly worse than the best attainable ones.

101.3.3 What kind of equilibrium is this?

The fixed point $C = B(C)$ has several names in the later literature, each emphasizing a different feature.

- It is a **restricted-perceptions equilibrium** (Evans and Honkapohja [2001], chapter 13): forecasts are optimal within a restricted class of models that excludes the truth.
- For given C, the best estimate c is a **pseudo-true value** in the sense of White [1982], the parameter of a misspecified model that comes closest to the truth. Unlike in White's setting, here the truth itself depends on the agents' parameter, through C.
- Up to constants, criterion (101.8) is the frequency-domain (Whittle) form of the Kullback–Leibler divergence of the perceived model from the true process. So the fixed point is also a **Berk–Nash equilibrium** in the sense of Esponda and Pouzo [2016], with the forecasting rule as the agents' action.

Learning, Approximation, and Equilibrium Computation computes the same kind of object, which it calls an approximate equilibrium, for a saving rule restricted to a polynomial family.

There the agents' fit is a least squares projection onto a basis; here it is a nonlinear fit of an ARMA model across the whole spectrum.

Because B works like a best-response map, we can ask whether agents who nudge their gain toward the best estimate would find the equilibrium.

The associated differential equation $dC/d\tau = B(C) - C$ is the E-stability condition of *One recursion, many learners*: the fixed point is E-stable when $B'(C^*) < 1$.

```
h = 1e-4
B_prime = (best_estimate(bray, C_star + h) - best_estimate(bray, C_star - h)) / (2 *␣
 ↪h)
print(f"B'(C*) = {B_prime:.3f}")
```

```
B'(C*) = 0.148
```

Here $0 < B'(C^*) < 1$, so the equilibrium is E-stable, and iterating B converges to it monotonically, which is how `solve_equilibrium` found it.

Why don't the agents eventually discover that their model is wrong?

Kalai and Lehrer [1993] applied a merging theorem of Blackwell and Dubins [1962] to establish conditions under which forecasts must eventually merge with rational expectations.

The present model shuts off that mechanism at the outset: the agents' model puts zero probability on events that have positive probability under the truth, violating the absolute continuity condition that merging requires.

With enough data, appropriate statistical tests would eventually tell agents how to improve their forecasts.

Kreps [1998] argues that sensible agents would eventually abandon a misspecified model that the data reject; Exercise 101.7.3 estimates how much data that would take here.

The answer lies in how far apart the true and approximating spectral densities are, which we now study.

101.4 Comparing the true and forecasting models

For the equilibrium C, we plot the equilibrium spectral densities of the true and approximating models.

```
G = bray.approx_spectrum(c_star, σ_bar**2)

half = bray.N // 2
fig, ax = plt.subplots(figsize=(8, 5))
ax.plot(bray.ω[:half], np.log(F_star[:half]), 'C0', label='true model', lw=2)
ax.plot(bray.ω[:half], np.log(G[:half]), 'C1--', label='forecasting model', lw=2)
ax.set_xlabel(r'angular frequency $\omega$')
ax.set_ylabel('log spectral density')
ax.legend()
plt.show()
```

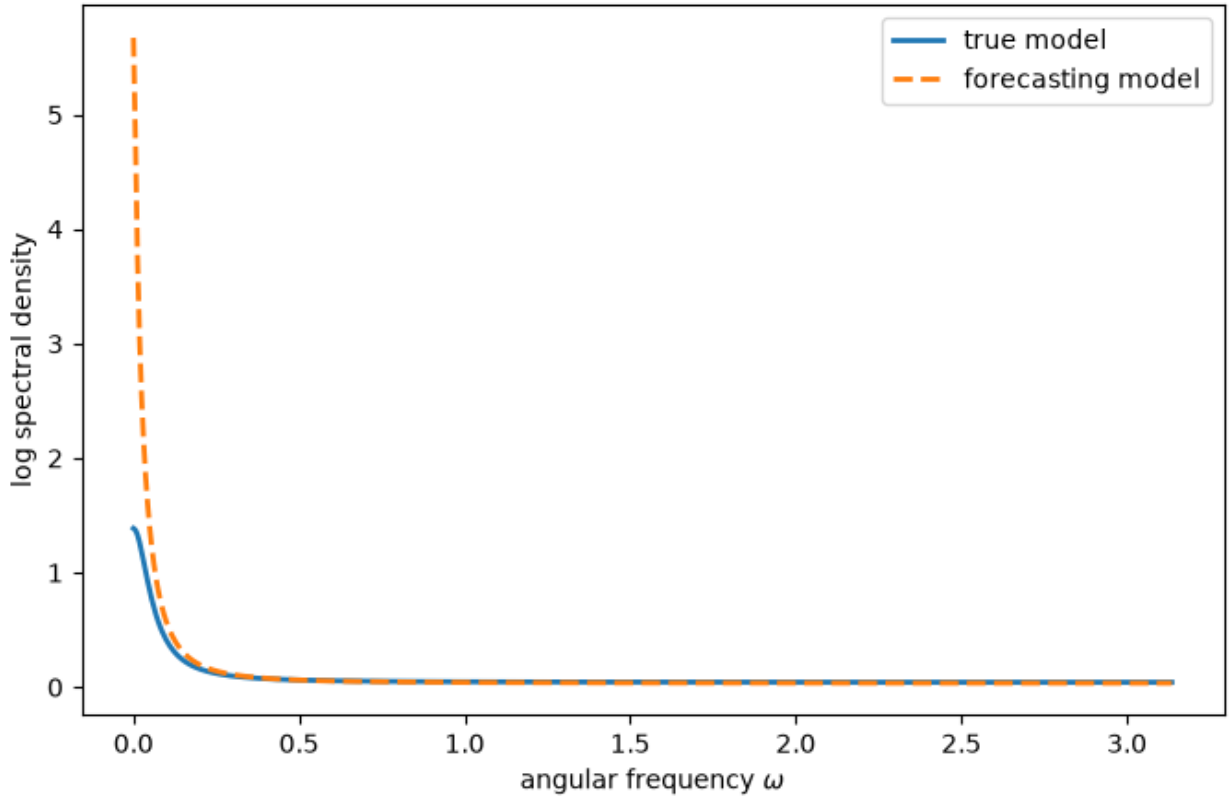


Fig. 101.1: Spectral densities of the true price process and of the agents' approximating model

In minimizing (101.8), the approximating model uses its near-unit-root to fit the mean.

The large gap between the spectral densities at low frequencies reflects how the approximating model fits a *first* moment (the mean ν) with features of *second* moments (a spike in the spectral density at frequency zero).

The true spectral density decreases sharply with frequency — Granger's "typical spectral shape" [Granger, 1966] — revealing substantial positive serial correlation in the price, because agents' belief that the price is subject to permanent shocks makes shocks persist.

101.4.1 Impulse responses

We compare the impulse response functions of the two models by feeding a unit shock through each moving-average representation.

The true process (101.5) and the perceived model (101.7) are driven by different shocks, u_t and ϵ_t.

To compare like with like, we normalize the true response by its impact effect $f(0) = 1/(1 - bC)$, so that both models respond to a unit one-step-ahead forecast error.

```
def impulse_response(num, den, T=25):
    "IRF of (1 - num L)/(1 - den L)."
    h = np.empty(T)
    h[0] = 1.0
    for k in range(1, T):
        h[k] = den * h[k - 1]
    h[1:] -= num * h[:-1]            # apply the numerator (1 - num L)
    return h

ψ = (1 - C_star) / (1 - bray.b * C_star)
irf_true = impulse_response(1 - C_star, ψ)              # f(L) / f(0)
irf_approx = impulse_response(1 - c_star, bray.ρ)       # g(L)

lags = np.arange(len(irf_true))
fig, axes = plt.subplots(1, 2, figsize=(11, 4.5))
for ax, start in zip(axes, [0, 1]):
    ax.plot(lags[start:], irf_true[start:], 'C0o-', ms=4, label='true model', lw=2)
    ax.plot(lags[start:], irf_approx[start:], 'C1s--', ms=4,
            label='approximating model', lw=2)
    ax.set_xlabel('lag')
    ax.set_ylabel('response')
axes[0].set_title('all lags')
axes[1].set_title('lags 1 and beyond')
axes[0].legend()
plt.tight_layout()
plt.show()

print(f"response at lag 1:  true {irf_true[1]:.3f},  approximating {irf_approx[1]:.3f}
 ↪")
```

```
response at lag 1:  true 0.039,  approximating 0.076
```

The impulse response of the true model affirms the serial correlation in the price.

From lag one on, the true response lies below the approximating one: the approximating model over-predicts the effects of a shock at every horizon, and its near unit root makes those effects die out much more slowly, at rate ρ rather than $\phi = (1 - C)/(1 - bC)$.

Chapter 6 of Sargent [1999] says that the approximating model under-predicts short-term consequences; that statement rests on comparing impact effects without the normalization, and so reflects only the different scales of the two shocks.

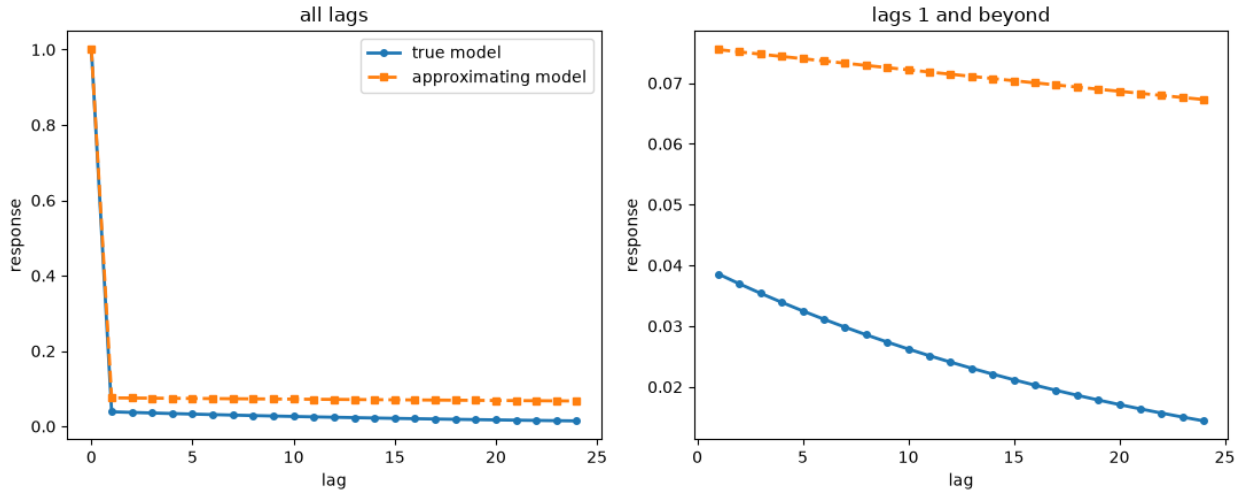


Fig. 101.2: Impulse responses of the true and approximating models to a unit one-step-ahead forecast error

101.5 How the equilibrium depends on ρ

Every number so far was computed with the perceived root $\rho = 0.995$.

The next cell recomputes the equilibrium for values of ρ closer to one, and then for the exact unit root $\rho = 1$ of the rule (101.2) itself.

We use a finer frequency grid, because the perceived spectral density has a spike at frequency zero whose width is about $1 - \rho$.

```
print(f"{'ρ':>7} {'C*':>8} {'weight sum':>11} {'RMSE of rule':>13}")
for ρ in [0.98, 0.99, 0.995, 0.999, 0.9995]:
    m = BrayModel(a=1.0, b=0.5, σ_u=1.0, ρ=ρ, N=2**14)
    C = solve_equilibrium(m, C0=0.1)
    rmse = np.sqrt(forecast_mse(m.true_spectrum(C), m.ν, C, 1.0, m.z))
    print(f"{ρ:>7} {C:>8.4f} {1 - (1 - ρ) / C:>11.3f} {rmse:>13.4f}")

m1 = BrayModel(a=1.0, b=0.5, σ_u=1.0, ρ=1.0, N=2**14)
C_grid = np.linspace(0.005, 0.4, 40)
ratio = max(best_estimate(m1, C) / C for C in C_grid)
print(f"\nρ = 1: largest value of B(C)/C on the grid = {ratio:.3f}")
print(f"ρ = 1: best c given C = {C_star:.4f} is {best_estimate(m1, C_star):.4f}")
```

```
      ρ       C*  weight sum  RMSE of rule
   0.98   0.1901       0.895        1.1352
   0.99   0.1245       0.920        1.0846
  0.995   0.0805       0.938        1.0530
  0.999   0.0284       0.965        1.0181
 0.9995   0.0180       0.972        1.0114

ρ = 1: largest value of B(C)/C on the grid = 0.365
ρ = 1: best c given C = 0.0805 is 0.0274
```

As ρ rises toward one, the equilibrium gain C^* falls toward zero.

With the exact unit root, $B(C) < C$ for every C on the grid, so iterating B drives C to zero, and the only fixed point is $C = 0$.

That is the rational expectations equilibrium: with $C = 0$ expectations stop moving, and the price is $p_t = \nu + u_t$.

Accordingly, the root mean square error of the rule in the last column falls toward $\sigma_u = 1$ as $\rho \to 1$.

The mechanism is the mean term in (101.9).

A perceived model with $\rho < 1$ forecasts with weights that sum to $1 - (1 - \rho)/c < 1$ (the column "weight sum"), so its errors have mean $\nu(1 - \rho)/c$.

A larger c shrinks that bias, and this pull toward a large c is what keeps C^* away from zero; as $\rho \to 1$ the bias, and the pull, vanish.

There is a tension in the construction.

The price (101.5) is generated by the rule (101.2), whose weights sum to one, but the gain is chosen by fitting the ρ-model.

An agent who chose c to minimize the actual mean square error of the rule he uses would, in effect, set $\rho = 1$, and the last line of output shows that he would pick a gain well below C^* (Exercise 101.7.2 confirms this by simulation).

So the equilibrium with $C^* > 0$, and the persistent serial correlation that comes with it, exists for each fixed $\rho < 1$ and disappears as $\rho \to 1$.

The issue goes back to footnotes 4 and 5 of chapter 6 of Sargent [1999], where $\rho < 1$ is introduced to keep the spectral density well defined.

It matters again in *Self-Confirming Equilibria*, where the better-than-Nash outcomes of an equilibrium with misspecified beliefs rest on it.

101.6 Lessons

The agents in this model are **boundedly rational**, in the sense of *A Peculiar Definition of Bounded Rationality*: *rational* describes their use of least squares, and *bounded* describes their model misspecification.

Under rational expectations there is only one model in play.

Under bounded rationality there must be at least two: the one used by the boundedly rational agents, and the true one.

These mutually influence each other — the boundedly rational agents use their model to approximate the true one, and the true one reflects the decisions of the agents — and both differ from the rational expectations model.

The peculiar way that the adaptive expectations model uses a unit root to mimic a constant foreshadows a version of the Phillips curve model, developed in *Self-Confirming Equilibria*, that will help vindicate econometric policy evaluation, and whose good outcomes also depend on $\rho < 1$.

The same trick — using a unit root to approximate a constant — drives the *escape dynamics* of *Adaptive Learning and Escape Dynamics*, where a learning government's estimated distributed-lag Phillips curve drifts toward the induction hypothesis and, believing it, cuts inflation toward Ramsey.

In the static model of *Escaping Nash Inflation*, an escape instead makes the perceived Phillips curve vertical, which trivially satisfies the induction property.

101.7 Exercises

Exercise 101.7.1

The equilibrium belief C depends on the feedback parameter b in (101.1).

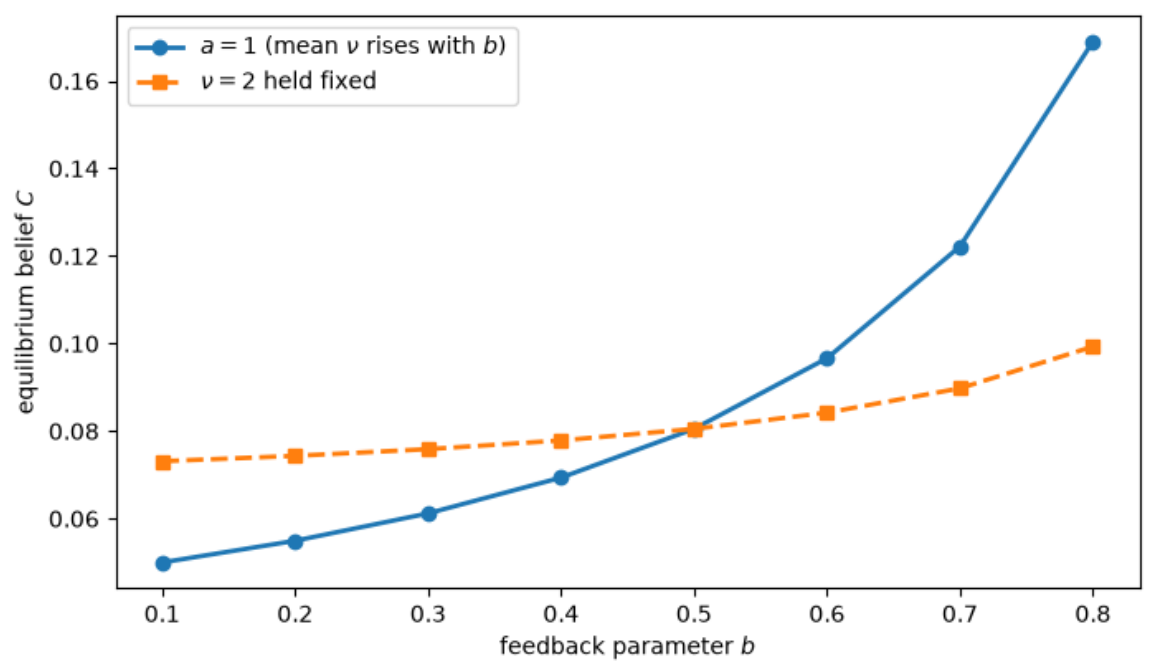
1. Compute and plot the equilibrium C as a function of $b \in \{0.1, 0.2, \ldots, 0.8\}$, holding $a = 1$ and $\sigma_u = 1$ fixed.
2. With a fixed, a change in b also changes the mean $\nu = a/(1-b)$ that the perceived model must fit. Repeat the calculation holding $\nu = 2$ fixed, by setting $a = 2(1-b)$.
3. How much of the effect of b works through the mean? How does the persistence $\phi = (1-C)/(1-bC)$ of the actual price change with b?

Solution

```
b_grid = np.arange(0.1, 0.85, 0.1)
C_a1 = np.array([solve_equilibrium(BrayModel(a=1.0, b=b)) for b in b_grid])
C_v2 = np.array([solve_equilibrium(BrayModel(a=2 * (1 - b), b=b)) for b in b_grid])
ψ_a1 = (1 - C_a1) / (1 - b_grid * C_a1)

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(b_grid, C_a1, 'o-', lw=2, label=r'$a = 1$ (mean $\nu$ rises with $b$)')
ax.plot(b_grid, C_v2, 's--', lw=2, label=r'$\nu = 2$ held fixed')
ax.set_xlabel('feedback parameter $b$')
ax.set_ylabel('equilibrium belief $C$')
ax.legend()
plt.show()

print(f"a = 1:   C from {C_a1[0]:.3f} to {C_a1[-1]:.3f},  "
      f"ν from {1 / (1 - b_grid[0]):.2f} to {1 / (1 - b_grid[-1]):.2f}")
print(f"ν = 2:   C from {C_v2[0]:.3f} to {C_v2[-1]:.3f}")
print(f"persistence ψ ranges from {ψ_a1.min():.3f} to {ψ_a1.max():.3f}")
```




```
a = 1:   C from 0.050 to 0.169,  ν from 1.11 to 5.00
ν = 2:   C from 0.073 to 0.099
persistence ψ ranges from 0.955 to 0.961
```

With $a = 1$, stronger feedback raises the equilibrium gain a great deal, but it also raises the mean ν that the perceived model must fit.

Holding ν fixed, the gain rises much less.

Most of the effect of b therefore works through the mean: a larger mean makes the bias $\nu(1-\rho)/c$ of the perceived model costlier, so agents choose a larger gain.

The persistence ϕ of the actual price barely changes, so stronger feedback does not make the price process noticeably less persistent.

Exercise 101.7.2

The best-estimate map is computed in the frequency domain, under the perceived model (101.7).

Check what the equilibrium implies in the time domain.

Simulate the market (101.1)–(101.2) with $C = C^*$ for one million periods.

For a grid of values of c, compute the mean square error of the forecast $p^e_{t+1}(c) = c\,p_t + (1-c)p^e_t(c)$ of p_{t+1}.

Which c minimizes it?

Compare it with C^* and with the exact-unit-root best estimate printed in *How the equilibrium depends on \rho*.

Solution

Substituting (101.2) into (101.1) gives $p_t = [a + b(1-C)p^e_t + u_t]/(1-bC)$, and then $p^e_{t+1} = \phi\,p^e_t + C(a +$

$u_t)/(1 - bC)$ with $\phi = (1 - C)/(1 - bC)$.

```
rng = np.random.default_rng(0)
T, burn = 1_000_000, 10_000
a, b, C = bray.a, bray.b, C_star
u = rng.standard_normal(T)

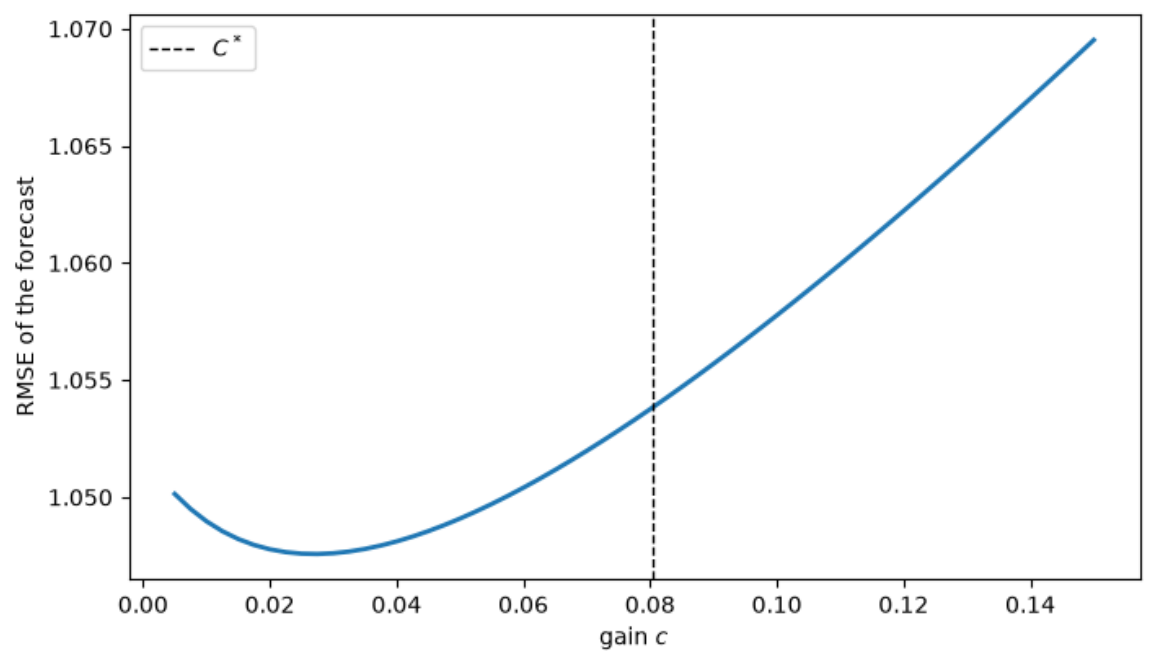
ψ = (1 - C) / (1 - b * C)
pe_next = lfilter([C / (1 - b * C)], [1, -ψ], a + u, zi=[ψ * bray.v])[0]
pe = np.concatenate([[bray.v], pe_next[:-1]])                    # p^e_t
p = (a + b * (1 - C) * pe + u) / (1 - b * C)

def rmse(c):
    "Root mean square error of the forecast c p_t + (1-c) p^e_t of p_{t+1}."
    f = lfilter([c], [1, -(1 - c)], p, zi=[(1 - c) * bray.v])[0]
    return np.sqrt(np.mean((p[burn + 1:] - f[burn:-1])**2))

c_grid = np.linspace(0.005, 0.15, 59)
errors = np.array([rmse(c) for c in c_grid])

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(c_grid, errors, lw=2)
ax.axvline(C_star, color='k', ls='--', lw=1, label='$C^*$')
ax.set_xlabel('gain $c$')
ax.set_ylabel('RMSE of the forecast')
ax.legend()
plt.show()

print(f"RMSE-minimizing c = {c_grid[errors.argmin()]:.3f},  C* = {C_star:.4f}")
```




```
RMSE-minimizing c = 0.027,  C* = 0.0805
```

The gain that minimizes the actual forecast error is far below C^* and close to the exact-unit-root best estimate of *How the equilibrium depends on \rho*.

An individual who judged the rule (101.2) by its actual forecast errors would not choose C^*.

C^* is a best response only within the perceived model (101.7), whose mean penalty does not apply to the rule that agents actually use, because that rule's errors have mean zero.

Exercise 101.7.3

How much data would agents need to detect that their model is wrong?

For Gaussian processes, the expected log-likelihood ratio per observation of the true model (101.5) against the perceived model (101.7) (at $c = B(C^*)$ and the fitted $\bar{\sigma}_\epsilon^2$) is the Kullback–Leibler divergence rate.

1. Using the concentrated criterion and Kolmogorov's formula, show that this rate equals $\log(\bar{\sigma}_\epsilon/\sigma_{\text{true}})$, with $\sigma_{\text{true}} = \sigma_u/(1 - bC^*)$.
2. Compute it, and the number of observations it takes for the expected log-likelihood ratio to reach 3 (a likelihood ratio of about 20).
3. Confirm the rate by simulating the log-likelihood ratio.

Solution

The Kullback–Leibler divergence rate of a mean-zero Gaussian model with spectral density G from a Gaussian process with mean ν and spectral density F is

$$\frac{1}{2}\left\{\frac{1}{N}\sum_j\left[\frac{F(\omega_j)}{G(\omega_j)} - 1 - \log\frac{F(\omega_j)}{G(\omega_j)}\right] + \frac{\nu^2}{G(0)}\right\}.$$

At the fitted $\bar{\sigma}_\epsilon^2$, the terms $\frac{1}{N}\sum_j F/G + \nu^2/G(0)$ add up to one.

By Kolmogorov's formula, the average of $\log G$ is $\log\bar{\sigma}_\epsilon^2$ and the average of $\log F$ is $\log\sigma_{\text{true}}^2$, the log one-step prediction variance of the true process.

So the rate is $\frac{1}{2}\log(\bar{\sigma}_\epsilon^2/\sigma_{\text{true}}^2) = \log(\bar{\sigma}_\epsilon/\sigma_{\text{true}})$.

```
kl = np.log(σ_bar / σ_true)
print(f"KL rate = {kl:.4f} nats per observation")
print(f"observations for an expected log-likelihood ratio of 3: {3 / kl:.0f}")

# simulate: true one-step errors are u_t/(1-bC); the perceived model's are g(L)^{-
↪1} p_t
rng = np.random.default_rng(1)
T, burn = 200_000, 10_000
a, b, C = bray.a, bray.b, C_star
u = rng.standard_normal(T)
ψ = (1 - C) / (1 - b * C)
pe_next = lfilter([C / (1 - b * C)], [1, -ψ], a + u, zi=[ψ * bray.v])[0]
pe = np.concatenate([[bray.v], pe_next[:-1]])
p = (a + b * (1 - C) * pe + u) / (1 - b * C)

e_true = u / (1 - b * C)
e_model = lfilter([1, -bray.ρ], [1, -(1 - c_star)], p)
llr = (-np.log(σ_true) - e_true**2 / (2 * σ_true**2)) \
    - (-np.log(σ_bar) - e_model**2 / (2 * σ_bar**2))
print(f"simulated mean log-likelihood ratio = {llr[burn:].mean():.4f}")
```

```
KL rate = 0.0152 nats per observation
observations for an expected log-likelihood ratio of 3: 198
simulated mean log-likelihood ratio = 0.0153
```

The divergence is about 0.015 nats per observation, so it takes roughly two hundred observations before the data favor the true model by a likelihood ratio of about 20, and that is for a statistician who already knows which alternative to test.

Exercise 101.7.4

Muth [1960] showed that the rule (101.2) is optimal for a random walk observed with noise.

Consider the local-level model $p_t = \mu_t + \eta_t$, $\mu_t = \mu_{t-1} + w_t$, with $\mathbb{V}[w_t] = q$ and $\mathbb{V}[\eta_t] = r$.

1. Show that the steady-state Kalman gain K for this model satisfies $q/r = K^2/(1-K)$, and that the Kalman filter's forecast of p_{t+1} then has the form (101.2) with $C = K$.
2. Set $q/r = C^{*2}/(1-C^*)$, iterate the Riccati equation of *Convergence* to its fixed point, and confirm that $K = C^*$.
3. What does the implied signal-to-noise ratio say about how agents perceive the price?

Solution

With $A = G = 1$, $Q = q$, and $R = r$, the Riccati equation of *A First Look at the Kalman Filter* for the one-step-ahead prediction variance Σ of μ_t is $\Sigma_{t+1} = \Sigma_t - \Sigma_t^2/(\Sigma_t + r) + q$.

At its fixed point, $\Sigma r/(\Sigma + r) = \Sigma - q$, and the gain is $K = \Sigma/(\Sigma + r)$.

Hence $(1-K)\Sigma = \Sigma - q$, so $\Sigma = q/K$, and substituting into $K = \Sigma/(\Sigma + r)$ gives $q/r = K^2/(1-K)$.

The forecast of p_{t+1} is the forecast of μ_{t+1}, which equals the filtered estimate of μ_t: $\hat{\mu}_{t+1} = \hat{\mu}_t + K(p_t - \hat{\mu}_t) = Kp_t + (1-K)\hat{\mu}_t$, which is (101.2) with $C = K$.

```
r = 1.0
q = r * C_star**2 / (1 - C_star)
Σ = 1.0
for _ in range(10_000):
    Σ = Σ - Σ**2 / (Σ + r) + q
K = Σ / (Σ + r)
print(f"q/r = {q / r:.5f},  steady-state Kalman gain K = {K:.4f},  C* = {C_star:.
↪4f}")
```

```
q/r = 0.00705,  steady-state Kalman gain K = 0.0805,  C* = 0.0805
```

The equilibrium beliefs imply a signal-to-noise ratio of less than one percent.

Agents perceive the price as almost entirely transitory noise around a level that drifts very slowly, and that slowly drifting level is how their model mimics the constant mean ν.

CHAPTER

TWO

SELF-CONFIRMING EQUILIBRIA

Contents

- *Self-Confirming Equilibria*
  - *Overview*
  - *Objects in the Phelps problem*
  - *The actual Phillips curve*
  - *Self-confirming equilibria*
  - *The special case solved by hand*
  - *Uniqueness and E-stability*
  - *Equilibrium with misspecified beliefs*
  - *Exercises*

In addition to what's in Anaconda, this lecture will use the following library:

```
!pip install quantecon
```

102.1 Overview

This lecture completes the *equilibrium* half of the study begun in *The Credibility Problem*.

It follows chapter 7 of Sargent [1999], after which the suite turns from fixed beliefs to beliefs that are learned in real time.

We seek models that depart minimally from the basic Kydland and Prescott [1977] model of *The Credibility Problem* but that also let a government's *beliefs* be shaped by the data its own policies generate.

The key equilibrium concept is a **self-confirming equilibrium**: the government has a possibly *wrong* model of the Phillips curve, but it fits that model to the data by least squares, and the data it observes confirm its beliefs.

Fudenberg and Levine [1993] use the term self-confirming equilibrium for a related idea in game theory.

We combine ideas from two literatures:

- King and Watson [1994] document how inferences about the Phillips curve depend on the *direction of fit* — whether one regresses unemployment on inflation (a *classical* identification) or inflation on unemployment (a *Keynesian* identification).

- A literature going back to Muth, Lucas, and Prescott [Lucas and Prescott, 1971, Muth, 1961], and exploited by Marcet and Sargent [1989] and Evans and Honkapohja [2001], formulates rational expectations equilibria as fixed points of mappings from beliefs to the population moments of statistical models; *Rational Expectations Equilibrium* describes the fixed-point view.

We build two self-confirming equilibria that are identical except for the direction of fit, and we find that the direction of minimization affects outcomes.

Both equilibria are unique and E-stable.

We then study an **equilibrium with misspecified beliefs**, combining the Phelps problem of *Adaptive Expectations and the Phelps Problem* with the optimal-misspecification machinery of *Optimal Misspecified Beliefs*.

It delivers outcomes better than Nash, approaching Ramsey as the government becomes patient, provided that the public's perceived model keeps the root ρ of *Optimal Misspecified Beliefs* strictly below one.

In the self-confirming equilibria only the government's model is wrong; in the equilibrium with misspecified beliefs only the public's is.

Let's import what we need:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
from scipy.optimize import minimize_scalar
```

102.2 Objects in the Phelps problem

Recall from *Adaptive Expectations and the Phelps Problem* the ingredients of the general Phelps problem.

The government believes in a reduced-form Phillips curve that it can fit in either direction:

$$\text{Classical:} \quad U_t = \gamma^\top X_{C,t} + \varepsilon_{C,t}, \qquad X_{C,t} = \begin{bmatrix} y_t & X_{t-1}^\top \end{bmatrix}^\top,$$

$$\text{Keynesian:} \quad y_t = \beta^\top X_{K,t} + \varepsilon_{K,t}, \qquad X_{K,t} = \begin{bmatrix} U_t & X_{t-1}^\top \end{bmatrix}^\top,$$

where X_{t-1} collects lagged U's, lagged y's, and a constant.

We partition $\gamma = \begin{bmatrix} \gamma_1 & \gamma_{-1}^\top \end{bmatrix}^\top$, where γ_1 is the coefficient on y_t and γ_{-1} collects the coefficients on X_{t-1}, and partition β the same way.

Solving the Phelps problem takes the government's beliefs γ as given and delivers a decision rule $h(\gamma)$ for inflation, $\hat{y}_t = h(\gamma) X_{t-1}$.

The two parameterizations are related by the inversion formulas

$$\gamma_1 = \beta_1^{-1}, \qquad \gamma_{-1} = -\beta_{-1}/\beta_1. \tag{102.1}$$

102.3 The actual Phillips curve

The *actual* Phillips curve extends the one used in earlier lectures to allow for serially correlated shocks:

$$U_t = U^* - \frac{\theta}{1-\rho_2 L}(y_t - x_t) + \frac{v_{1t}}{1-\rho_1 L}, \tag{102.2}$$

with $|\rho_1| < 1$ and $|\rho_2| < 1$.

Here $v_t = (v_{1t}, v_{2t})^\top$ is a vector white noise with a *diagonal* covariance matrix: the Phillips curve shock v_{1t}, with variance σ_1^2, is uncorrelated with the surprise in inflation $v_{2t} \equiv y_t - x_t$, with variance σ_2^2.

For much of this lecture we set $\rho_1 = \rho_2 = 0$ to make theoretical points, which reduces (102.2) to

$$U_t = U^* - \theta(y_t - x_t) + v_{1t}.$$

102.4 Self-confirming equilibria

A self-confirming equilibrium reconciles the government's beliefs with the environment those beliefs generate.

Definition (Self-confirming equilibrium)

A self-confirming equilibrium is a fixed belief vector γ, a government decision rule $h = h(\gamma)$, and a stationary process for (y_t, U_t, x_t) such that

(a) inflation solves the Phelps problem, $y_t = hX_{t-1} + v_{2t}$;

(b) the public optimally forecasts inflation, $x_t = hX_{t-1}$;

© unemployment is generated by the actual Phillips curve (102.2); and

(d) the government's beliefs satisfy the least squares orthogonality conditions $\mathbb{E}\left[\left(U_t - \gamma^\top X_{C,t}\right) X_{C,t}^\top\right] = 0$ (the **classical** direction of fit).

Condition (d) makes the government's beliefs depend on moment matrices that, through (a)–©, themselves depend on the government's beliefs.

The government's beliefs imply behavior that produces data whose moments *confirm* those beliefs.

A distinct self-confirming equilibrium results from replacing (d) with the **Keynesian** direction of fit:

> (d′) the government fits the Keynesian Phillips curve, $\mathbb{E}\left[\left(y_t - \beta^\top X_{K,t}\right) X_{K,t}^\top\right] = 0$, then recovers γ from the inversion formulas (102.1).

Because the government's beliefs affect the whole probability distribution of the data, the direction of minimization affects outcomes.

Note

Computing a self-confirming equilibrium in general means finding a fixed point of a map $\gamma = T(h(\gamma))$ (classical) or $\beta = S(h(\gamma(\beta)))$ (Keynesian).

The moments in the orthogonality conditions are obtained from the state-space representation of the system by solving a discrete Lyapunov equation.

In practice one iterates a relaxation algorithm $\beta_{j+1} = \kappa\beta_j + (1-\kappa)S(\beta_j)$, the damped iteration of *A Peculiar Definition of Bounded Rationality*, which resembles the least squares learning recursion of *The Credibility Problem*; *Uniqueness and E-stability* uses it.

102.5 The special case solved by hand

When $\rho_1 = \rho_2 = 0$, the government's problem collapses to a sequence of static problems, and each self-confirming equilibrium can be computed by hand.

We take $X_{t-1} = 1$, so the government believes $U_t = \gamma_{-1} + \gamma_1 y_t + \varepsilon_{C,t}$; *Uniqueness and E-stability* shows that adding lags of U and y to X_{t-1} changes nothing.

Minimizing $\mathbb{E}\left[U_t^2 + y_t^2\right]$ under this belief gives the static Phelps rule

$$\hat{y} = h(\gamma) = -\frac{\gamma_1\gamma_{-1}}{1+\gamma_1^2}. \tag{102.3}$$

The actual Phillips curve and the diagonal covariance matrix of v_t imply the second moments

$$\mathbb{V}[U_t] = \theta^2\sigma_2^2 + \sigma_1^2, \qquad \mathbb{V}[y_t] = \sigma_2^2, \qquad \operatorname{cov}(U_t, y_t) = -\theta\sigma_2^2. \tag{102.4}$$

Classical direction of fit (U on y): the slope is

$$\gamma_1 = \frac{\operatorname{cov}(U_t, y_t)}{\mathbb{V}[y_t]} = -\theta.$$

Mean unemployment is U^* and mean inflation is $\bar{x} = h(\gamma)$, and the requirement that the means lie on the regression line, $U^* = \gamma_{-1} + \gamma_1\bar{x}$, gives the intercept $\gamma_{-1} = (\gamma_1^2 + 1)U^*$.

Keynesian direction of fit (y on U): the slope is

$$\beta_1 = \frac{\operatorname{cov}(U_t, y_t)}{\mathbb{V}[U_t]} = \frac{-\theta\sigma_2^2}{\sigma_1^2 + \theta^2\sigma_2^2},$$

with intercept $\beta_{-1} = -\frac{\beta_1^2+1}{\beta_1}U^*$, from which the implied classical coefficients follow by inversion, $\gamma_1 = \beta_1^{-1}$ and $\gamma_{-1} = \frac{\beta_1^2+1}{\beta_1^2}U^*$.

```
class SelfConfirmingStatic:
    """
    The two static self-confirming equilibria (ρ1 = ρ2 = 0), one for
    each direction of fit.  γ1 is the slope; γ_int is the intercept γ_{-1}.
    """

    def __init__(self, θ=1.0, U_star=5.0, σ1=0.3, σ2=0.3):
        self.θ, self.U_star, self.σ1, self.σ2 = θ, U_star, σ1, σ2

    def classical(self):
        "Perceived Phillips curve U = γ_{-1} + γ_1 y under classical fit."
        θ, U_star = self.θ, self.U_star
        γ1 = -θ
        γ_int = (γ1**2 + 1) * U_star
        y_bar = -γ_int * γ1 / (γ1**2 + 1)       # mean (= Nash) inflation
        return γ1, γ_int, y_bar
```

(continues on next page)

(continued from previous page)

```
    def keynesian(self):
        "Perceived Phillips curve under Keynesian fit, inverted to y."
        θ, U_star, σ1, σ2 = self.θ, self.U_star, self.σ1, self.σ2
        β1 = -θ * σ2**2 / (σ1**2 + θ**2 * σ2**2)
        β_int = -(β1**2 + 1) / β1 * U_star
        y1 = 1 / β1
        y_int = (β1**2 + 1) / β1**2 * U_star
        y_bar = β_int / (β1**2 + 1)             # mean inflation
        return y1, y_int, y_bar
```

We compute both equilibria for the parameters of chapter 7 of Sargent [1999].

```
sce = SelfConfirmingStatic(θ=1.0, U_star=5.0, σ1=0.3, σ2=0.3)

y1_C, y0_C, y_C = sce.classical()
y1_K, y0_K, y_K = sce.keynesian()

print("Classical direction of fit")
print(f"  y_1 = {y1_C:.1f},  y_(-1) = {y0_C:.1f},  mean inflation = {y_C:.1f}")
print("Keynesian direction of fit")
print(f"  y_1 = {y1_K:.1f},  y_(-1) = {y0_K:.1f},  mean inflation = {y_K:.1f}")
```

```
Classical direction of fit
  y_1 = -1.0,  y_(-1) = 10.0,  mean inflation = 5.0
Keynesian direction of fit
  y_1 = -2.0,  y_(-1) = 25.0,  mean inflation = 10.0
```

These reproduce the numerical example of chapter 7 of Sargent [1999].

Under the classical direction of fit, mean inflation is the Nash value $\theta U^* = 5$.

Under the Keynesian direction of fit, the government estimates a *flatter* Phillips curve and believes the tradeoff is more favorable than it is.

Its mean inflation is

$$\bar{y}_K = -\frac{U^*}{\beta_1} = \theta U^* \left(1 + \frac{\sigma_1^2}{\theta^2 \sigma_2^2}\right), \tag{102.5}$$

which exceeds the Nash rate whenever $\sigma_1 > 0$.

With $\sigma_1 = \theta \sigma_2$, as here, it is twice the Nash rate, 10.

Let's draw the two self-confirming Phillips curves, reproducing Figure 7.1.

```
fig, ax = plt.subplots(figsize=(7, 6))

U_grid = np.linspace(0, 12, 100)

# perceived Phillips curves U = y_{-1} + y_1 y  =>  y = (U - y_{-1}) / y_1
ax.plot(U_grid, (U_grid - y0_C) / y1_C, 'C0', label='P: classical fit', lw=2)
ax.plot(U_grid, (U_grid - y0_K) / y1_K, 'C1', label='Q: Keynesian fit', lw=2)

ax.plot(sce.U_star, y_C, 'C0o')
ax.annotate('Nash', (sce.U_star, y_C), (sce.U_star + 0.4, y_C - 0.6))
ax.plot(sce.U_star, y_K, 'C1o')
ax.annotate('Keynesian mean', (sce.U_star, y_K),
            (sce.U_star + 0.4, y_K + 0.2))
```

(continues on next page)

(continued from previous page)

```
ax.set_xlim(0, 12)
ax.set_ylim(0, 12)
ax.set_xlabel('unemployment $U$')
ax.set_ylabel('inflation $y$')
ax.legend()
plt.show()
```

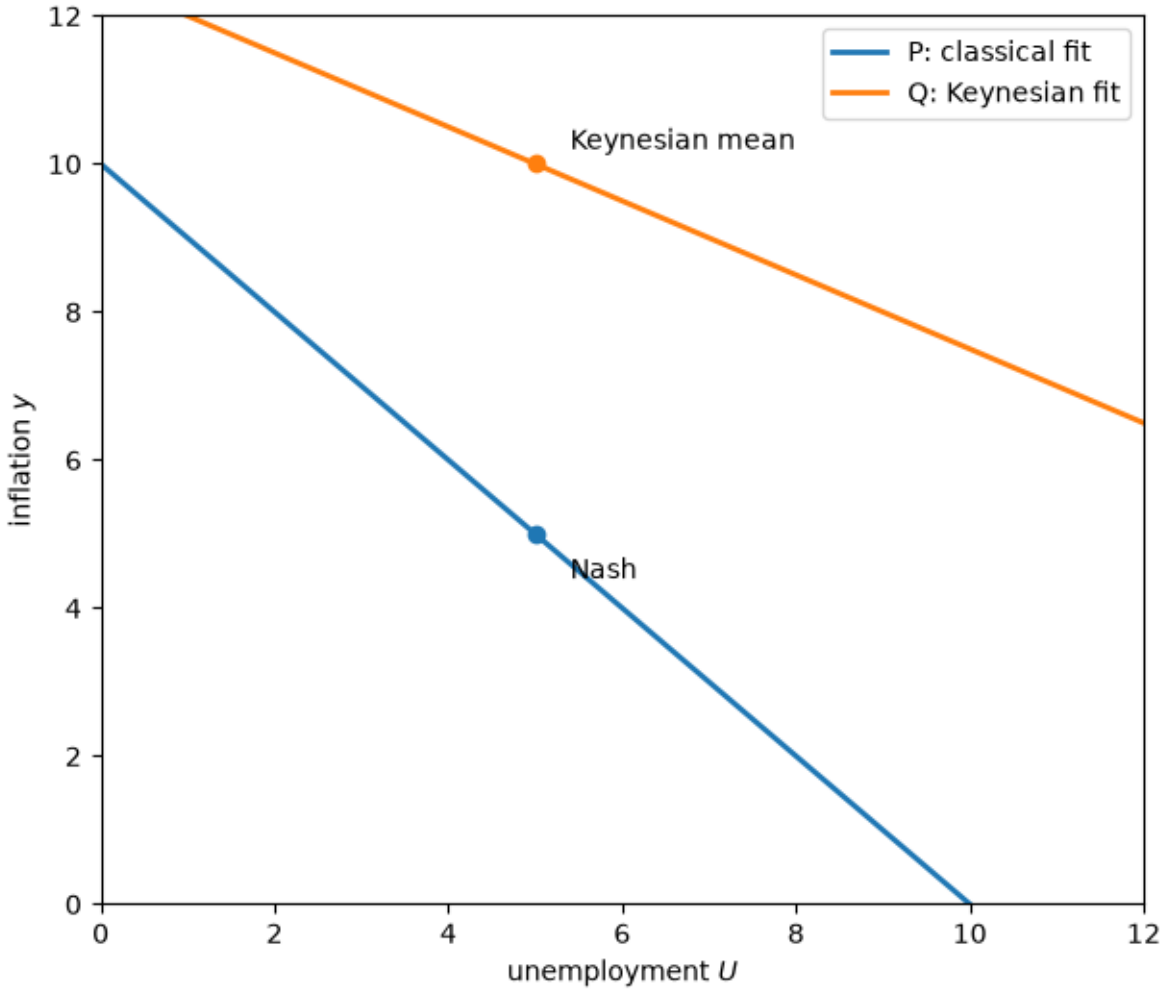


Fig. 102.1: The two self-confirming Phillips curves, one for each direction of fit

Curve P is the self-confirming Phillips curve estimated with the classical direction of fit; curve Q with the Keynesian direction.

Because the Keynesian fit makes the government believe the tradeoff is more exploitable, average inflation is higher.

102.5.1 Why the classical equilibrium is Nash

The classical self-confirming equilibrium reproduces the Nash outcome of *The Credibility Problem* for a simple reason.

Its perceived intercept is $\gamma_{-1} = (1 + \theta^2)U^* = U^* + \theta\bar{x}$, with $\bar{x} = \theta U^*$.

So the perceived curve $U_t = U^* + \theta\bar{x} - \theta y_t$ is the true Phillips curve evaluated at $x_t = \bar{x}$.

The government's model is right about everything it sees in equilibrium.

Its only mistake is to treat $\bar{x}$ as invariant to its own rule, which is exactly the assumption that defines the Nash government, which takes the public's expectations as given.

That is why the time-consistent, or "consistent", equilibrium of Kydland and Prescott [1977] is a self-confirming equilibrium (*A government learning the Phillips curve*).

Readers of Kalai and Lehrer [1993] might have foreseen this.

Learning converges to the truth when decision makers' probability models encompass the Nash equilibrium, and the classical identification scheme satisfies that encompassing condition.

The Keynesian scheme does not: it gives the government a mistaken interpretation of the contemporaneous covariance between innovations to inflation and unemployment, and that worsens outcomes relative to Nash.

Both conclusions rest on the orthogonality conditions in the actual Phillips curve: v_{1t} is uncorrelated with v_{2t}, and v_{2t} is orthogonal to $\widehat{y}_t$.

Different orthogonality conditions would change the quality of the outcomes in both equilibria, which is how, though our model of the government differs from theirs, we join King and Watson's inquiry.

102.5.2 Why not Ramsey?

Both self-confirming equilibria give a mean outcome *worse* than the Ramsey outcome.

This signifies a breakdown of the hypotheses of the proposition in *Adaptive Expectations and the Phelps Problem* that solving the Phelps problem eventually sustains nearly the Ramsey outcome.

What fails is the **induction hypothesis**.

Because $\rho_1 = \rho_2 = 0$, the self-confirming equilibria are serially uncorrelated, so lagged inflation drops out of the empirical Phillips curves.

Deactivating the induction hypothesis makes the government in effect solve a one-period problem and shuts down the intertemporal channel that promotes better outcomes.

The two equilibria illustrate the assertion of Lucas [1972] and Sargent [1971], discussed in *Testing the natural-rate hypothesis*, that rational expectations and the natural-rate hypothesis do not imply the induction hypothesis: both hold here, and the induction hypothesis fails.

102.6 Uniqueness and E-stability

Two properties of these equilibria matter for the lectures that follow: each is unique, and each is E-stable, so that least squares learning can converge to it.

Both are easy to establish in the static case.

Uniqueness. In the classical fit, the slope $T_1 = \text{cov}(U_t, y_t)/\mathbb{V}[y_t] = -\theta$ comes from moments (102.4) that do not depend on policy: the government's beliefs move the mean of inflation, but not its variance or its covariance with unemployment.

Given the slope, the intercept must satisfy

$$\gamma_{-1} = T_{-1}(\gamma) = U^* + \theta h(\gamma) = U^* + \frac{\theta^2}{1+\theta^2}\gamma_{-1},$$

a linear equation with slope $\theta^2/(1+\theta^2) < 1$, and so with a unique solution.

The Keynesian case is the same, with slope $1/(1+\beta_1^2)$ in the equation for the intercept.

E-stability. Write the classical equilibrium as a fixed point $\gamma = T(\gamma)$ of the map from beliefs to the population regression coefficients they generate.

The associated differential equation $d\gamma/d\tau = T(\gamma) - \gamma$ is the E-stability condition of *One recursion, many learners*, and *Least Squares Learning in Self-Referential Models* shows that its stability governs whether least squares learning converges.

Because T_1 does not depend on beliefs, the Jacobian of $T - I$ at the equilibrium is triangular, with eigenvalues -1 (slope) and $\theta^2/(1+\theta^2) - 1 = -1/(1+\theta^2)$ (intercept).

For the Keynesian fit $\beta = S(\beta)$, the eigenvalues are -1 and $1/(1+\beta_1^2) - 1 = -\beta_1^2/(1+\beta_1^2)$.

Let's confirm these by differentiating the two maps numerically.

```
def T_classical(γ, θ=1.0, U_star=5.0):
    "Static classical T-map: beliefs (γ_1, γ_{-1}) ↦ regression of U on (y, 1)."
    γ1, γ_int = γ
    y_hat = -γ1 * γ_int / (1 + γ1**2)
    return np.array([-θ, U_star + θ * y_hat])

def S_keynesian(β, θ=1.0, U_star=5.0, σ1=0.3, σ2=0.3):
    "Static Keynesian S-map: beliefs (β_1, β_{-1}) ↦ regression of y on (U, 1)."
    β1, β_int = β
    y_hat = β_int / (1 + β1**2)
    slope = -θ * σ2**2 / (σ1**2 + θ**2 * σ2**2)
    return np.array([slope, y_hat - slope * U_star])

def jacobian(f, x, ε=1e-6):
    "Central-difference Jacobian of f at x."
    J = np.empty((len(x), len(x)))
    for j in range(len(x)):
        dx = np.zeros(len(x))
        dx[j] = ε
        J[:, j] = (f(x + dx) - f(x - dx)) / (2 * ε)
    return J

for θ in [0.5, 1.0, 2.0]:
    γ_sce = np.array([-θ, (1 + θ**2) * 5.0])
    eig = np.linalg.eigvals(jacobian(lambda g: T_classical(g, θ), γ_sce) - np.eye(2))
    print(f"classical, θ = {θ}:  eigenvalues {np.sort(eig).round(4)},"
          f"  -1/(1+θ²) = {-1 / (1 + θ**2):.4f}")

β1 = -0.3**2 / (0.3**2 + 0.3**2)
β_sce = np.array([β1, -(β1**2 + 1) / β1 * 5.0])
eig = np.linalg.eigvals(jacobian(S_keynesian, β_sce) - np.eye(2))
print(f"Keynesian, θ = 1:  eigenvalues {np.sort(eig).round(4)},"
      f"  -β1²/(1+β1²) = {-β1**2 / (1 + β1**2):.4f}")
```

```
classical, θ = 0.5:  eigenvalues [-1.  -0.8],  -1/(1+θ²) = -0.8000
classical, θ = 1.0:  eigenvalues [-1.  -0.5],  -1/(1+θ²) = -0.5000
classical, θ = 2.0:  eigenvalues [-1.  -0.2],  -1/(1+θ²) = -0.2000
Keynesian, θ = 1:  eigenvalues [-1.  -0.2],  -β1²/(1+β1²) = -0.2000
```

Both equilibria are E-stable.

The classical eigenvalue $-1/(1+\theta^2)$ is the slope $\mathcal{M}$ of the mean dynamics in *Least squares learning converges to Nash*, where the public learned expected inflation by least squares.

At the suite's $\theta = 1$ it equals $-1/2$, exactly the boundary below which decreasing-gain least squares converges at the $\sqrt{t}$ rate.

The Keynesian eigenvalue, -0.2 here, lies above that boundary, so least squares converges more slowly, like $t^{-0.2}$.

These rates describe convergence near the equilibrium.

They do not govern how often a learning government escapes from it, which depends instead on the large-deviation analysis of *A primer on recursive algorithms*; Exercise 102.8.2 shows the difference.

Lags. Do these results depend on $X_{t-1} = 1$?

Let the government's regression include one lag of each variable, so that $X_{t-1} = \begin{bmatrix} U_{t-1} & y_{t-1} & 1 \end{bmatrix}^\top$ and $\gamma = \begin{bmatrix} \gamma_1 & \gamma_{-1}^\top \end{bmatrix}^\top$ has four elements.

Now the government's rule $h(\gamma)$ comes from a genuinely dynamic Phelps problem.

With $\rho_1 = \rho_2 = 0$ the population regression is still exact.

Because $x_t = hX_{t-1}$ and v_{1t} is orthogonal to every regressor,

$$U_t = -\theta y_t + \left[\theta h(\gamma) + \begin{bmatrix} 0 & 0 & U^* \end{bmatrix}\right] X_{t-1} + v_{1t},$$

so $T_1(\gamma) = -\theta$ and $T_{-1}(\gamma) = \theta h(\gamma)^\top + \begin{bmatrix} 0 & 0 & U^* \end{bmatrix}^\top$.

We compute $h(\gamma)$ with the general Phelps solver of *Adaptive Expectations and the Phelps Problem*, repeated here so that this lecture runs on its own.

We check that the Nash beliefs $\gamma_N = \begin{bmatrix} -\theta & 0 & 0 & (1+\theta^2)U^* \end{bmatrix}^\top$ are a fixed point, compute the eigenvalues of the Jacobian of $T - I$, and run the relaxation algorithm from random starting beliefs.

```
def phelps_general(γ, δ):
    """
    The general Phelps problem of phillips_adaptive for the believed curve
    U_t = γ[0] y_t + γ[1] U_{t-1} + γ[2] y_{t-1} + γ[3].
    Returns h, with ŷ_t = h @ [U_{t-1}, y_{t-1}, 1].
    """
    γ_y = γ[0]
    c = np.asarray(γ[1:], dtype=float).reshape(-1, 1)
    R = 0.5 * (c @ c.T)
    Q = np.array([[0.5 * (γ_y**2 + 1)]])
    N = 0.5 * γ_y * c.T
    A = np.vstack([c.T, [0.0, 0.0, 0.0], [0.0, 0.0, 1.0]])
    B = np.array([[γ_y], [1.0], [0.0]])
    lq = qe.LQ(Q, R, A, B, N=N, beta=min(δ, 1 - 1e-7))
    P, F, d = lq.stationary_values()
    return -F[0]

def T_one_lag(γ, δ, θ=1.0, U_star=5.0):
    "Regression of U_t on (y_t, U_{t-1}, y_{t-1}, 1) when ŷ_t = h(γ) X_{t-1}."
    h = phelps_general(γ, δ)
    return np.concatenate([[-θ], θ * h + np.array([0.0, 0.0, U_star])])

for θ in [0.5, 1.0, 2.0]:
    γ_N = np.array([-θ, 0.0, 0.0, (1 + θ**2) * 5.0])
    for δ in [0.9, 0.98]:
```

(continues on next page)

(continued from previous page)

```
        J = jacobian(lambda g: T_one_lag(g, δ, θ), y_N)
        eig = np.sort(np.linalg.eigvals(J - np.eye(4)).real)
        gap = np.max(np.abs(T_one_lag(y_N, δ, θ) - y_N))
        print(f"θ = {θ}, δ = {δ}:  |T(y_N) - y_N| = {gap:.1e},  "
              f"eigenvalues {eig.round(4)}")

rng = np.random.default_rng(0)
κ, δ = 0.5, 0.96
y_N = np.array([-1.0, 0.0, 0.0, 10.0])
limits = []
for trial in range(10):
    y = y_N + np.array([0.0, rng.uniform(-0.9, 0.9),
                        rng.uniform(-1.5, 1.5), rng.uniform(-5, 5)])
    for _ in range(2_000):
        y_new = κ * y + (1 - κ) * T_one_lag(y, δ)
        if np.max(np.abs(y_new - y)) < 1e-10:
            break
        y = y_new
    limits.append(y_new)
dist = max(np.max(np.abs(g - y_N)) for g in limits)
print(f"\nrelaxation from 10 random starts: largest distance of the limit from y_N =
↪{dist:.1e}")
```

```
θ = 0.5, δ = 0.9:  |T(y_N) - y_N| = 0.0e+00,  eigenvalues [-1.  -0.8 -0.8 -0.8]
θ = 0.5, δ = 0.98:  |T(y_N) - y_N| = 0.0e+00,  eigenvalues [-1.  -0.8 -0.8 -0.8]
θ = 1.0, δ = 0.9:  |T(y_N) - y_N| = 0.0e+00,  eigenvalues [-1.  -0.5 -0.5 -0.5]
θ = 1.0, δ = 0.98:  |T(y_N) - y_N| = 0.0e+00,  eigenvalues [-1.  -0.5 -0.5 -0.5]
θ = 2.0, δ = 0.9:  |T(y_N) - y_N| = 0.0e+00,  eigenvalues [-1.  -0.2 -0.2 -0.2]
θ = 2.0, δ = 0.98:  |T(y_N) - y_N| = 0.0e+00,  eigenvalues [-1.  -0.2 -0.2 -0.2]

relaxation from 10 random starts: largest distance of the limit from y_N = 3.0e-10
```

The Nash beliefs are a fixed point for every θ and δ, and the eigenvalues are -1 and $-1/(1+\theta^2)$, the latter repeated three times, whatever the discount factor.

Relaxation from every starting point converges to the Nash beliefs, with zero weights on the lags.

Adding lags gives the government room to discover the induction hypothesis, but in a self-confirming equilibrium the lag weights are zero and the hypothesis stays deactivated.

Adaptive Learning and Escape Dynamics finds the same pair of eigenvalues with two lags of each variable (*Mean dynamics near the self-confirming equilibrium*).

102.7 Equilibrium with misspecified beliefs

To explore how imputing a *different* wrong model might *improve* on the Nash outcome, we now let the *public*, not the government, make a subtle specification error.

This links adaptive expectations, the induction hypothesis, and the Phelps problem.

The government knows the correct model,

$$U_t = U^* - \theta(y_t - x_t) + v_{1t}, \qquad x_t = Cy_{t-1} + (1-C)x_{t-1}, \quad C \in (0,1), \tag{102.6}$$

where the public has constant-gain adaptive expectations with a parameter C that it *tunes to fit the data*.

Taking x_t as a state variable, the government solves the Phelps problem: it maximizes $-\mathbb{E}_0 \sum_{t=0}^{\infty} \delta^t \left[(U^* - \theta(y_t - x_t))^2 + y_t^2\right]$ by choice of a feedback rule $y_t = f_1 + f_2 x_t + v_{2t}$.

This is exactly the LQ Phelps problem of *Adaptive Expectations and the Phelps Problem* with adaptation parameter $\lambda = 1 - C$; the code below repeats it so that this lecture runs on its own.

```
def phelps_policy(C, δ, θ=1.0, U_star=5.0):
    "Government feedback rule y_t = f1 + f2 x_t, with public gain C."
    λ = 1 - C
    a = np.array([[U_star], [θ]])
    R = 0.5 * (a @ a.T)
    Q = np.array([[0.5 * (θ**2 + 1)]])
    N = -0.5 * θ * a.T
    A = np.array([[1.0, 0.0], [0.0, λ]])
    B = np.array([[0.0], [1 - λ]])
    lq = qe.LQ(Q, R, A, B, N=N, beta=min(δ, 1 - 1e-7))
    P, F, d = lq.stationary_values()
    return -F[0, 0], -F[0, 1]
```

The government's behavior makes the *actual* inflation rate

$$y_t = \frac{f_1}{1 - f_2} + \frac{1 - (1 - C)L}{1 - (1 - C(1 - f_2))L} v_{2t} = \nu + f(L) v_{2t}, \tag{102.7}$$

with mean $\nu = f_1/(1 - f_2)$ and spectrum $F(\omega; C) = |f(e^{i\omega})|^2 \sigma_2^2$.

The mean ν is the steady state of the Phelps rule, so the closed form $y_\infty = \theta U^*(1 - \delta)/(1 - \delta\lambda)$ derived in *Adaptive Expectations and the Phelps Problem*, with $\lambda = 1 - C$, gives it:

$$\nu = \theta U^* \frac{1 - \delta}{1 - \delta + \delta C}. \tag{102.8}$$

Mean inflation is below the Nash rate θU^* whenever $C > 0$; it falls toward zero as $\delta \to 1$ with C fixed, and it rises to θU^* as $C \to 0$.

Everything therefore depends on the equilibrium value of C.

Given C, the public seeks the best-fitting misspecified model of the same form as in *Optimal Misspecified Beliefs*, the stationary ARMA(1,1) model

$$y_t = \frac{1 - (1 - c)L}{1 - \rho L} \epsilon_t = g(L) \epsilon_t,$$

whose root ρ slightly below one (not to be confused with the roots ρ_1, ρ_2 of (102.2)) lets it fit the mean ν with near-unit-root persistence.

The induced best-estimate map $c = B(C)$ defines an **equilibrium under forecast misspecification** as a fixed point $C = B(C)$.

This makes both the true and the approximating models equilibrium outcomes, and turns the adaptation parameter C into an outcome rather than a free parameter.

The best-estimate code is the same as in *Optimal Misspecified Beliefs*.

```
def forecast_mse(F, ν, c, ρ, z):
    """
    Mean square one-step forecast error of the perceived model
    g(L) = (1 - (1-c)L)/(1 - ρL) when the true process has mean ν and
    spectral density F on the grid z = exp(iω).  ρ = 1 is allowed.
    """
```

(continues on next page)

(continued from previous page)

```
    W = np.abs((1 - ρ * z) / (1 - (1 - c) * z))**2
    return np.mean(F * W) + ν**2 * W[0]

class MisspecifiedPhillips:
    """
    Equilibrium with misspecified public beliefs: a fixed point of the
    best-estimate map, with the government solving the Phelps problem.
    """

    def __init__(self, θ=1.0, U_star=5.0, σ2=1.0, δ=0.97, ρ=0.995, N=1024):
        self.θ, self.U_star, self.σ2 = θ, U_star, σ2
        self.δ, self.ρ, self.N = δ, ρ, N
        ω = 2 * np.pi * np.arange(N) / N
        self.z = np.exp(1j * ω)
        self.ω = ω

    def true_process(self, C):
        "Return the mean ν, spectrum F, and policy (f1, f2) given belief C."
        f1, f2 = phelps_policy(C, self.δ, self.θ, self.U_star)
        ν = f1 / (1 - f2)
        ψ = 1 - C * (1 - f2)
        f = (1 - (1 - C) * self.z) / (1 - ψ * self.z)
        F = np.abs(f)**2 * self.σ2**2
        return ν, F, (f1, f2)

    def best_estimate(self, C):
        "The best-estimate map c = B(C)."
        ν, F, _ = self.true_process(C)
        return minimize_scalar(lambda c: forecast_mse(F, ν, c, self.ρ, self.z),
                               bounds=(1e-4, 0.99), method='bounded',
                               options={'xatol': 1e-10}).x

    def solve(self, C0=0.1, tol=1e-10, maxit=500):
        "Iterate the best-estimate map to a fixed point."
        C = C0
        for _ in range(maxit):
            C_new = self.best_estimate(C)
            if abs(C_new - C) < tol:
                break
            C = C_new
        ν, _, (f1, f2) = self.true_process(C_new)
        return C_new, ν, (f1, f2)
```

The size of the inflation surprise.

Chapter 7 of Sargent [1999] states $\sigma_1 = \sigma_2 = 0.3$ for this example.

Its reported numbers are reproduced with $\sigma_2 = 1$, which is also consistent with its statement that the true model's one-step prediction error has standard deviation 1.

We use $\sigma_2 = 1$.

The choice matters because the fitting criterion is not scale-invariant.

The criterion of *Optimal Misspecified Beliefs* equals σ_2^2 times a function of c that depends on σ_2 only through the ratio ν/σ_2 of the mean that the perceived model must fit to the size of the surprises.

A smaller σ_2 makes the mean relatively larger, which pushes the public toward a larger gain and so, by (102.8), lowers mean inflation.

```
mp = MisspecifiedPhillips(δ=0.97)
C_star, ν_star, (f1, f2) = mp.solve()

print(f"equilibrium gain     C = {C_star:.4f}")
print(f"policy rule          y = {f1:.4f} + {f2:.4f} x")
print(f"mean inflation       ν = {ν_star:.3f}   (formula: "
      f"{mp.θ * mp.U_star * (1 - mp.δ) / (1 - mp.δ + mp.δ * C_star):.3f})")
print(f"Nash inflation   θ U*  = {mp.θ * mp.U_star:.3f}")

C_03, ν_03, _ = MisspecifiedPhillips(δ=0.97, σ2=0.3).solve()
print(f"\nwith σ2 = 0.3:  C = {C_03:.4f},  mean inflation = {ν_03:.3f}")
```

```
equilibrium gain     C = 0.0673
policy rule          y = 1.0242 + 0.3495 x
mean inflation       ν = 1.574   (formula: 1.574)
Nash inflation   θ U*  = 5.000

with σ2 = 0.3:  C = 0.1138,  mean inflation = 1.069
```

With $\sigma_2 = 1$, these are the values that chapter 7 of Sargent [1999] reports.

The most important outcome is that the implied mean inflation rate is substantially *below* the Nash value of 5.

The induction hypothesis embedded in the adaptive expectations scheme, together with a high discount factor, delivers this improvement.

Because C is now an equilibrium outcome, the mechanism is sharper than in *Adaptive Expectations and the Phelps Problem*, where $\lambda = 1 - C$ could be manipulated independently of δ.

As in *Optimal Misspecified Beliefs*, the equilibrium is E-stable.

```
h = 1e-4
B_prime = (mp.best_estimate(C_star + h) - mp.best_estimate(C_star - h)) / (2 * h)
print(f"B'(C*) = {B_prime:.3f}")
```

```
B'(C*) = -0.305
```

Here $B'(C^*)$ is negative but greater than -1, so the iteration of B converges, oscillating around C^*.

One reason for the negative slope is that a higher gain makes the government disinflate further, which lowers the mean that the public's model must fit and so lowers the gain the public wants.

102.7.1 Approaching Ramsey

As the government becomes more patient, the equilibrium mean inflation rate falls toward the Ramsey value of zero.

```
δ_grid = np.array([0.95, 0.96, 0.97, 0.98, 0.99, 0.995])
C_vals, ν_vals = [], []
for δ in δ_grid:
    C, ν, _ = MisspecifiedPhillips(δ=δ).solve()
    C_vals.append(C)
    ν_vals.append(ν)

fig, axes = plt.subplots(1, 2, figsize=(11, 4.5))
axes[0].plot(δ_grid, ν_vals, 'o-')
axes[0].axhline(mp.θ * mp.U_star, color='k', ls='--', lw=1, label='Nash')
axes[0].set_xlabel(r'discount factor $\delta$')
```

(continues on next page)

(continued from previous page)

```
axes[0].set_ylabel('mean inflation')
axes[0].legend()

axes[1].plot(δ_grid, C_vals, 'o-', color='C1', lw=2)
axes[1].set_xlabel(r'discount factor $\delta$')
axes[1].set_ylabel('equilibrium gain $C$')

plt.tight_layout()
plt.show()

print(f"{'δ':>6} {'C':>8} {'ν':>7}")
for δ, C, ν in zip(δ_grid, C_vals, ν_vals):
    print(f"{δ:>6} {C:>8.4f} {ν:>7.2f}")
ν_book = mp.θ * mp.U_star * 0.01 / (0.01 + 0.99 * 0.019)
print(f"\nformula (sc_nu) at the book's δ = 0.99 entry, C = 0.019:  ν = {ν_book:.2f}")
```

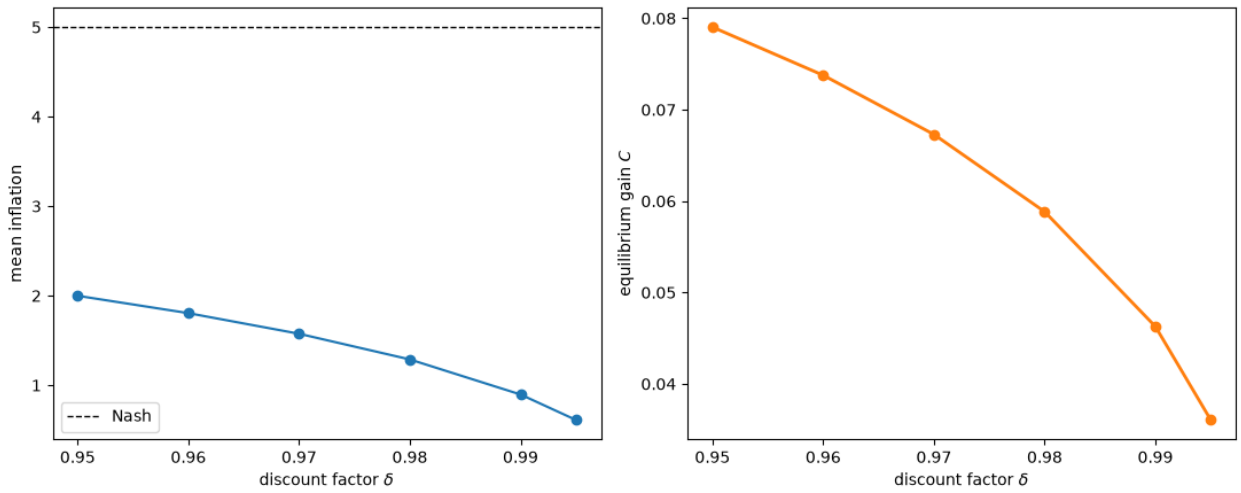


Fig. 102.2: Mean inflation and the equilibrium gain as the government becomes more patient

```
     δ        C       ν
  0.95   0.0790    2.00
  0.96   0.0738    1.80
  0.97   0.0673    1.57
  0.98   0.0588    1.29
  0.99   0.0463    0.90
 0.995   0.0361    0.61

formula (sc_nu) at the book's δ = 0.99 entry, C = 0.019:  ν = 1.74
```

Mean inflation lies far below the Nash value at every discount factor and declines toward the Ramsey value of zero as $\delta \to 1$.

The values of C and ν match footnote 14 of chapter 7 of Sargent [1999] for $\delta = 0.95, 0.96$, and 0.98.

The book's entry for $\delta = 0.99$ ($C = 0.019$, mean inflation 0.25) is inconsistent with (102.8), which, as printed above, gives a much higher mean inflation at $C = 0.019$.

102.7.2 The role of ρ

Every number above was computed with the public's perceived root $\rho = 0.995$.

How the equilibrium depends on \rho showed that in Bray's model the equilibrium gain vanishes as $\rho \to 1$.

The next cell asks what happens here, for both values of σ_2, again using a finer frequency grid.

```
print(f"{'ρ':>7} {'C (σ2=1)':>10} {'ν (σ2=1)':>10} {'C (σ2=.3)':>11} {'ν (σ2=.3)':>11}
 ↪")
for ρ in [0.99, 0.995, 0.999, 0.9995]:
    row = []
    for σ2 in [1.0, 0.3]:
        C, ν, _ = MisspecifiedPhillips(δ=0.97, σ2=σ2, ρ=ρ, N=2**14).solve()
        row += [C, ν]
    print(f"{ρ:>7} {row[0]:>10.4f} {row[1]:>10.3f} {row[2]:>11.4f} {row[3]:>11.3f}")

m1 = MisspecifiedPhillips(δ=0.97, ρ=1.0, N=2**14)
ratio = max(m1.best_estimate(C) / C for C in np.linspace(0.005, 0.3, 30))
print(f"\nρ = 1: largest value of B(C)/C on the grid = {ratio:.3f}")
```

```
      ρ   C (σ2=1)    ν (σ2=1)    C (σ2=.3)    ν (σ2=.3)
   0.99     0.0918       1.260       0.1527       0.842
  0.995     0.0673       1.574       0.1138       1.069
  0.999     0.0316       2.474       0.0558       1.783
 0.9995     0.0223       2.905       0.0404       2.169

ρ = 1: largest value of B(C)/C on the grid = 0.311
```

Mean inflation rises toward the Nash value as $\rho \to 1$, for both values of σ_2.

With the exact unit root, $B(C) < C$ for every C, so the equilibrium gain is zero, the public's expectations stop moving, and by (102.8) mean inflation is the Nash rate θU^*.

So the better-than-Nash outcomes hold for each fixed $\rho < 1$, and vanish as $\rho \to 1$.

For each fixed $\rho < 1$ they approach Ramsey as $\delta \to 1$, but for each fixed $\delta < 1$ they disappear as $\rho \to 1$ (Exercise 102.8.3 studies the two limits).

The reason is the one given in *How the equilibrium depends on \rho*: the pull toward a large gain comes from the mean bias that a perceived model with $\rho < 1$ would leave, and it disappears with the unit root that the public's forecasting rule actually embodies.

102.7.3 Spectra and impulse responses

Let's compare the true and approximating inflation processes at the equilibrium, as in Figures 7.2 and 7.3 of Sargent [1999].

```
ν_star, F, _ = mp.true_process(C_star)
c_star = mp.best_estimate(C_star)
σ_ε2 = forecast_mse(F, ν_star, c_star, mp.ρ, mp.z)
H = np.abs((1 - (1 - c_star) * mp.z) / (1 - mp.ρ * mp.z))**2
G = H * σ_ε2

half = mp.N // 2
fig, ax = plt.subplots(figsize=(8, 5))
ax.plot(mp.ω[:half], np.log(F[:half]), 'C0', label='true model', lw=2)
ax.plot(mp.ω[:half], np.log(G[:half]), 'C1--', label='approximating model', lw=2)
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel(r'angular frequency $\omega$')
ax.set_ylabel('log spectral density')
ax.legend()
plt.show()
```

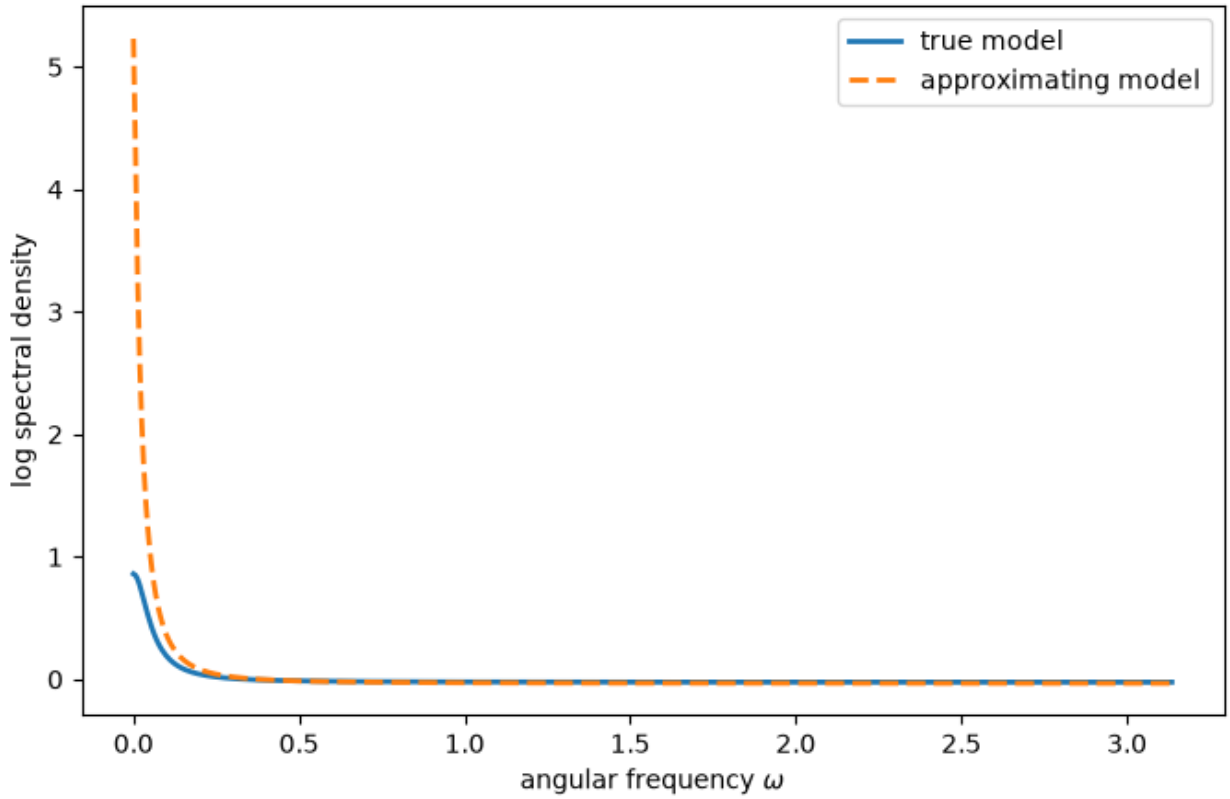


Fig. 102.3: Spectral densities of the true and approximating inflation processes at the equilibrium

The true and approximating spectral densities match well at all but the lowest frequencies.

The true inflation rate is only moderately serially correlated, and — as in the Bray model of *Optimal Misspecified Beliefs* — the approximating model uses a near unit root to simulate a mean, capturing first moments with second moments.

```
def ima_impulse(num, den, T=25):
    "IRF of (1 - num L)/(1 - den L)."
    h = np.empty(T)
    h[0] = 1.0
    for k in range(1, T):
        h[k] = den * h[k - 1]
    h[1:] -= num * h[:-1]
    return h

ψ = 1 - C_star * (1 - f2)
irf_true = ima_impulse(1 - C_star, ψ)
irf_approx = ima_impulse(1 - c_star, mp.ρ)
```

(continues on next page)

(continued from previous page)

```
lags = np.arange(len(irf_true))
fig, axes = plt.subplots(1, 2, figsize=(11, 4.5))
for ax, start in zip(axes, [0, 1]):
    ax.plot(lags[start:], irf_true[start:], 'C0o-', ms=4, label='true model', lw=2)
    ax.plot(lags[start:], irf_approx[start:], 'C1s--', ms=4,
            label='approximating model', lw=2)
    ax.set_xlabel('lag')
    ax.set_ylabel('response')
axes[0].set_title('all lags')
axes[1].set_title('lags 1 and beyond')
axes[0].legend()
plt.tight_layout()
plt.show()
```

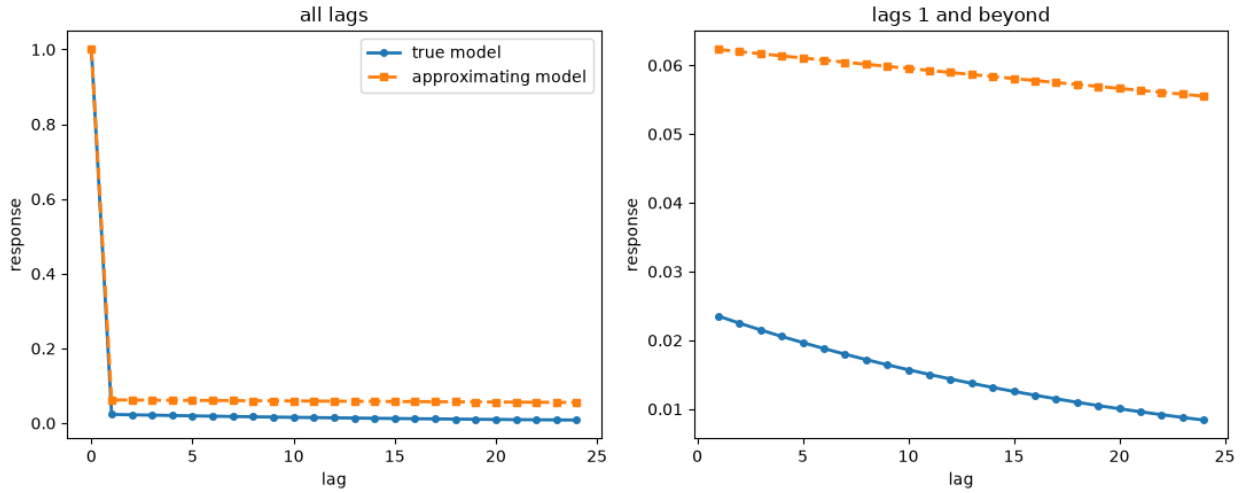


Fig. 102.4: Impulse responses of the true and approximating inflation processes

Here both models respond to a unit one-step-ahead forecast error, since $f(0) = g(0) = 1$.

The near unit root of the approximating model shows up as a response that dies out only at rate ρ, which over the horizon plotted looks like a nonzero asymptote.

102.7.4 Grounds for optimism

After the disappointments of the self-confirming equilibria, the equilibrium with forecast misspecification is heartening: it supports better-than-Nash outcomes.

The equilibrium concept is not self-confirming, but it has that spirit — it embodies a type of self-confirmation with a *wrong* model.

The approximation error is small.

```
print(f"one-step RMSE of the approximating model / true innovation s.d. = "
      f"{np.sqrt(σ_ε2) / mp.σ2:.4f}")
```

```
one-step RMSE of the approximating model / true innovation s.d. = 1.0155
```

The approximating model's one-step forecast errors are only slightly larger than the true innovations, so there is a nearly self-confirming model with much better than Nash outcomes.

But the optimism is conditional.

The good outcomes rest on the public's perceived model having a root ρ strictly below one; with the exact unit root that its forecasting rule embodies, the equilibrium gain is zero and the outcome is Nash.

What the example shows is that a small, hard-to-detect misspecification in the public's model can matter a great deal for outcomes.

The next lectures in this suite — *Adaptive Learning and Escape Dynamics*, *Escaping Nash Inflation*, and *Priors, Escapes, and Learning Cycles* — build adaptive, real-time versions of these models in which the government estimates its Phillips curve recursively from the most recent data.

There the classical (Nash) self-confirming equilibrium becomes the attractor of the learning dynamics.

With a constant gain, the government escapes from it recurrently toward Ramsey and is pulled back, and better-than-Nash outcomes reappear during the escapes.

Those escapes run through the *government's* beliefs, not through a misspecification of the public's.

102.8 Exercises

Exercise 102.8.1

The gap between the two static self-confirming equilibria is driven by the variances σ_1^2 (of the Phillips curve shock) and σ_2^2 (of the inflation surprise).

The classical equilibrium always delivers mean inflation equal to the Nash value θU^*, but the Keynesian mean inflation depends on the ratio σ_1/σ_2.

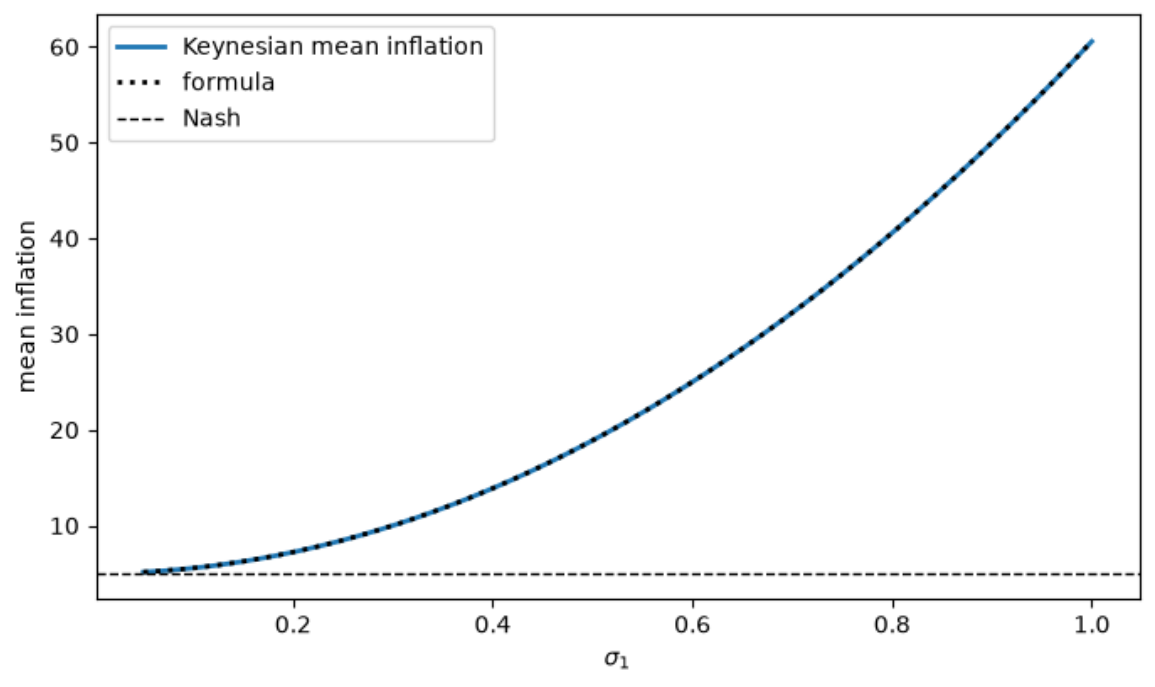
Fixing $\sigma_2 = 0.3$, $\theta = 1$, $U^* = 5$, plot Keynesian mean inflation as a function of $\sigma_1 \in [0.05, 1.0]$, and compare it with formula (102.5).

What happens as $\sigma_1 \to 0$, and why?

Solution

```
σ1_grid = np.linspace(0.05, 1.0, 50)
y_keynes = [SelfConfirmingStatic(σ1=σ1, σ2=0.3).keynesian()[2]
            for σ1 in σ1_grid]
y_formula = 5.0 * (1 + σ1_grid**2 / 0.3**2)

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(σ1_grid, y_keynes, lw=2, label='Keynesian mean inflation')
ax.plot(σ1_grid, y_formula, 'k:', lw=2, label='formula')
ax.axhline(5.0, color='k', ls='--', lw=1, label='Nash')
ax.set_xlabel(r'$\sigma_1$')
ax.set_ylabel('mean inflation')
ax.legend()
plt.show()
```



The computed values lie on the formula $\theta U^*(1 + \sigma_1^2/(\theta^2\sigma_2^2))$.

As $\sigma_1 \to 0$ the Phillips curve becomes nearly deterministic, the Keynesian slope $\beta_1 \to -1/\theta$, and the two directions of fit agree, so Keynesian mean inflation approaches the Nash value.

As σ_1 grows, unemployment becomes noisier, the Keynesian regression of y on U flattens, and the government perceives a more exploitable tradeoff, raising mean inflation quadratically in σ_1.

Exercise 102.8.2

Self-confirming equilibria are rest points of least squares learning.

Here is what decreasing-gain learning does in the static model, in which the government re-estimates its Phillips curve every period.

Each period the government sets $\hat{y}_t$ by (102.3) at its current beliefs, the public expects $x_t = \hat{y}_t$, (U_t, y_t) are generated by the actual Phillips curve with $\rho_1 = \rho_2 = 0$, and the government updates its regression by recursive least squares with gain $1/(t + t_0)$.

The value $t_0 = 100$ represents a prior sample of 100 observations, and the regression's moment matrix starts at its population value in the self-confirming equilibrium.

1. For each direction of fit, simulate 500 paths of 20,000 periods, starting at the respective self-confirming equilibrium.
2. For the Keynesian fit, show that the root mean square deviation of $\hat{y}_t$ from its equilibrium value 10 shrinks roughly like $t^{-0.2}$, as its E-stability eigenvalue predicts.
3. For the classical fit, compute the fraction of paths whose target ends more than one point away from the Nash value 5.
4. Explain the difference, using the moments of (U_t, y_t) when the target $\hat{y}_t$ moves over time.

Solution

```
def rls_static(direction, n_paths=500, T=20_000, t0=100,
               θ=1.0, U_star=5.0, σ1=0.3, σ2=0.3, seed=0):
    """
    Decreasing-gain recursive least squares in the static model.
    Beliefs are (intercept, slope).  Returns the targets ŷ_t, shape
    (n_paths, T), and the target in the self-confirming equilibrium.
    """
    rng = np.random.default_rng(seed)
    if direction == 'classical':                    # regress U on (1, y)
        y_sce = θ * U_star
        b0 = np.array([(1 + θ**2) * U_star, -θ])
        R0 = np.array([[1.0, y_sce], [y_sce, y_sce**2 + σ2**2]])
    else:                                           # regress y on (1, U)
        β1 = -θ * σ2**2 / (σ1**2 + θ**2 * σ2**2)
        y_sce = -U_star / β1
        b0 = np.array([(1 + β1**2) * y_sce, β1])
        R0 = np.array([[1.0, U_star],
                       [U_star, U_star**2 + θ**2 * σ2**2 + σ1**2]])
    b = np.tile(b0, (n_paths, 1))
    R = np.tile(R0, (n_paths, 1, 1))
    targets = np.empty((n_paths, T))
    ones = np.ones(n_paths)

    for t in range(T):
        intercept, slope = b[:, 0], b[:, 1]
        if direction == 'classical':
            y_hat = -slope * intercept / (1 + slope**2)
        else:
            y_hat = intercept / (1 + slope**2)
        targets[:, t] = y_hat

        v1 = σ1 * rng.standard_normal(n_paths)
        v2 = σ2 * rng.standard_normal(n_paths)
        y = y_hat + v2
        U = U_star - θ * v2 + v1
        if direction == 'classical':
            X, lhs = np.column_stack([ones, y]), U
        else:
            X, lhs = np.column_stack([ones, U]), y

        g = 1 / (t + 1 + t0)
        R = R + g * (X[:, :, None] * X[:, None, :] - R)
        err = lhs - np.sum(X * b, axis=1)
        b = b + g * np.linalg.solve(R, (X * err[:, None])[:, :, None])[:, :, 0]

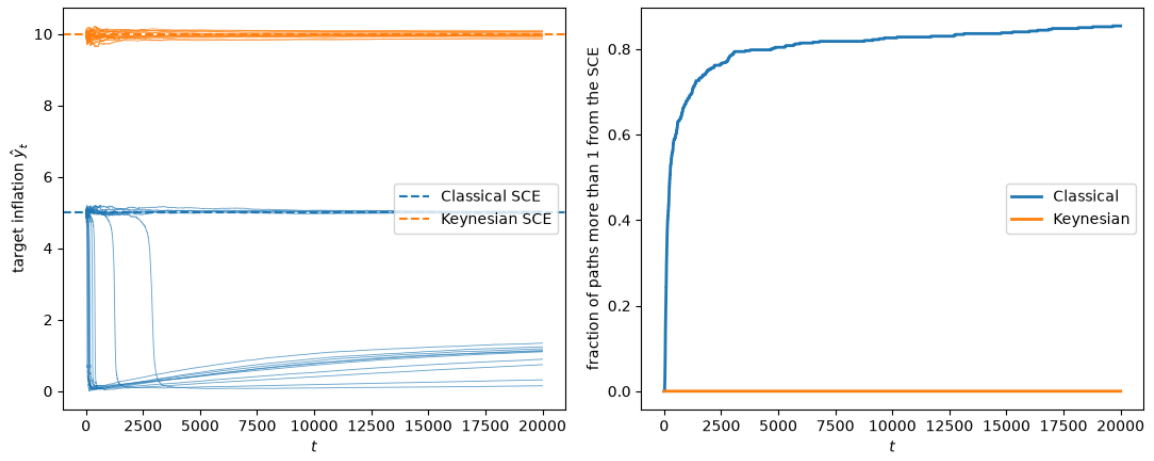
    return targets, y_sce

results = {d: rls_static(d) for d in ['classical', 'keynesian']}
t_axis = np.arange(1, 20_001)

fig, axes = plt.subplots(1, 2, figsize=(11, 4.5))
for d, color in [('classical', 'C0'), ('keynesian', 'C1')]:
    targets, y_sce = results[d]
    axes[0].plot(t_axis, targets[:15].T, color=color, lw=0.6, alpha=0.7)
    axes[0].axhline(y_sce, color=color, ls='--', lw=1.5, label=f'{d.capitalize()}␣
↪SCE')
```

```
    far = np.mean(np.abs(targets - y_sce) > 1.0, axis=0)
    axes[1].plot(t_axis, far, color=color, lw=2, label=d.capitalize())
axes[0].set_xlabel('$t$')
axes[0].set_ylabel(r'target inflation $\hat y_t$')
axes[0].legend()
axes[1].set_xlabel('$t$')
axes[1].set_ylabel('fraction of paths more than 1 from the SCE')
axes[1].legend()
plt.tight_layout()
plt.show()

targets, y_sce = results['keynesian']
rms = np.sqrt(np.mean((targets - y_sce)**2, axis=0))
keep = slice(200, None)
slope = np.polyfit(np.log(t_axis[keep] + 100), np.log(rms[keep]), 1)[0]
print(f"Keynesian: slope of log RMS deviation on log t = {slope:.2f}  (eigenvalue -
↪0.2)")
for d in ['classical', 'keynesian']:
    targets, y_sce = results[d]
    print(f"{d}: fraction of paths more than 1 from the SCE at the end = "
          f"{np.mean(np.abs(targets[:, -1] - y_sce) > 1.0):.2f}")
```




```
Keynesian: slope of log RMS deviation on log t = -0.17  (eigenvalue -0.2)
classical: fraction of paths more than 1 from the SCE at the end = 0.85
keynesian: fraction of paths more than 1 from the SCE at the end = 0.00
```

The Keynesian learner converges slowly, at roughly the rate its eigenvalue predicts.

The classical learner, whose eigenvalue $-1/2$ promises faster convergence, does much worse: a large fraction of its paths escape toward Ramsey and are still far from Nash after 20,000 periods.

The reason lies in what the data show when the target moves.

Since only surprises move unemployment, a change in $\hat{y}_t$ moves y_t but not U_t.

In a regression of U on y, that variation adds to the variance of the regressor but not to its covariance with U, so it flattens the estimated slope toward zero.

A flatter perceived curve makes the Phelps rule (102.3) recommend lower inflation, which moves the target further and flattens the slope again: a self-reinforcing escape toward Ramsey.

In a regression of y on U, the same variation adds only to the variance of the dependent variable, and leaves the slope alone.

E-stability is a statement about small deviations from the equilibrium; escapes are large deviations, studied in *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation*.

With a decreasing gain, least squares still converges to the equilibrium asymptotically, but an escape that starts early can last a long time, because the falling gain also slows the return (*Least squares learning: a decreasing gain*).

Exercise 102.8.3

Formula (102.8) shows that mean inflation in the equilibrium with misspecified beliefs depends on δ and on the equilibrium gain C, which depends on δ and on the public's perceived root ρ.

1. Compute the equilibrium C and mean inflation ν for $\delta \in \{0.97, 0.99, 0.999\}$ and $\rho \in \{0.995, 0.9995, 1\}$, with $\sigma_2 = 1$.
2. What happens as $\delta \to 1$ with ρ fixed, and as $\rho \to 1$ with δ fixed?
3. What does this imply for the claim that the equilibrium approaches Ramsey as the government becomes patient?

Solution

```
def v_closed(C, δ, θ=1.0, U_star=5.0):
    "Mean inflation (sc_nu) given the public's gain C."
    return θ * U_star * (1 - δ) / (1 - δ + δ * C)

print(f"{'δ':>6} {'ρ':>7} {'C':>8} {'ν':>7}")
for δ in [0.97, 0.99, 0.999]:
    for ρ in [0.995, 0.9995]:
        C, ν, _ = MisspecifiedPhillips(δ=δ, ρ=ρ, N=2**14).solve()
        print(f"{δ:>6} {ρ:>7} {C:>8.4f} {ν:>7.3f}")
    m1 = MisspecifiedPhillips(δ=δ, ρ=1.0, N=2**14)
    ratio = max(m1.best_estimate(C) / C for C in np.linspace(0.005, 0.3, 20))
    print(f"{δ:>6} {1:>7} {0:>8.4f} {v_closed(0.0, δ):>7.3f}"
          f"   (largest B(C)/C = {ratio:.2f})")
```

```
     δ       ρ        C       ν
  0.97   0.995   0.0673   1.574
  0.97  0.9995   0.0223   2.905
  0.97       1   0.0000   5.000   (largest B(C)/C = 0.31)
  0.99   0.995   0.0463   0.896
  0.99  0.9995   0.0166   1.894
  0.99       1   0.0000   5.000   (largest B(C)/C = 0.23)
 0.999   0.995   0.0199   0.239
 0.999  0.9995   0.0075   0.587
 0.999       1   0.0000   5.000   (largest B(C)/C = 0.04)
```

For fixed $\rho < 1$, mean inflation falls toward zero as δ rises, because the equilibrium gain falls much more slowly than $1 - \delta$ in (102.8).

For fixed δ, mean inflation rises toward the Nash value as $\rho \to 1$, because the equilibrium gain falls to zero, and at $\rho = 1$ it is exactly θU^* for every $\delta < 1$.

The two limits do not commute.

The claim that the equilibrium approaches Ramsey as the government becomes patient holds for each fixed $\rho < 1$, but it is not uniform in ρ: the closer the public's perceived root is to the unit root that its rule embodies, the more

patient the government must be to get close to Ramsey.

CHAPTER

THREE

ADAPTIVE LEARNING AND ESCAPE DYNAMICS

Contents

- *Adaptive Learning and Escape Dynamics*
  - *Overview*
  - *A primer on recursive algorithms*
  - *The adaptive model*
  - *Mean dynamics near the self-confirming equilibrium*
  - *Least squares learning: a decreasing gain*
  - *Constant gain and escape dynamics*
  - *The escape route and the induction hypothesis*
  - *Relation to equilibria under forecast misspecification*
  - *Role of the discount factor*
  - *Anticipated utility*
  - *Conclusions*
  - *Exercises*

103.1 Overview

This lecture is the culmination of the theory built by Sargent [1999].

It follows chapter 8, the most ambitious chapter of the book.

Everything after it in this suite builds on it:

- *Escaping Nash Inflation* and *Priors, Escapes, and Learning Cycles* sharpen its analytics;
- *Econometric Policy Evaluation* follows the book's chapter 9 in taking the adaptive models to data;
- *Drifts and Volatilities* and *Drifts and Volatilities after 2000* put its central claim to an empirical test; and
- *The Lost Conquest: Fed Policy in the 2020s* carries its tools to a later episode.

In *Self-Confirming Equilibria* a government held *fixed* beliefs about the Phillips curve — beliefs that were confirmed by the data those beliefs generated.

Here we make the government a real-time econometrician.

Each period it

- re-estimates its Phillips curve by recursive least squares from the data seen so far, and
- sets inflation at the first-period recommendation of the *Phelps problem* for its *current* estimate.

We ask whether such an adaptive government converges to a self-confirming equilibrium.

The answer depends on a single parameter — the **gain** that governs how fast old data are discounted:

- With a *decreasing* gain that implements least squares, the mean dynamics pull the economy toward a self-confirming equilibrium *asymptotically*, and in the limit we get nothing new: the system is stuck near the Nash outcome.
- With a *constant* gain, agents discount past data, convergence is arrested, and **new outcomes emerge** as the system recurrently *escapes* the self-confirming equilibrium toward the Ramsey (zero-inflation) outcome, in spontaneous stabilizations that resemble the arrival of Volcker.

The word "asymptotically" matters.

In a finite sample a falling gain also slows the journey *back* to the self-confirming equilibrium, so a least squares government whose prior is not tight can escape early and then stay near Ramsey for a very long time (*Least squares learning: a decreasing gain*).

These escapes are the heart of the *vindication of econometric policy evaluation* story from *The Rise and Fall of U.S. Inflation*: an adaptive government, learning a Solow-Tobin distributed-lag version of the natural-rate hypothesis, is led by chance observations to stabilize inflation.

This lecture replicates, analyzes, and reinterprets simulations like those of Christopher Sims and Heetaik Chung [Chung, 1990, Sims, 1988].

Their governments ran **Kalman filters** on Phillips-curve coefficients that they believed drift as random walks (see *A First Look at the Kalman Filter*).

Sims's government fit a *static* Phillips curve, and Chung's fit a distributed lag.

In their simulations the government escaped from Nash once and stayed near Ramsey, apparently indefinitely.

Chung's filter had a gain that did not vanish, which let him estimate the government's learning rule from post-war U.S. data (*The econometricians did return the compliment*; Exercise **??** poses a stripped-down version of his problem).

Following Sargent [1999], we instead give the government **constant-gain recursive least squares**, which produces *recurrent* escapes: the mean dynamics keep pulling the beliefs back toward the self-confirming equilibrium, and the economy keeps escaping again.

Sargent and Williams [2005] trace the difference to the government's prior about coefficient drift, and *Priors, Escapes, and Learning Cycles* reproduces their account.

Because Sims's government fit a static curve, the induction-hypothesis mechanism studied below does not apply to his simulations (chapter 8, footnote 15 of Sargent [1999]).

Here we study the escapes by *simulation*; *Escaping Nash Inflation* then characterizes them analytically as a second, deterministic ODE, and *Priors, Escapes, and Learning Cycles* asks how the government's prior about coefficient drift reshapes both forces.

Let's import what we need:

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.integrate import solve_ivp
from scipy.linalg import solve_discrete_are, solve_discrete_lyapunov
```

103.2 A primer on recursive algorithms

This section is a self-contained primer.

It develops the two analytical objects — **mean dynamics** and **escape routes** — that organize everything that follows, and it explains the sense in which one and the same recursion can be read either as an *algorithm* for computing a self-confirming equilibrium or as a *model* of a government adapting in real time.

The material is technical, and readers who want the punchline can skip to the simulations below and refer back as needed.

Readers arriving from the Bounded Rationality lectures have seen the scalar version of the mean dynamics, together with the T-map and E-stability, in *One recursion, many learners*, and *Least Squares Learning in Self-Referential Models* proves the convergence theorems for linear self-referential models.

This primer holds the general large-deviation theory for the whole suite; *Escaping Nash Inflation* applies it to a static version of the model.

103.2.1 Beliefs, moment conditions, and self-confirming equilibria

A self-confirming equilibrium under the classical identification is pinned down by the government's beliefs about some population moments and the regression coefficients they imply.

Under the classical identification these beliefs are measured by the triple $(\gamma,\ EX_C X_C^\top,\ EUX_C)$, where γ is the vector of Phillips-curve coefficients.

In the adaptive models of this lecture the *time-t values* of these objects are among the economy's state variables; they disappear as state variables in a self-confirming equilibrium only because there they are constants.

A self-confirming equilibrium under the classical identification satisfies the moment conditions

$$\begin{aligned} \mathbb{E}\, R_{XC}^{-1}(\gamma)\left[U_t X_{Ct}^\top - (X_{Ct}X_{Ct}^\top)\,\gamma\right] &= 0, \\ \mathbb{E}\, X_{Ct}X_{Ct}^\top - R_{XC}(\gamma) &= 0, \end{aligned} \tag{103.1}$$

where the mathematical expectation is taken with respect to a distribution of (U_t, X_{Ct}) that depends on γ through the solution $h(\gamma)$ of the Phelps problem.

Self-reference surfaces precisely in this dependence of the distribution on γ: the government's beliefs shape its policy, which shapes the data on which the beliefs are then checked.

The first line of (103.1) is the least squares normal equations for γ, pre-multiplied by the inverse second-moment matrix; the second line defines R_{XC} as the second-moment matrix of the regressors.

It is convenient to assemble all the unknowns into a single vector

$$\phi = \begin{bmatrix} \gamma \\ \mathrm{col}(R_{XC}) \end{bmatrix}, \tag{103.2}$$

where $\mathrm{col}(R_{XC})$ stacks the columns of R_{XC}.

The moment conditions (103.1) then take the compact form

$$\mathbb{E}\left[F(\phi,\zeta)\right] = 0, \qquad b(\phi) \equiv \mathbb{E}\left[F(\phi,\zeta)\right], \tag{103.3}$$

where ζ is a random vector and the expectation is over its distribution (which, again, depends on ϕ).

A self-confirming equilibrium is a set of beliefs ϕ_f that is a zero of b,

$$b(\phi_f) = 0. \tag{103.4}$$

The rest of this section describes recursive algorithms for finding such a zero, and a change of perspective that converts each computational algorithm into a model of real-time adaptation.

103.2.2 Iteration

The simplest algorithm computes a sequence $\{\phi_k\}$ of estimates from

$$\phi_{k+1} = \phi_k + a\, b(\phi_k), \tag{103.5}$$

where the distribution used to evaluate the expectation defining $b(\phi_k)$ in (103.3) is itself evaluated at the current estimate ϕ_k, and $a > 0$ is a step size.

This is the relaxation algorithm used to compute self-confirming equilibria in *Self-Confirming Equilibria*.

Each step requires evaluating the mathematical expectation $b(\phi) = \mathbb{E}[F(\phi, \zeta)]$ — which is exactly why we needed the moment (Lyapunov) formulas there.

103.2.3 Stochastic approximation

A random version of (103.5) is obtained by replacing the mean $b(\phi_n)$ by a single random draw $F(\phi_n, \zeta_n)$ and letting the *step size* do the averaging:

$$\phi_{n+1} = \phi_n + a_n F(\phi_n, \zeta_n), \qquad a_n > 0, \quad \sum_{n=0}^{\infty} a_n = +\infty. \tag{103.6}$$

To study the limiting behavior of (103.6), define **artificial time**

$$t_n = \sum_{k=0}^{n} a_k, \tag{103.7}$$

form the sampled process $\phi(t_n) = \phi_n$, and interpolate (typically piecewise-linearly) to obtain a continuous-time process $\phi^o(t)$.

One then approximates $\phi^o(t)$ by a continuous-time process as $n \to \infty$ and uses it to characterize the tail of the original sequence.

Different rates of decrease of the gain sequence $\{a_n\}$ produce different approximating processes, because they change the mapping (103.7) from real time n to artificial time t_n.

Note

Recursive stochastic approximation originates with Robbins and Monro [1951], who devised (103.6) to find the root of a regression function observed with noise, and with Kiefer and Wolfowitz [1952], who adapted it to find the maximum of a regression function (the "K-W" algorithms referred to below).

The "ODE method" for analyzing such recursions — approximating the interpolated process by the solution of a differential equation — is due to Ljung [1977]; book-length treatments are Benveniste *et al.* [1990] and Kushner and Yin [2003].

Its use to study learning in self-referential macroeconomic models is developed by Marcet and Sargent [1989] and, comprehensively, by Evans and Honkapohja [2001].

103.2.4 Mean dynamics

The classic stochastic approximation algorithms of Kushner and Clark [1978] and Ljung [1977] set the gain to decrease like $a_n \sim 1/n$ (at least for $t \geq N$ for some $N > 0$).

This permits strong statements about the almost sure convergence of (103.6) to a zero of $b(\phi)$.

For $a_n \sim 1/n$, as $n \to \infty$ the interpolated process $\phi^o(t)$ approaches the solution of the ordinary differential equation

$$\frac{d\phi^o(t)}{dt} = b\left(\phi^o(t)\right), \tag{103.8}$$

which we call the **mean dynamics** — the vector generalization of the scalar least-squares-learning ODE of *Least squares learning converges to Nash*.

A law of large numbers makes the random term in the continuous-time approximation vanish fast enough that the mean dynamics (103.8) describe the *tail* behavior of the stochastic process (103.6).

Consequently:

- if the algorithm converges (almost surely), it converges to a zero of the mean dynamics, $b(\phi) = 0$ — a self-confirming equilibrium; and
- the ODE (103.8) carries the information about local and global stability of the algorithm.

Local stability is governed by the eigenvalues of the Jacobian of b at a rest point: if all eigenvalues have negative real parts the rest point is locally stable.

This is the condition called **E-stability** (*One recursion, many learners*).

For least squares the Jacobian has a simple structure, established by the Jacobian-reduction proposition (*Proposition 138.4.2*) of *Least Squares Learning in Self-Referential Models*:

- the block belonging to the second-moment matrix contributes repeated eigenvalues equal to -1; and
- the remaining eigenvalues are those of $DT - I$, where the map T sends a belief to the population regression coefficients that the belief induces.

Under a decreasing gain the eigenvalue of largest real part also governs the *rate* of convergence.

A necessary condition for the usual $\sqrt{t}$ rate, with asymptotically normal estimates, is that every eigenvalue have real part below $-\frac{1}{2}$.

Mean dynamics near the self-confirming equilibrium computes these eigenvalues for our model; one of them, repeated five times, sits exactly at the $-\frac{1}{2}$ boundary.

That boundary concerns the asymptotic distribution of a decreasing-gain estimator.

It does not govern how often a constant-gain government escapes; the large-deviation objects introduced below do.

Note

Brock and Hommes [1997] build models whose global behavior is driven by stable mean dynamics far from rational expectations equilibria together with local instability of adaptation near them — a complementary mechanism for generating endogenous fluctuations from learning.

103.2.5 Constant gain and convergence in distribution

We are equally interested in versions of (103.6) with a *constant* gain $a_n = \epsilon > 0$ for all n.

Limit theorems for constant-gain algorithms use a weaker notion of convergence — convergence *in distribution* — than the almost-sure convergence available when $a_n \sim 1/n$.

They concern small-noise limits, taken as $\epsilon \to 0$ and $n\epsilon \to +\infty$ simultaneously.

Again using artificial time (103.7), form the family of processes

$$\phi^{\epsilon}_{n+1} = \phi^{\epsilon}_n + \epsilon\, F(\phi^{\epsilon}_n, \zeta_n), \tag{103.9}$$

interpolate to obtain $\phi^{\epsilon}(t)$, and study its small-ϵ limit.

Dupuis and Kushner, Kushner and Yin [2003], and others verified conditions under which, as $\epsilon \to 0$ and $\epsilon n \to \infty$, the process ϕ^{ϵ}_n converges *in distribution* to the zeros of the same mean dynamics (103.8); the restrictions on the mean dynamics needed for convergence match those from the classic $a_n \sim 1/n$ theory.

Unlike the decreasing-gain algorithm, a constant-gain algorithm does not settle down: (γ, R_{XC}) converges to a *stationary stochastic process*, perpetually fluctuating around — and occasionally far from — the self-confirming equilibrium.

103.2.6 Escape routes and the theory of large deviations

The feature of the constant-gain apparatus that matters most for this lecture is not convergence toward ϕ_f but the *excursions away from it*.

We are as interested in movements away from a self-confirming equilibrium as in those toward one, because the recurrent stabilizations in the simulations are precisely such excursions.

The **theory of large deviations** characterizes these excursions through three objects.

First, the log moment generating function of (an averaged version of) the innovation process $F(\phi_n, \zeta_n)$: for a vector α conformable to F,

$$H(\alpha, \phi) = \log \mathbb{E} \exp\left(\alpha^{\top} F(\phi, \zeta)\right), \tag{103.10}$$

where the expectation is over the distribution of ζ.

Note

Equation (103.10) is a heuristic shorthand.

The object that actually enters the theory is a *time-averaged* limit; Dupuis and Kushner [1987] and Kushner and Yin [2003] assume that for each block width $\Delta > 0$ the following limit exists uniformly in ϕ_i, α_i on any compact set: $$\sum_{i=0}^{T/\Delta-1} \Delta\, H(\alpha_i, \phi_i) = \lim_{N\to\infty} \frac{\Delta}{N} \log \mathbb{E} \exp \sum_{i=0}^{T/\Delta-1} \alpha_i^{\top} \sum_{j=iN}^{iN+N-1} F(\phi_i, \zeta_j).$$ The inner sum averages the innovations over a block of length N; the double limit lets us treat serially dependent innovations.

Second, the **Legendre transform** of H, which plays the role of a rate function:

$$L(\xi, \phi) = \sup_{\alpha} \left[\alpha^{\top}\xi - H(\alpha, \phi)\right]. \tag{103.11}$$

Third, the **action functional**, which measures the "cost" of a candidate escape path $\phi(\cdot)$:

$$S(T, \phi) = \begin{cases} \displaystyle\int_0^T L\big(\tfrac{d}{ds}\phi(s),\, \phi(s)\big)\, ds & \text{if } \phi(s) \text{ is absolutely continuous and } \phi(0) = \phi_f, \\ \infty & \text{otherwise.} \end{cases} \tag{103.12}$$

Dupuis and Kushner turn the search for the most likely escape into a *deterministic control problem*.

Let D be a compact set containing ϕ_f, with boundary ∂D, and let $C[0,T]$ be the continuous functions on $[0,T]$.

The escape route is the path $\tilde{\phi}(\cdot)$ that solves

$$\inf_{T>0} \inf_{\phi \in A} S(T,\phi), \qquad A = \{\phi(\cdot) \in C[0,T] : \phi(T) \in \partial D\}. \tag{103.13}$$

Assuming the minimizer $\tilde{\phi}(\cdot)$ is unique, and letting t_D^ϵ be the first time the constant-gain process $\phi^\epsilon(t)$ leaves D, Dupuis and Kushner [1987] show that for every $\eta > 0$

$$\lim_{\epsilon \to 0} \mathbb{P}\left(\left|\phi^\epsilon(t_D^\epsilon) - \tilde{\phi}(T)\right| > \eta\right) = 0. \tag{103.14}$$

In words: *conditional on escaping the set D, the system leaves it near the terminal point of the least-action path* — so the escape has a deterministic direction and shape, even though it is triggered by chance.

This is the sense in which the escapes below, though they require no large shock, "seem purposeful": they follow the least-action route dictated by (103.13).

The same theory delivers the *mean time* to escape.

Let $\bar{S}$ be the minimized value of (103.13); then $\epsilon \log \mathbb{E}\, t_D^\epsilon \to \bar{S}$ as $\epsilon \to 0$.

So escapes become exponentially rare as the gain shrinks.

That is a small-gain statement.

When $\bar{S}$ is small relative to ϵ, as it is at the gains used in this suite, escapes are *frequent*: mean escape times grow roughly like $1/\epsilon$, not like $\exp(\bar{S}/\epsilon)$ (*Simulating escapes* measures them), and in the simulations below the government spends most of its time near the Ramsey outcome.

A crucial contrast with the mean dynamics: the mean dynamics (103.8) do *not* depend on the noise around them, whereas the escape routes *do* — the noise not only adds random fluctuations around (103.8), it also carves out this second family of paths.

Note

The mathematical foundations are the large-deviation theory for randomly perturbed dynamical systems of Freidlin and Wentzell [1998] (especially their chapter 4), specialized to stochastic approximation by Dupuis and Kushner [1987] and Dupuis and Kushner [1989]; a weak-convergence treatment is Dupuis and Ellis [1997].

103.2.7 A tractable action functional

The escape-route calculation promises cheap information about central tendencies of the algorithm, but the action functional (103.12) is generally hard to compute.

An important special case simplifies it dramatically.

Suppose the innovation is additive and Gaussian,

$$F(\phi, \zeta) = b(\phi) + \sigma(\phi)\,\zeta,$$

where ζ_n is stationary and Gaussian but not necessarily serially uncorrelated, and define $\Omega = \sum_j \mathbb{E}\,\zeta_t \zeta_{t-j}^\top$.

Then the action functional takes the quadratic form

$$S(T,\phi) = \frac{1}{2}\int_0^T \left(\tfrac{d}{ds}\phi - b(\phi)\right)^\top \left[\sigma(\phi)\,\Omega\,\sigma(\phi)^\top\right]^+ \left(\tfrac{d}{ds}\phi - b(\phi)\right) w(s)\,ds, \tag{103.15}$$

where $(\cdot)^+$ is the Moore-Penrose generalized inverse (used to handle possible stochastic singularity of $\sigma\Omega\sigma^\top$).

The weight $w(s)$ depends on the gain: $w(s) = \exp(s)$ when $p = 1$ in $a_n = a_0/n^p$, and $w(s) = 1$ when $p < 1$.

Read (103.15) as a cost that penalizes departures of the *realized drift* $\frac{d}{ds}\phi$ from the *mean drift* $b(\phi)$, weighting each direction by the inverse of the local noise covariance $\sigma\Omega\sigma^\top$.

The least-action escape therefore threads the beliefs through regions where the mean dynamics are weak and the noise is large — which, in our model, is the direction of the *induction hypothesis*.

Our model does not satisfy the assumptions behind (103.15) exactly.

Its innovation F contains products such as $X_{C,t}X_{C,t}^\top$ and $X_{C,t}U_t$, so it is quadratic in the shocks rather than additive and Gaussian.

Note

Cho *et al.* [2002], the published treatment of the model of this lecture, reduce the escape problem to minimizing a quadratic cost of the form (103.15), with a weighting matrix that is a fourth-moment matrix computed from Lyapunov equations.

The reduction draws on an early version of Williams's work, [Williams, 2001].

Williams [2019], the published version, notes that the early version misapplied a large-deviation theorem, a discrepancy that matters only for dynamic models like the one in this lecture, so the CWS calculations for the dynamic model rest on that shortcut.

They solve the control problem (103.13) *numerically* for the Nash self-confirming equilibrium of the dynamic model and find (their Figure 3) that along the dominant escape path the sum of weights on inflation rises from -1 to nearly zero, the value that activates the induction hypothesis — that is, the escape heads toward the Ramsey outcome.

Sargent and Williams [2005] replace recursive least squares with a Kalman filter whose drift covariance V is fixed, and study how V reshapes the escape (*Priors, Escapes, and Learning Cycles*).

The constant-gain algorithm of this lecture is a Kalman filter whose drift covariance is proportional to the current P_{t-1} (see below); a fixed V reproduces it only asymptotically.

Kasa [2004] applies the same large-deviation machinery to recurrent currency crises.

103.2.8 From computation to adaptation

The preceding recursions were introduced as *algorithms* to approximate a self-confirming equilibrium.

The same mathematics tells us what happens when we instead *modify* our self-confirming-equilibrium models to incorporate real-time adaptation — simply by reading ϕ_n as the government's time-n beliefs rather than as the n-th iterate of a solver.

Two facts organize everything below:

1. gain sequences that implement least squares (decreasing like $1/t$) make the mean dynamics pull the economy *toward* self-confirming equilibria, at least asymptotically; while
2. gain sequences that fall off more slowly — in the limit, constant gains that discount the past — *arrest* that pull and increase the frequency with which the escape dynamics influence outcomes.

Note

A brief intellectual history.

Lucas and Prescott [1971] dismissed iterating on the moment conditions (103.4) as a computational strategy, but Townsend [1983] used it.

DeCanio [1979] and Evans [1985], Evans [1989] used the map from perceived to actual laws of motion to define expectational stability.

Woodford [1990] and Marcet and Sargent [1989], Marcet and Sargent [1989] used the mean dynamics (103.8) to establish conditions for the convergence of least squares learning to rational expectations in models with self-reference, both requiring continuity of $b(\phi)$.

In-Koo Cho studied problems with *discontinuous* $b(\phi)$ inherited from discontinuous decision rules (trigger strategies in credibility and search problems); to make least squares learning approach rational expectations he used gains satisfying $\frac{1}{\sqrt{n}} < a_n < \frac{1}{\log n}$, which yield a *diffusion* approximation to (103.6) that promotes enough experimentation to discover an equilibrium (footnote 8 of chapter 8 of Sargent [1999] prints the two bounds the other way round, which describes an empty set because $\log n < \sqrt{n}$).

Kandori *et al.* [1993] use related mathematics to select long-run equilibria in games via mutation, and Roger Myerson applied an escape-route calculation to a voting problem.

The modern synthesis of these learning methods is Evans and Honkapohja [2001].

103.3 The adaptive model

We now build the classical adaptive model.

103.3.1 Government beliefs and behavior

The government believes in a distributed-lag Phillips curve

$$U_t = \gamma^\top X_{C,t} + \varepsilon_{C,t}, \qquad X_{C,t} = \begin{bmatrix} y_t & U_{t-1} & U_{t-2} & y_{t-1} & y_{t-2} & 1 \end{bmatrix}^\top. \tag{103.16}$$

Arriving at time t with an estimate γ_{t-1}, it sets the systematic part of inflation by solving the Phelps problem *as if* γ_{t-1} will govern the Phillips curve forever:

$$y_t = h(\gamma_{t-1}) X_{t-1} + v_{2t}, \qquad X_{t-1} = \begin{bmatrix} U_{t-1} & U_{t-2} & y_{t-1} & y_{t-2} & 1 \end{bmatrix}^\top. \tag{103.17}$$

It then updates its beliefs by **recursive least squares** (RLS):

$$\begin{aligned} \gamma_t &= \gamma_{t-1} + g_t R_{XC,t}^{-1} X_{C,t} \left(U_t - \gamma_{t-1}^\top X_{C,t} \right), \\ R_{XC,t} &= R_{XC,t-1} + g_t \left(X_{C,t} X_{C,t}^\top - R_{XC,t-1} \right), \end{aligned} \tag{103.18}$$

where $\{g_t\}$ is the gain sequence.

Least squares sets $g_t = 1/t$; a constant-gain algorithm sets $g_t = g_0 > 0$ and discounts past observations, which is sensible if the government suspects the Phillips curve wanders over time.

Readers arriving from *A government learning the Phillips curve* should note a change of letters: there, and throughout the Bounded Rationality lectures (*Notation across the series*), γ is the gain and α the beliefs, whereas here, following the book, γ denotes the beliefs and g_t, or $1-\lambda$ below, the gain (see also *A note on notation*).

The public is assumed to know the government's rule, so its inflation forecast is $x_t = h(\gamma_{t-1}) X_{t-1}$, the systematic part of (103.17).

Unemployment is generated by the actual Phillips curve of *Self-Confirming Equilibria* with $\rho_1 = \rho_2 = 0$:

$$U_t = U^* - \theta(y_t - x_t) + v_{1t} = U^* - \theta v_{2t} + v_{1t}.$$

103.3.2 The Phelps problem with lags

Given a belief γ, the decision rule $h(\gamma)$ solves an LQ control problem.

Write the believed Phillips curve as $U_t = \gamma_0 y_t + c^\top s_t$, where γ_0 is the coefficient on current inflation and c collects the coefficients on the state $s_t = X_{t-1}$.

The government minimizes $\mathbb{E}\sum_t \delta^t (U_t^2 + y_t^2)$, so the per-period loss is $s_t^\top (cc^\top) s_t + (\gamma_0^2 + 1) y_t^2 + 2\gamma_0\, y_t\, c^\top s_t$, and the state evolves as $s_{t+1} = A s_t + B y_t$ with

$$s_{t+1} = \begin{bmatrix} U_t \\ U_{t-1} \\ y_t \\ y_{t-1} \\ 1 \end{bmatrix}, \qquad U_t = c^\top s_t + \gamma_0 y_t .$$

This is the Phelps problem of *Adaptive Expectations and the Phelps Problem*, now with a distributed-lag Phillips curve.

We solve the discounted LQ problem with `scipy`'s discrete algebraic Riccati equation.

To fold the discount factor into an undiscounted problem we scale A and B by $\sqrt{\delta}$.

The cross-product term $2\gamma_0\, y_t\, c^\top s_t$ in the loss (`N` in the code) is a per-period cost, like the other two terms, so it is not scaled.

```
class AdaptivePhillips:
    """
    Classical adaptive Phillips curve model: the government re-estimates a
    distributed-lag Phillips curve by recursive least squares and each period
    acts on the first-period recommendation of the Phelps problem.
    """

    def __init__(self, θ=1.0, U_star=5.0, σ1=0.3, σ2=0.3, δ=0.98):
        self.θ, self.U_star, self.σ1, self.σ2, self.δ = θ, U_star, σ1, σ2, δ

        # classical self-confirming belief: U = -θ y + (θ²+1)U*
        self.γ_sce = np.array([-θ, 0.0, 0.0, 0.0, 0.0, (θ**2 + 1) * U_star])

        # self-confirming moment matrix M = E[X_C X_C'] and residual variance
        self.M = self._sce_moments()
        self.σC2 = σ1**2                        # var(U | X_C) at the SCE

    def _sce_moments(self):
        "E[X_C X_C'] at the serially-uncorrelated classical SCE."
        θ, σ1, σ2 = self.θ, self.σ1, self.σ2
        μU, μy = self.U_star, self.θ * self.U_star
        Σ = {('U', 'U'): θ**2 * σ2**2 + σ1**2, ('y', 'y'): σ2**2,
             ('U', 'y'): -θ * σ2**2, ('y', 'U'): -θ * σ2**2}
        # regressors: (time, type) for [y_t, U_{t-1}, U_{t-2}, y_{t-1}, y_{t-2}, 1]
        regs = [(0, 'y'), (1, 'U'), (2, 'U'), (1, 'y'), (2, 'y'), (None, 'c')]
        mean = {'U': μU, 'y': μy, 'c': 1.0}
        M = np.zeros((6, 6))
        for i, (ti, tyi) in enumerate(regs):
            for j, (tj, tyj) in enumerate(regs):
                if tyi == 'c' or tyj == 'c' or ti != tj:
                    M[i, j] = mean[tyi] * mean[tyj]
                else:
                    M[i, j] = Σ[(tyi, tyj)] + mean[tyi] * mean[tyj]
        return M
```

(continues on next page)

(continued from previous page)

```
    def phelps_h(self, γ):
        "Government decision rule ŷ_t = h(γ)·X_{t-1} for belief γ."
        δ, γ0, c = self.δ, γ[0], γ[1:]
        R = np.outer(c, c)
        Q = np.array([[γ0**2 + 1.0]])
        N = (γ0 * c).reshape(1, -1)
        A = np.zeros((5, 5)); B = np.zeros((5, 1))
        A[0, :] = c; B[0, 0] = γ0          # U_t
        A[1, 0] = 1.0                       # U_{t-1}
        B[2, 0] = 1.0                       # y_t
        A[3, 2] = 1.0                       # y_{t-1}
        A[4, 4] = 1.0                       # constant
        sb = np.sqrt(δ)
        Ad, Bd = sb * A, sb * B             # discounting folded into A, B
        P = solve_discrete_are(Ad, Bd, R, Q, s=N.T)   # cross term unscaled
        F = np.linalg.solve(Q + Bd.T @ P @ Bd, Bd.T @ P @ Ad + N)
        return -F.ravel()
```

We use the Kalman-filter implementation of RLS from Appendix A of Sargent [1999].

The government treats γ as a hidden state that follows a random walk, $\tilde{\gamma}_t = \tilde{\gamma}_{t-1} + w_t$, and observes it through the regression $U_t = \tilde{\gamma}_t^\top X_{C,t} + \varepsilon_{C,t}$; the Kalman filter of *A First Look at the Kalman Filter* then updates its estimate γ_t and the estimate's covariance matrix P_t.

A forgetting factor $\lambda \in (0, 1]$ maps to the gain:

- $\lambda = 1$ sets the drift covariance $\mathbb{E} w_t w_t^\top$ to zero and gives least squares ($g_t \to 1/t$);
- $\lambda < 1$ sets $\mathbb{E} w_t w_t^\top = (1/\lambda - 1) P_{t-1}$, a drift covariance proportional to the government's current uncertainty, and gives a constant gain $g_0 = 1 - \lambda$.

The prior is initialized as if the government had already seen T_0 periods of self-confirming-equilibrium data, through $P_0 = (\sigma_C^2 / T_0)\, M^{-1}$; larger T_0 (`T_prior` in the code) means a tighter prior.

The function also returns the final beliefs and covariance matrix, which we use below.

```
def simulate(model, λ, T_prior, n=1000, seed=0):
    "Simulate the adaptive system. λ=1 is least squares; λ<1 is constant gain."
    rng = np.random.default_rng(seed)
    θ, U_star, σ1, σ2 = model.θ, model.U_star, model.σ1, model.σ2

    γ = model.γ_sce.copy()
    P = (model.σC2 / T_prior) * np.linalg.inv(model.M)
    R2 = model.σC2
    g0 = 1 - λ

    U1 = U2 = U_star
    y1 = y2 = θ * U_star
    y_path, U_path, sumweights, constant = (np.empty(n) for _ in range(4))

    for t in range(n):
        h = model.phelps_h(γ)
        X_lag = np.array([U1, U2, y1, y2, 1.0])
        yhat = h @ X_lag

        v2, v1 = σ2 * rng.standard_normal(), σ1 * rng.standard_normal()
        y = yhat + v2
```

(continues on next page)

(continued from previous page)

```
        U = U_star - θ * (y - yhat) + v1

        ψ = np.array([y, U1, U2, y1, y2, 1.0])            # X_C,t
        denom = R2 + ψ @ P @ ψ
        gain = P @ ψ / denom
        γ = γ + gain * (U - γ @ ψ)
        R1 = (g0 / (1 - g0)) * P if λ < 1 else 0.0         # constant vs decreasing
        P = P - np.outer(P @ ψ, ψ @ P) / denom + R1

        y_path[t], U_path[t] = y, U
        sumweights[t] = γ[0] + γ[3] + γ[4]                # weights on current+lagged y
        constant[t] = γ[5]
        U2, U1 = U1, U
        y2, y1 = y1, y

    return dict(y=y_path, U=U_path, sumweights=sumweights, constant=constant,
                γ=γ, P=P)
```

Let's check the decision rule at the self-confirming belief.

```
model = AdaptivePhillips()
h_sce = model.phelps_h(model.γ_sce)
print("decision rule at the self-confirming belief:")
print(f"  h(γ_sce) = {np.round(h_sce, 3) + 0.0}  (a constant rule of "
      f"{h_sce[-1]:.1f} = Nash inflation)")
```

```
decision rule at the self-confirming belief:
  h(γ_sce) = [0. 0. 0. 0. 5.]  (a constant rule of 5.0 = Nash inflation)
```

At the self-confirming belief the Phelps rule is a constant equal to the Nash inflation rate, and the beliefs reproduce themselves — the adaptive system's rest point is the self-confirming equilibrium of *Self-Confirming Equilibria*.

We follow Sargent [1999] in setting the true data-generating parameters to the classical example at the end of *Self-Confirming Equilibria*: $U^* = 5$, $\theta = 1$, $\sigma_1 = \sigma_2 = 0.3$, $\rho_1 = \rho_2 = 0$, $\delta = 0.98$.

The classical self-confirming equilibrium has serially uncorrelated (U, y) fluctuating around means $(5, 5)$.

103.4 Mean dynamics near the self-confirming equilibrium

Before simulating, we compute the mean dynamics (103.8) for this model and their behavior near the self-confirming equilibrium.

Substituting the government's rule (103.17) and the public's forecast into the true Phillips curve gives

$$U_t = -\theta y_t + \theta\, h(\gamma) X_{t-1} + U^* + v_{1t}.$$

The shock v_{1t} is independent of every regressor in $X_{C,t}$.

So when the government uses the rule $h(\gamma)$, the population regression of U_t on $X_{C,t}$ is exactly

$$T(\gamma) = \begin{bmatrix} -\theta & \theta h_1(\gamma) & \theta h_2(\gamma) & \theta h_3(\gamma) & \theta h_4(\gamma) & U^* + \theta h_5(\gamma) \end{bmatrix}^\top, \tag{103.19}$$

where $h_j(\gamma)$ is the j-th entry of $h(\gamma)$.

This is the T-map of *One recursion, many learners* for our model: it sends the government's perceived Phillips curve to the actual one that its policy generates.

Applied to the recursive least squares algorithm (103.18), stochastic approximation delivers the mean dynamics (equation (89) of the book)

$$\frac{d\gamma}{d\tau} = R_{XC}^{-1} M(\gamma)\,[T(\gamma) - \gamma]\,, \qquad \frac{dR_{XC}}{d\tau} = M(\gamma) - R_{XC}, \tag{103.20}$$

where τ is artificial time and $M(\gamma) = \mathbb{E} X_{C,t} X_{C,t}^\top$ is computed from the stationary distribution of the system when the government's beliefs are frozen at γ.

A rest point has $T(\gamma) = \gamma$ and $R_{XC} = M(\gamma)$: it is a self-confirming equilibrium.

At a rest point the Jacobian of (103.20) is block triangular.

Its eigenvalues are therefore those of $DT - I$, together with repeated eigenvalues of -1 from the R_{XC} block (*Proposition 138.4.2* of *Least Squares Learning in Self-Referential Models*).

We compute DT at the self-confirming belief by finite differences.

```
def T_map(model, γ):
    "Population regression of U_t on X_C,t when the government uses h(γ)."
    θ, h = model.θ, model.phelps_h(γ)
    return np.concatenate([[-θ], θ * h[:4], [model.U_star + θ * h[4]]])


def jacobian_T(model, ε=1e-5):
    "Finite-difference Jacobian of T at the self-confirming belief."
    J = np.empty((6, 6))
    for j in range(6):
        d = np.zeros(6)
        d[j] = ε
        J[:, j] = (T_map(model, model.γ_sce + d)
                   - T_map(model, model.γ_sce - d)) / (2 * ε)
    return J


DT_I = jacobian_T(model) - np.eye(6)
print("DT - I at the self-confirming belief:")
print(np.round(DT_I, 3) + 0.0)
print("eigenvalues:", np.round(np.sort(np.linalg.eigvals(DT_I).real), 4))
print(f"-1/(1 + θ²) = {-1 / (1 + model.θ**2)}")
```

```
DT - I at the self-confirming belief:
[[-1.     0.     0.     0.     0.     0.   ]
 [ 0.    -0.5    0.     0.     0.     0.   ]
 [ 0.     0.    -0.5    0.     0.     0.   ]
 [ 0.     0.     0.    -0.5    0.     0.   ]
 [ 0.     0.     0.     0.    -0.5    0.   ]
 [ 0.     2.45   2.401 -2.45  -2.401 -0.5  ]]
eigenvalues: [-1.  -0.5 -0.5 -0.5 -0.5 -0.5]
-1/(1 + θ²) = -0.5
```

The matrix $DT - I$ is lower triangular, so its eigenvalues are its diagonal entries.

- The first, -1, belongs to the coefficient on current inflation: T assigns that coefficient $-\theta$ whatever the government believes.
- The other five all equal $-1/(1+\theta^2) = -\frac{1}{2}$.

The second group has a simple source.

At the self-confirming equilibrium the lags carry no weight, and to first order the Phelps rule responds to a small believed coefficient c_j on a lagged variable as the static first-order condition $y_t = \theta\, c^\top X_{t-1}/(1+\theta^2)$ prescribes.

That makes the diagonal entry of DT equal to $\theta^2/(1+\theta^2)$.

The off-diagonal entries all sit in the last row: a believed coefficient on a lag also shifts the constant in the rule, through the dynamic part of the Phelps problem.

The repeated eigenvalue $-1/(1+\theta^2)$ is the same number as the slope $\mathcal{M}$ of the scalar mean dynamics in *Least squares learning converges to Nash*, where the public learned a single number, and as the classical E-stability eigenvalue of *Uniqueness and E-stability*.

Sargent [1999] reports a repeated eigenvalue of $-.5$ "of multiplicity 6", together with repeated eigenvalues of -1 for R_{XC}.

The computation shows that the multiplicity is five; the sixth eigenvalue for the beliefs, the one for the coefficient on current inflation, is -1.

All eigenvalues are negative, so the self-confirming equilibrium is E-stable.

At $\theta = 1$ the repeated eigenvalue sits exactly on the $-\frac{1}{2}$ boundary, so least squares converges more slowly than the usual $\sqrt{t}$ rate, as in the scalar case of *Least squares learning converges to Nash*.

As noted in the primer, that boundary concerns the asymptotic distribution of the least squares estimates; how often a constant-gain government escapes is governed instead by the minimized action $\bar{S}$.

Exercise 103.12.3 asks how the eigenvalues change with θ.

103.5 Least squares learning: a decreasing gain

First, least squares (a decreasing gain), with a tight prior worth $T_0 = 5000$ periods of data.

```
ls = simulate(model, λ=1.0, T_prior=5000, n=1000, seed=1)

fig, ax = plt.subplots(figsize=(9, 4.5))
ax.plot(ls['y'], lw=0.8)
ax.axhline(5, color='k', ls='--', lw=1, label='self-confirming (Nash)')
ax.axhline(0, color='C2', ls=':', lw=1, label='Ramsey')
ax.set_xlabel('$t$')
ax.set_ylabel('inflation $y_t$')
ax.legend()
plt.show()
```

In Fig. 103.1 inflation hugs the self-confirming value of 5, and the simulation looks like a draw from the self-confirming equilibrium itself.

This is the outcome the book describes: least squares gets us nothing new, and the government is stuck near Nash.

But the result depends on the prior.

The book's Figure 8.1 uses a looser prior, $T_0 = 800$, and shows 400 periods in which inflation stays near 5.

Let's compare eight paths of 1000 periods under each prior.

To see through the noise we plot a 20-period moving average of inflation, and we call the first date at which that average falls below 1 the government's *first escape*.

```
def moving_average(y, k=20):
    return np.convolve(y, np.ones(k) / k, mode='valid')
```

(continues on next page)

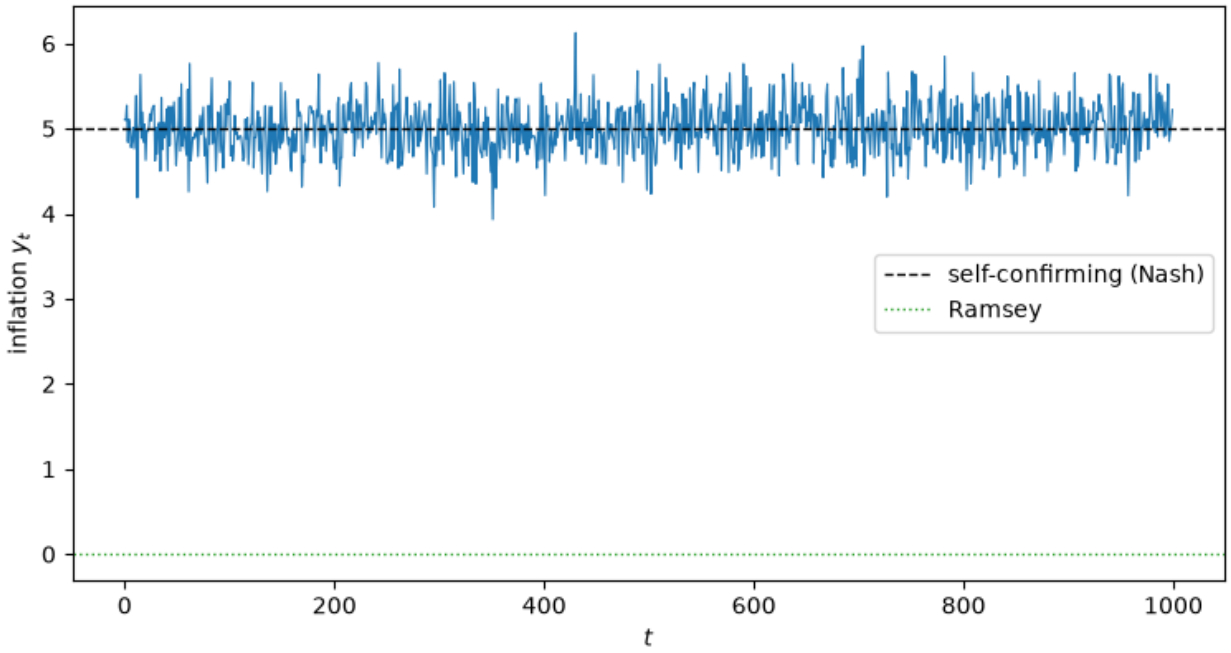


Fig. 103.1: Classical adaptive model under least squares with a tight prior: inflation hugs the self-confirming value

(continued from previous page)

```
def first_escape(y, k=20, level=1.0):
    "First date at which the k-period moving average of y falls below level."
    below = moving_average(y, k) < level
    return int(np.argmax(below)) if below.any() else None


def returns_to_nash(y, k=20, level=4.0):
    "Whether the moving average climbs back above level after the first escape."
    t = first_escape(y, k)
    return t is not None and bool((moving_average(y, k)[t:] > level).any())


ls_runs = {T0: [simulate(model, λ=1.0, T_prior=T0, n=1000, seed=s)
                for s in range(8)]
           for T0 in (5000, 800)}

fig, axes = plt.subplots(1, 2, figsize=(11, 4), sharey=True)
for ax, (T0, runs) in zip(axes, ls_runs.items()):
    for sim in runs:
        ax.plot(np.arange(19, 1000), moving_average(sim['y']), lw=0.8)
    ax.axhline(5, color='k', ls='--', lw=1)
    ax.axhline(0, color='C2', ls=':', lw=1)
    ax.set_title(rf'$T_0 = {T0}$')
    ax.set_xlabel('$t$')
axes[0].set_ylabel('20-period moving average of $y_t$')
plt.tight_layout()
plt.show()

for T0, runs in ls_runs.items():
```

(continues on next page)

(continued from previous page)

```
    escapes = [first_escape(sim['y']) for sim in runs]
    print(f"T_0 = {T0}: {sum(e is not None for e in escapes)} of {len(runs)} "
          f"paths escape; {sum(returns_to_nash(sim['y']) for sim in runs)} "
          f"of them return above 4")
```

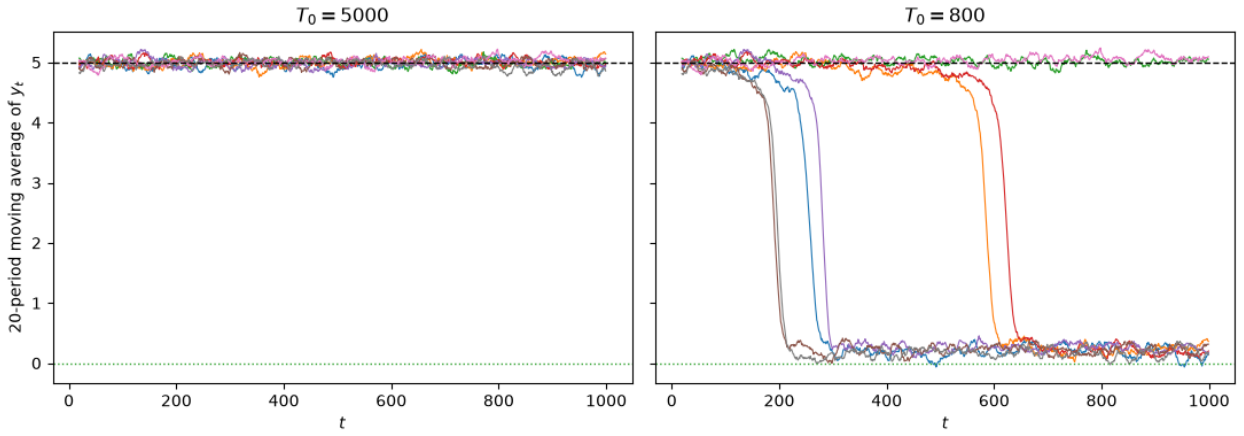


Fig. 103.2: Least squares under a tight and a looser prior: escapes under the looser prior stick

```
T_0 = 5000: 0 of 8 paths escape; 0 of them return above 4
T_0 = 800: 6 of 8 paths escape; 0 of them return above 4
```

Under the looser prior, Fig. 103.2 shows least squares escaping toward Ramsey, and every path that escapes stays there.

Why, if least squares converges to the self-confirming equilibrium?

Because the convergence theorems are about the limit.

With a prior worth T_0 periods, the least squares gain is $1/(T_0 + t)$, so after t periods artificial time (103.7) has advanced by about $\tau = \log\left((T_0 + t)/T_0\right)$.

The clock slows as it runs: each additional unit of τ costs $e \approx 2.7$ times as many periods as the one before.

A constant gain g_0, by contrast, advances τ by g_0 every period.

That explains both panels.

A tight prior slows the clock from the start: 1000 periods supply only $\log(6000/5000) \approx 0.18$ units of artificial time when $T_0 = 5000$, against $\log(1800/800) \approx 0.81$ when $T_0 = 800$.

And once the government has escaped, the slowing clock also delays the return.

To see how much artificial time the return needs, we integrate the mean dynamics (103.20) from the beliefs $(\gamma_t, R_{XC,t})$ that one escaped path holds at $t = 1000$.

We obtain $M(\gamma)$ from the stationary distribution of (y_t, U_t) under the frozen rule $h(\gamma)$, using a discrete Lyapunov equation.

Under least squares, $R_{XC,t} = \sigma_C^2 P_t^{-1}/(T_0 + t)$.

```
def stationary_moments(model, γ):
    """
    Mean inflation and M(γ) = E[X_C X_C'] in the stationary distribution
    of the true system when the government's beliefs are frozen at γ.
    """
```

(continues on next page)

(continued from previous page)

```
    θ, U_star, σ1, σ2 = model.θ, model.U_star, model.σ1, model.σ2
    h = model.phelps_h(γ)
    # z_t = (y_t, y_{t-1}, y_{t-2}, U_t, U_{t-1}, U_{t-2}) obeys
    # z_t = c + A z_{t-1} + C (v_1t/σ1, v_2t/σ2)
    A = np.zeros((6, 6))
    A[0, [0, 1, 3, 4]] = h[2], h[3], h[0], h[1]   # the rule h(γ)
    A[1, 0] = A[2, 1] = A[4, 3] = A[5, 4] = 1.0
    c = np.array([h[4], 0, 0, U_star, 0, 0])
    C = np.zeros((6, 2))
    C[0] = 0, σ2                                  # y_t = h X_{t-1} + v_2t
    C[3] = σ1, -θ * σ2                            # U_t = U* - θ v_2t + v_1t
    μ = np.linalg.solve(np.eye(6) - A, c)
    Σ = solve_discrete_lyapunov(A, C @ C.T)
    # X_C = (y_t, U_{t-1}, U_{t-2}, y_{t-1}, y_{t-2}, 1)
    idx = [0, 4, 5, 1, 2]
    M = np.ones((6, 6))
    M[:5, :5] = Σ[np.ix_(idx, idx)] + np.outer(μ[idx], μ[idx])
    M[:5, 5] = M[5, :5] = μ[idx]
    return μ[0], M


def mean_dynamics(τ, x, model):
    "Right side of the mean dynamics for the stacked vector (γ, col R_XC)."
    γ, R = x[:6], x[6:].reshape(6, 6)
    _, M = stationary_moments(model, γ)
    dγ = np.linalg.solve(R, M @ (T_map(model, γ) - γ))
    return np.concatenate([dγ, (M - R).ravel()])


def halfway(τ, x, model):
    "Event: mean inflation climbs back to half the Nash rate."
    return stationary_moments(model, x[:6])[0] - 0.5 * model.θ * model.U_star


halfway.direction = 1

T0, n = 800, 1000
sim = next(s for s in ls_runs[T0] if first_escape(s['y']) is not None)
R_n = model.σC2 * np.linalg.inv(sim['P']) / (T0 + n)
x0 = np.concatenate([sim['γ'], R_n.ravel()])

τ_grid = [0, 1, 2, 5, 10, 20, 30]
sol = solve_ivp(mean_dynamics, (0, τ_grid[-1]), x0, t_eval=τ_grid,
                events=halfway, args=(model,), rtol=1e-6, atol=1e-8)

print(f"{'τ':>4} {'sum of weights':>15} {'mean inflation':>15} "
      f"{'periods, least squares':>23} {'periods, gain 0.025':>20}")
for τ, x in zip(sol.t, sol.y.T):
    γ = x[:6]
    print(f"{τ:4.0f} {γ[0] + γ[3] + γ[4]:15.2f} "
          f"{stationary_moments(model, γ)[0]:15.2f} "
          f"{(T0 + n) * (np.exp(τ) - 1):23.1e} {τ / 0.025:20,.0f}")

τ_half = sol.t_events[0][0]
print(f"\nmean inflation is back to {0.5 * model.θ * model.U_star} at τ = "
      f"{τ_half:.1f}: after {(T0 + n) * (np.exp(τ_half) - 1):.1e} more periods "
      f"of least squares, or {τ_half / 0.025:.0f} periods at a gain of 0.025")
```

```
   τ  sum of weights  mean inflation  periods, least squares  periods, gain 0.025
   0           0.00            0.21                 0.0e+00                    0
   1          -0.00            0.25                 3.1e+03                   40
   2          -0.01            0.30                 1.2e+04                   80
   5          -0.11            1.07                 2.7e+05                  200
  10          -0.27            2.10                 4.0e+07                  400
  20          -0.52            3.35                 8.7e+11                  800
  30          -0.72            4.10                 1.9e+16                1,200

mean inflation is back to 2.5 at τ = 12.6: after 5.6e+08 more periods of least␣
↪squares, or 506 periods at a gain of 0.025
```

The last two columns count periods after $t = 1000$.

The mean dynamics do carry the beliefs back toward the self-confirming equilibrium: the sum of weights on inflation falls back toward -1, and mean inflation climbs back toward 5.

So least squares converges, but only asymptotically.

Near the induction hypothesis the mean dynamics are weak, and the return needs many units of artificial time.

A constant gain of 0.025 supplies them within several hundred periods.

Least squares supplies them only over a span of time that dwarfs any historical sample, because a falling gain slows the return as well as the departure.

In a finite sample, then, a decreasing gain can turn an early escape into an escape that sticks.

Paths of this kind, one escape followed by a long stay near Ramsey, resemble what Sims and Chung found with their Kalman filters, although the mechanism here differs from theirs; *Priors, Escapes, and Learning Cycles* (*Sims's nonconvergence*) explains their outcome through the prior in their filters.

Exercise 103.12.2 measures how often least squares escapes under three priors.

103.6 Constant gain and escape dynamics

Now give the government a *constant* gain $g_0 = 1 - \lambda = 0.025$, so it discounts past data.

We use the book's Figure 8.2 settings, $\lambda = 0.975$ and $T_0 = 300$; other lectures in this suite use other gains, each taken from its own source.

```
cg = simulate(model, λ=0.975, T_prior=300, n=1000, seed=1)

fig, ax = plt.subplots(figsize=(9, 4.5))
ax.plot(cg['y'], lw=0.8)
ax.axhline(5, color='k', ls='--', lw=1, label='self-confirming (Nash)')
ax.axhline(0, color='C2', ls=':', lw=1, label='Ramsey')
ax.set_xlabel('$t$')
ax.set_ylabel('inflation $y_t$')
ax.legend()
plt.show()
```

In Fig. 103.3 inflation starts near the self-confirming value of 5 and within a few dozen periods drops almost to zero.

Over the next several hundred periods it drifts back up, but only part of the way toward 5, before a second stabilization sends it back toward zero.

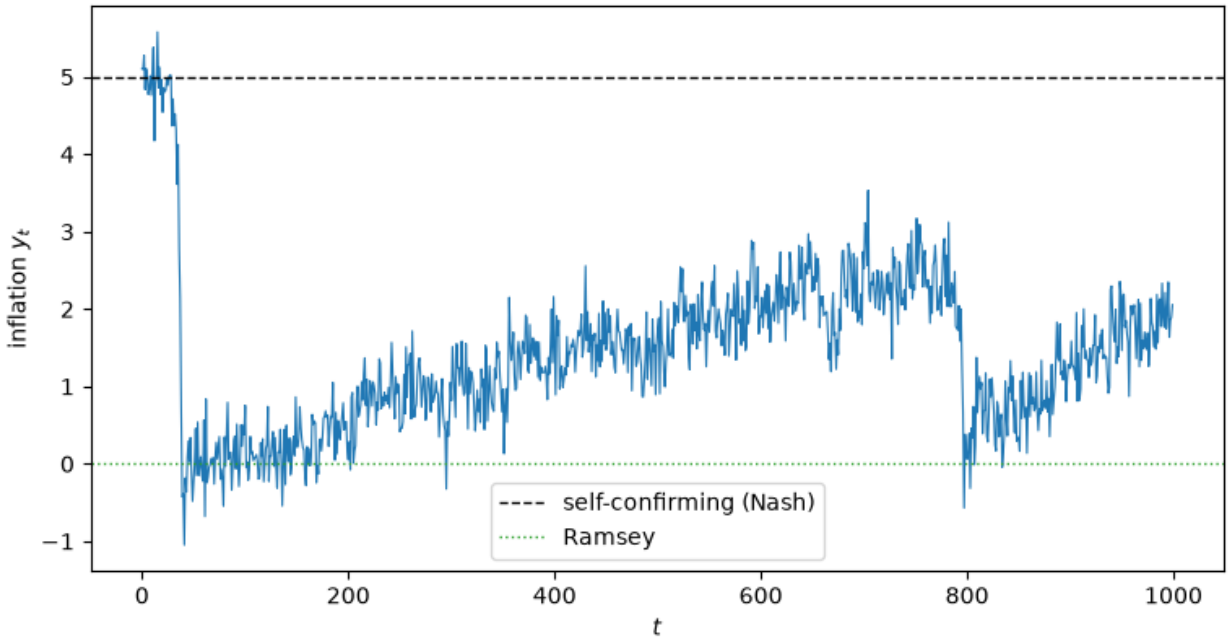


Fig. 103.3: Classical adaptive model under a constant gain: recurrent escapes toward Ramsey

The mean dynamics that pull the system toward the self-confirming equilibrium are opposed by a recurrent force that sends inflation close to the Ramsey outcome.

The drift back up is the slow return computed in *Least squares learning: a decreasing gain*, now run on a clock that does not slow down: the mean dynamics pull the beliefs back, and the noise launches a new escape before they get there.

Sargent [1999] checked that no large shocks trigger these stabilizations: they are an *endogenous* feature of the constant-gain learning dynamics.

One path can mislead, so we repeat the simulation for 24 seeds.

For each path we record the date of the first escape, the fraction of periods with inflation below 2, and whether the moving average of inflation ever climbs back above 4 after the first escape.

We also record the lowest value that the sum of weights on current and lagged inflation reaches after the first escape; we return to that sum below.

```
cg_runs = [simulate(model, λ=0.975, T_prior=300, n=1000, seed=s)
           for s in range(24)]

firsts = [first_escape(sim['y']) for sim in cg_runs]
frac_low = [(sim['y'] < 2).mean() for sim in cg_runs]
low_sw = [sim['sumweights'][t + 20:].min()
          for sim, t in zip(cg_runs, firsts) if t is not None]

print(f"paths that escape: {sum(t is not None for t in firsts)} of {len(cg_runs)}")
escaped = [t for t in firsts if t is not None]
print(f"first escape: median t = {np.median(escaped):.0f}, latest t = {max(escaped)}")
print(f"fraction of periods with y < 2: {min(frac_low):.0%} to "
      f"{max(frac_low):.0%} across seeds (mean {np.mean(frac_low):.0%})")
print(f"paths that return above 4 after escaping: "
      f"{sum(returns_to_nash(sim['y']) for sim in cg_runs)}")
print(f"lowest sum of weights after the first escape: {min(low_sw):.2f}")
```

```
paths that escape: 24 of 24
first escape: median t = 56, latest t = 145
fraction of periods with y < 2: 74% to 95% across seeds (mean 83%)
paths that return above 4 after escaping: 0
lowest sum of weights after the first escape: -0.71
```

Every path escapes early, inflation is below 2 in most periods, and no path's moving average climbs back above 4 within 1000 periods.

At this gain escapes are *frequent*, not rare.

The large-deviation statement that escapes become exponentially rare is a small-gain limit; *Simulating escapes* measures how escape times grow as the gain shrinks.

103.7 The escape route and the induction hypothesis

Why does the system escape toward Ramsey rather than in some other direction?

The answer is the **induction hypothesis** of *Adaptive Expectations and the Phelps Problem*: when the sum of the weights on current and lagged inflation in the estimated Phillips curve approaches zero, the Phelps problem advises the government to *reduce* inflation.

Let's plot inflation together with that sum of weights.

```
fig, axes = plt.subplots(2, 1, figsize=(9, 7), sharex=True)

axes[0].plot(cg['y'], lw=0.8)
axes[0].axhline(0, color='C2', ls=':', lw=1)
axes[0].set_ylabel('inflation $y_t$')

axes[1].plot(cg['sumweights'], lw=0.8, color='C1')
axes[1].axhline(-1, color='k', ls='--', lw=1, label='self-confirming value')
axes[1].axhline(0, color='C3', ls=':', lw=1, label='induction hypothesis')
axes[1].set_xlabel('$t$')
axes[1].set_ylabel('sum of weights on $y$')
axes[1].legend()

plt.tight_layout()
plt.show()
```

In Fig. 103.4 the first stabilization coincides with the sum of weights jumping from its self-confirming value of -1 to about zero.

The sum of weights never goes back to -1, here or in any of the 24 runs above.

As inflation drifts back up, the sum of weights falls only part of the way toward -1, and the next stabilization coincides with its rising back toward zero.

When it reaches zero, the induction hypothesis is (temporarily) satisfied, the Phelps problem calls for near-Ramsey inflation, and the resulting data briefly *reinforce* the induction hypothesis — a situation that is not self-confirming in the technical sense but is nonetheless self-reinforcing.

We can see the escape route directly by plotting the joint path of the constant and the sum of weights in the estimated Phillips curve.

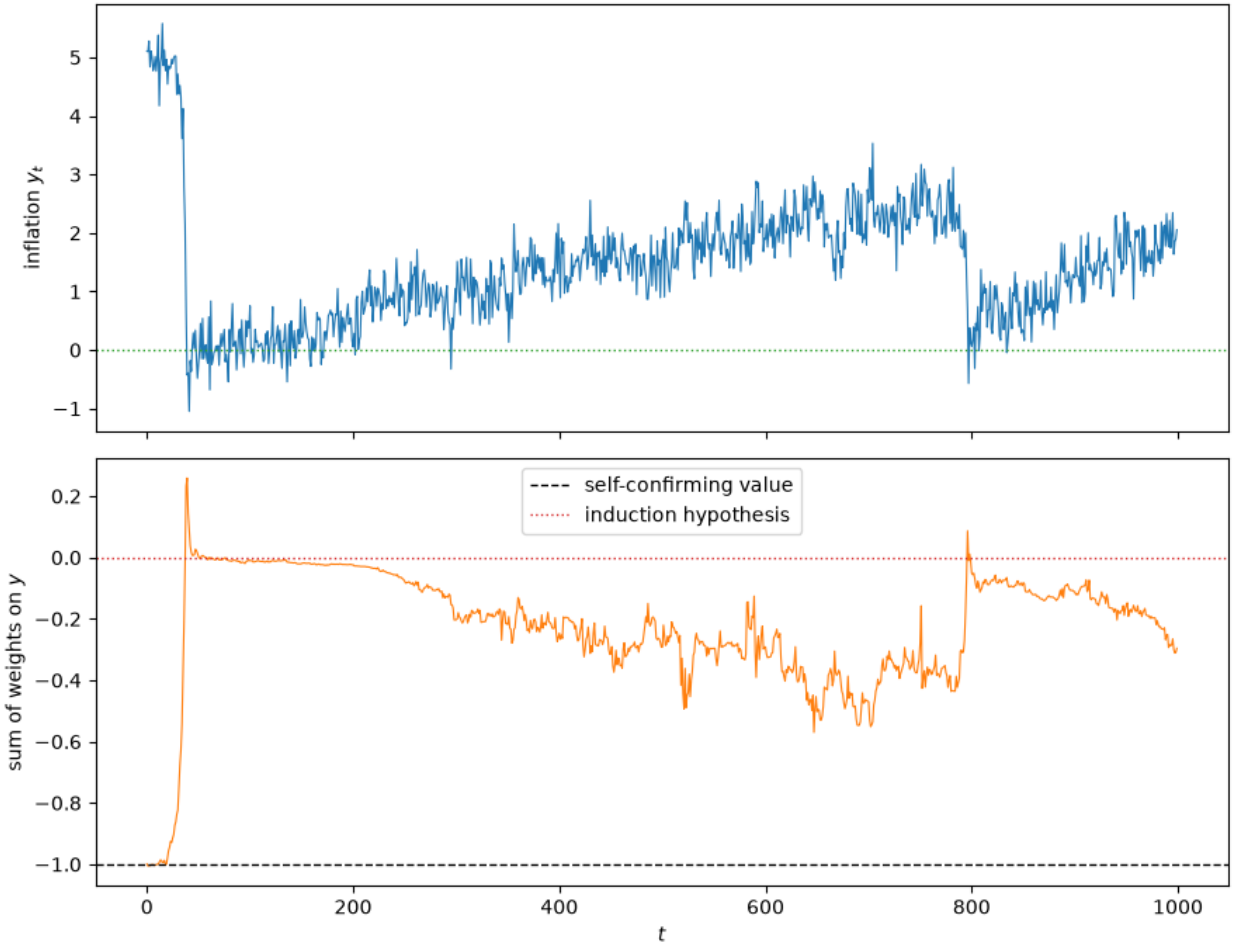


Fig. 103.4: Stabilizations coincide with the sum of weights on inflation rising toward zero

```
fig, ax = plt.subplots(figsize=(8, 6))
sc = ax.scatter(cg['constant'], cg['sumweights'], c=np.arange(len(cg['y'])),
                cmap='viridis', s=6)
ax.axhline(0, color='C3', ls=':', lw=1.5, label='induction hypothesis')
ax.axhline(-1, color='k', ls='--', lw=1, label='self-confirming value')
ax.set_xlabel('constant in estimated Phillips curve')
ax.set_ylabel('sum of weights on $y$')
ax.legend()
plt.colorbar(sc, label='time $t$')
plt.show()
```

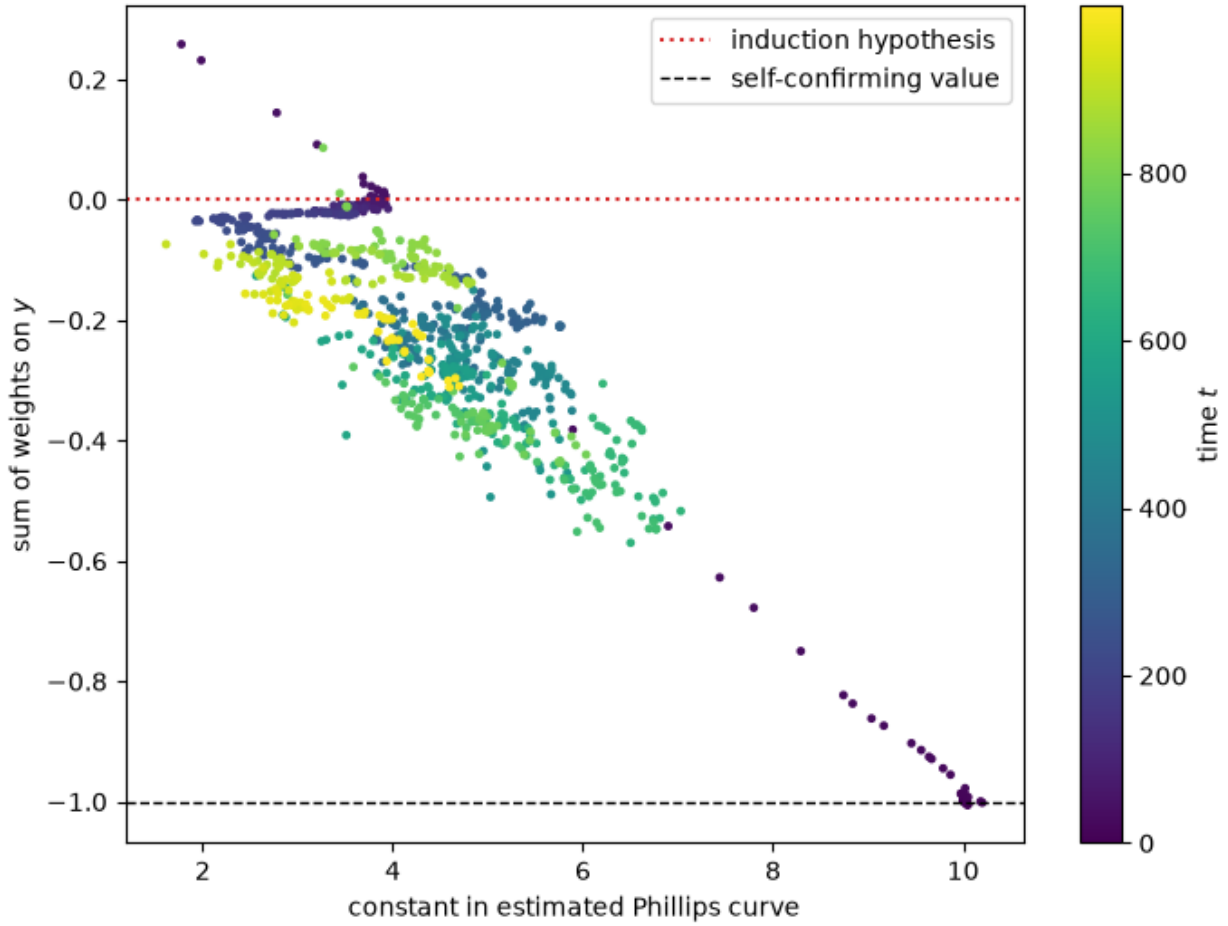


Fig. 103.5: The escape route in belief space, coloured by time

In Fig. 103.5 the beliefs leave the self-confirming point (constant 10, sum of weights -1) early and do not come back to it.

Afterwards they move back and forth below the induction line: as inflation drifts up, the sum of weights falls and the constant rises, and each stabilization sends the beliefs back up toward the induction line (sum of weights $= 0$).

That line is the end of the escape route along which the government learns a Solow-Tobin version of the natural-rate hypothesis and stabilizes inflation.

103.8 Relation to equilibria under forecast misspecification

The near-Ramsey episodes are reminiscent of the equilibria with optimal misspecified forecasts of *Optimal Misspecified Beliefs* and *Self-Confirming Equilibria*.

There, a forecasting model *without a constant but with a unit root* could closely approximate a true model that *includes* a constant.

Here an approximation with a similar flavor operates during the near-Ramsey episodes: the government's estimated Phillips curve, by driving the sum of weights toward zero, uses the induction hypothesis to approximate a constant — except that the approximated model is not fixed but changes as the government's own beliefs feed back through the Phelps problem.

103.9 Role of the discount factor

Sargent [1999] reports two findings about the government's discount factor δ:

- the recurrent stabilizations toward Ramsey depend on δ being near one; and
- lowering δ raises the inflation rate observed during the low-inflation episodes.

The book does not report the simulations behind them, so let's run some.

We simulate 16 seeds for each of three discount factors, reusing the first 16 constant-gain runs above for $\delta = 0.98$.

```
δ_runs = {0.98: cg_runs[:16]}
for δ in (0.90, 0.95):
    m = AdaptivePhillips(δ=δ)
    δ_runs[δ] = [simulate(m, λ=0.975, T_prior=300, n=1000, seed=s)
                 for s in range(16)]

for δ in sorted(δ_runs):
    runs = δ_runs[δ]
    firsts = [first_escape(sim['y']) for sim in runs]
    eig = np.linalg.eigvals(jacobian_T(AdaptivePhillips(δ=δ)) - np.eye(6)).real
    print(f"δ = {δ}: {sum(t is not None for t in firsts)} of {len(runs)} escape, "
          f"median first escape t = {np.median([t for t in firsts if t is not None]):.
↪0f}, "
          f"mean inflation for t ≥ 200 = {np.mean([sim['y'][200:].mean() for sim in␣
↪runs]):.2f}, "
          f"eig(DT − I) = {np.unique(np.round(eig, 4))}")
```

```
δ = 0.9: 16 of 16 escape, median first escape t = 56, mean inflation for t ≥ 200 =␣
↪1.70, eig(DT − I) = [-1.  -0.5]
δ = 0.95: 16 of 16 escape, median first escape t = 55, mean inflation for t ≥ 200␣
↪= 1.44, eig(DT − I) = [-1.  -0.5]
δ = 0.98: 16 of 16 escape, median first escape t = 54, mean inflation for t ≥ 200␣
↪= 1.24, eig(DT − I) = [-1.  -0.5]
```

Over this range of discount factors the second finding holds, but the first does not.

The timing of the first escape hardly depends on δ; what changes is how low inflation goes once the government has escaped.

That makes sense.

At the self-confirming equilibrium the Phelps rule is a constant, Nash inflation, for every δ, and the eigenvalues computed in *Mean dynamics near the self-confirming equilibrium* do not depend on δ either, so the forces that govern the escape are much the same.

Once the induction hypothesis is active, though, the government believes it faces a vertical long-run Phillips curve and a short-run tradeoff, and *Adaptive Expectations and the Phelps Problem* showed that a discounted Phelps problem of that kind stops short of zero inflation, more so the less patient the government is.

```
fig, ax = plt.subplots(figsize=(9, 4.5))
for δ in sorted(δ_runs):
    ax.plot(np.arange(19, 1000), moving_average(δ_runs[δ][1]['y']),
            lw=0.9, label=rf'$\delta = {δ}$')
ax.axhline(0, color='k', lw=0.5)
ax.set_xlabel('$t$')
ax.set_ylabel('20-period moving average of $y_t$')
ax.legend()
plt.show()
```

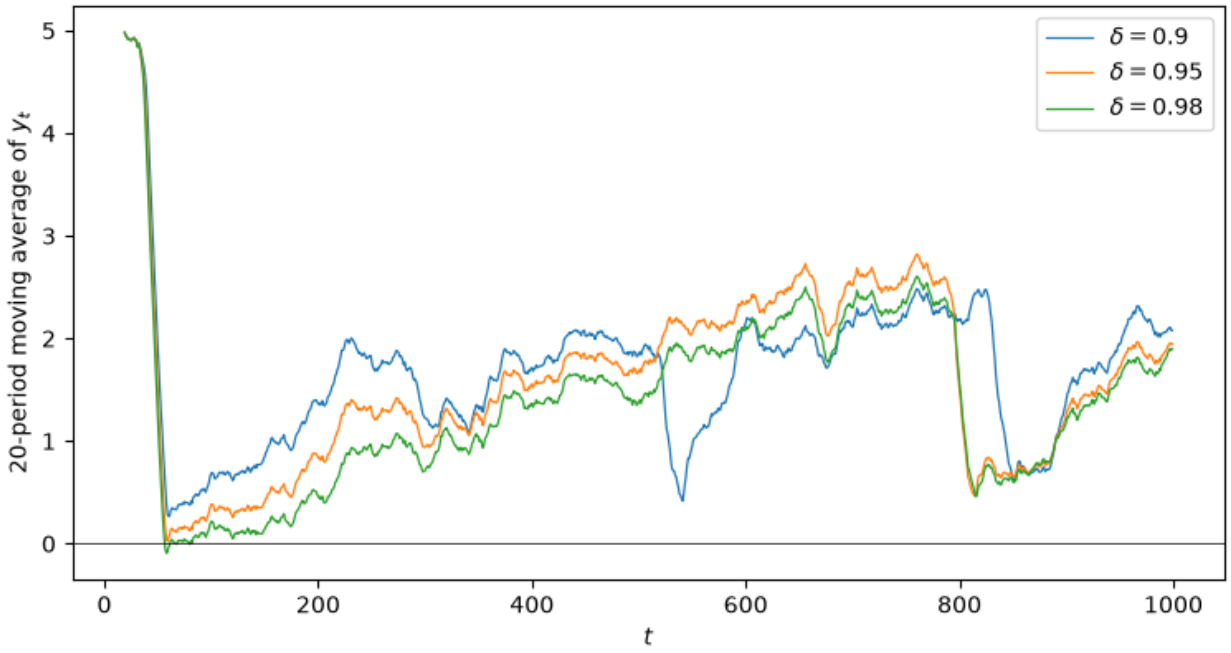


Fig. 103.6: The government's discount factor shifts the level of inflation after an escape, not the escape's timing

Fig. 103.6 shows one seed: the three governments escape at about the same date, and right after the escape the less patient ones settle at higher inflation.

103.10 Anticipated utility

The adaptive model is an example of what David Kreps calls an **anticipated utility** model [Kreps, 1998].

The government adapts a temporarily misspecified model — a Phillips curve with fixed coefficients — to incorporate the most recent observations, and reoptimizes along the way.

In forming decisions at t, it acts as if its current estimate γ_{t-1} will govern the Phillips curve forever, using the same policy functional $h(\cdot)$ that would be optimal if γ were truly time-invariant.

This is a small departure from rational expectations: calendar time enters only through the drifting beliefs γ_t.

The government treats coefficient drift in two different ways.

When it *estimates*, it allows for drift: the constant gain discounts old data precisely because the government suspects that the Phillips curve wanders, and in the Kalman-filter form of (103.18) it attributes a positive covariance to the drift.

When it *decides*, it ignores drift: it acts as if γ_{t-1} will never change.

Anticipated utility thus stretches a property of linear-quadratic problems.

Certainty equivalence (*Certainty Equivalence*) lets a decision maker who faces *additive* uncertainty replace random variables by their expectations.

The anticipated-utility government does the same with its uncertain coefficients γ, which enter multiplicatively, where certainty equivalence no longer holds.

A Bayesian decision maker would instead take into account, when choosing y_t, that its beliefs will change (*Rational Learning and Rational Expectations*), and a robust one would guard against misspecification of its model.

Yet, as the simulations show, this modest departure from rationality is enough to generate a rich account of the rise and fall of U.S. inflation, and to supply underpinnings for the vindication of econometric policy evaluation.

103.11 Conclusions

For long stretches, an adaptive government learns to generate *better than Nash* outcomes.

These results come from the recurrent dynamics of adaptation: the mean dynamics that under least squares pull the system toward a self-confirming equilibrium continue to operate, but under a constant gain the noise lets the system recurrently escape toward the Ramsey outcome.

Starting from a self-confirming equilibrium, the adaptive algorithm gradually makes the government put enough weight on the induction hypothesis that chance observations eventually promote a stabilization.

With a decreasing gain the same forces operate on a clock that keeps slowing down: least squares converges to the self-confirming equilibrium only asymptotically, and in a finite sample an early escape can stick.

Adaptation makes the government's beliefs a hidden state that imparts serial correlation into inflation — so that an outside forecaster would do well to use a random-coefficients model, or to make the constant adjustments that Lucas noted in his Critique [Lucas, 1976].

(Unemployment stays serially uncorrelated in this model, because the public knows the government's rule and only inflation surprises move unemployment.)

In this sense the adaptive models contain the underpinnings for vindicating econometric policy evaluation — the second of the two stories of *The Rise and Fall of U.S. Inflation*.

That last observation is a testable one, and it is worth recording what the model predicts before anyone looks.

If the government's beliefs really are a hidden drifting state, then a reduced-form description of post-war inflation and unemployment ought to show *drifting coefficients*, not merely drifting shock variances.

The model also predicts how the persistence of inflation should move with its level.

At the self-confirming equilibrium inflation is serially uncorrelated, while after a constant-gain escape it is persistent (Exercise 103.12.4), so low inflation should come with high persistence.

Drifts and Volatilities fits a drifting-coefficients model to the data and reports the verdict: persistence rose and fell with the level of inflation, and it fell during the Volcker disinflation, the opposite of this prediction.

Before that, the suite continues with the theory.

Escaping Nash Inflation characterizes the escape route analytically, *Priors, Escapes, and Learning Cycles* shows how the government's prior about coefficient drift shapes both the mean dynamics and the escapes, and *Econometric Policy Evaluation* follows the book's chapter 9 in estimating the adaptive models on U.S. data.

103.12 Exercises

Exercise 103.12.1

Everything above hinges on the constant gain, so it is worth seeing how much.

Simulate the classical adaptive model with $\lambda \in \{0.99, 0.975, 0.95\}$ and $T_0 = 300$, for 24 seeds each, and measure how often inflation escapes toward Ramsey.

Use two statistics that measure *frequency*:

- the median date of the first escape, defined as in *Least squares learning: a decreasing gain*; and
- the number of escape episodes per 1000 periods, where an episode starts when the sum of weights on current and lagged inflation rises above -0.1, and the sum must fall back below -0.3 before another episode can start.

Also compute the fraction of periods with $y_t < 2$.

How does a larger gain (smaller λ, faster discounting of the past) affect each statistic?

Why is the fraction of periods with low inflation a poor measure of how often the government escapes?

Note

A more ambitious version of this exercise is to build the **Keynesian** adaptive model, in which the government fits the Phillips curve in the reverse direction, regressing inflation on unemployment.
The regressors are $X_{K,t} = [U_t \quad U_{t-1} \quad U_{t-2} \quad y_{t-1} \quad y_{t-2} \quad 1]^\top$; the government estimates β in $y_t = \beta^\top X_{K,t} + \varepsilon_{K,t}$, then inverts to γ using the inversion formulas $\gamma_1 = \beta_1^{-1}$, $\gamma_{-1} = -\beta_{-1}/\beta_1$ of *Self-Confirming Equilibria* before solving the Phelps problem.
The Keynesian self-confirming inflation rate is 10.
In the book's simulation ($\lambda = 0.99$, $T_0 = 300$, Figures 8.3a–8.3d), the Keynesian system "spends much time away from the self-confirming equilibrium, and, like the classical system, recurrently escapes to the Ramsey outcome" [Sargent, 1999].
Its Jacobian has a repeated eigenvalue of $-.2$, which makes the mean dynamics slow and, in the book's words, "sustains long episodes near the Ramsey outcome".

Solution

We count episodes with a small function and reuse the 24 runs with $\lambda = 0.975$ from above.

```
def escape_episodes(sw, up=-0.1, down=-0.3):
    "Upcrossings of `up` by the sum of weights, re-armed only below `down`."
    count, armed = 0, True
    for v in sw:
        if armed and v > up:
            count, armed = count + 1, False
        elif not armed and v < down:
            armed = True
    return count


gain_runs = {0.975: cg_runs}
for λ in (0.99, 0.95):
    gain_runs[λ] = [simulate(model, λ=λ, T_prior=300, n=1000, seed=s)
                    for s in range(24)]

print(f"{'λ':>6} {'median first escape':>20} {'episodes per 1000':>18} "
      f"{'fraction y < 2':>15}")
for λ in (0.99, 0.975, 0.95):
    runs = gain_runs[λ]
    firsts = [first_escape(sim['y']) for sim in runs]
    episodes = [escape_episodes(sim['sumweights']) * 1000 / len(sim['y'])
                for sim in runs]
    print(f"{λ:>6} {np.median([t for t in firsts if t is not None]):20.0f} "
          f"{np.mean(episodes):18.2f} "
          f"{np.mean([(sim['y'] < 2).mean() for sim in runs]):15.2f}")
```

```
     λ  median first escape  episodes per 1000  fraction y < 2
  0.99                   90               1.00            0.89
 0.975                   56               2.62            0.83
  0.95                   38               6.08            0.85
```

A larger gain makes escapes come sooner and more often: the median first escape comes earlier, and the number of episodes per 1000 periods rises, as λ falls.

The fraction of periods with $y_t < 2$ barely moves, and it need not move monotonically.

The reason is that a larger gain also speeds the return trip, since the clock of *Least squares learning: a decreasing gain* runs faster.

Each near-Ramsey spell is therefore shorter, and inflation within it is noisier.

The fraction of time spent near Ramsey mixes how often the government escapes with how long it stays, and as the gain rises those two move in opposite directions.

Exercise 103.12.2

Fig. 103.1 and Fig. 103.2 show that whether least squares escapes depends on the tightness of the prior.

Simulate the least squares system ($\lambda = 1$) for $T_0 \in \{300, 800, 5000\}$, with 24 seeds and 1000 periods each.

1. For each T_0, report how many paths escape, the median date of the first escape, and how many of the escaped paths return above 4.
2. For the paths that escape, compute the artificial time $\tau = \log\left((T_0 + t)/T_0\right)$ at the first escape. Use the answer to explain why the tight prior rarely escapes within 1000 periods.
3. Explain why no escaped path returns, although least squares converges to the self-confirming equilibrium. Is the eigenvalue $-\frac{1}{2}$ of *Mean dynamics near the self-confirming equilibrium*, which sits at the boundary of $\sqrt{t}$

convergence, the reason?

Solution

```
ls_prior_runs = {T0: [simulate(model, λ=1.0, T_prior=T0, n=1000, seed=s)
                      for s in range(24)]
                 for T0 in (300, 800, 5000)}

for T0, runs in ls_prior_runs.items():
    firsts = np.array([t for t in (first_escape(sim['y']) for sim in runs)
                       if t is not None])
    line = f"T_0 = {T0:>4}: {len(firsts):>2} of {len(runs)} escape"
    if len(firsts):
        τ = np.log((T0 + firsts) / T0)
        line += (f", median first escape t = {np.median(firsts):.0f} "
                 f"(τ = {np.median(τ):.2f}), "
                 f"{sum(returns_to_nash(sim['y']) for sim in runs)} return")
    else:
        line += f"; τ at t = 1000 is only {np.log((T0 + 1000) / T0):.2f}"
    print(line)
```

```
T_0 =  300: 21 of 24 escape, median first escape t = 132 (τ = 0.36), 0 return
T_0 =  800: 15 of 24 escape, median first escape t = 270 (τ = 0.29), 0 return
T_0 = 5000:  0 of 24 escape; τ at t = 1000 is only 0.18
```

A looser prior escapes more often and sooner, and no escaped path returns.

Measured in artificial time, first escapes under the two looser priors come after a few tenths of a unit.

With $T_0 = 5000$, 1000 periods supply only $\log(6000/5000) \approx 0.18$ units of artificial time, fewer than most escaping paths under the looser priors needed.

No escaped path returns because, as *Least squares learning: a decreasing gain* showed, the return needs many units of artificial time, and the least squares clock slows down exponentially.

The eigenvalue $-\frac{1}{2}$ is not the reason.

It governs how fast decreasing-gain estimates converge, and their asymptotic distribution, once they are near the self-confirming equilibrium.

It says nothing special about the probability of leaving the self-confirming equilibrium, which the large-deviation objects of the primer govern, or about the speed of the return from far away, which depends on the global mean dynamics; those are slow near the induction hypothesis.

Priors, Escapes, and Learning Cycles shows how, in a Kalman filter with a gain that does not vanish, the government's prior produces escapes that stick, and the bullets on decreasing and constant gains in *One recursion, many learners* summarize the general theory.

Exercise 103.12.3

Mean dynamics near the self-confirming equilibrium computed the eigenvalues of $DT - I$ at $\theta = 1$.

1. Check that (103.19) is exact away from the self-confirming equilibrium. Freeze the government's beliefs at those that one of the constant-gain runs holds at $t = 1000$, simulate the economy for 50,000 periods under the frozen rule, regress U_t on $X_{C,t}$, and compare the coefficients with `T_map`.

2. Compute the eigenvalues of $DT - I$ at the self-confirming equilibrium for $\theta \in \{0.5, 1, 2\}$, holding the other parameters fixed, and compare them with $-1/(1+\theta^2)$.
3. For which of these θ do least squares estimates converge at the usual $\sqrt{t}$ rate? Compare with *Least squares learning converges to Nash* and *Uniqueness and E-stability*.

Solution

For the first part we freeze the rule and simulate.

```
y_frozen = cg_runs[1]['y']
h = model.phelps_h(y_frozen)
rng = np.random.default_rng(0)
N = 50_000
v1 = model.σ1 * rng.standard_normal(N)
v2 = model.σ2 * rng.standard_normal(N)
U = model.U_star - model.θ * v2 + v1
y = np.zeros(N)
y[:2] = h[4]
for t in range(2, N):
    y[t] = h @ [U[t-1], U[t-2], y[t-1], y[t-2], 1.0] + v2[t]

X_C = np.column_stack([y[2:], U[1:-1], U[:-2], y[1:-1], y[:-2], np.ones(N - 2)])
coef = np.linalg.lstsq(X_C, U[2:], rcond=None)[0]
print("regression of U on X_C:", np.round(coef, 3))
print("T(y) from the formula: ", np.round(T_map(model, y_frozen), 3))
```

```
regression of U on X_C: [-0.998  0.007  0.056  0.044  0.124  6.186]
T(y) from the formula:  [-1.     0.011  0.054  0.049  0.126  6.169]
```

The two agree up to sampling error, because v_{1t} is independent of every regressor.

Now the eigenvalues.

```
for θ in (0.5, 1.0, 2.0):
    m = AdaptivePhillips(θ=θ)
    eig = np.sort(np.linalg.eigvals(jacobian_T(m) - np.eye(6)).real)
    print(f"θ = {θ}: eig(DT - I) = {np.round(eig, 4)}, "
          f"-1/(1 + θ²) = {-1 / (1 + θ**2):.4f}")
```

```
θ = 0.5: eig(DT - I) = [-1.  -0.8 -0.8 -0.8 -0.8 -0.8], -1/(1 + θ²) = -0.8000
θ = 1.0: eig(DT - I) = [-1.  -0.5 -0.5 -0.5 -0.5 -0.5], -1/(1 + θ²) = -0.5000
θ = 2.0: eig(DT - I) = [-1.  -0.2 -0.2 -0.2 -0.2 -0.2], -1/(1 + θ²) = -0.2000
```

For every θ the eigenvalues are -1 once and $-1/(1+\theta^2)$ five times.

The $\sqrt{t}$ rate requires $-1/(1+\theta^2) < -\frac{1}{2}$, that is, $\theta < 1$.

So least squares converges at the usual rate when $\theta = 0.5$, sits on the boundary when $\theta = 1$, and converges more slowly when $\theta = 2$, with the slow mean dynamics of a repeated eigenvalue -0.2.

These are the same numbers as the slope $\mathcal{M} = -1/(1+\theta^2)$ of the scalar mean dynamics in *Least squares learning converges to Nash* and the classical E-stability eigenvalues of *Uniqueness and E-stability*.

A steeper Phillips curve makes the static Phelps rule respond more strongly to a believed coefficient on a lag, which pushes the diagonal of DT, $\theta^2/(1+\theta^2)$, toward one and weakens the pull of the mean dynamics.

Exercise 103.12.4

The conclusions claim that inflation is serially uncorrelated at the self-confirming equilibrium but persistent after a constant-gain escape.

Check the claim, measuring persistence by the slope coefficient in a least squares regression of y_t on a constant and y_{t-1}.

1. For the least squares economy with $T_0 = 5000$, which stays near the self-confirming equilibrium, compute it over the first 200 periods of each path of Exercise 103.12.2.
2. For the constant-gain economies of Exercise 103.12.1, compute it over the 200 periods that begin 20 periods after the first escape.
3. Do the same for the least squares paths with $T_0 = 800$ that escape.
4. What relation between the level of inflation and its persistence does the model predict? Compare with *Drifts and Volatilities*.

Solution

```
def ar1(z):
    "Slope of a least squares regression of z_t on a constant and z_{t-1}."
    X = np.column_stack([np.ones(len(z) - 1), z[:-1]])
    return np.linalg.lstsq(X, z[1:], rcond=None)[0][1]


def after_escape(runs, start=20, length=200):
    "AR(1) coefficients over a window that starts `start` periods after the first␣
 ↪escape."
    out = []
    for sim in runs:
        t = first_escape(sim['y'])
        if t is not None and t + start + length <= len(sim['y']):
            out.append(ar1(sim['y'][t + start:t + start + length]))
    return np.array(out)


def summarize(label, values):
    print(f"{label:28s} median {np.median(values):+.2f}, "
          f"range {values.min():+.2f} to {values.max():+.2f} ({len(values)} paths)
 ↪")


summarize("least squares, T_0 = 5000:",
          np.array([ar1(sim['y'][:200]) for sim in ls_prior_runs[5000]]))
for λ in (0.99, 0.975, 0.95):
    summarize(f"constant gain, λ = {λ}:", after_escape(gain_runs[λ]))
summarize("least squares, T_0 = 800:", after_escape(ls_prior_runs[800]))
```

```
least squares, T_0 = 5000:   median +0.01, range -0.14 to +0.21 (24 paths)
constant gain, λ = 0.99:     median +0.10, range -0.11 to +0.29 (24 paths)
constant gain, λ = 0.975:    median +0.38, range +0.06 to +0.70 (24 paths)
constant gain, λ = 0.95:     median +0.73, range +0.58 to +0.88 (24 paths)
least squares, T_0 = 800:    median +0.06, range -0.02 to +0.21 (15 paths)
```

Near the self-confirming equilibrium the AR(1) coefficient is centered on zero.

After a constant-gain escape inflation is persistent, and more so the larger the gain; at $\lambda = 0.99$ the effect is small.

After a least squares escape that sticks, persistence is small again.

So the persistence comes from the motion of the government's beliefs, which a constant gain keeps alive, not from low inflation as such.

Under a constant gain the model therefore predicts that persistence should be high when inflation is low and low when inflation is high, with the transition from high to low inflation raising persistence.

Drifts and Volatilities (*Inflation persistence*) finds the reverse in U.S. data: persistence rose and fell with the level of inflation, and it fell during the Volcker disinflation.

CHAPTER

FOUR

ESCAPING NASH INFLATION

Contents

- *Escaping Nash Inflation*
  - *Overview*
  - *The model*
  - *The self-confirming equilibrium*
  - *Adaptation and the mean dynamics*
  - *Escape dynamics as a control problem*
  - *The changed-measure escape path*
  - *The race of four ODEs*
  - *Mean dynamics reinforce the escape*
  - *The experimentation trap*
  - *Simulating escapes*
  - *Escape frequency and model richness*
  - *Escaping volatile inflation*
  - *Exercises*

104.1 Overview

> If an unlikely event occurs, it is very likely to occur in the most likely way.
>
> – Michael Harrison

This lecture is the analytical companion to *Adaptive Learning and Escape Dynamics*.

It follows Cho *et al.* [2002] (CWS), who give a *deterministic characterization* of the escape dynamics that chapter 8 of Sargent [1999] found by simulation.

In *Adaptive Learning and Escape Dynamics* we watched a constant-gain government recurrently escape the Nash self-confirming equilibrium and spend long spells near the Ramsey outcome, but we described those escapes only informally and by simulation.

CWS show that the escapes are governed by their own ordinary differential equation, obtained from the **theory of large deviations**.

The picture that emerges has two deterministic pieces:

- the **mean dynamics**, an ODE that pulls the government's beliefs *toward* the self-confirming equilibrium; and
- the **escape dynamics**, a second ODE that, driven by a "most likely unlikely" sequence of shocks, pushes beliefs *away* from it, toward the beliefs that support the Ramsey outcome.

The remarkable finding is that the escape has a *dominant path*: conditional on escaping, the government's beliefs follow a nearly deterministic route along which it temporarily learns a version of the natural-rate hypothesis and cuts inflation.

These escapes recur because the government uses constant-gain least squares, as in Sargent [1999] and CWS.

The governments of Sims [1988] and Chung [1990] instead ran Kalman filters on drifting coefficients; in their simulations a government that escaped did so once and then stayed near Ramsey, apparently indefinitely, and *Priors, Escapes, and Learning Cycles* traces the difference to the government's prior.

The large-deviation results are statements about small gains.

Simulating escapes simulates the model and finds that, at the gains used in this suite, escapes are frequent rather than rare.

We work with the analytically tractable **static** model, in which the government regresses unemployment on current inflation and a constant.

Adaptive Learning and Escape Dynamics simulates CWS's richer *dynamic* model, in which the government's regression also includes lags; we compare the two near the end of the lecture.

A government learning the Phillips curve presents the same static model, with the same calibration, in condensed form.

Readers arriving from there should note that it writes the beliefs as α (intercept first) and the gain as γ, while here γ denotes the beliefs and ε the gain.

The next lecture, *Priors, Escapes, and Learning Cycles*, asks how the government's *prior* about parameter drift reshapes both the mean dynamics and the escape dynamics.

Let's start with our imports:

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.integrate import solve_ivp
from scipy.linalg import solve_discrete_are
```

104.2 The model

The true economy is the natural-rate model of Kydland and Prescott [1977],

$$U_n = u - \theta(\pi_n - \hat{x}_n) + \sigma_1 W_{1n}, \qquad \pi_n = x_n + \sigma_2 W_{2n}, \qquad \hat{x}_n = x_n, \tag{104.1}$$

with $\theta, u > 0$ and $W_n = (W_{1n}, W_{2n})^\top$ IID standard Gaussian.

The government does not know (104.1).

In the static model it fits a non-expectational Phillips curve by regressing unemployment on inflation and a constant,

$$U_n = \gamma_1 \pi_n + \gamma_{-1} + \eta_n, \tag{104.2}$$

with beliefs $\gamma = (\gamma_1, \gamma_{-1})$ (slope and intercept), and it treats η_n as exogenous.

Note

We keep the ordering of Cho *et al.* [2002]: the slope comes first, and the regressors are $\Phi = (\pi, 1)$.

This also matches the γ_1, γ_{-1} convention of *Self-Confirming Equilibria*.

Priors, Escapes, and Learning Cycles and *A government learning the Phillips curve* follow Sargent and Williams [2005] and put the intercept first.

Each lecture follows its source paper; the warning box in *Priors, Escapes, and Learning Cycles* translates between the two.

Believing (104.2), the government solves the *Phelps problem*, whose static best response sets inflation to the constant

$$x(\gamma) = -\frac{\gamma_{-1}\,\gamma_1}{1+\gamma_1^2}. \tag{104.3}$$

Three beliefs are worth naming, following CWS:

- **Belief 1 (Nash):** $\gamma_1 = -\theta$ with an intercept that makes $x = \theta u$, the time-consistent outcome of Kydland and Prescott [1977].
- **Belief 2 (Ramsey):** $\gamma_1 = 0$, so the government perceives no tradeoff and sets $x = 0$.
- **Belief 3 (induction):** in CWS's dynamic model, coefficients on current and lagged inflation that sum to zero, which for a patient government also sends inflation to 0.

In the static model there is only one coefficient on inflation, so Belief 3 coincides with Belief 2: a vertical perceived Phillips curve "trivially satisfies the induction property" (CWS, section 6.2).

```
class EscapeModel:
    "CWS 2002 static model. γ = (γ₁ slope, γ₋₁ intercept), regressors Φ = (π, 1)."

    def __init__(self, θ=1.0, u=5.0, σ1=0.3, σ2=0.3):
        self.θ, self.u, self.σ1, self.σ2 = θ, u, σ1, σ2

    def γ_sce(self):
        "Self-confirming beliefs (slope -θ, intercept u(1+θ²))."
        return np.array([-self.θ, self.u * (1 + self.θ**2)])

    def x(self, γ):
        γ1, γm1 = γ
        return -γm1 * γ1 / (1 + γ1**2)

    def M(self, γ):
        "E[ΦΦ'] with Φ = (π, 1)."
        x = self.x(γ)
        return np.array([[x**2 + self.σ2**2, x], [x, 1.0]])

    def g_bar(self, γ):
        "Mean-dynamics forcing E[Φ(U - Φ'γ)] = M(T(γ) - γ)."
        x = self.x(γ)
        E_ΦU = np.array([x * self.u - self.θ * self.σ2**2, self.u])
        return E_ΦU - self.M(γ) @ γ
```

104.3 The self-confirming equilibrium

A self-confirming equilibrium is a belief that reproduces itself: the population regression coefficients of (104.2), computed on the data the government generates by acting on γ, equal γ.

Writing $T(\gamma)$ for those population coefficients and $\bar{M}(\gamma) = \mathbb{E}\,\Phi\Phi^\top$ for the population moment matrix of the regressors, CWS show

$$\bar{g}(\gamma) \equiv \mathbb{E}\left[\Phi(U - \Phi^\top\gamma)\right] = \bar{M}(\gamma)\,(T(\gamma) - \gamma)\,,$$

so a self-confirming equilibrium solves $\bar{g}(\gamma) = 0$.

Under a constant policy x, only the surprise $\sigma_2 W_{2n}$ moves inflation, so $T(\gamma) = (-\theta,\ u + \theta x(\gamma))$.

For the static model the equilibrium is therefore the intersection of the line $\gamma_1 = -\theta$ with the parabola $\gamma_{-1} = u(1+\gamma_1^2)$, a unique point that supports the Nash outcome of *The Credibility Problem*.

It is the same self-confirming equilibrium constructed in *Self-Confirming Equilibria*, written here in the static (constant-plus-slope) special case that *Priors, Escapes, and Learning Cycles* also uses, and it is what Kydland and Prescott [1977] call a *consistent* equilibrium (*A government learning the Phillips curve*).

```
model = EscapeModel(θ=1.0, u=5.0, σ1=0.3, σ2=0.3)
γ_sce = model.γ_sce()

print(f"self-confirming beliefs γ = {γ_sce}  (slope -θ, intercept u(1+θ²))")
print(f"self-confirming inflation x = {model.x(γ_sce):.2f}  (= Nash = θu)")
print(f"check g_bar = {model.g_bar(γ_sce)}")
```

```
self-confirming beliefs γ = [-1. 10.]  (slope -θ, intercept u(1+θ²))
self-confirming inflation x = 5.00  (= Nash = θu)
check g_bar = [0. 0.]
```

The figure plots the two loci whose intersection is the equilibrium.

```
γ1_grid = np.linspace(-2, 1, 200)
fig, ax = plt.subplots(figsize=(7, 6))
ax.plot(γ1_grid, model.u * (1 + γ1_grid**2), label=r'$\gamma_{-1} = u(1+\gamma_1^2)$',
 ↪ lw=2)
ax.axvline(-model.θ, color='C1', ls='--', label=r'$\gamma_1 = -\theta$')
ax.plot(γ_sce[0], γ_sce[1], 'ko', ms=8)
ax.annotate('SCE (Nash)', γ_sce, (γ_sce[0] + 0.1, γ_sce[1] + 1))
ax.set_xlabel(r'slope $\gamma_1$')
ax.set_ylabel(r'intercept $\gamma_{-1}$')
ax.set_ylim(0, 25)
ax.legend()
plt.show()
```

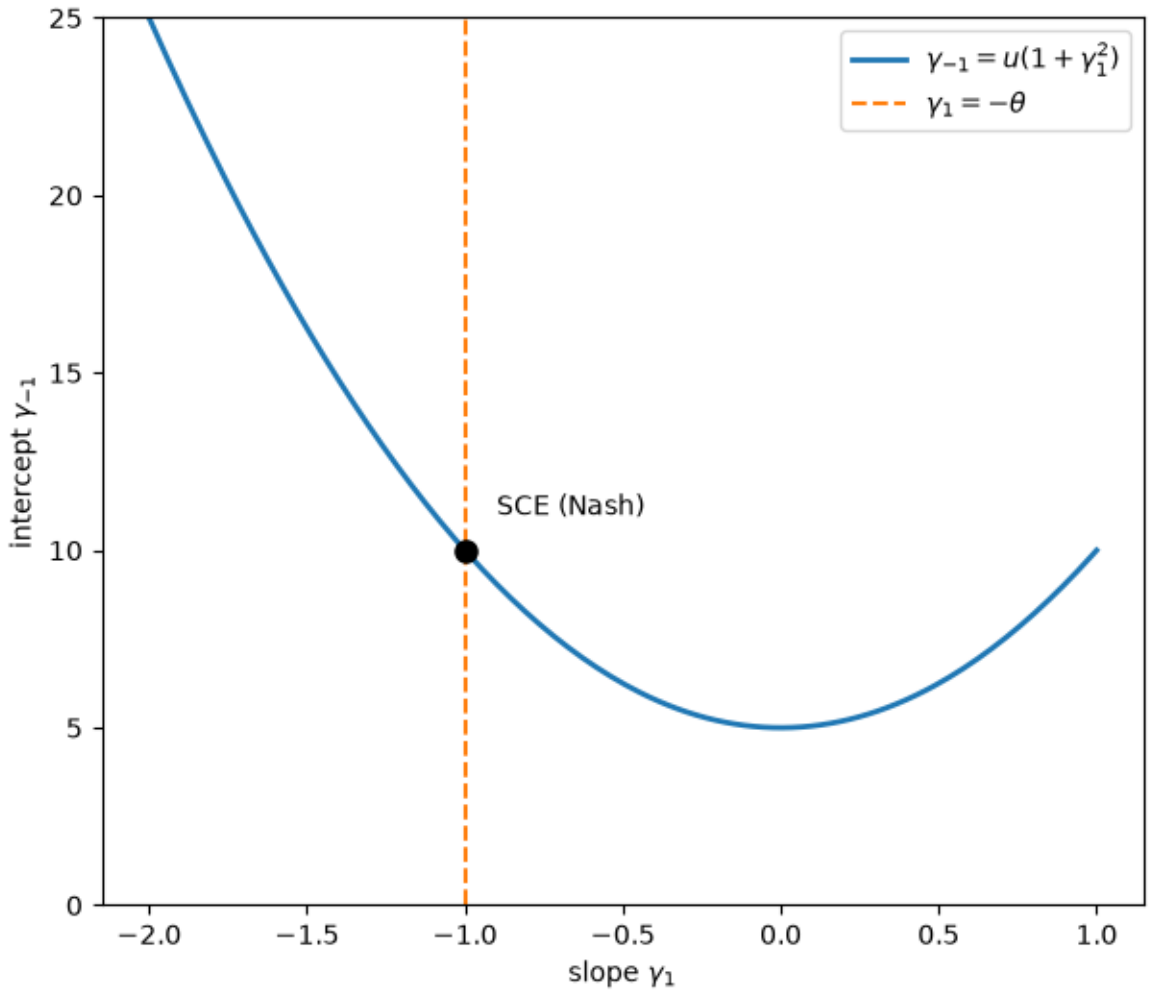


Fig. 104.1: The unique self-confirming equilibrium in belief space

104.4 Adaptation and the mean dynamics

We make the government adaptive.

Each period it sets $x_n = x(\gamma_n)$, observes (π_n, U_n), and updates its estimates by constant-gain recursive least squares,

$$R_n = R_{n-1} + \varepsilon\left(\Phi_n\Phi_n^\top - R_{n-1}\right), \qquad \gamma_{n+1} = \gamma_n + \varepsilon R_n^{-1}\Phi_n\left(U_n - \Phi_n^\top\gamma_n\right), \tag{104.4}$$

where $\Phi_n = (\pi_n, 1)^\top$ and $\varepsilon > 0$ is the gain.

This is the recursion of `constant_gain` in *A government learning the Phillips curve*, with the regressors in CWS's order.

CWS's equations (4.24)–(4.25) date R one period earlier, which makes no difference as $\varepsilon \to 0$.

Because the government acts on its current estimate as if it were permanent, this is an *anticipated-utility* model in the sense of Kreps [1998].

The literature on least squares learning in self-referential models [Evans and Honkapohja, 2001, Marcet and Sargent, 1989, Woodford, 1990] shows that, as $\varepsilon \to 0$ with time measured as $t = n\varepsilon$, the beliefs are approximated by the **mean-dynamics** ODE

$$\dot\gamma = R^{-1}\bar g(\gamma), \qquad \dot R = \bar M(\gamma) - R. \tag{104.5}$$

One recursion, many learners derives this kind of ODE for a scalar recursion and connects its stability to E-stability.

> **Note**
>
> **Notation across lectures.** *A primer on recursive algorithms* stacks the beliefs as $\phi = (\gamma, \mathrm{col}(R))$ and writes the mean dynamics as $\dot\phi = b(\phi)$.
>
> In the static model $b(\phi) = \left(R^{-1}\bar g(\gamma),\ \mathrm{col}(\bar M(\gamma) - R)\right)$, so (104.5) is that ODE written out.
>
> The escape dynamics below add a forcing v to the γ block only.
>
> *Priors, Escapes, and Learning Cycles* replaces the scaled inverse moment matrix εR^{-1} by the covariance matrix P of a Kalman filter, and orders γ intercept first.

A rest point of (104.5) is a self-confirming equilibrium, and CWS show that this ODE is *globally stable* about it.

Locally, with $R = \bar M(\bar\gamma)$, the γ block has eigenvalues -1 and $-1/(1+\theta^2)$: these are the E-stability eigenvalues of *Uniqueness and E-stability* and *Mean dynamics near the self-confirming equilibrium*.

```
R_bar = model.M(γ_sce)                          # moment matrix at the SCE
h = 1e-6
J = np.column_stack([(model.g_bar(γ_sce + h * e) - model.g_bar(γ_sce - h * e)) / (2 *␣
 ↪h)
                     for e in np.eye(2)])
print("eigenvalues of R̄⁻¹ Dḡ at the SCE:", np.linalg.eigvals(np.linalg.solve(R_bar,␣
 ↪J)).round(4))
print("-1/(1+θ²) =", -1 / (1 + model.θ**2))
```

```
eigenvalues of R̄⁻¹ Dḡ at the SCE: [-1.  -0.5]
-1/(1+θ²) = -0.5
```

So under the mean dynamics alone, the adaptive government is drawn to Nash inflation.

But the route can be circuitous.

The next figure starts the beliefs part of the way toward Ramsey, at $\gamma(0) = \bar{\gamma} + (0.4, -3)$, in two ways.

In the first, the moment matrix R starts at its Nash-era value $\bar{M}(\bar{\gamma})$; in the second, it starts at the moments $\bar{M}(\gamma(0))$ that the new beliefs would generate.

```
def mean_ode(t, z, model):
    y, R = z[:2], z[2:].reshape(2, 2)
    return np.concatenate([np.linalg.solve(R, model.g_bar(y)),
                           (model.M(y) - R).ravel()])

y0 = y_sce + np.array([0.4, -3.0])
fig, ax = plt.subplots(figsize=(8, 4.5))
starts = [(R_bar, r'$R(0) = \bar M(\bar\gamma)$ (Nash-era moments)', 'R(0) = M(SCE)'),
          (model.M(y0), r'$R(0) = \bar M(\gamma(0))$', 'R(0) = M(y(0))')]
for R0, label, short in starts:
    sol = solve_ivp(mean_ode, [0, 60], np.concatenate([y0, R0.ravel()]), args=(model,
 ↪),
                    max_step=0.1, rtol=1e-9, atol=1e-11)
    x_path = -sol.y[1] * sol.y[0] / (1 + sol.y[0]**2)
    ax.plot(sol.t, x_path, label=label)
    k = np.argmin(x_path)
    print(f"{short:15s}: x(0) = {x_path[0]:.2f}, min x = {x_path[k]:.2f} "
          f"at t = {sol.t[k]:.2f}, x(60) = {x_path[-1]:.2f}")
ax.axhline(model.θ * model.u, color='k', ls='--', lw=1, label='Nash')
ax.set_xlabel('time $t = n\\varepsilon$')
ax.set_ylabel('inflation $x$')
ax.legend()
plt.show()
```

```
R(0) = M(SCE)  : x(0) = 3.09, min x = 1.68 at t = 1.33, x(60) = 5.00
R(0) = M(y(0)) : x(0) = 3.09, min x = 3.09 at t = 0.00, x(60) = 5.00
```

Both paths return to Nash, but they get there differently.

When R starts at the Nash-era moments, inflation first *falls* further, from about 3.1 to about 1.7, before it climbs back.

When R starts at $\bar{M}(\gamma(0))$, inflation rises monotonically.

The dip is the circuitous route of CWS (section 6.4) made visible, and it comes from R lagging behind $\bar{M}(\gamma)$.

We return to this lag in *Mean dynamics reinforce the escape*, where it turns out to drive the escape.

The mean dynamics alone cannot explain the recurrent visits to low inflation in the simulations of *Adaptive Learning and Escape Dynamics*.

For that we need a second force.

104.5 Escape dynamics as a control problem

Although the impact of the noise vanishes as $\varepsilon \to 0$, for a fixed positive gain unusual sequences of shocks can push beliefs a long way from the self-confirming equilibrium.

A primer on recursive algorithms sets out the general theory.

It builds the log moment-generating function H of the innovations, its Legendre transform L, and an **action functional** $S(T, \gamma) = \int_0^T L\, ds$ that measures how improbable a path of beliefs is.

The **dominant escape path** minimizes the action among paths that leave a neighbourhood G of $\bar{\gamma}$.

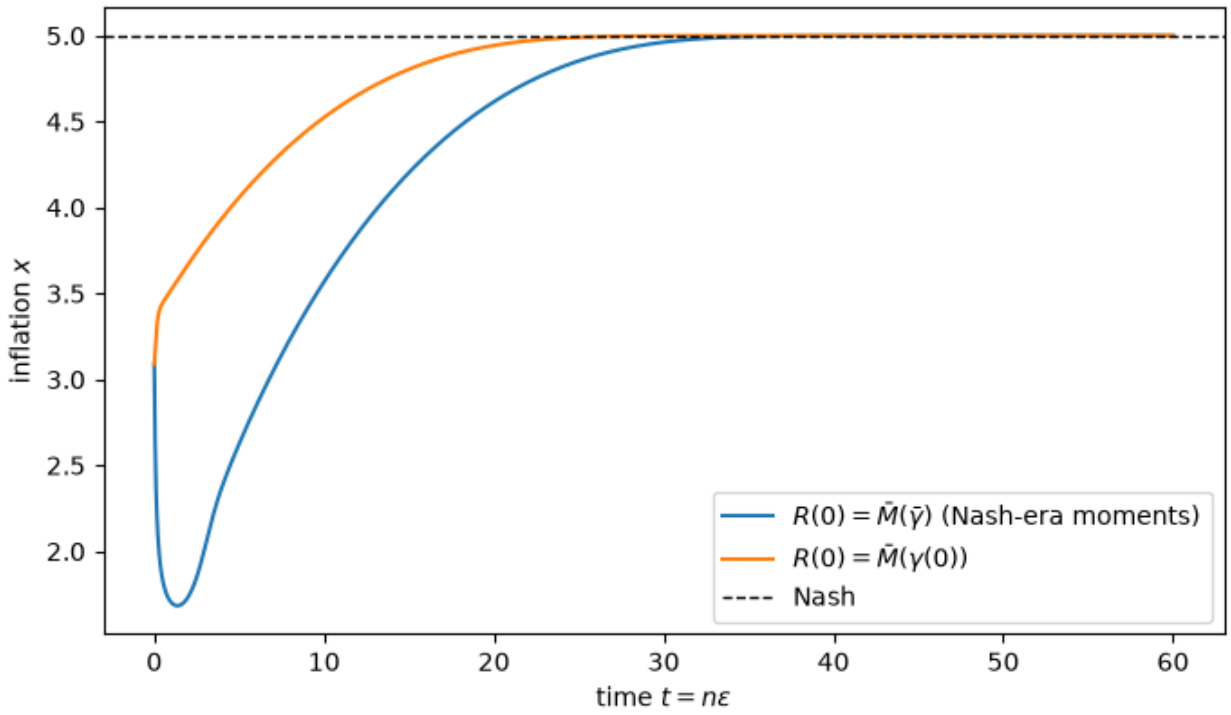


Fig. 104.2: Mean dynamics from a perturbed start: when R starts at its Nash-era value, inflation first dips

Here we need only the operational version.

CWS draw on Noah Williams's dissertation [Williams, 2001], published in corrected form as Williams [2019], to reduce the problem to a clean control problem.

The H-functional becomes a quadratic form with a normalizing matrix Q (a fourth-moment matrix obtained from Lyapunov equations), and the escape path solves

$$\bar{S} = \inf_{v(\cdot),\, T} \frac{1}{2} \int_0^T v(s)^\top Q(\gamma(s), R(s))^{-1} v(s)\, ds \tag{104.6}$$

subject to the *perturbed* mean dynamics

$$\dot{\gamma} = R^{-1} \bar{g}(\gamma) + v, \qquad \dot{R} = \bar{M}(\gamma) - R, \qquad \gamma(0) = \bar{\gamma},\ \gamma(T) \notin G. \tag{104.7}$$

Read (104.6) as a least squares problem: v is the extra "forcing" that the mean dynamics would need to escape, Q plays the role of a covariance matrix, and the least-cost forcing is the most likely unusual shock sequence.

Williams [2019] (p. 892) reports that the 2001 version computed H by misapplying a theorem of Worms (1999).

The discrepancy matters only for dynamic models; for the static model CWS computed H directly.

So the static-model results below stand, while CWS's rate for the dynamic model rests on the earlier shortcut.

Let τ^ε be the first time, measured as $t = n\varepsilon$, at which the beliefs leave G.

Two consequences (CWS, Theorem 5.3) tie the control problem to the stochastic model:

- $\varepsilon \log \mathbb{P}(\text{escape before } T) \to -\bar{S}$ and $\varepsilon \log \mathbb{E}\, \tau^\varepsilon \to \bar{S}$ as $\varepsilon \to 0$; and
- conditional on escaping, beliefs follow the dominant escape path with probability approaching one.

The first is a statement about the limit $\varepsilon \to 0$.

Williams [2019] (footnote 16) writes it as $\log \mathbb{E}\, \tau^{\varepsilon} \approx S_0 + \bar{S}/\varepsilon$ for some constant S_0 that the theory does not pin down.

So $\bar{S}$ governs how fast escape times grow as the gain shrinks, not their level.

For the static model CWS (Table 1) find $\bar{S} \approx 5 \times 10^{-4}$.

At the gains of *Adaptive Learning and Escape Dynamics* and *A government learning the Phillips curve* (0.025 and 0.02), $e^{\bar{S}/\varepsilon} \approx 1.02$, so the large-deviation factor by itself does not make escapes rare.

Simulating escapes measures how often they occur.

The escape dynamics, like the mean dynamics, are *deterministic*.

104.6 The changed-measure escape path

Solving (104.6) requires the matrix Q and a two-point boundary-value problem.

For the static model with binomial shocks, CWS (section 7) identify the escape route by a different and simpler argument.

They change the probability measure on the shocks to the most likely unusual distribution that pushes beliefs away from $\bar{\gamma}$, and average the least squares update under that measure.

The resulting escape forcing is strikingly simple:

$$v = R^{-1} \begin{bmatrix} \sigma_1 \sigma_2 \\ 0 \end{bmatrix}, \tag{104.8}$$

so the changed-measure escape path solves

$$\dot{\gamma} = R^{-1} \left(\bar{g}(\gamma) + \begin{bmatrix} \sigma_1 \sigma_2 \\ 0 \end{bmatrix} \right), \qquad \dot{R} = \bar{M}(\gamma) - R. \tag{104.9}$$

This path shares the dominant escape path's exit point, but it gets there faster.

CWS (footnote 14) warn that the transformed measure underestimates the time to escape.

In their Figure 7, the changed-measure path is shorter than even the shortest of 1,000 simulated escapes.

Let's integrate (104.9) from the self-confirming equilibrium until beliefs leave a circle of radius 5, the escape set CWS use for the static model.

```
def escape_ode(t, z, model):
    γ, R = z[:2], z[2:].reshape(2, 2)
    force = np.array([model.σ1 * model.σ2, 0.0])
    return np.concatenate([np.linalg.solve(R, model.g_bar(γ) + force),
                           (model.M(γ) - R).ravel()])

def left_circle(t, z, *args):
    "Terminal event: beliefs leave the radius-5 circle around the SCE."
    return 5.0 - np.linalg.norm(z[:2] - γ_sce)
left_circle.terminal = True
left_circle.direction = -1

z0 = np.concatenate([γ_sce, R_bar.ravel()])
esc = solve_ivp(lambda t, z: escape_ode(t, z, model), [0, 200], z0,
                events=left_circle, max_step=0.02, rtol=1e-9, atol=1e-11)

slope, intercept = esc.y[0], esc.y[1]
infl = -intercept * slope / (1 + slope**2)
```

The figure plots the two coefficients and the inflation rate they imply.

```
fig, axes = plt.subplots(1, 2, figsize=(12, 5))

axes[0].plot(esc.t, intercept, label='intercept $\\gamma_{-1}$')
axes[0].plot(esc.t, slope, label='slope $\\gamma_1$')
axes[0].axhline(0, color='C3', ls=':', lw=1, label='vertical curve ($\\gamma_1 = 0$)')
axes[0].set_xlabel('time $t = n\\varepsilon$')
axes[0].set_ylabel('coefficient')
axes[0].set_title('Changed-measure escape path (cf. CWS Figures 4 and 7)')
axes[0].legend()

axes[1].plot(esc.t, infl)
axes[1].axhline(model.θ * model.u, color='k', ls='--', lw=1, label='Nash')
axes[1].axhline(0, color='C2', ls=':', lw=1, label='Ramsey')
axes[1].set_xlabel('time $t = n\\varepsilon$')
axes[1].set_ylabel('inflation $x$')
axes[1].set_title('Inflation along the escape')
axes[1].legend()

plt.tight_layout()
plt.show()
```

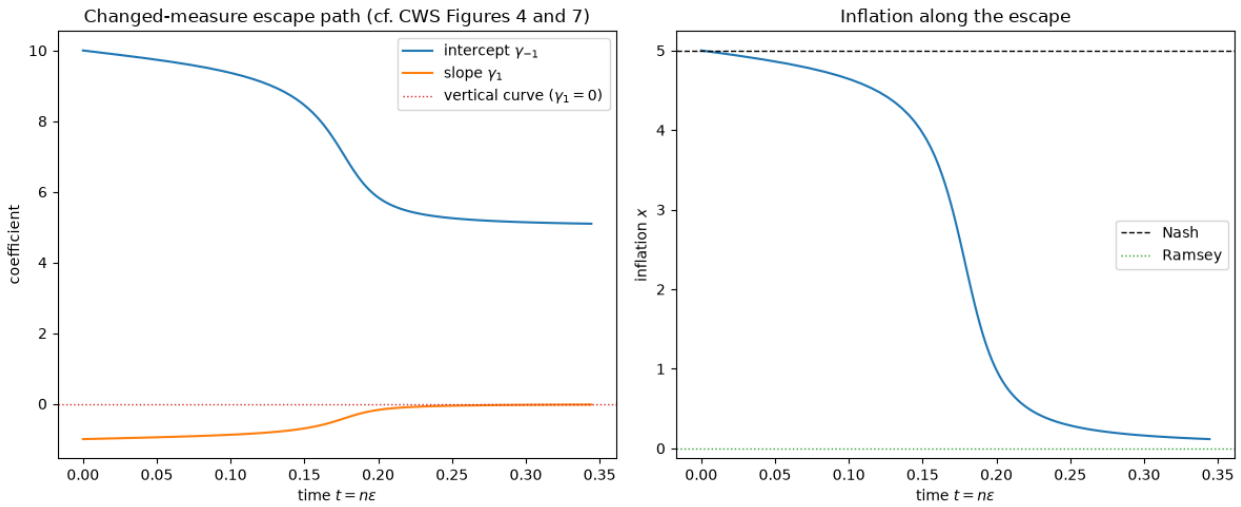


Fig. 104.3: The changed-measure escape path and the inflation rate along it

Here are the start and end points of the path.

```
print(f"exit time : t = {esc.t[-1]:.3f}")
print(f"slope     : {slope[0]:.2f} → {slope[-1]:.3f}   (CWS Table 1, radius 5: -0.025)
↪")
print(f"intercept : {intercept[0]:.1f} → {intercept[-1]:.2f}")
print(f"inflation : {infl[0]:.2f} → {infl[-1]:.2f}   (Nash {model.θ*model.u:.0f} →␣
↪Ramsey 0)")
```

```
exit time : t = 0.344
slope     : -1.00 → -0.023   (CWS Table 1, radius 5: -0.025)
```

(continues on next page)

(continued from previous page)

```
intercept : 10.0 → 5.10
inflation : 5.00 → 0.12   (Nash 5 → Ramsey 0)
```

Along the escape path the slope rises from its self-confirming value of -1 toward zero, and inflation falls from the Nash value toward Ramsey.

The path leaves the circle at a slope of about -0.023, close to the -0.025 that CWS (Table 1) report for the dominant path at the same radius.

It leaves at $t \approx 0.34$, while the dominant path in CWS's Figure 4 needs until about $t = 3$: the two paths share a route but not a timetable.

In the static model, a slope near zero means a vertical perceived Phillips curve, which trivially satisfies the induction hypothesis.

The dynamic model that *Adaptive Learning and Escape Dynamics* simulates escapes differently.

There the *sum* of the coefficients on current and lagged inflation rises toward zero, while the coefficient on current inflation stays negative, so a short-run Phillips curve survives while the long-run curve becomes vertical (CWS, section 6.2 and Figure 3).

In both models the government, by chance, generates enough inflation *experiments* to discover a good-enough version of the natural-rate hypothesis, and acts on it.

104.7 The race of four ODEs

Where does the escape forcing (104.8) come from?

With binomial shocks $W_{in} \in \{-1, +1\}$, each period brings one of four shock realizations $w \in \{-1, 1\}^2$, each with probability $1/4$.

An escape needs a run of updates that push beliefs consistently in one direction.

CWS show that the most efficient runs draw repeatedly from just two realizations whose updates lie in the same half-space, each with probability $1/2$; four such pairs exist.

Averaging the update over a pair gives a *conditional* mean ODE,

$$\dot{\gamma} = R^{-1}\left(\bar{g}(\gamma) + f_j(\gamma)\right), \qquad \dot{R} = \bar{M}(\gamma) - R, \tag{104.10}$$

with forcings

$$f_1 = \begin{bmatrix}\sigma_1\sigma_2\\ 0\end{bmatrix}, \quad f_2 = \begin{bmatrix}-\sigma_1\sigma_2\\ 0\end{bmatrix}, \quad f_3 = \sigma_1\begin{bmatrix}x(\gamma)\\ 1\end{bmatrix}, \quad f_4 = -\sigma_1\begin{bmatrix}x(\gamma)\\ 1\end{bmatrix}$$

for the pairs $\{(1,1),(-1,-1)\}$, $\{(1,-1),(-1,1)\}$, $\{(1,1),(1,-1)\}$ and $\{(-1,1),(-1,-1)\}$.

R follows the unconditional mean dynamics because $\bar{M}$ depends only on W_2, which takes each value with probability $1/2$ within every pair.

Each conditional ODE is a candidate escape route.

Every shock realization has probability $1/4$, so the probability of a run falls geometrically with its length, and the candidate that reaches the boundary first gives the most likely exit point (CWS, section 7.2).

We first compare the candidates' speeds at the self-confirming equilibrium, and then run the race.

```
def cond_ode(t, z, model, f):
    "Conditional mean dynamics for one shock pair with forcing f(y)."
    y, R = z[:2], z[2:].reshape(2, 2)
    return np.concatenate([np.linalg.solve(R, model.g_bar(y) + f(y, model)),
                           (model.M(y) - R).ravel()])

pairs = {
    "1: {(1,1),(-1,-1)}":  lambda y, m: np.array([m.σ1 * m.σ2, 0.0]),
    "2: {(1,-1),(-1,1)}":  lambda y, m: np.array([-m.σ1 * m.σ2, 0.0]),
    "3: {(1,1),(1,-1)}":   lambda y, m: m.σ1 * np.array([m.x(y), 1.0]),
    "4: {(-1,1),(-1,-1)}": lambda y, m: -m.σ1 * np.array([m.x(y), 1.0]),
}

race = {}
print(f"{'pair':22s} {'speed at SCE':>12s} {'exit time':>10s} {'exit y':>17s} {'exit x
↪':>7s}")
for name, f in pairs.items():
    speed = np.linalg.norm(np.linalg.solve(R_bar, f(y_sce, model)))
    s = solve_ivp(cond_ode, [0, 50], z0, args=(model, f), events=left_circle,
                  max_step=0.02, rtol=1e-9, atol=1e-11)
    race[name] = s
    y_end = s.y[:2, -1]
    if s.status == 1:
        print(f"{name:22s} {speed:12.2f} {s.t[-1]:10.3f}   ({y_end[0]:+.3f}, {y_
↪end[1]:5.2f}) {model.x(y_end):7.2f}")
    else:
        print(f"{name:22s} {speed:12.2f} {'never':>10s}   (rests at ({y_end[0]:+.2f},
↪{y_end[1]:.2f}))")
```

```
pair                   speed at SCE  exit time            exit y  exit x
1: {(1,1),(-1,-1)}             5.10      0.344   (-0.023,  5.10)    0.12
2: {(1,-1),(-1,1)}             5.10      8.633   (-1.468, 14.98)    6.97
3: {(1,1),(1,-1)}              0.30      never   (rests at (-1.00, 10.60))
4: {(-1,1),(-1,-1)}            0.30      0.473   (-0.043,  5.09)    0.22
```

At the self-confirming equilibrium itself, pairs 1 and 2 tie: their initial velocities are equal and opposite, and much faster than those of pairs 3 and 4.

Speed at the starting point therefore cannot settle the race, so we integrate the four ODEs.

Pair 1 reaches the boundary first, near the Ramsey belief.

Pair 4 also exits near the Ramsey belief, but later.

Pair 2, the mirror image of pair 1, heads toward *higher* inflation and needs far longer to reach the boundary.

Pair 3 never exits: it settles at a rest point inside the circle.

So pair 1 wins the race, and the escape forcing it induces is the $R^{-1}(\sigma_1\sigma_2, 0)^\top$ of (104.8).

CWS describe the winner's velocity as substantially larger than the other three; the table shows that this holds along the route, not at its start.

104.8 Mean dynamics reinforce the escape

Why does the Ramsey-ward escape succeed so quickly while its mirror image crawls?

Near the self-confirming equilibrium the mean dynamics point *toward* it, opposing any escape.

But CWS (section 6.4, Figures 8–9) show that once beliefs have moved a little way along the escape route, the mean dynamics themselves point *toward* Ramsey, reinforcing the escape.

The reason is that R lags $\bar{M}(\gamma)$.

During an escape the beliefs γ move fast, while R adjusts toward $\bar{M}(\gamma)$ only at rate one, so R still holds roughly the moments of the Nash era, $\bar{R} = \bar{M}(\bar{\gamma})$.

At the self-confirming equilibrium inflation barely varies, so $\det \bar{R} = \sigma_2^2$ is small and $\bar{R}^{-1}$ is large.

Premultiplying the drift $\bar{g}(\gamma) = \bar{M}(\gamma)(T(\gamma) - \gamma)$ by $\bar{R}^{-1}$ instead of $\bar{M}(\gamma)^{-1}$ both rotates and amplifies it.

CWS (section 6.5) describe the same self-reinforcement in words, adding that "overweighting recent observations helps this process along."

To see this, we plot the normalized mean-dynamics drift $R^{-1}\bar{g}(\gamma)$ in belief space, once with R held at its Nash-era value $\bar{R}$ and once with R caught up at $\bar{M}(\gamma)$.

```
gs = np.linspace(-1.2, 0.1, 16)       # slope
gi = np.linspace(4.5, 10.5, 16)       # intercept
GS, GI = np.meshgrid(gs, gi)

fig, axes = plt.subplots(1, 2, figsize=(13, 5.5), sharey=True)
titles = [r'$R = \bar M(\bar\gamma)$: $R$ lags', r'$R = \bar M(\gamma)$: $R$ caught up
 ↪']
for ax, lagging, title in zip(axes, [True, False], titles):
    DS, DI = np.zeros_like(GS), np.zeros_like(GI)
    for i in range(GS.shape[0]):
        for j in range(GS.shape[1]):
            y = np.array([GS[i, j], GI[i, j]])
            R = R_bar if lagging else model.M(y)
            d = np.linalg.solve(R, model.g_bar(y))
            DS[i, j], DI[i, j] = d / np.linalg.norm(d)
    ax.quiver(GS, GI, DS, DI, angles='xy', alpha=0.7)
    ax.plot(*y_sce, 'ko', ms=8, label='SCE (Nash)')
    ax.plot(0, model.u, 'C2s', ms=8, label='Ramsey belief')
    ax.plot(slope, intercept, 'C3', lw=2, label='escape path')
    ax.set_xlabel(r'slope $\gamma_1$')
    ax.set_title(title)
axes[0].set_ylabel(r'intercept $\gamma_{-1}$')
axes[0].legend(loc='lower left')
plt.tight_layout()
plt.show()
```

In the left panel, with R at its Nash-era value, the arrows along the escape route sweep toward the Ramsey belief.

In the right panel, where R has caught up, they point back toward slope $-\theta$ everywhere.

There the drift is $T(\gamma) - \gamma$, whose slope component $-\theta - \gamma_1$ is negative whenever $\gamma_1 > -\theta$.

So the reinforcement is a property of the *joint* dynamics of γ and R; it is the same lag that produced the dip in Fig. 104.2.

Along the escape path itself R moves only part of the way toward $\bar{M}(\gamma)$, because the escape takes only a fraction of a unit of time; Exercise 104.13.2 measures the alignment with the path's own R.

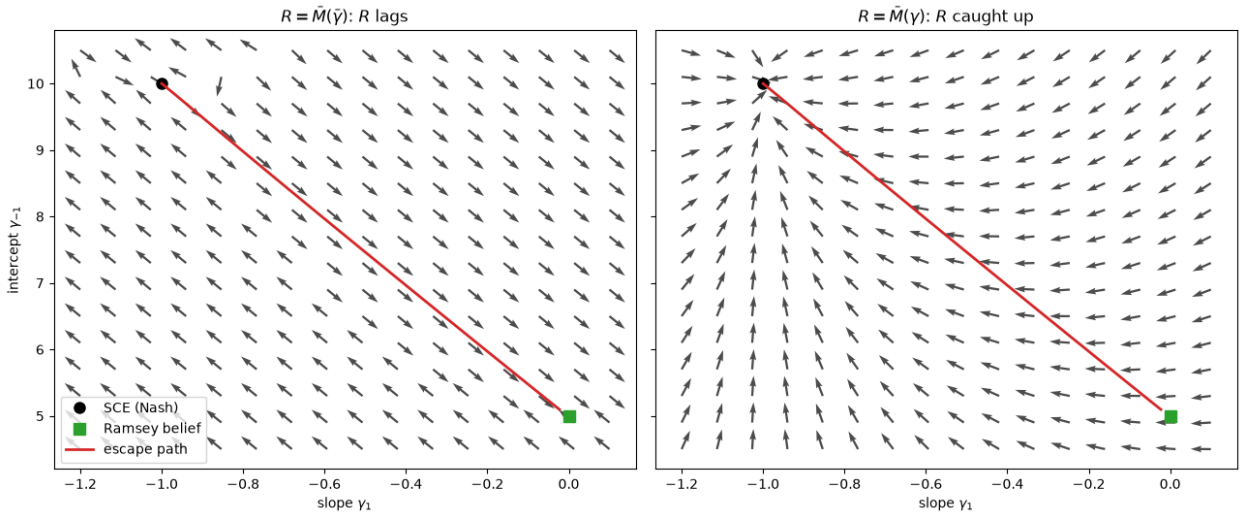


Fig. 104.4: Mean-dynamics drift with R lagging (left) and with R caught up (right), and the escape path

The escape dynamics only need to *start* the departure; the mean dynamics finish the job.

This is the sense in which the mean dynamics trace a *circuitous* route: pushed off the equilibrium along the escape path, the system travels near Ramsey before R catches up, the residual short-run Phillips curve is rediscovered, and the mean dynamics eventually carry beliefs back to Nash.

104.9 The experimentation trap

The escape has a compelling behavioural interpretation.

Within its approximating model, the government can only detect the natural-rate hypothesis if there is enough *dispersion* in inflation.

But inside a self-confirming equilibrium the government sets a constant systematic inflation rate, so it generates no such dispersion: it is caught in an **experimentation trap**.

Only an unusual run of shocks makes the government vary inflation enough to steepen its estimated Phillips curve; a steeper perceived curve leads it (through the best response) to cut inflation, which generates further influential observations that steepen the curve further.

As CWS note, overweighting recent observations helps this process along.

This self-reinforcing process halts when the perceived Phillips curve is vertical, and inflation is near Ramsey.

The system cannot stay there forever: in truth there *is* a short-run Phillips curve, which the government eventually rediscovers, rekindling the mean dynamics that carry it back to Nash.

104.10 Simulating escapes

The large-deviation results describe the limit $\varepsilon \to 0$.

To see what they imply at the gains this suite actually uses, we simulate the static model under the constant-gain recursion (104.4).

Every path starts at the self-confirming equilibrium, with $R_0 = \bar{M}(\bar{\gamma})$.

```
def simulate_path(model, ε, n_periods, seed=0):
    "One path of the static model under constant-gain RLS (en_rls), started at the␣
↪SCE."
    rng = np.random.default_rng(seed)
    θ, u, σ1, σ2 = model.θ, model.u, model.σ1, model.σ2
    γ = model.γ_sce()
    R = model.M(γ)
    x_path, γ_path = np.empty(n_periods), np.empty((n_periods, 2))
    for n in range(n_periods):
        x = model.x(γ)
        w1, w2 = rng.standard_normal(2)
        π = x + σ2 * w2
        U = u - θ * σ2 * w2 + σ1 * w1
        Φ = np.array([π, 1.0])
        R = R + ε * (np.outer(Φ, Φ) - R)
        γ = γ + ε * np.linalg.solve(R, Φ * (U - Φ @ γ))
        x_path[n], γ_path[n] = x, γ
    return x_path, γ_path
```

We use $\varepsilon = 0.02$, the gain of *A government learning the Phillips curve*; *Adaptive Learning and Escape Dynamics* uses 0.025, the book's $\lambda = .975$.

```
ε = 0.02
x_path, γ_path = simulate_path(model, ε, 4_000, seed=1)

fig, axes = plt.subplots(2, 1, figsize=(10, 6), sharex=True)
axes[0].plot(x_path, lw=0.8)
axes[0].axhline(model.θ * model.u, color='k', ls='--', lw=1, label='Nash')
axes[0].axhline(0, color='C2', ls=':', lw=1, label='Ramsey')
axes[0].set_ylabel('target inflation $x_n$')
axes[0].legend(loc='upper right')
axes[1].plot(γ_path[:, 0], lw=0.8, color='C1')
axes[1].axhline(-model.θ, color='k', ls='--', lw=1)
axes[1].set_ylabel(r'slope $\gamma_{1n}$')
axes[1].set_xlabel('period $n$')
plt.tight_layout()
plt.show()
```

Like the long simulation in *A government learning the Phillips curve*, the path is a recurrent sawtooth: abrupt escapes toward zero inflation, each followed by a slow climb back toward Nash, often interrupted by the next escape.

A longer run shows how little time the government spends near Nash.

```
x_long, _ = simulate_path(model, ε, 20_000, seed=2)
print(f"ε = {ε}, 20,000 periods:")
print(f"  share of periods with x > 4: {np.mean(x_long > 4):.2f}")
print(f"  share of periods with x < 2: {np.mean(x_long < 2):.2f}")
```

(continues on next page)

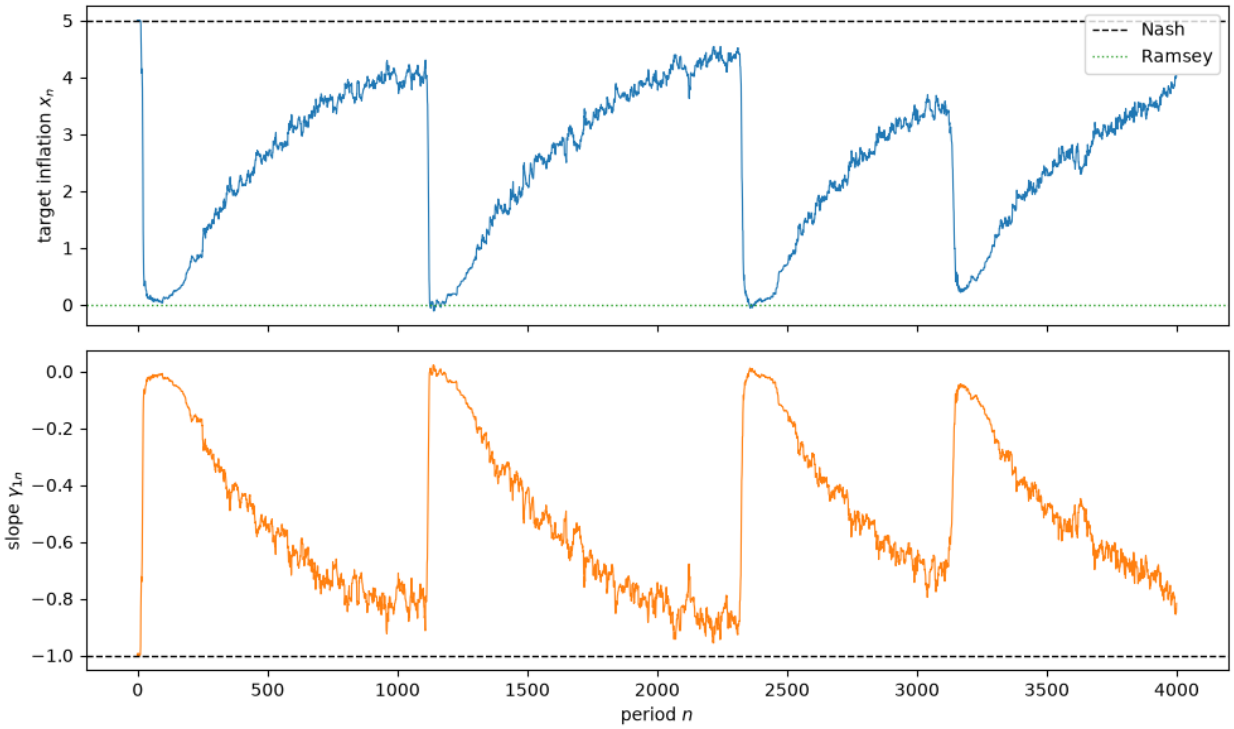


Fig. 104.5: A simulated path of the static model with gain 0.02

(continued from previous page)

```
print(f"  mean target inflation:       {x_long.mean():.2f}   (Nash {model.θ * model.
↪u:.0f})")
```

```
ε = 0.02, 20,000 periods:
  share of periods with x > 4: 0.10
  share of periods with x < 2: 0.34
  mean target inflation:       2.47   (Nash 5)
```

Target inflation is above 4 in only a small minority of periods.

To connect with the theorem, we measure the first exit time from the radius-5 circle, starting at the self-confirming equilibrium, for many independent paths.

```
def first_exits(model, ε, n_paths=400, radius=5.0, seed=0):
    """
    Simulate n_paths independent copies of (en_rls) from the SCE.  Return each
    path's first exit period from the circle of the given radius and its beliefs
    at exit (NaN if it has not exited after 100/ε periods).
    """
    rng = np.random.default_rng(seed)
    θ, u, σ1, σ2 = model.θ, model.u, model.σ1, model.σ2
    γ_bar = model.γ_sce()
    γ = np.tile(γ_bar, (n_paths, 1))
    R = np.tile(model.M(γ_bar), (n_paths, 1, 1))
    τ = np.full(n_paths, np.nan)
    γ_exit = np.full((n_paths, 2), np.nan)
    alive = np.ones(n_paths, dtype=bool)
```

(continues on next page)

(continued from previous page)

```
    n = 0
    while alive.any() and n < 100 / ε:
        n += 1
        x = -y[:, 1] * y[:, 0] / (1 + y[:, 0]**2)
        w1, w2 = rng.standard_normal((2, n_paths))
        π = x + σ2 * w2
        U = u - θ * σ2 * w2 + σ1 * w1
        Φ = np.stack([π, np.ones(n_paths)], axis=1)
        R += ε * (Φ[:, :, None] * Φ[:, None, :] - R)
        e = U - (Φ * y).sum(axis=1)
        y += ε * np.linalg.solve(R, (Φ * e[:, None])[:, :, None])[:, :, 0]
        out = alive & (np.linalg.norm(y - y_bar, axis=1) > radius)
        τ[out], y_exit[out] = n, y[out]
        alive &= ~out
    return τ, y_exit


S_bar = 4.987e-4        # CWS Table 1: static model, normal shocks, radius 5

print(f"{'ε':>6s} {'mean':>7s} {'median':>7s} {'ε·mean':>7s} {'exp(S̄/ε)':>9s} "
      f"{'x at exit < 1':>13s} {'not exited':>10s}")
for ε in [0.1, 0.05, 0.03, 0.02, 0.01, 0.005]:
    τ, y_exit = first_exits(model, ε)
    done = ~np.isnan(τ)
    x_exit = -y_exit[done, 1] * y_exit[done, 0] / (1 + y_exit[done, 0]**2)
    print(f"{ε:6.3f} {np.mean(τ[done]):7.1f} {np.median(τ[done]):7.0f} "
          f"{ε * np.mean(τ[done]):7.2f} {np.exp(S_bar / ε):9.3f} "
          f"{np.mean(x_exit < 1):13.2f} {np.sum(~done):10d}")
print(f"\nexp(S̄/ε) at ε = 0.0002: {np.exp(S_bar / 0.0002):.1f}")
```

```
     ε    mean  median  ε·mean exp(S̄/ε) x at exit < 1 not exited
 0.100    15.8      11    1.58     1.005          1.00          0
 0.050    27.6      22    1.38     1.010          1.00          0
 0.030    46.1      36    1.38     1.017          1.00          0
 0.020    77.1      59    1.54     1.025          1.00          0
 0.010   173.6     135    1.74     1.051          1.00          0
 0.005   405.5     329    2.03     1.105          1.00          0

exp(S̄/ε) at ε = 0.0002: 12.1
```

Mean first exit times run from about 16 periods at $\varepsilon = 0.1$ to about 400 at $\varepsilon = 0.005$.

At $\varepsilon = 0.01$ and 0.005 they are close to the 165.7 and 409.6 periods that Sargent and Williams [2005] (Table 1) report for constant-gain least squares.

Every path exits on the Ramsey side, near the exit point of the escape path.

The product $\varepsilon\times$ mean stays between about 1.4 and 2, so mean exit times scale roughly like $1/\varepsilon$.

The factor $e^{\bar{S}/\varepsilon}$, by contrast, never exceeds 1.11 in the table.

At these gains the large-deviation factor plays almost no role; exit times are set by the constant S_0, that is, by how long beliefs take to wander off the equilibrium and travel the route.

At the gains this suite uses, escapes are frequent, not rare.

Starting from the self-confirming equilibrium, a government with gain ε escapes on average after roughly $1.5/\varepsilon$ to $2/\varepsilon$ periods, and between escapes it seldom gets all the way back to Nash.

Rarity is a small-gain statement.

Only for gains of the order of $\bar{S}$ itself does the exponential factor bite: at $\varepsilon = 0.0002$, five times smaller than the smallest gain in Table 1 of Sargent and Williams [2005], $e^{\bar{S}/\varepsilon} \approx 12$.

The simulated exit times also bear out CWS's footnote 14.

Measured in units of $t = n\varepsilon$, they are several times the 0.34 that the changed-measure path takes; Exercise 104.13.5 compares the two in detail.

104.11 Escape frequency and model richness

The minimized action $\bar{S}$ sets the exponential rate at which escape times grow as the gain shrinks.

A striking finding of CWS is that escapes are *more frequent* when the government's model is richer.

CWS's Table 1 gives the dynamic model of *Adaptive Learning and Escape Dynamics*, with lagged unemployment and inflation, a rate $\bar{S} \approx 9.7 \times 10^{-6}$, about fifty times smaller than the static model's 5.0×10^{-4}, even though the two share the same self-confirming equilibrium.

That dynamic-model number needs a caveat: it was computed with the 2001 shortcut that Williams [2019] corrected, which matters for dynamic models.

The comparison does not rest on that number alone, though.

In CWS's simulations (their Figure 5), mean escape times are shorter in the dynamic model at every gain, and Exercise 104.13.6 finds the same for median first-escape times.

A richer model lets the government detect the subtler distributed-lag ("induction-hypothesis") version of the natural-rate hypothesis, so it escapes toward Ramsey more readily.

Note

The escape dynamics inherit the same "near determinism" that makes the mean dynamics useful: for small gains, simulated escapes cluster around the deterministic escape route, in the dynamic model (CWS, Figure 6) and in the static model (Exercise 104.13.5; CWS, Figure 7).

The next lecture, *Priors, Escapes, and Learning Cycles*, shows that the government's *prior* about how its coefficients drift reshapes both dynamics, and can even make the escape a deterministic *cycle*.

104.12 Escaping volatile inflation

The escape delivers a fall in the *level* of inflation.

Ellison and Yates [2007] extend the model to explain a second post-war fact: inflation *volatility* rose and fell together with the level.

Their device is to give the government a reason to stabilize.

Following Phelps and Taylor [1977], they add an unemployment shock W_3 that the government, but not the price-setting private sector, can react to.

Now a government that believes in an exploitable Phillips curve is tempted to *lean against* W_3 by varying inflation, so the perceived effectiveness of policy, $|\gamma_1|$, drives the volatility of inflation as well as its level.

In their model the expected inflation volatility a private agent faces is

$$\mathbb{E}(\sigma_\pi \mid \gamma) = \left[\sigma_2^2 + \left(\frac{\gamma_1}{1+\gamma_1^2}\right)^2 \sigma_3^2\right]^{1/2}. \tag{104.11}$$

At the self-confirming equilibrium $\gamma_1 = -\theta$, the government believes policy is effective and leans against W_3 aggressively, so inflation is *volatile*.

Along an escape, $\gamma_1 \to 0$: the government stops believing it can exploit the Phillips curve, abandons stabilization, and the volatility term collapses to the control error σ_2.

As an approximation, we apply (104.11) to the belief path computed above for $\sigma_3 = 0$.

In Ellison and Yates's model a positive σ_3 also changes the escape itself, as we discuss below, but the approximation shows the level and volatility of inflation escaping *in tandem*.

```
σ3 = 0.9                                    # size of the stabilizable shock
infl_vol = np.sqrt(model.σ2**2 + (slope / (1 + slope**2))**2 * σ3**2)

fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))
axes[0].plot(esc.t, infl)
axes[0].axhline(model.θ * model.u, color='k', ls='--', lw=1, label='Nash')
axes[0].axhline(0, color='C2', ls=':', lw=1, label='Ramsey')
axes[0].set_xlabel('time $t = n\\varepsilon$'); axes[0].set_ylabel('inflation level')
axes[0].set_title('level escapes'); axes[0].legend()

axes[1].plot(esc.t, infl_vol, 'C1')
axes[1].axhline(model.σ2, color='k', ls='--', lw=1, label=r'control error $\sigma_2$')
axes[1].set_xlabel('time $t = n\\varepsilon$'); axes[1].set_ylabel('inflation␣
↪volatility')
axes[1].set_title('volatility escapes too'); axes[1].legend()

plt.tight_layout()
plt.show()
```

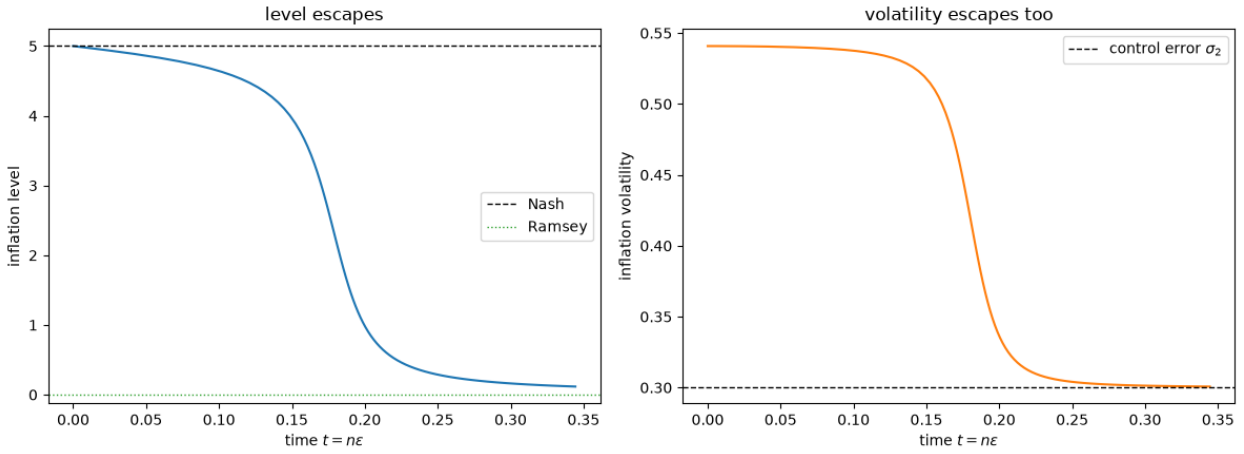


Fig. 104.6: The level and volatility of inflation escape in tandem

Both fall as beliefs escape, because both spring from the same source: the government's belief in an exploitable tradeoff.

Ellison and Yates [2007] draw a further, subtler lesson about the *timing* of escapes.

A larger stabilizable shock σ_3 makes an escape *harder to trigger*: to create the illusion that inflation moves while unemployment stays put, an unusual sequence of shocks must now offset not only the control errors but also the government's own stabilizing reaction to W_3.

The more shocks the government can offset, the more complex the escape-triggering sequence, and the longer the wait.

Ellison and Yates conclude that an escape to low inflation is more likely in a period with relatively few shocks that policy can offset.

It is tempting, and it is our extrapolation rather than theirs, to read the calm of the mid-1980s as such a period, and to link it to the disinflation.

The timing does not fit easily.

In the estimates of *Drifts and Volatilities*, core inflation fell in 1980–82, while most of the decline in the variance of shocks came in the mid-1980s.

The disinflation came first, during a period of unusually large shocks.

That lecture builds a statistical model with room for both *shrinking shocks* and *shifting dynamics* and lets the data apportion them.

Before turning to data, *Priors, Escapes, and Learning Cycles* returns to the static model and asks how the government's prior about drifting coefficients changes everything above.

104.13 Exercises

Exercise 104.13.1

At the self-confirming equilibrium the moment matrix is $\bar{R} = \bar{M}(\bar{\gamma})$, with $\det \bar{R} = \sigma_2^2$.

1. Show that the escape velocity at the self-confirming equilibrium is $\bar{R}^{-1}(\sigma_1\sigma_2, 0)^\top = (\sigma_1/\sigma_2)(1, -\bar{x})^\top$, where $\bar{x} = \theta u$. So with $\sigma_1 = \sigma_2$ it does not depend on the size of the shocks.
2. Integrate the escape ODE (104.9) for $\sigma_1 = \sigma_2 = \sigma \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$ and report the exit time from the radius-5 circle.
3. Exit times nonetheless rise with σ. Explain why. Hint: write $\bar{M}(\gamma) = \bar{R} + \Delta(\gamma)$ and look at how Δ and $\bar{R}^{-1}$ depend on σ_2.
4. Now vary one shock at a time: $\sigma_2 \in \{0.2, 0.3, 0.45\}$ with $\sigma_1 = 0.3$, and $\sigma_1 \in \{0.2, 0.3, 0.45\}$ with $\sigma_2 = 0.3$.

Solution

For part 1, $\bar{R} = \begin{bmatrix} \bar{x}^2 + \sigma_2^2 & \bar{x} \\ \bar{x} & 1 \end{bmatrix}$ has determinant σ_2^2 and inverse $\sigma_2^{-2} \begin{bmatrix} 1 & -\bar{x} \\ -\bar{x} & \bar{x}^2 + \sigma_2^2 \end{bmatrix}$, so $\bar{R}^{-1}(\sigma_1\sigma_2, 0)^\top = (\sigma_1/\sigma_2)(1, -\bar{x})^\top$.

For parts 2 and 4, we integrate the escape ODE for each pair of shock sizes.

```
def exit_time(σ1, σ2):
    m = EscapeModel(θ=1.0, u=5.0, σ1=σ1, σ2=σ2)
    y0 = m.y_sce()

    def leave(t, z):
        return 5.0 - np.linalg.norm(z[:2] - y0)
    leave.terminal = True
```

```
    leave.direction = -1

    v0 = np.linalg.solve(m.M(y0), np.array([σ1 * σ2, 0.0]))
    s = solve_ivp(lambda t, z: escape_ode(t, z, m), [0, 500],
                  np.concatenate([y0, m.M(y0).ravel()]),
                  events=leave, max_step=0.02, rtol=1e-9, atol=1e-11)
    return v0, s.t[-1]

for σ in [0.1, 0.2, 0.3, 0.4, 0.5]:
    v0, T = exit_time(σ, σ)
    print(f"σ1 = σ2 = {σ}: velocity at SCE = {v0.round(3)}, exit time = {T:.3f}")
print()
for σ2 in [0.2, 0.3, 0.45]:
    print(f"σ1 = 0.3, σ2 = {σ2}: exit time = {exit_time(0.3, σ2)[1]:.3f}")
for σ1 in [0.2, 0.3, 0.45]:
    print(f"σ1 = {σ1}, σ2 = 0.3: exit time = {exit_time(σ1, 0.3)[1]:.3f}")
```

```
σ1 = σ2 = 0.1: velocity at SCE = [ 1. -5.], exit time = 0.083
σ1 = σ2 = 0.2: velocity at SCE = [ 1. -5.], exit time = 0.201
σ1 = σ2 = 0.3: velocity at SCE = [ 1. -5.], exit time = 0.344
σ1 = σ2 = 0.4: velocity at SCE = [ 1. -5.], exit time = 0.503
σ1 = σ2 = 0.5: velocity at SCE = [ 1. -5.], exit time = 0.671

σ1 = 0.3, σ2 = 0.2: exit time = 0.171
σ1 = 0.3, σ2 = 0.3: exit time = 0.344
σ1 = 0.3, σ2 = 0.45: exit time = 0.740
σ1 = 0.2, σ2 = 0.3: exit time = 0.411
σ1 = 0.3, σ2 = 0.3: exit time = 0.344
σ1 = 0.45, σ2 = 0.3: exit time = 0.285
```

With $\sigma_1 = \sigma_2$ the starting velocity is the same for every σ, yet the exit time rises steadily with σ.

The speed comes from the R-lag amplification of *Mean dynamics reinforce the escape*.

During the escape R stays near $\bar{R}$, so the mean drift is $\bar{R}^{-1}\bar{M}(\gamma)(T(\gamma)-\gamma) = (T(\gamma)-\gamma) + \bar{R}^{-1}\Delta(\gamma)(T(\gamma)-\gamma)$.

Here $\Delta(\gamma) = \begin{bmatrix} x^2 - \bar{x}^2 & x - \bar{x} \\ x - \bar{x} & 0 \end{bmatrix}$ depends only on how far inflation has moved, not on σ_2, while $\bar{R}^{-1}$ is of order $1/\sigma_2^2$.

So the smaller the control error, the more strongly the lagging moment matrix amplifies any move away from Nash, and the faster beliefs travel the route.

That is why σ_2 matters so much in part 4, while a larger σ_1 speeds the escape only through the initial push σ_1/σ_2.

None of this concerns the *frequency* of escapes, which *Simulating escapes* measures; a noisier economy travels the deterministic route more slowly once an escape is under way.

Exercise 104.13.2

Make the reinforcement in the vector-field plot quantitative.

At several points *along* the escape path, evaluate the mean-dynamics drift $R^{-1}\bar{g}(\gamma)$, with the path's own R, and measure its cosine alignment with the direction from the current beliefs toward the Ramsey belief $(0, u)$.

A cosine near $+1$ means the mean dynamics are pushing beliefs *toward* Ramsey, reinforcing the escape.

Solution

We evaluate the drift at five points along the escape path.

```
y_ramsey = np.array([0.0, model.u])            # Belief 2

def cosine_toward_ramsey(y, R):
    drift = np.linalg.solve(R, model.g_bar(y))
    to_ramsey = y_ramsey - y
    denom = np.linalg.norm(drift) * np.linalg.norm(to_ramsey)
    return np.nan if denom < 1e-9 else drift @ to_ramsey / denom

for frac in [0.0, 0.25, 0.5, 0.75, 0.95]:
    k = min(int(frac * len(esc.t)), len(esc.t) - 1)
    y, R = esc.y[:2, k], esc.y[2:, k].reshape(2, 2)
    print(f"frac {frac:.2f}:  y = ({y[0]:+.2f}, {y[1]:.1f}),  "
          f"cos(drift, →Ramsey) = {cosine_toward_ramsey(y, R):+.2f},  "
          f"with R = M(y): {cosine_toward_ramsey(y, model.M(y)):+.2f}")
```

```
frac 0.00:  y = (-1.00, 10.0),  cos(drift, →Ramsey) = +nan,  with R = M(y): +nan
frac 0.25:  y = (-0.65, 8.2),  cos(drift, →Ramsey) = +1.00,  with R = M(y): -0.
↪93
frac 0.50:  y = (-0.23, 6.1),  cos(drift, →Ramsey) = +1.00,  with R = M(y): -0.
↪47
frac 0.75:  y = (-0.07, 5.3),  cos(drift, →Ramsey) = +1.00,  with R = M(y): -0.
↪25
frac 0.95:  y = (-0.03, 5.1),  cos(drift, →Ramsey) = +1.00,  with R = M(y): -0.
↪25
```

Right at the self-confirming equilibrium the drift vanishes (the cosine is undefined), so the mean dynamics neither help nor hinder.

But once beliefs have moved even slightly along the escape route, the mean-dynamics drift points almost exactly toward the Ramsey belief (cosine $\approx +1$): the mean dynamics reinforce the escape all the way to Ramsey.

The last column replaces the path's R by $\bar{M}(\gamma)$, and the cosine turns negative: the reinforcement comes from R lagging behind.

The opposition that CWS emphasize is confined to a *tiny* neighbourhood of the equilibrium; the escape dynamics only need to nudge beliefs out of it, after which the mean dynamics finish the job.

Exercise 104.13.3

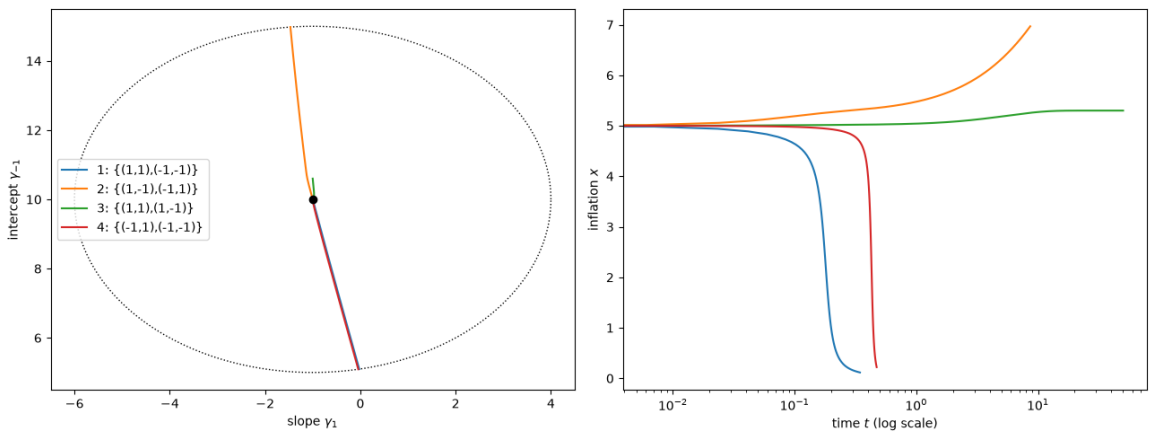
Look more closely at the race of four ODEs.

1. Plot the four conditional paths `race` in the (γ_1, γ_{-1}) plane, together with the radius-5 circle, and plot inflation along each.
2. CWS give the rest points of the conditional ODEs for pairs 3 and 4 as $(-\theta,\ (u \pm \sigma_1)(1+\theta^2))$, both well inside the circle. Pair 3 settles there. Why does pair 4 escape instead?
3. Rerun the race with R held equal to $\bar{M}(\gamma)$, as if the moment matrix adjusted instantly. Which pair wins now, and where does it exit?

Solution

We first plot the four paths computed in the text.

```
fig, axes = plt.subplots(1, 2, figsize=(13, 5))
ψ = np.linspace(0, 2 * np.pi, 200)
axes[0].plot(y_sce[0] + 5 * np.cos(ψ), y_sce[1] + 5 * np.sin(ψ), 'k:', lw=1)
for name, s in race.items():
    axes[0].plot(s.y[0], s.y[1], label=name)
    axes[1].plot(s.t, -s.y[1] * s.y[0] / (1 + s.y[0]**2), label=name)
axes[0].plot(*y_sce, 'ko')
axes[0].set_xlabel(r'slope $\gamma_1$')
axes[0].set_ylabel(r'intercept $\gamma_{-1}$')
axes[1].set_xscale('log')
axes[1].set_xlabel('time $t$ (log scale)')
axes[1].set_ylabel('inflation $x$')
axes[0].legend()
plt.tight_layout()
plt.show()
```



Next we rerun the race with $R = \bar{M}(\gamma)$ imposed, which reduces each conditional ODE to $\dot{\gamma} = T(\gamma) - \gamma + \bar{M}(\gamma)^{-1} f_j(\gamma)$.

```
def cond_ode_no_lag(t, y, model, f):
    return np.linalg.solve(model.M(y), model.g_bar(y) + f(y, model))

def left_circle_y(t, y, *args):
    return 5.0 - np.linalg.norm(y - y_sce)
left_circle_y.terminal = True
left_circle_y.direction = -1

for name, f in pairs.items():
    s = solve_ivp(cond_ode_no_lag, [0, 50], y_sce, args=(model, f), events=left_
 ↪circle_y,
                  max_step=0.02, rtol=1e-9, atol=1e-11)
    y_end = s.y[:, -1]
    status = f"exits at t = {s.t[-1]:.3f}" if s.status == 1 else "never exits"
    print(f"{name:22s} {status:22s} y_end = ({y_end[0]:+.3f}, {y_end[1]:.2f}), "
          f"x = {model.x(y_end):.2f}")
```

```
1: {(1,1),(-1,-1)}     exits at t = 3.941     y_end = (-0.019, 5.10), x = 0.10
2: {(1,-1),(-1,1)}     exits at t = 1.194     y_end = (-1.697, 14.95), x = 6.54
3: {(1,1),(1,-1)}      never exits            y_end = (-1.000, 10.60), x = 5.30
4: {(-1,1),(-1,-1)}    never exits            y_end = (-1.000, 9.40), x = 4.70
```

With $R = \bar{M}(\gamma)$ the rest point of each conditional ODE is a fixed point of a shifted T-map.

For pair 4 it is $(-\theta, (u-\sigma_1)(1+\theta^2))$; the Jacobian of the drift there again has eigenvalues -1 and $-1/(1+\theta^2)$, so the path converges to it.

With the true, lagging R, the same forcing first lowers the intercept and inflation a little; the lagging moment matrix then amplifies the move into a full escape to Ramsey.

Without the lag, the race changes completely.

Pair 1 still heads for Ramsey but only drifts there slowly, while pair 2 now exits first, toward *higher* inflation.

So the R lag decides both the direction and the speed of the escape: it is what makes the Ramsey-ward route the most likely one.

Exercise 104.13.4

The theorem says $\varepsilon \log \mathbb{E}\,\tau^\varepsilon \to \bar{S}$, where $\tau^\varepsilon = n\varepsilon$ is the exit time on the ODE time scale; equivalently $\log \mathbb{E}\,\tau^\varepsilon \approx S_0 + \bar{S}/\varepsilon$.

1. Use `first_exits` to compute the mean exit period for $\varepsilon \in \{0.1, 0.05, 0.02, 0.01, 0.005, 0.0025, 0.001\}$.
2. Plot $\log(\varepsilon \times \text{mean exit period})$ against $1/\varepsilon$, and compare the slope with CWS's $\bar{S}$.
3. Compare the mean exit periods with Sargent and Williams [2005] (Table 1), who report 20.6, 165.7, 409.6, 2414.8 and 4486.7 periods at $\varepsilon = 0.05, 0.01, 0.005, 0.002, 0.001$.

Solution

We compute the mean exit period on a finer grid of small gains.

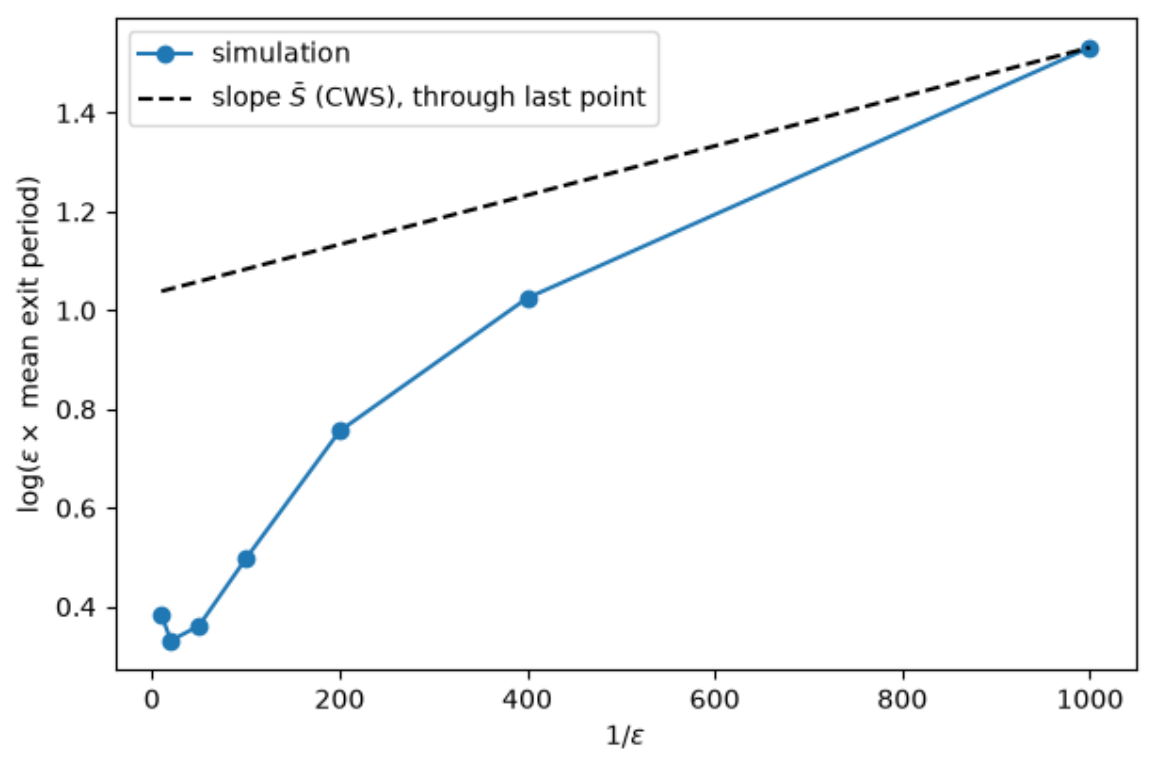
```
gains = np.array([0.1, 0.05, 0.02, 0.01, 0.005, 0.0025, 0.001])
mean_exit = np.array([np.nanmean(first_exits(model, ε, seed=1)[0]) for ε in gains])
log_scaled = np.log(gains * mean_exit)

for ε, m, l in zip(gains, mean_exit, log_scaled):
    print(f"ε = {ε:6.4f}: mean exit period = {m:7.0f},  ε·mean = {ε * m:5.2f}")
local_slope = np.diff(log_scaled) / np.diff(1 / gains)
print("\nlocal slopes of log(ε·mean) in 1/ε:", local_slope.round(5))
print(f"CWS S̄ = {S_bar:.5f}")

fig, ax = plt.subplots(figsize=(7, 4.5))
ax.plot(1 / gains, log_scaled, 'o-', label='simulation')
ax.plot(1 / gains, log_scaled[-1] + S_bar * (1 / gains - 1 / gains[-1]), 'k--',
        label=r'slope $\bar S$ (CWS), through last point')
ax.set_xlabel(r'$1/\varepsilon$')
ax.set_ylabel(r'$\log(\varepsilon \times$ mean exit period$)$')
ax.legend()
plt.show()
```

```
ε = 0.1000: mean exit period =      15,  ε·mean =  1.47
ε = 0.0500: mean exit period =      28,  ε·mean =  1.39
ε = 0.0200: mean exit period =      72,  ε·mean =  1.43
ε = 0.0100: mean exit period =     164,  ε·mean =  1.64
ε = 0.0050: mean exit period =     426,  ε·mean =  2.13
ε = 0.0025: mean exit period =    1114,  ε·mean =  2.79
ε = 0.0010: mean exit period =    4629,  ε·mean =  4.63

local slopes of log(ε·mean) in 1/ε: [-0.00525  0.001    0.00272  0.00259  0.
↪00135  0.00085]
CWS S̄ = 0.00050
```



The scaled mean $\varepsilon\times$ mean stays between about 1.4 and 4.6 over a hundredfold range of gains.

Over that range $\bar{S}/\varepsilon$ rises only from 0.005 to 0.5, so the level of $\log \mathbb{E}\,\tau^{\varepsilon}$ is dominated by the constant S_0.

At large gains the curve even slopes down, and only at the smallest gains does its slope approach the order of magnitude of $\bar{S}$: the asymptotic regime of the theorem lies at gains far smaller than those this suite uses.

The simulated means agree closely with SW's at $\varepsilon = 0.01$, 0.005 and 0.001.

Exercise 104.13.5

CWS's Figure 7 compares simulated escapes with the dominant and changed-measure paths.

1. Simulate 50 paths at $\varepsilon = 0.005$ from the self-confirming equilibrium, recording beliefs until each first leaves the radius-5 circle.
2. Plot the slope and intercept of each path against $\log n$, together with the changed-measure path at $n = t/\varepsilon$.
3. Compare the exit points, the total time to exit, and the duration of the final run, measured from the last period in which the slope is below -0.9.

Solution

We record the full belief history of each path until it exits.

```
def escape_histories(model, ε, n_paths=50, n_max=10_000, radius=5.0, seed=0):
    rng = np.random.default_rng(seed)
    θ, u, σ1, σ2 = model.θ, model.u, model.σ1, model.σ2
    y_bar = model.y_sce()
```

```
    γ = np.tile(γ_bar, (n_paths, 1))
    R = np.tile(model.M(γ_bar), (n_paths, 1, 1))
    hist = np.full((n_max + 1, n_paths, 2), np.nan)
    hist[0] = γ
    τ = np.zeros(n_paths, dtype=int)
    alive = np.ones(n_paths, dtype=bool)
    for n in range(1, n_max + 1):
        x = -γ[:, 1] * γ[:, 0] / (1 + γ[:, 0]**2)
        w1, w2 = rng.standard_normal((2, n_paths))
        π = x + σ2 * w2
        U = u - θ * σ2 * w2 + σ1 * w1
        Φ = np.stack([π, np.ones(n_paths)], axis=1)
        R += ε * (Φ[:, :, None] * Φ[:, None, :] - R)
        e = U - (Φ * γ).sum(axis=1)
        γ += ε * np.linalg.solve(R, (Φ * e[:, None])[:, :, None])[:, :, 0]
        hist[n, alive] = γ[alive]
        out = alive & (np.linalg.norm(γ - γ_bar, axis=1) > radius)
        τ[out] = n
        alive &= ~out
        if not alive.any():
            break
    return hist[:n + 1], τ

ε = 0.005
hist, τ = escape_histories(model, ε, seed=3)
print(f"all exited: {np.all(τ > 0)}")
print(f"simulated exit periods: min {τ.min()}, median {np.median(τ):.0f}, max {τ.
↪max()}")
print(f"changed-measure path:   exits at n = t/ε = {esc.t[-1] / ε:.0f}")

γ_end = hist[τ, np.arange(len(τ))]
print(f"exit slope: simulated mean {γ_end[:, 0].mean():+.3f}, "
      f"changed-measure {slope[-1]:+.3f}")

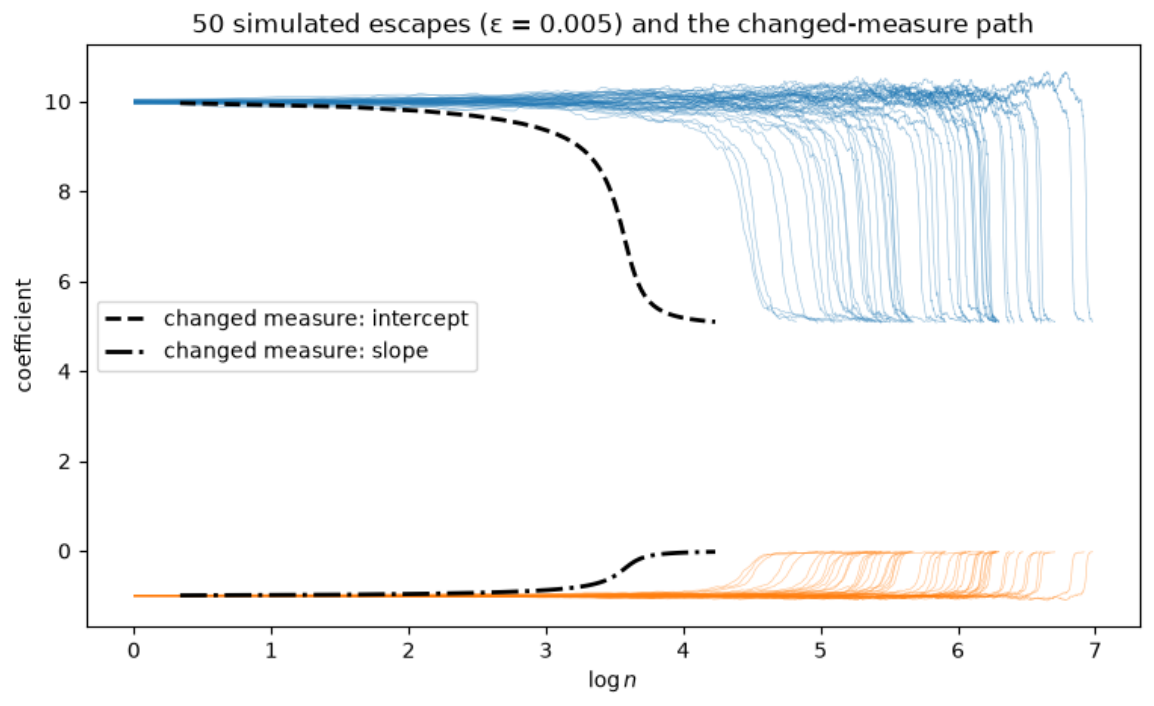
final_run = np.array([(τ[i] - np.where(hist[:τ[i] + 1, i, 0] < -0.9)[0].max()) * ε
                      for i in range(len(τ))])
k_cm = np.where(slope < -0.9)[0].max()
print(f"final run (in t = nε): simulated median {np.median(final_run):.2f}, "
      f"changed-measure {esc.t[-1] - esc.t[k_cm]:.2f}")
```

```
all exited: True
simulated exit periods: min 124, median 366, max 1072
changed-measure path:   exits at n = t/ε = 69
exit slope: simulated mean -0.024, changed-measure -0.023
final run (in t = nε): simulated median 0.35, changed-measure 0.26
```

Next we plot the paths on CWS's log time scale.

```
fig, ax = plt.subplots(figsize=(9, 5))
for i in range(len(τ)):
    n_i = np.arange(1, τ[i] + 1)
    ax.plot(np.log(n_i), hist[1:τ[i] + 1, i, 1], color='C0', lw=0.4, alpha=0.5)
    ax.plot(np.log(n_i), hist[1:τ[i] + 1, i, 0], color='C1', lw=0.4, alpha=0.5)
n_cm = esc.t[1:] / ε
ax.plot(np.log(n_cm), intercept[1:], 'k--', lw=2, label='changed measure: intercept
↪')
ax.plot(np.log(n_cm), slope[1:], 'k-.', lw=2, label='changed measure: slope')
ax.set_xlabel(r'$\log n$')
ax.set_ylabel('coefficient')
```

```
ax.set_title(f'50 simulated escapes (ε = {ε}) and the changed-measure path')
ax.legend()
plt.show()
```



The simulated escapes leave the circle at almost exactly the changed-measure path's exit point, as Theorem 5.3 predicts for the dominant path.

But the changed-measure path exits before even the fastest simulated escape, as CWS found.

Most of the difference is waiting time: a simulated path spends a long spell near the self-confirming equilibrium before its final run.

Once that run is under way, it takes a time of the same order as the changed-measure path.

Exercise 104.13.6

CWS find escapes more frequent in their dynamic model than in the static one.

Check this by simulation.

In the dynamic model of *Adaptive Learning and Escape Dynamics* (which writes y for inflation), the government regresses U_n on $(\pi_n, U_{n-1}, U_{n-2}, \pi_{n-1}, \pi_{n-2}, 1)$ and sets inflation by solving the discounted Phelps problem with $\delta = 0.98$.

1. Simulate the dynamic model with the recursion (104.4), now with six regressors, starting at its self-confirming equilibrium with R_0 equal to the self-confirming moment matrix.
2. Define the first escape as the first period in which the 20-period average of inflation falls below 1.
3. Report the median first escape over many paths for both models, for $\varepsilon \in \{0.002, 0.005, 0.01, 0.02\}$.

Solution

The government's decision rule uses the same Phelps solver as *Adaptive Learning and Escape Dynamics*.

The state is $(U_{n-1}, U_{n-2}, \pi_{n-1}, \pi_{n-2}, 1)$, and the cross term of the discounted Riccati equation enters unscaled.

```
def phelps_rule(γ, δ=0.98):
    "Rule ĥ_n = h·(U_{n-1}, U_{n-2}, π_{n-1}, π_{n-2}, 1) for beliefs γ."
    γ0, c = γ[0], γ[1:]
    A, B = np.zeros((5, 5)), np.zeros((5, 1))
    A[0, :], B[0, 0] = c, γ0            # U_n
    A[1, 0] = 1.0                        # U_{n-1}
    B[2, 0] = 1.0                        # π_n
    A[3, 2] = 1.0                        # π_{n-1}
    A[4, 4] = 1.0                        # constant
    Q_x, Q_u, N = np.outer(c, c), np.array([[γ0**2 + 1.0]]), (γ0 * c).reshape(1, -
↪1)
    Ad, Bd = np.sqrt(δ) * A, np.sqrt(δ) * B
    P = solve_discrete_are(Ad, Bd, Q_x, Q_u, s=N.T)
    F = np.linalg.solve(Q_u + Bd.T @ P @ Bd, Bd.T @ P @ Ad + N)
    return -F.ravel()


def dynamic_sce(model):
    "SCE beliefs and moments of (π_n, U_{n-1}, U_{n-2}, π_{n-1}, π_{n-2}, 1)."
    θ, u, σ1, σ2 = model.θ, model.u, model.σ1, model.σ2
    γ = np.array([-θ, 0, 0, 0, 0, (1 + θ**2) * u])
    mean = np.array([θ * u, u, u, θ * u, θ * u, 1.0])
    lag = [0, 1, 2, 1, 2, None]                    # date of each regressor
    var = {('π', 'π'): σ2**2, ('U', 'U'): θ**2 * σ2**2 + σ1**2,
           ('π', 'U'): -θ * σ2**2, ('U', 'π'): -θ * σ2**2}
    kind = ['π', 'U', 'U', 'π', 'π', 'c']
    M = np.outer(mean, mean)
    for i in range(6):
        for j in range(6):
            if lag[i] is not None and lag[i] == lag[j]:
                M[i, j] += var[(kind[i], kind[j])]
    return γ, M


def first_escape_dynamic(model, ε, seed, n_max=10_000, window=20):
    rng = np.random.default_rng(seed)
    θ, u, σ1, σ2 = model.θ, model.u, model.σ1, model.σ2
    γ, R = dynamic_sce(model)
    U1 = U2 = u
    π1 = π2 = θ * u
    recent = np.full(window, θ * u)
    for n in range(n_max):
        π_hat = phelps_rule(γ) @ np.array([U1, U2, π1, π2, 1.0])
        w1, w2 = rng.standard_normal(2)
        π = π_hat + σ2 * w2
        U = u - θ * σ2 * w2 + σ1 * w1
        Φ = np.array([π, U1, U2, π1, π2, 1.0])
        R = R + ε * (np.outer(Φ, Φ) - R)
        γ = γ + ε * np.linalg.solve(R, Φ * (U - Φ @ γ))
        U2, U1, π2, π1 = U1, U, π1, π
        recent[n % window] = π
        if n >= window and recent.mean() < 1:
            return n + 1
```

```
    return np.nan


def first_escape_static(model, ε, n_paths=200, n_max=20_000, window=20, seed=0):
    rng = np.random.default_rng(seed)
    θ, u, σ1, σ2 = model.θ, model.u, model.σ1, model.σ2
    y = np.tile(model.y_sce(), (n_paths, 1))
    R = np.tile(model.M(model.y_sce()), (n_paths, 1, 1))
    recent = np.full((n_paths, window), θ * u)
    τ = np.full(n_paths, np.nan)
    alive = np.ones(n_paths, dtype=bool)
    for n in range(n_max):
        x = -y[:, 1] * y[:, 0] / (1 + y[:, 0]**2)
        w1, w2 = rng.standard_normal((2, n_paths))
        π = x + σ2 * w2
        U = u - θ * σ2 * w2 + σ1 * w1
        Φ = np.stack([π, np.ones(n_paths)], axis=1)
        R += ε * (Φ[:, :, None] * Φ[:, None, :] - R)
        e = U - (Φ * y).sum(axis=1)
        y += ε * np.linalg.solve(R, (Φ * e[:, None])[:, :, None])[:, :, 0]
        recent[:, n % window] = π
        hit = alive & (n >= window) & (recent.mean(axis=1) < 1)
        τ[hit] = n + 1
        alive &= ~hit
        if not alive.any():
            break
    return τ


print("check: dynamic rule at the SCE =", phelps_rule(dynamic_sce(model)[0]).
↪round(4))
```

```
check: dynamic rule at the SCE = [-0. -0. -0. -0.  5.]
```

Now we compare median first escapes.

```
print(f"{'ε':>6s} {'static median':>14s} {'dynamic median':>15s}")
for ε in [0.002, 0.005, 0.01, 0.02]:
    τ_s = first_escape_static(model, ε)
    τ_d = np.array([first_escape_dynamic(model, ε, seed) for seed in range(24)])
    print(f"{ε:6.3f} {np.nanmedian(τ_s):14.0f} {np.nanmedian(τ_d):15.0f}"
          f"   (not escaped: {np.isnan(τ_s).sum()}/{len(τ_s)}, {np.isnan(τ_d).
↪sum()}/{len(τ_d)})")
```

```
     ε  static median  dynamic median
 0.002            948             200   (not escaped: 0/200, 0/24)
 0.005            286             102   (not escaped: 0/200, 0/24)
 0.010            121              52   (not escaped: 0/200, 0/24)
 0.020             62              32   (not escaped: 0/200, 0/24)
```

At every gain the dynamic model escapes sooner, and the gap widens as the gain falls, as CWS's smaller $\bar{S}$ for the dynamic model suggests.

The government with lags can find the induction hypothesis while keeping a short-run Phillips curve, so it needs a less unusual run of shocks to begin the escape.

The comparison holds even though the two models share the same self-confirming equilibrium.

CHAPTER

FIVE

PRIORS, ESCAPES, AND LEARNING CYCLES

Contents

- *Priors, Escapes, and Learning Cycles*
  - *Overview*
  - *The static model*
  - *The self-confirming equilibrium*
  - *Drifting beliefs and the Kalman filter*
  - *Mean dynamics and E-stability*
  - *Learning cycles*
  - *Escape dynamics and the direction of escape*
  - *Sims's nonconvergence*
  - *Conclusion*
  - *Exercises*

105.1 Overview

> The Bourbons remember everything and learn nothing.
>
> – Charles Maurice de Talleyrand

This lecture is a sequel to *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation*.

It follows Sargent and Williams [2005], which generalizes the adaptive model of chapter 8 of Sargent [1999] and of Cho *et al.* [2002] (CWS), whose escape dynamics we derived in *Escaping Nash Inflation*.

In *Adaptive Learning and Escape Dynamics* the government estimated its Phillips curve by constant-gain recursive least squares.

That is a very particular way to learn.

In large samples, a constant-gain least squares learner behaves as if it believed that the coefficients of its model follow a random walk with one specific innovation covariance matrix.

Here we free the government to hold *any* prior about how its coefficients drift, encoded in a covariance matrix V, and we ask how that prior affects the two forces that drive the model:

- the **mean dynamics**, which pull the government's beliefs *toward* a self-confirming (Nash) equilibrium; and
- the **escape dynamics**, which recurrently push beliefs *away* from it, toward the low-inflation Ramsey outcome.

We will find three things.

1. Some priors make the self-confirming equilibrium *unstable*, so that the mean dynamics themselves produce recurrent disinflations: a **learning cycle**, born in a Hopf bifurcation, rather than a stochastic escape.
2. The prior shapes the *direction* and the *speed* of escapes, but in every case we study an escape heads toward the Ramsey outcome.
3. The prior resolves a puzzle in the literature: why the simulated governments of Sims [1988] and Chung [1990] escape Nash inflation once and stay near Ramsey, apparently indefinitely, while those of Sargent [1999] and CWS escape only to be pulled back, again and again.

Throughout we work with the analytically tractable **static** model of *Escaping Nash Inflation* and *A government learning the Phillips curve*, in which the government regresses unemployment on inflation and a constant.

Let's start with our imports:

```
import matplotlib.pyplot as plt
import numpy as np
from scipy.linalg import sqrtm
from scipy.integrate import solve_ivp
from scipy.optimize import brentq
```

105.2 The static model

The true economy is a version of the natural-rate model of Kydland and Prescott [1977]:

$$\begin{aligned} U_n &= u - (\pi_n - \hat{x}_n) + \sigma_1 W_{1n}, \qquad u > 0, \\ \pi_n &= x_n + \sigma_2 W_{2n}, \\ \hat{x}_n &= x_n, \end{aligned} \tag{105.1}$$

where U_n is unemployment, π_n is inflation, x_n is the systematic part of inflation set by the government, $\hat{x}_n$ is the public's (rational) forecast, and $W_n = (W_{1n}, W_{2n})^\top$ is IID standard Gaussian noise.

Since $\pi_n - \hat{x}_n = \sigma_2 W_{2n}$, the true unemployment rate is $U_n = u - \sigma_2 W_{2n} + \sigma_1 W_{1n}$ — it fluctuates around the natural rate u regardless of systematic policy.

The government does not know this.

In the **static model** it fits a non-expectational Phillips curve by regressing unemployment on current inflation and a constant,

$$U_n = a + b\,\pi_n + \eta_n, \tag{105.2}$$

with belief vector $\gamma = (a, b)$ (intercept and slope), and it treats η_n as an exogenous shock.

> **Warning**
>
> *Escaping Nash Inflation* studies this same static model but orders the belief vector the other way round, as $\gamma = (\gamma_1, \gamma_{-1})$ with the *slope* first and regressors $\Phi = (\pi, 1)$, following Cho *et al.* [2002].
>
> Here we put the *intercept* first, with $\Phi = (1, \pi)$, following Sargent and Williams [2005], as *A government learning the Phillips curve* also does.

The two are the same model in different coordinates: $a = \gamma_{-1}$ and $b = \gamma_1$, so the self-confirming belief $(2u, -1)$ of this lecture is the $(-1, u(1+\theta^2))$ of the last one, with $\theta = 1$.

Matrices such as M, V, and P have their rows and columns correspondingly transposed.

A government learning the Phillips curve uses our calibration but other letters, as the Bounded Rationality lectures do (*Notation across the series*): there the beliefs are α and the gain is γ.

Here, following Sargent and Williams [2005], the beliefs are γ, the gain is ε, and α denotes the drifting coefficients that the government thinks generate its data.

Believing (105.2), the government solves the Phelps problem — minimize $\hat{\mathbb{E}} \sum_n \delta^n (U_n^2 + \pi_n^2)$ — whose static best response sets inflation to the constant

$$x(\gamma) = -\frac{a\, b}{1 + b^2}. \tag{105.3}$$

```
class StaticPhillips:
    "The static Sargent-Williams model: government regresses U on (1, π)."

    def __init__(self, u=5.0, σ1=0.3, σ2=0.3):
        self.u, self.σ1, self.σ2 = u, σ1, σ2
        self.σ = σ1                          # govt regression error std = σ1

    def x(self, γ):
        "Government best response (systematic inflation) given beliefs γ."
        a, b = γ
        return -a * b / (1 + b**2)

    def M(self, γ):
        "Second moment matrix E[ΦΦ'] of regressors Φ = (1, π), given γ."
        x = self.x(γ)
        return np.array([[1.0, x], [x, x**2 + self.σ2**2]])

    def g_bar(self, γ):
        "Least squares moment E[Φ(U - Φ'γ)] under the data γ generates."
        u, σ2 = self.u, self.σ2
        x = self.x(γ)
        E_ΦU = np.array([u, x * u - σ2**2])
        return E_ΦU - self.M(γ) @ γ
```

Three belief vectors are worth naming.

- **Belief 1 (Nash):** $b = -1$ with an intercept that makes the government set $x = u$, the time-consistent outcome of Kydland and Prescott [1977].
- **Belief 2 (Ramsey):** $b = 0$, so the government perceives *no* tradeoff and sets $x = 0$.
- **Belief 3 (induction):** in a dynamic version, coefficients on current and lagged inflation summing to zero, which for a patient government also sends inflation toward 0.

105.3 The self-confirming equilibrium

A self-confirming equilibrium is a belief $\bar{\gamma}$ that reproduces itself: the data generated when the government acts on $\bar{\gamma}$ make the population regression coefficients equal $\bar{\gamma}$, i.e. $\bar{g}(\bar{\gamma}) = 0$.

For the static model this is easy to solve by hand.

The slope is $b = \mathrm{cov}(U, \pi)/\mathbb{V}[\pi] = -\sigma_2^2/\sigma_2^2 = -1$, and matching means gives the intercept $a = u + x(\bar{\gamma})$.

Substituting the best response (105.3) with $b = -1$ gives $x = a/2$, so $a = u + a/2$, i.e. $a = 2u$.

```
model = StaticPhillips(u=5.0, σ1=0.3, σ2=0.3)
γ_sce = np.array([2 * model.u, -1.0])

print(f"self-confirming beliefs  γ = {γ_sce}  (intercept 2u, slope -1)")
print(f"self-confirming inflation x = {model.x(γ_sce):.2f}  (= Nash = u)")
print(f"check g_bar(γ_sce) = {model.g_bar(γ_sce)}")
```

```
self-confirming beliefs  γ = [10. -1.]  (intercept 2u, slope -1)
self-confirming inflation x = 5.00  (= Nash = u)
check g_bar(γ_sce) = [0. 0.]
```

The self-confirming equilibrium inflation rate equals the Nash (time-consistent) outcome $u = 5$, even though the government's model is misspecified (*Why the classical equilibrium is Nash* explains why).

The government's non-expectational Phillips curve is observationally equivalent to the truth *along* the equilibrium path, but wrong *off* it — which is exactly what its adaptive behavior repeatedly probes.

105.4 Drifting beliefs and the Kalman filter

We now make the government adaptive.

Following Sargent and Williams [2005], the government believes that its data come from a Phillips curve whose coefficients *drift* as a random walk,

$$U_n = \Phi_n^\top \alpha_n + \eta_n, \qquad \alpha_{n+1} = \alpha_n + \Lambda_{n+1}, \qquad \mathbb{E}\,\eta_n^2 = \sigma^2, \qquad \mathrm{cov}(\Lambda_n) = V, \tag{105.4}$$

with regressors $\Phi_n = (1, \pi_n)^\top$ and IID shocks η_n and Λ_n.

The covariance matrix V is the government's **prior about parameter drift**, the object we set free, and σ^2 is the variance that the government attributes to its regression error η_n.

The government's estimate γ_n is its best forecast of α_n from data through $n-1$, which it computes with the Kalman filter.

This is the linear state-space model of *A First Look at the Kalman Filter*, with hidden state α_n, transition matrix I, state-noise covariance V, a time-varying observation matrix $\Phi_n^\top$, and measurement-noise variance σ^2.

Following Sargent and Williams [2005], we measure the covariance of the estimation error in units of σ^2, writing $\mathrm{cov}(\alpha_n - \gamma_n) = \sigma^2 P_n$.

Then the Kalman filter is

$$\begin{aligned} \gamma_{n+1} &= \gamma_n + \frac{P_n \Phi_n}{1 + \Phi_n^\top P_n \Phi_n}\left(U_n - \Phi_n^\top \gamma_n\right), \\ P_{n+1} &= P_n - \frac{P_n \Phi_n \Phi_n^\top P_n}{1 + \Phi_n^\top P_n \Phi_n} + \sigma^{-2} V. \end{aligned} \tag{105.5}$$

The normalization is what puts a 1 in the denominators and the scaled drift $\sigma^{-2}V$ in the update of P_n.

For large n, Benveniste *et al.* [1990] show that the filter is well approximated by

$$\begin{aligned} \gamma_{n+1} &= \gamma_n + P_n \Phi_n \left(U_n - \Phi_n^\top \gamma_n\right), \\ P_{n+1} &= P_n - P_n M(\gamma_n) P_n + \sigma^{-2} V, \end{aligned} \tag{105.6}$$

in which the second moment matrix $M(\gamma_n)$ replaces $\Phi_n \Phi_n^\top$ in the update of P_n.

Sargent and Williams [2005] base their analysis on (105.6), and report that for the small drift they study it gives simulations nearly indistinguishable from those of (105.5).

We simulate with the exact filter (105.5) throughout, because the approximation can be unstable when the drift is large, as it is in Sims's example below.

For a fixed γ, the matrix P_n settles at the solution of the algebraic Riccati equation

$$-PM(\gamma)P + \sigma^{-2}V = 0, \tag{105.7}$$

the counterpart for this model of the algebraic Riccati equation in *Convergence*.

Following Sargent and Williams [2005], we index the size of the drift by a scalar ε, writing $V = \varepsilon^2 \hat{V}$ for a fixed matrix $\hat{V}$.

Equation (105.7) then makes P proportional to ε/σ, so ε plays the role of the constant gain of *Adaptive Learning and Escape Dynamics*.

Most of our calculations use $\varepsilon = 0.05$, the largest gain in Tables 1 and 2 of Sargent and Williams [2005]; when we compare with constant-gain least squares we use the gain 0.02 of *A government learning the Phillips curve*, and Sims's example carries its own, much larger, drift.

Constant-gain recursive least squares is a special case.

If the government's prior is $V = \varepsilon^2 V^*$, with $V^* \equiv \sigma^2 M(\bar{\gamma})^{-1}$, and $\sigma = \sigma_1$, then (105.7) gives $P = \varepsilon M(\bar{\gamma})^{-1}$, and (105.6) becomes the constant-gain recursive least squares algorithm of *Adaptive Learning and Escape Dynamics*, with gain ε and with the estimated moment matrix replaced by its limit.

The connection is asymptotic, not exact.

In large samples the two algorithms merge, but their transient responses differ: Sargent and Williams [2005] (section 3.4) show that after an escape the Kalman filter returns to the self-confirming equilibrium much faster than recursive least squares does (Exercise 105.10.3).

```
def solve_riccati(V, M, σ):
    "Symmetric positive-definite P solving P M P = σ^{-2} V."
    W = V / σ**2
    Mh = sqrtm(M).real
    Mh_inv = np.linalg.inv(Mh)
    return Mh_inv @ sqrtm(Mh @ W @ Mh).real @ Mh_inv

M_sce = model.M(γ_sce)
V_star = model.σ**2 * np.linalg.inv(M_sce)      # the RLS prior (ε = 1)
P_star = solve_riccati(V_star, M_sce, model.σ)

print("RLS prior gives P = M^{-1}?", np.allclose(P_star, np.linalg.inv(M_sce)))
```

```
RLS prior gives P = M^{-1}? True
```

The next function simulates the economy when the government uses the exact filter (105.5).

```
def kalman_path(model, V, σ, T, seed, γ0=None):
    """
    Simulate T periods in which the government uses the exact Kalman filter
    with prior (V, σ), while the truth uses model.σ1.

    Beliefs start at γ0 (default: the self-confirming equilibrium), with P at
    its Riccati value.  Returns inflation π, unemployment U, systematic
    inflation x, and beliefs γ.
    """
    rng = np.random.default_rng(seed)
    u, σ1, σ2 = model.u, model.σ1, model.σ2
    γ = np.array([2 * u, -1.0]) if γ0 is None else np.array(γ0, dtype=float)
    P = solve_riccati(V, model.M(γ), σ)
    W = V / σ**2
    π, U, x = np.empty(T), np.empty(T), np.empty(T)
    beliefs = np.empty((T, 2))
    for n in range(T):
        beliefs[n] = γ
        x[n] = model.x(γ)
        w1, w2 = rng.standard_normal(2)
        π[n] = x[n] + σ2 * w2
        U[n] = u - σ2 * w2 + σ1 * w1                # truth uses σ1
        Φ = np.array([1.0, π[n]])
        PΦ = P @ Φ
        denom = 1 + Φ @ PΦ
        γ = γ + PΦ / denom * (U[n] - Φ @ γ)
        P = P - np.outer(PΦ, PΦ) / denom + W
    return π, U, x, beliefs
```

105.5 Mean dynamics and E-stability

As in *Adaptive Learning and Escape Dynamics*, the beliefs are organized by mean dynamics, now a *joint* ordinary differential equation in (γ, P), with time measured in periods:

$$\dot{\gamma} = P\,\bar{g}(\gamma), \qquad \dot{P} = \sigma^{-2}V - PM(\gamma)P. \tag{105.8}$$

A rest point of (105.8) has $\bar{g}(\gamma) = 0$ and $P =$ the Riccati solution — a self-confirming equilibrium.

Local stability turns on the Jacobian of $\bar{g}$ at the self-confirming equilibrium.

For the static model this can be computed in closed form,

$$\frac{\partial \bar{g}}{\partial \gamma}(\bar{\gamma}) = -\begin{bmatrix} \frac{1}{2} & u \\ \frac{1}{2}u & u^2 + \sigma_2^2 \end{bmatrix}. \tag{105.9}$$

```
def jacobian(model, γ, h=1e-6):
    "Numerical Jacobian of g_bar at γ."
    J = np.zeros((2, 2))
    for j in range(2):
        gp, gm = γ.copy(), γ.copy()
        gp[j] += h; gm[j] -= h
        J[:, j] = (model.g_bar(gp) - model.g_bar(gm)) / (2 * h)
    return J

J = jacobian(model, γ_sce)
print("∂g/∂γ at SCE =\n", J.round(3))
```

```
∂g/∂γ at SCE =
 [[ -0.5   -5.  ]
 [ -2.5  -25.09]]
```

Under constant-gain recursive least squares the moment-matrix block of the corresponding differential equation decouples, and stability is governed by the eigenvalues of $M^{-1}\,\partial\bar{g}/\partial\gamma$.

Writing $\bar{g}(\gamma) = M(\gamma)\,[T(\gamma) - \gamma]$, where $T(\gamma) = M(\gamma)^{-1}\,\mathbb{E}[\Phi U]$ maps the government's perceived Phillips curve into the population regression that its policy generates, shows that at the self-confirming equilibrium $M^{-1}\,\partial\bar{g}/\partial\gamma = DT - I$.

Negative real parts of these eigenvalues are the **E-stability** condition.

The idea goes back to DeCanio [1979] and Evans [1985], Evans [1989]; Marcet and Sargent [1989] and Woodford [1990] showed that it governs the convergence of least squares learning, and Evans and Honkapohja [2001] systematized the theory (*One recursion, many learners*).

```
eig_rls = np.linalg.eigvals(np.linalg.inv(M_sce) @ J)
print(f"E-stability eigenvalues (RLS): {eig_rls.round(3)}")
print("both negative ⇒ the SCE is E-stable")
```

```
E-stability eigenvalues (RLS): [-0.5 -1. ]
both negative ⇒ the SCE is E-stable
```

Both eigenvalues are negative: they are -1 and $-\frac{1}{2}$.

They are the static model's version of numbers met earlier in the suite.

With $\theta = 1$, as here, $-\frac{1}{2} = -1/(1+\theta^2)$ is the slope $\mathcal{M}$ of the mean dynamics of expected inflation in *Least squares learning converges to Nash*, and *Mean dynamics near the self-confirming equilibrium* finds the eigenvalues -1 and $-1/(1+\theta^2)$ again in the dynamic model of *Adaptive Learning and Escape Dynamics*.

At $\theta = 1$ the second eigenvalue sits exactly on the boundary $\mathcal{M} = -\frac{1}{2}$ that *The Credibility Problem* identifies for $\sqrt{t}$ asymptotic normality of decreasing-gain least squares.

That boundary concerns the rate of convergence; how often beliefs escape is governed instead by the large-deviation rate discussed below.

With a decreasing gain, then, least squares beliefs converge to the Nash self-confirming equilibrium asymptotically [Marcet and Sargent, 1989], although in a finite sample a falling gain also slows the return, so that an early escape can stick (*Least squares learning: a decreasing gain*).

With a constant gain, beliefs hover near the self-confirming equilibrium and escape from it recurrently.

But the reduction to $M^{-1}\,\partial\bar{g}/\partial\gamma$ relied on the special RLS prior.

Under a *general* prior V, the P block of (105.8) does **not** decouple, and local stability is governed instead by the eigenvalues of $\bar{P}\,\partial\bar{g}/\partial\gamma$ — where $\bar{P}$ is the Riccati solution for that prior.

Prior beliefs about parameter drift now matter for whether the self-confirming equilibrium is even stable.

105.6 Learning cycles

Here is the paper's most striking result: some priors make the self-confirming equilibrium *unstable*, and the mean dynamics then cycle.

We illustrate it, as Sargent and Williams [2005] do, by *tightening the government's prior on the slope coefficient*, starting from the RLS prior V^* and shrinking the slope-related entries by a factor $\lambda \in [0, 1]$:

$$V(\lambda) = \begin{bmatrix} V^*_{11} & \sqrt{\lambda}\, V^*_{12} \\ \sqrt{\lambda}\, V^*_{12} & \lambda\, V^*_{22} \end{bmatrix}. \tag{105.10}$$

For each λ we solve the Riccati equation with $\varepsilon = 0.05$ and look at the largest real part among the eigenvalues of $\bar{P}(\lambda)\, \partial\bar{g}/\partial\gamma$: where it is positive, the self-confirming equilibrium is unstable.

```
def V_tighten_slope(λ, V_star):
    V = V_star.copy()
    V[0, 1] = V[1, 0] = np.sqrt(λ) * V_star[0, 1]
    V[1, 1] = λ * V_star[1, 1]
    return V

ε = 0.05                                    # gain (sets the timescale)

def eig_PJ(V, σ=model.σ):
    "Eigenvalues of P̄ ∂g/∂γ at the SCE for the prior (V, σ)."
    return np.linalg.eigvals(solve_riccati(V, M_sce, σ) @ J)

λ_grid = np.linspace(0.01, 0.999, 200)
max_re = np.array([eig_PJ(ε**2 * V_tighten_slope(λ, V_star)).real.max()
                   for λ in λ_grid])

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(λ_grid, max_re)
ax.axhline(0, color='k', lw=0.8)
ax.set_xlabel(r'prior-tightening parameter $\lambda$')
ax.set_ylabel('max real part of eigenvalue')
plt.show()
```

Over an intermediate range of λ the maximum real part becomes *positive*: the self-confirming equilibrium loses stability.

The figure reproduces Fig. 4 of Sargent and Williams [2005].

Let's locate the two points where stability changes, and look at the eigenvalues there and deep inside the unstable range.

```
def max_real(λ):
    return eig_PJ(ε**2 * V_tighten_slope(λ, V_star)).real.max()

sign_changes = np.where(np.diff(np.sign(max_re)) != 0)[0]
for i in sign_changes:
    λ_c = brentq(max_real, λ_grid[i], λ_grid[i + 1])
    eigs = eig_PJ(ε**2 * V_tighten_slope(λ_c, V_star))
    print(f"stability changes at λ = {λ_c:.3f}; eigenvalues there: {eigs.round(4)}")
print(f"eigenvalues at λ = 0.7: {eig_PJ(ε**2 * V_tighten_slope(0.7, V_star)).round(4)}
 ↪")
```

```
stability changes at λ = 0.261; eigenvalues there: [-0.+0.0253j -0.-0.0253j]
stability changes at λ = 0.956; eigenvalues there: [0.+0.035j 0.-0.035j]
eigenvalues at λ = 0.7: [0.0055 0.1897]
```

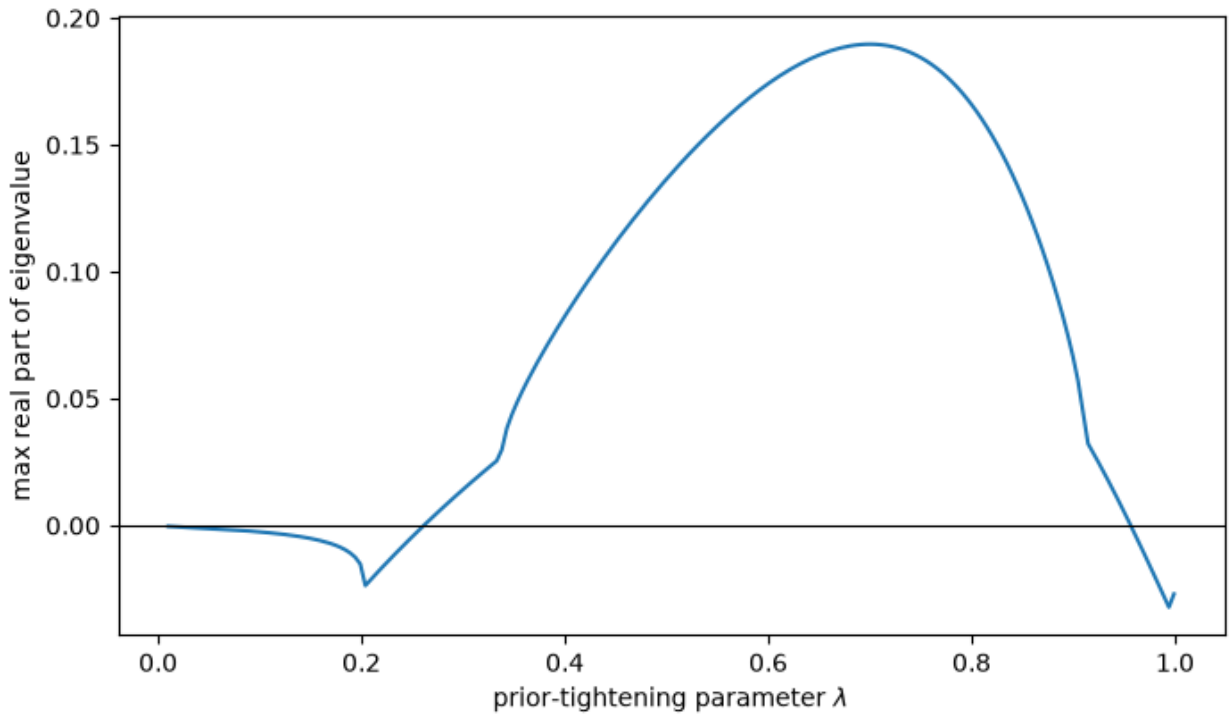


Fig. 105.1: Stability of the self-confirming equilibrium as the slope prior tightens

At each end of the unstable range a complex pair of eigenvalues crosses the imaginary axis.

That is the situation the Hopf bifurcation theorem addresses (see Perko [1996]): a stable limit cycle bifurcates from the self-confirming equilibrium as the real parts cross zero (Theorem 4.1 of Sargent and Williams [2005]).

Bullard [1994] calls such a cycle a *learning equilibrium*: the underlying model has no cycles, and learning creates them.

Deep inside the unstable range, at $\lambda = 0.7$, both eigenvalues are real and positive, so the self-confirming equilibrium is an unstable node.

The Hopf theorem describes cycles near the bifurcation points, so at $\lambda = 0.7$ we find the cycle numerically, as Sargent and Williams [2005] do in their Fig. 5, by integrating the joint mean dynamics (105.8) for (γ, P) from a point near the self-confirming equilibrium.

```
def mean_dynamics(t, z, V, σ):
    "The joint ODE pp_ode for (γ, P), stacked into one vector."
    γ, P = z[:2], z[2:].reshape(2, 2)
    γ_dot = P @ model.g_bar(γ)
    P_dot = V / σ**2 - P @ model.M(γ) @ P
    return np.concatenate([γ_dot, P_dot.ravel()])

λ = 0.7
V = ε**2 * V_tighten_slope(λ, V_star)
P_bar = solve_riccati(V, M_sce, model.σ)

z0 = np.concatenate([γ_sce + np.array([0.3, 0.05]), P_bar.ravel()])
sol = solve_ivp(mean_dynamics, [0, 3000], z0, args=(V, model.σ),
                max_step=1.0, rtol=1e-9, atol=1e-11)
t_path, a_path, b_path = sol.t, sol.y[0], sol.y[1]
x_path = model.x((a_path, b_path))
```

(continues on next page)

(continued from previous page)

```
# period: time between upward crossings of x through u/2, once the cycle has settled
up = t_path[1:][(x_path[:-1] < model.u / 2) & (x_path[1:] >= model.u / 2)]
period = np.diff(up[up > 1000]).mean()
mask = t_path > t_path[-1] - 2 * period        # the last two cycles

print(f"period of the cycle ≈ {period:.0f} periods")
print(f"along the cycle: intercept in [{a_path[mask].min():.2f}, {a_path[mask].max():.
↪2f}], "
      f"slope in [{b_path[mask].min():.2f}, {b_path[mask].max():.2f}], "
      f"inflation in [{x_path[mask].min():.2f}, {x_path[mask].max():.2f}]")
```

```
period of the cycle ≈ 391 periods
along the cycle: intercept in [5.03, 11.41], slope in [-1.16, -0.06], inflation in␣
↪[0.28, 5.65]
```

Here are the last two cycles, in time and in belief space.

```
fig, axes = plt.subplots(1, 2, figsize=(12, 5))

axes[0].plot(t_path[mask], a_path[mask], label='intercept', lw=2)
axes[0].plot(t_path[mask], b_path[mask], label='slope', lw=2)
axes[0].set_xlabel('time (periods)')
axes[0].set_ylabel('coefficient')
axes[0].legend()

axes[1].plot(a_path[mask], b_path[mask], lw=2)
axes[1].plot(*y_sce, 'kx', ms=10, label='SCE')
axes[1].set_xlabel('intercept')
axes[1].set_ylabel('slope')
axes[1].legend()

plt.tight_layout()
plt.show()
```

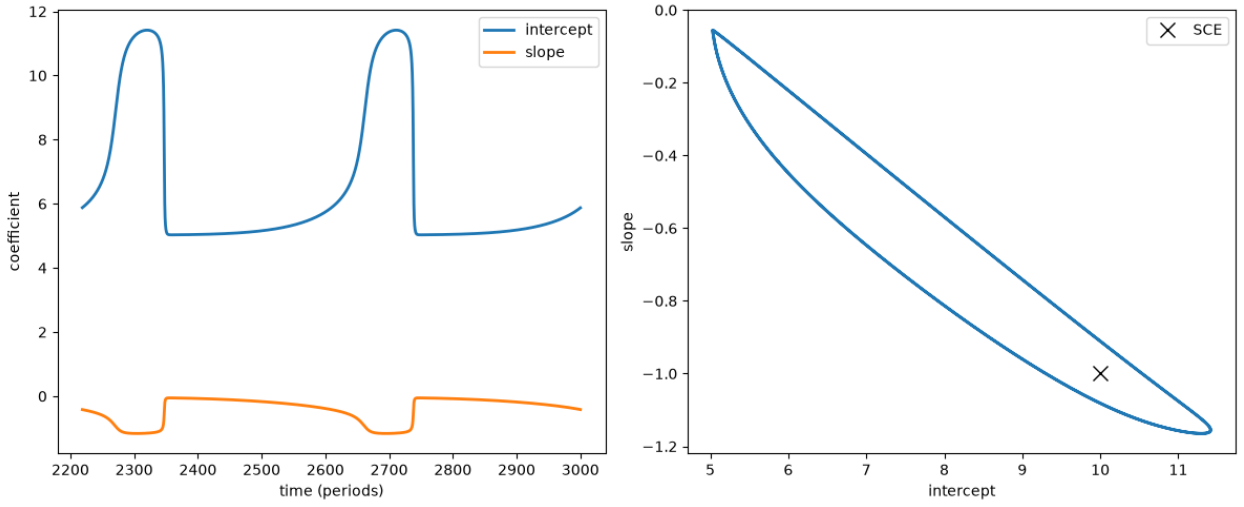


Fig. 105.2: Mean dynamics at $\lambda = 0.7$, the coefficients over two cycles (left) and the limit cycle in belief space (right)

The beliefs settle into a closed orbit around the self-confirming equilibrium, which matches Fig. 5 of Sargent and Williams [2005].

The intercept falls from near $2u$ to near u while the slope climbs from about -1 toward 0, and then both snap back.

Because inflation is a function of the coefficients through the best response (105.3), the cycle in beliefs shows up as a cycle in inflation between the Nash and Ramsey outcomes.

```
fig, ax = plt.subplots(figsize=(9, 4.5))
ax.plot(t_path[mask], x_path[mask])
ax.axhline(model.u, color='k', ls='--', lw=1, label='Nash')
ax.axhline(0, color='C2', ls=':', lw=1, label='Ramsey')
ax.set_xlabel('time (periods)')
ax.set_ylabel('inflation $x$')
ax.legend()
plt.show()
```

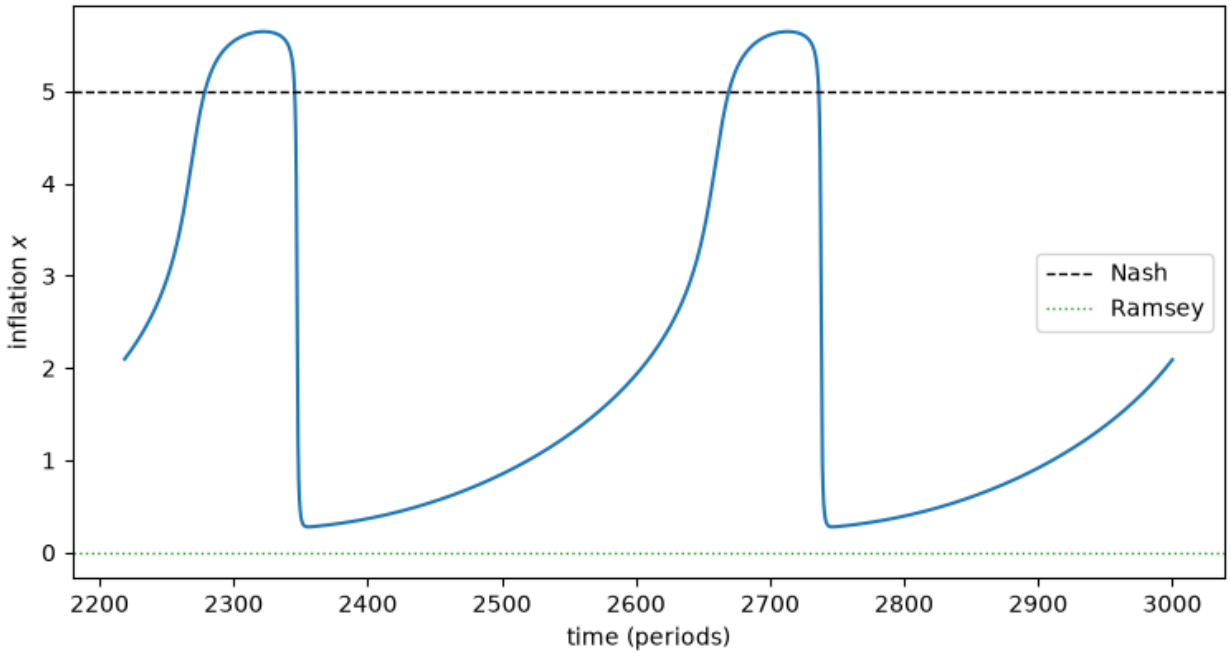


Fig. 105.3: Inflation oscillates between Nash and Ramsey along the learning cycle

This is a different mechanism from the escapes of *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation*.

There the self-confirming equilibrium is stable, and a disinflation needs an unusual run of shocks to push beliefs away from it; here the mean dynamics themselves carry beliefs from Nash toward Ramsey and back.

The difference is sharpest as the gain shrinks.

As $\varepsilon \to 0$ the cycle persists, with a period proportional to $1/\varepsilon$, while escapes from a stable self-confirming equilibrium become exponentially rare.

But "rare" is a statement about small gains.

For the baseline prior, Sargent and Williams [2005] compute a large-deviation rate $\bar{S} \approx 3 \times 10^{-4}$ (their Table 1), so the factor $\exp(\bar{S}/\varepsilon)$ that makes escapes rare is close to one at the gains of 0.02 to 0.05 used in this lecture and elsewhere in the suite, and escapes are frequent; *Simulating escapes* shows how escape times grow as the gain shrinks.

At those gains, escapes from a stable self-confirming equilibrium recur on the same time scale as this cycle turns over, and Exercise 105.10.4 shows that a simulated time series can hardly tell the two apart.

105.7 Escape dynamics and the direction of escape

For priors that keep the self-confirming equilibrium stable, disinflations are escapes.

The theory of large deviations (*A primer on recursive algorithms*), in the form developed in Noah Williams's dissertation [Williams, 2001], later published as Williams [2019], characterizes the most likely escape as the solution of a control problem: apply the least-cost perturbation to beliefs that pushes them a fixed distance from $\bar{\gamma}$.

The minimized cost $\bar{S}$ is the rate at which escapes become rare as the gain shrinks: if τ is the time of the first escape, then $\varepsilon \log \mathbb{E}\tau \to \bar{S}$ as $\varepsilon \to 0$.

The *instantaneous* escape direction has a beautifully simple characterization: it is the eigenvector associated with the largest eigenvalue of the belief-innovation covariance $Q(\bar{\gamma}, \bar{P}) = \hat{V}$ — the prior itself.

```
w, vecs = np.linalg.eigh(V_star)            # baseline prior V*
v_escape = vecs[:, np.argmax(w)]
v_escape = v_escape / np.linalg.norm(v_escape)

# scale the escape so the slope moves from -1 up to 0
t_star = 1.0 / v_escape[1]
terminal = y_sce + t_star * v_escape

print(f"escape direction  ≈ {v_escape.round(3)}   (∝ [-u, 1])")
print(f"terminal beliefs  ≈ {terminal.round(2)}   (Belief 2 = [u, 0] = Ramsey)")
```

```
escape direction  ≈ [-0.981  0.195]   (∝ [-u, 1])
terminal beliefs  ≈ [4.98 0.  ]   (Belief 2 = [u, 0] = Ramsey)
```

For the baseline prior the escape direction is nearly proportional to $(-u, 1)$, exactly so in the limit $\sigma \to 0$, and it carries beliefs from Nash $(2u, -1)$ toward $(u, 0)$ — Belief 2, which supports the Ramsey outcome.

So an escape, when it happens, is a movement toward zero inflation.

Different priors bend the *path* and change the *rate* of escape.

Tightening the slope prior can destabilize the equilibrium entirely (the cycles above).

Tightening the intercept prior keeps it stable but makes escapes much *slower*: in Table 2 of Sargent and Williams [2005] it raises $\bar{S}$ from 3.1×10^{-4} to 0.28, and the mean escape time at $\varepsilon = 0.05$ from 23 to 167 periods (Exercise 105.10.1).

Yet in every case that Sargent and Williams [2005] study, the escapes lead to zero inflation, by different routes and supported by different beliefs (Exercise 105.10.2).

105.8 Sims's nonconvergence

We can now resolve the puzzle described in *The Rise and Fall of U.S. Inflation*.

Sims [1988] and Chung [1990] gave their governments Kalman filters on drifting coefficients; Sims's government fit a static Phillips curve like ours, and Chung's a distributed lag.

In their simulations the government escapes from Nash inflation once and then *stays* near the Ramsey outcome, apparently indefinitely.

Sargent [1999] and CWS instead used constant-gain recursive least squares, and their governments escape and are pulled back, again and again, producing the sawtooth of *A government learning the Phillips curve*.

Sargent and Williams [2005] trace the difference to the government's prior: the drift covariance V and the noise variance σ^2 that its filter assumes.

105.8.1 The variance of the regression error

So far, following CWS, the government's σ has equaled the true σ_1.

That is natural: in a self-confirming equilibrium the regression (105.2) and the truth (105.1) coincide, and with $\sigma = \sigma_1$ the government correctly splits the variation it sees in unemployment between its regressors and the shocks.

Sims instead set σ well below σ_1, and he did not let the drift in the coefficients shrink toward zero.

A government that attributes too little variance to its regression error must attribute the variation it sees to something else, and the only thing else available is drift in its own coefficients.

The Riccati equation (105.7) makes this precise: $\bar{P}$ solves $\bar{P}M\bar{P} = \sigma^{-2}V$, so $\bar{P}$ is proportional to $1/\sigma$.

Understating σ by some factor therefore raises the government's gain by the same factor, exactly as raising ε would.

```
Φ_sce = np.array([1.0, model.x(y_sce)])
P_by_σ = {}
for σ_govt in (model.σ1, model.σ1 / 3):
    P_by_σ[σ_govt] = solve_riccati(ε**2 * V_star, M_sce, σ_govt)
    P = P_by_σ[σ_govt]
    gain = (P @ Φ_sce)[0] / (1 + Φ_sce @ P @ Φ_sce)
    print(f"σ = {σ_govt:.2f}: steady-state weight on a forecast error "
          f"(intercept) = {gain:.4f}")

ratio = P_by_σ[model.σ1 / 3] / P_by_σ[model.σ1]
print(f"P(σ1/3) / P(σ1), element by element: {ratio.round(6).ravel()}")
```

```
σ = 0.30: steady-state weight on a forecast error (intercept) = 0.0476
σ = 0.10: steady-state weight on a forecast error (intercept) = 0.1304
P(σ1/3) / P(σ1), element by element: [3. 3. 3. 3.]
```

Cutting σ to a third of σ_1 triples $\bar{P}$ exactly.

The weight that the filter puts on each forecast error rises almost as much: the 1 in the denominator of (105.5) keeps the ratio a little below three at this gain, and the ratio tends to three as ε shrinks.

At small gains, a larger gain changes the speed of learning but not the local balance between the mean dynamics and the noise: Sargent and Williams [2005] (section 5) show that the local strength of the pull toward the self-confirming equilibrium, measured as drift per unit of standard deviation, does not depend on the prior.

What a large gain can change is whether the mean dynamics describe the beliefs at all, because they are a small-gain approximation.

105.8.2 Replicating Sims's example

Sargent and Williams [2005] simulate a specification analogous to one of Sims's benchmark cases, his Chart 5, which in our notation has $u = 6$, $\sigma_1 = 0.5$, $\sigma_2 = 0.4$, $V = 0.01\,I$ and $\sigma = 0.1$, so that σ_1 is five times σ; they treat it as a benchmark with $\varepsilon = 1$.

We start the government at the self-confirming equilibrium $(2u, -1) = (12, -1)$, with P at its Riccati value, and run 10,000 periods, as Sims and Sargent–Williams did.

We compare it with the variation $\sigma = \sigma_1$, holding V fixed.

This variation leaves $\sigma^{-2}V = 0.04\,I$, the same as shrinking the drift by setting $\varepsilon = 0.2$ with $\sigma = 0.1$, so the two give *identical* paths, as Sargent and Williams [2005] note.

We also need a way to count escapes and returns along a path of inflation.

```
def escapes_and_returns(π, lo, hi, k=50):
    """
    Count escapes and returns along a path of inflation that starts near Nash.

    An escape is a fall of the k-period moving average of inflation below lo;
    a return is a later rise back above hi.  Also returns the period by which
    the first escape is complete (None if there is none).
    """
    ma = np.convolve(π, np.ones(k) / k, mode='valid')
    near_nash, n_esc, n_ret, first = True, 0, 0, None
    for n, m in enumerate(ma):
        if near_nash and m < lo:
            near_nash, n_esc = False, n_esc + 1
            if first is None:
                first = n + k - 1
        elif not near_nash and m > hi:
            near_nash, n_ret = True, n_ret + 1
    return n_esc, n_ret, first
```

Now we simulate ten paths of Sims's specification and of the variation, counting an escape when the 50-period average of inflation falls below $u/3 = 2$ and a return when it climbs back above $5u/6 = 5$.

```
sims = StaticPhillips(u=6.0, σ1=0.5, σ2=0.4)
V_sims, σ_sims = 0.01 * np.eye(2), 0.1
T_sims, seeds = 10_000, range(10)
lo, hi = sims.u / 3, 5 * sims.u / 6

runs = {s: kalman_path(sims, V_sims, σ_sims, T_sims, s) for s in seeds}
runs_σ1 = {s: kalman_path(sims, V_sims, sims.σ1, T_sims, s) for s in seeds}

same = all(np.allclose(kalman_path(sims, 0.04 * V_sims, σ_sims, T_sims, s)[0],
                       runs_σ1[s][0]) for s in (0, 1))
print(f"σ = σ1 and ε = 0.2 give identical paths: {same}\n")

print(f"{'':6}{'Sims: σ = 0.1':^42}| {'σ = σ1':^24}")
print(f"{'seed':>4} {'escapes':>8} {'by period':>10} {'returns':>8} "
      f"{'mean π after':>13} | {'escapes':>8} {'mean π':>8}")
first_escape = {}
for s in seeds:
    π = runs[s][0]
    n_esc, n_ret, first = escapes_and_returns(π, lo, hi)
    first_escape[s] = first
    after = π[first:].mean() if first is not None else np.nan
```

(continues on next page)

(continued from previous page)

```
    n_esc1 = escapes_and_returns(runs_σ1[s][0], lo, hi)[0]
    print(f"{s:>4} {n_esc:>8} {str(first):>10} {n_ret:>8} {after:>13.2f} | "
          f"{n_esc1:>8} {runs_σ1[s][0].mean():>8.2f}")
```

```
σ = σ1 and ε = 0.2 give identical paths: True

                    Sims: σ = 0.1                |          σ = σ1
seed  escapes  by period  returns  mean π after |  escapes   mean π
   0        1       4807        0          0.61 |        0     6.14
   1        1       3488        0          0.57 |        0     6.15
   2        1        307        0          0.59 |        0     5.83
   3        1        417        0          0.65 |        0     5.89
   4        1        598        0          0.63 |        0     5.52
   5        1       4509        0          0.60 |        0     6.07
   6        1       1289        0          0.67 |        0     5.94
   7        1       1284        0          0.59 |        0     5.82
   8        1       1133        0          0.59 |        0     6.03
   9        1       1433        0          0.57 |        0     5.61
```

Every one of Sims's paths escapes, at very different dates, and none returns: after its escape each path settles into a noisy regime with mean inflation close to the 0.61 that Sargent and Williams [2005] report.

With $\sigma = \sigma_1$, or equivalently $\varepsilon = 0.2$, no path escapes, and inflation stays near the self-confirming value $u = 6$.

The next figure shows the first 1,000 periods of the path whose escape comes earliest, the analog of Fig. 14 of Sargent and Williams [2005].

```
s_show = min(seeds, key=lambda s: first_escape[s])

fig, ax = plt.subplots(figsize=(9.5, 4.5))
ax.plot(runs[s_show][0][:1000], lw=0.8, label=r"Sims: $\sigma = 0.1$")
ax.plot(runs_σ1[s_show][0][:1000], lw=0.8, ls='--',
        label=r"$\sigma = \sigma_1$ or $\varepsilon = 0.2$")
ax.axhline(sims.u, color='k', ls='--', lw=1)
ax.axhline(0, color='C2', ls=':', lw=1)
ax.set_xlabel('$n$')
ax.set_ylabel('inflation')
ax.legend(frameon=False, loc='lower right')
plt.show()
```

105.8.3 Why the mean dynamics do not reassert themselves

The mean dynamics (105.8) average the data over many periods while holding beliefs, and hence the government's policy $x(\gamma)$, fixed.

That is a good approximation when beliefs move slowly, but not under Sims's prior, which lets beliefs move a lot from one period to the next.

Then the government's own policy injects variation into inflation, and that variation does not move unemployment, because only surprises do.

The regression of U on π then reads the induced variation as evidence of a nearly vertical Phillips curve: with policy variance $\mathbb{V}[x]$, the population slope is $-\sigma_2^2/(\sigma_2^2 + \mathbb{V}[x])$ rather than -1.

A flat perceived curve supports low inflation, so the low-inflation regime sustains itself.

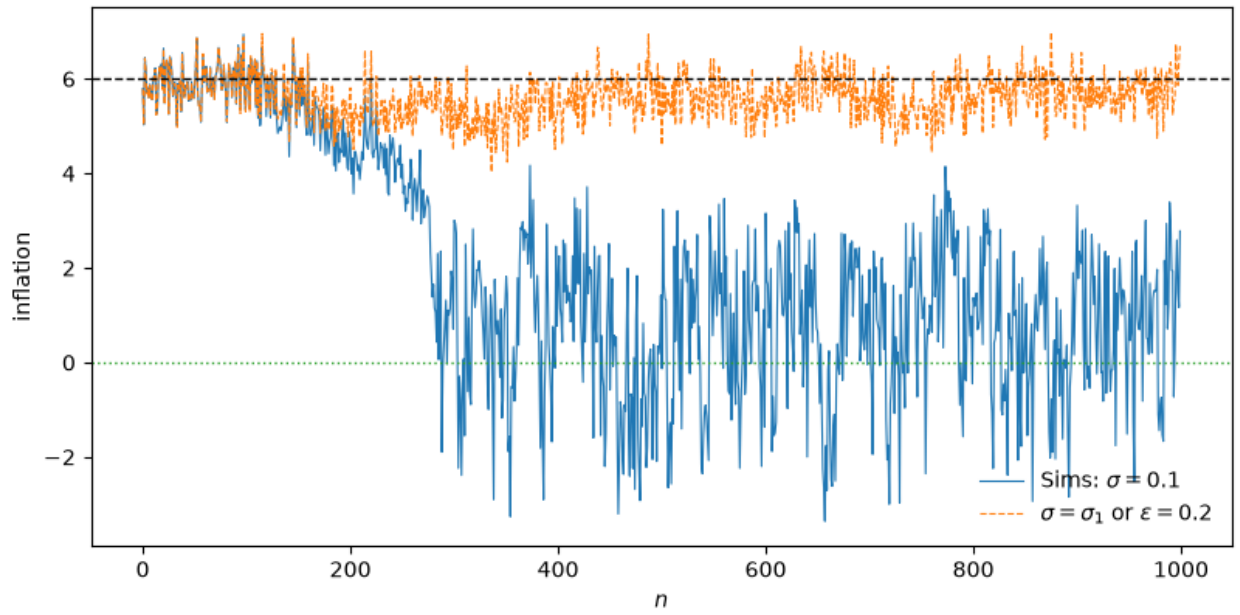


Fig. 105.4: Inflation under Sims's specification and under σ = σ1 (equivalently ε = 0.2), first 1,000 periods

Let's check this on the simulated paths after their escapes.

```
var_x, slope, theory = [], [], []
for s in seeds:
    π, U, x, _ = runs[s]
    n0 = first_escape[s] + 200                    # well after the escape
    var_x.append(x[n0:].var())
    slope.append(np.cov(U[n0:], π[n0:])[0, 1] / π[n0:].var())
    theory.append(-sims.σ2**2 / (sims.σ2**2 + var_x[-1]))

print(f"variance of systematic inflation x after the escape: "
      f"{min(var_x):.2f} to {max(var_x):.2f}")
print(f"sample slope of U on π after the escape:  {min(slope):.3f} to {max(slope):.3f}
↪")
print(f"-σ2²/(σ2² + var x):                       {min(theory):.3f} to {max(theory):.
↪3f}")
```

```
variance of systematic inflation x after the escape: 2.10 to 2.21
sample slope of U on π after the escape:  -0.073 to -0.062
-σ2²/(σ2² + var x):                       -0.071 to -0.068
```

The government's beliefs, and therefore its policy, keep moving after the escape, and the data it generates confirm a Phillips curve whose slope is close to zero, the Ramsey belief.

That is the sense in which Sargent and Williams [2005] say that the results about mean and escape dynamics are "approximations for small parameter drift": with drift as large as Sims's, the mean dynamics fail to reassert themselves.

105.8.4 Constant gains and the sawtooth

Now give the government of the same economy constant-gain recursive least squares instead, with the gain 0.02 of *A government learning the Phillips curve* (the economy of that section has the calibration of the rest of this lecture).

```
def rls_path(model, gain, T, seed):
    "Constant-gain recursive least squares, as in olg_phillips, from the SCE."
    rng = np.random.default_rng(seed)
    u, σ1, σ2 = model.u, model.σ1, model.σ2
    γ = np.array([2 * u, -1.0])
    R = model.M(γ)
    π = np.empty(T)
    for n in range(T):
        w1, w2 = rng.standard_normal(2)
        π[n] = model.x(γ) + σ2 * w2
        U = u - σ2 * w2 + σ1 * w1
        Φ = np.array([1.0, π[n]])
        R = R + gain * (np.outer(Φ, Φ) - R)
        γ = γ + gain * np.linalg.solve(R, Φ * (U - Φ @ γ))
    return π

runs_cg = {s: rls_path(sims, 0.02, T_sims, s) for s in seeds}
counts_cg = [escapes_and_returns(runs_cg[s], lo, hi)[:2] for s in seeds]
print("constant gain 0.02, (escapes, returns) by seed:", counts_cg)
```

```
constant gain 0.02, (escapes, returns) by seed: [(4, 3), (4, 3), (6, 5), (3, 2),␣
↪(2, 1), (6, 5), (3, 2), (3, 2), (5, 4), (3, 2)]
```

Under constant-gain least squares every path escapes and returns, most of them several times.

The count of returns is conservative: many climbs back toward Nash turn into new escapes before the 50-period average reaches 5.

The figure compares 50-period moving averages of inflation under the two learning rules.

```
def moving_average(π, k=50):
    return np.convolve(π, np.ones(k) / k, mode='valid')

fig, axes = plt.subplots(2, 1, figsize=(9.5, 7), sharex=True, sharey=True)
for s in seeds:
    axes[0].plot(moving_average(runs[s][0]), lw=0.6)
for s in (0, 1, 2):
    axes[1].plot(moving_average(runs_cg[s]), lw=0.8, label=f'seed {s}')
titles = [r"Sims's Kalman filter ($\sigma = 0.1$), ten seeds",
          'constant-gain least squares (gain 0.02), three seeds']
for ax, title in zip(axes, titles):
    ax.axhline(sims.u, color='k', ls='--', lw=1)
    ax.axhline(0, color='C2', ls=':', lw=1)
    ax.set_ylabel('inflation, 50-period average')
    ax.set_title(title)
axes[1].legend(frameon=False, fontsize=8, ncol=3, loc='lower right')
axes[1].set_xlabel('$n$')
plt.tight_layout()
plt.show()
```

The same economy produces both of the histories in the literature.

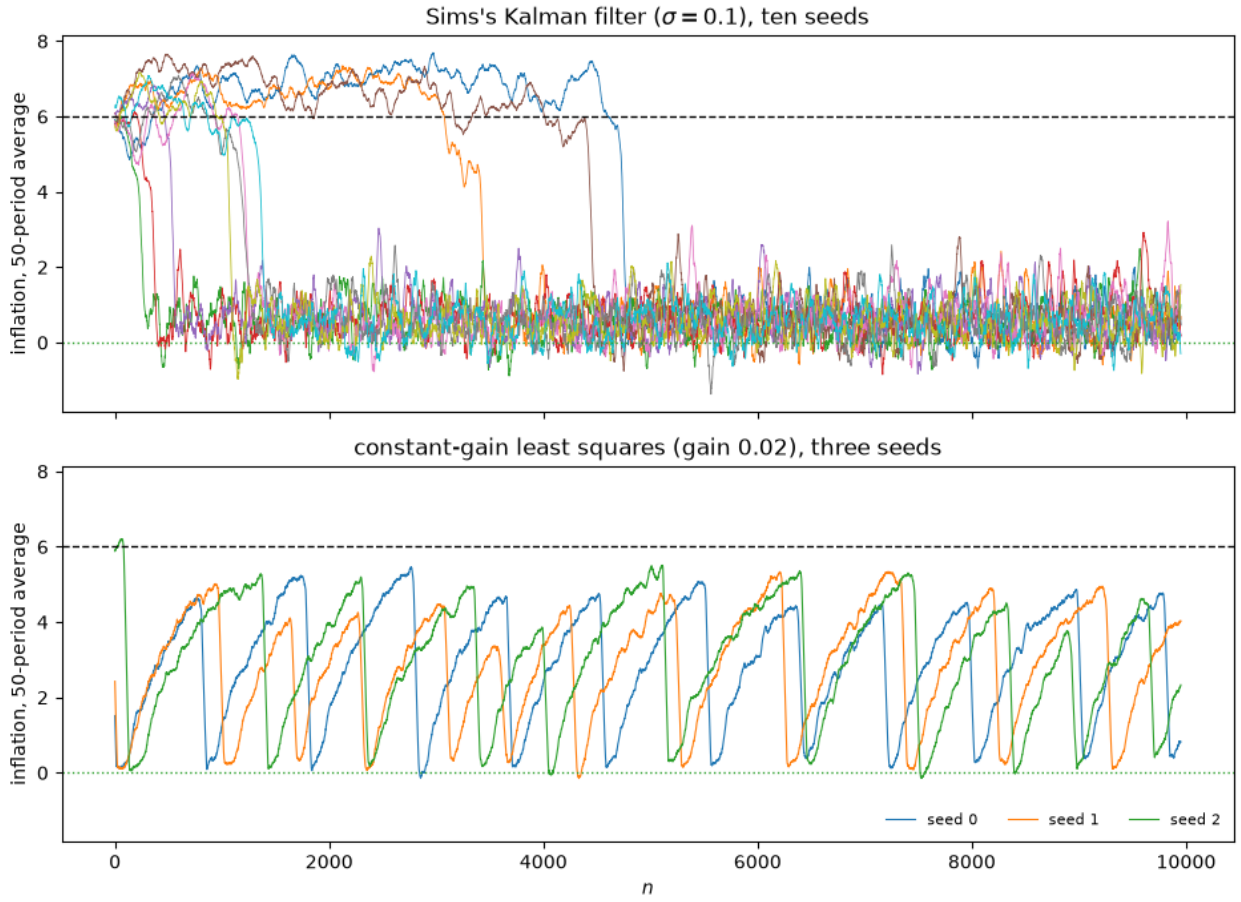


Fig. 105.5: One escape under Sims's Kalman filter, recurrent escapes under constant-gain least squares (50-period moving averages)

With Sims's prior the government escapes once and stays near Ramsey; with constant-gain least squares the mean dynamics pull it back toward Nash after every escape, the sawtooth of Sargent [1999] and CWS.

With Sims's prior and $\sigma = \sigma_1$, the government never escapes in 10,000 periods.

A decreasing gain gives a third route to an escape that sticks, for a different reason: least squares then converges to the self-confirming equilibrium asymptotically, but in a finite sample the falling gain slows the return (*Least squares learning: a decreasing gain*).

As Sargent and Williams [2005] put it, one can read the difference between Sims's results and theirs in two equivalent ways: either Sims allowed too much parameter drift to permit convergence, or he did not let the government attribute enough variation to its regression error.

Here the two readings are exactly equivalent, because the filter (105.5) depends on ε and σ only through ε/σ; Exercise 105.10.6 draws the econometric implication.

105.9 Conclusion

Free parameters are dangerous, as Arthur Goldberger and Robert Lucas warned.

This lecture's government carries free parameters in the covariance matrix V of the drift in its beliefs.

But those parameters buy something.

The empirical sequel Sargent *et al.* [2006] gives the government the Kalman filter of this lecture, estimates V from post-war U.S. data, and so reverse-engineers a sequence of policy makers' evolving subjective models of the Phillips curve that rationalize the actual rise and fall of American inflation — the vindication story of *The Rise and Fall of U.S. Inflation*, made quantitative.

Its authors draw a guarded lesson, though: their estimates "teach us to deemphasize the empirical relevance of both the mean dynamics and the escape dynamics" and to focus instead on the short-term impact of shocks on government beliefs, and their estimated gap between Nash and Ramsey inflation is small, which supports Alan Blinder's skepticism about whether that gap matters much for policy.

Related evidence that the reduced-form coefficients and the shock volatilities both drifted appears in Cogley and Sargent [2005] (*Drifts and Volatilities*), whose model holds the covariance of the coefficient drift constant.

Estimating V is possible because a Kalman filter with a fixed prior has a gain that never vanishes, so the drift parameters stay identified by the whole sample; *The econometricians did return the compliment* tells how econometricians came to exploit this, and Exercise **??** and Exercise 105.10.6 pose stripped-down versions of the estimation problem that Chung [1990] first solved.

105.9.1 Are the estimated beliefs realistic?

That empirical success came with a challenge that is really a challenge about the *prior*.

To fit the data, Sargent *et al.* [2006] estimated a large drift covariance V — a government so open to new data that its beliefs about the monetary transmission mechanism lurch from month to month.

Primiceri [2006], Sims [2007], and, candidly, Sargent [2008] in his presidential address objected that such volatile beliefs are *unrealistic*: the imputed government forecasts unemployment poorly and holds views no real central bank would.

Carboni and Ellison [2009] answer the objection by disciplining the prior with data the Federal Reserve actually produced.

They re-estimate the model subject to the requirement that its beliefs reproduce the unemployment forecasts published in the Fed's *Greenbook* — imposing on the learning model a cross-equation restriction of the kind familiar from rational-expectations econometrics.

Imposing it shrinks the estimated drift variance of the constant term about fiftyfold, from 30.57 to 0.63, and removes the wild belief swings, yet leaves the low-frequency conquest story intact: a *stable* evolution of Federal Reserve beliefs still explains the rise and fall of inflation.

The volatile beliefs, they conclude, "appear to matter only for 'overfitting' high frequency inflation movements", which Sargent [1999] never set out to explain.

The restricted model's perceived tradeoff peaks at about a tenth of its unrestricted size, so the Greenbook evidence supports the account of Sargent [1999] more than the estimates of Sargent *et al.* [2006].

The moral for this lecture is direct: the prior V is not a nuisance parameter to be maximized over freely — it is an economic object, and pinning it down with independent evidence on what policy makers actually believed is what makes the vindication story credible.

The broader message is that *how* an adaptive government learns — the prior it brings to the drift in its own beliefs — is not a technical detail.

It determines whether the economy hovers near Nash with recurrent escapes, cycles between Nash and Ramsey, or escapes to Ramsey and stays there.

The next lecture, *Econometric Policy Evaluation*, takes the adaptive models to data, as chapter 9 of Sargent [1999] did following Chung [1990], by estimating them by maximum likelihood on U.S. data.

Drifts and Volatilities and *Drifts and Volatilities after 2000* then fit a drifting-coefficient, stochastic-volatility VAR to post-war data and ask how much of the Great Inflation was drifting beliefs and how much was simply bad luck, and *The Lost Conquest: Fed Policy in the 2020s* closes the suite by carrying these tools into the present, to rationalize the Federal Reserve's response to the inflation of the 2020s.

105.10 Exercises

Exercise 105.10.1

The learning cycle above tightened the prior on the *slope* coefficient.

Tightening the prior on the *intercept* coefficient instead — using

$$V(\lambda) = \begin{bmatrix} \lambda V_{11}^* & \sqrt{\lambda} V_{12}^* \\ \sqrt{\lambda} V_{12}^* & V_{22}^* \end{bmatrix}$$

— does *not* destabilize the self-confirming equilibrium.

1. Verify this by plotting the maximum real part of the eigenvalues of $\bar{P}(\lambda)\, \partial \bar{g}/\partial \gamma$ against λ for this intercept-tightening family, and confirm it stays negative.
2. Sargent and Williams [2005] find that this prior nonetheless changes the *speed* of escapes. With $\varepsilon = 0.05$, simulate 100 paths from the self-confirming equilibrium under the baseline prior and under their intercept-tightened prior with $\lambda = 0.63$. For each path, record the first period in which beliefs are farther from $\bar{\gamma}$ than $\sqrt{26}$, the distance from the self-confirming equilibrium to Belief 2. Are escapes faster or slower under the tighter prior?

Solution

Here is the stability calculation.

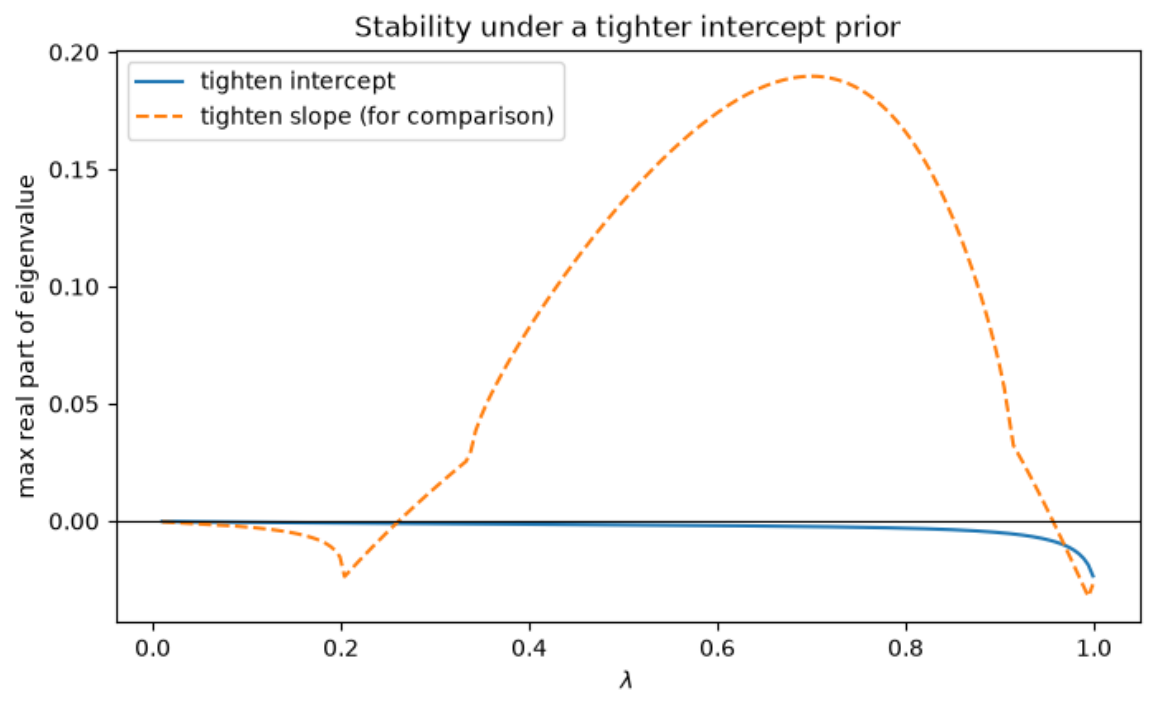
```
def V_tighten_intercept(λ, V_star):
    V = V_star.copy()
```

```
    V[0, 0] = λ * V_star[0, 0]
    V[0, 1] = V[1, 0] = np.sqrt(λ) * V_star[0, 1]
    return V

max_re_int = [eig_PJ(ε**2 * V_tighten_intercept(λ, V_star)).real.max()
              for λ in λ_grid]

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(λ_grid, max_re_int, label='tighten intercept')
ax.plot(λ_grid, max_re, ls='--', label='tighten slope (for comparison)')
ax.axhline(0, color='k', lw=0.8)
ax.set_xlabel(r'$\lambda$')
ax.set_ylabel('max real part of eigenvalue')
ax.set_title('Stability under a tighter intercept prior')
ax.legend()
plt.show()

print(f"largest value over the grid: {max(max_re_int):.2e}")
```




```
largest value over the grid: -1.58e-04
```

Tightening the intercept prior keeps the maximum real part negative for all λ: the self-confirming equilibrium stays stable, so there is no learning cycle.

Now the escape times.

The function below stops each path as soon as beliefs leave the ball.

```
def first_exit(model, V, σ, radius, seed, T_max=50_000):
    "First period in which |y - y_sce| > radius under the exact Kalman filter."
    rng = np.random.default_rng(seed)
    u, σ1, σ2 = model.u, model.σ1, model.σ2
    y_bar = np.array([2 * u, -1.0])
    y, P, W = y_bar.copy(), solve_riccati(V, model.M(y_bar), σ), V / σ**2
    for n in range(T_max):
```

```
        w1, w2 = rng.standard_normal(2)
        π = model.x(y) + σ2 * w2
        U = u - σ2 * w2 + σ1 * w1
        Φ = np.array([1.0, π])
        PΦ = P @ Φ
        denom = 1 + Φ @ PΦ
        y = y + PΦ / denom * (U - Φ @ y)
        P = P - np.outer(PΦ, PΦ) / denom + W
        if np.linalg.norm(y - y_bar) > radius:
            return n + 1
    return np.nan

radius = np.sqrt(26)
for name, V_hat in [("baseline V*", V_star),
                    ("tighter intercept, λ = 0.63", V_tighten_intercept(0.63, V_
 ↪star))]:
    times = np.array([first_exit(model, ε**2 * V_hat, model.σ, radius, s)
                      for s in range(100)])
    print(f"{name:28s}: mean escape time {np.nanmean(times):6.1f} periods "
          f"(median {np.nanmedian(times):.0f}, censored {np.isnan(times).sum()})")
```

```
baseline V*                  : mean escape time   23.0 periods (median 20,␣
 ↪censored 0)
tighter intercept, λ = 0.63 : mean escape time  259.5 periods (median 220,␣
 ↪censored 0)
```

Escapes are much *slower* under the tighter intercept prior.

The baseline time is close to the 22.9 periods that Sargent and Williams [2005] report in their Table 2 at this gain; the intercept-tightened time is of the same order as their 167.3 periods, many times the baseline.

A government that is surer that its intercept is not drifting puts less weight on the forecast errors that would carry its beliefs along the escape route, and Sargent and Williams [2005] (section 6.3) show that priors that make $\bar{P}^{-1}$ larger raise the rate $\bar{S}$ that governs how rare escapes are.

So this prior changes the escape dynamics substantially, without overturning the mean dynamics.

Exercise 105.10.2

The instantaneous escape direction is the dominant eigenvector of the prior covariance $\hat{V}$.

1. Compute the escape direction, and the terminal beliefs implied by following it in a straight line until the slope reaches 0, for the baseline prior V^* and for the intercept-tightened prior of Exercise 105.10.1 with $\lambda = 0.63$, the case of Fig. 9 of Sargent and Williams [2005]. Why would the slope-tightened prior $V(0.5)$ be a poor choice for this comparison?
2. With $\varepsilon = 0.05$, simulate five escapes under each prior, stopping each path when the slope first reaches 0, and plot the belief paths in the (intercept, slope) plane together with the two instantaneous directions.
3. Where do the escapes end, and how well does the instantaneous direction predict their route?

Solution

The slope-tightened prior $V(0.5)$ lies inside the unstable range of Fig. 105.1, so it has no stable self-confirming equilibrium to escape from.

The intercept-tightened prior keeps the equilibrium stable (Exercise 105.10.1).

```
def escape_direction(V):
    "Dominant eigenvector of V, oriented toward a higher slope, and its terminal␣
 ↪point."
    w, vecs = np.linalg.eigh(V)
    v = vecs[:, np.argmax(w)]
    v = v / np.linalg.norm(v)
    if v[1] < 0:
        v = -v
    return v, y_sce + v / v[1]

priors = {"baseline": V_star,
          "tighter intercept (λ = 0.63)": V_tighten_intercept(0.63, V_star)}
for name, V_hat in priors.items():
    v, term = escape_direction(V_hat)
    stable = eig_PJ(ε**2 * V_hat).real.max() < 0
    print(f"{name:30s}: direction {v.round(3)}, straight-line terminal "
          f"{term.round(2)}, SCE stable: {stable}")
```

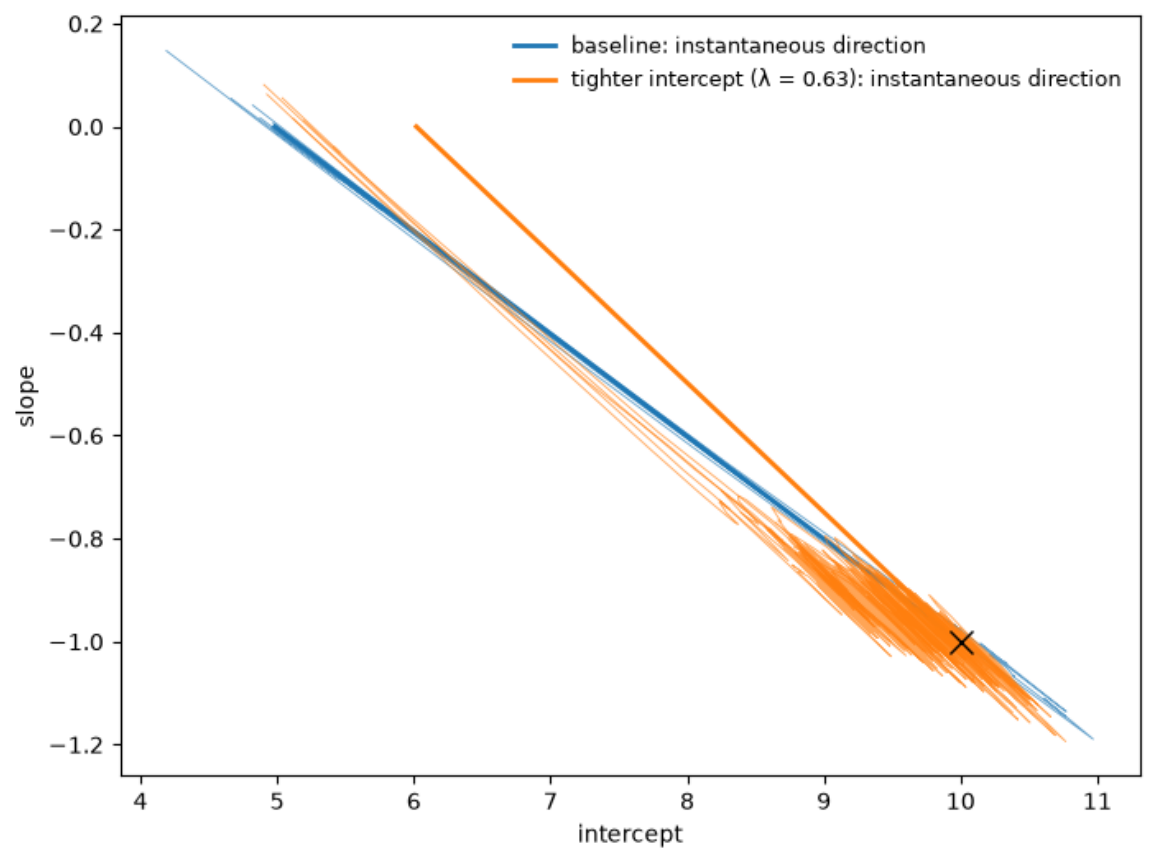
```
baseline                       : direction [-0.981  0.195], straight-line␣
 ↪terminal [4.98 0.  ], SCE stable: True
tighter intercept (λ = 0.63)  : direction [-0.97   0.244], straight-line␣
 ↪terminal [6.02 0.  ], SCE stable: True
```

The directions differ only by a few degrees, but because the escape covers a long distance, the straight-line terminal intercepts differ by about one unit.

Now the simulated escapes.

```
fig, ax = plt.subplots(figsize=(8, 6))
for (name, V_hat), color in zip(priors.items(), ('C0', 'C1')):
    v, term = escape_direction(V_hat)
    ax.plot([y_sce[0], term[0]], [y_sce[1], term[1]], color=color, lw=2,
            label=f'{name}: instantaneous direction')
    ends = []
    for s in range(5):
        _, _, _, beliefs = kalman_path(model, ε**2 * V_hat, model.σ, 5_000, s)
        n_end = np.argmax(beliefs[:, 1] >= 0)
        ax.plot(beliefs[:n_end + 1, 0], beliefs[:n_end + 1, 1], color=color,
                lw=0.6, alpha=0.7)
        ends.append((n_end, beliefs[n_end, 0]))
    print(f"{name:30s}: (period, intercept) when the slope reaches 0: "
          f"{[(int(n), round(a, 2)) for n, a in ends]}")
ax.plot(*y_sce, 'kx', ms=10)
ax.set_xlabel('intercept')
ax.set_ylabel('slope')
ax.legend(frameon=False, fontsize=9)
plt.show()
```

```
baseline                       : (period, intercept) when the slope reaches 0:␣
 ↪[(10, np.float64(4.87)), (35, np.float64(4.66)), (13, np.float64(4.82)), (33,␣
 ↪np.float64(4.97)), (8, np.float64(4.19))]
tighter intercept (λ = 0.63)  : (period, intercept) when the slope reaches 0:␣
 ↪[(142, np.float64(5.19)), (767, np.float64(4.92)), (193, np.float64(5.04)),␣
 ↪(287, np.float64(5.12)), (267, np.float64(4.9))]
```



Under the baseline prior the simulated escapes run close to the straight line, as in Fig. 7 of Sargent and Williams [2005], and end near Belief 2, $(u, 0)$.

Under the tighter intercept prior the escapes take much longer to start: beliefs first wander around the self-confirming equilibrium, and when they do escape they run along a line close to the baseline's, not along their own instantaneous direction, ending at slope 0 with an intercept close to the baseline's rather than at the straight-line terminal point.

So the instantaneous direction predicts the route well only in the baseline case, where, as $\sigma \to 0$, the mean dynamics themselves keep pushing beliefs along it once an escape has started (section 6.2 of Sargent and Williams [2005]); for other priors the whole dynamic control problem matters.

In both cases the escape ends at a slope of 0, where the government perceives no exploitable tradeoff and (105.3) sets inflation to zero whatever the intercept.

The destination is robust; the route and its timing depend on the shape of the prior.

Exercise 105.10.3

Sargent and Williams [2005] (section 3.4) stress that constant-gain recursive least squares is only *asymptotically* a special case of the Kalman filter with $V = \varepsilon^2 V^*$.

With $\varepsilon = 0.05$ and $V = \varepsilon^2 V^*$, start both algorithms' mean dynamics at the Ramsey belief $\gamma_0 = (u, 0)$, with the Kalman P at its Riccati value for $M(\gamma_0)$ and the RLS moment matrix at $R = M(\gamma_0)$.

The RLS mean dynamics are

$$\dot{\gamma} = \varepsilon R^{-1} \bar{g}(\gamma), \qquad \dot{R} = \varepsilon \left[M(\gamma) - R \right].$$

Integrate both systems for 800 periods and plot the intercept, the slope, and inflation $x(\gamma)$.

Which algorithm returns to the self-confirming equilibrium faster, and do they end at the same place?

Solution

The Kalman mean dynamics are the function `mean_dynamics` from the lecture.

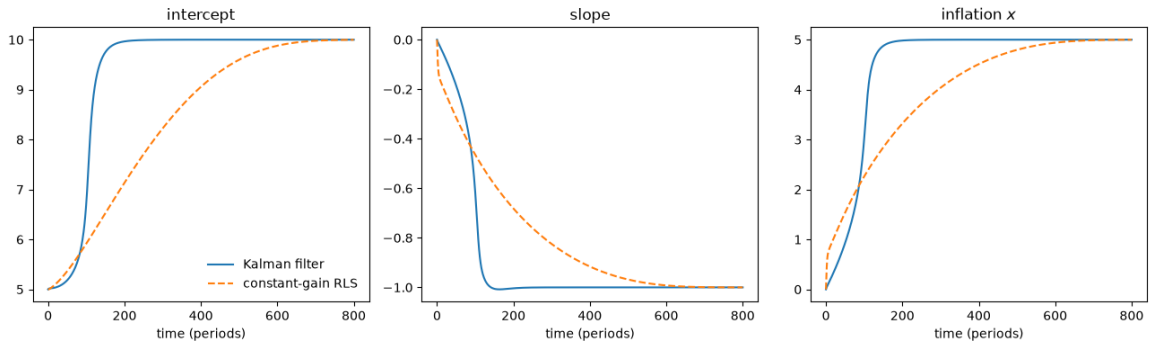
```
def rls_mean_dynamics(t, z):
    "Mean dynamics of constant-gain RLS for (γ, R), stacked into one vector."
    γ, R = z[:2], z[2:].reshape(2, 2)
    γ_dot = ε * np.linalg.solve(R, model.g_bar(γ))
    R_dot = ε * (model.M(γ) - R)
    return np.concatenate([γ_dot, R_dot.ravel()])

γ0 = np.array([model.u, 0.0])                    # Belief 2 (Ramsey)
M0 = model.M(γ0)
V = ε**2 * V_star
t_eval = np.linspace(0, 800, 801)

kal = solve_ivp(mean_dynamics, [0, 800],
                np.concatenate([γ0, solve_riccati(V, M0, model.σ).ravel()]),
                args=(V, model.σ), t_eval=t_eval, rtol=1e-8, atol=1e-10)
rls = solve_ivp(rls_mean_dynamics, [0, 800], np.concatenate([γ0, M0.ravel()]),
                t_eval=t_eval, rtol=1e-8, atol=1e-10)

fig, axes = plt.subplots(1, 3, figsize=(13, 4))
for out, label, ls in [(kal, 'Kalman filter', '-'), (rls, 'constant-gain RLS', '--
↪')]:
    a, b = out.y[0], out.y[1]
    for ax, series in zip(axes, (a, b, model.x((a, b)))):
        ax.plot(out.t, series, ls=ls, label=label)
for ax, title in zip(axes, ('intercept', 'slope', 'inflation $x$')):
    ax.set_title(title)
    ax.set_xlabel('time (periods)')
axes[0].legend(frameon=False)
plt.tight_layout()
plt.show()

for out, label in [(kal, 'Kalman filter'), (rls, 'constant-gain RLS')]:
    x = model.x((out.y[0], out.y[1]))
    print(f"{label:18s}: x first exceeds 0.8u at period {out.t[np.argmax(x > 0.8 *␣
↪model.u)]:.0f}; "
          f"beliefs at period 800 = {out.y[:2, -1].round(3)}")
```




```
Kalman filter     : x first exceeds 0.8u at period 111; beliefs at period 800 =␣
↪[10. -1.]
constant-gain RLS : x first exceeds 0.8u at period 294; beliefs at period 800 =␣
↪[ 9.998 -1.   ]
```

Both algorithms end at the self-confirming equilibrium $(2u, -1)$, but the Kalman filter gets there much faster, with an abrupt reversal of the slope rather than a gradual one, as in Figs. 1 and 2 of Sargent and Williams [2005].

The difference comes from how the gain matrices respond to the changing moments of the data: in the Kalman filter P responds through the term $PM(\gamma)P$, while the RLS moment matrix R only follows $M(\gamma)$ with a lag.

Because the two algorithms share their rest point and its local stability, they agree asymptotically; their transient responses, which matter after every escape, differ.

Exercise 105.10.4

This exercise compares the learning cycle with recurrent escapes, as Sargent and Williams [2005] do in their Fig. 6.

1. With $\varepsilon = 0.05$, simulate 2,000 periods of the exact Kalman filter from the self-confirming equilibrium with the slope-tightened prior $V(0.7)$, and plot inflation together with the mean-dynamics path $x(t)$ computed in the lecture. Repeat the simulation with the baseline prior V^*, under which the self-confirming equilibrium is stable.
2. For $\varepsilon \in \{0.05, 0.02, 0.005\}$ and both priors, count disinflations (the 20-period average of inflation falling below 1, after having been above 4) per 100 units of εn, averaging over three seeds.
3. What distinguishes the learning cycle from recurrent escapes at these gains? What would distinguish them as $\varepsilon \to 0$?

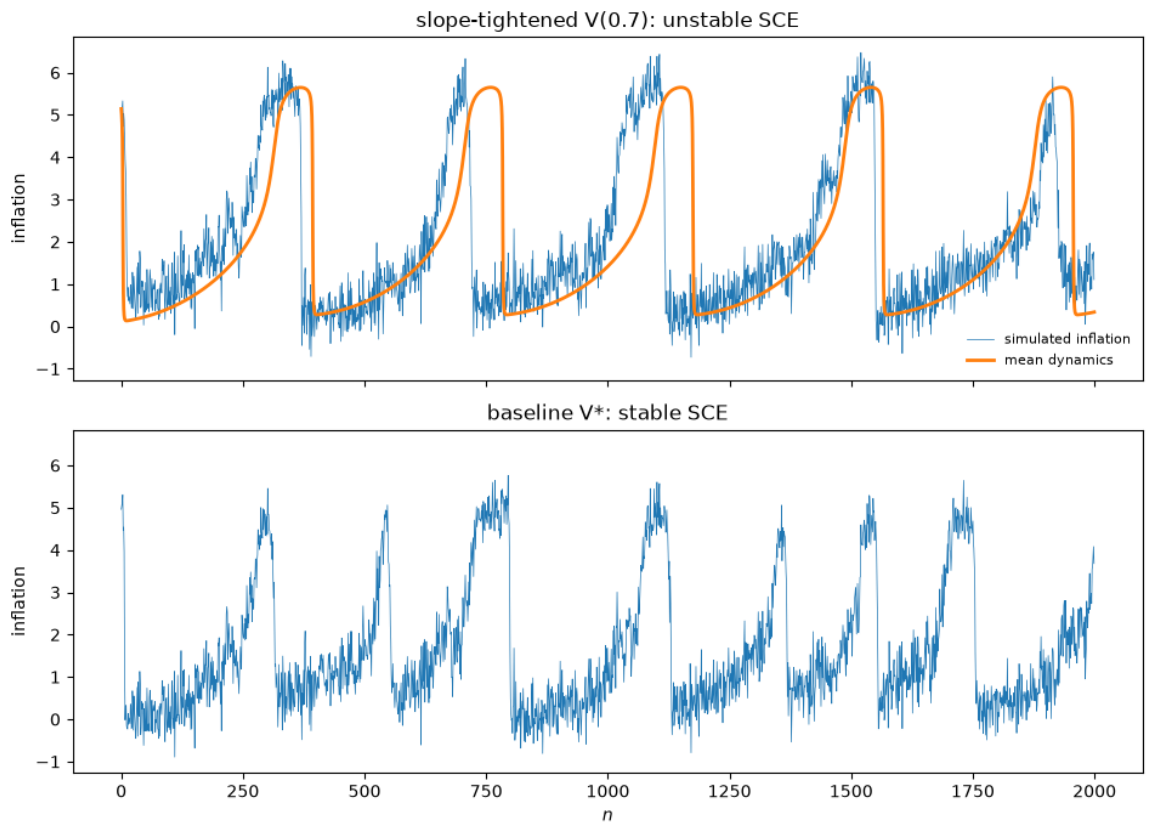
Solution

Here are the two simulations.

```
priors = {"slope-tightened V(0.7): unstable SCE": V_tighten_slope(0.7, V_star),
          "baseline V*: stable SCE": V_star}

fig, axes = plt.subplots(2, 1, figsize=(9.5, 7), sharex=True, sharey=True)
for ax, (name, V_hat) in zip(axes, priors.items()):
    π = kalman_path(model, ε**2 * V_hat, model.σ, 2_000, 0)[0]
    ax.plot(π, lw=0.5, label='simulated inflation')
    ax.set_title(name)
    ax.set_ylabel('inflation')
keep = t_path <= 2_000
```

```
axes[0].plot(t_path[keep], x_path[keep], lw=2, label='mean dynamics')
axes[0].legend(frameon=False, loc='lower right', fontsize=8)
axes[1].set_xlabel('$n$')
plt.tight_layout()
plt.show()
```



Both simulated series show recurrent disinflations.

Under the slope-tightened prior the simulation tracks the deterministic cycle of the mean dynamics closely; under the baseline prior the disinflations are escapes, but they come about as often.

Now the counts.

```
def disinflations_per_100(V_hat, ε_, seeds=range(3)):
    "Disinflations per 100 units of εn under the exact Kalman filter."
    T = int(100 / ε_)
    counts = [escapes_and_returns(kalman_path(model, ε_**2 * V_hat, model.σ, T,␣
 ↪s)[0],
                                  lo=1.0, hi=4.0, k=20)[0] for s in seeds]
    return np.mean(counts)

print(f"{'ε':>6} | {'V(0.7): cycle':>14} | {'V*: escapes':>12}")
for ε_ in (0.05, 0.02, 0.005):
    print(f"{ε_:>6} | {disinflations_per_100(priors['slope-tightened V(0.7): unstable SCE'], ε_):>14.1f}"
          f" | {disinflations_per_100(V_star, ε_):>12.1f}")
```

```
     ε |  V(0.7): cycle |  V*: escapes
  0.05 |            5.3 |          7.7
  0.02 |            6.0 |          7.0
 0.005 |            6.0 |          6.7
```

At every gain in the table, both priors produce disinflations at a similar rate per unit of εn, so the stable prior's escapes are not rare at these gains.

For the cycle that is expected: its period is proportional to $1/\varepsilon$, about 20 units of εn.

For the stable prior it reflects the small large-deviation rate $\bar{S} \approx 3 \times 10^{-4}$ that Sargent and Williams [2005] compute: the mean time to escape grows like $\exp(\bar{S}/\varepsilon)$ only once ε is small relative to $\bar{S}$, and until then it scales roughly like $1/\varepsilon$, just as the period of the cycle does.

The stable prior's rate per unit of εn does fall as ε shrinks, but slowly.

A simulated time series at these gains therefore cannot easily tell a learning cycle from recurrent escapes.

The distinction is asymptotic: as $\varepsilon \to 0$ the cycle's rate per unit of εn stays put, while the stable prior's escapes eventually become exponentially rare.

Exercise 105.10.5

This exercise varies the replication of Sims's example in *Sims's nonconvergence*.

1. The time series in Fig. 14 of Sargent and Williams [2005] appear to start away from the self-confirming equilibrium, at low inflation. Start Sims's specification and the $\sigma = \sigma_1$ variation at the Ramsey belief $(u, 0)$ instead, with P at its Riccati value, and report mean inflation over the first 1,000 and the last 2,000 of 10,000 periods for five seeds.
2. Starting from the self-confirming equilibrium, scale Sims's drift covariance to $V = \varepsilon^2 \times 0.01\, I$ for $\varepsilon \in \{0.2, 0.3, 0.4, 0.6\}$, keeping $\sigma = 0.1$. For each ε, count escapes and returns over ten seeds of 10,000 periods, as in the lecture. How does the history change as the drift grows?

Solution

First, the paths that start at the Ramsey belief.

```
y_ramsey = np.array([sims.u, 0.0])
for name, σ_govt in [("Sims, σ = 0.1", σ_sims), ("σ = σ1", sims.σ1)]:
    early, late = [], []
    for s in range(5):
        π = kalman_path(sims, V_sims, σ_govt, T_sims, s, y0=y_ramsey)[0]
        early.append(π[:1000].mean())
        late.append(π[-2000:].mean())
    print(f"{name:14s}: mean π, first 1,000 periods {np.round(early, 2)}")
    print(f"{'':14s}  mean π, last 2,000 periods  {np.round(late, 2)}")
```

```
Sims, σ = 0.1 : mean π, first 1,000 periods [0.87 0.62 0.58 0.55 0.51]
                mean π, last 2,000 periods  [0.5  0.65 0.69 0.8  0.62]
σ = σ1        : mean π, first 1,000 periods [4.45 4.77 3.16 2.73 3.08]
                mean π, last 2,000 periods  [5.96 6.09 5.45 5.66 5.74]
```

Started at low inflation, Sims's government never converges to the self-confirming equilibrium: inflation stays low throughout.

With $\sigma = \sigma_1$, inflation climbs toward the self-confirming value $u = 6$ during the first thousand periods and then stays near it, which is what Fig. 14 of Sargent and Williams [2005] shows.

Now the scan over the size of the drift.

```
for ε_ in (0.2, 0.3, 0.4, 0.6):
    counts = [escapes_and_returns(
                  kalman_path(sims, ε_**2 * V_sims, σ_sims, T_sims, s)[0], lo,␣
↪hi)[:2]
              for s in seeds]
    n_escape = sum(c[0] > 0 for c in counts)
    n_return = sum(c[1] > 0 for c in counts)
    print(f"ε = {ε_}: paths that escape {n_escape}/10, paths that return {n_return}
↪/10, "
          f"(escapes, returns) by seed {counts}")
```

```
ε = 0.2: paths that escape 0/10, paths that return 0/10, (escapes, returns) by␣
↪seed [(0, 0), (0, 0), (0, 0), (0, 0), (0, 0), (0, 0), (0, 0), (0, 0), (0, 0),␣
↪(0, 0)]
ε = 0.3: paths that escape 8/10, paths that return 5/10, (escapes, returns) by␣
↪seed [(1, 1), (0, 0), (2, 2), (2, 1), (1, 0), (1, 0), (1, 0), (2, 1), (0, 0),␣
↪(3, 3)]
ε = 0.4: paths that escape 10/10, paths that return 1/10, (escapes, returns) by␣
↪seed [(1, 0), (1, 0), (1, 0), (1, 0), (1, 0), (1, 0), (2, 1), (1, 0), (1, 0),␣
↪(1, 0)]
ε = 0.6: paths that escape 10/10, paths that return 0/10, (escapes, returns) by␣
↪seed [(1, 0), (1, 0), (1, 0), (1, 0), (1, 0), (1, 0), (1, 0), (1, 0), (1, 0),␣
↪(1, 0)]
```

The history changes with the size of the drift.

With the smallest drift no path escapes in 10,000 periods: the mean dynamics hold beliefs near the self-confirming equilibrium.

With somewhat larger drift, paths escape and some return, the beginnings of a sawtooth.

With larger drift still, almost every path escapes once and never returns, as in Sims's own case $\varepsilon = 1$ in the lecture.

The two kinds of history in the literature are two ends of a continuum indexed by the drift that the government entertains, relative to the noise it attributes to its regression.

Exercise 105.10.6

Exercise **??** asked an econometrician to estimate the gain of a constant-gain government from the data that its policy generates.

Here the government runs the exact Kalman filter of this lecture with prior $V = \varepsilon^2 V^*$, $\varepsilon = 0.02$, and $\sigma = \sigma_1$.

Simulate 5,000 periods from the self-confirming equilibrium.

Suppose an econometrician knows everything about the economy and the government except ε and the government's σ.

1. Show that the candidate pairs (ε, σ) and $(c\varepsilon, c\sigma)$ imply exactly the same history of targets x_n for any $c > 0$.
2. Compute the Gaussian log likelihood of inflation, $\ell(\varepsilon, \sigma) = -\sum_n (\pi_n - x_n(\varepsilon, \sigma))^2/(2\sigma_2^2)$ up to a constant, on a grid of (ε, σ), and plot its contours.
3. What do the data identify? Relate your answer to the two equivalent interpretations of Sims's example in *Sims's nonconvergence*.

Solution

The government's beliefs are a deterministic function of the data and of the filter's parameters, so the econometrician can rerun the filter for each candidate pair, as in Exercise **??**.

The function below does this for many candidates at once.

```
def reconstruct_targets(π, U, ε_vals, σ_vals):
    """
    Targets x_n implied by the exact Kalman filter with prior (ε² V*, σ),
    reconstructed from data (π, U), for each candidate pair (ε, σ).
    """
    K = len(ε_vals)
    γ = np.tile(γ_sce, (K, 1))
    P = np.array([solve_riccati(e**2 * V_star, M_sce, s)
                  for e, s in zip(ε_vals, σ_vals)])
    W = np.array([e**2 * V_star / s**2 for e, s in zip(ε_vals, σ_vals)])
    targets = np.empty((len(π), K))
    for n in range(len(π)):
        targets[n] = -γ[:, 0] * γ[:, 1] / (1 + γ[:, 1]**2)
        Φ = np.array([1.0, π[n]])
        PΦ = P @ Φ
        denom = 1 + PΦ @ Φ
        γ = γ + PΦ / denom[:, None] * (U[n] - γ @ Φ)[:, None]
        P = P - PΦ[:, :, None] * PΦ[:, None, :] / denom[:, None, None] + W
    return targets

ε_true, σ_true = 0.02, model.σ1
π_data, U_data, x_data, _ = kalman_path(model, ε_true**2 * V_star, σ_true, 5_000,␣
↪1)

check = reconstruct_targets(π_data, U_data, [ε_true, 2 * ε_true, ε_true],
                            [σ_true, 2 * σ_true, σ_true / 2])
print(f"largest error at the true (ε, σ):          {np.max(np.abs(check[:, 0] - x_
↪data)):.1e}")
print(f"largest difference, (2ε, 2σ) vs (ε, σ):    {np.max(np.abs(check[:, 1] -␣
↪check[:, 0])):.1e}")
print(f"largest difference, (ε, σ/2) vs (ε, σ):    {np.max(np.abs(check[:, 2] -␣
↪check[:, 0])):.1e}")
```

```
largest error at the true (ε, σ):          0.0e+00
largest difference, (2ε, 2σ) vs (ε, σ):    0.0e+00
largest difference, (ε, σ/2) vs (ε, σ):    1.5e+00
```

Doubling both ε and σ leaves the targets unchanged, while changing σ alone does not.

The reason is that the filter (105.5) uses V and σ only through $\sigma^{-2}V = (\varepsilon/\sigma)^2 V^*$, and the Riccati value of P_0 depends on them in the same way.

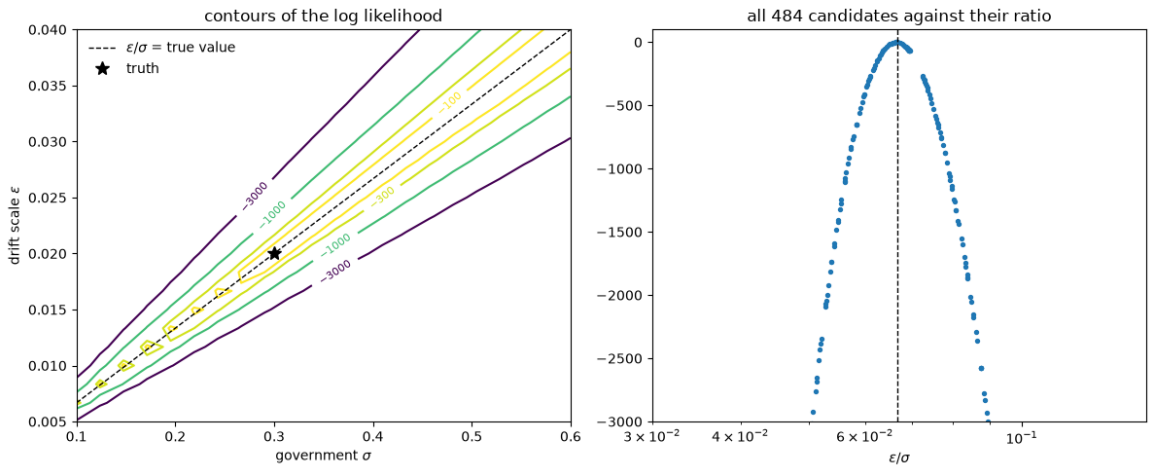
Now the likelihood on a grid.

```
ε_grid = np.linspace(0.005, 0.04, 22)
σ_grid = np.linspace(0.1, 0.6, 22)
EE, SS = np.meshgrid(ε_grid, σ_grid)
targets = reconstruct_targets(π_data, U_data, EE.ravel(), SS.ravel())
loglik = -np.sum((π_data[:, None] - targets)**2, axis=0) / (2 * model.σ2**2)
loglik = (loglik - loglik.max()).reshape(EE.shape)

fig, axes = plt.subplots(1, 2, figsize=(12, 5))
cs = axes[0].contour(SS, EE, loglik, levels=[-3000, -1000, -300, -100])
axes[0].clabel(cs, fontsize=8)
s_line = np.linspace(0.1, 0.6, 2)
axes[0].plot(s_line, (ε_true / σ_true) * s_line, 'k--', lw=1,
             label=r'$\varepsilon / \sigma$ = true value')
axes[0].plot(σ_true, ε_true, 'k*', ms=10, label='truth')
axes[0].set_xlabel(r'government $\sigma$')
axes[0].set_ylabel(r'drift scale $\varepsilon$')
axes[0].set_title('contours of the log likelihood')
axes[0].legend(frameon=False, loc='upper left')

axes[1].scatter((EE / SS).ravel(), loglik.ravel(), s=8)
axes[1].axvline(ε_true / σ_true, color='k', ls='--', lw=1)
axes[1].set_xscale('log')
axes[1].set_xlim(0.03, 0.15)
axes[1].set_ylim(-3000, 100)
axes[1].set_xlabel(r'$\varepsilon / \sigma$')
axes[1].set_title(f'all {loglik.size} candidates against their ratio')
plt.tight_layout()
plt.show()

ratio_hat = (EE.ravel() / SS.ravel())[np.argmax(loglik)]
print(f"true ε/σ = {ε_true / σ_true:.4f}; ratio at the grid maximum = {ratio_hat:.
↪4f}")
```




```
true ε/σ = 0.0667; ratio at the grid maximum = 0.0664
```

The contours of the likelihood are rays through the origin, and plotted against ε/σ the log likelihoods of all the candidates fall on a single curve: the data pin down ε/σ, the government's effective gain, but say nothing about ε and σ separately.

That is the econometric counterpart of the two equivalent interpretations of Sims's example: "too much drift" and "too little error variance" are observationally equivalent descriptions of the same learning rule.

Separating them would require evidence on how *uncertain* the government says it is, for example its predictive variance $\sigma^2(1 + \Phi_n^\top P_n \Phi_n)$, and not only on what it does or forecasts.

What the data do identify, they keep identifying as the sample grows, because the filter's gain never vanishes and so keeps shaping beliefs (*The econometricians did return the compliment*).

CHAPTER

SIX

ECONOMETRIC POLICY EVALUATION

Contents

- *Econometric Policy Evaluation*
  - *Overview*
  - *The model*
  - *The likelihood function*
  - *Data*
  - *Beliefs and recommendations*
  - *Estimates*
  - *Recommendations and misses*
  - *Misses favor vindication*
  - *What came later*
  - *Exercises*

106.1 Overview

This lecture follows chapter 9 of Sargent [1999], "Econometric Policy Evaluation".

The preceding lectures built the adaptive governments of the book's chapter 8 and studied them by simulation (*Adaptive Learning and Escape Dynamics*, *Escaping Nash Inflation*, *Priors, Escapes, and Learning Cycles*).

Chapter 9 takes the same models to U.S. data, estimating them by maximum likelihood under both the classical and the Keynesian identification schemes, following Chung [1990].

It is the book's only structural confrontation of the vindication story of *The Rise and Fall of U.S. Inflation* with data.

The question is counterfactual: what would a policy adviser of the 1960s and 1970s have recommended if he had fit a distributed-lag Phillips curve, discounted old data, and fed the estimates into the Phelps problem of *Adaptive Expectations and the Phelps Problem*?

The estimated model answers directly, because in it the fitted value of inflation *is* the government's recommendation.

Once the gain and the discount factor are fixed, the likelihood has no free parameters describing beliefs: we compute the government's beliefs and recommendations from the data, as the government itself would have.

We reproduce the book's main findings:

- both models fit inflation badly, and the classical model fits much worse;
- the misses are systematic, with recommended inflation below actual inflation through the 1970s, which is the sense in which "the pattern of misses from the estimated models favors vindication";
- the Keynesian government's estimated Phillips curve satisfied the induction hypothesis by the mid-1970s, and its plans called for low inflation years before Volcker arrived;
- the classical government recommended low inflation all along, which is why its recommendations "would have been even more timely."

The book does not report its lag lengths or how it converted its data to quarterly frequency, and we say where these choices matter.

Let's start with some imports.

```
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from scipy.linalg import solve_discrete_are
from scipy.optimize import minimize
from scipy.signal import lfilter
from scipy.stats import chi2
```

106.2 The model

The model combines the actual Phillips curve of *Self-Confirming Equilibria* with the adaptive government of *Adaptive Learning and Escape Dynamics*; the notation is that of *A note on notation*.

106.2.1 The actual Phillips curve

With U_t unemployment, y_t inflation and x_t the public's expectation of inflation, the economy obeys a natural-rate Phillips curve with a geometric distributed lag in surprise inflation $v_{2t} = y_t - x_t$ and a serially correlated disturbance:

$$U_t = U^* - \frac{\theta}{1-\rho_2 L}\, v_{2t} + \frac{1}{1-\rho_1 L}\, v_{1t}, \tag{106.1}$$

where L is the lag operator, $|\rho_1| < 1$ and $|\rho_2| < 1$.

Multiplying through by $(1-\rho_1 L)(1-\rho_2 L)$ gives the form that the book numbers (48):

$$U_t = U^*(1-\rho_1)(1-\rho_2) + (\rho_1+\rho_2)U_{t-1} - \rho_1\rho_2 U_{t-2} + (1-\rho_2 L)v_{1t} - \theta(1-\rho_1 L)v_{2t}. \tag{106.2}$$

The shock vector $v_t = (v_{1t}, v_{2t})^\top$ is Gaussian white noise with an unrestricted covariance matrix V_v.

Chapters 7 and 8 of the book set $\rho_1 = \rho_2 = 0$ and made V_v diagonal to make theoretical points; chapter 9 frees both.

106.2.2 The government

Let $X_{t-1} = \begin{bmatrix} U_{t-1} & \cdots & U_{t-m_U} & y_{t-1} & \cdots & y_{t-m_y} & 1 \end{bmatrix}^\top$ collect information dated $t-1$ and earlier.

As in *Adaptive Expectations and the Phelps Problem*, the government fits one of two reduced-form Phillips curves, which differ only in their direction of fit:

$$\text{classical:} \quad U_t = \gamma^\top X_{C,t} + \varepsilon_{C,t}, \qquad \text{Keynesian:} \quad y_t = \beta^\top X_{K,t} + \varepsilon_{K,t}, \tag{106.3}$$

where $X_{C,t} = \begin{bmatrix} y_t & X_{t-1}^\top \end{bmatrix}^\top$ and $X_{K,t} = \begin{bmatrix} U_t & X_{t-1}^\top \end{bmatrix}^\top$.

It re-estimates its curve every period by constant-gain recursive least squares, which for the classical scheme is

$$\begin{aligned} \gamma_t &= \gamma_{t-1} + g\, R_{XC,t}^{-1} X_{C,t} \left(U_t - \gamma_{t-1}^\top X_{C,t}\right), \\ R_{XC,t} &= R_{XC,t-1} + g\left(X_{C,t} X_{C,t}^\top - R_{XC,t-1}\right), \end{aligned} \tag{106.4}$$

with a constant gain $g = 1 - \lambda$; the Keynesian scheme updates β_t and $R_{XK,t}$ in the same way.

A Keynesian government converts its estimate into the classical form by inverting the Phillips curve: if β_1 is the coefficient on U_t and β_{-1} collects the rest, then $\gamma(\beta) = \begin{bmatrix} \beta_1^{-1} & -\beta_{-1}^\top / \beta_1 \end{bmatrix}^\top$.

Each period the government acts as if its current estimate will govern the Phillips curve forever and sets inflation by solving the Phelps problem, minimizing $\mathbb{E}\sum_t \delta^t (U_t^2 + y_t^2)$ subject to its estimated curve, which yields the rule

$$y_t = h(\gamma_{t-1}) X_{t-1} + v_{2t}, \tag{106.5}$$

where γ_{t-1} is the classical estimate or $\gamma(\beta_{t-1})$, and v_{2t} is a control error.

Following Sims [1988] and Chung [1990], the public knows the rule, so $x_t = h(\gamma_{t-1}) X_{t-1}$, and the control error v_{2t} is also the inflation surprise in (106.1).

106.3 The likelihood function

To simulate the model, we would draw shocks $\{v_t\}$ and compute histories of (y_t, U_t) and of the government's beliefs.

To estimate it, we reverse the procedure: starting from initial beliefs, we use the data $\{y_t, U_t\}$ to solve for the shocks that would have produced them,

$$\begin{aligned} v_{2t} &= y_t - h(\gamma_{t-1}) X_{t-1}, \\ v_{1t} - \theta v_{2t} &= U_t - (\rho_1 + \rho_2) U_{t-1} + \rho_1 \rho_2 U_{t-2} + \rho_2 v_{1,t-1} - \rho_1 \theta v_{2,t-1} - U^*(1-\rho_1)(1-\rho_2), \end{aligned} \tag{106.6}$$

while updating γ_t (or β_t) by (106.4); these are the book's equations (105).

Given the past, the map from v_t to (U_t, y_t) is triangular with unit diagonal, so its Jacobian is one and the Gaussian log likelihood of a sample of length T is

$$\log L = -T \log 2\pi - \frac{T}{2} \log\det V_v - \frac{1}{2} \sum_{t=1}^{T} v_t^\top V_v^{-1} v_t. \tag{106.7}$$

As the book suggests, we concentrate out V_v: its maximizer is $\hat{V}_v = T^{-1} \sum_t v_t v_t^\top$, and then $\log L = -T \log 2\pi - \frac{T}{2} \log\det \hat{V}_v - T$.

> **Note**
>
> The book's appendix to chapter 9 writes the log of a sum of densities where it means a sum of log densities, and gives the constant as $-\frac{T}{2}\log 2\pi$, which is right for a scalar observation; a bivariate density has $-T\log 2\pi$.
>
> Neither misprint affects the estimates, but the second shifts the level of $\log L$.

The book estimates $(\theta, U^*, \rho_1, \rho_2, V_v)$, and it fixes the discount factor at $\delta = .999$ and the forgetting factor at $\lambda = .985$, so the gain is $g = .015$.

Notice what is *not* on the list: nothing describes beliefs.

With λ, δ and the initial beliefs fixed, the government's beliefs γ_t, its recommendations $h(\gamma_{t-1})X_{t-1}$ and therefore its misses v_{2t} are computed from the data alone, whatever the values of the parameters being estimated.

This has two consequences.

1. Because $\det \hat{V}_v$ equals the mean square of v_{2t} times the mean square of the part of v_{1t} not explained by v_{2t}, the parameters $(\theta, U^*, \rho_1, \rho_2)$ are estimated from the unemployment equation, taking the government's misses as given.
2. The comparison between the two identification schemes is mostly a comparison of the size of their misses, that is, of how far the recommendations were from what actually happened.

This is a cross-equation restriction of the kind familiar from rational expectations econometrics: the government's rule is pinned down by its econometrics and the Phelps problem, not left free.

106.4 Data

Chapter 9 uses the measures of chapter 1: the consumer price index (all items) and the unemployment rate of white men aged 20 and over, the series plotted in *The Rise and Fall of U.S. Inflation*.

Although chapter 1 plots monthly data, chapter 9's estimates appear to be quarterly:

- the sample period is dated in quarters (it is printed as "1995I–1997I", evidently a misprint for 1955I–1997I, since the figures begin in 1955);
- Figure 9.2b refers to "the first quarter of 1960";
- quarterly data reproduce the book's estimates and figures closely, while monthly data do not, for a reason we explain *below*.

We therefore convert the monthly series to quarterly by taking the last month of each quarter, and measure inflation as the annualized percentage change in the CPI.

Because we sample one month per quarter, we use the seasonally adjusted index rather than the unadjusted one plotted in *The Rise and Fall of U.S. Inflation*.

```
def read_fred(series_id):
    "Read a monthly series from FRED."
    url = f'https://fred.stlouisfed.org/graph/fredgraph.csv?id={series_id}'
    s = pd.read_csv(url, index_col=0, parse_dates=True)[series_id]
    return pd.to_numeric(s, errors='coerce')

cpi_monthly = read_fred('CPIAUCSL')      # CPI, all items, seasonally adjusted
u_monthly = read_fred('LNS14000028')     # unemployment rate, white men 20+
```

(continues on next page)

(continued from previous page)

```
def make_data(how='last', end='1997-03'):
    "Inflation (annualized percent) and unemployment at quarterly or monthly␣
 ↪frequency."
    if how == 'monthly':
        cpi, u, scale = cpi_monthly, u_monthly, 1200
    elif how == 'last':                          # last month of each quarter
        cpi, u, scale = (cpi_monthly.resample('QS').last(),
                         u_monthly.resample('QS').last(), 400)
    else:                                        # quarterly averages
        cpi, u, scale = (cpi_monthly.resample('QS').mean(),
                         u_monthly.resample('QS').mean(), 400)
    y = scale * np.log(cpi).diff()
    data = pd.concat([y.rename('y'), u.rename('U')], axis=1).dropna()
    return data.loc[:end]

data = make_data()
print(f"{data.index[0]:%Y-%m} to {data.index[-1]:%Y-%m}: {len(data)} quarters")
```

```
1954-01 to 1997-01: 173 quarters
```

The unemployment series begins in 1954, and the government's regression needs a few years of data before its estimates mean anything.

So we treat 1954–1957 as a training sample for the government's beliefs and evaluate the likelihood from 1958Q1 through 1997Q1.

The book does not report its lag lengths m_U and m_y; the simulations of its chapter 8 used two lags of each.

We use one lag of unemployment and two of inflation, which among the lag lengths we tried puts our estimates closest to the book's; the robustness table *below* shows what changes with two lags of each.

106.5 Beliefs and recommendations

The Phelps problem is the discounted LQ problem of *Adaptive Learning and Escape Dynamics*, generalized to any number of lags.

Write the believed Phillips curve as $U_t = \gamma_0 y_t + c^\top s_t$ with state $s_t = X_{t-1}$, so that the per-period loss is $s_t^\top c c^\top s_t + (\gamma_0^2 + 1) y_t^2 + 2\gamma_0 y_t c^\top s_t$.

We use the same solver as *Adaptive Learning and Escape Dynamics*: in the discounted problem only A and B are multiplied by $\sqrt{\delta}$, and the cross term N is passed to the Riccati solver unscaled.

```
def lq_matrices(y, mU, my):
    """
    Believed law of motion s_{t+1} = A s_t + B y_t for the state
    s_t = [U_{t-1},...,U_{t-mU}, y_{t-1},...,y_{t-my}, 1]
    when the believed Phillips curve is U_t = y[0] y_t + y[1:] @ s_t.
    """
    y0, c = y[0], y[1:]
    n = mU + my + 1
    A, B = np.zeros((n, n)), np.zeros((n, 1))
    if mU > 0:
        A[0, :], B[0, 0] = c, y0                 # U_t
        for j in range(1, mU):
            A[j, j - 1] = 1.0
```

(continues on next page)

(continued from previous page)

```
    if my > 0:
        B[mU, 0] = 1.0                              # y_t
        for j in range(mU + 1, mU + my):
            A[j, j - 1] = 1.0
    A[-1, -1] = 1.0                                 # constant
    return A, B

def phelps_h(γ, δ, mU, my):
    "Phelps rule y_t = h @ s_t for the believed Phillips curve γ."
    γ0, c = γ[0], γ[1:]
    A, B = lq_matrices(γ, mU, my)
    R = np.outer(c, c)
    Q = np.array([[γ0**2 + 1.0]])
    N = (γ0 * c).reshape(1, -1)
    Ad, Bd = np.sqrt(δ) * A, np.sqrt(δ) * B
    P = solve_discrete_are(Ad, Bd, R, Q, s=N.T)
    F = np.linalg.solve(Q + Bd.T @ P @ Bd, Bd.T @ P @ Ad + N)
    return -F.ravel()

def γ_of_β(β):
    "Invert a Keynesian Phillips curve into classical form."
    return np.concatenate([[1 / β[0]], -β[1:] / β[0]])
```

With a diffuse initial condition, $R_{XC,-1} = 0$, the recursion (106.4) delivers *discounted least squares*: γ_t minimizes $\sum_{s \le t} \lambda^{t-s}(U_s - \gamma^\top X_{C,s})^2$, so an observation k quarters old gets weight λ^k.

The function below computes the beliefs this way from the first quarter of data, and rebuilds, quarter by quarter, the recommendation $h(\gamma_{t-1})X_{t-1}$ that the government would have made.

```
def government(data, scheme, λ=0.985, δ=0.999, mU=1, my=2, n_init=8):
    """
    Beliefs and recommendations of a government that fits the classical ('C')
    or Keynesian ('K') Phillips curve by discounted least squares and solves
    the Phelps problem, reconstructed from data on (y, U).
    """
    U, y = data['U'].to_numpy(), data['y'].to_numpy()
    T, k, m = len(U), mU + my + 2, max(mU, my)
    S, b, zz, w = np.zeros((k, k)), np.zeros(k), 0.0, 0.0
    coef = np.full((T, k), np.nan)
    S_path, σ2 = np.full((T, k, k), np.nan), np.full(T, np.nan)
    states, yhat = np.full((T, k - 1), np.nan), np.full(T, np.nan)
    for t in range(m, T):
        s = np.concatenate([U[t - mU:t][::-1], y[t - my:t][::-1], [1.0]])
        states[t] = s
        if t - m > n_init:                          # recommendation from y_{t-1}
            β = coef[t - 1]
            γ = β if scheme == 'C' else γ_of_β(β)
            yhat[t] = phelps_h(γ, δ, mU, my) @ s
        x = np.r_[y[t], s] if scheme == 'C' else np.r_[U[t], s]
        z = U[t] if scheme == 'C' else y[t]
        S, b = λ * S + np.outer(x, x), λ * b + x * z
        zz, w = λ * zz + z**2, λ * w + 1
        if t - m >= n_init:
            coef[t] = np.linalg.solve(S, b)
            S_path[t] = S
            σ2[t] = (zz - b @ coef[t]) / (w - k)  # discounted residual variance
    lags_y = coef[:, 1 + mU:1 + mU + my].sum(axis=1)
```

(continues on next page)

(continued from previous page)

```
    out = data.copy()
    out['yhat'] = yhat
    out['miss'] = out['y'] - out['yhat']           # v_2t
    out['sum_y'] = lags_y + coef[:, 0] if scheme == 'C' else lags_y
    return dict(frame=out, coef=coef, S=S_path, σ2=σ2, states=states,
                scheme=scheme, λ=λ, δ=δ, mU=mU, my=my)

gov = {s: government(data, s) for s in 'CK'}
```

106.6 Estimates

The function `loglik` evaluates (106.7) with V_v concentrated out.

It uses data from two quarters before the likelihood sample to supply U_{t-1} and U_{t-2}, and, as the book does, it sets the pre-sample shocks to zero.

```
def residuals(params, U, v2):
    "Solve the actual Phillips curve for v_1 given the misses v_2."
    θ, U_star, ρ1, ρ2 = params
    U0, U1, U2 = U[2:], U[1:-1], U[:-2]
    v2 = v2[2:]
    v2_lag = np.r_[0.0, v2[:-1]]
    rhs = (θ * (v2 - ρ1 * v2_lag) + U0 - (ρ1 + ρ2) * U1 + ρ1 * ρ2 * U2
           - U_star * (1 - ρ1) * (1 - ρ2))
    v1 = lfilter([1.0], [1.0, -ρ2], rhs)           # v1_t = ρ2 v1_{t-1} + rhs_t
    return v1, v2

def loglik(params, U, v2):
    "Gaussian log likelihood with V_v concentrated out."
    v = np.vstack(residuals(params, U, v2))
    T = v.shape[1]
    V = v @ v.T / T
    return -T * np.log(2 * np.pi) - T / 2 * np.linalg.slogdet(V)[1] - T, V

def estimate(g, start='1958'):
    "Maximum likelihood estimates of (θ, U*, ρ1, ρ2) for a government g."
    f = g['frame']
    i0 = f.index.get_loc(pd.Timestamp(start)) - 2
    U, v2 = f['U'].to_numpy()[i0:], f['miss'].to_numpy()[i0:]

    def neg(p):
        return -loglik(p, U, v2)[0] if max(abs(p[2]), abs(p[3])) < 1 else 1e10

    starts = [(0.2, 4.0, ρ1, ρ2) for ρ1 in (0.8, 0.95) for ρ2 in (0.0, 0.5, 0.9)]
    fits = [minimize(neg, x0, method='Nelder-Mead',
                     options=dict(xatol=1e-8, fatol=1e-9, maxiter=20_000))
            for x0 in starts]
    best = min(fits, key=lambda r: r.fun)
    return dict(params=best.x, loglik=-best.fun, U=U, v2=v2,
                V=loglik(best.x, U, v2)[1], T=len(U) - 2)

def std_errors(est, h=1e-4):
    "Standard errors from the numerical Hessian of the concentrated likelihood."
    x, k = est['params'], len(est['params'])
```

(continues on next page)

(continued from previous page)

```
    f = lambda p: loglik(p, est['U'], est['v2'])[0]
    E, H = h * np.eye(k), np.empty((k, k))
    for i in range(k):
        for j in range(k):
            H[i, j] = (f(x + E[i] + E[j]) - f(x + E[i] - E[j])
                       - f(x - E[i] + E[j]) + f(x - E[i] - E[j])) / (4 * h**2)
    return np.sqrt(np.diag(np.linalg.inv(-H)))

est = {s: estimate(gov[s]) for s in 'CK'}
se = {s: std_errors(est[s]) for s in 'CK'}
```

Here are our estimates next to the book's Tables 3 and 4.

```
book = {'C': [0.1192, 1.6483, 0.9145, 0.5754, -560],
        'K': [0.163, 2.745, 0.932, 0.667, -333]}
names = ['θ', 'U*', 'ρ1', 'ρ2']
rows = []
for s, label in [('C', 'classical'), ('K', 'Keynesian')]:
    for i, name in enumerate(names):
        p, e = est[s]['params'][i], se[s][i]
        rows.append([label, name, book[s][i], p, e, abs(p - book[s][i]) / e])
    rows.append([label, 'log L', book[s][4], est[s]['loglik'], np.nan, np.nan])
table = pd.DataFrame(rows, columns=['scheme', 'parameter', 'book', 'ours',
                                    'std. error', '|ours - book| / s.e.'])
print(f"likelihood sample: {est['C']['T']} quarters, 1958Q1-1997Q1")
table.set_index(['scheme', 'parameter']).round(3)
```

```
likelihood sample: 157 quarters, 1958Q1-1997Q1
```

scheme	parameter	book	ours	std. error	|ours - book| / s.e.
classical	θ	0.119	0.115	0.042	0.097
	U*	1.648	0.424	1.158	1.057
	ρ1	0.914	0.915	0.025	0.034
	ρ2	0.575	0.695	0.081	1.480
	log L	-560.000	-587.608	NaN	NaN
Keynesian	θ	0.163	0.169	0.056	0.106
	U*	2.745	2.786	0.604	0.069
	ρ1	0.932	0.922	0.026	0.403
	ρ2	0.667	0.720	0.071	0.745
	log L	-333.000	-457.271	NaN	NaN

The book's estimates of θ and ρ_1 lie within a fraction of a standard error of ours under both schemes; its values of U^* and ρ_2 are within about one and a half of our standard errors.

The impact effect θ of an inflation surprise is small, between a tenth and a fifth of a percentage point of unemployment per percentage point of annualized quarterly inflation.

The Phillips-curve disturbance is persistent (ρ_1 above .9), and the effect of a surprise decays faster (ρ_2 near .7).

Two parameters are weakly determined: the estimate of ρ_2 moves noticeably if the likelihood sample starts two quarters later, after the 1958 recession,

```
for s, label in [('C', 'classical'), ('K', 'Keynesian')]:
    p = estimate(gov[s], start='1958-07')['params']
    print(f"{label:10s} from 1958Q3: θ = {p[0]:.3f}, U* = {p[1]:.2f}, "
          f"ρ1 = {p[2]:.3f}, ρ2 = {p[3]:.3f}")
```

```
classical  from 1958Q3: θ = 0.114, U* = 0.34, ρ1 = 0.924, ρ2 = 0.519
Keynesian  from 1958Q3: θ = 0.138, U* = 2.60, ρ1 = 0.938, ρ2 = 0.612
```

and the natural rate U^* has a large standard error.

The Keynesian scheme has a much higher likelihood.

Our gap between the two schemes is smaller than the book's; the levels of $\log L$ are not comparable, because of the misprinted constant and different data vintages, but the sign of the gap is the same.

The residuals show where the gap comes from, and why U^* is hard to pin down:

```
for s, label in [('C', 'classical'), ('K', 'Keynesian')]:
    v1, v2 = residuals(est[s]['params'], est[s]['U'], est[s]['v2'])
    rms = lambda v: np.sqrt(np.mean(v**2))
    print(f"{label:10s}: misses v2: root mean square {rms(v2):.2f}, mean {v2.mean():.
 ↪2f};  "
          f"v1: root mean square {rms(v1):.2f}, mean {v1.mean():.2f};  "
          f"corr(v1, v2) = {np.corrcoef(v1, v2)[0, 1]:.2f}")
print(f"average unemployment over the sample: {est['C']['U'][2:].mean():.2f}")
```

```
classical : misses v2: root mean square 6.82, mean 6.08;  v1: root mean square 0.
 ↪69, mean 0.52;  corr(v1, v2) = 0.58
Keynesian : misses v2: root mean square 3.11, mean 1.55;  v1: root mean square 0.
 ↪51, mean 0.19;  corr(v1, v2) = 0.68
average unemployment over the sample: 4.44
```

The classical government's misses are more than twice as large, in root mean square, as the Keynesian government's.

Both sets of misses have a positive mean, because the recommendations lay below actual inflation for long stretches.

Because v_1 is correlated with v_2, it can have a positive mean too, and a low U^* combined with a positive mean of v_1 fits nearly as well as a higher U^*; that is presumably why both the book and we find values of U^* below the average unemployment rate.

106.7 Recommendations and misses

The fitted value of inflation in this model is not a forecast made by an outside econometrician: it is the inflation rate that the government's own procedure recommended, given the data it had seen.

The next two figures reproduce the book's Figures 9.1 and 9.2a: each shows the recommendation against actual inflation, and the sum of the coefficients on inflation in the government's estimated Phillips curve.

The induction hypothesis (*Testing the natural-rate hypothesis*) holds when that sum is 0 in the classical regression and 1 in the Keynesian regression.

```
def plot_government(g, ref, title, sum_label, start='1958'):
    f = g['frame'].loc[start:]
    fig, axes = plt.subplots(2, 1, figsize=(9, 7), sharex=True)
    axes[0].plot(f.index, f['y'], 'C0--', lw=1, label='actual')
    axes[0].plot(f.index, f['yhat'], 'C1', lw=1.8, label='recommended')
    axes[0].axhline(0, color='k', lw=0.5)
    axes[0].set_ylabel('inflation (percent)')
    axes[0].legend(frameon=False)
    axes[1].plot(f.index, f['sum_y'], 'C2', lw=1.5)
    axes[1].axhline(ref, color='k', ls=':', lw=1, label='induction hypothesis')
    axes[1].set_ylabel(sum_label)
```

(continues on next page)

(continued from previous page)

```
    axes[1].legend(frameon=False)
    axes[0].set_title(title)
    plt.tight_layout()
    plt.show()
```

First the classical government.

```
plot_government(gov['C'], 0, 'Classical identification',
                'sum on current and lagged $y$')
```

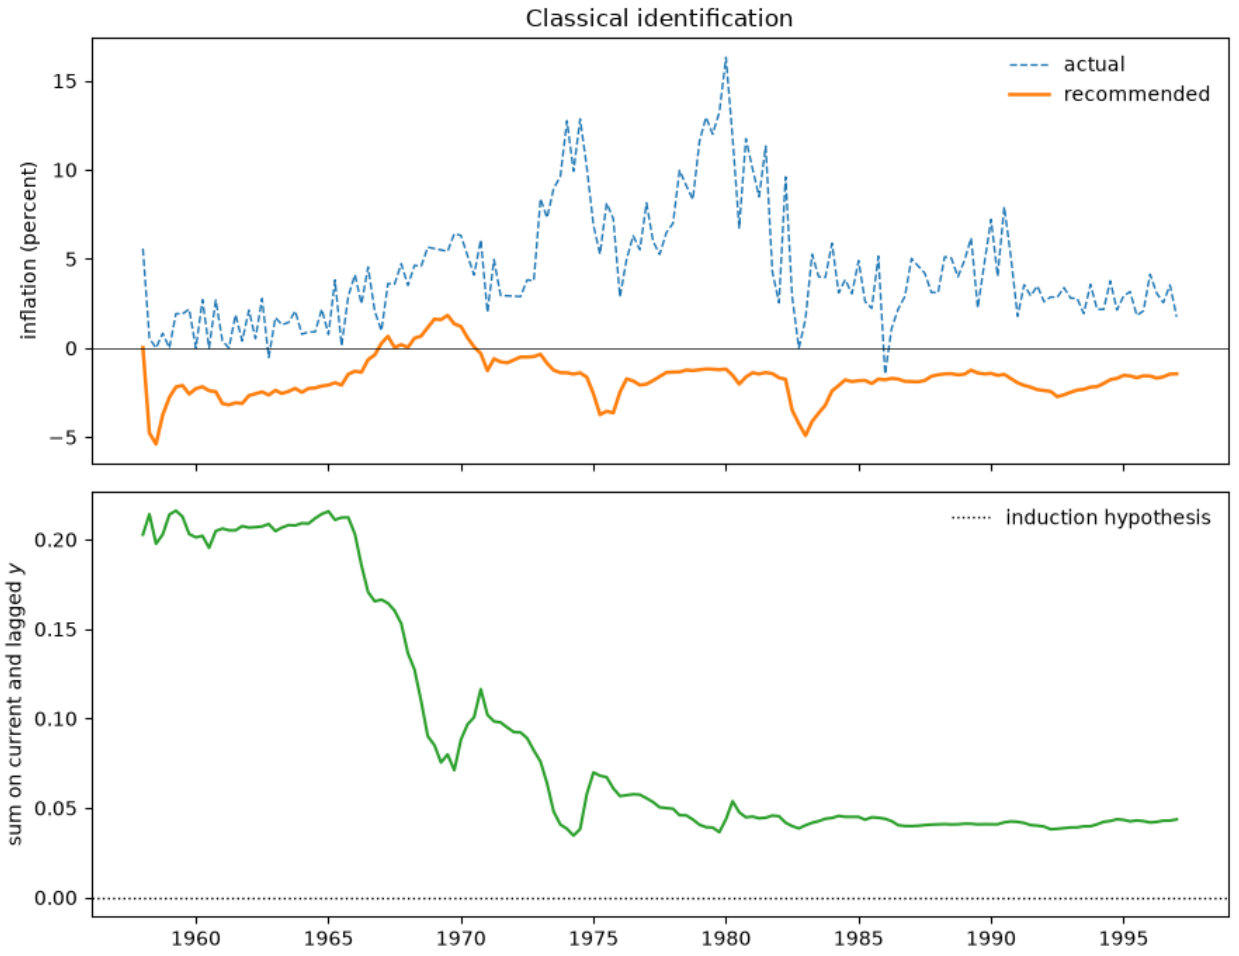


Fig. 106.1: Classical government: recommended and actual inflation, and the sum of coefficients on current and lagged inflation (book Figure 9.1)

The classical government recommends inflation below zero almost throughout, rising above it only briefly around 1970.

As the book says, “the classical model fails even to match the acceleration in inflation leading up to 1970.”

Its regression of unemployment on current and lagged inflation finds coefficients on inflation whose sum is small, and close to zero from the early 1970s on.

Such a government sees almost no exploitable tradeoff, so the Phelps problem tells it to aim for low inflation.

```
plot_government(gov['K'], 1, 'Keynesian identification', 'sum on lagged $y$')
```

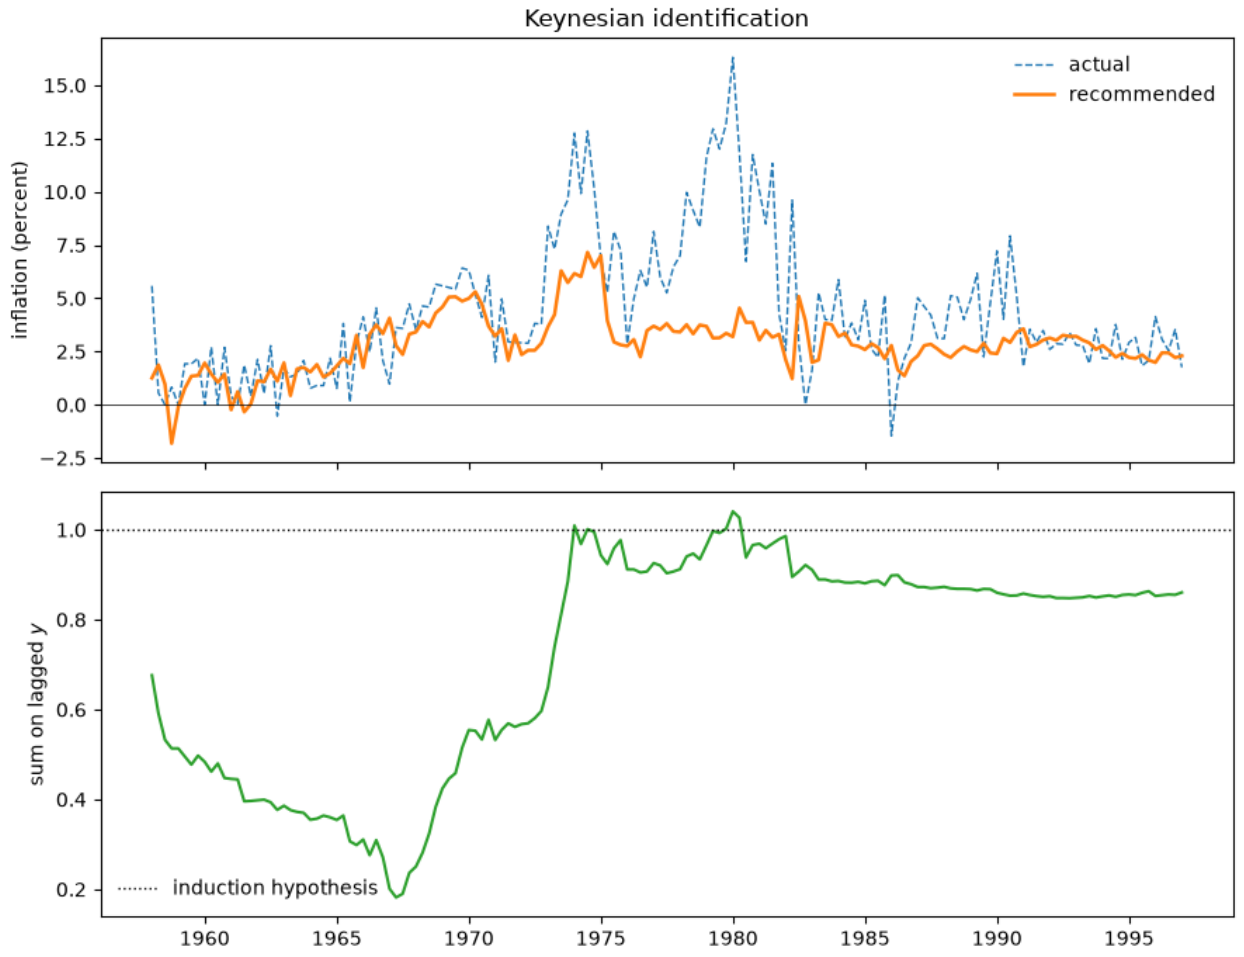


Fig. 106.2: Keynesian government: recommended and actual inflation, and the sum of coefficients on lagged inflation (book Figure 9.2a)

The Keynesian recommendations reflect the inflation pattern better, partly matching the acceleration of the late 1960s.

From the early 1970s, however, they lie well below actual inflation, and they stay below it for most of the next two decades.

The lower panel shows why.

The sum of the weights on lagged inflation climbs from well below one in the 1960s to about one by the mid-1970s and stays near it.

By then the Keynesian government had learned the Solow–Tobin version of the natural-rate hypothesis, and the Phelps problem told it to bring inflation down.

106.7.1 The induction hypothesis emerges

The book's Figure 9.2b shows how the data gradually removed the government's doubt.

We approximate the government's uncertainty about β_t by $\hat{\sigma}_t^2 \left(\sum_{s \le t} \lambda^{t-s} X_{K,s} X_{K,s}^\top\right)^{-1}$, with $\hat{\sigma}_t^2$ the discounted residual variance.

This is the covariance matrix that a Kalman filter implementing constant-gain least squares would attach to its estimate.

The figure plots the 95 and 99 percent confidence ellipses for the sum of the weights on lagged inflation and the constant in the first quarter of 1960, 1965, …, 1985.

```
def ellipse(center, cov, level, n=200):
    "Points on a confidence ellipse for a bivariate normal estimate."
    w, Q = np.linalg.eigh(cov)
    a = np.linspace(0, 2 * np.pi, n)
    circle = np.vstack([np.cos(a), np.sin(a)])
    return center[:, None] + Q @ (np.sqrt(chi2.ppf(level, 2) * w)[:, None] * circle)

g = gov['K']
k = g['coef'].shape[1]
D = np.zeros((2, k))
D[0, 1 + g['mU']:1 + g['mU'] + g['my']] = 1.0      # sum of weights on lagged y
D[1, -1] = 1.0                                      # constant

fig, axes = plt.subplots(2, 3, figsize=(10, 6.5), sharex=True, sharey=True)
for ax, year in zip(axes.ravel(), range(1960, 1990, 5)):
    t = g['frame'].index.get_loc(pd.Timestamp(f'{year}-01-01'))
    center = D @ g['coef'][t]
    cov = g['σ2'][t] * D @ np.linalg.inv(g['S'][t]) @ D.T
    for level, ls in [(0.95, '-'), (0.99, '--')]:
        e = ellipse(center, cov, level)
        ax.plot(e[0], e[1], 'C0', ls=ls, lw=1)
    ax.plot(*center, 'C0o', ms=3)
    ax.axvline(1, color='k', ls=':', lw=1)
    ax.set_title(f'{year}Q1')
for ax in axes[1]:
    ax.set_xlabel('sum on lagged $y$')
for ax in axes[:, 0]:
    ax.set_ylabel('constant')
plt.tight_layout()
plt.show()
```

The pattern resembles the book's figure and the confidence ellipses from the simulations of its chapter 8.

In the 1960s the ellipses lie away from the induction hypothesis (the dotted line) but are large enough to leave room for doubt.

During the 1970s the centers move onto the induction hypothesis and stay there, and the ellipses become smaller than those of the early 1960s, though not as small as the book's.

In the book's words, the point estimates move, "causing the centers of the ellipsoids to shift and doubt to dwindle."

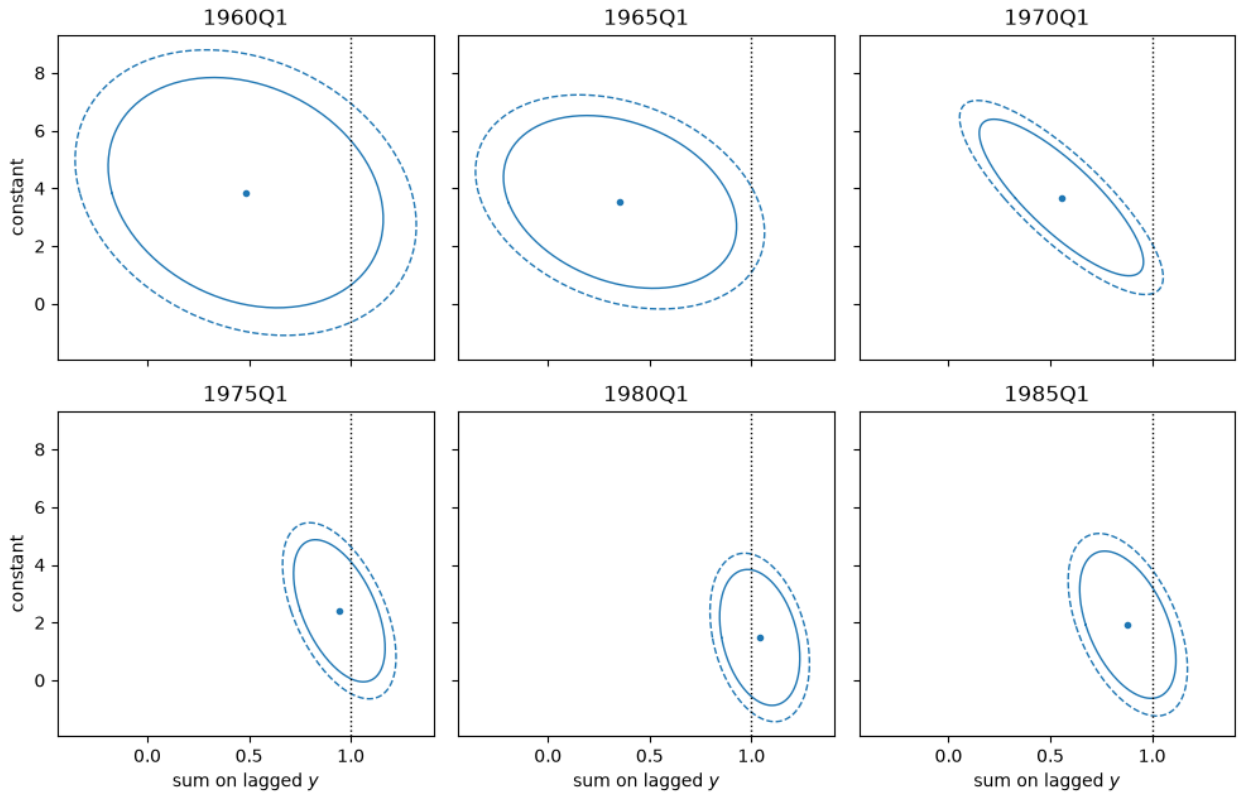


Fig. 106.3: Keynesian government: 95% and 99% confidence ellipses for the sum of weights on lagged inflation and the constant (book Figure 9.2b)

106.8 Misses favor vindication

Judged by period-by-period fit, the adaptive models are disappointing: there is a long string of under-predictions of inflation during the 1970s, even for the Keynesian model.

Sargent reads the misses differently:

> But such misses do not necessarily fail to vindicate econometric policy evaluation. To the contrary, the pattern of misses from the estimated models favors vindication. Remembering that the fitted value from the adaptive model is the government's recommendation as time passes extracts vindication from the econometric shortcomings.

The misses say that the recommendations were systematically *lower* than the inflation that actually occurred: the government's procedure was telling it to disinflate while actual policy did not.

The one-step recommendation $h(\gamma_{t-1})X_{t-1}$ mixes two things: where the government wanted inflation to go, and how fast it planned to get there from the inflation it inherited.

To separate them, we compute the steady state of the government's plan, the long-run inflation rate toward which the Phelps rule would steer the economy if the current beliefs were correct.

```
def planned_long_run(γ, δ, mU, my):
    "Steady-state inflation of the government's plan under its own beliefs."
    h = phelps_h(γ, δ, mU, my)
    A, B = lq_matrices(γ, mU, my)
    M = A + B @ h[None, :]                          # believed closed loop
    n = M.shape[0] - 1                              # last state is the constant
    s_bar = np.linalg.solve(np.eye(n) - M[:n, :n], M[:n, -1])
    return h @ np.r_[s_bar, 1.0]

def long_run_path(g):
    path = np.full(len(g['frame']), np.nan)
    for t in range(1, len(path)):
        β = g['coef'][t - 1]
        if not np.isnan(β).any():
            γ = β if g['scheme'] == 'C' else γ_of_β(β)
            path[t] = planned_long_run(γ, g['δ'], g['mU'], g['my'])
    return pd.Series(path, index=g['frame'].index)

for s in 'CK':
    gov[s]['frame']['long_run'] = long_run_path(gov[s])
```

The next figure puts the Keynesian government's plan next to its one-step recommendation, the classical recommendation, and actual inflation.

```
fK, fC = gov['K']['frame'].loc['1958':], gov['C']['frame'].loc['1958':]
fig, ax = plt.subplots(figsize=(9, 4.8))
ax.plot(fK.index, fK['y'], 'C0--', lw=1, label='actual inflation')
ax.plot(fK.index, fK['yhat'], 'C1', lw=1.5, label='Keynesian recommendation')
ax.plot(fK.index, fK['long_run'], 'C3', lw=2, label='Keynesian planned long-run␣
↪inflation')
ax.plot(fC.index, fC['yhat'], 'C2', lw=1, alpha=0.8, label='classical recommendation')
ax.axvline(pd.Timestamp('1979-08-01'), color='k', ls=':', lw=1)
ax.text(pd.Timestamp('1980-06-01'), -4, 'Volcker', fontsize=9)
ax.axhline(0, color='k', lw=0.5)
ax.set_ylabel('percent')
ax.legend(frameon=False, fontsize=9, loc='upper left')
plt.show()
```

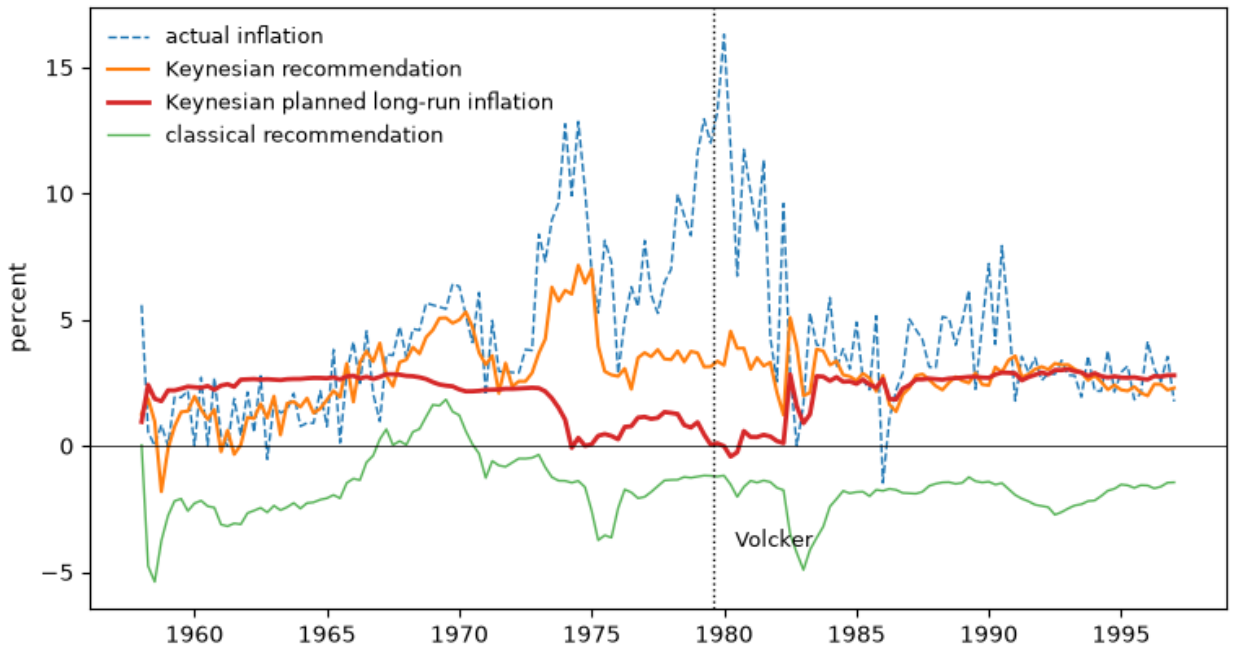


Fig. 106.4: Actual inflation, the Keynesian government's recommendation, and the long-run inflation rate its plan aimed for

Averages over five eras summarize the figure.

```
eras = [('1958', '1965'), ('1966', '1972'), ('1973', '1979'),
        ('1980', '1984'), ('1985', '1997')]
summary = pd.DataFrame(
    {f'{a}-{b}': [fK['y'].loc[a:b].mean(),
                  fC['yhat'].loc[a:b].mean(),
                  fK['yhat'].loc[a:b].mean(),
                  fK['long_run'].loc[a:b].median()]
     for a, b in eras},
    index=['actual inflation', 'classical recommendation',
           'Keynesian recommendation', 'Keynesian long-run plan (median)'])
summary.round(2)
```

```
                                1958-1965  1966-1972  1973-1979  1980-1984  \
actual inflation                     1.40       4.12       8.47       6.32
classical recommendation            -2.56       0.08      -1.72      -2.36
Keynesian recommendation             1.18       3.60       4.14       3.23
Keynesian long-run plan (median)     2.62       2.41       0.80       0.74

                                1985-1997
actual inflation                     3.39
classical recommendation            -1.82
Keynesian recommendation             2.60
Keynesian long-run plan (median)     2.71
```

From 1973 through 1990 the recommendations were mostly below actual inflation:

```
for s, label in [('C', 'classical'), ('K', 'Keynesian')]:
    m = gov[s]['frame']['miss'].loc['1973':'1990']
    print(f"{label:10s}: mean miss 1973-1990 = {m.mean():5.2f}, "
          f"share of quarters with recommendation below actual = {(m > 0).mean():.0%}
↪")
```

```
classical : mean miss 1973-1990 =  8.26, share of quarters with recommendation␣
↪below actual = 100%
Keynesian : mean miss 1973-1990 =  3.04, share of quarters with recommendation␣
↪below actual = 86%
```

The computation bears out the book's interpretation:

> The econometric estimates tell us that unreconstructed Keynesian Phillips curve fitters would have detected the adverse shift in the empirical Phillips curve and through the induction hypothesis combined with the Phelps problem would have recommended lowering the inflation rate. Those quantitative policy evaluators would not have concurred with the loosely argued recommendations current in the late 1970's that long lags in expectations made it too costly to disinflate. Our results say that recommendations under classical identification would have been even more timely.

Through the 1960s the Keynesian government aimed at a long-run inflation rate of two to three percent.

Once its estimated sum of weights on lagged inflation approached one, its target fell to about one percent or less, and it stayed there through the early 1980s, while actual inflation ran far higher.

Its one-step recommendations in the late 1970s, around four percent, called for cutting inflation well below the rates it inherited.

After the Volcker disinflation, the sum of weights drifted down a little and the planned long-run rate returned to two to three percent, close to actual inflation.

The classical government's recommendations were close to zero even during the acceleration of the late 1960s, and below zero after it: that is the sense in which they "would have been even more timely."

Chapter 1 of the book draws the moral: "The econometric estimates present the mystery of why the stabilization was postponed."

106.8.1 How robust are the misses?

Because the book does not report its lag lengths or its conversion to quarterly data, the table varies the lags, the quarterly sampling, and the frequency.

The monthly columns use $\lambda^{1/3}$ and $\delta^{1/3}$, which keep the same memory and the same discounting in calendar time.

```
specs = [('end of quarter', 'last', 1, 2), ('end of quarter', 'last', 2, 2),
         ('quarterly average', 'avg', 1, 2), ('quarterly average', 'avg', 2, 2),
         ('monthly', 'monthly', 1, 2), ('monthly', 'monthly', 2, 2)]
rows = []
for label, how, mU, my in specs:
    d = make_data(how)
    λ, δ = (0.985**(1/3), 0.999**(1/3)) if how == 'monthly' else (0.985, 0.999)
    g = {s: government(d, s, λ=λ, δ=δ, mU=mU, my=my) for s in 'CK'}
    e = {s: estimate(g[s]) for s in 'CK'}
    fk = g['K']['frame']
    fk['long_run'] = long_run_path(g['K'])
    β_U = pd.Series(g['K']['coef'][:, 0], index=fk.index)
    m = fk['miss'].loc['1973':'1990']
```

(continues on next page)

(continued from previous page)

```
    rows.append([label, f'({mU}, {my})',
                 e['K']['loglik'] - e['C']['loglik'],
                 e['K']['params'][0],
                 g['C']['frame']['yhat'].loc['1973':'1979'].mean(),
                 β_U.loc['1973':'1979'].mean(),
                 fk['yhat'].loc['1973':'1979'].mean(),
                 (m > 0).mean(),
                 fk['sum_y'].loc['1975':'1984'].mean(),
                 fk['long_run'].loc['1973':'1984'].median()])
robust = pd.DataFrame(rows, columns=[
    'data', '(m_U, m_y)', 'log L: K minus C', 'K estimate of θ',
    'C recommendation 1973-79', 'K coefficient on U_t, 1973-79',
    'K recommendation 1973-79', 'K share of misses > 0, 1973-90',
    'K sum on lagged y, 1975-84', 'K long-run plan 1973-84'])
robust.set_index(['data', '(m_U, m_y)']).round(2).T
```

data	end of quarter		quarterly average	
(m_U, m_y)	(1, 2)	(2, 2)	(1, 2)	
log L: K minus C	130.34	152.63	140.07	
K estimate of θ	0.17	0.18	0.19	
C recommendation 1973-79	-1.72	-1.31	-1.70	
K coefficient on U_t, 1973-79	-0.74	-0.38	-0.68	
K recommendation 1973-79	4.14	6.00	4.39	
K share of misses > 0, 1973-90	0.86	0.74	0.93	
K sum on lagged y, 1975-84	0.94	0.96	0.96	
K long-run plan 1973-84	0.80	0.74	0.69	

data	monthly		
(m_U, m_y)	(2, 2)	(1, 2)	(2, 2)
log L: K minus C	180.05	438.71	430.68
K estimate of θ	-0.02	0.02	0.02
C recommendation 1973-79	-1.02	-1.85	-1.71
K coefficient on U_t, 1973-79	0.05	0.35	0.28
K recommendation 1973-79	8.01	6.75	6.66
K share of misses > 0, 1973-90	0.61	0.46	0.48
K sum on lagged y, 1975-84	0.97	0.80	0.80
K long-run plan 1973-84	0.71	5.39	5.40

Some findings hold in every column: the Keynesian scheme fits better, and the classical government recommends inflation below zero in the 1970s.

With quarterly data, whatever the lags or the sampling, the Keynesian government's sum of weights on lagged inflation is close to one after 1975, and its long-run plan calls for inflation of about one percent or less in 1973–1984.

The *one-step* Keynesian recommendations are more fragile, and so are the estimates of θ.

The reason is the coefficient β_1 on current unemployment in the Keynesian regression, which the government inverts to get its short-run tradeoff $1/\beta_1$.

With one lag of unemployment it is clearly negative in 1973–1979; with two lags it is about half as large with end-of-quarter sampling and about zero with quarterly averages.

Inverting the curve then divides by a number near zero, so the believed unemployment dynamics become erratic and the government's short-run response to the inflation it inherits swings with them; that is why the one-step misses shrink in those columns.

Monthly data fail for a more basic reason.

Two lags of monthly inflation span only two months, and their coefficients sum to well short of one, so the monthly

Keynesian government never activates the induction hypothesis and its long-run plan rises with inflation instead of falling.

A distributed lag long enough to capture the induction hypothesis needs to span quarters, not months.

106.8.2 A positive, not a normative, conclusion

Sargent does not turn these estimates into an endorsement.

The simulations of *Adaptive Learning and Escape Dynamics* show that the mean dynamics eventually pull an adaptive government back toward the self-confirming equilibrium, so a long spell of good advice is no guarantee of lasting good policy.

Chapter 10 of the book closes:

> In the end, though our simulations and econometric evidence bolster the vindication of econometric policy evaluation story, we hope that this is the wrong story. We hope instead that policy makers somehow have learned a correct rational expectations version of the natural rate hypothesis and to commit themselves to low inflation.

106.9 What came later

Chapter 9 fixed the gain.

Exercise **??** poses Chung's problem of estimating the gain itself on simulated data from the Bounded Rationality suite, where the gain is called γ and the beliefs α; Exercise 106.10.1 below takes up the same question on the data of this lecture.

The econometricians did return the compliment explains why a gain that does not vanish makes such learning parameters estimable.

Sargent *et al.* [2006] are the modern counterpart of this chapter.

Their government runs a Kalman filter (*A First Look at the Kalman Filter*) on drifting Phillips-curve coefficients, the Sims–Chung specification, and, as in chapter 9, the systematic part of inflation solves the Phelps problem at the government's current estimates.

Unlike chapter 9, they estimated the drift covariance that sets the government's gain, jointly with the parameters of the true model, using a Bayesian Markov chain Monte Carlo algorithm.

Priors, Escapes, and Learning Cycles discusses their estimates and Carboni and Ellison [2009]'s refinement, which disciplines the government's beliefs with the Federal Reserve's Greenbook forecasts.

The Lost Conquest: Fed Policy in the 2020s asks Chung's question about the Federal Reserve of the 2020s.

Here the government's beliefs left no free parameters in the likelihood.

The next lecture, *Drifts and Volatilities*, takes the opposite tack: it fits a reduced-form vector autoregression whose coefficients drift and whose shocks have stochastic volatility.

It asks whether the drifting beliefs that the vindication story puts at the center of post-war inflation leave the footprints in the data that the story predicts.

106.10 Exercises

Exercise 106.10.1

The book fixed $\lambda = .985$.

Compute the profile likelihood of the gain: for λ on a grid from .95 to 1, rebuild the government's beliefs and recommendations, re-estimate $(\theta, U^*, \rho_1, \rho_2)$, and plot the maximized log L against λ for both identification schemes.

Note that $\lambda = 1$ is least squares, with a gain that falls like $1/t$.

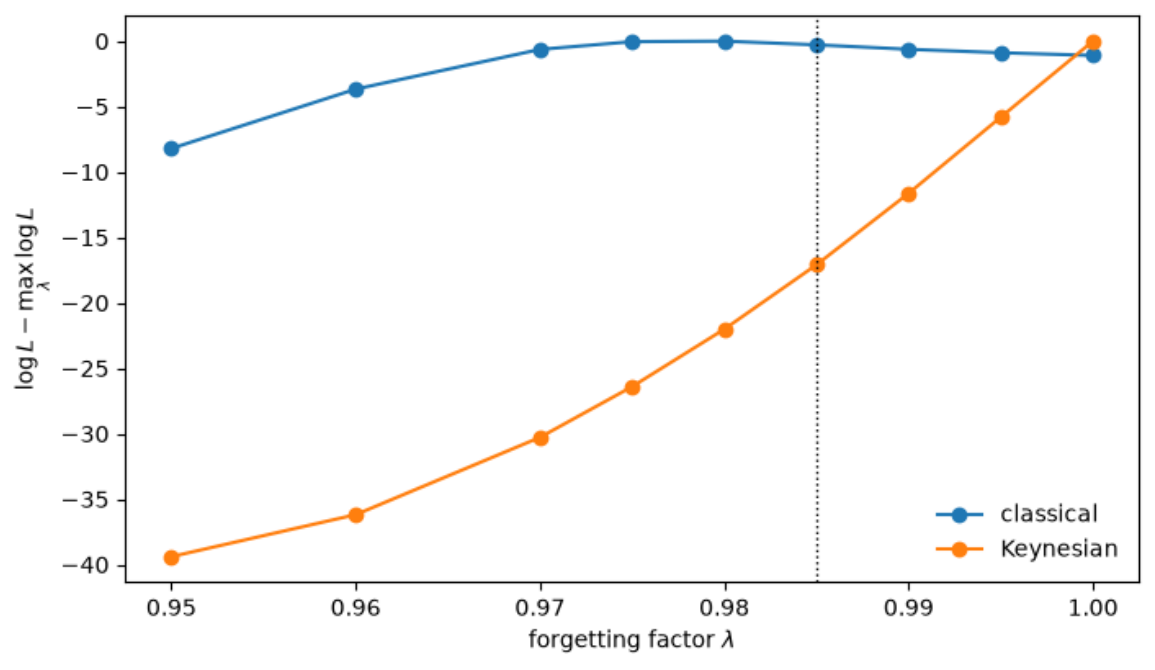
Which values of λ do the data favor, and how does the answer bear on the book's "mystery of why the stabilization was postponed"?

Solution

```
λ_grid = [0.95, 0.96, 0.97, 0.975, 0.98, 0.985, 0.99, 0.995, 1.0]
profile = {s: [] for s in 'CK'}
mean_K = []
for λ in λ_grid:
    for s in 'CK':
        g = government(data, s, λ=λ)
        profile[s].append(estimate(g)['loglik'])
        if s == 'K':
            mean_K.append(g['frame']['yhat'].loc['1973':'1979'].mean())

fig, ax = plt.subplots(figsize=(8, 4.5))
for s, label in [('C', 'classical'), ('K', 'Keynesian')]:
    ax.plot(λ_grid, np.array(profile[s]) - max(profile[s]), 'o-', label=label)
ax.axvline(0.985, color='k', ls=':', lw=1)
ax.set_xlabel(r'forgetting factor $\lambda$')
ax.set_ylabel(r'$\log L - \max_\lambda \log L$')
ax.legend(frameon=False)
plt.show()

print(" λ      log L (C)  log L (K)  K recommendation 1973-79")
for λ, lc, lk, m in zip(λ_grid, profile['C'], profile['K'], mean_K):
    print(f"{λ:5.3f}  {lc:9.1f}  {lk:9.1f}  {m:8.2f}")
```



λ	log L (C)	log L (K)	K recommendation 1973–79
0.950	−595.5	−479.6	3.12
0.960	−591.0	−476.4	3.22
0.970	−588.0	−470.5	3.49
0.975	−587.4	−466.6	3.68
0.980	−587.3	−462.2	3.90
0.985	−587.6	−457.3	4.14
0.990	−588.0	−451.8	4.39
0.995	−588.2	−446.0	4.64
1.000	−588.4	−440.2	4.86

The classical likelihood is nearly flat for λ above .97, because the classical government recommends low inflation whatever its memory.

The Keynesian likelihood rises steadily as λ approaches one: the data prefer the government that discounts the past least.

The printed averages show why.

A government with a longer memory gives more weight to the 1950s and 1960s, when its regression showed a long-run tradeoff, so it activates the induction hypothesis later and recommends higher inflation in the 1970s, closer to the inflation that actually occurred.

Because the likelihood can reward a gain only by bringing the recommendations closer to actual inflation, it favors the slowest learner.

That is another way to state the book's mystery: actual policy behaved as if it learned the natural-rate hypothesis later than the book's constant-gain government did.

In the simulated economy of Exercise **??** the government really does follow its recommendations, and the likelihood recovers the true gain; here a preference for slow learning partly measures how far actual policy lagged behind the procedure.

Exercise 106.10.2

Sims [1988] gave his government a *static* Phillips curve, $U_t = \gamma_0 y_t + \gamma_c + \varepsilon_t$, with no lags.

Estimate the static model ($m_U = m_y = 0$) under both identification schemes, and compare it with the distributed-lag model of the lecture.

For each of the four models, report the maximized log L, the average recommendation in 1973–1979, and the correlation between recommended and actual inflation.

Does the static classical model fail, and is it the lags or the direction of fit that matters more?

Solution

```
rows = []
for mU, my, lag_label in [(0, 0, 'static'), (1, 2, 'distributed lag')]:
    for s, label in [('C', 'classical'), ('K', 'Keynesian')]:
        g = government(data, s, mU=mU, my=my)
        f = g['frame'].loc['1958':]
        rows.append([lag_label, label, estimate(g)['loglik'],
                     f['yhat'].loc['1973':'1979'].mean(),
                     f[['y', 'yhat']].corr().iloc[0, 1]])
pd.DataFrame(rows, columns=['lags', 'scheme', 'log L', 'recommendation 1973-79',
                            'corr(recommended, actual)']
             ).set_index(['lags', 'scheme']).round(2)
```

```
                                log L  recommendation 1973-79  \
lags            scheme
static          classical -537.58                    0.23
                Keynesian -473.11                    4.05
distributed lag classical -587.61                   -1.72
                Keynesian -457.27                    4.14

                                corr(recommended, actual)
lags            scheme
static          classical                         -0.24
                Keynesian                          0.07
distributed lag classical                          0.31
                Keynesian                          0.58
```

The static classical model does fail: it recommends inflation near zero in the 1970s, and its recommendations are negatively correlated with actual inflation.

But adding lags does not rescue the classical scheme: the distributed-lag classical government recommends even lower inflation, and its likelihood is lower than the static one's.

The direction of fit matters more than the lags, and both Keynesian models fit far better than either classical model.

In a regression of unemployment on inflation, noise in quarterly inflation attenuates the coefficient on inflation, so a classical government sees only a weak tradeoff; inverting a regression of inflation on unemployment yields a steep one.

For the Keynesian scheme the lags matter too: the distributed-lag model has the higher likelihood, and its recommendations track actual inflation far more closely.

Only a distributed lag lets the government discover the induction hypothesis; a static Keynesian government has no lagged inflation in its model, so it cannot learn the natural-rate hypothesis in the Solow–Tobin form.

The book takes the two identification schemes, and their names, from King and Watson [1994], and it notes that the better fit of the Keynesian direction is what led the U.S. Phillips-curve literature to adopt it.

Exercise 106.10.3

Take the Keynesian government's beliefs at the end of 1979Q2, just before Paul Volcker became chairman of the Federal Reserve.

Compute the paths of inflation and unemployment that it *planned* for the next twenty quarters: the paths implied by its Phelps rule and its own estimated Phillips curve, with all future shocks set to zero, starting from the state in 1979Q3.

Do the same for the classical government, and plot both plans against what actually happened.

How does the Keynesian plan compare with the actual disinflation, and how large a rise in unemployment did each government expect to pay for it?

Solution

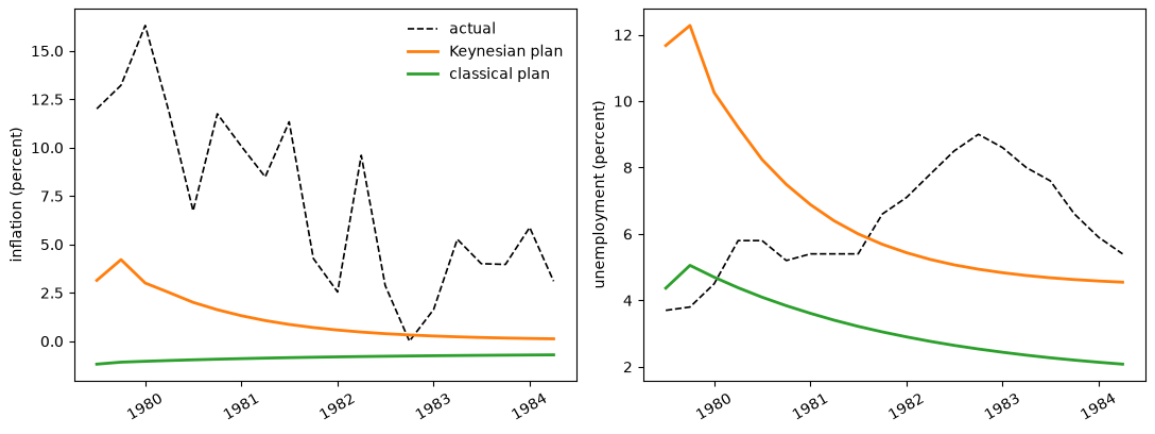
```
def planned_paths(g, date, H=20):
    "Inflation and unemployment the government planned from `date` on."
    t = g['frame'].index.get_loc(pd.Timestamp(date))
    β = g['coef'][t - 1]
    γ = β if g['scheme'] == 'C' else γ_of_β(β)
    h = phelps_h(γ, g['δ'], g['mU'], g['my'])
    A, B = lq_matrices(γ, g['mU'], g['my'])
    s = g['states'][t].copy()
    y_plan, U_plan = [], []
    for _ in range(H):
        y = h @ s
        y_plan.append(y)
        U_plan.append(γ[0] * y + γ[1:] @ s)
        s = A @ s + B[:, 0] * y
    index = g['frame'].index[t:t + H]
    return pd.DataFrame({'y': y_plan, 'U': U_plan}, index=index[:H])

date = '1979-07-01'
plans = {s: planned_paths(gov[s], date) for s in 'CK'}
actual = data.loc[date:].iloc[:20]

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))
for ax, var, ylabel in [(axes[0], 'y', 'inflation'), (axes[1], 'U', 'unemployment
 ↪')]:
    ax.plot(actual.index, actual[var], 'k--', lw=1.2, label='actual')
    ax.plot(plans['K'].index, plans['K'][var], 'C1', lw=2, label='Keynesian plan')
    ax.plot(plans['C'].index, plans['C'][var], 'C2', lw=2, label='classical plan')
    ax.set_ylabel(f'{ylabel} (percent)')
    ax.tick_params(axis='x', rotation=30)
axes[0].legend(frameon=False)
plt.tight_layout()
plt.show()

print(f"1979Q2: inflation {data.loc['1979-04-01', 'y']:.1f}, "
      f"unemployment {data.loc['1979-04-01', 'U']:.1f}")
paths = {'Keynesian plan': plans['K'], 'classical plan': plans['C'], 'actual':␣
```

```
 ↪actual}
print(f"{'':16s} {'y, 1st qtr':>10s} {'y, 4th qtr':>10s} {'y, 20th qtr':>11s} {
 ↪'peak U':>7s}")
for label, p in paths.items():
    print(f"{label:16s} {p['y'].iloc[0]:10.1f} {p['y'].iloc[3]:10.1f} "
          f"{p['y'].iloc[-1]:11.1f} {p['U'].max():7.1f}")
```




```
1979Q2: inflation 12.9, unemployment 3.4
                 y, 1st qtr y, 4th qtr y, 20th qtr  peak U
Keynesian plan          3.1        2.5         0.1    12.3
classical plan         -1.2       -1.0        -0.7     5.1
actual                 12.0       11.8         3.1     9.0
```

On the eve of Volcker's appointment the Keynesian government planned to cut inflation at once, from double digits to a few percent, and then to lower it slowly toward zero.

It expected the cut to be costly: unemployment in its plan jumps to well above its 1979 level and then declines slowly.

This is the pattern of the Phelps problem in chapter 5 of the book (*Adaptive Expectations and the Phelps Problem*), where the government "engineers a major recession and immediately brings inflation down more than half way toward its eventual limiting value."

The classical government planned a mild deflation and expected it to cost little, because its estimated curve showed hardly any short-run effect of inflation on unemployment.

Actual policy moved more slowly than either plan: inflation stayed high for more than a year and reached its 1980s level only in 1982–1983, while the unemployment rate of white men rose to its peak at the end of 1982.

Both governments' econometrics said disinflate now, and the Keynesian government recommended it even though it expected disinflation to be costly.

Neither would have accepted the argument, current in the late 1970s, that long lags in expectations made disinflation too costly to attempt.

CHAPTER

SEVEN

DRIFTS AND VOLATILITIES

Contents

- *Drifts and Volatilities*
  - *Overview*
  - *Bad policy or bad luck?*
  - *A VAR with drifting coefficients and stochastic volatility*
  - *The data*
  - *Priors*
  - *Simulating the posterior*
  - *What the data say*
  - *Why classical tests missed the drift*
  - *Bad policy or bad luck? A verdict*
  - *Appendix: the sampler in detail*
  - *Exercises*

107.1 Overview

The lectures in this section have told a story about how a government's *model* of the Phillips curve, and the *policy* it induces, can drift over time.

In *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation* a government that refits an approximating Phillips curve by constant-gain least squares escapes recurrently from a bad *self-confirming equilibrium* and is then drawn back toward it.

Priors, Escapes, and Learning Cycles shows how the government's prior about drifting coefficients shapes those dynamics, and *Econometric Policy Evaluation* confronts the adaptive models with postwar U.S. data by maximum likelihood.

That confrontation was *structural*: it estimated a particular model of a learning government.

This lecture steps back to a *reduced-form* question that every one of those models has to answer.

It studies Cogley and Sargent [2005], which asks a deceptively simple question:

> When we look at postwar U.S. time series on inflation, unemployment, and interest rates, do we see evidence that the dynamics have *drifted*?

Tim Cogley and Thomas Sargent began this work as an empirical companion to the *Conquest* book [Sargent, 1999] and to Cho *et al.* [2002], whose recurrent escapes generate nonlinearities that an econometrician would see as drifting coefficients.

It answers searching comments by Sims [2001] and Stock [2001] on an earlier paper [Cogley and Sargent, 2001], and it confronts evidence from Sims [1999] and Bernanke and Mihov [1998] that the coefficients of vector autoregressions had *not* changed.

That evidence bears on a question that organizes this whole section:

Was the Great Inflation of the 1970s and its conquest in the 1980s a story of bad policy, or of bad luck?

To let the data speak to that question we need a statistical model flexible enough to accommodate *both* answers.

That model is a *Bayesian vector autoregression whose coefficients drift as random walks and whose shock variances evolve as stochastic volatilities*.

It is the very specification that Sims [1988] and Chung [1990] handed their model governments: a Kalman filter on regression coefficients that are believed to drift as random walks (*The econometricians did return the compliment*).

There the government ran the filter to learn the Phillips curve; here the econometrician runs it to learn about the government.

Fitting the model requires a Markov chain Monte Carlo algorithm that combines the *Kalman filter*, the forward-filter/backward-sample smoother of Carter and Kohn [1994], and the stochastic-volatility sampler of Jacquier *et al.* [1994].

Companion-form vector autoregressions are treated in *Linear State Space Models* and *Vector Autoregressions for Subsets of Variables*, and the Kalman filter in *A First Look at the Kalman Filter*.

No other QuantEcon lecture covers the Kalman *smoother*, so this lecture introduces it inside the sampler, and Exercise **??** shows why it matters here.

We fit the model to Cogley and Sargent's sample, which ends in 2000; *Drifts and Volatilities after 2000* re-estimates it with another quarter-century of data.

Let's start with some imports and the URL for the data.

```
import time

import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from IPython.display import display
from scipy import linalg
from scipy.special import expit
from threadpoolctl import threadpool_limits

# The matrices in this lecture are small, so extra BLAS threads only add overhead.
threadpool_limits(1)

data_url = 'https://github.com/QuantEcon/data-lectures/raw/main/lectures/NEWQDATA.csv'
```

107.2 Bad policy or bad luck?

Two respectable views compete to explain the American Great Inflation.

The **bad policy** view is the one dramatized throughout this section and in the *Conquest* book [Sargent, 1999].

Something about Arthur Burns's *model* of the economy, his *patience*, or his inability to *commit* to a better rule led the Federal Reserve to administer monetary policy in a way that produced the greatest peacetime inflation in U.S. history, while an improved model, more patience, or greater discipline led Paul Volcker to conquer it [DeLong, 1997, Taylor, 1997].

The two stories that open *The Rise and Fall of U.S. Inflation*, triumph and vindication, are both versions of this view: they disagree about how the Fed's model improved, not about whether policy changed.

On this view, what changed between the 1970s and the 1980s was the *systematic part* of policy, namely the way the Fed's interest-rate setting responded to inflation and unemployment.

The **bad luck** view rejects both stories.

What distinguished the Burns and Volcker eras was not their models or policies, but the *shocks* that hit the economy.

On this view the *coefficients* of a reduced-form description of the economy were essentially constant, and what changed was the *size* of the disturbances, namely the *volatility*.

Classical tests supported the second view.

Sims [1999] and Bernanke and Mihov [1998] could not reject the hypothesis that VAR coefficients were time invariant, while finding that innovation variances had shifted.

Sims and Zha [2006] reached a similar verdict by a different route.

Comparing Markov-switching models by their posterior fit, they found that the best fit lets only the variances of the disturbances switch.

When the systematic coefficients were also allowed to switch, the best fit changed only the monetary policy rule, and the estimated policy regimes differed too little to account for the rise and fall of inflation [Sargent, 2008].

Markov switching and the random-walk volatilities used below describe changing volatility differently.

A hidden Markov chain permits *jumps*, but without an absorbing state it switches forever among the same few configurations, and with one it permits a permanent shift only once.

A random-walk volatility must adjust continuously, but it permits *recurrent, permanent* shifts, such as a Great Inflation followed by a Great Moderation.

Primiceri [2005] fit a time-varying VAR close to the one below, with drifting coefficients and stochastic volatility and a prior that puts more weight on slow drift.

He found that systematic monetary policy had changed, but that the changes had small effects on the rest of the economy, which puts him nearer the bad-luck side [Sargent, 2008].

How can we discriminate?

A constant-coefficient, constant-volatility VAR can generate unusually large realized shocks, but it cannot represent systematic changes in their variance over time.

A model with drifting coefficients but constant volatility can also mistake changing volatility for coefficient drift, which is what Stock [2001] suspected of the 2001 paper.

So Cogley and Sargent build a model that has room for *both* channels at once, and they let a Bayesian posterior sort out how much of each the data call for.

107.3 A VAR with drifting coefficients and stochastic volatility

Let the variables be ordered as nominal interest, transformed unemployment, and inflation,

$$y_t = \begin{bmatrix} i_t & u_t & \pi_t \end{bmatrix}^\top .$$

Here u_t is not the raw unemployment rate but its logit, a transformation we define in the data section below.

The measurement equation is a VAR with two lags and date-specific coefficients,

$$y_t = X_t^\top \theta_t + \varepsilon_t, \qquad X_t^\top = I_3 \otimes \begin{bmatrix} 1 & y_{t-1}^\top & y_{t-2}^\top \end{bmatrix}. \tag{107.1}$$

Each equation has an intercept and six lag coefficients, so θ_t contains $3(1 + 2 \times 3) = 21$ elements.

A two-lag VAR can be rewritten as a one-lag system by stacking y_t and y_{t-1} into a single vector; the matrix that multiplies this stacked vector is the **companion matrix**, and the rewritten system is the VAR in **companion form** (*Linear State Space Models*, *Vector Autoregressions for Subsets of Variables*).

The following function builds the companion matrix from the stacked coefficients.

```
n_variables = 3
n_lags = 2
n_regressors = 1 + n_variables * n_lags
n_coefficients = n_variables * n_regressors


def companion_matrix(θ):
    """Return the intercept and companion matrix for one coefficient vector."""
    rows = np.asarray(θ, dtype=float).reshape(n_variables, n_regressors)
    intercept = np.r_[rows[:, 0], np.zeros(n_variables)]
    companion = np.zeros((n_variables * n_lags, n_variables * n_lags))
    companion[:n_variables] = rows[:, 1:]
    companion[n_variables:, :n_variables] = np.eye(n_variables)
    return intercept, companion
```

The coefficient vector follows a driftless random walk,

$$\theta_t = \theta_{t-1} + v_t, \qquad v_t \sim N(0, Q). \tag{107.2}$$

The drift covariance Q is a close relative of the drift covariance V in *Priors, Escapes, and Learning Cycles*.

There V is the *government's* prior about how fast the Phillips curve drifts, and it sets the gain of the government's Kalman filter; here Q is a parameter of the *econometrician's* model, to be estimated.

For a given drift covariance the Kalman gain does not die out, so the filter behaves like the constant-gain least squares of *Adaptive Learning and Escape Dynamics*, a stochastic-approximation algorithm of the kind studied in *One recursion, many learners*.

Exercise **??** asks how well such a gain can be estimated from data, and Exercise **??** computes the gain that this model's estimates imply.

When Sargent *et al.* [2006] estimated V by requiring a learning government's policy to fit U.S. inflation, they needed a large V, a government that "discounts past data heavily" [Sargent, 2008]; Carboni and Ellison [2009] shrank it by asking the government's beliefs to match the Fed staff's Greenbook forecasts.

The companion system is stable when every companion-matrix eigenvalue lies strictly inside the unit circle.

For an $AR(1)$, this is $|\rho| < 1$; equivalently, the zero of $1 - \rho z$ lies outside the unit circle.

Cogley and Sargent rule out explosive paths with the truncated prior

$$p(\theta^T, Q) \propto I(\theta^T) f(\theta^T \mid Q) f(Q), \tag{107.3}$$

where $I(\theta^T) = 1$ if the companion matrix is stable at every date and $I(\theta^T) = 0$ otherwise.

Their reason is economic.

A Fed that chooses its rule purposefully, with a loss function that penalizes the variance of inflation, would never choose a rule that leaves a unit root in inflation, because that would deliver an infinite loss.

The roots of the VAR are not determined by the Phillips curve alone; they also depend on the policy rule.

Under an arbitrary rule they can lie inside, on, or outside the unit circle, but under optimal or nearly optimal rules they are stable: when a shock moves inflation away from its target, a badly chosen rule may let it drift, while a well-chosen rule pulls it back.

The Phelps problem of *Adaptive Expectations and the Phelps Problem* is an example: the optimal rule steers inflation to a rest point even when the Phillips curve itself has the Solow–Tobin unit-sum form.

The restriction also tilts the marginal prior for Q toward values that are less likely to generate explosive coefficient paths.

The code below tests a whole trajectory for stability.

```
def companion_roots(θ_path):
    """Return the companion roots at every date of a (21, T) path."""
    rows = np.asarray(θ_path, dtype=float).reshape(n_variables, n_regressors, -1)
    companions = np.zeros((rows.shape[2], 6, 6))
    companions[:, :n_variables] = rows[:, 1:].transpose(2, 0, 1)
    companions[:, n_variables:, :n_variables] = np.eye(n_variables)
    return np.linalg.eigvals(companions)


def is_stable(θ_path):
    """Test whether every companion root lies strictly inside the unit circle."""
    return bool(np.max(np.abs(companion_roots(θ_path))) < 1)
```

The reduced-form innovation covariance changes over time according to

$$\varepsilon_t = R_t^{1/2}\xi_t, \qquad \xi_t \sim N(0, I_3), \qquad R_t = B^{-1}H_tB^{-1\prime}, \tag{107.4}$$

where

$$B = \begin{bmatrix} 1 & 0 & 0 \\ \beta_{21} & 1 & 0 \\ \beta_{31} & \beta_{32} & 1 \end{bmatrix}, \qquad H_t = \operatorname{diag}(h_{1t}, h_{2t}, h_{3t}).$$

The diagonal elements h_{it} let the size of each orthogonalized shock wax and wane over time.

The matrix B orthogonalizes ε_t, but it is not an identification scheme.

The next two functions construct B and R_t.

```
def b_matrix(β):
    """Construct B from β_21, β_31, and β_32."""
    matrix = np.eye(n_variables)
    matrix[1, 0], matrix[2, 0], matrix[2, 1] = β
    return matrix


def innovation_covariance(h, β):
    """Construct R_t = B^{-1} H_t B^{-1}' from one vector of variances."""
    inverse = np.linalg.inv(b_matrix(β))
    return inverse @ np.diag(h) @ inverse.T
```

Each diagonal volatility is a geometric random walk,

$$\log h_{it} = \log h_{i,t-1} + \sigma_i \eta_{it}, \qquad \eta_{it} \sim N(0,1). \tag{107.5}$$

The standardized measurement innovations, coefficient innovations, and volatility innovations are mutually independent.

Setting $Q = 0$ produces constant coefficients with drifting volatility, while holding H_t fixed produces drifting coefficients with constant volatility.

The posterior contains the full paths θ^T and H^T together with Q, $\beta = (\beta_{21}, \beta_{31}, \beta_{32})$, and $\sigma = (\sigma_1, \sigma_2, \sigma_3)$.

107.4 The data

We use Cogley and Sargent [2005]'s quarterly U.S. dataset, which ends in 2000Q4.

Inflation is the log difference of the seasonally adjusted CPI for all urban consumers, point sampled in the third month of each quarter.

Unemployment is the quarterly average of the seasonally adjusted civilian unemployment rate and enters the VAR as $0.01 \log[u/(1-u)]$, a logit transformation that maps a bounded rate into an unconstrained variable.

The nominal interest rate starts from the three-month Treasury-bill rate i, in annual percent and averaged over daily observations in the first month of each quarter, and enters as $i/(400 - i)$, the quarterly return on a bill bought at a discount of $i/400$ of its face value.

The data start in 1948Q2 because the first inflation observation is already differenced, so two VAR lags make 1948Q4 the first usable regression date.

Observations through 1958Q4 calibrate the prior, and 1959Q1–2000Q4 form the estimation sample.

```
def prepare_data(source):
    """Transform a quarterly table and construct the VAR(2) data."""
    table = pd.read_csv(source) if isinstance(source, str) else source
    ur = table['ur'].to_numpy(dtype=float)
    raw_y = np.column_stack([
        table['y3'].to_numpy(dtype=float),        # nominal interest
        0.01 * np.log(ur / (1 - ur)),             # logit unemployment
        table['dp'].to_numpy(dtype=float),        # inflation
    ])
    regressors = np.ones((len(raw_y) - n_lags, n_regressors))
    for lag in range(1, n_lags + 1):
        left = 1 + n_variables * (lag - 1)
        regressors[:, left:left + n_variables] = raw_y[n_lags - lag:-lag]
    targets = raw_y[n_lags:]
    dates = table['date'].to_numpy(dtype=float)[n_lags:]
    n_training = 4 * 11 - n_lags - 1           # 1948Q4--1958Q4
    return {
        'raw_dates': table['date'].to_numpy(dtype=float),
        'raw_y': raw_y,
        'prior_y': targets[:n_training],
        'prior_x': regressors[:n_training],
        'dates': dates[n_training:],
        'y': targets[n_training:],
        'x': regressors[n_training:],
    }


data = prepare_data(data_url)
```

Let's view the data in familiar economic units.

```
dates_raw = data['raw_dates']
fig, axes = plt.subplots(3, 1, figsize=(9, 7), sharex=True)
axes[0].plot(dates_raw, 400 * data['raw_y'][:, 2], lw=2)
axes[0].set_ylabel('inflation (annual %)')
axes[1].plot(dates_raw, 100 * expit(100 * data['raw_y'][:, 1]), lw=2)
axes[1].set_ylabel('unemployment (%)')
interest = data['raw_y'][:, 0]
axes[2].plot(dates_raw, 400 * interest / (1 + interest), lw=2)
axes[2].set_ylabel('interest (annual %)')
axes[2].set_xlabel('year')
plt.tight_layout()
plt.show()
```

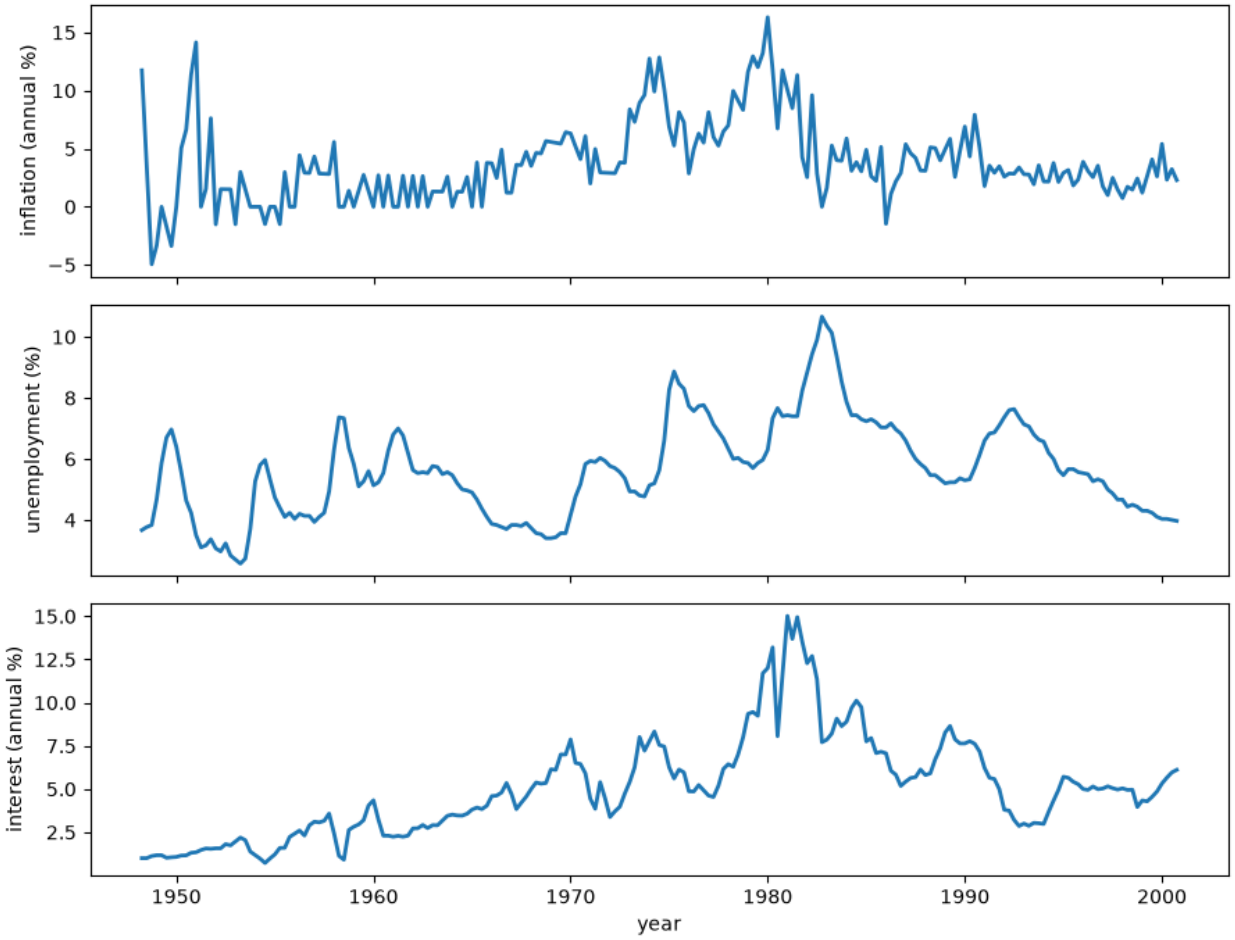


Fig. 107.1: Observed π_t, u_t, and i_t, 1948Q2–2000Q4

Inflation and nominal interest rates rise together through the 1970s and peak around 1980, unemployment peaks only after disinflation begins, and all three series are calmer in the 1990s.

These observations locate the episode but cannot tell whether changed systematic dynamics or unusually large shocks produced it, which is the distinction the model is built to examine.

107.5 Priors

The priors are independent across parameters and deliberately weak so that, in Cogley and Sargent's phrase, "the data are free to speak."

They are calibrated from a time-invariant VAR fitted to the training sample.

The initial coefficient prior is a stable, truncated Gaussian,

$$p(\theta_0) \propto I(\theta_0) N(\bar{\theta}, \bar{P}),$$

where $\bar{\theta}$ and $\bar{P}$ are the point estimate and asymptotic covariance of a constant-coefficient seemingly unrelated regression fitted to 1948–1958.

Because all three equations have identical regressors, the SUR estimates equal equation-by-equation OLS estimates.

The coefficient-drift covariance has the inverse-Wishart prior

$$Q \sim IW_{21}\left(T_0, T_0 \bar{Q}\right), \qquad T_0 = 22, \qquad \bar{Q} = \gamma^2 \bar{P}, \qquad \gamma^2 = 3.5 \times 10^{-4}, \tag{107.6}$$

where the first argument is the degrees of freedom and the second the scale matrix.

Here γ^2 is Cogley and Sargent's prior scale, not a learning gain; in the Bounded Rationality lectures γ denotes the gain (*Notation across the series*).

Because T_0 is only one greater than the dimension of θ_t, the prior is proper but has no finite mean.

The matrix $\bar{Q}$ is therefore a conservative scale calibration that implies little drift, not the expectation of Q, and the weakness of the prior lets the data move far away from it.

The remaining priors are

$$\begin{aligned} \log h_{i0} &\sim N(\log \bar{R}_{ii}, 10), \\ \beta &\sim N(0, 10000\, I_3), \\ \sigma_i^2 &\sim IG\left(\frac{1}{2}, \frac{0.01^2}{2}\right), \end{aligned}$$

where $\bar{R}$ is the residual covariance from the training-sample regression.

The next function gathers these hyperparameters.

```
def calibrate_prior(model_data, γ_squared=3.5e-4):
    """Calibrate every prior from a constant-coefficient training-sample VAR."""
    y, x = model_data['prior_y'], model_data['prior_x']
    xx_inverse = np.linalg.inv(x.T @ x)
    coefficients = xx_inverse @ x.T @ y          # SUR = equation-by-equation OLS
    residuals = y - x @ coefficients
    r_bar = np.cov(residuals, rowvar=False, ddof=1)
    p_bar = np.kron(r_bar, xx_inverse)
    T0 = n_coefficients + 1
    return {
        'θ_mean': coefficients.T.reshape(-1),
        'θ_covariance': p_bar,
        'q_center': γ_squared * p_bar,           # Q bar
        'q_scale': T0 * γ_squared * p_bar,       # T0 Q bar
        'q_degrees_freedom': T0,
        'log_h_mean': np.log(np.diag(r_bar)),
        'log_h_variance': 10.0,
        'β_variance': 10_000.0,
```

(continues on next page)

(continued from previous page)

```
        'σ_degrees_freedom': 1.0,
        'σ_scale': 0.01**2,
        'γ_squared': γ_squared,
    }


prior = calibrate_prior(data)
trace_q_bar = np.trace(prior['q_center'])
assert is_stable(prior['θ_mean'])

pd.Series({
    'T0': prior['q_degrees_freedom'],
    'γ squared': prior['γ_squared'],
    'tr(Q bar)': trace_q_bar,
}, name='value').to_frame()
```

```
                value
T0          22.000000
γ squared    0.000350
tr(Q bar)    0.003004
```

The training-sample estimate $\bar{\theta}$ is itself stable, so the chain can start there.

107.6 Simulating the posterior

We simulate the posterior by cycling through the five blocks that Cogley and Sargent [2005] use; one pass through all five is a *sweep*.

1. **The coefficient path** $\theta_0, \ldots, \theta_T$, given Q and the covariances R^T: a Kalman filter runs forward, and the Carter–Kohn backward pass then draws θ_T and each θ_t conditional on the draw of θ_{t+1}. This is the Kalman *smoother* run as a sampler.
2. **The drift covariance** Q, given the path: an inverse-Wishart draw whose scale adds the T outer products $v_t v_t^\top$ of the increments to $T_0 \bar{Q}$.
3. **The volatility scales** σ_i^2, given the volatility paths: inverse-gamma draws.
4. **The covariance parameters** β, given the VAR residuals and H^T: two Gaussian regressions.
5. **The volatility paths** H^T: the single-site Metropolis step of Jacquier *et al.* [1994], one date at a time.

Any fixed order of valid conditional updates leaves the posterior invariant.

The order needs care only in samplers that represent the volatility likelihood with mixture indicators, as in Kim *et al.* [1998]; Del Negro and Primiceri [2015] show that Primiceri [2005] drew those indicators in the wrong place in the cycle.

The Jacquier–Polson–Rossi step has no such indicators.

107.6.1 Imposing stability

The restricted posterior assigns zero density to paths that are explosive at any date, including θ_0.

Cogley and Sargent simulate the *unrestricted* posterior and keep only the sweeps whose coefficient path is stable.

That is valid because the restricted posterior is the unrestricted posterior conditioned on stability, but if a is the share of unrestricted sweeps that are stable, it takes about $1/a$ complete sweeps to produce one that is kept.

We impose stability inside block 1 instead, by redrawing the backward pass until it delivers a stable path.

An accepted draw comes from the Gaussian Carter–Kohn distribution conditioned on stability, which is exactly the restricted full conditional, so the step is an exact Gibbs update.

Because the forward filter does not change between redraws, each redraw costs only a backward pass.

In rare sweeps the conditioning values make a stable draw very unlikely; if 1,000 passes all fail, the step keeps the current path.

The probability of that event depends only on the other blocks, not on the current path, so the step is a mixture of an exact draw and staying put, and it still leaves the restricted conditional invariant.

The two schemes are linked by an identity.

Collect the other blocks in $\Omega = (Q, H^T, \beta, \sigma)$ and let $p(\Omega)$ be the probability that one unrestricted Carter–Kohn draw is stable.

The restricted posterior of Ω is the unrestricted posterior $\pi_U(\Omega)$ reweighted by $p(\Omega)$, and the number of backward passes in a sweep is geometric with mean $1/p(\Omega)$, so at stationarity the mean number of passes per sweep is

$$\int \frac{1}{p(\Omega)} \frac{p(\Omega)\,\pi_U(\Omega)}{a}\, d\Omega = \frac{1}{a}.$$

Our sampler therefore pays in cheap backward passes what the discard scheme pays in complete sweeps, and the average number of passes estimates $1/a$, where a is Cogley and Sargent's share of stable sweeps.

Because rare sweeps need very many passes, the average is noisy, and the cap at 1,000 passes biases it down, so the implied a is an upper bound.

Another exact option is the elliptical slice sampler of Murray *et al.* [2010], which moves along an ellipse through the current path and a fresh Carter–Kohn draw and shrinks toward the current stable path.

Its new path stays close to the current one, whereas a redraw is independent of it, but in our experiments the slow mixing of both samplers came mainly from the dependence between Q and the coefficient path, which neither step addresses.

Appendix: the sampler in detail writes out the recursions.

The code is long, so it is collapsed; expand the cells to read it.

The first cell implements block 1.

The second cell implements blocks 2–5.

The last cell composes the five blocks, after a warm-up that lets the volatilities settle with the coefficients held at $\bar{\theta}$.

Passing `fixed_q` holds Q at that value instead of drawing it; a matrix of zeros nests the constant-coefficient model.

Cogley and Sargent ran 100,000 sweeps, discarded the first half, and kept every tenth of the rest.

We run 40,000 sweeps, discard the first 4,000, and keep every 20th of the rest for the posterior summaries, while recording $\mathrm{tr}(Q)$ and the number of backward passes at every sweep.

```
burn = 4_000
posterior = run_sampler(
    data['y'], data['x'], prior,
    n_sweeps=40_000, burn=burn, thin=20, seed=42,
)
passes = posterior['passes'][burn:]

pd.Series({
    'run time (minutes)': posterior['minutes'],
    'retained draws': len(posterior['β']),
    'mean backward passes per sweep': passes.mean(),
    'share of sweeps stable at the first pass': np.mean(passes == 1),
    'implied share a of stable unrestricted sweeps': 1 / passes.mean(),
    'sweeps that kept the current path': np.sum(passes > 1_000),
}, name='value').to_frame().round(3)
```

```
                                                value
run time (minutes)                             14.696
retained draws                               1800.000
mean backward passes per sweep                  4.686
share of sweeps stable at the first pass        0.458
implied share a of stable unrestricted sweeps   0.213
sweeps that kept the current path               4.000
```

The implied share a is well below one half, so the discard scheme would throw away most of its sweeps, each of which costs a forward filter and four other blocks, while our sampler pays for each failure with a backward pass alone.

Successive sweeps remain strongly correlated, especially in the direction of Q, so *Monte Carlo error* measures the Monte Carlo error in the numbers quoted below.

107.7 What the data say

We summarize the posterior by its mean coefficient path $\mathbb{E}(\theta_t \mid T)$ and mean covariance path $\mathbb{E}(R_t \mid T)$, and then interpret them in the context of the question we asked.

107.7.1 The rate and structure of drift

The trace of Q measures the total rate of coefficient drift, with $\mathrm{tr}(Q) = 0$ corresponding to constant coefficients.

The histogram shows $\mathrm{tr}(Q)$ at every sweep after the burn-in, together with the prior scale $\mathrm{tr}(\bar{Q})$.

```
trace_q = posterior['trace_q'][burn:]
fig, ax = plt.subplots()
ax.hist(trace_q, bins=50, histtype='step', lw=2)
ax.axvline(trace_q_bar, color='C1', lw=2, label=r'prior $\mathrm{tr}(\bar Q)$')
ax.set_xlabel(r'$\mathrm{tr}(Q)$')
ax.set_ylabel('sweeps')
ax.legend()
plt.show()
```

The next table compares summaries of Q with the numbers that Cogley and Sargent report (their Table 1 and p. 272).

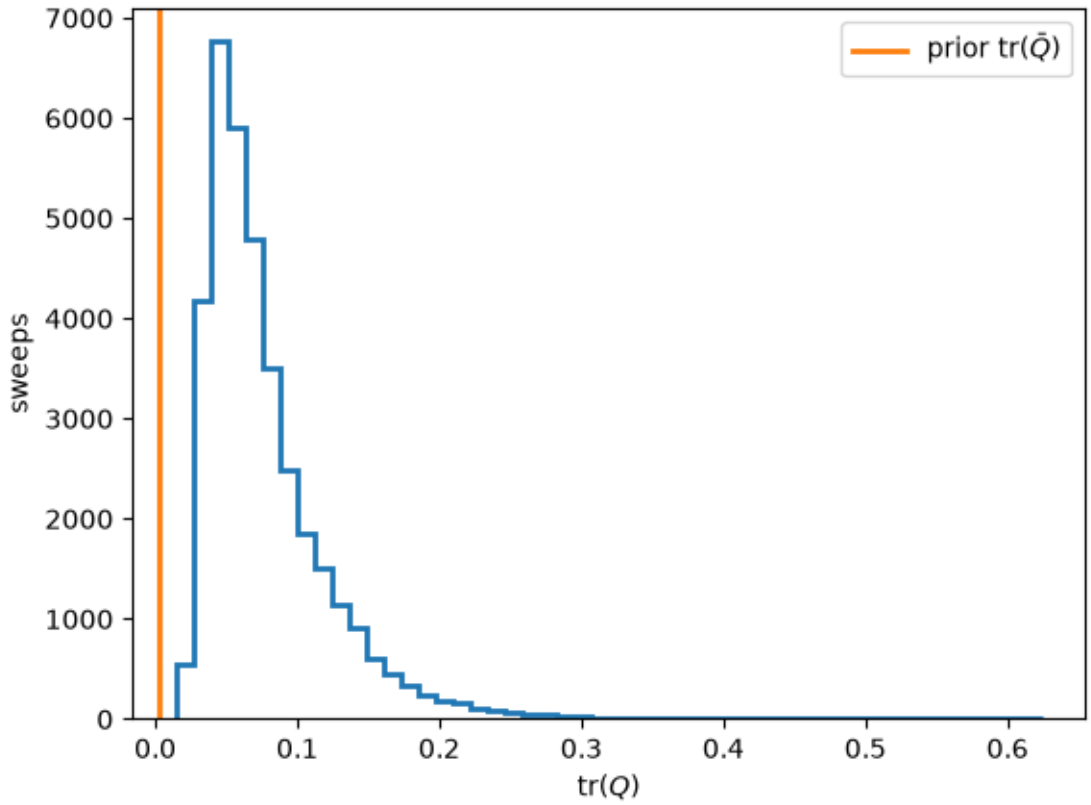


Fig. 107.2: Posterior tr(Q) and prior tr($\bar{Q}$)

```
q_mean = posterior['Q'].mean(axis=0)
q_eigenvalues = np.linalg.eigvalsh(q_mean)[::-1]
retained_trace_q = np.trace(posterior['Q'], axis1=1, axis2=2)

pd.DataFrame({
    'this lecture': [
        trace_q.mean(),
        q_eigenvalues[0],
        retained_trace_q.min() / trace_q_bar,
    ],
    'Cogley and Sargent': [0.047, 0.023, 3.0],
}, index=[
    'posterior mean of tr(Q)',
    'largest eigenvalue of posterior-mean Q',
    'smallest retained tr(Q) / tr(Q bar)',
]).round(3)
```

```
                                    this lecture  Cogley and Sargent
posterior mean of tr(Q)                    0.077               0.047
largest eigenvalue of posterior-mean Q     0.034               0.023
smallest retained tr(Q) / tr(Q bar)        6.145               3.000
```

The whole posterior lies far to the right of the conservative prior scale, so within this model the data call for much more coefficient drift than the prior anticipated.

The posterior mean of $\mathrm{tr}(Q)$ and the largest eigenvalue are larger than Cogley and Sargent's.

The posterior of $\mathrm{tr}(Q)$ has a long right tail that the chain explores slowly, so its mean carries noticeable Monte Carlo error (*Monte Carlo error*), but it was above Cogley and Sargent's value in every run we made with other seeds as well.

This is not a formal test of $Q = 0$, because the continuous prior assigns no mass to $Q = 0$; Cogley and Sargent avoid Bayes factors because their prior on Q is so diffuse.

Variation attributed to θ_t is variation in the *systematic* part of the reduced form; the VAR does not by itself say whether policy or private behavior caused it.

Cogley and Sargent estimated every ordering of the three variables, and their posterior means show that the ordering changes magnitudes but does not remove drift:

Ordering	Stable $\mathrm{tr}(Q)$	Stable $\max(\lambda)$	Unrestricted $\mathrm{tr}(Q)$	Unrestricted $\max(\lambda)$
(i, π, u)	0.055	0.025	0.056	0.027
(i, u, π)	0.047	0.023	0.059	0.031
(π, i, u)	0.064	0.031	0.082	0.044
(π, u, i)	0.062	0.031	0.088	0.051
(u, i, π)	0.057	0.026	0.051	0.028
(u, π, i)	0.055	0.024	0.072	0.035

Here $\max(\lambda)$ is the largest eigenvalue of the posterior mean of Q.

The ordering matters because the first two variables share sources of stochastic volatility with the others, while the third has its own; following the paper, we use (i, u, π), the ordering with the least drift.

For that ordering, removing the stability restriction raises the posterior mean drift rate, so the restriction does not manufacture the drift.

Diagonalizing the posterior mean of Q shows that the drift is low dimensional.

```
q_shares = np.cumsum(q_eigenvalues) / q_eigenvalues.sum()
pd.DataFrame({
    'eigenvalue': q_eigenvalues[:3],
    'cumulative share': q_shares[:3],
    'Cogley and Sargent': [0.485, 0.832, 0.945],
}, index=pd.Index([1, 2, 3], name='principal component')).round(3)
```

principal component	eigenvalue	cumulative share	Cogley and Sargent
1	0.034	0.448	0.485
2	0.029	0.824	0.832
3	0.009	0.936	0.945

As Cogley and Sargent found, only a few linear combinations of the 21 coefficients drift: the first principal component accounts for more than two-fifths of total drift and the first three for the large majority.

The posterior mean coefficient paths show where the drift lives.

```
θ_mean = posterior['θ'].mean(axis=0)
dates = data['dates']
coefficient_labels = ('constant', r'$i_{t-1}$', r'$u_{t-1}$', r'$\pi_{t-1}$',
                      r'$i_{t-2}$', r'$u_{t-2}$', r'$\pi_{t-2}$')
equation_labels = ('interest', 'unemployment', 'inflation')

fig, axes = plt.subplots(1, 3, figsize=(10, 4), sharex=True)
for equation, ax in enumerate(axes):
    rows = slice(7 * equation, 7 * equation + 7)
    lines = ax.plot(dates, θ_mean[rows].T, lw=2)
    ax.axhline(0, color='0.65', lw=1)
    ax.set_xlabel('year')
    ax.set_title(f'{equation_labels[equation]} equation')
fig.legend(lines, coefficient_labels, loc='lower center', ncol=4)
plt.tight_layout(rect=(0, 0.17, 1, 1))
plt.show()
```



Fig. 107.3: Posterior mean VAR coefficients $E(\theta_t \mid T)$

The interest and unemployment equations are comparatively stable, whereas coefficients of the inflation equation, especially the one on lagged interest, move strongly through the 1970s and turn around near 1980.

Individual lag coefficients should not be given structural interpretations, because paired lags can offset one another; the dynamic summaries below are more informative.

107.7.2 The evolution of volatility

We first ask how the *size* of the shocks changed.

The posterior mean $\mathbb{E}(R_t \mid T)$ averages $B^{-1}H_tB^{-1\prime}$ over the retained draws.

```
def mean_innovation_covariance(h_draws, β_draws):
    """E(R_t | T): the average of B^{-1} H_t B^{-1}' over retained draws."""
    b_inverse = np.linalg.inv(np.stack([b_matrix(β) for β in β_draws]))
    total = np.einsum('dij,dtj,dkj->tik', b_inverse, h_draws, b_inverse,
                      optimize=True)
    return total / len(β_draws)


r_mean = mean_innovation_covariance(posterior['h'], posterior['β'])
```

The next plot shows the innovation standard deviations, in basis points at quarterly rates, and the innovation correlations.

```
variances = ((0, 'Nominal interest'), (2, 'Inflation'), (1, 'Unemployment'))
correlations = ((0, 1, 'Interest--unemployment'),
                (0, 2, 'Interest--inflation'),
                (2, 1, 'Inflation--unemployment'))
fig, axes = plt.subplots(3, 2, figsize=(9, 8), sharex=True)
for row, (index, label) in enumerate(variances):
    axes[row, 0].plot(dates, 1e4 * np.sqrt(r_mean[:, index, index]), lw=2)
    axes[row, 0].set_title(label)
for row, (left, right, label) in enumerate(correlations):
    scale = np.sqrt(r_mean[:, left, left] * r_mean[:, right, right])
    axes[row, 1].plot(dates, r_mean[:, left, right] / scale, lw=2)
    axes[row, 1].set_title(label)
axes[1, 0].set_ylabel(r'standard deviation $\times 10^4$')
axes[1, 1].set_ylabel('correlation')
axes[-1, 0].set_xlabel('year')
axes[-1, 1].set_xlabel('year')
plt.tight_layout()
plt.show()
```

The interest-rate and inflation innovation standard deviations spike around 1980 and fall back within a few years, while unemployment innovation volatility trends down over the whole sample, punctuated by countercyclical bumps.

All three innovation correlations are largest in magnitude around the Volcker disinflation, so both the size and the composition of reduced-form shocks changed.

These movements give the changing-shocks, or bad-luck, explanation an important role, although the interest-rate innovation is not itself a structural monetary-policy shock.

Indeed, its positive correlation with inflation innovations and negative correlation with unemployment innovations look more like a policy *reaction* than like policy shocks, which conventional wisdom says should move inflation and unemployment the other way.

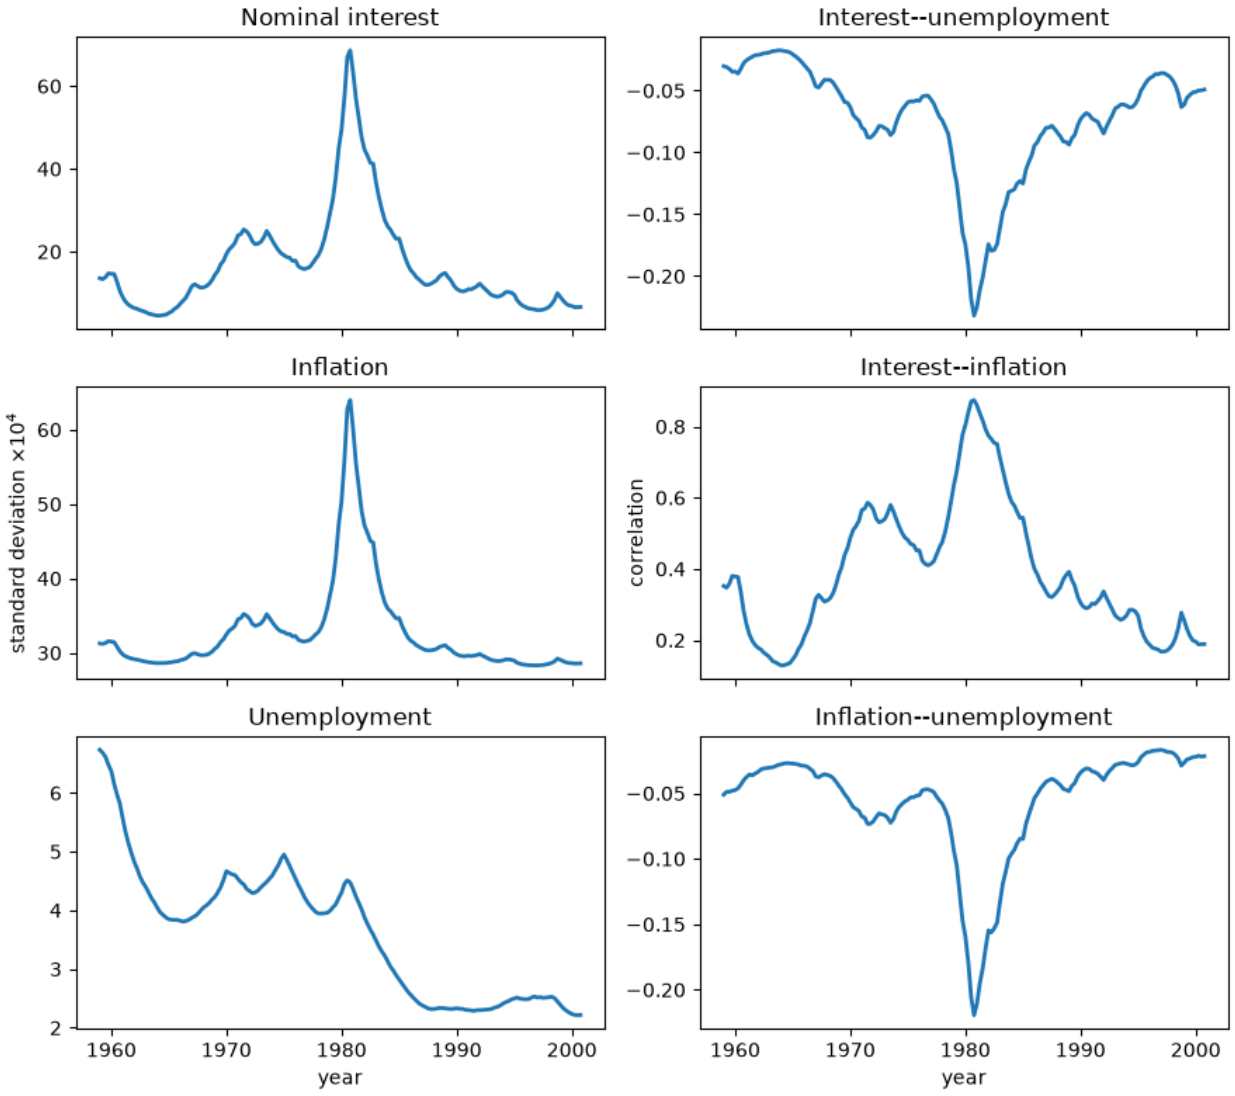


Fig. 107.4: Standard deviations and correlations implied by $E(R_t \mid T)$

Following Whittle [1953], the log determinant of $\mathbb{E}(R_t \mid T)$ summarizes the total one-step uncertainty entering the system at each date.

```
years = np.floor(dates + 1e-8).astype(int)
logdet_r = np.linalg.slogdet(r_mean)[1]

fig, ax = plt.subplots()
ax.plot(dates, logdet_r, lw=2)
ax.set_xlabel('year')
ax.set_ylabel(r'$\log |E(R_t\mid T)|$')
plt.show()
```

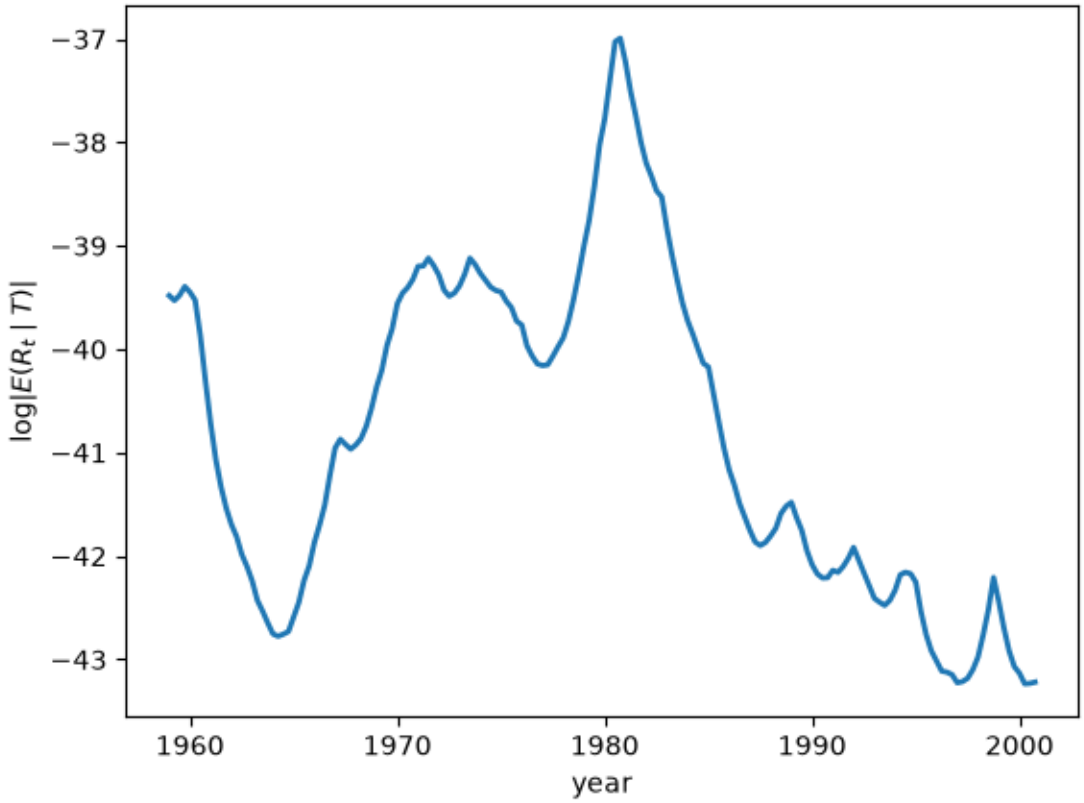


Fig. 107.5: Generalized innovation variance $\log|E(R_t \mid T)|$

The table below gives the numbers behind the figure.

```
def period_mean(values, first, last):
    """Average a quarterly series over the years first, ..., last."""
    return values[(years >= first) & (years <= last)].mean()


pd.Series({
    'mean, 1960--1964': period_mean(logdet_r, 1960, 1964),
    'peak': logdet_r.max(),
    'date of peak': dates[np.argmax(logdet_r)],
    'mean, 1990--2000': period_mean(logdet_r, 1990, 2000),
}, name='log |E(R_t | T)|').to_frame().round(2)
```

```
                   log |E(R_t | T)|
mean, 1960--1964             -41.61
peak                         -36.99
date of peak                1980.75
mean, 1990--2000             -42.59
```

Total innovation variance rises in two steps, from the mid-1960s to the early 1970s and again in the late 1970s, to a peak at the end of 1980, and most of the subsequent decline comes in the first half of the 1980s.

Its 1990s average is at or below its early-1960s average, so the Great Moderation documented by Kim and Nelson [1999] and McConnell and Perez-Quiros [2000] looks like a return to the stability of the early 1960s.

Cogley and Sargent read this as a *return* to stability, and they note how short-lived the earlier calm proved to be.

This is the clearest aggregate evidence for changing luck, but it cannot explain the coefficient-based changes in inflation dynamics examined next.

107.7.3 Core inflation and the natural rate

To study the systematic dynamics, write the VAR at date t in companion form as

$$z_t = \mu_{t|T} + A_{t|T} z_{t-1} + \begin{bmatrix} \varepsilon_t \\ 0 \end{bmatrix},$$

where $z_t = (y_t^\top, y_{t-1}^\top)^\top$ and $\mu_{t|T}$ and $A_{t|T}$ are built from $\mathbb{E}(\theta_t \mid T)$.

At date t, the local mean m_t is the point to which the companion system would converge if its coefficients stayed at their date-t values and future innovations were zero:

$$m_t = (I - A_{t|T})^{-1} \mu_{t|T}, \qquad \bar{\pi}_t = 4 s_\pi m_t, \qquad \bar{u}_t = \frac{\exp(100 s_u m_t)}{1 + \exp(100 s_u m_t)}. \tag{107.7}$$

Here s_π and s_u select inflation and unemployment from m_t, the factor four annualizes inflation, and the inverse logit returns unemployment to a rate.

Core inflation $\bar{\pi}_t$ is the long-horizon inflation forecast implied by freezing the date-t coefficients, and the natural rate $\bar{u}_t$ is the corresponding long-run unemployment anchor.

They are steady states of frozen local systems, not unconditional means of the drifting process.

Freezing the current coefficients and projecting forward is the *anticipated-utility* device [Kreps, 1998] of the learning governments of *Adaptive Learning and Escape Dynamics*, *Escaping Nash Inflation*, and *The Lost Conquest: Fed Policy in the 2020s*, which act as if their current beliefs will never be revised.

```
def local_means(θ_path):
    """Core inflation (annual rate) and the natural unemployment rate."""
    θ_path = np.asarray(θ_path, dtype=float).reshape(n_coefficients, -1)
    core = np.empty(θ_path.shape[1])
    natural = np.empty(θ_path.shape[1])
    for t, θ in enumerate(θ_path.T):
        intercept, companion = companion_matrix(θ)
        m = np.linalg.solve(np.eye(6) - companion, intercept)
        core[t], natural[t] = 4 * m[2], expit(100 * m[1])
    return core, natural


core_inflation, natural_rate = local_means(θ_mean)
```

We plot the fourth-quarter value in each year.

```
annual = np.array([np.flatnonzero(years == year)[-1]
                   for year in np.unique(years)])

fig, ax = plt.subplots()
ax.plot(dates[annual], 100 * core_inflation[annual], 'o-', lw=2,
        markersize=3, label='core inflation')
ax.plot(dates[annual], 100 * natural_rate[annual], '+-', lw=2,
        markersize=5, label='natural rate')
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()
```

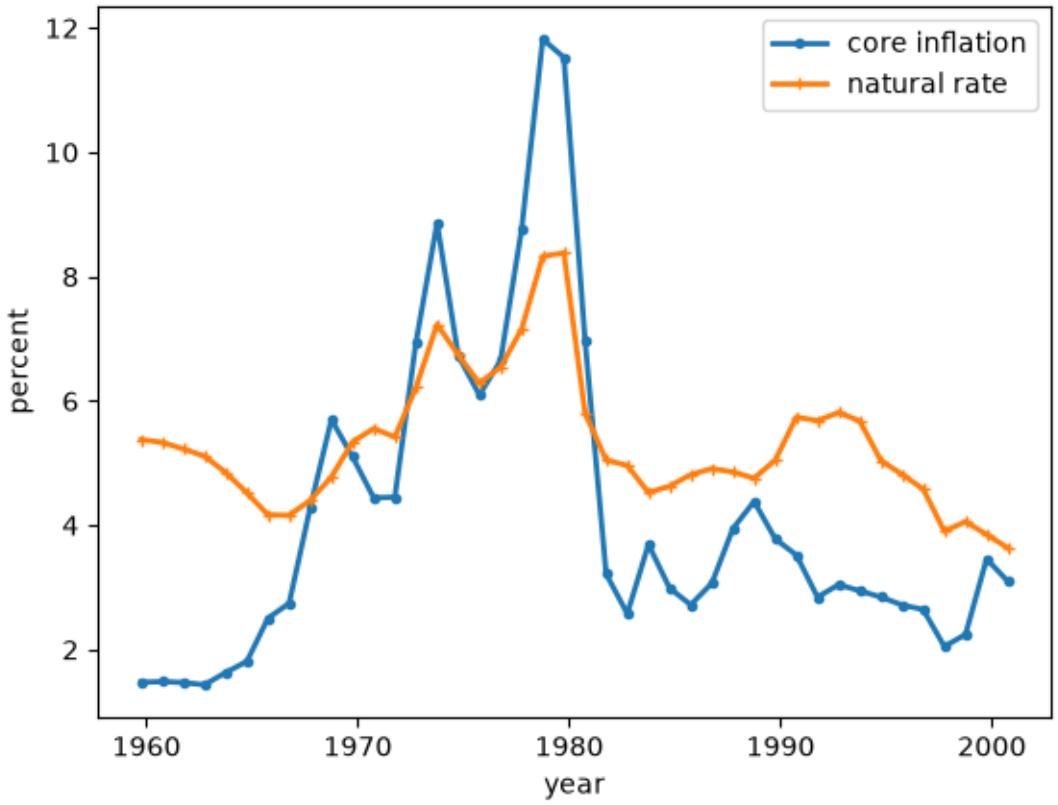


Fig. 107.6: Core inflation $\bar{\pi}_t$ and the natural rate $\bar{u}_t$

The table below gives the numbers behind the figure.

```
pd.Series({
    'core inflation, 1960--1964 mean': 100 * period_mean(core_inflation, 1960, 1964),
    'core inflation, peak': 100 * core_inflation[annual].max(),
    'year of the core peak': years[annual][np.argmax(core_inflation[annual])],
    'core inflation, 1985--2000 mean': 100 * period_mean(core_inflation, 1985, 2000),
    'natural rate, 1960': 100 * natural_rate[annual][0],
    'natural rate, peak': 100 * natural_rate[annual].max(),
    'year of the natural-rate peak': years[annual][np.argmax(natural_rate[annual])],
    'natural rate, 2000': 100 * natural_rate[-1],
```

(continues on next page)

(continued from previous page)

```
    'correlation of core and natural rate': np.corrcoef(core_inflation, natural_
↪rate)[0, 1],
}, name='estimate').to_frame().round(2)
```

	estimate
core inflation, 1960--1964 mean	1.53
core inflation, peak	11.82
year of the core peak	1978.00
core inflation, 1985--2000 mean	3.09
natural rate, 1960	5.38
natural rate, peak	8.38
year of the natural-rate peak	1979.00
natural rate, 2000	3.64
correlation of core and natural rate	0.84

Core inflation climbs from about 1.5 percent in the early 1960s to a peak in the late 1970s and then falls to about 3 percent, the pattern Cogley and Sargent report (about 1.5, 8, and 2.5 to 3.5 percent), although our peak is higher.

The natural rate rises from about 5 percent to a peak in the late 1970s and falls below 4 percent by 2000.

Because both lines are determined by the fitted coefficients rather than by realized shocks, their persistent shifts point to a changing systematic component.

The two anchors are highly correlated, more so than the 0.748 that Cogley and Sargent report: the model's long-run inflation and unemployment rates rose and fell together, which is not the same as current inflation and unemployment moving together.

Because $(I - A_t)^{-1}$ amplifies small coefficient changes when the largest root is close to one, long-run means are intrinsically more sensitive than short-horizon forecasts.

107.7.4 Inflation persistence

We now ask whether the *systematic* dynamics drifted on top of the moving volatilities.

The main summary is inflation persistence, measured by the normalized spectrum of inflation at frequency zero.

The spectral density of inflation at date t is

$$f_{\pi\pi}(\omega, t) = \frac{1}{2\pi} s_\pi (I - A_{t|T} e^{-i\omega})^{-1} \mathcal{R}_t (I - A_{t|T}^\top e^{i\omega})^{-1} s_\pi^\top, \tag{107.8}$$

where $\mathcal{R}_t$ embeds $\mathbb{E}(R_t \mid T)$ in the companion system.

Low-frequency power depends on both the autoregressive coefficients and the innovation covariance, so Cogley and Sargent normalize by the variance of inflation at date t,

$$g_{\pi\pi}(\omega, t) = \frac{f_{\pi\pi}(\omega, t)}{\int_{-\pi}^{\pi} f_{\pi\pi}(\omega, t) d\omega}, \tag{107.9}$$

which measures autocorrelation rather than autocovariance.

They interpret $g_{\pi\pi}(0, t)$ as a measure of inflation persistence.

Statistics at frequency zero are long-run statistics: *Two Illustrations of the Quantity Theory of Money* computes long-run slopes as ratios of spectral densities at frequency zero, and Exercise **??** computes such a slope of inflation on unemployment from this model.

The normalization removes a common scale factor from R_t, but it can still depend on the relative variances and covariances in R_t.

The next function evaluates both spectra at frequencies measured in cycles per quarter.

```
def inflation_spectrum(θ, covariance, frequencies):
    """Inflation power f(ω) and normalized power g(ω) at one date."""
    _, companion = companion_matrix(θ)
    innovation = np.zeros((6, 6))
    innovation[:3, :3] = covariance
    variance = linalg.solve_discrete_lyapunov(companion, innovation)[2, 2]
    power = np.empty(len(frequencies))
    for j, frequency in enumerate(frequencies):
        transfer = np.linalg.inv(
            np.eye(6) - companion * np.exp(-2j * np.pi * frequency))
        power[j] = np.real(transfer @ innovation @ transfer.conj().T)[2, 2]
    return power / (2 * np.pi), power / (2 * np.pi * variance)


def persistence_path(θ_path, r_path):
    """g_ππ(0, t) at every date."""
    return np.array([inflation_spectrum(θ, r, [0.0])[1][0]
                     for θ, r in zip(θ_path.T, r_path)])


inflation_persistence = persistence_path(θ_mean, r_mean)
white_noise = 1 / (2 * np.pi)
```

For white noise the spectrum is flat, so $g_{\pi\pi}(0) = 1/(2\pi) \approx 0.16$; that is the benchmark against which to read the path.

```
fig, ax = plt.subplots()
ax.plot(dates[annual], inflation_persistence[annual], 'o-', lw=2, markersize=3)
ax.axhline(white_noise, color='0.45', ls='--', lw=1,
           label=r'white noise, $1/(2\pi)$')
ax.set_xlabel('year')
ax.set_ylabel(r'$g_{\pi\pi}(0,t)$')
ax.legend()
plt.show()
```

The table below gives the numbers behind the figure.

```
pd.Series({
    '1960--1964 mean': period_mean(inflation_persistence, 1960, 1964),
    '1970--1979 mean': period_mean(inflation_persistence, 1970, 1979),
    'peak': inflation_persistence[annual].max(),
    'year of the peak': years[annual][np.argmax(inflation_persistence[annual])],
    '1985--2000 mean': period_mean(inflation_persistence, 1985, 2000),
}, name='g(0, t)').to_frame().round(3)
```

```
                    g(0, t)
1960--1964 mean       0.180
1970--1979 mean       3.839
peak                  6.886
year of the peak   1979.000
1985--2000 mean       0.350
```

Persistence is close to the white-noise value in the early 1960s, sweeps upward in the late 1960s, stays high through the 1970s with a peak in 1979, and falls sharply after 1980, back toward the white-noise value for much of the 1980s and 1990s.

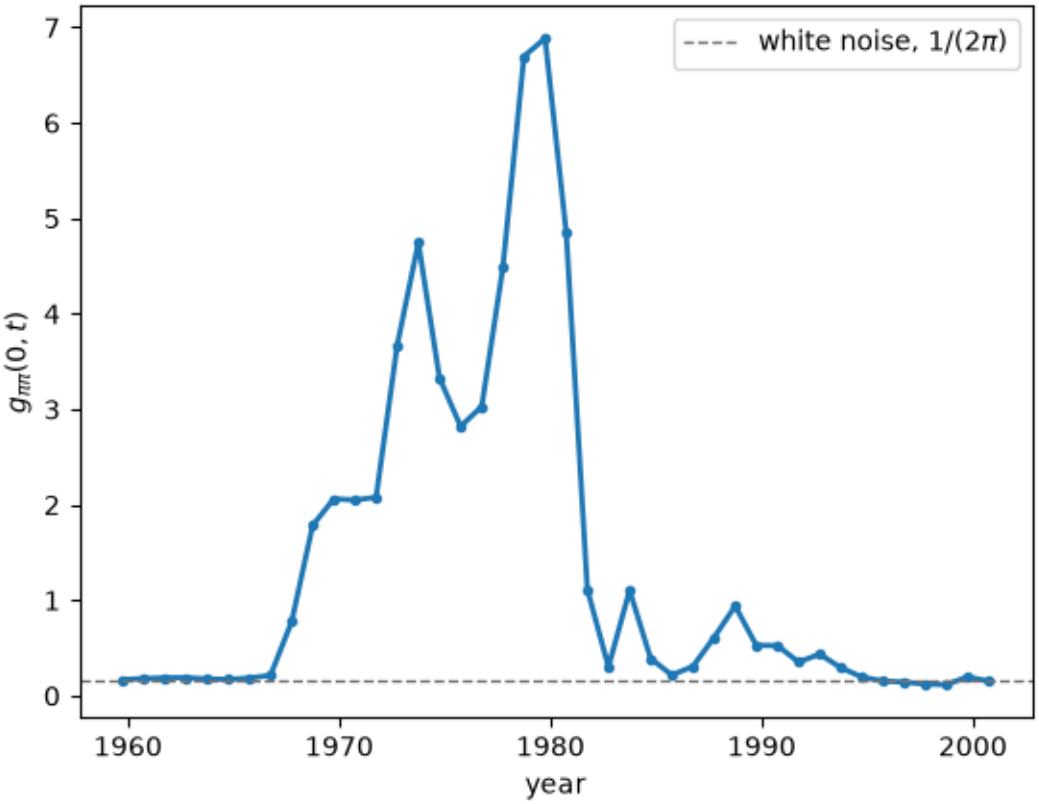


Fig. 107.7: Normalized zero-frequency spectrum $g_{\pi\pi}(0,t)$

The post-1980 fall cannot be explained by a proportional rescaling of all innovations, although the normalized statistic can still depend on the composition of R_t.

For comparison, an $AR(1)$ with coefficient ρ has $g(0) = (1+\rho)/[2\pi(1-\rho)]$, so values between 2 and 10 correspond to ρ between about 0.85 and 0.97; Exercise 107.11.1 converts the whole path into such an "equivalent $AR(1)$" coefficient.

The zero-frequency path omits the rest of the frequency distribution, which the next heatmap shows.

```
spectrum_frequencies = np.linspace(0, 0.5, 41)
normalized_spectrum = np.column_stack([
    inflation_spectrum(θ_mean[:, t], r_mean[t], spectrum_frequencies)[1]
    for t in range(len(dates))
])

fig, ax = plt.subplots(figsize=(8, 4))
image = ax.pcolormesh(dates, spectrum_frequencies,
                      np.log10(normalized_spectrum), shading='auto')
ax.set_xlabel('year')
ax.set_ylabel('cycles per quarter')
fig.colorbar(image, ax=ax, label=r'$\log_{10} g_{\pi\pi}(\omega,t)$')
plt.show()
```

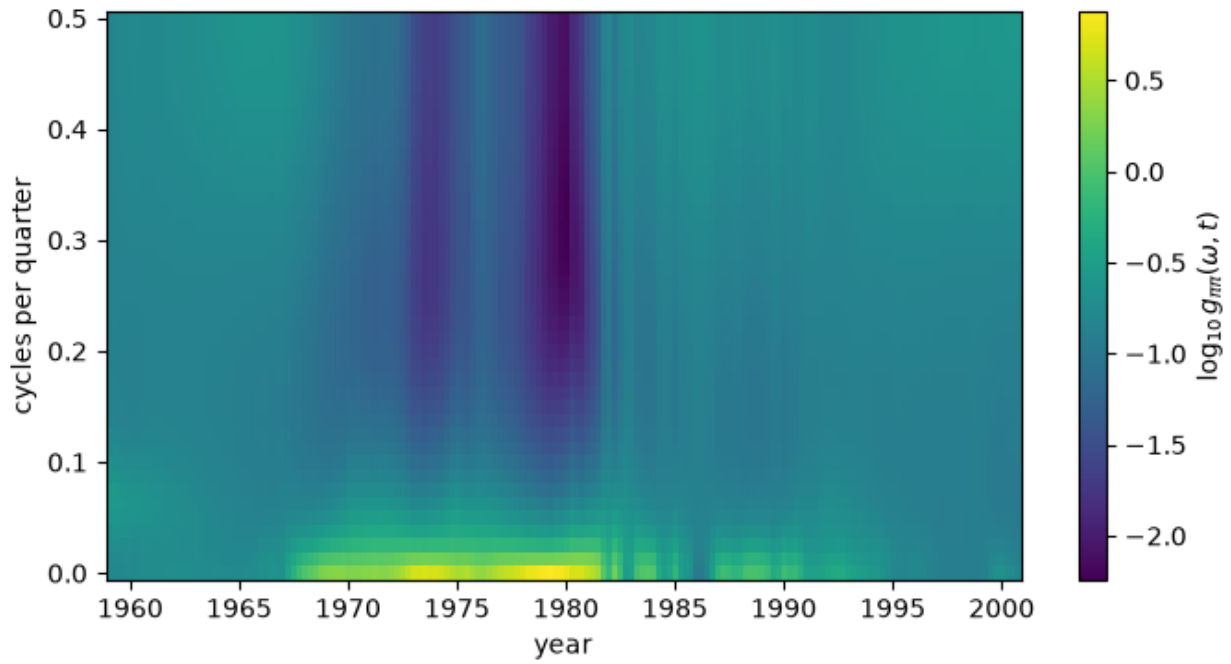


Fig. 107.8: Normalized inflation spectrum $g_{\pi\pi}(\omega, t)$

A ridge of low-frequency power runs from the late 1960s to the early 1980s and then recedes.

Because the ridge survives normalization, the Great Inflation was not only a period of large shocks: inflation shocks were also propagated more persistently.

Point estimates do not reveal how strongly the data locate these paths, so we compute the annual features from every retained draw.

```
def posterior_feature_draws(draws, indices):
    """Core inflation, the natural rate, and g(0) at selected dates, by draw."""
    n_draws = len(draws['β'])
    core = np.empty((len(indices), n_draws))
    natural = np.empty_like(core)
    persistence = np.empty_like(core)
    for d in range(n_draws):
        θ_d = draws['θ'][d][:, indices]
        core[:, d], natural[:, d] = local_means(θ_d)
        for row, t in enumerate(indices):
            covariance = innovation_covariance(draws['h'][d][t], draws['β'][d])
            persistence[row, d] = inflation_spectrum(
                θ_d[:, row], covariance, [0.0])[1][0]
    return {'core': 100 * core, 'natural': 100 * natural,
            'persistence': persistence}


feature_draws = posterior_feature_draws(posterior, annual)
```

The figure plots, for each feature, the posterior median and a 90 percent band across draws.

```
feature_specs = (('core', 'core inflation (%)'),
                 ('natural', 'natural rate (%)'),
                 ('persistence', r'$g_{\pi\pi}(0,t)$'))
fig, axes = plt.subplots(3, 1, figsize=(9, 8), sharex=True)
for ax, (key, ylabel) in zip(axes, feature_specs):
    lower, median, upper = np.quantile(feature_draws[key], (0.05, 0.5, 0.95),
                                       axis=1)
    line, = ax.plot(dates[annual], median, lw=2)
    ax.fill_between(dates[annual], lower, upper, color=line.get_color(),
                    alpha=0.2)
    ax.set_ylabel(ylabel)
axes[-1].set_xlabel('year')
plt.tight_layout()
plt.show()
```

The bands are pointwise, not bands for whole paths.

The bands are wide and skewed to the right in the 1970s, so the rise and the post-1980 fall are clearer than their magnitudes.

As Cogley and Sargent note, uncertainty about core inflation grows with persistence, because core inflation is a long-horizon forecast.

The next plot compares the timing of core inflation and persistence, as in Cogley and Sargent's Figure 11.

```
fig, ax = plt.subplots()
ax.plot(dates[annual], 100 * core_inflation[annual], 'o-', lw=2,
        markersize=3, label='core inflation (%)')
ax.plot(dates[annual], inflation_persistence[annual], 'x-', lw=2,
        markersize=4, label=r'$g_{\pi\pi}(0,t)$')
ax.set_xlabel('year')
ax.legend()
plt.show()

core_persistence_correlation = np.corrcoef(
    core_inflation, inflation_persistence)[0, 1]
```

(continues on next page)

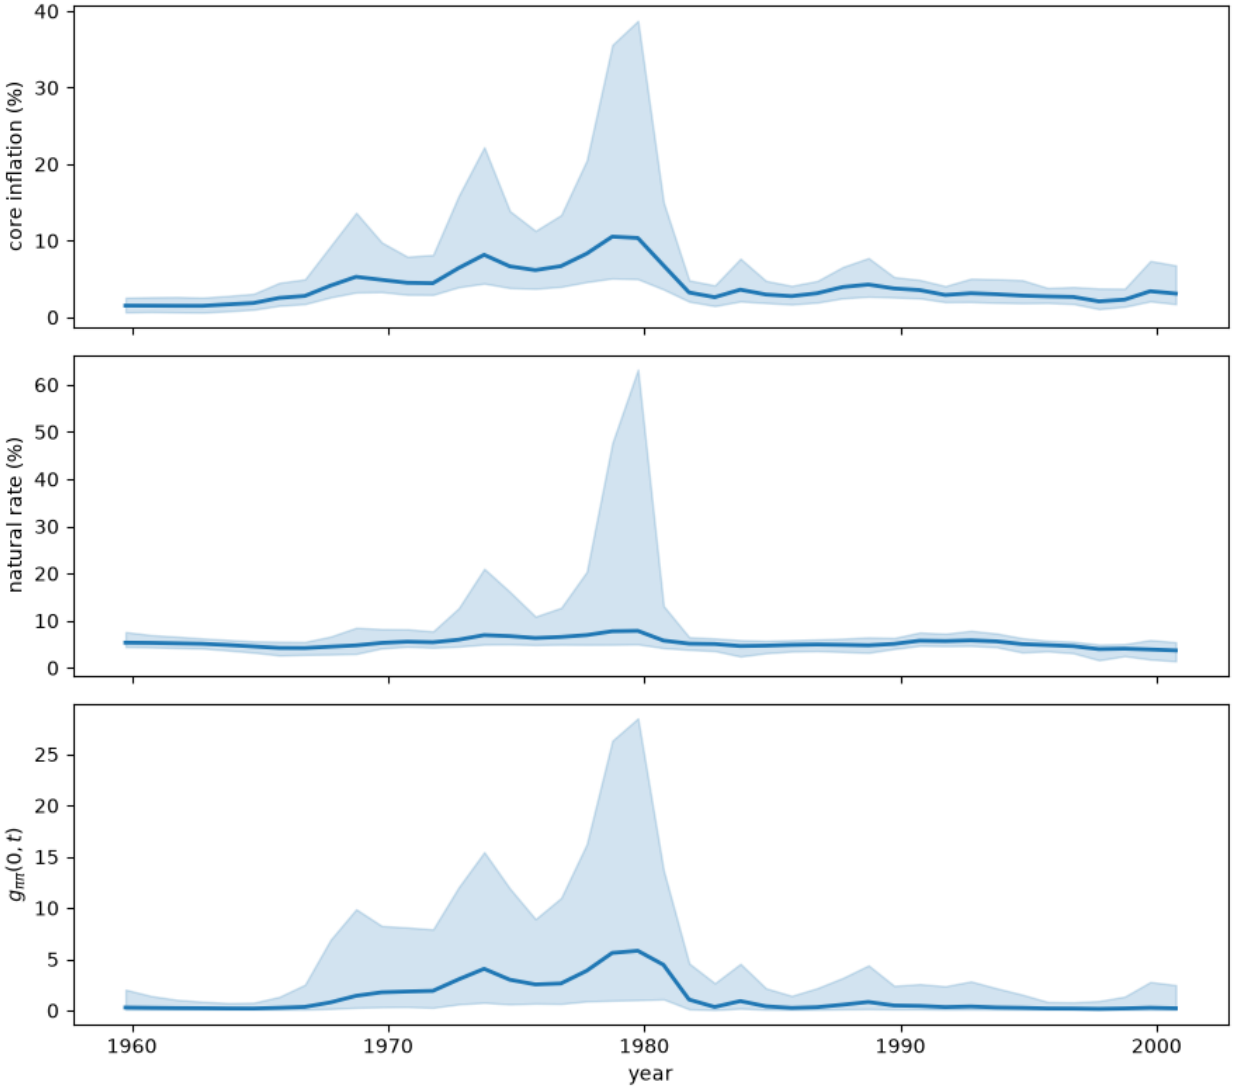


Fig. 107.9: Posterior medians and pointwise 90 percent intervals for $\bar{\pi}_t$, $\bar{u}_t$, and $g_{\pi\pi}(0,t)$

(continued from previous page)

```
print(f'correlation of core inflation and persistence: '
      f'{core_persistence_correlation:.3f}')
```

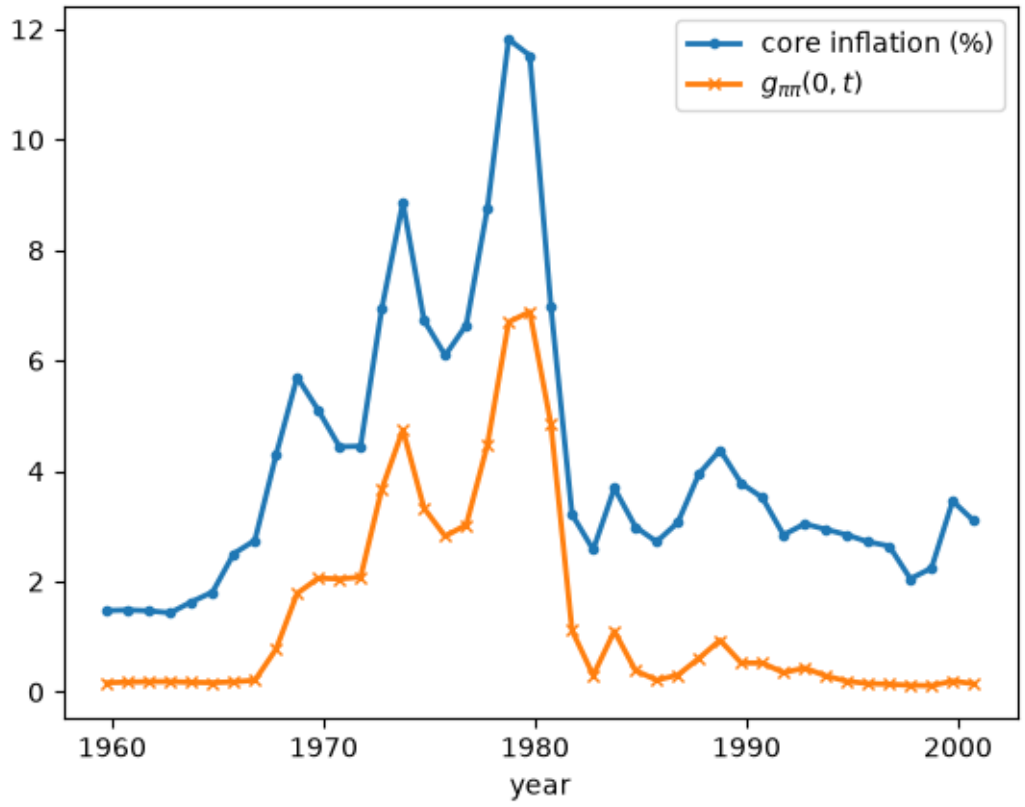


Fig. 107.10: Core inflation $\bar{\pi}_t$ and persistence $g_{\pi\pi}(0,t)$

```
correlation of core inflation and persistence: 0.965
```

Core inflation and persistence rise together through the late 1960s and 1970s and fall together during the Volcker disinflation.

The printed quarterly correlation is high and close to the 0.92 that Cogley and Sargent report.

It summarizes common timing rather than a causal relation: both measures are nonlinear functions of the same fitted coefficients.

Bad policy or bad luck? A verdict explains why this positive correlation troubled the escape-route models of *Escaping Nash Inflation*.

107.7.5 Monetary policy activism

Cogley and Sargent summarize systematic policy with the forward-looking Taylor rule of Clarida *et al.* [2000], with interest-rate smoothing,

$$i_t = \psi_0 + \psi_1 \mathbb{E}_t \bar{\pi}_{t,t+h_\pi} + \psi_2 \mathbb{E}_t \bar{u}_{t,t+h_u} + \psi_3 i_{t-1} + e_t, \tag{107.10}$$

where $\bar{\pi}_{t,t+h_\pi}$ and $\bar{u}_{t,t+h_u}$ are average inflation and unemployment over the next h_π and h_u quarters.

They define the activism coefficient $\mathcal{A}_t = \psi_1/(1-\psi_3)$ and call policy active when $\mathcal{A}_t \geq 1$, so that the nominal rate rises more than one-for-one with expected inflation in the long run.

The long-run response $\psi_1(1+\psi_3+\psi_3^2+\cdots)$ exists only when $|\psi_3| < 1$.

At each date, population two-stage least squares projections implied by the local VAR deliver the coefficients.

The first stage projects the Fed's forecasts on lagged information, namely the VAR's forecasts made at $t-1$, because the Fed does not see current-quarter inflation and unemployment when it sets the rate.

The second stage projects i_t on those fitted forecasts and i_{t-1}.

The benchmark horizons are $h_\pi = 4$ quarters and $h_u = 2$ quarters, reflecting Federal Reserve folk wisdom about policy lags.

```
def policy_rule_coefficients(θ, covariance, h_pi=4, h_u=2):
    """Population 2SLS coefficients (ψ_1, ψ_2, ψ_3) implied by a local VAR."""
    _, companion = companion_matrix(θ)
    innovation = np.zeros((6, 6))
    innovation[:3, :3] = covariance
    stationary = linalg.solve_discrete_lyapunov(companion, innovation)
    power, π_load, u_load = np.eye(6), np.zeros(6), np.zeros(6)
    for horizon in range(1, max(h_pi, h_u) + 1):
        power = power @ companion
        if horizon <= h_pi:
            π_load += power[2] / h_pi
        if horizon <= h_u:
            u_load += power[1] / h_u
    loadings = np.vstack((π_load, u_load, np.eye(6)[0]))
    return np.linalg.solve(loadings @ stationary @ loadings.T,
                           loadings @ stationary @ companion.T[:, 0])


rules = np.array([policy_rule_coefficients(θ_mean[:, t], r_mean[t])
                  for t in range(len(dates))])
ψ1, ψ3 = rules[:, 0], rules[:, 2]
activism = np.where(np.abs(ψ3) < 1, ψ1 / (1 - ψ3), np.nan)
print(f'share of dates with |ψ_3| < 1: {np.mean(np.abs(ψ3) < 1):.3f}')
```

```
share of dates with |ψ_3| < 1: 1.000
```

The figure plots $\mathcal{A}_t$ at the posterior mean, as in Cogley and Sargent's Figure 12, leaving blank any date with $|\psi_3| \geq 1$.

```
fig, ax = plt.subplots()
ax.plot(dates, activism, lw=2)
ax.axhline(1, color='0.45', lw=1)
ax.set_xlabel('year')
ax.set_ylabel(r'$\mathcal{A}_t$')
plt.show()
```

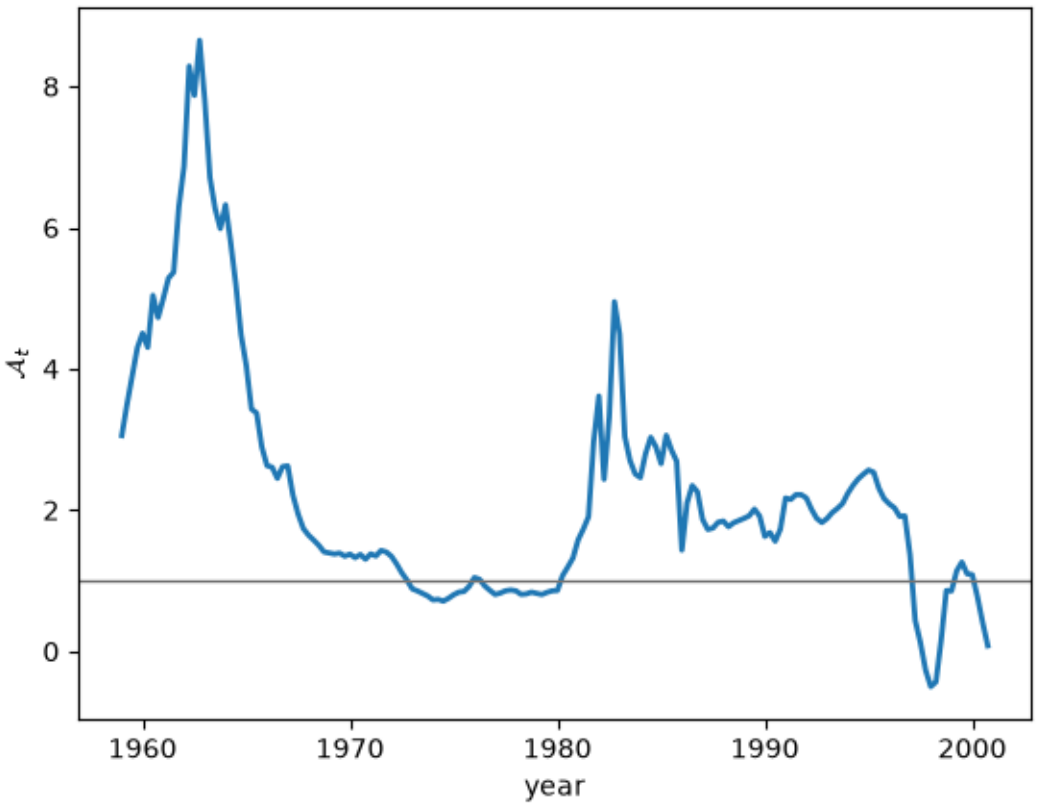


Fig. 107.11: Activism coefficient $\mathcal{A}_t = \psi_{1t}/(1 - \psi_{3t})$ at the posterior mean

Judged by point estimates, policy was active in the early 1960s, was mostly passive from the early 1970s until 1980, and turned active at the start of the Volcker disinflation, staying active through the mid-1990s.

Estimates at the very end of the sample are erratic, for a reason explained next.

The rule's coefficients are weakly identified when inflation is close to white noise, because lagged variables are then poor instruments for future inflation, so a path built from posterior-mean inputs understates the uncertainty.

We therefore compute $\mathcal{A}_t$ from every retained draw in 1975, 1985, and 1995 and compare the posterior probabilities with Cogley and Sargent's.

```
selected_years = (1975, 1985, 1995)
selected_dates = [np.flatnonzero(years == year)[-1] for year in selected_years]
n_draws = len(posterior['β'])
activism_draws = {year: np.empty(n_draws) for year in selected_years}
inside = {year: np.empty(n_draws, dtype=bool) for year in selected_years}
for d in range(n_draws):
    for year, t in zip(selected_years, selected_dates):
        covariance = innovation_covariance(posterior['h'][d][t],
                                           posterior['β'][d])
        ψ = policy_rule_coefficients(posterior['θ'][d][:, t], covariance)
        activism_draws[year][d] = ψ[0] / (1 - ψ[2])
        inside[year][d] = abs(ψ[2]) < 1

activism_events = (
    activism_draws[1975] > 1,
    activism_draws[1985] > 1,
    activism_draws[1995] > 1,
    activism_draws[1985] > activism_draws[1975],
    activism_draws[1995] > activism_draws[1975],
)
pd.DataFrame({
    'this lecture': [event.mean() for event in activism_events],
    'Cogley and Sargent': [0.208, 0.919, 0.941, 0.923, 0.943],
}, index=['P(A_1975 > 1)', 'P(A_1985 > 1)', 'P(A_1995 > 1)',
          'P(A_1985 > A_1975)', 'P(A_1995 > A_1975)']).round(3)
```

	this lecture	Cogley and Sargent
P(A_1975 > 1)	0.285	0.208
P(A_1985 > 1)	0.736	0.919
P(A_1995 > 1)	0.693	0.941
P(A_1985 > A_1975)	0.744	0.923
P(A_1995 > A_1975)	0.699	0.943

Like Cogley and Sargent, these probabilities use every draw; the share of draws with $|\psi_3| < 1$, for which $\mathcal{A}_t$ has its long-run interpretation, is

```
pd.Series({year: inside[year].mean() for year in selected_years},
          name='share of draws with |ψ_3| < 1').to_frame().round(3)
```

	share of draws with |ψ_3| < 1
1975	1.000
1985	0.961
1995	0.973

The posterior favors passive policy in 1975 and active policy in 1985 and 1995, and it says that activism probably rose.

But our probabilities of active policy in 1985 and 1995, and of a rise since 1975, are noticeably lower than Cogley and Sargent's, in every run we made.

That is consistent with the larger drift covariance in our posterior, which locates the coefficients at any date less precisely.

107.7.6 The disinflation and the calm

Escaping Nash Inflation ends by asking whether the calm of the mid-1980s and the disinflation arrived together, because a model in which calm makes escapes easier would predict that they should.

The fitted model separates the two: core inflation measures the disinflation, and the generalized innovation variance measures the calm.

The table dates the peak of each and the quarter in which each had covered half the distance from its peak to its post-1985 average.

```
def peak_and_halfway(series, first_after=1985):
    """Date of the peak, and of the first later date halfway to the later mean."""
    peak = np.argmax(series)
    later_mean = series[years >= first_after].mean()
    halfway = 0.5 * (series[peak] + later_mean)
    after = peak + np.flatnonzero(series[peak:] <= halfway)[0]
    return dates[peak], dates[after]


pd.DataFrame(
    [peak_and_halfway(core_inflation), peak_and_halfway(logdet_r)],
    index=['core inflation', 'log |E(R_t | T)|'],
    columns=['peak', 'halfway down'],
).round(2)
```

	peak	halfway down
core inflation	1979.25	1980.25
log |E(R_t | T)|	1980.75	1984.00

Core inflation fell first: it was halfway down a year after its peak, while total innovation variance peaked later and covered half of its decline only around 1984.

The disinflation preceded most of the calm, so the timing gives no support to the idea that a calm made the escape to low inflation possible, although a reduced-form VAR cannot settle questions of causation.

107.8 Why classical tests missed the drift

How can a model with this much drift coexist with the classical tests that failed to reject time invariance?

Cogley and Sargent's answer, in their section 3.7, is that those tests have little power against drift of the kind their posterior describes.

They first applied Bernanke and Mihov's test, the sup-LM test of Andrews [1993] for a single break at an unknown date, to their data, and like Bernanke and Mihov they could not reject.

A failure to reject is evidence against drift only if the test would probably have rejected had drift been present.

To measure that probability, they drew (θ^T, H^T, B) from their posterior, simulated an artificial sample from each draw, applied the test, and counted rejections at the 5 percent level in 10,000 replications.

They did the same for the test of Nyblom [1989] and Hansen [1992], which is designed for parameters that drift as random walks, and for the sup-Wald version of Andrews's test.

Rejection frequency under the posterior	Nominal interest	Unemployment	Inflation	VAR
Andrews sup-LM (Table 4)	0.136	0.172	0.112	0.252
Nyblom–Hansen (Table 5)	0.076	0.170	0.086	0.234
Andrews sup-Wald (Table 6)	0.173	0.269	0.711	0.296

On the actual data the sup-LM and Nyblom–Hansen tests fail to reject for every equation and for the VAR as a whole.

Yet even when θ drifts as the posterior says it does, these tests fail to reject about three times as often as they reject, or more.

The one case with better-than-even power is the sup-Wald test for the inflation equation, and on the actual data that test rejects time invariance at the 1 percent level; it also rejects for the VAR as a whole at 5 percent.

Tests aimed at the parameters of the Clarida *et al.* [2000] policy rule, or at the first principal component of θ_t, also turn out to have low power.

So the failures to reject by Bernanke and Mihov [1998] and Sims [1999] are what one should expect even if the coefficients drifted as much as they do here.

Low power is not surprising.

The drifting-coefficient model nests time invariance at $Q = 0$: for Q near zero power is close to size, for very large Q the tests reject almost surely, and in between lies a range of economically meaningful drift that the tests usually miss.

Relatedly, Stock and Watson [1998] show that when the drift variance is small its maximum likelihood estimate is often exactly zero, so a classical estimate of "no drift" is weak evidence as well.

Cogley and Sargent close with a deeper reason, clearest in continuous time [Anderson *et al.*, 2003].

The volatility of a diffusion can be recovered from its quadratic variation over an arbitrarily short stretch of a path, but its drift can be learned only slowly, over long spans of data.

Movements in the *systematic* part of a time series are therefore much harder to detect than movements in its *volatility*, which is exactly the asymmetry between the two halves of the bad-luck evidence: strong evidence of changing variances and weak evidence against changing coefficients.

Exercise 107.11.3 computes the power of the Nyblom–Hansen test in a simple autoregression with a drifting coefficient.

107.9 Bad policy or bad luck? A verdict

The Bayesian VAR delivers a nuanced answer to the question that opened this lecture.

- *Volatilities drifted:* the size of the shocks changed enormously, with a spike around 1980 and a return to calm in the mid-1980s, so the bad-luck story captures something real.
- *Coefficients drifted too:* core inflation and inflation persistence rose through the 1970s and fell in the 1980s, and the estimated policy rule moved from passive to active; Exercise 107.11.2 shows that freezing the coefficients cannot reproduce the post-1980 fall in persistence.
- *Classical tests are not counterevidence:* they have little power against drift of this kind.

The reduced-form VAR cannot by itself prove that changes in Federal Reserve beliefs caused the coefficient drift, because private behavior and other omitted mechanisms can also change reduced-form dynamics.

There is one more twist, and it loops us back to the theory of this section.

The escape-route models of *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation* make a sharp prediction about persistence.

At the high-inflation self-confirming equilibrium the government's inflation target barely moves, so inflation is nearly serially uncorrelated; in the static model it is white noise.

Measured persistence rises only *along the transition* from high to low inflation, when beliefs, and with them mean inflation, shift rapidly within any window of data.

These models therefore predict that persistence should peak during a disinflation and be low when inflation is high, an inverse relation between the level and the persistence of inflation in the words of Cogley and Sargent [2005] (footnote 28).

The data show the opposite: persistence was highest when core inflation was highest, and both fell together after 1980.

Exercise 107.11.4 checks the prediction in the model of *Escaping Nash Inflation*.

The positive correlation helped motivate a different mechanism, a policymaker who is *reluctant to disinflate*.

In Cogley and Sargent [2005], a Fed that averages over competing Phillips-curve models wants to reduce inflation in the 1970s but moves very slowly, because a model that warns of large unemployment costs still carries weight, and that gradualism makes deviations from target persistent precisely when inflation is highest.

In Primiceri [2006], policymakers who underestimated the natural rate and the persistence of inflation and overestimated the cost of disinflation are likewise slow to act in the 1970s and then change course.

The friendly debate with Sims, Zha, Bernanke, and Mihov thus did more than adjudicate a historical question: it sharpened the models of learning and drift that run through this section.

Cogley and Sargent's sample ends in 2000.

Drifts and Volatilities after 2000 re-estimates the model with data through the financial crisis, the zero lower bound, the pandemic, and the 2021–2022 inflation surge, and *The Lost Conquest: Fed Policy in the 2020s* then returns to a learning Fed to interpret the 2020s.

107.10 Appendix: the sampler in detail

107.10.1 The coefficient block

Conditional on Q and R^T, the measurement equation (107.1) and the random walk (107.2) form a linear Gaussian state-space model, and the forward recursion is the Kalman filter of *A First Look at the Kalman Filter*, started from $\theta_{0|0} = \bar{\theta}$ and $P_{0|0} = \bar{P}$:

$$
\begin{aligned}
P_{t|t-1} &= P_{t-1|t-1} + Q, \\
K_t &= P_{t|t-1} X_t \left(X_t^\top P_{t|t-1} X_t + R_t \right)^{-1}, \\
\theta_{t|t} &= \theta_{t-1|t-1} + K_t \left(y_t - X_t^\top \theta_{t-1|t-1} \right), \\
P_{t|t} &= P_{t|t-1} - K_t X_t^\top P_{t|t-1}.
\end{aligned}
$$

The backward pass draws $\theta_T \sim N(\theta_{T|T}, P_{T|T})$ and then, for $t = T-1, \dots, 0$,

$$
\theta_t \mid \theta_{t+1} \sim N\left(\theta_{t|t} + J_t(\theta_{t+1} - \theta_{t|t}),\; P_{t|t} - J_t P_{t+1|t} J_t^\top \right), \qquad J_t = P_{t|t} P_{t+1|t}^{-1}.
$$

Setting the random draws to zero turns the backward pass into the Rauch–Tung–Striebel smoother, which delivers $\mathbb{E}(\theta_t \mid T)$ given Q and R^T.

The gains J_t and the Cholesky factors of the conditional covariances depend only on Q and R^T, so `backward_factors` computes them once per sweep and each redraw reuses them.

Drawing θ_0 supplies all T increments to the update of Q.

Because the stability test runs after every backward pass, the sampler uses `stable_path`, which applies the Schur–Cohn criterion to the characteristic polynomials of all the companion matrices at once instead of computing their eigenvalues: the roots lie inside the unit circle exactly when a recursion of "reflection coefficients" stays inside $(-1, 1)$, and the test gives the same answer as `is_stable` several times faster.

When $Q = 0$, $P_{t+1|t} = P_{t|t}$, so $J_t = I$ and the conditional covariance is zero: the backward pass returns a constant path, which is how the sampler nests the constant-coefficient model.

The redraw step is rejection sampling.

Given Ω, the backward passes are independent draws from the Gaussian $N(\hat{z}, C)$ of the whole path, and the first draw that lands in the stable set $\mathcal{S}$ has density

$$\frac{N(z; \hat{z}, C)\, \mathbb{1}_{\mathcal{S}}(z)}{\mathbb{P}\{z \in \mathcal{S} \mid \Omega\}},$$

which is the restricted full conditional; the normalizing probability is $p(\Omega)$ in the identity of *Simulating the posterior*, and it never has to be computed.

107.10.2 The volatility block

Given its neighbors and the orthogonalized residual ε^*_{it}, the i-th component of $B\varepsilon_t$, the conditional density of h_{it} is proportional to

$$h_{it}^{-1/2} \exp\left(-\frac{(\varepsilon^*_{it})^2}{2h_{it}}\right) \times \frac{1}{h_{it}} \exp\left(-\frac{\left(\log h_{it} - \frac{1}{2}(\log h_{i,t-1} + \log h_{i,t+1})\right)^2}{\sigma_i^2}\right).$$

The proposal draws $\log h_{it}$ from the normal factor, with mean $\frac{1}{2}(\log h_{i,t-1} + \log h_{i,t+1})$ and variance $\sigma_i^2/2$, so the Metropolis acceptance ratio involves only the likelihood factor.

The last date has a single neighbor, and $\log h_{i0}$ has a conjugate normal conditional.

107.10.3 Monte Carlo error

Successive sweeps are correlated, so the chain's draws carry less information than their number suggests.

The trace of $\mathrm{tr}(Q)$ over all sweeps, including the burn-in, shows how slowly the chain moves in the direction of Q.

The reason is the dependence between Q and the coefficient path: given the path, its T increments pin Q down tightly, and given Q, the roughness of the next path is pinned down too, so the two can move only gradually together.

```
fig, ax = plt.subplots(figsize=(9, 3.5))
ax.plot(posterior['trace_q'], lw=0.5)
ax.axvline(burn, color='0.45', ls='--', lw=1, label='end of burn-in')
ax.set_xlabel('sweep')
ax.set_ylabel(r'$\mathrm{tr}(Q)$')
ax.legend()
plt.show()
```

The integrated autocorrelation time τ measures how many sweeps are worth one independent draw; we estimate it with Sokal's adaptive window.

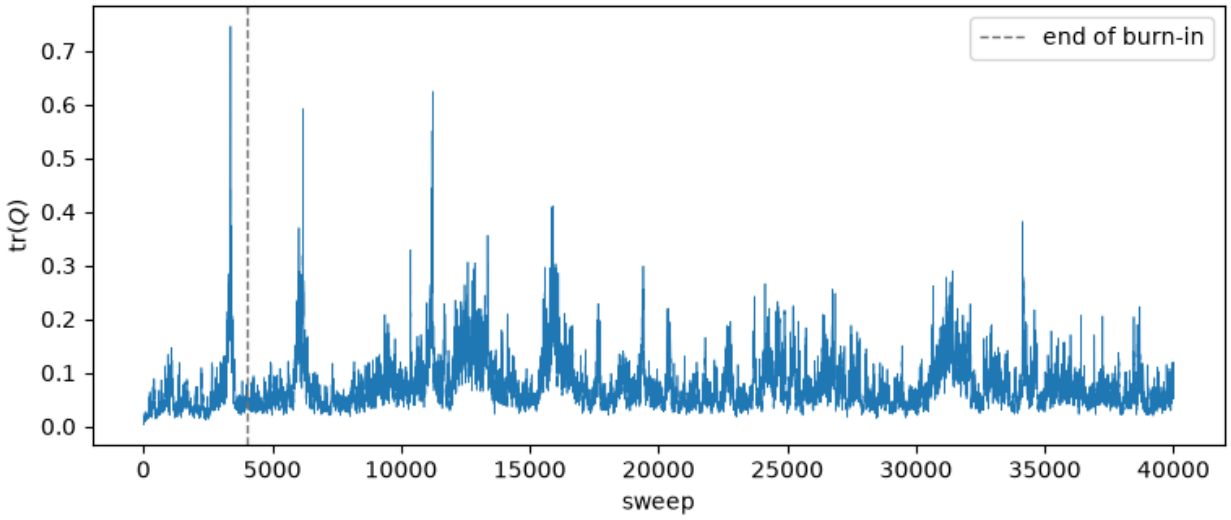


Fig. 107.12: tr(Q) at every sweep

```
def integrated_autocorrelation_time(x, c=5):
    """Sokal's adaptive-window estimate of the integrated autocorrelation time."""
    x = np.asarray(x) - np.mean(x)
    n = len(x)
    spectrum = np.fft.rfft(x, 2 * n)
    acf = np.fft.irfft(spectrum * np.conj(spectrum))[:n]
    τ = 2 * np.cumsum(acf / acf[0]) - 1
    window = np.flatnonzero(np.arange(n) >= c * τ)
    return τ[window[0]] if window.size else τ[-1]   # chain too short for the window


τ_trace_q = integrated_autocorrelation_time(trace_q)
pd.Series({
    'sweeps after burn-in': len(trace_q),
    'autocorrelation time of tr(Q)': τ_trace_q,
    'effective sample size of tr(Q)': len(trace_q) / τ_trace_q,
}, name='value').to_frame().round(1)
```

```
                                value
sweeps after burn-in          36000.0
autocorrelation time of tr(Q)   416.0
effective sample size of tr(Q)   86.5
```

For the other headline numbers we use batch means: split the retained draws into ten consecutive batches, recompute each statistic on every batch, and estimate the Monte Carlo standard error of the full-run value by the standard deviation across batches divided by $\sqrt{10}$.

```
def headline_statistics(s):
    """Statistics quoted in the lecture, from the retained draws in slice s."""
    θ_s = posterior['θ'][s].mean(axis=0)
    r_s = mean_innovation_covariance(posterior['h'][s], posterior['β'][s])
    core, natural = local_means(θ_s)
    persistence = persistence_path(θ_s, r_s)
    eigenvalues = np.linalg.eigvalsh(posterior['Q'][s].mean(axis=0))[::-1]
```

(continues on next page)

(continued from previous page)

```
    return np.r_[
        np.trace(posterior['Q'][s], axis1=1, axis2=2).mean(),
        eigenvalues[0] / eigenvalues.sum(),
        100 * core[annual].max(),
        np.corrcoef(core, natural)[0, 1],
        np.corrcoef(core, persistence)[0, 1],
        persistence[annual].max(),
        [event[s].mean() for event in activism_events],
    ]


n_batches = 10
batch = n_draws // n_batches
batch_values = np.array([headline_statistics(slice(b * batch, (b + 1) * batch))
                         for b in range(n_batches)])
pd.DataFrame({
    'estimate': headline_statistics(slice(None)),
    'Monte Carlo s.e.': batch_values.std(axis=0, ddof=1) / np.sqrt(n_batches),
    'Cogley and Sargent': [0.047, 0.485, 8.0, 0.748, 0.92, np.nan,
                           0.208, 0.919, 0.941, 0.923, 0.943],
}, index=['posterior mean of tr(Q)', 'share of first principal component',
          'peak core inflation (%)', 'corr(core, natural rate)',
          'corr(core, persistence)', 'peak g(0)', 'P(A_1975 > 1)',
          'P(A_1985 > 1)', 'P(A_1995 > 1)', 'P(A_1985 > A_1975)',
          'P(A_1995 > A_1975)']).round(3)
```

	estimate	Monte Carlo s.e.	\
posterior mean of tr(Q)	0.077	0.004	
share of first principal component	0.448	0.018	
peak core inflation (%)	11.818	0.837	
corr(core, natural rate)	0.837	0.027	
corr(core, persistence)	0.965	0.007	
peak g(0)	6.886	0.302	
P(A_1975 > 1)	0.285	0.034	
P(A_1985 > 1)	0.736	0.051	
P(A_1995 > 1)	0.693	0.054	
P(A_1985 > A_1975)	0.744	0.050	
P(A_1995 > A_1975)	0.699	0.053	

	Cogley and Sargent
posterior mean of tr(Q)	0.047
share of first principal component	0.485
peak core inflation (%)	8.000
corr(core, natural rate)	0.748
corr(core, persistence)	0.920
peak g(0)	NaN
P(A_1975 > 1)	0.208
P(A_1985 > 1)	0.919
P(A_1995 > 1)	0.941
P(A_1985 > A_1975)	0.923
P(A_1995 > A_1975)	0.943

These standard errors are rough, because each batch spans only a few autocorrelation times of $\mathrm{tr}(Q)$.

Runs with other seeds differ from this one by more than the standard errors suggest for $\mathrm{tr}(Q)$, the peak of core inflation, and the activism probabilities.

The statements in the text are those that held in all of our runs: the structure of drift, the timing of the rise and fall of

core inflation and persistence, their high correlation, and the shift from passive to active policy, but not exact magnitudes.

107.11 Exercises

Exercise 107.11.1

Consider a stationary $AR(1)$, $\pi_t = \rho \pi_{t-1} + w_t$ with $|\rho| < 1$.

1. Use (107.8), with the companion matrix replaced by the scalar ρ, to show that $g(0) = (1+\rho)/[2\pi(1-\rho)]$, and invert the formula to express ρ as a function of $g(0)$.
2. Apply the inversion draw by draw to the persistence draws in `feature_draws['persistence']`, and plot the posterior median and pointwise 90 percent band of the resulting "equivalent $AR(1)$" coefficient ρ_t.
3. Which ρ corresponds to white noise, and which to $g(0) = 1$? How did the equivalent ρ_t change between the early 1960s, the 1970s, and the 1990s?

Solution

For the $AR(1)$, $f(0) = \sigma_w^2/[2\pi(1-\rho)^2]$ and the variance is $\sigma_w^2/(1-\rho^2)$, so

$$g(0) = \frac{1-\rho^2}{2\pi(1-\rho)^2} = \frac{1+\rho}{2\pi(1-\rho)}, \qquad \rho = \frac{2\pi g(0) - 1}{2\pi g(0) + 1}.$$

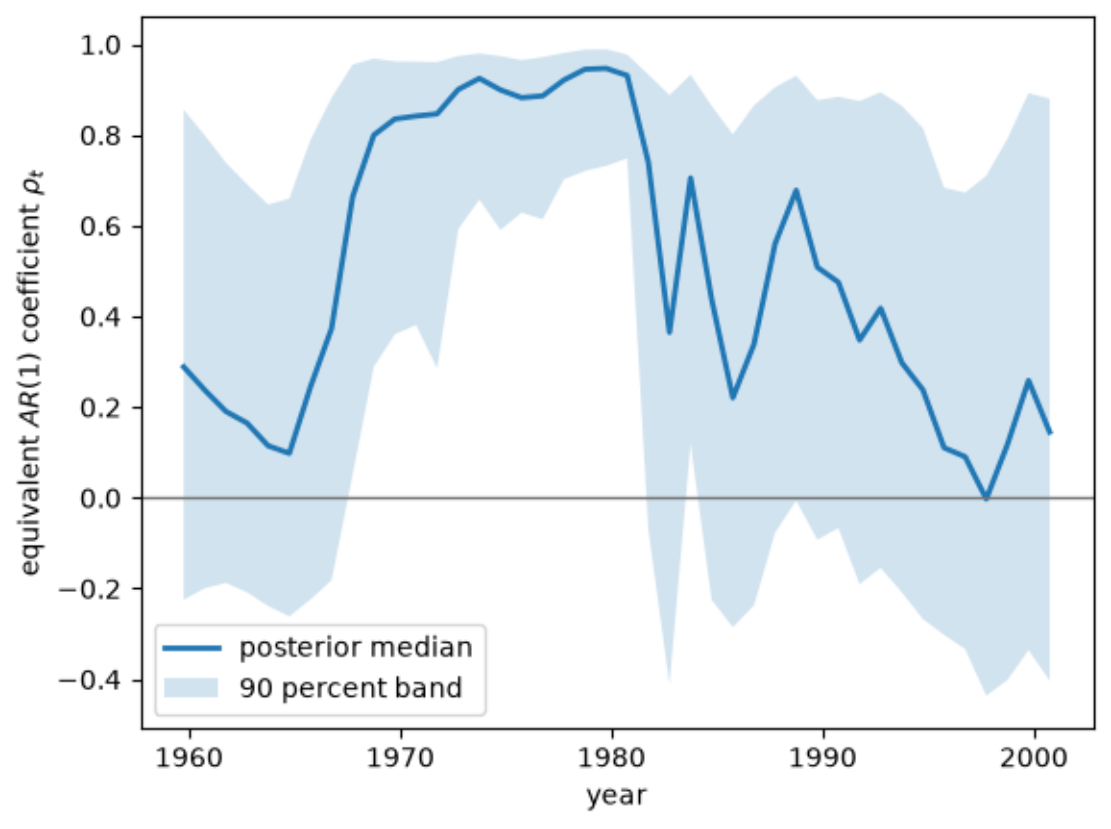
White noise, $g(0) = 1/(2\pi)$, gives $\rho = 0$, and $g(0) = 1$ gives $\rho = (2\pi - 1)/(2\pi + 1) \approx 0.73$.

The inversion is increasing in $g(0)$, so it maps posterior quantiles of $g_{\pi\pi}(0, t)$ into posterior quantiles of ρ_t.

```
def equivalent_ar1(g):
    """AR(1) coefficient with the same normalized power at frequency zero."""
    return (2 * np.pi * g - 1) / (2 * np.pi * g + 1)


ρ_draws = equivalent_ar1(feature_draws['persistence'])
lower, median, upper = np.quantile(ρ_draws, (0.05, 0.5, 0.95), axis=1)

fig, ax = plt.subplots()
ax.plot(dates[annual], median, lw=2, label='posterior median')
ax.fill_between(dates[annual], lower, upper, alpha=0.2, label='90 percent band')
ax.axhline(0, color='0.45', lw=1)
ax.set_xlabel('year')
ax.set_ylabel(r'equivalent $AR(1)$ coefficient $\rho_t$')
ax.legend()
plt.show()
```



The table below gives the numbers behind the figure.

```
pd.Series({
    period: np.median(ρ_draws[(years[annual] >= first) & (years[annual] <= last)])
    for period, (first, last) in {'1960--1964': (1960, 1964),
                                  '1970--1979': (1970, 1979),
                                  '1990--2000': (1990, 2000)}.items()
}, name='median equivalent ρ').to_frame().round(2)
```

```
           median equivalent ρ
1960--1964                0.16
1970--1979                0.91
1990--2000                0.22
```

The equivalent coefficient is small in the early 1960s, about 0.9 in the 1970s, and below one half in the 1990s, although the band in the 1970s is wide.

Exercise 107.11.2

A formal contrast between drifting and constant coefficients requires refitting the model with $Q = 0$, the special case in which the VAR coefficients are constant and only the stochastic volatilities H_t move.

The sampler supports this: passing `fixed_q=np.zeros((n_coefficients, n_coefficients))` to `run_sampler` holds Q at zero.

Fit the constant-coefficient model to the same sample and plot its $g_{\pi\pi}(0, t)$ against the drifting-coefficient path of Fig. 107.7.

Can drifting volatility alone account for the rise and fall of measured persistence?

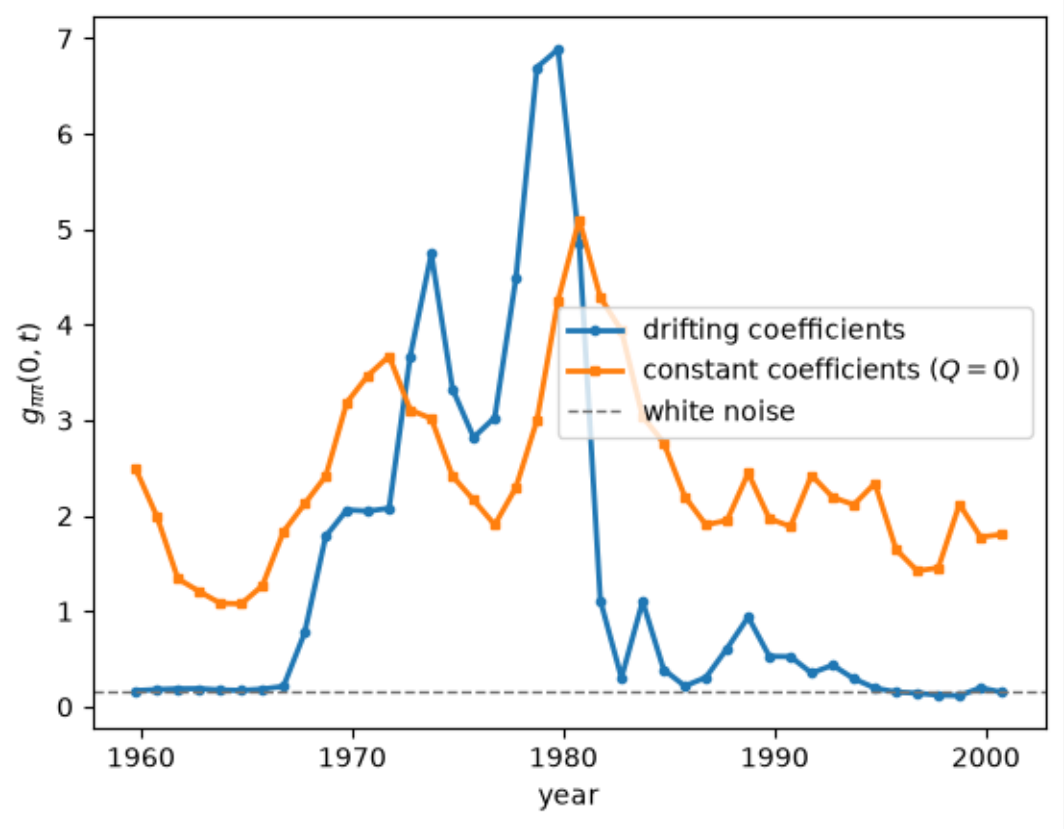
Solution

With $Q = 0$ the coefficient path is constant (*Appendix: the sampler in detail*), so $A_{t|T}$ no longer moves and $g_{\pi\pi}(0, t)$ can vary only through the composition of R_t.

The constant-coefficient chain mixes much faster, so a shorter run suffices.

```
constant_posterior = run_sampler(
    data['y'], data['x'], prior,
    n_sweeps=4_000, burn=1_000, thin=10, seed=42,
    fixed_q=np.zeros((n_coefficients, n_coefficients)),
)
constant_θ = constant_posterior['θ'].mean(axis=0)
constant_r = mean_innovation_covariance(constant_posterior['h'],
                                        constant_posterior['β'])
constant_persistence = persistence_path(constant_θ, constant_r)

fig, ax = plt.subplots()
ax.plot(dates[annual], inflation_persistence[annual], 'o-', lw=2,
        markersize=3, label='drifting coefficients')
ax.plot(dates[annual], constant_persistence[annual], 's-', lw=2,
        markersize=3, label=r'constant coefficients ($Q=0$)')
ax.axhline(white_noise, color='0.45', ls='--', lw=1, label='white noise')
ax.set_xlabel('year')
ax.set_ylabel(r'$g_{\pi\pi}(0,t)$')
ax.legend(loc='center right')
plt.show()
```



The table below gives the numbers behind the figure.

```
after_1981 = years[annual] >= 1982
pd.DataFrame({
    label: {
        '1960--1964 mean': period_mean(path, 1960, 1964),
        '1970--1979 mean': period_mean(path, 1970, 1979),
        'peak': path[annual].max(),
        'year of the peak': years[annual][np.argmax(path[annual])],
        '1990--2000 mean': period_mean(path, 1990, 2000),
        'share of 1982--2000 years below one': np.mean(path[annual][after_1981] <␣
↪1),
    }
    for label, path in (('drifting coefficients', inflation_persistence),
                        ('constant coefficients', constant_persistence))
}).round(2)
```

```
                                    drifting coefficients  \
1960--1964 mean                                      0.18
1970--1979 mean                                      3.84
peak                                                 6.89
year of the peak                                  1979.00
1990--2000 mean                                      0.26
share of 1982--2000 years below one                  0.95

                                    constant coefficients
1960--1964 mean                                      1.49
1970--1979 mean                                      2.88
peak                                                 5.10
year of the peak                                  1980.00
1990--2000 mean                                      1.92
share of 1982--2000 years below one                  0.00
```

With constant coefficients, measured persistence still moves, because the composition of R_t changes even after its scale is normalized out, and the constant-coefficient path peaks around 1980, when inflation innovations are large relative to the others.

But it misses both ends of the story.

It puts persistence in the early 1960s far above the white-noise value, and after 1980 it stays above one, with 1990s persistence not far below its 1970s level, while the drifting-coefficient path is below one in most or all of those years.

Drifting volatility alone accounts for some of the movement in measured persistence but neither its low level in the early 1960s nor its post-1980 collapse, which is why both channels are needed to answer the bad-policy-or-bad-luck question.

Exercise 107.11.3

For a regression $z_t = x_t^\top b + u_t$ estimated by least squares with residuals $\hat{u}_t$, the statistic of Nyblom [1989] and Hansen [1992] for constancy of all coefficients is

$$L = \frac{1}{T^2} \sum_{t=1}^{T} S_t^\top V^{-1} S_t, \qquad S_t = \sum_{s \le t} x_s \hat{u}_s, \qquad V = \frac{1}{T} \sum_{s=1}^{T} x_s x_s^\top \hat{u}_s^2.$$

Apply it to the $AR(1)$ with intercept $y_t = 0.5 + \rho_t y_{t-1} + w_t$, $w_t \sim N(0,1)$, with $T = 168$ and $\rho_0 = 0.6$, using 2,000 replications per case.

1. Simulate the 5 percent critical value when ρ_t is constant, and compare it with Hansen's asymptotic value of 0.749 for two coefficients.
2. Compute the rejection frequency when ρ_t is a random walk, kept inside $[-0.97, 0.97]$, with innovation variance $q \in \{10^{-4}, 4 \times 10^{-4}, 10^{-3}, 4 \times 10^{-3}\}$.
3. Compute the rejection frequency when ρ is constant but the standard deviation of w_t triples halfway through the sample.

Relate your results to the power that Cogley and Sargent report in *Why classical tests missed the drift*.

Solution

We simulate all replications at once.

```
def simulate_ar1_panel(n, T, q, rng, volatility_break=False):
    """n replications of y_t = 0.5 + ρ_t y_{t-1} + s_t w_t."""
    y = np.zeros((n, T + 1))
    ρ = np.full(n, 0.6)
    for t in range(1, T + 1):
        ρ = np.clip(ρ + np.sqrt(q) * rng.standard_normal(n), -0.97, 0.97)
        scale = 3.0 if volatility_break and t > T // 2 else 1.0
        y[:, t] = 0.5 + ρ * y[:, t - 1] + scale * rng.standard_normal(n)
    return y


def nyblom_hansen(y):
    """Nyblom-Hansen L statistic for the intercept and slope, by replication."""
    X = np.stack((np.ones_like(y[:, :-1]), y[:, :-1]), axis=2)
    z = y[:, 1:]
    b = np.linalg.solve(np.einsum('nti,ntj->nij', X, X),
```

```
                        np.einsum('nti,nt->ni', X, z)[:, :, None])[:, :, 0]
    scores = X * (z - np.einsum('nti,ni->nt', X, b))[:, :, None]
    S = np.cumsum(scores, axis=1)
    V = np.einsum('nti,ntj->nij', scores, scores) / z.shape[1]
    return np.einsum('nti,nij,ntj->n', S, np.linalg.inv(V), S) / z.shape[1]**2


rng = np.random.default_rng(1234)
n, T = 2_000, 168
critical_value = np.quantile(nyblom_hansen(simulate_ar1_panel(n, T, 0.0, rng)), 0.
↪95)

rows = {}
for q in (1e-4, 4e-4, 1e-3, 4e-3):
    L = nyblom_hansen(simulate_ar1_panel(n, T, q, rng))
    rows[f'random-walk ρ, q = {q:g}'] = {
        'sd of the change in ρ over the sample': np.sqrt(T * q),
        'rejection frequency': np.mean(L > critical_value),
    }
L = nyblom_hansen(simulate_ar1_panel(n, T, 0.0, rng, volatility_break=True))
rows['constant ρ, shock sd triples'] = {
    'sd of the change in ρ over the sample': 0.0,
    'rejection frequency': np.mean(L > critical_value),
}
print(f'simulated 5 percent critical value: {critical_value:.3f}')
pd.DataFrame(rows).T.round(3)
```

```
simulated 5 percent critical value: 0.731
```

```
                              sd of the change in ρ over the sample  \
random-walk ρ, q = 0.0001                                      0.130
random-walk ρ, q = 0.0004                                      0.259
random-walk ρ, q = 0.001                                       0.410
random-walk ρ, q = 0.004                                       0.820
constant ρ, shock sd triples                                   0.000

                              rejection frequency
random-walk ρ, q = 0.0001                   0.100
random-walk ρ, q = 0.0004                   0.210
random-walk ρ, q = 0.001                    0.292
random-walk ρ, q = 0.004                    0.428
constant ρ, shock sd triples                0.032
```

The simulated critical value is close to Hansen's asymptotic value.

Power rises with q but stays below one half even when the standard deviation of the change in ρ over the sample is about 0.8, and drift that moves ρ by a few tenths over the sample is detected only a fifth to a third of the time, the same order as the 0.08 to 0.23 that Cogley and Sargent find for their posterior.

A tripling of the shock standard deviation with constant ρ is rejected no more often than the test's 5 percent size, so the test does not mistake volatility for drift; it is simply weak against drift.

Exercise 107.11.4

Do escape-route models predict the positive correlation between mean inflation and persistence that we found in *Inflation persistence*?

Simulate the static model of *Escaping Nash Inflation*, with $\theta = 1$, $u = 5$, and $\sigma_1 = \sigma_2 = 0.3$, under constant-gain least squares with gain $\varepsilon = 0.02$ for 40,000 periods, starting at the self-confirming equilibrium.

In rolling windows of 40 periods, fit an $AR(1)$ to inflation, compute its $g(0) = (1 + \hat{\rho})/[2\pi(1 - \hat{\rho})]$, and record mean inflation.

1. Compare the median $g(0)$ in windows that sit at Nash (mean inflation above 4 in both halves), near Ramsey (below 1 in both halves), and that span a fall (above 3.5 in the first half and below 1.5 in the second).
2. Compute the correlation between rolling mean inflation and rolling $g(0)$, and compare it with the correlation between core inflation and persistence in the data.

Solution

The simulation repeats `simulate_path` from *Escaping Nash Inflation*, with beliefs $\gamma = (\gamma_1, \gamma_{-1})$ about the slope and intercept of the perceived Phillips curve $U = \gamma_1 \pi + \gamma_{-1}$.

```
def simulate_escapes(ε, n_periods, θ=1.0, u=5.0, σ1=0.3, σ2=0.3, seed=0):
    """Inflation in the static CWS model under constant-gain least squares."""
    rng = np.random.default_rng(seed)
    γ = np.array([-θ, u * (1 + θ**2)])                    # SCE beliefs
    x = -γ[1] * γ[0] / (1 + γ[0]**2)
    R = np.array([[x**2 + σ2**2, x], [x, 1.0]])
    π = np.empty(n_periods)
    for n in range(n_periods):
        x = -γ[1] * γ[0] / (1 + γ[0]**2)                   # Phelps choice
        w1, w2 = rng.standard_normal(2)
        π[n] = x + σ2 * w2
        U = u - θ * σ2 * w2 + σ1 * w1
        Φ = np.array([π[n], 1.0])
        R = R + ε * (np.outer(Φ, Φ) - R)
        γ = γ + ε * np.linalg.solve(R, Φ * (U - Φ @ γ))
    return π


π_sim = simulate_escapes(0.02, 40_000, seed=2)
windows = np.lib.stride_tricks.sliding_window_view(π_sim, 40)[::2]
centered = windows - windows.mean(axis=1, keepdims=True)
ρ_hat = np.clip(np.sum(centered[:, 1:] * centered[:, :-1], axis=1)
                / np.sum(centered[:, :-1]**2, axis=1), -0.99, 0.99)
g0 = (1 + ρ_hat) / (2 * np.pi * (1 - ρ_hat))
first, second = windows[:, :20].mean(axis=1), windows[:, 20:].mean(axis=1)

groups = {
    'at Nash': (first > 4) & (second > 4),
    'near Ramsey': (first < 1) & (second < 1),
    'spanning a fall': (first > 3.5) & (second < 1.5),
}
display(pd.DataFrame({
    'windows': {label: mask.sum() for label, mask in groups.items()},
    'median g(0)': {label: np.median(g0[mask]) for label, mask in groups.items()},
}).round(2))
print(f'correlation of rolling mean inflation and rolling g(0): '
      f'{np.corrcoef(windows.mean(axis=1), g0)[0, 1]:.3f}')
print(f'correlation of core inflation and persistence in the data: '
      f'{core_persistence_correlation:.3f}')
```

```
                 windows  median g(0)
at Nash             1849         0.16
near Ramsey         3125         0.17
spanning a fall      135         9.73
```

```
correlation of rolling mean inflation and rolling g(0): 0.033
correlation of core inflation and persistence in the data: 0.965
```

Measured persistence is close to the white-noise value $1/(2\pi)$ both at Nash and near Ramsey, and it is high only in windows that span a fall, where the shifting mean inside the window looks like persistence.

Over the whole simulation, rolling mean inflation and rolling persistence are nearly uncorrelated, far from the strong positive correlation in the data.

The escape-route model ties persistence to *changes* in inflation, not to its *level*: high inflation at the self-confirming equilibrium comes with white-noise inflation, and persistence peaks while inflation falls, the inverse relation that Cogley and Sargent [2005] describe.

CHAPTER

EIGHT

DRIFTS AND VOLATILITIES AFTER 2000

Contents

- *Drifts and Volatilities after 2000*
  - *Overview*
  - *The model and the sampler*
  - *The extended sample*
  - *Fitting the extended sample*
  - *Did coefficient drift continue?*
  - *Volatility after the Great Moderation*
  - *Core inflation, the natural rate, and persistence*
  - *What the additional observations change*
  - *Exercises*

108.1 Overview

Drifts and Volatilities fit the Bayesian vector autoregression of Cogley and Sargent [2005], with drifting coefficients and stochastic volatilities, to quarterly U.S. data on interest, unemployment, and inflation for 1959–2000.

It found that both the volatilities and the coefficients drifted: shocks were large around 1980, and core inflation and inflation persistence rose through the 1970s and fell together after 1980.

Cogley and Sargent's sample ends in 2000, so it misses the financial crisis, the long stretch at the zero lower bound, the pandemic, and the 2021–2022 inflation surge.

This lecture re-estimates the same model, with the same priors and the same 1948–1958 training sample, on data extended through the latest complete quarter, and asks two questions.

1. Do the new observations change what the model says about 1959–2000?
2. Did the recent inflation surge bring back 1970s-style persistence?

The code repeats that of *Drifts and Volatilities*, where each piece is explained; it is collapsed below so that this lecture runs on its own.

We begin with imports and the addresses of the data.

```
import time

import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from scipy import linalg
from scipy.special import expit
from threadpoolctl import threadpool_limits

# The matrices in this lecture are small, so extra BLAS threads only add overhead.
threadpool_limits(1)

archive_url = 'https://github.com/QuantEcon/data-lectures/raw/main/lectures/NEWQDATA.
↪csv'
fred_url = ('https://fred.stlouisfed.org/graph/fredgraph.csv?'
            'id=CPIAUCSL%2CUNRATE%2CTB3MS')
```

108.2 The model and the sampler

The model is the one in *A VAR with drifting coefficients and stochastic volatility*: a VAR(2) in the nominal interest rate, the logit of unemployment, and inflation, $y_t = X_t^\top \theta_t + \varepsilon_t$, whose coefficients follow a random walk $\theta_t = \theta_{t-1} + v_t$ with $v_t \sim N(0, Q)$, restricted to stable paths, and whose innovation covariance $R_t = B^{-1} H_t B^{-1\prime}$ has random-walk log volatilities.

The first cell builds the model, the data, and the priors.

The second cell is the sampler of *Simulating the posterior*: a Carter–Kohn draw of the coefficient path, redrawn until it is stable, followed by the updates of Q, the volatility scales σ, the covariance parameters β, and the volatility paths.

The third cell computes the summaries of *What the data say*: the posterior mean of R_t, core inflation and the natural rate, and the normalized inflation spectrum $g_{\pi\pi}(\omega, t)$.

108.3 The extended sample

We keep Cogley and Sargent's archived data through 2000Q4 and append newer observations from FRED: seasonally adjusted CPI, seasonally adjusted unemployment, and the three-month Treasury-bill rate.

The transformations and within-quarter timing are unchanged: CPI comes from the third month of each quarter, unemployment is a three-month average, and the interest rate comes from the first month.

Keeping the archive for 1948–2000 means that revisions to old data are not mistaken for information in the new quarters; the first appended inflation rate uses revised CPI levels for both 2000Q4 and 2001Q1.

As the BLS notes, the October 2025 unemployment rate was never collected, because of the federal government shutdown, so 2025Q4 lacks one of its three months.

Rather than end the sample at 2025Q3 forever, we fill isolated missing months by linear interpolation between their neighbors, and we print the months filled.

The sample then ends at the last quarter for which all three series are complete, found automatically.

```
archive = pd.read_csv(archive_url)
fred = pd.read_csv(fred_url, parse_dates=['observation_date'],
                   index_col='observation_date')
```

(continues on next page)

(continued from previous page)

```
unrate = fred['UNRATE'].dropna()
unrate = unrate.reindex(pd.date_range(unrate.index[0], unrate.index[-1],
                                      freq='MS'))
filled_months = list(unrate.index[unrate.isna()].strftime('%Y-%m'))
unrate = unrate.interpolate(limit=1, limit_area='inside')


def by_quarter(series, months):
    """Observations from the given months of each quarter, indexed by quarter."""
    picked = series[series.index.month.isin(months)].dropna()
    picked.index = picked.index.to_period('Q')
    return picked


month_counts = unrate.groupby(unrate.index.to_period('Q')).count()
table = pd.concat({
    'interest': by_quarter(fred['TB3MS'], (1, 4, 7, 10)),
    'cpi': by_quarter(fred['CPIAUCSL'], (3, 6, 9, 12)),
    'unemployment': unrate.groupby(unrate.index.to_period('Q')).mean()[
        month_counts == 3],
}, axis=1)
table['y3'] = table['interest'] / (400 - table['interest'])
table['ur'] = table['unemployment'] / 100
table['dp'] = np.log(table['cpi']).diff()
table['date'] = table.index.year + (table.index.quarter - 1) / 4

extension = table.loc['2001Q1':, ['date', 'y3', 'ur', 'dp']]
incomplete = extension.index[extension.isna().any(axis=1)]
if len(incomplete):
    extension = extension.loc[:incomplete[0] - 1]
spliced = pd.concat((archive, extension.reset_index(drop=True)),
                    ignore_index=True)
assert np.allclose(np.diff(spliced['date']), 0.25)
sample_end = extension.index[-1]

print('months filled by interpolation:', filled_months)
print('last quarter of the sample:', sample_end)
```

```
months filled by interpolation: ['2025-10']
last quarter of the sample: 2026Q3
```

The next figure shows the data, with a dashed line at the end of Cogley and Sargent's sample.

```
shown = spliced[spliced['date'] >= 1959]
fig, axes = plt.subplots(3, 1, figsize=(9, 7), sharex=True)
axes[0].plot(shown['date'], 400 * shown['dp'], lw=2)
axes[0].set_ylabel('inflation (annual %)')
axes[1].plot(shown['date'], 100 * shown['ur'], lw=2)
axes[1].set_ylabel('unemployment (%)')
axes[2].plot(shown['date'], 400 * shown['y3'] / (1 + shown['y3']), lw=2)
axes[2].set_ylabel('interest (annual %)')
axes[2].set_xlabel('year')
for ax in axes:
    ax.axvline(2000.75, color='0.6', ls='--', lw=1)
plt.tight_layout()
plt.show()
```

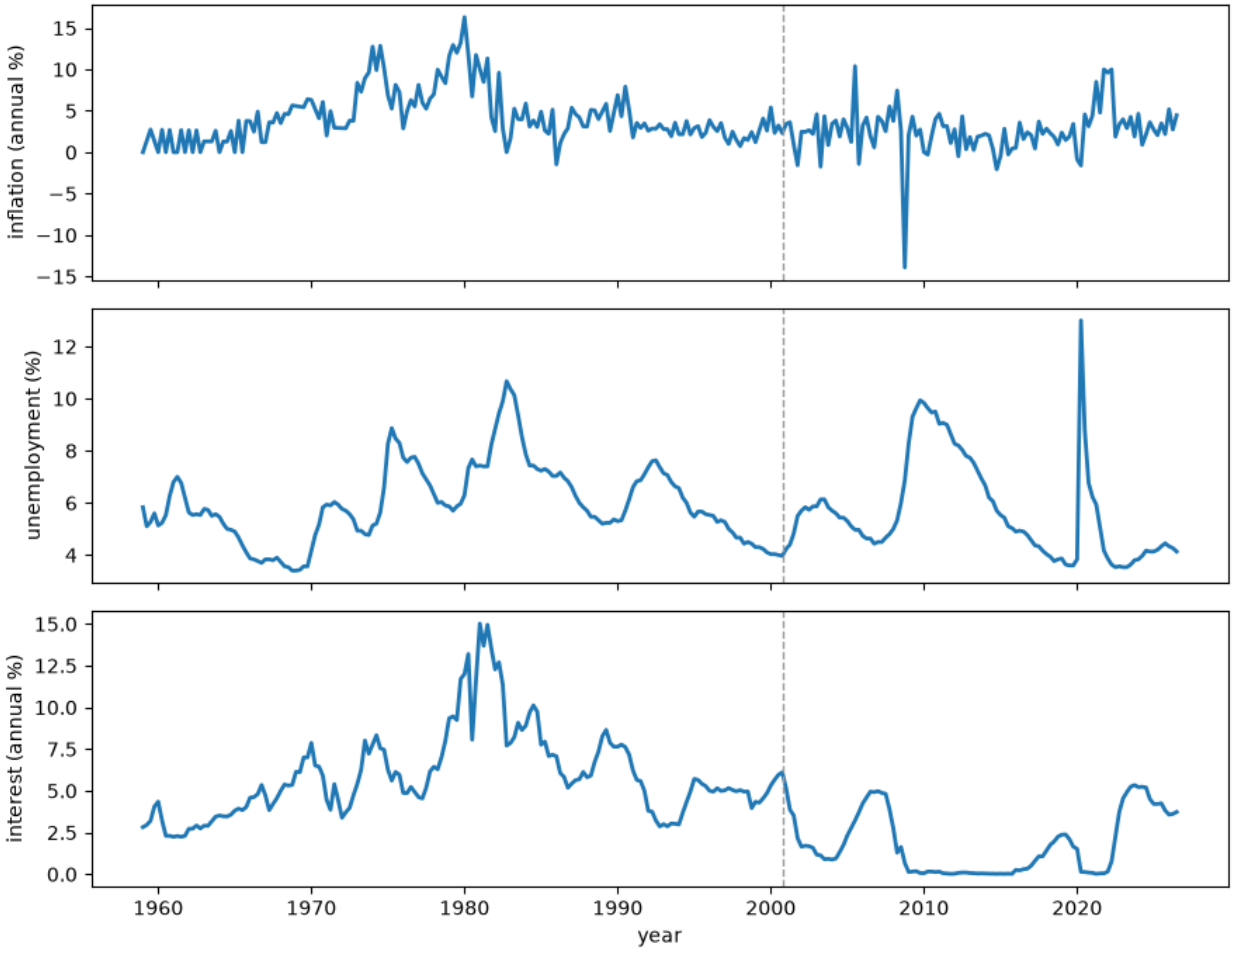


Fig. 108.1: Observed π_t, u_t, and i_t in the extended sample

The extension adds the financial-crisis recession, a pandemic spike in quarterly unemployment, a short inflation surge in 2021–2022, and two long stretches of near-zero interest rates.

Because inflation is annualized from quarterly changes, isolated movements look large, but the recent surge is much shorter than the sustained rise of the 1970s.

108.4 Fitting the extended sample

The sample is now much longer, and stable coefficient paths are much harder to draw over it, so a sweep costs several times more than in *Drifts and Volatilities*.

To bound the cost, a sweep tries at most 50 backward passes before keeping the current path, which still leaves the posterior invariant (the table reports how often that happens), and we run 10,000 sweeps, discard the first 1,000, and keep every 5th of the rest.

```
data = prepare_data(spliced)
prior = calibrate_prior(data)
burn = 1_000
posterior = run_sampler(data['y'], data['x'], prior,
                        n_sweeps=10_000, burn=burn, thin=5, seed=42)

dates = data['dates']
years = np.floor(dates + 1e-8).astype(int)
annual = np.array([np.flatnonzero(years == year)[-1]
                   for year in np.unique(years)])
trace_q = posterior['trace_q'][burn:]
τ = integrated_autocorrelation_time(trace_q)
passes = posterior['passes'][burn:]

pd.Series({
    'run time (minutes)': posterior['minutes'],
    'quarters in the estimation sample': len(dates),
    'retained draws': len(posterior['β']),
    'mean backward passes per sweep': passes.mean(),
    'sweeps that kept the current path': np.sum(passes > 50),
    'autocorrelation time of tr(Q)': τ,
    'effective sample size of tr(Q)': len(trace_q) / τ,
}, name='value').to_frame().round(3)
```

```
                                    value
run time (minutes)                 12.291
quarters in the estimation sample  271.000
retained draws                    1800.000
mean backward passes per sweep      18.675
sweeps that kept the current path 1834.000
autocorrelation time of tr(Q)      992.580
effective sample size of tr(Q)       9.067
```

As in *Drifts and Volatilities*, the chain moves slowly in the direction of Q, and with fewer sweeps the effective sample size of $\mathrm{tr}(Q)$ is even smaller, so magnitudes that depend on Q carry large Monte Carlo error; the text relies only on features that held in runs with other seeds.

108.5 Did coefficient drift continue?

The histogram shows $\text{tr}(Q)$ at every sweep after the burn-in.

```
fig, ax = plt.subplots()
ax.hist(trace_q, bins=50, histtype='step', lw=2)
ax.axvline(np.trace(prior['q_center']), color='C1', lw=2,
           label=r'prior $\mathrm{tr}(\bar Q)$')
ax.set_xlabel(r'$\mathrm{tr}(Q)$')
ax.set_ylabel('sweeps')
ax.legend()
plt.show()
```

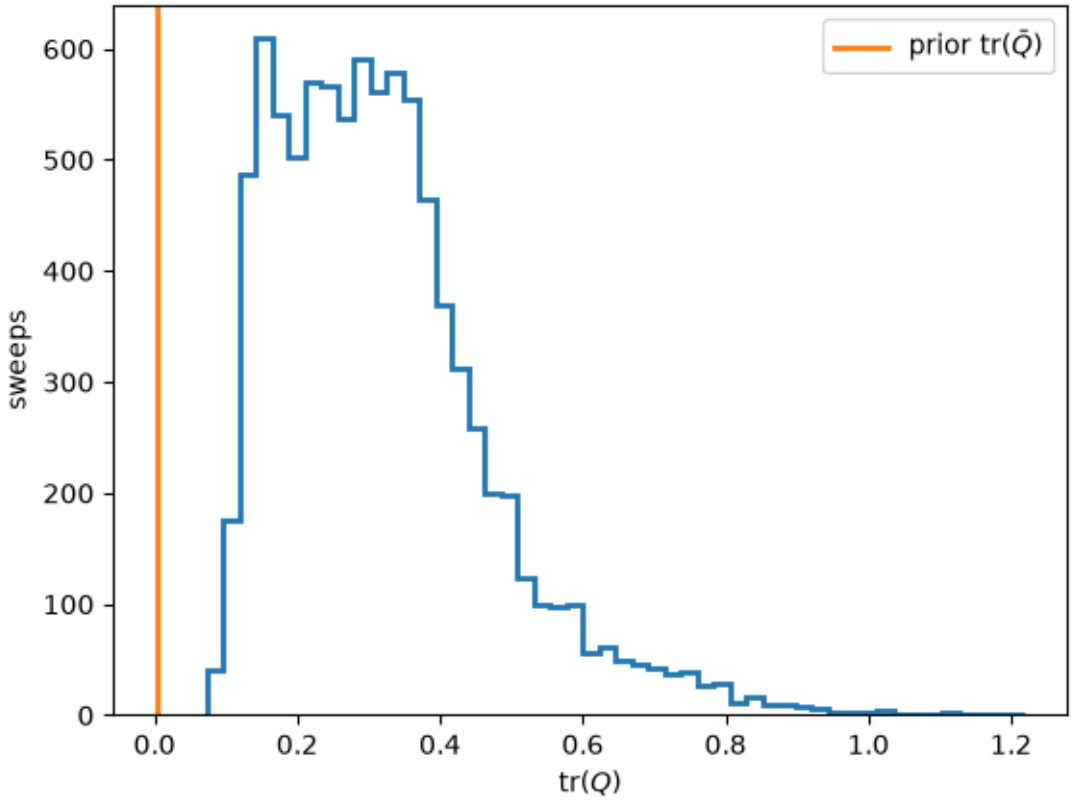


Fig. 108.2: Posterior $\text{tr}(Q)$ in the extended sample and prior $\text{tr}(\bar{Q})$

The table below gives the numbers behind the figure.

```
q_eigenvalues = np.linalg.eigvalsh(posterior['Q'].mean(axis=0))[::-1]
pd.Series({
    'posterior mean of tr(Q)': trace_q.mean(),
    'smallest tr(Q) / tr(Q bar)': trace_q.min() / np.trace(prior['q_center']),
    'share of the first principal component': q_eigenvalues[0] / q_eigenvalues.sum(),
    'share of the first three': q_eigenvalues[:3].sum() / q_eigenvalues.sum(),
}, name='value').to_frame().round(3)
```

```
                                          value
posterior mean of tr(Q)                   0.318
smallest tr(Q) / tr(Q bar)               24.456
share of the first principal component    0.639
share of the first three                  0.976
```

The drift rate is several times larger than in the 1959–2000 fit of *Drifts and Volatilities*, and the drift is concentrated in even fewer directions.

The chain explores this posterior very slowly: in runs with other seeds the posterior mean of $\mathrm{tr}(Q)$ differed by a factor of more than two, although it was always several times the 1959–2000 value.

One likely contributor is the pandemic: in 2020Q2 the logit of unemployment jumped by dozens of innovation standard deviations, far more than a random-walk volatility can absorb within a quarter, so part of the jump is attributed to the coefficients.

Because Q is a single drift covariance for the whole sample, this histogram cannot by itself say whether drift slowed or sped up after 2000; for that we look at what the coefficients imply.

108.6 Volatility after the Great Moderation

The innovation standard deviations and correlations come from the posterior mean $\mathbb{E}(R_t \mid T)$, as in *Drifts and Volatilities*.

```
r_mean = mean_innovation_covariance(posterior['h'], posterior['β'])
sd = 1e4 * np.sqrt(np.diagonal(r_mean, axis1=1, axis2=2))


def correlation(left, right):
    return r_mean[:, left, right] / np.sqrt(r_mean[:, left, left]
                                            * r_mean[:, right, right])


fig, axes = plt.subplots(3, 2, figsize=(9, 8), sharex=True)
for row, (index, label) in enumerate(((0, 'Nominal interest'),
                                      (2, 'Inflation'), (1, 'Unemployment'))):
    axes[row, 0].plot(dates, sd[:, index], lw=2)
    axes[row, 0].set_title(label)
for row, (left, right, label) in enumerate(((0, 1, 'Interest--unemployment'),
                                            (0, 2, 'Interest--inflation'),
                                            (2, 1, 'Inflation--unemployment'))):
    axes[row, 1].plot(dates, correlation(left, right), lw=2)
    axes[row, 1].set_title(label)
for ax in axes.flat:
    ax.axvline(2000.75, color='0.6', ls='--', lw=1)
axes[1, 0].set_ylabel(r'standard deviation $\times 10^4$')
axes[1, 1].set_ylabel('correlation')
axes[-1, 0].set_xlabel('year')
axes[-1, 1].set_xlabel('year')
plt.tight_layout()
plt.show()
```

The table below gives the numbers behind the figure.

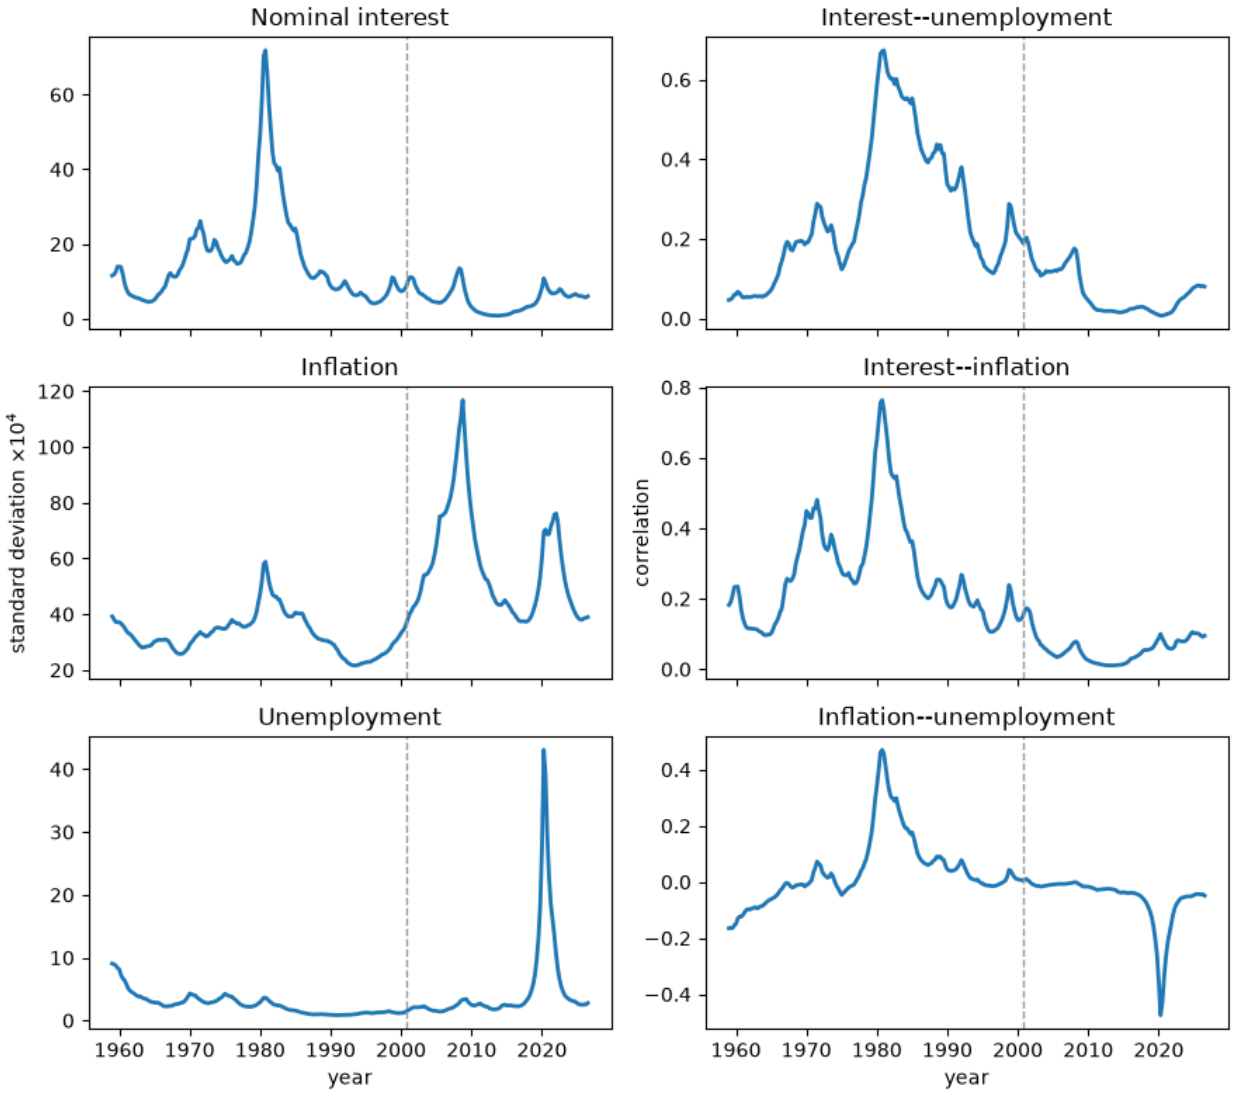


Fig. 108.3: Standard deviations and correlations implied by $E(R_t \mid T)$, extended sample

```
pd.DataFrame({
    'peak date': [dates[np.argmax(sd[:, i])] for i in range(3)],
    'standard deviation at the end': sd[-1],
}, index=['interest', 'unemployment', 'inflation']).round(2)
```

	peak date	standard deviation at the end
interest	1980.75	6.02
unemployment	2020.25	2.77
inflation	2008.75	38.97

The Volcker years still dominate interest-rate innovations, the financial crisis produces the largest inflation innovations, and the pandemic produces by far the largest unemployment innovations, together with a sharp swing in the inflation–unemployment correlation.

Recent episodes thus changed the size and composition of reduced-form shocks dramatically, which is direct evidence for a bad-luck component, although reduced-form innovations are not structural shocks.

The generalized innovation variance $\log|\mathbb{E}(R_t \mid T)|$ summarizes the total one-step uncertainty.

```
logdet_r = np.linalg.slogdet(r_mean)[1]
fig, ax = plt.subplots()
ax.plot(dates, logdet_r, lw=2)
ax.axvline(2000.75, color='0.6', ls='--', lw=1)
ax.set_xlabel('year')
ax.set_ylabel(r'$\log |E(R_t\mid T)|$')
plt.show()
```

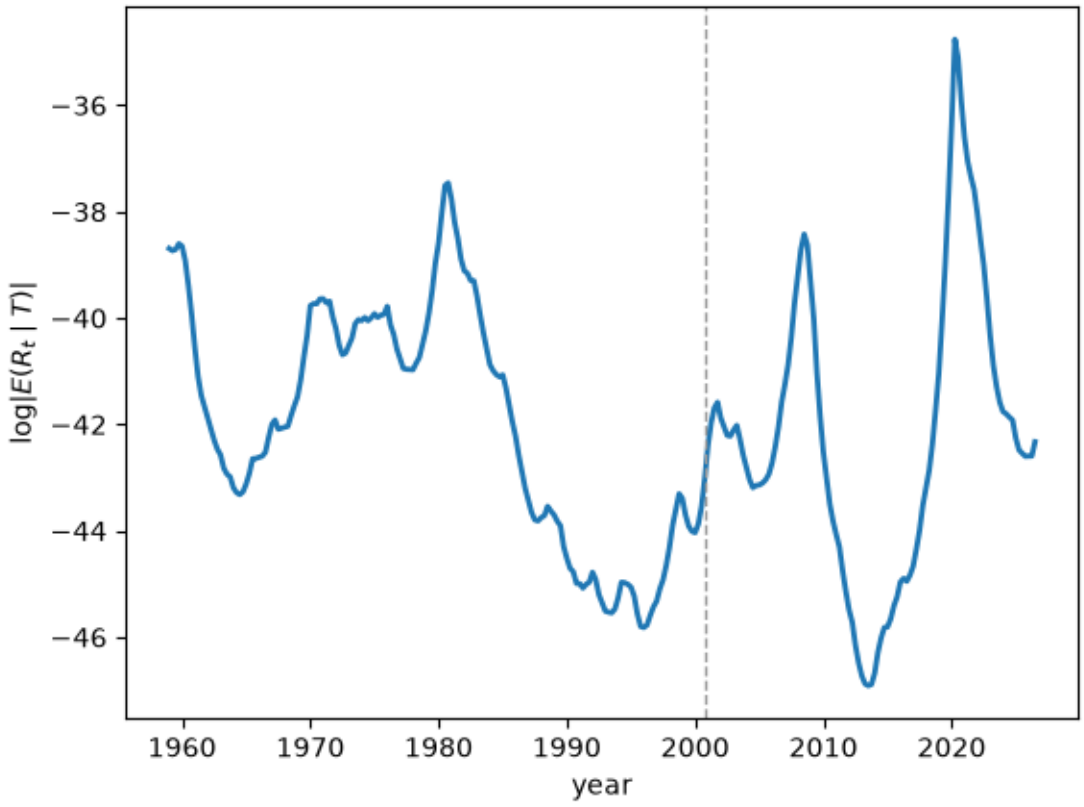


Fig. 108.4: Generalized innovation variance $\log|E(R_t \mid T)|$, extended sample

The table below gives the numbers behind the figure.

```
def period_mean(values, first, last):
    """Average a quarterly series over the years first, ..., last."""
    return values[(years >= first) & (years <= last)].mean()


display_periods = {'1980--1982': (1980, 1982), '1990--2000': (1990, 2000),
                   '2010--2015': (2010, 2015), '2020--2021': (2020, 2021)}


def extreme(values, first, last, statistic):
    return statistic(values[(years >= first) & (years <= last)])


pd.DataFrame({
    'interest innovation sd (basis points)': [
        period_mean(sd[:, 0], *p) for p in display_periods.values()],
    'interest-inflation correlation': [
        period_mean(correlation(0, 2), *p) for p in display_periods.values()],
    'log |E(R_t | T)|, max or min': [
        extreme(logdet_r, 1980, 1982, np.max), extreme(logdet_r, 1990, 2000, np.min),
        extreme(logdet_r, 2010, 2015, np.min), extreme(logdet_r, 2020, 2021, np.max)],
}, index=display_periods.keys()).round(3)
```

```
           interest innovation sd (basis points)  \
1980--1982                                 52.647
1990--2000                                  7.133
2010--2015                                  1.333
2020--2021                                  8.145

           interest-inflation correlation  log |E(R_t | T)|, max or min
1980--1982                           0.645                       -37.459
1990--2000                           0.172                       -45.819
2010--2015                           0.015                       -46.911
2020--2021                           0.075                       -34.759
```

Joint innovation variance rises in the financial crisis, falls to a trough in the early 2010s deeper than its 1990s low, and jumps in 2020 above even its 1981 peak before falling back.

The trough is not a second Great Moderation.

At the zero lower bound the Treasury-bill rate barely moved, so its innovations nearly vanished and lost their correlation with inflation innovations, as the table shows, and the log determinant collapses with them.

For the same reason a short-rate measure of policy activism carries little information after 2008, and we do not compute one here.

108.7 Core inflation, the natural rate, and persistence

We compute core inflation, the natural rate, and $g_{\pi\pi}(0, t)$ from every retained draw, as in *Inflation persistence*.

```
feature_draws = posterior_feature_draws(posterior, annual)
medians = {key: np.median(value, axis=1) for key, value in feature_draws.items()}

fig, axes = plt.subplots(3, 1, figsize=(9, 8), sharex=True)
for ax, (key, ylabel) in zip(axes, (('core', 'core inflation (%)'),
                                    ('natural', 'natural rate (%)'),
                                    ('persistence', r'$g_{\pi\pi}(0,t)$'))):
    lower, upper = np.quantile(feature_draws[key], (0.05, 0.95), axis=1)
    line, = ax.plot(dates[annual], medians[key], lw=2)
    ax.fill_between(dates[annual], lower, upper, color=line.get_color(),
                    alpha=0.2)
    ax.axvline(2000.75, color='0.6', ls='--', lw=1)
    ax.set_ylabel(ylabel)
axes[2].axhline(1 / (2 * np.pi), color='0.45', ls=':', lw=1)
axes[-1].set_xlabel('year')
plt.tight_layout()
plt.show()
```

The dotted line in the bottom panel is the white-noise value $1/(2\pi)$.

The table summarizes the posterior medians by period.

```
annual_years = years[annual]


def median_summary(key, first, last, statistic=np.mean):
    """Apply a statistic to the annual posterior medians over first, ..., last."""
    mask = (annual_years >= first) & (annual_years <= last)
    return statistic(medians[key][mask])


core_peak = np.argmax(medians['core'] * (annual_years <= 2000))
pd.Series({
    'core inflation, 1970s peak (%)': medians['core'][core_peak],
    'year of the core peak': annual_years[core_peak],
    'natural rate, 1970--1979 range (%)': (
        f"{median_summary('natural', 1970, 1979, np.min):.1f}--"
        f"{median_summary('natural', 1970, 1979, np.max):.1f}"),
    'g(0), 1970--1979 mean': median_summary('persistence', 1970, 1979),
    'last year before 2001 with g(0) > 1': annual_years[
        (annual_years <= 2000) & (medians['persistence'] > 1)].max(),
    'core inflation, 2001--2019 range (%)': (
        f"{median_summary('core', 2001, 2019, np.min):.1f}--"
        f"{median_summary('core', 2001, 2019, np.max):.1f}"),
    'core inflation, 2020 onward mean (%)': median_summary('core', 2020, 2100),
    'g(0), 2001--2019 mean': median_summary('persistence', 2001, 2019),
    'g(0), 2020 onward mean': median_summary('persistence', 2020, 2100),
}, name='posterior median').to_frame()
```

```
                                  posterior median
core inflation, 1970s peak (%)             23.596626
```

(continues on next page)

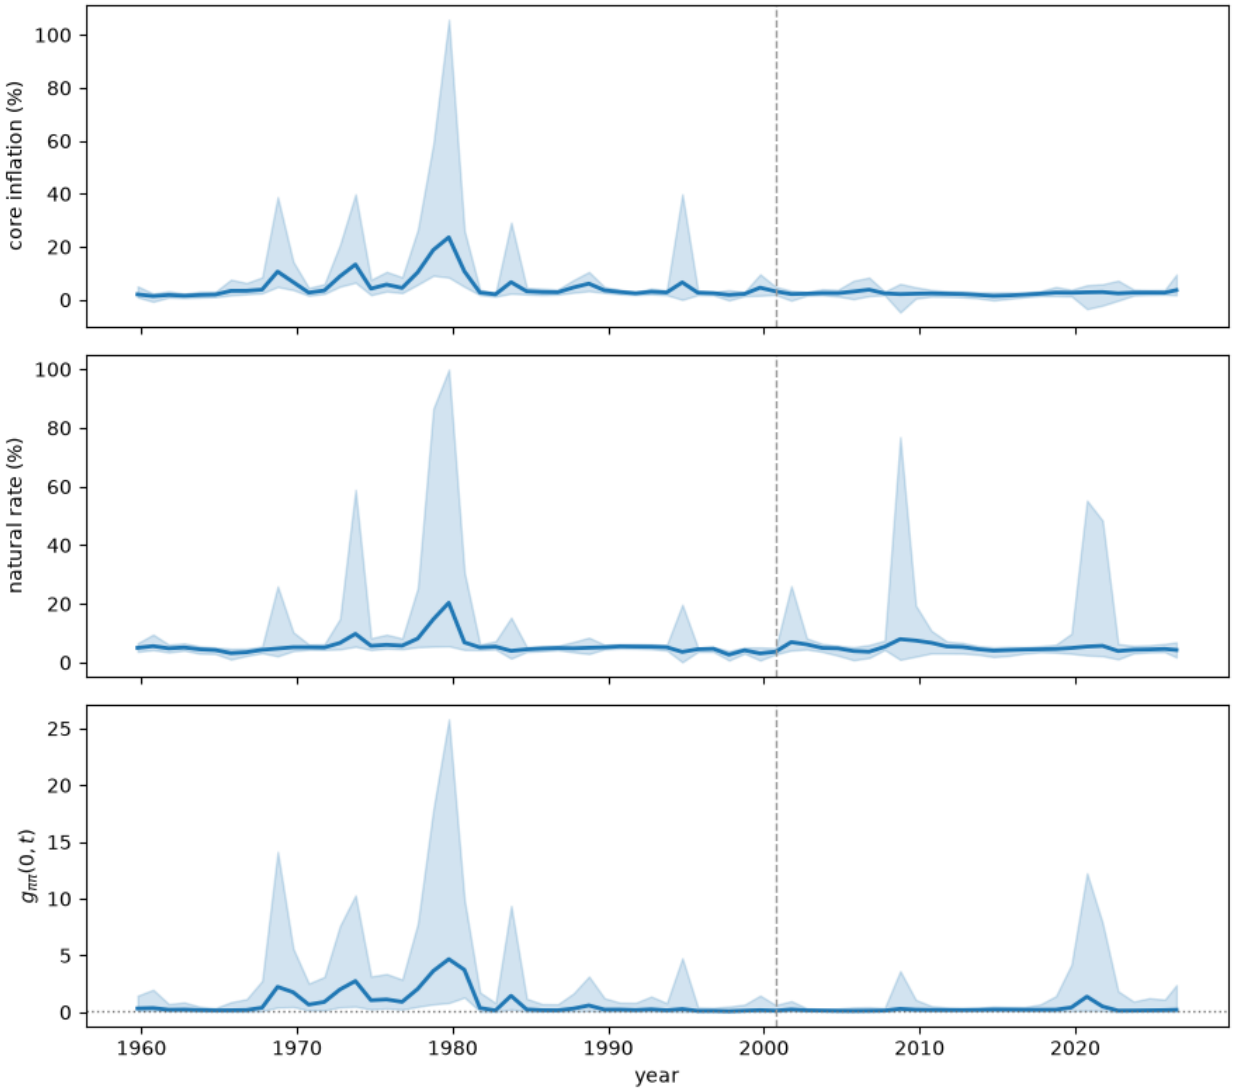


Fig. 108.5: Posterior medians and pointwise 90 percent intervals for $\bar{\pi}_t$, $\bar{u}_t$, and $g_{\pi\pi}(0,t)$, extended sample

(continued from previous page)

```
year of the core peak                          1979
natural rate, 1970--1979 range (%)          5.1--20.4
g(0), 1970--1979 mean                        1.973974
last year before 2001 with g(0) > 1             1983
core inflation, 2001--2019 range (%)         1.5--3.8
core inflation, 2020 onward mean (%)         2.859929
g(0), 2001--2019 mean                        0.201675
g(0), 2020 onward mean                       0.385413
```

Adding a quarter-century of data changes the estimated history of 1959–2000 a great deal.

With the larger drift rate, many draws put a root of the late-1970s VAR close to the unit circle, so the medians of core inflation and of the natural rate around 1979 are higher than in *Drifts and Volatilities*, in some runs far higher, their bands are enormous, and persistence stays above one until 1980 or later.

How large these numbers are depends on how far the chain has wandered into the tail of Q, and they differed a lot across runs with other seeds; what survives in all runs is the shape: core inflation and persistence rise through the 1970s and collapse in the early 1980s.

After 2000 the picture is calm and robust.

Core inflation stays below about 4 percent, and persistence stays close to the white-noise value, apart from a brief rise around 2020–2021 that is far smaller than in the 1970s and soon reverses.

A burst of large inflation innovations kept measured inflation high for several quarters in 2021–2022 without the strong propagation that the model finds in the 1970s.

The normalized spectrum over all frequencies shows the same thing.

```
θ_mean = posterior['θ'].mean(axis=0)
frequencies = np.linspace(0, 0.5, 41)
normalized = np.column_stack([
    inflation_spectrum(θ_mean[:, t], r_mean[t], frequencies)[1]
    for t in range(len(dates))
])
fig, ax = plt.subplots(figsize=(8, 4))
image = ax.pcolormesh(dates, frequencies, np.log10(normalized), shading='auto')
ax.axvline(2000.75, color='0.9', ls='--', lw=1)
ax.set_xlabel('year')
ax.set_ylabel('cycles per quarter')
fig.colorbar(image, ax=ax, label=r'$\log_{10} g_{\pi\pi}(\omega,t)$')
plt.show()
```

The normalized spectrum has a low-frequency ridge from the late 1960s to the early 1980s and, after 2000, only a brief one around 2020: the recent inflation surge shows up mainly as large shocks, not as lasting stronger propagation.

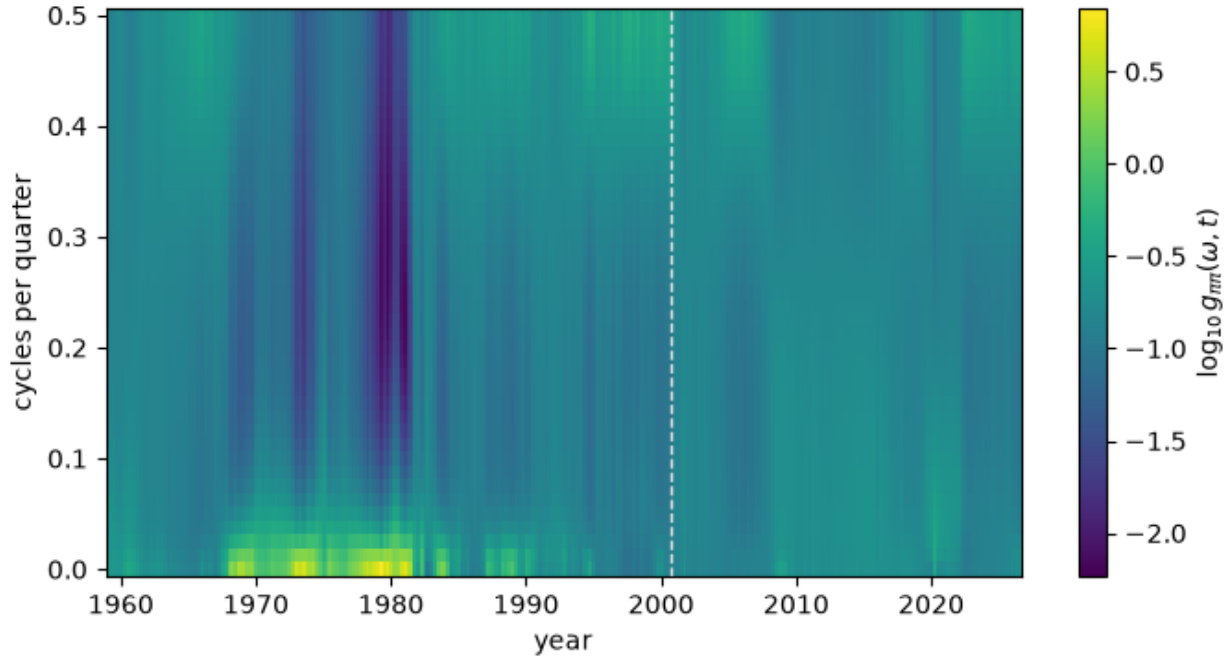


Fig. 108.6: Normalized inflation spectrum $g_{\pi\pi}(\omega, t)$, extended sample

108.7.1 Two measures of persistence

The Lost Conquest: Fed Policy in the 2020s gives the Fed a learning model in which perceived persistence is the coefficient on lagged inflation in a regression of *year-over-year* inflation on its own lag, sampled quarterly.

Year-over-year rates overlap: consecutive values share three of their four quarters.

Even if quarterly inflation were white noise, year-over-year inflation sampled quarterly would have a first-order autocorrelation of $3/4$, so a persistence measure based on it is pushed toward one and falls later and less than a measure based on quarterly rates.

The next cell compares least-squares $AR(1)$ coefficients of the two inflation rates in our data, period by period; Exercise **??** reruns that lecture's learning Fed on quarterly inflation.

```
cpi_inflation = pd.Series(400 * spliced['dp'].to_numpy(), index=spliced['date'])
year_over_year = cpi_inflation.rolling(4).mean()


def ar1_coefficient(series, first, last):
    """Least-squares AR(1) slope over the years first, ..., last."""
    z = series[(series.index >= first) & (series.index < last + 1)].dropna()
    X = np.column_stack((np.ones(len(z) - 1), z.to_numpy()[:-1]))
    return np.linalg.lstsq(X, z.to_numpy()[1:], rcond=None)[0][1]


periods = ((1960, 1969), (1970, 1979), (1985, 1999), (2000, 2019), (2020, 2100))
pd.DataFrame({
    'quarterly inflation': [ar1_coefficient(cpi_inflation, *p) for p in periods],
    'year-over-year inflation': [ar1_coefficient(year_over_year, *p)
                                 for p in periods],
}, index=['1960s', '1970s', '1985--1999', '2000--2019', '2020 onward']).round(2)
```

```
             quarterly inflation  year-over-year inflation
1960s                       0.43                      1.02
1970s                       0.80                      1.01
1985--1999                  0.37                      0.89
2000--2019                 -0.05                      0.68
2020 onward                 0.43                      0.87
```

The $AR(1)$ coefficient of year-over-year inflation is near one in the 1960s and 1970s, and it stays high after 1985 and even in 2000–2019, when the coefficient for quarterly inflation is about zero.

The overlap both inflates the year-over-year measure and blunts its decline, which is why the perceived persistence in *The Lost Conquest: Fed Policy in the 2020s* falls later and less than the quarterly measures here and in *Drifts and Volatilities*.

108.8 What the additional observations change

The extra quarter-century strengthens the bad-luck evidence for the recent period.

- It adds large volatility episodes, the financial crisis in inflation and the pandemic in unemployment, and a zero-lower-bound period in which the short rate says little about policy.
- The recent inflation surge did not bring back 1970s-style propagation: persistence after 2000 stays near the white-noise value, and core inflation stays low.
- The drift rate, and with it the estimated history of the 1970s, becomes much harder to pin down: the posterior puts weight on far faster drift than the 1959–2000 fit, and a chain short enough for a lecture cannot explore that tail reliably, so magnitudes that depend on Q should not be taken at face value.

The shape of the earlier drift documented in *Drifts and Volatilities* survives: persistence and core inflation still rise in the 1970s and fall in the early 1980s.

The next and last lecture, *The Lost Conquest: Fed Policy in the 2020s*, returns to a learning Fed: it asks whether drifting beliefs about persistence and the slope of the Phillips curve, of the kind that *Adaptive Learning and Escape Dynamics* and *Priors, Escapes, and Learning Cycles* model, help explain the Fed's slow response to the inflation of 2021–2022, and so whether the conquest of American inflation was lost.

108.9 Exercises

Exercise 108.9.1

Sims [2001] pointed out that *filtered* estimates of drifting coefficients wander even when the true coefficients are constant, which is why Cogley and Sargent [2005] report *smoothed* estimates.

Simulate $y_t = \rho_t y_{t-1} + w_t$, $w_t \sim N(0,1)$, for $T = 168$ quarters in two cases: (a) $\rho_t = 0.5$ for all t, and (b) ρ_t a random walk that starts at 0.5 with innovation variance 4×10^{-4}.

In each case run the scalar Kalman filter of *A First Look at the Kalman Filter* that treats ρ_t as a random walk with innovation variance $q = 4 \times 10^{-4}$, starting from mean 0 and variance 1, and then the smoother of *Appendix: the sampler in detail*.

Compare the standard deviations over time of the true, filtered, and smoothed paths, ignoring the first 20 quarters.

What does the smoothed path look like if the filter assumes $q = 0$?

At the last date of a sample the smoothed and filtered estimates coincide.

What does that imply for the estimates in this lecture that bear on the 2020s?

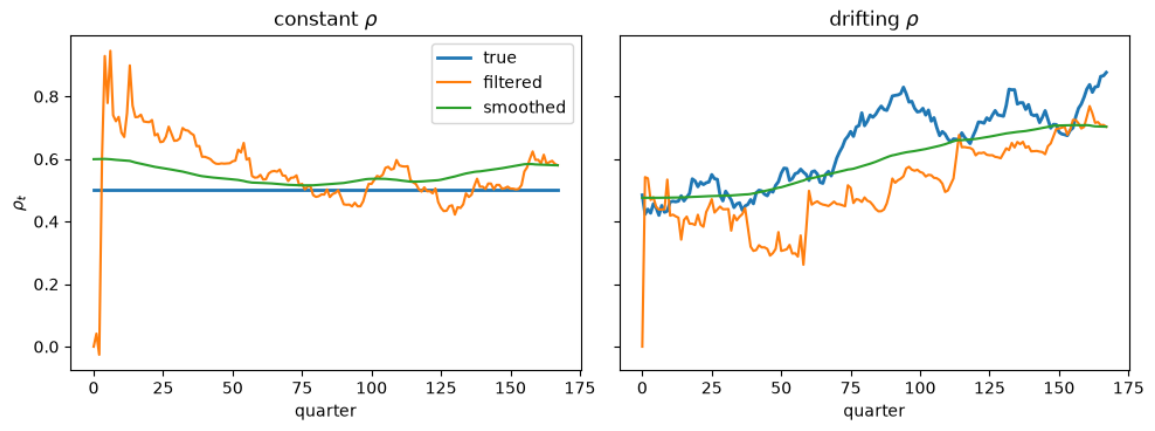
Solution

Here is one solution.

```
def tvp_ar1(T, ρ0, q, rng):
    """Simulate y_t = ρ_t y_{t-1} + w_t with ρ_t a random walk."""
    ρ = ρ0 + np.r_[0, np.cumsum(np.sqrt(q) * rng.standard_normal(T))]
    y = np.zeros(T + 1)
    for t in range(1, T + 1):
        y[t] = ρ[t] * y[t - 1] + rng.standard_normal()
    return y, ρ


def filter_and_smooth(y, q, mean0=0.0, var0=1.0):
    """Scalar Kalman filter and Rauch-Tung-Striebel smoother for ρ_t."""
    T = len(y) - 1
    m, P, P_pred = np.empty(T + 1), np.empty(T + 1), np.empty(T + 1)
    m[0], P[0] = mean0, var0
    for t in range(1, T + 1):
        P_pred[t] = P[t - 1] + q
        K = P_pred[t] * y[t - 1] / (y[t - 1]**2 * P_pred[t] + 1)
        m[t] = m[t - 1] + K * (y[t] - y[t - 1] * m[t - 1])
        P[t] = P_pred[t] - K * y[t - 1] * P_pred[t]
    s = m.copy()
    for t in range(T - 1, -1, -1):
        s[t] = m[t] + P[t] / P_pred[t + 1] * (s[t + 1] - m[t])
    return m, s


rng = np.random.default_rng(3)
fig, axes = plt.subplots(1, 2, figsize=(10, 3.8), sharey=True)
rows = {}
for ax, (q_true, title) in zip(axes, ((0.0, r'constant $\rho$'),
                                      (4e-4, r'drifting $\rho$'))):
    y, ρ = tvp_ar1(168, 0.5, q_true, rng)
    filtered, smoothed = filter_and_smooth(y, q=4e-4)
    _, flat = filter_and_smooth(y, q=0.0)
    for path, label in ((ρ, 'true'), (filtered, 'filtered'),
                        (smoothed, 'smoothed')):
        ax.plot(path[1:], lw=2 if label == 'true' else 1.5, label=label)
    ax.set_title(title)
    ax.set_xlabel('quarter')
    rows[title] = {'true': ρ[20:].std(), 'filtered': filtered[20:].std(),
                   'smoothed': smoothed[20:].std(),
                   'smoothed, q = 0': flat[20:].std()}
axes[0].set_ylabel(r'$\rho_t$')
axes[0].legend()
plt.tight_layout()
plt.show()
```



The table below gives the numbers behind the figure.

```
pd.DataFrame(rows).T.round(3)
```

	true	filtered	smoothed	smoothed, q = 0
constant ρ	0.000	0.070	0.022	0.0
drifting ρ	0.117	0.127	0.081	0.0

When ρ is constant, the filtered estimate wanders, with a standard deviation over time several times that of the smoothed estimate, which is nearly flat; that is Sims's point.

When ρ drifts, the filtered path adds transient noise to the drift, while the smoothed path is smoother than the truth.

With $q = 0$ the smoother returns an exactly flat path, as in Cogley and Sargent's footnote 16: the backward pass has $J_t = 1$ and no noise, so every date receives the end-of-sample estimate.

Near the end of a sample the smoother has few later observations to draw on, so its estimates behave more like filtered ones, and at the last date they are the filtered estimate.

The estimates for the 2020s in this lecture therefore carry some of the transient wandering that Sims warned about, and they will be revised as new quarters arrive, which is one more reason to read the endpoint bands in *Core inflation, the natural rate, and persistence* with care.

Exercise 108.9.2

A Kalman filter on coefficients that drift as random walks behaves like constant-gain least squares, and this exercise asks how large the implied gain is.

1. For the local-level model $y_t = \mu_t + w_t$, $\mu_t = \mu_{t-1} + v_t$, with $\mathbb{V}[w_t] = 1$ and $\mathbb{V}[v_t] = q$, show that the steady-state predicted variance is $p = (q + \sqrt{q^2 + 4q})/2$ and the steady-state Kalman gain is $K = p/(1+p) \approx \sqrt{q}$ for small q. Check this numerically by iterating the Riccati equation for $q \in \{10^{-4}, 10^{-3}, 10^{-2}\}$.
2. Run the Kalman filter of *Appendix: the sampler in detail* for the 21 coefficients with Q and R_t fixed at their posterior means. At each date compute w_t, the weight that the updated fitted value of inflation puts on the newest inflation forecast error, which is the $(3,3)$ element of $X_t^\top K_t$.
3. Under constant-gain least squares, that weight is the gain times $x_t^\top R^{-1} x_t$, where R is the regressors' second-moment matrix, and $x_t^\top R^{-1} x_t$ averages k, the number of regressors. Convert w_t into an equivalent gain $w_t/7$ and compare its average by decade with the gains used in this section: 0.025 in *Adaptive Learning and Escape Dynamics*, 0.02 in *A government learning the Phillips curve*, and 0.03 in *The Lost Conquest: Fed Policy in the 2020s*.

Solution

In steady state the predicted variance satisfies $p = p - p^2/(p+1) + q$, so $p^2 - qp - q = 0$ and $p = (q + \sqrt{q^2 + 4q})/2 \approx \sqrt{q}$, which gives $K = p/(1+p) \approx \sqrt{q}$.

```
for q in (1e-4, 1e-3, 1e-2):
    p = 1.0
    for _ in range(10_000):
        p = p - p**2 / (p + 1) + q
    formula = (q + np.sqrt(q**2 + 4 * q)) / 2
    print(f'q = {q:g}: iterated gain {p / (1 + p):.4f}, '
          f'formula {formula / (1 + formula):.4f}, sqrt(q) {np.sqrt(q):.4f}')
```

```
q = 0.0001: iterated gain 0.0100, formula 0.0100, sqrt(q) 0.0100
q = 0.001: iterated gain 0.0311, formula 0.0311, sqrt(q) 0.0316
q = 0.01: iterated gain 0.0951, formula 0.0951, sqrt(q) 0.1000
```

For the VAR we run the forward recursion with the posterior means of Q and R_t and record the weight on the newest inflation surprise.

```
q_bar = posterior['Q'].mean(axis=0)
X = design_matrices(data['x'])
P = prior['θ_covariance'].copy()
weight = np.empty(len(dates))
for t in range(len(dates)):
    predicted = P + q_bar
    px = predicted @ X[t].T
    gain = np.linalg.solve(X[t] @ px + r_mean[t], px.T).T
    weight[t] = (X[t] @ gain)[2, 2]
    P = predicted - gain @ px.T

decades = {f'{d}s': (d, d + 9) for d in (1960, 1970, 1980, 1990, 2000, 2010)}
decades['2020 onward'] = (2020, 2100)
pd.DataFrame({
    'weight on newest surprise': [period_mean(weight, *p) for p in decades.
 ↪values()],
    'equivalent gain per regressor': [period_mean(weight, *p) / 7
                                      for p in decades.values()],
}, index=decades.keys()).round(3)
```

	weight on newest surprise	equivalent gain per regressor
1960s	0.383	0.055
1970s	0.538	0.077
1980s	0.660	0.094
1990s	0.477	0.068
2000s	0.208	0.030
2010s	0.191	0.027
2020 onward	0.312	0.045

The steady-state gain is close to $\sqrt{q}$, as the formula says.

For the VAR, the weight on the newest inflation surprise is largest before 2000, when the equivalent gain per regressor is roughly two to five times the 0.025 of *Adaptive Learning and Escape Dynamics*; after 2000 it falls to between about one and two times that value, close to the gains of 0.02 to 0.03 used in this section.

These gains inherit the uncertainty about Q noted above, because the gain grows with the drift covariance.

Exercise 108.9.3

Two Illustrations of the Quantity Theory of Money interprets a long-run slope as a ratio of spectral densities at frequency zero, $h(0) = S_{yz}(0)/S_z(0)$.

Compute the analogous low-frequency slope of inflation on unemployment implied by the local VAR at each date,

$$h_{\pi u}(0,t) = \frac{\operatorname{Re} f_{\pi u}(0,t)}{f_{uu}(0,t)},$$

draw by draw, and convert it into percentage points of annual inflation per percentage point of the unemployment rate, using the natural rate $\bar{u}_t$ to undo the logit transformation.

Plot the posterior median and a 90 percent band.

When was the low-frequency slope negative, as a long-run Phillips tradeoff would require?

Solution

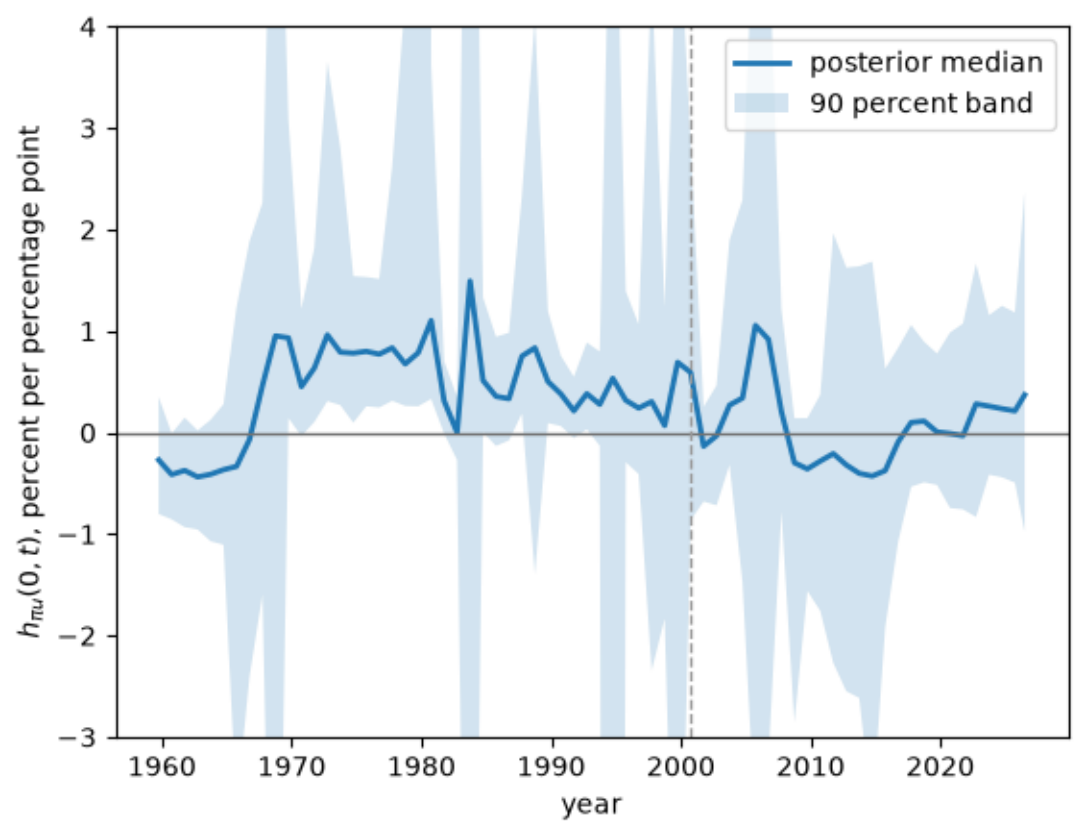
At frequency zero the companion system gives $2\pi f(0,t) = (I - A_t)^{-1}\mathcal{R}_t(I - A_t)^{-\top}$.

Inflation is a quarterly log change, so it is multiplied by 400; the logit satisfies $du_t = 10^{-4}\, dU/[\bar{u}_t(1-\bar{u}_t)]$ when dU is measured in percentage points.

```
def low_frequency_slope(θ, covariance):
    """Low-frequency slope of annual inflation on the unemployment rate (pp/pp)."""
    intercept, companion = companion_matrix(θ)
    inverse = np.linalg.inv(np.eye(6) - companion)
    innovation = np.zeros((6, 6))
    innovation[:3, :3] = covariance
    S = inverse @ innovation @ inverse.T
    logit_u = 100 * (inverse @ intercept)[1]
    density = expit(logit_u) * expit(-logit_u)      # u(1 - u) at the natural rate
    if density == 0:                                # natural rate at 0 or 100%
        return np.nan
    return 400 * S[2, 1] / S[1, 1] * 1e-4 / density


slopes = np.array([
    [low_frequency_slope(posterior['θ'][d][:, t],
                         innovation_covariance(posterior['h'][d][t],
                                               posterior['β'][d]))
     for t in annual]
    for d in range(len(posterior['β']))
])
lower, median, upper = np.nanquantile(slopes, (0.05, 0.5, 0.95), axis=0)

fig, ax = plt.subplots()
ax.plot(dates[annual], median, lw=2, label='posterior median')
ax.fill_between(dates[annual], lower, upper, alpha=0.2, label='90 percent band')
ax.axhline(0, color='0.45', lw=1)
ax.axvline(2000.75, color='0.6', ls='--', lw=1)
ax.set_xlabel('year')
ax.set_ylabel(r'$h_{\pi u}(0,t)$, percent per percentage point')
ax.set_ylim(-3, 4)      # a few draws near the unit-root boundary explode
ax.legend()
plt.show()
```



The table below gives the numbers behind the figure.

```
pd.DataFrame({'5%': lower, 'median': median, '95%': upper},
             index=annual_years).loc[[1962, 1970, 1975, 1980, 1985, 1990, 2000,
                                      2010, 2020, annual_years[-1]]].round(2)
```

	5%	median	95%
1962	−0.95	−0.44	0.02
1970	−0.03	0.45	1.22
1975	0.26	0.80	1.53
1980	0.33	1.10	3.62
1985	−0.13	0.36	0.94
1990	0.06	0.38	0.76
2000	−0.86	0.59	2.85
2010	−1.76	−0.28	0.37
2020	−0.74	−0.01	0.99
2026	−0.97	0.37	2.36

The low-frequency slope is positive in the 1970s, with bands that exclude zero in the second half of the decade: at low frequencies, inflation and unemployment rose together, the opposite of an exploitable long-run tradeoff.

It is negative in the early 1960s, and close to zero and imprecise after 2000.

The vertical axis is cut off because in a few years some draws lie near the unit-root boundary, where the slope explodes.

These are reduced-form statistics: as Sargent [1971] and Lucas [1972] argued, a low-frequency slope is not by itself a structural tradeoff that policy could exploit.

CHAPTER

NINE

THE LOST CONQUEST: FED POLICY IN THE 2020S

Contents

- *The Lost Conquest: Fed Policy in the 2020s*
  - *Overview*
  - *The three elements*
  - *The Fed's drifting-coefficients beliefs*
  - *The Fed's Phelps problem*
  - *Why the Fed was slow: a counterfactual*
  - *Real-time data and the delay*
  - *How beliefs shape the Phelps rule*
  - *From policy to measured beliefs*
  - *Was the conquest lost?*
  - *Exercises*

109.1 Overview

The previous lecture, *Drifts and Volatilities after 2000*, carried the drifting-coefficients evidence of Cogley and Sargent [2005] past 2000: it asked how the reduced-form dynamics of inflation, and the volatilities of the shocks hitting them, moved in the decades after the *Conquest* was written.

This final lecture of the suite takes the policy maker's point of view instead.

It follows Sargent and Williams [2025], who impute drifting beliefs to the Federal Reserve and ask what policy those beliefs recommended during the inflation of the 2020s.

Their artificial Fed is built from tools that earlier lectures developed:

- a *Phelps control problem*;
- the *anticipated-utility* assumption of Kreps [1998], met in *Adaptive Learning and Escape Dynamics*;
- a *drifting-coefficients* model estimated by *constant-gain recursive least squares*, the constant-gain special case of the Kalman filter of *Priors, Escapes, and Learning Cycles*.

The puzzle is a live one.

After the COVID-19 pandemic, U.S. inflation surged to its highest rate since the early 1980s, the episode at the end of the post-1999 data in *The Rise and Fall of U.S. Inflation*.

That lecture plots CPI inflation; here, following the paper, we use the price index for personal consumption expenditures (PCE), the Fed's preferred measure.

Yet for more than a year the Federal Reserve did not respond.

It first raised its policy rate in March 2022 and began to tighten in earnest only in the second quarter of 2022, with inflation already at its peak.

Why was the Fed so slow?

Sargent and Williams [2025] build an *artificial Fed* that, each quarter,

- re-estimates a drifting-coefficients Phillips curve by constant-gain recursive least squares, and
- solves a linear-quadratic Phelps problem, treating its latest estimates as if they would hold forever, to set its interest-rate instrument.

Control engineers call this *model predictive control*.

It is the *estimate-then-optimize* loop that drove the learning government of *Adaptive Learning and Escape Dynamics*, but now the instrument is the policy rate and the beliefs concern the *persistence* of inflation and the *slope* of the Phillips curve.

Three elements shape the artificial Fed's decisions:

1. **Declining inflation persistence.** The Fed had learned that inflation shocks fade faster than they used to, so the surge looked *transitory*.
2. **A flatter Phillips curve.** The Fed had learned that inflation responds weakly to slack, so disinflation looked *costly*.
3. **Real-time mismeasurement of the output gap.** In real time the Fed perceived more slack than the revised data now show.

We reproduce all three from public data.

The main finding, which is also the paper's, fits in one sentence:

> drifting beliefs make the response *muted*; real-time mismeasurement makes it *late*.

With today's revised data, the artificial Fed responds to the 2021 surge mildly, but *earlier* than the actual Fed did.

Only with the output gap that the Fed saw in real time does its response arrive at about the time the Fed's did.

The last part of the lecture is our own extension.

It asks where the benign beliefs came from: can aggressive past policy make inflation *look* transitory and the Phillips curve *look* flat?

We compute the self-confirming fixed point of that loop and find that, unlike the government of *Escaping Nash Inflation*, this Fed has no escape route.

Note

Readers arriving from the Bounded Rationality suite will find γ in its usual role as the learning gain.

There the beliefs are called α (*Notation across the series*, *A government learning the Phillips curve*); here they are θ_t.

Let's import what we need:

```
import time
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from scipy.linalg import solve_discrete_are
from scipy.optimize import brentq
```

109.2 The three elements

The first two elements are among the most documented facts in modern macroeconomics.

Declining persistence. From the 1970s into the 1980s inflation was highly persistent: a shock raised inflation for years.

Cogley and Sargent [2001] and Cogley and Sargent [2005] documented the decline in inflation persistence during the Volcker–Greenspan conquest.

Drifts and Volatilities reproduces their evidence (*Inflation persistence*), and *Drifts and Volatilities after 2000* carries it past 2000 (*Core inflation, the natural rate, and persistence*).

Stock and Watson [2007] found that the permanent component of inflation shrank in the 1980s and 1990s.

Benati and Surico [2008] and Carlstrom *et al.* [2009] linked the decline in persistence to more aggressive monetary policy, a link we take up in *From policy to measured beliefs*.

Pivetta and Reis [2007] cautioned that part of the apparent decline may reflect smaller shocks rather than a structural break; that is the distinction between drifting coefficients and drifting volatilities that *Drifts and Volatilities* draws.

Former Fed Chair Janet Yellen summarized the evidence in 2019: "the slope of the Phillips curve … has diminished very significantly since the 1960s … And second, inflation has become much less persistent, because the impact of lagged inflation on current inflation has declined considerably."

A flatter Phillips curve. Since the 1990s, estimates of the Phillips curve's slope have trended toward zero; the "missing disinflation" after the Great Recession is the leading example.

Policy makers knew what this implied.

Bernanke [2022] (p. 220) writes: "A flat Phillips curve means that inflation is a less reliable indicator of economic overheating. Should inflation get too high, the costs, in terms of unemployment, of bringing inflation back down to target could be higher than in the past."

Notice that a flat Phillips curve points policy in opposite directions here and in the *Conquest*.

In the Lucas-style models of Sargent [1999], a flatter curve means that surprise inflation does less to real activity, so the Phelps problem recommends *lower* inflation.

The induction hypothesis of *Adaptive Expectations and the Phelps Problem* (*Testing the natural-rate hypothesis*), under which inflation has no long-run real effect, is the limiting case, and it is how the *Conquest*'s adaptive government brings inflation down.

In the statistical model of the 2020s, the Fed controls inflation by moving the output gap.

There a flatter curve means a higher sacrifice ratio (*The sacrifice ratio*): each point of disinflation costs more output, so the Phelps problem counsels patience (Sargent and Williams [2025], footnote 5).

Real-time uncertainty. The third element, emphasized by Orphanides [2001], is that the output gap is badly mismeasured in *real time*, especially at business-cycle turning points.

Through 2020–2023 the real-time gap was persistently *below* the later revised measure, so the Fed perceived more slack than there was.

We measure this in *Real-time data and the delay*.

109.3 The Fed's drifting-coefficients beliefs

We give the artificial Fed a backward-looking Phillips curve with drifting coefficients,

$$\pi_t = \alpha_{0,t} + \rho_t \, \pi_{t-1} + \kappa_t \, x_t + \varepsilon_t^{\pi}, \tag{109.1}$$

where π_t is four-quarter PCE inflation, x_t is the CBO output gap, ρ_t is the *perceived persistence*, and κ_t is the *perceived slope*.

ρ_t plays the role that the sum of the coefficients on lagged inflation plays in the *Conquest*'s Phelps problems (*Testing the natural-rate hypothesis*; Sargent and Williams [2025], footnote 11).

The Fed updates $\theta_t = (\alpha_{0,t}, \rho_t, \kappa_t)^\top$ by constant-gain recursive least squares,

$$\theta_t = \theta_{t-1} + \gamma \, M_t^{-1} X_t \left(\pi_t - X_t^\top \theta_{t-1}\right), \qquad M_t = M_{t-1} + \gamma \left(X_t X_t^\top - M_{t-1}\right), \tag{109.2}$$

with regressors $X_t = (1, \pi_{t-1}, x_t)^\top$ and moment matrix M_t.

The estimate θ_t uses data through quarter t, and it is what the Fed brings to its decision in quarter t.

The moment matrix is updated before the coefficients, the timing of Evans and Honkapohja [2001].

The paper's display writes the lagged matrix M_{t-1} in the coefficient update.

The two timings differ by terms of order γ^2 unless the regressors move far outside their past range.

In 2022 they did, and the timing then changes the recommendations noticeably, so it is worth saying which one the code uses.

The gain $\gamma = 0.03$ is the paper's choice for quarterly data; it discounts past observations so that the estimates can track drift.

Constant-gain least squares is the special case of the Kalman filter of *Priors, Escapes, and Learning Cycles* (see also *A First Look at the Kalman Filter*) that Sargent and Williams [2005] obtain with a particular prior about how the coefficients drift.

Which gain best rationalizes the Fed's decisions is a version of Chung's problem, which we pose in Exercise 109.10.4.

We read the PCE price index, real GDP, CBO potential output, and the federal funds rate directly from FRED.

If FRED lacks an isolated month, as happened for some series around the federal government shutdown of October 2025, the code fills it by linear interpolation and says so.

```
FRED = 'https://fred.stlouisfed.org/graph/fredgraph.csv?id='

def fred(series_id):
    "Read one series directly from FRED."
    s = pd.read_csv(FRED + series_id, index_col=0, parse_dates=True).iloc[:, 0]
    s.index.name = None
    return pd.to_numeric(s, errors='coerce')      # FRED marks gaps with '.'

def quarterly_mean(monthly):
    "Quarterly averages of a monthly series, filling isolated missing months."
    monthly = monthly.asfreq('MS')
    missing = monthly.index[monthly.isna()]
    if len(missing) > 0:
        print(f"{monthly.name}: interpolating",
              [f"{d:%Y-%m}" for d in missing])
    filled = monthly.interpolate(limit=1, limit_area='inside')
    q = filled.resample('QS')
    return q.mean()[q.count() == 3]                # complete quarters only
```

(continues on next page)

(continued from previous page)

```
def qtr(date):
    "Label a date by its quarter, e.g. 2022Q2."
    return f"{date.year}Q{date.quarter}"

pce = quarterly_mean(fred('PCEPI'))
fedfunds = quarterly_mean(fred('FEDFUNDS'))
gdp, potential = fred('GDPC1'), fred('GDPPOT')      # real GDP, CBO potential

inflation = 100 * (pce / pce.shift(4) - 1)           # four-quarter PCE inflation
gap = 100 * (gdp / potential - 1)                    # output gap, percent

data = pd.concat([inflation.rename('pi'), gap.rename('x'),
                  fedfunds.rename('i')], axis=1, sort=True).dropna().loc['1960':]
print(f"sample: {qtr(data.index[0])} to {qtr(data.index[-1])}, "
      f"{len(data)} quarters")

peak = data['pi']['2020':].idxmax()
higher = data['pi'][:peak].iloc[:-1]
last_higher = higher[higher > data['pi'][peak]].index[-1]
print(f"PCE inflation peaked at {data['pi'][peak]:.2f}% in {qtr(peak)}, "
      f"the highest since {qtr(last_higher)}")
```

```
sample: 1960Q1 to 2026Q3, 267 quarters
PCE inflation peaked at 6.88% in 2022Q2, the highest since 1981Q4
```

Following the paper, we freeze beliefs during 2020 and 2021, setting the gain to zero, because the pandemic observations were extreme outliers that would otherwise whipsaw the estimates; updating resumes in 2022.

The freeze is not innocuous.

The 2020 quarters are the true outliers, but it is the freeze of 2021 that drives the results below.

If the Fed had learned from the 2021 data, its perceived Phillips curve would have steepened and its recommended rate would have jumped to several percent by the end of 2021 (Exercise 109.10.3).

Setting the gain to zero in 2021 is itself the judgment that the 2021 surge was transitory.

We start the recursion at $\theta = (0.5, 0.9, 0.05)$ rather than at the paper's 1949–1959 presample regression; the belief paths after the mid-1980s hardly depend on this choice.

We compute the beliefs once and reuse them below.

```
def estimate_beliefs(data, γ=0.03, freeze=(2020, 2021)):
    "Constant-gain RLS for the drifting Phillips curve: paths of α₀, ρ, κ."
    π, x = data['pi'].values, data['x'].values
    θ = np.array([0.5, 0.9, 0.05])            # [intercept, persistence, slope]
    M = np.diag([1.0, 10.0, 5.0])             # moment matrix
    rows = []
    for t in range(1, len(data)):
        γ_t = 0.0 if data.index[t].year in freeze else γ
        X = np.array([1.0, π[t - 1], x[t]])
        M = M + γ_t * (np.outer(X, X) - M)                    # M_t first,
        θ = θ + γ_t * np.linalg.solve(M, X * (π[t] - X @ θ))  # then θ_t
        rows.append(θ)
    return pd.DataFrame(rows, index=data.index[1:],
                        columns=['alpha0', 'rho', 'kappa'])
```

(continues on next page)

(continued from previous page)

```
beliefs = estimate_beliefs(data)
```

Here are the perceived persistence and slope.

```
fig, axes = plt.subplots(2, 1, figsize=(10, 6), sharex=True)
axes[0].plot(beliefs['rho'])
axes[0].axhline(1, color='k', lw=0.5, ls=':')
axes[0].set_ylabel(r'persistence $\rho_t$')

axes[1].plot(beliefs['kappa'], color='C1')
axes[1].axhline(0, color='k', lw=0.5, ls=':')
axes[1].set_ylabel(r'slope $\kappa_t$')
axes[1].set_xlabel('year')

plt.tight_layout()
plt.show()

ρ, κ = beliefs['rho'], beliefs['kappa']
after85 = ρ['1985':]
print(f"ρ_t over 1985-2001: mean {ρ['1985':'2001'].mean():.3f}, "
      f"range {ρ['1985':'2001'].min():.3f} to {ρ['1985':'2001'].max():.3f}")
print(f"first quarter after 1985 with ρ_t < 0.96: "
      f"{qtr(after85[after85 < 0.96].index[0])}")
print(f"trough of ρ_t: {ρ.min():.3f} in {qtr(ρ.idxmin())}")
print(f"κ_t: {κ['1980':'1984'].mean():.3f} (1980-84 average), "
      f"{κ['2009':'2009'].max():.3f} (2009 peak), "
      f"{κ['2019-10-01']:.3f} (2019Q4)")
print(f"2019Q4: ρ = {ρ['2019-10-01']:.3f};   2022Q2: ρ = {ρ['2022-04-01']:.3f}")
```

```
ρ_t over 1985-2001: mean 0.989, range 0.976 to 1.010
first quarter after 1985 with ρ_t < 0.96: 2006Q3
trough of ρ_t: 0.814 in 2011Q1
κ_t: 0.131 (1980-84 average), 0.198 (2009 peak), 0.039 (2019Q4)
2019Q4: ρ = 0.858;   2022Q2: ρ = 0.981
```

Perceived **persistence** ρ_t jumps in the mid-1970s and then stays close to one for a quarter of a century, through 2001.

It drifts down only slowly in the 2000s and falls steeply after the 2008 recession, to a trough in 2011.

Sargent and Williams [2025] describe a decline after the mid-1980s "with an especially precipitous fall following the 2008 recession"; in this replication the decline is mainly that post-2008 fall.

When updating resumes in 2022, ρ_t jumps back toward one, just as the Fed abandoned the "transitory" characterization and began to tighten.

Perceived **slope** κ_t declines from its early-1980s level, spikes during the 2008–2009 recession, when the economy briefly experienced deflation, and is lower still by 2019: the Phillips curve flattens.

A word on the scale of ρ_t.

Because π_t is a four-quarter rate, consecutive observations share three quarters of data.

Even if quarterly inflation were serially uncorrelated, a regression of π_t on π_{t-1} would give a coefficient near 0.75 (Exercise 109.10.5).

Against that floor, the 2019 value of ρ_t means *less persistent*, not *transitory*:

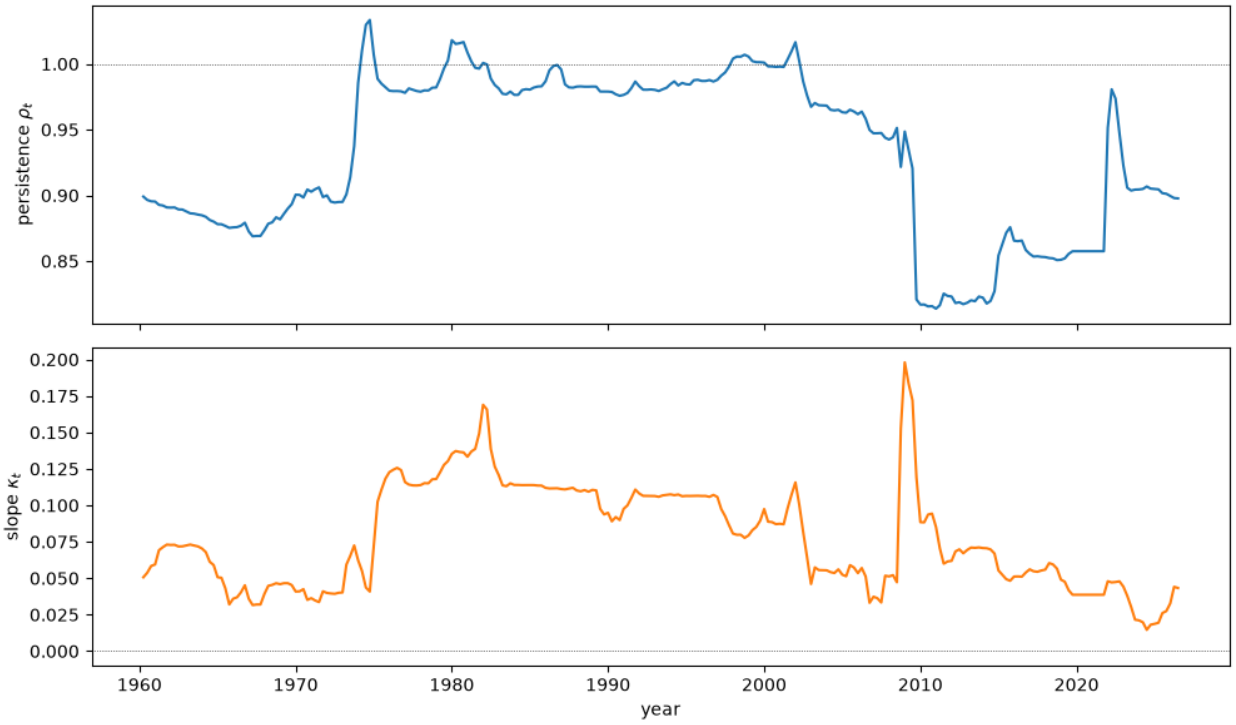


Fig. 109.1: The Fed's perceived inflation persistence and Phillips curve slope

```
ρ_2019 = ρ['2019-10-01']
print(f"half-life of an inflation shock at ρ = {ρ_2019:.3f}: "
      f"{np.log(0.5) / np.log(ρ_2019):.1f} quarters")
```

```
half-life of an inflation shock at ρ = 0.858: 4.5 quarters
```

With quarterly annualized inflation, which has no overlap, the same algorithm finds persistence falling earlier and much further.

That measure is closer in spirit to the one in *Drifts and Volatilities*; Exercise 109.10.5 reconciles the two.

Still, by 2019 the Fed's model said that inflation was less persistent than it used to be and only weakly linked to slack.

Fed into a Phelps problem, these beliefs counsel patience.

109.4 The Fed's Phelps problem

Each quarter the Fed sets its policy rate by solving a linear-quadratic Phelps problem, treating its *current* estimates as if they will hold forever.

That is the anticipated-utility assumption of Kreps [1998] that we met in *Adaptive Learning and Escape Dynamics*.

The belief Phillips curve (109.1) is paired with an IS curve

$$x_t = b_0 + b_1 x_{t-1} + g\,(i_{t-1} - \pi_{t-1}) + \varepsilon_t^x,$$

whose coefficients are estimated once, by OLS over the whole sample.

Together they give linear dynamics for the state $s_t = (1, \pi_t, x_t, i_{t-1})^\top$,

$$s_{t+1} = A_t s_t + B_t\, i_t + C\varepsilon_{t+1},$$

whose matrices depend on the time-t beliefs.

The Fed minimizes

$$\mathbb{E}_t \sum_{j=t}^{\infty} \beta^{j-t}\Big[(\pi_j - \pi^*)^2 + \lambda_x\, x_j^2 + \eta\,(i_j - i_{j-1})^2\Big], \tag{109.3}$$

where $\pi^* = 2$ is the target and the last term penalizes abrupt rate changes.

We use the paper's parameter values: $\beta = 0.95$, $\lambda_x = 0.2$, and $\eta = 0.5$.

The paper says these values were "loosely tuned to provide a reasonable match to the observed policy outcomes."

Writing the period loss as $s^\top Q s + 2 s^\top N i + R\, i^2$ with $R = \eta$ turns this into a discounted linear-quadratic regulator with a cross term (*LQ Control: Foundations*).

The matrix P_t of the expected discounted loss $s^\top P_t s$ (plus a constant) solves the discounted Riccati equation

$$P = Q + \beta A^\top P A - (\beta A^\top P B + N)(R + \beta B^\top P B)^{-1}(\beta B^\top P A + N^\top), \tag{109.4}$$

and the decision rule is $i_t = -F_t s_t$ with $F_t = (R + \beta B^\top P_t B)^{-1}(\beta B^\top P_t A + N^\top)$.

Written out, the rule is a *smoothed Taylor rule*,

$$i_t = \phi_{0,t} + \phi_{\pi,t}\,\pi_t + \phi_{x,t}\,x_t + \phi_{i,t}\,i_{t-1},$$

whose coefficients move as beliefs move.

`scipy.linalg.solve_discrete_are` solves the undiscounted version of (109.4).

Passing $\sqrt{\beta}A$ and $\sqrt{\beta}B$ absorbs the discount factor.

The cross term then enters unscaled, as `s=N`, because $\beta A^\top P B + N = (\sqrt{\beta}A)^\top P(\sqrt{\beta}B) + N$.

Following the paper (Sargent and Williams [2025], section 3.3), we impose the zero lower bound only ex post, $i_t^{\text{opt}} = \max\{-F_t s_t, 0\}$.

The Fed's Phelps problem does not anticipate that the bound might bind.

```
# IS curve, estimated once by OLS over the full sample
π_d, x_d, i_d = data['pi'].values, data['x'].values, data['i'].values
n = len(data)
Z_is = np.column_stack([np.ones(n - 1), x_d[:-1], i_d[:-1] - π_d[:-1]])
b0, b1, g = np.linalg.lstsq(Z_is, x_d[1:], rcond=None)[0]
print(f"IS curve: b0 = {b0:.3f}, b1 = {b1:.3f}, g = {g:.4f}; "
      f"long-run effect of the real rate g/(1-b1) = {g / (1 - b1):.2f}")

β, π_star, λ_x, η = 0.95, 2.0, 0.2, 0.5

def phelps_rule(θ, η=η):
    "Coefficients (ϕ₀, ϕ_π, ϕ_x, ϕ_i) of the Phelps rule i = -F s under beliefs θ."
    α0, ρ, κ = θ
    A = np.array([[1, 0, 0, 0],
                  [α0 + κ * b0, ρ - κ * g, κ * b1, 0],
                  [b0, -g, b1, 0],
                  [0, 0, 0, 0]], float)
    B = np.array([[0], [κ * g], [g], [1]], float)
```

(continues on next page)

(continued from previous page)

```
    c_π = np.array([-π_star, 1, 0, 0.])       # π - π*
    c_x = np.array([0, 0, 1, 0.])             # x
    e_i = np.array([0, 0, 0, 1.])             # i_{t-1}
    Q = np.outer(c_π, c_π) + λ_x * np.outer(c_x, c_x) + η * np.outer(e_i, e_i)
    N = (-η * e_i).reshape(4, 1)
    R = np.array([[η]])
    # √β absorbs discounting in A and B; the cross term N enters unscaled
    P = solve_discrete_are(np.sqrt(β) * A, np.sqrt(β) * B, Q, R, s=N)
    F = np.linalg.solve(R + β * B.T @ P @ B, β * B.T @ P @ A + N.T)
    return -F.ravel()

def phelps_rules(beliefs, η=η):
    "The Phelps rule at every date, given a path of beliefs."
    return pd.DataFrame([phelps_rule(θ, η) for θ in beliefs.values],
                        index=beliefs.index,
                        columns=['ψ_0', 'ψ_π', 'ψ_x', 'ψ_i'])

def recommended_rate(data, rules):
    "Recommended rate max(-F_t s_t, 0): the zero lower bound imposed ex post."
    π, x, i = data['pi'].values, data['x'].values, data['i'].values
    s = np.column_stack([np.ones(len(data) - 1), π[1:], x[1:], i[:-1]])
    return pd.Series(np.maximum((rules.values * s).sum(axis=1), 0.0),
                     index=rules.index)

rules = phelps_rules(beliefs)
optimal = recommended_rate(data, rules)
```

```
IS curve: b0 = 0.013, b1 = 0.907, g = -0.0220; long-run effect of the real rate g/
↪(1-b1) = -0.24
```

The estimated interest sensitivity g is small.

Policy endogeneity, which we meet again in *From policy to measured beliefs*, is one reason a regression can understate how much interest rates matter: a Fed that leans against expected booms weakens the correlation between the real rate and subsequent output.

```
fig, ax = plt.subplots(figsize=(10, 4.5))
window = slice('1991', None)
ax.plot(optimal[window], 'C0', label="Phelps problem's recommended rate")
ax.plot(data['i'][window], 'C3', lw=1, label='actual federal funds rate')
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()

def fit(rate, start='1991'):
    "Correlation and root mean squared error against the actual funds rate."
    r = rate[start:]
    actual = data['i'][r.index]
    return np.corrcoef(r, actual)[0, 1], np.sqrt(((r - actual)**2).mean())

no_change = data['i'].shift(1).dropna()               # i_t = i_{t-1}
for name, rate in [('Phelps rule', optimal), ('no-change rule', no_change)]:
    c, rmse = fit(rate)
    print(f"{name:>15}: correlation {c:.3f}, RMSE {rmse:.2f} "
          f"(1991-{data.index[-1].year})")
```

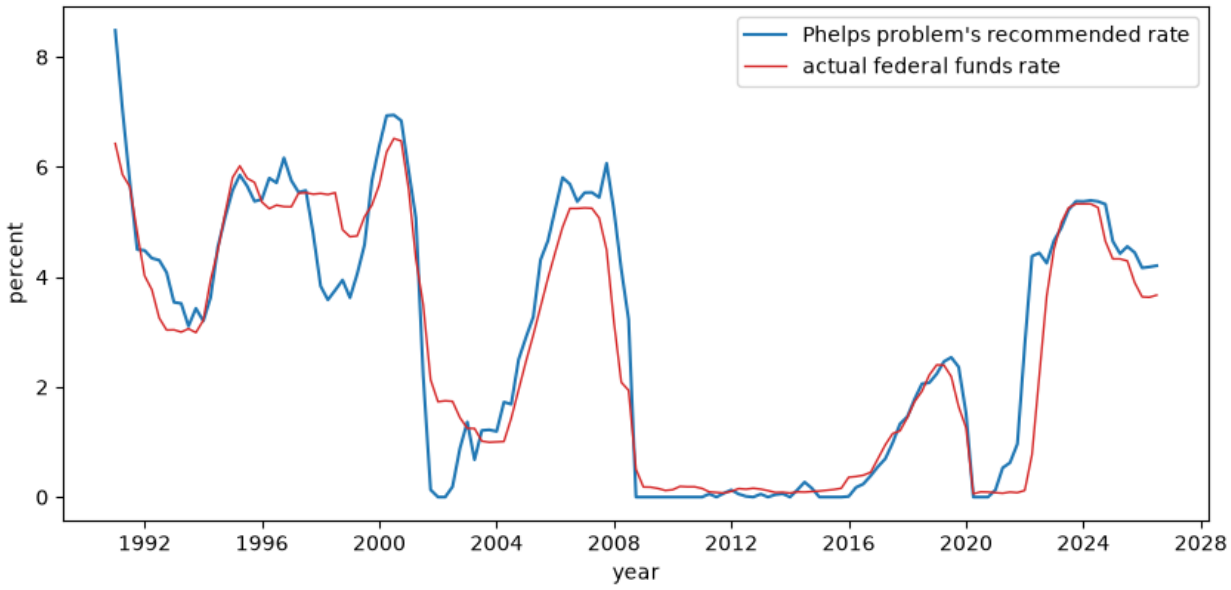


Fig. 109.2: The belief-driven Phelps rule against the actual federal funds rate

```
    Phelps rule: correlation 0.944, RMSE 0.79 (1991–2026)
no-change rule: correlation 0.979, RMSE 0.45 (1991–2026)
```

The subjectively optimal rate tracks the level and turning points of actual policy over three decades: the mid-1990s tightening, the 2001 easing, the zero-bound era after 2008, and the normalization that began in 2015.

It misses some episodes, falling well below the actual rate in 1998 and again in 2002.

The correlation needs a benchmark.

Because the recommendation conditions on the actual lagged rate i_{t-1}, a naive rule that simply sets $i_t = i_{t-1}$ does *better* on both measures, as the output shows.

The paper makes the same point about one-step-ahead fit.

What the Phelps rule offers is not a better forecast of the funds rate but an interpretation of it, in terms of beliefs and a loss function.

Now look at 2021.

Beliefs are frozen, but the state is not: as inflation and the output gap rise, the rule recommends a rising rate.

```
print(pd.DataFrame({'Phelps rule': optimal, 'actual': data['i']})
        .loc['2021':'2022'].rename(index=qtr).round(2).T.to_string())
```

```
             2021Q1  2021Q2  2021Q3  2021Q4  2022Q1  2022Q2  2022Q3  2022Q4
Phelps rule    0.13    0.53    0.62    0.96    2.76    4.38    4.44    4.25
actual         0.08    0.07    0.09    0.08    0.12    0.77    2.19    3.65
```

With today's revised data, the artificial Fed starts raising its rate in 2021 and is well ahead of the actual Fed in early 2022.

So with current-vintage data the drifting-beliefs model *leads* the Fed rather than mirroring it, as the paper also finds when it uses revised data (Sargent and Williams [2025], section 3.5).

Drifting beliefs make the recommended response mild, but they do not make it late.

We return to the delay in *Real-time data and the delay*.

109.4.1 How the rule drifted

The paper's Figure 8 shows the mechanism directly: the coefficients of the Phelps rule, and the long-run response of the policy rate to inflation, $\phi_{\pi,t}/(1-\phi_{i,t})$.

We compute them with today's data; the paper's figure uses real-time data.

A long-run response above one is the Taylor principle.

```
long_run = rules['ψ_π'] / (1 - rules['ψ_i'])

fig, axes = plt.subplots(3, 1, figsize=(10, 8), sharex=True)
axes[0].plot(rules['ψ_π'], label=r'$\phi_{\pi,t}$')
axes[0].plot(rules['ψ_x'], label=r'$\phi_{x,t}$')
axes[0].legend()
axes[0].set_ylabel('response')
axes[1].plot(rules['ψ_i'], color='C2')
axes[1].set_ylabel(r'smoothing $\phi_{i,t}$')
axes[2].plot(long_run, color='C3')
axes[2].axhline(1, color='k', lw=0.5, ls=':')
axes[2].set_ylabel(r'$\phi_{\pi,t}/(1-\phi_{i,t})$')
axes[2].set_xlabel('year')
plt.tight_layout()
plt.show()

for date in ['2000-01-01', '2019-10-01']:
    print(f"long-run inflation response in {qtr(pd.Timestamp(date))}: "
          f"{long_run[date]:.2f}")
peak_lr = long_run['2021':'2023'].idxmax()
print(f"peak in 2021-2023: {long_run[peak_lr]:.2f} in {qtr(peak_lr)}")
```

```
long-run inflation response in 2000Q1: 6.65
long-run inflation response in 2019Q4: 1.30
peak in 2021-2023: 6.05 in 2022Q2
```

As persistence and slope drifted down, the rule's long-run response to inflation fell from well above one in 2000 to just above one by 2019, while interest-rate smoothing rose.

The 2019 rule barely satisfied the Taylor principle.

When updating resumed in 2022 and ρ_t jumped, the long-run response jumped with it, before falling back as inflation receded.

Part of the 2022 tightening in Fig. 109.2 therefore comes from a change in the rule itself, not only from the rise in inflation and the gap.

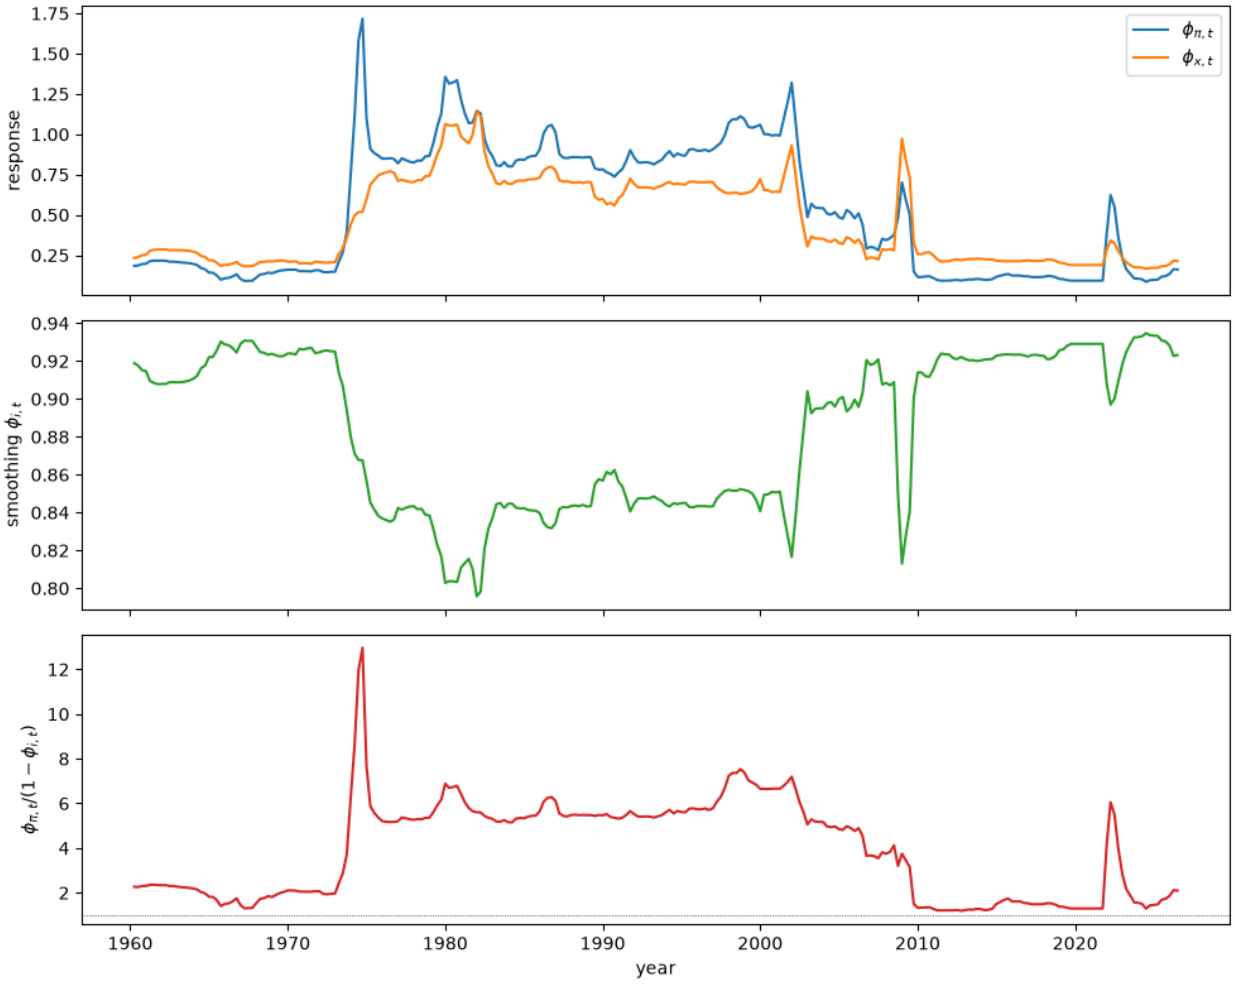


Fig. 109.3: Coefficients of the Phelps rule and its long-run response to inflation

109.5 Why the Fed was slow: a counterfactual

To isolate the role of the drifting beliefs, we recompute the Phelps recommendations holding beliefs fixed at their values in the first quarter of 2000.

```
θ_2000 = beliefs.loc['2000-01-01'].values
print(f"beliefs in 2000Q1: α₀ = {θ_2000[0]:.3f}, ρ = {θ_2000[1]:.3f}, "
      f"κ = {θ_2000[2]:.3f}")
```

```
beliefs in 2000Q1: α₀ = -0.025, ρ = 1.001, κ = 0.097
```

In 2000 the Fed's model had a unit root: $\rho \approx 1$, so inflation shocks were perceived as permanent.

Its perceived slope was more than twice the 2019 value.

```
rules_2000 = pd.DataFrame(np.tile(phelps_rule(θ_2000), (len(beliefs), 1)),
                          index=beliefs.index, columns=rules.columns)
counterfactual = recommended_rate(data, rules_2000)

fig, ax = plt.subplots(figsize=(10, 4.5))
w = slice('2015', None)
ax.plot(optimal[w], 'C0', label='baseline (drifting beliefs)')
ax.plot(counterfactual[w], 'C1--', label='counterfactual (beliefs frozen at 2000)')
ax.plot(data['i'][w], 'C3', lw=1, alpha=0.7, label='actual funds rate')
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()

print(f"counterfactual rate in 2021Q4: {counterfactual['2021-10-01']:.2f}%  "
      f"(actual {data['i']['2021-10-01']:.2f}%)")
```

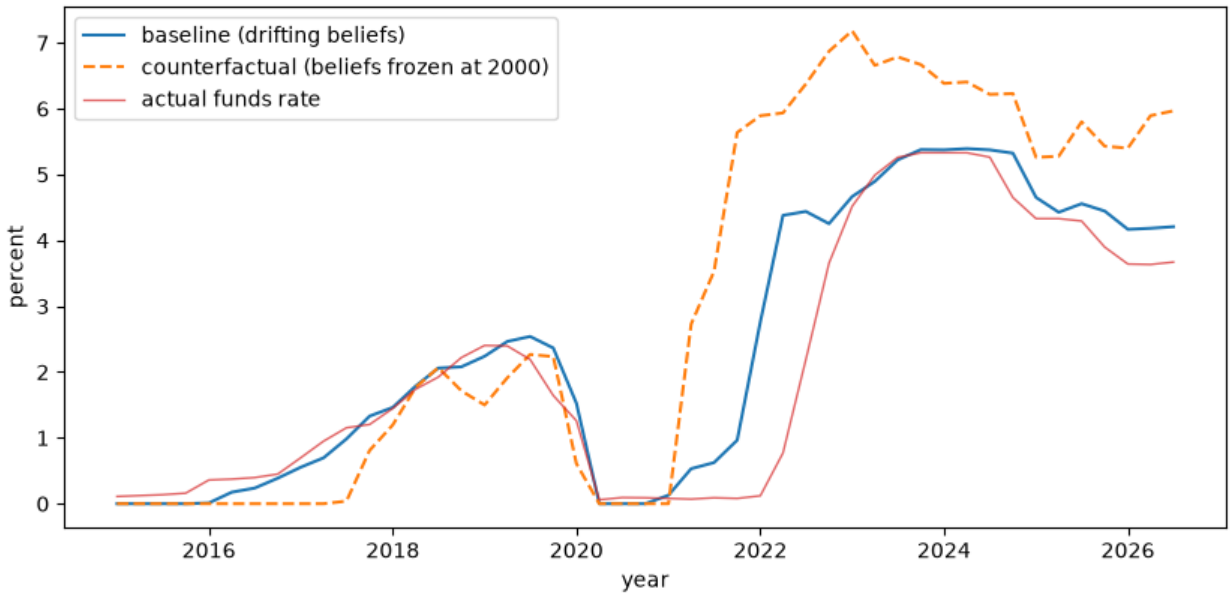


Fig. 109.4: Counterfactual: a Fed that had not updated its beliefs since 2000

```
counterfactual rate in 2021Q4: 5.64%  (actual 0.08%)
```

A Fed with year-2000 beliefs would have tightened immediately and sharply in 2021, driving the funds rate above 5% by the end of 2021, before the actual Fed had moved at all.

The objectives are the same in both paths; only the beliefs differ.

Perceiving inflation as less persistent and the Phillips curve as flatter, the Fed's own Phelps problem told it to respond mildly.

This is the 2020s counterpart of Cogley and Sargent [2005]'s account of the 1970s.

There a Bayesian Fed ran a race among three Phelps problems, and the costly disinflation predicted by the worst-case Samuelson–Solow model made it gradual, even after the data favored a natural-rate model.

109.6 Real-time data and the delay

The paper's conclusion turns on the third element.

In its words, "Only by incorporating real-time uncertainty about the economy, which tends to be elevated during business cycle turning points, do we rationalize the observed delay in the policy response" (Sargent and Williams [2025], section 3.5).

We now rebuild the output gap that the Fed saw in real time.

ALFRED, the St. Louis Fed's archive of past data vintages, keeps every vintage of real GDP (`GDPC1`) and of CBO potential output (`GDPPOT`).

For each quarter since 2019, we compute the gap from the two series as they stood 45 days after the quarter ended, by which time the advance GDP estimate is normally out.

Two kinds of quarters need care.

- Some GDP releases were delayed by the federal government shutdown of late 2025. For any quarter missing from its 45-day vintage, we use the first later vintage that contains it.
- In September 2023 the Bureau of Economic Analysis re-expressed real GDP in 2017 dollars, while CBO's potential output stayed in 2012 dollars until early 2024. A gap computed from such a vintage compares series in different units. We detect this by the ratio of the two series before 2019 and skip vintages in which they are not on the same base.

ALFRED serves up to twelve series per request, so the code asks for six vintages of both series at a time, and the whole history takes only a few requests.

```
ALFRED = 'https://alfred.stlouisfed.org/graph/alfredgraph.csv'

def alfred_vintages(series, dates, per_request=6, tries=3):
    "Vintages of several series at several dates, a few requests in all."
    frames = []
    for j in range(0, len(dates), per_request):
        vd = [f"{d:%Y-%m-%d}" for d in dates[j:j + per_request]]
        url = (f"{ALFRED}?id={','.join(s for s in series for _ in vd)}"
               f"&vintage_date={','.join(vd * len(series))}")
        for attempt in range(tries):
            try:
                frames.append(pd.read_csv(url, index_col=0, parse_dates=True))
                break
            except Exception as err:
```

(continues on next page)

(continued from previous page)

```
                if attempt == tries - 1:
                    raise RuntimeError(
                        "Could not download data vintages from ALFRED "
                        f"(alfred.stlouisfed.org): {err}. This section "
                        "needs internet access to ALFRED.") from err
                time.sleep(5)
    print(f"{len(frames)} requests to ALFRED")
    return pd.concat(frames, axis=1).apply(pd.to_numeric, errors='coerce')

def real_time_gap(quarters, lag_days=45):
    "Output gap for each quarter from the first suitable vintage after it ends."
    # ALFRED relabels a vintage date it treats as still in the future, as it
    # does with today's date for some hours after midnight in St. Louis
    latest = (pd.Timestamp.now(tz='America/Chicago').tz_localize(None).normalize()
              - pd.Timedelta(days=1))
    vintages = sorted({min(q + pd.offsets.QuarterEnd(0)
                           + pd.Timedelta(days=lag_days), latest)
                       for q in quarters})
    V = alfred_vintages(['GDPC1', 'GDPPOT'], vintages)
    base = slice('2015-01-01', '2018-10-01')        # a pre-pandemic window
    gap, late = {}, []
    for q in quarters:
        target = min(q + pd.offsets.QuarterEnd(0) + pd.Timedelta(days=lag_days),
                     latest)
        for v in vintages:
            if v <= q + pd.offsets.QuarterEnd(0):
                continue
            y, y_pot = V[f'GDPC1_{v:%Y%m%d}'], V[f'GDPPOT_{v:%Y%m%d}']
            if pd.isna(y.get(q)) or pd.isna(y_pot.get(q)):
                continue                                # not yet released
            if abs(np.log(y[base] / y_pot[base]).mean()) > 0.04:
                continue                                # different base years
            gap[q] = 100 * (y[q] / y_pot[q] - 1)
            if v != target:
                late.append(f"{qtr(q)} (vintage {v:%Y-%m-%d})")
            break
    print("quarters taken from a later vintage:", late)
    return pd.Series(gap)

x_real_time = real_time_gap(data.loc['2019':].index)
```

```
6 requests to ALFRED
quarters taken from a later vintage: ['2023Q3 (vintage 2024-02-14)', '2025Q3␣
↪(vintage 2026-02-14)', '2025Q4 (vintage 2026-05-15)']
```

Here is the real-time gap against the gap in today's data.

```
x_now = data['x'][x_real_time.index]
fig, ax = plt.subplots(figsize=(10, 4.5))
ax.plot(x_now, 'C0', label="today's data")
ax.plot(x_real_time, 'C1--', marker='o', ms=3, label='real time (ALFRED)')
ax.axhline(0, color='k', lw=0.5, ls=':')
ax.set_xlabel('year')
ax.set_ylabel('percent of potential')
ax.legend()
plt.show()
```

(continues on next page)

(continued from previous page)

```
diff = x_real_time - x_now
below = diff['2020':'2023'] < 0
print(f"real-time gap below today's gap in {below.sum()} of {len(below)} "
      f"quarters of 2020-2023; largest shortfall {-diff.min():.2f} pp "
      f"in {qtr(diff.idxmin())}")
```

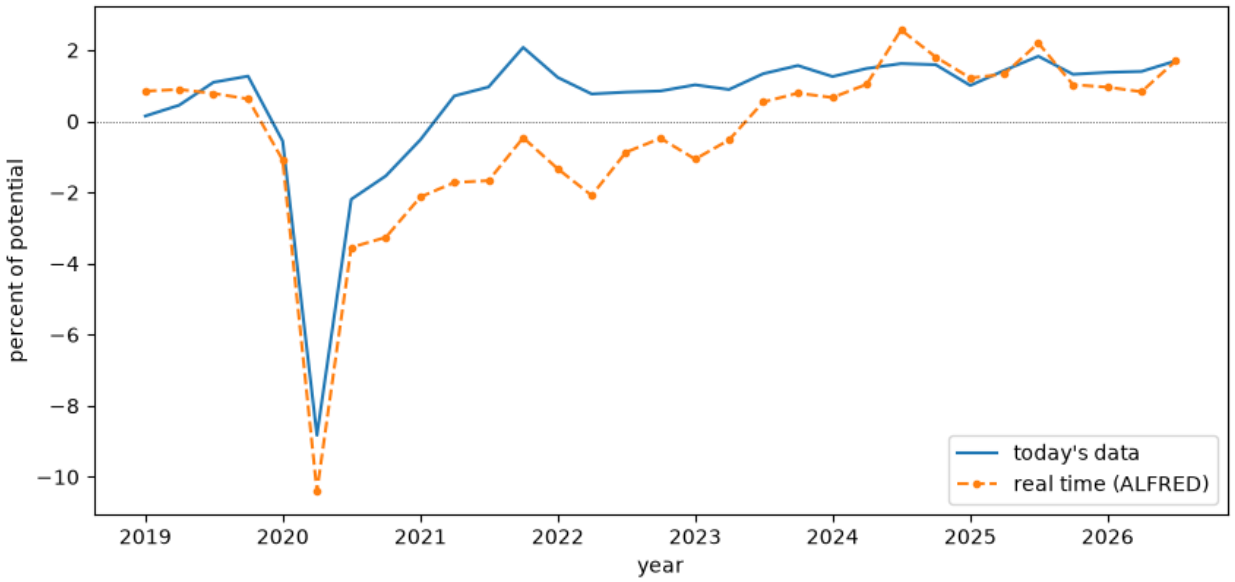


Fig. 109.5: The output gap in real time and in today's data

```
real-time gap below today's gap in 16 of 16 quarters of 2020-2023; largest␣
↪shortfall 2.85 pp in 2022Q2
```

From 2020 through 2023 the real-time gap was below the revised gap in every quarter, by up to about three percentage points.

In real time, the 2021 boom looked like a slow recovery from a deep recession.

We now give the artificial Fed the real-time gap from 2019 on, re-estimate its beliefs, and recompute its recommendations, holding the IS curve fixed.

```
data_rt = data.copy()
data_rt.loc[x_real_time.index, 'x'] = x_real_time.values
beliefs_rt = estimate_beliefs(data_rt)
optimal_rt = recommended_rate(data_rt, phelps_rules(beliefs_rt))

fig, ax = plt.subplots(figsize=(10, 4.5))
w = slice('2019', None)
ax.plot(optimal[w], 'C0', label="Phelps rule, today's data")
ax.plot(optimal_rt[w], 'C1--', label='Phelps rule, real-time gap')
ax.plot(data['i'][w], 'C3', lw=1, label='actual funds rate')
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()
```

(continues on next page)

(continued from previous page)

```
table = pd.DataFrame({"today's data": optimal, 'real time': optimal_rt,
                      'actual': data['i']}).loc['2021':'2023-06']
print(table.rename(index=qtr).round(2).T.to_string())
print("\nmean rate in 2021:", table.loc['2021'].mean().round(2).to_dict())
for name in table:
    liftoff = table.index[table[name] > 0.5][0]
    print(f"first quarter above 0.5%, {name:>12}: {qtr(liftoff)}")
```

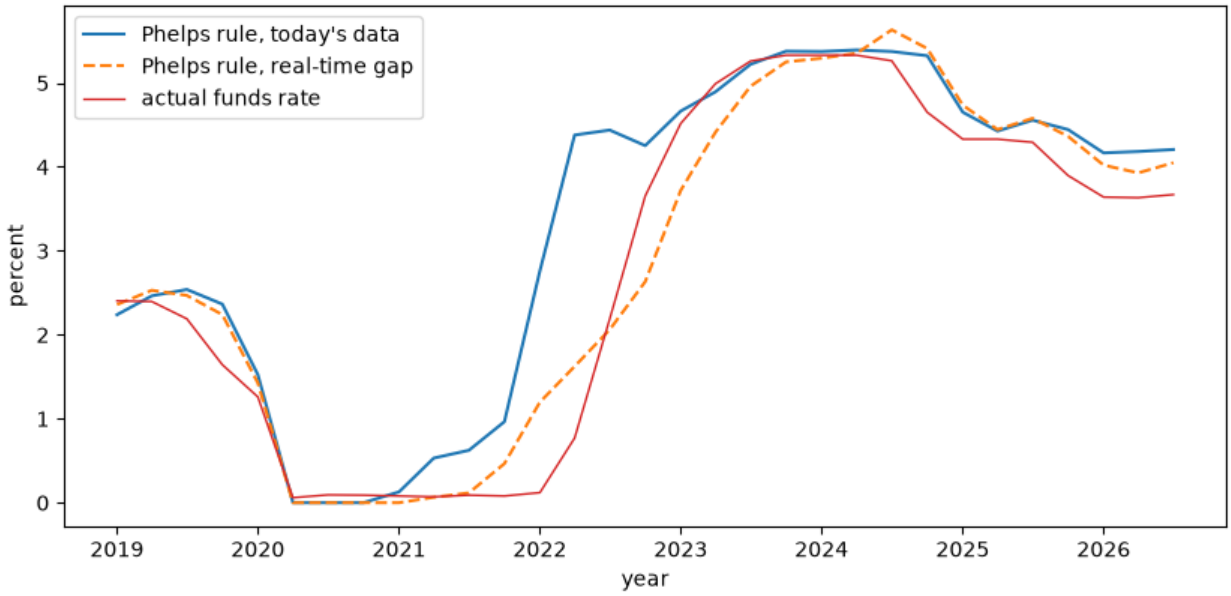


Fig. 109.6: Recommended rates with the revised and the real-time output gap

```
             2021Q1  2021Q2  2021Q3  2021Q4  2022Q1  2022Q2  2022Q3  2022Q4 ␣
 ↪2023Q1  2023Q2
today's data    0.13    0.53    0.62    0.96    2.76    4.38    4.44    4.25    4.
 ↪66    4.89
real time       0.00    0.06    0.12    0.47    1.20    1.62    2.06    2.63    3.
 ↪72    4.41
actual          0.08    0.07    0.09    0.08    0.12    0.77    2.19    3.65    4.
 ↪52    4.99

mean rate in 2021: {"today's data": 0.56, 'real time': 0.16, 'actual': 0.08}
first quarter above 0.5%, today's data: 2021Q2
first quarter above 0.5%,    real time: 2022Q1
first quarter above 0.5%,       actual: 2022Q2
```

With the gap that the Fed saw in real time, the recommended rate stays below half a percent through 2021.

It starts to rise at the turn of 2022, about when the Fed lifted off.

It runs ahead of the actual rate in the first half of 2022 and falls behind it from the third quarter on, as the Fed tightened in large steps.

That is the paper's "slightly lagging" pattern.

With today's data, by contrast, the rule lifts off about a year before the Fed did.

Putting the pieces together:

- the counterfactual shows that **drifting beliefs make the response muted**: with its 2000 beliefs the Fed would have tightened fast and far;
- this comparison shows that **real-time mismeasurement makes it late**: with drifting beliefs but accurate data, the Fed would have started tightening in 2021.

Both are needed to rationalize what the Fed did.

109.7 How beliefs shape the Phelps rule

Why do lower perceived persistence and a flatter perceived slope make the rule less aggressive?

Sargent and Williams [2025] answer analytically, in their Appendix A, for a stripped-down version of the Phelps problem.

The Fed believes that

$$\begin{aligned} \pi_{t+1} &= \rho\,\pi_t + \kappa\,x_{t+1} + \varepsilon^{\pi}_{t+1}, \\ x_{t+1} &= b_1 x_t + g\,(i_t - \pi_t) + \varepsilon^{x}_{t+1}, \end{aligned} \tag{109.5}$$

with $|b_1| < 1$, $\kappa > 0$, and $g < 0$.

These are the paper's equations (6) and (7), whose a_1 and k are our ρ and κ.

The Fed minimizes $\mathbb{E}\sum_t \beta^t(\pi_t^2 + \lambda_x x_t^2)$, with no smoothing penalty, no intercepts, and a zero target.

Taking the real rate $r_t = i_t - \pi_t$ as the control gives a regulator with state (π_t, x_t) whose Riccati equation has a closed-form solution.

The optimal rule is $i_t = \phi_\pi(\rho,\kappa)\,\pi_t + \phi_x x_t$ with

$$\phi_\pi(\rho,\kappa) = 1 - \frac{\rho}{g\kappa}\,\frac{Z}{Z+\lambda_x}, \qquad \phi_x = -\frac{b_1}{g},$$

where

$$Z = \frac{-\Delta + \sqrt{\Delta^2 + 4\lambda_x\kappa^2}}{2}, \qquad \Delta = \lambda_x(1-\beta\rho^2) - \kappa^2.$$

The paper proves two comparative statics.

Proposition 109.7.1 (Lower perceived persistence ⇒ less aggressive policy)

For any $\beta \in (0,1)$, $\lambda_x > 0$, $\kappa > 0$, and $g < 0$, the inflation coefficient $\phi_\pi(\rho,\kappa)$ is strictly increasing in $\rho \geq 0$.

Proposition 109.7.2 (Flatter perceived slope ⇒ less aggressive policy)

$\partial\phi_\pi/\partial\kappa > 0$ whenever a sufficient condition (the paper's inequality (11)) holds; it holds for small and moderate κ.

Let's verify both numerically, and check the closed form against `solve_discrete_are` along the way.

For the IS beliefs we use illustrative values, $b_1 = 0.9$ and $g = -0.5$.

b_1 is close to the estimate above, but g is much larger in absolute value, so that ϕ_π is of a familiar size.

Since $\phi_\pi - 1$ is proportional to $1/g$, the value of g does not affect the signs of the comparative statics.

```
def ψ_π(ρ, κ, b1=0.9, g=-0.5, λ_x=0.2, β=0.95):
    "Closed-form inflation coefficient of the simplified Phelps rule."
    Δ = λ_x * (1 - β * ρ**2) - κ**2
    Z = (-Δ + np.sqrt(Δ**2 + 4 * λ_x * κ**2)) / 2
    return 1 - (ρ / (g * κ)) * Z / (Z + λ_x)

def ψ_π_riccati(ρ, κ, b1=0.9, g=-0.5, λ_x=0.2, β=0.95, η_r=1e-9):
    "The same coefficient from solve_discrete_are, with control cost η_r r²."
    A = np.array([[ρ, κ * b1], [0, b1]])
    B = np.array([[κ * g], [g]])
    Q, R = np.diag([1, λ_x]), np.array([[η_r]])
    P = solve_discrete_are(np.sqrt(β) * A, np.sqrt(β) * B, Q, R)
    F = np.linalg.solve(R + β * B.T @ P @ B, β * B.T @ P @ A)
    return 1 - F[0, 0]

grid = [(ρ_, κ_) for ρ_ in np.linspace(0.05, 0.99, 12)
        for κ_ in np.linspace(0.01, 1.0, 12)]
gap_cf = max(abs(ψ_π(*p) - ψ_π_riccati(*p)) / ψ_π(*p) for p in grid)
print(f"largest relative difference, closed form vs Riccati: {gap_cf:.1e}")
```

```
largest relative difference, closed form vs Riccati: 1.5e-08
```

The closed form checks out.

Now the two comparative statics.

```
ρ_grid = np.linspace(0, 0.999, 200)
κ_grid = np.linspace(0.005, 1.0, 400)

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
for κ_ in [0.04, 0.1, 0.3]:
    axes[0].plot(ρ_grid, ψ_π(ρ_grid, κ_), label=rf'$\kappa = {κ_}$')
axes[0].set_xlabel(r'perceived persistence $\rho$')
axes[0].set_ylabel(r'$\phi_\pi$')
axes[0].legend()

for ρ_ in [0.5, 0.86, 0.99]:
    axes[1].plot(κ_grid, ψ_π(ρ_, κ_grid), label=rf'$\rho = {ρ_}$')
axes[1].set_xlabel(r'perceived slope $\kappa$')
axes[1].legend()
plt.tight_layout()
plt.show()

increasing = all(np.all(np.diff(ψ_π(ρ_grid, κ_)) > 0)
                 for κ_ in [0.01, 0.04, 0.1, 0.3, 1.0, 3.0])
print("ϕ_π increasing in ρ for every κ tried:", increasing)
for ρ_ in [0.5, 0.86, 0.99]:
    κ_peak = κ_grid[np.argmax(ψ_π(ρ_, κ_grid))]
    print(f"ρ = {ρ_}: ϕ_π increases in κ up to κ = {κ_peak:.3f}, then falls")
print(f"largest κ_t estimated in 2010-2019: {κ['2010':'2019'].max():.3f}")
```

```
ϕ_π increasing in ρ for every κ tried: True
ρ = 0.5: ϕ_π increases in κ up to κ = 0.392, then falls
ρ = 0.86: ϕ_π increases in κ up to κ = 0.244, then falls
ρ = 0.99: ϕ_π increases in κ up to κ = 0.117, then falls
largest κ_t estimated in 2010-2019: 0.094
```

Persistence matters everywhere.

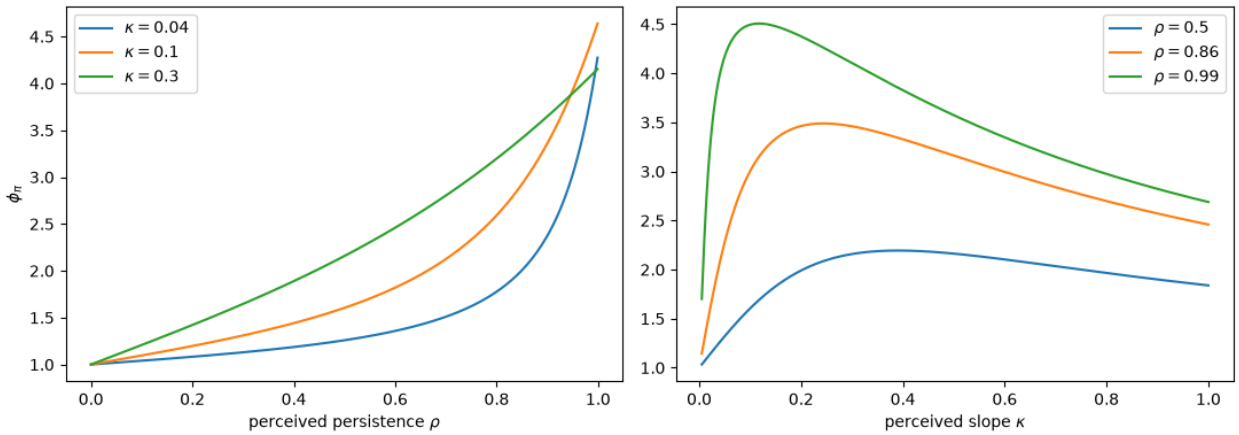


Fig. 109.7: The inflation coefficient of the simplified Phelps rule

The higher ρ, the more today's inflation predicts future inflation, and the harder the Fed leans against it.

The slope matters too: a flatter curve makes each point of disinflation cost more output, so the Fed responds less.

That holds for small slopes, including all those the Fed estimated in the 2010s.

For large slopes the effect reverses, and the peak moves to smaller κ as ρ rises (Exercise 109.10.2).

So *Proposition 109.7.2* should be read as a statement about small slopes.

Together, the two propositions explain the drift of the rule in Fig. 109.3: as ρ_t and κ_t fell, so did the rule's response to inflation.

109.8 From policy to measured beliefs

So far we have taken the Fed's beliefs from the data.

Where did they come from?

In the extended abstract of their project, Sargent and Williams [2025] suggest that "the perceptions of the low persistence of inflation and a flat Phillips curve were based on the past Fed policy, which had reacted strongly to shocks and stabilized inflation."

They frame the outcome as a self-confirming equilibrium in which the Fed is caught in a lack-of-experimentation trap.

This section is the lecture's own extension of that idea; the January 2026 version of the paper does not contain it.

We put a small New Keynesian model on the table as *nature's* model.

We compute what the Fed's regression would measure under different policies, and then close the loop with the Phelps rule of the previous section.

In the terms of Koopmans [1947], the Fed's drifting-coefficients model is a descriptive "Kepler stage" model.

Nature's model, by contrast, is a "Newton stage" model whose parameters are invariant to the Fed's rule.

109.8.1 Nature's model

Suppose that nature runs

$$\begin{aligned} \pi_t &= \beta_f \, \mathbb{E}_t \pi_{t+1} + \omega_b \, \pi_{t-1} + \kappa \, x_t + u_t, \\ x_t &= \mathbb{E}_t x_{t+1} - \sigma \left(i_t - \mathbb{E}_t \pi_{t+1} - r_t \right), \\ i_t &= \phi \, \pi_t, \end{aligned} \qquad \beta_f = \frac{\beta_n}{1 + \beta_n \omega}, \quad \omega_b = \frac{\omega}{1 + \beta_n \omega}, \tag{109.6}$$

where u_t is a cost-push shock and r_t a natural-rate (demand) shock.

The two shocks are IID and mutually uncorrelated, with standard deviations σ_u and σ_r.

The first equation is the hybrid Phillips curve that arises when firms that cannot reset their prices index them to lagged inflation at rate $\omega \in [0, 1]$.

Its weights on expected and lagged inflation sum to $(\beta_n + \omega)/(1 + \beta_n \omega) \leq 1$.

With this normalization the Taylor principle $\phi > 1$ delivers a determinate equilibrium, which the code below checks by counting roots (*Stability in Linear Rational Expectations Models*; [Blanchard and Kahn, 1980]).

Nature's discount factor $\beta_n = 0.99$ is distinct from the Fed's $\beta = 0.95$.

We set $\omega = 1$, a structural slope $\kappa = 0.1$, and $\sigma = 1$.

In the minimum-state-variable equilibrium, $\mathbb{E}_t \pi_{t+1} = \lambda \, \pi_t$, and

$$\pi_t = \lambda \, \pi_{t-1} + a \, u_t + b \, r_t, \qquad x_t = \psi \, \pi_{t-1} + c \, u_t + d \, r_t.$$

Here λ is the stable root of the cubic

$$\beta_f \lambda^3 - (1 + \beta_f + \kappa\sigma)\lambda^2 + (1 + \omega_b + \kappa\sigma\phi)\lambda - \omega_b = 0,$$

$\psi = -\sigma\lambda(\phi - \lambda)/(1 - \lambda)$, and the impact coefficients (a, b, c, d) are computed in the code.

λ is the coefficient of a univariate AR(1): it is what an econometrician who regressed inflation on its own lag would call persistence.

109.8.2 What the Fed measures

The Fed's regression (109.1) also contains x_t.

Since $x_t = \psi \, \pi_{t-1} + v_t$, with $v_t = c \, u_t + d \, r_t$ uncorrelated with π_{t-1}, its population coefficients are

$$\hat{\kappa} = \frac{\operatorname{cov}(\pi_t, v_t)}{\mathbb{V}[v_t]}, \qquad \hat{\rho} = \lambda - \hat{\kappa} \, \psi.$$

So the Fed's measured persistence $\hat{\rho}$ is *not* λ.

When only demand shocks hit ($\sigma_u = 0$), we can substitute $\mathbb{E}_t \pi_{t+1} = \lambda \pi_t$ into the Phillips curve.

That gives the exact relation $\pi_t = [\omega_b \, \pi_{t-1} + \kappa \, x_t]/(1 - \beta_f \lambda)$, so

$$\hat{\rho} = \frac{\omega_b}{1 - \beta_f \, \lambda(\phi)}, \qquad \hat{\kappa} = \frac{\kappa}{1 - \beta_f \, \lambda(\phi)}.$$

The Fed recovers the structural lag weight and slope, each divided by $1 - \beta_f \lambda(\phi)$, a factor that depends on policy.

That is the Lucas critique in one line.

A more aggressive rule lowers λ, and with it both $\hat{\rho}$ and $\hat{\kappa}$: aggressive policy makes inflation look less persistent and the Phillips curve look flatter.

The persistence result is in the spirit of Benati and Surico [2008] and Carlstrom *et al.* [2009]; the slope result in this form is Bullard [2018]'s.

Cost-push shocks push $\hat{\kappa}$ down further and can make it negative.

That is the identification problem stressed by McLeay and Tenreyro [2019]: when the Fed offsets cost-push shocks, high inflation coincides with a low output gap.

```
def nk_weights(ω, β_n=0.99):
    "Weights (β_f, ω_b) of the normalized hybrid Phillips curve."
    return β_n / (1 + β_n * ω), ω / (1 + β_n * ω)

def nk_roots(ψ, ω=1.0, κ=0.1, σ=1.0, β_n=0.99):
    "Roots of the cubic whose stable root is λ."
    β_f, ω_b = nk_weights(ω, β_n)
    return np.roots([β_f, -(1 + β_f + κ * σ), 1 + ω_b + κ * σ * ψ, -ω_b])

def fed_regression(ψ, ω=1.0, κ=0.1, σ=1.0, β_n=0.99, σ_u=0.0, σ_r=1.0):
    "λ, and the Fed's population estimates (ρ̂, κ̂), when nature's rule is ψ."
    β_f, ω_b = nk_weights(ω, β_n)
    roots = nk_roots(ψ, ω, κ, σ, β_n)
    λ = roots[np.abs(roots) < 1].real.item()        # determinacy: exactly one
    m = -σ * (ψ - λ) / (1 - λ)                       # x_t = m π_t + σ r_t
    ψ = m * λ
    den = 1 - β_f * λ - κ * m
    a, b = 1 / den, κ * σ / den                      # π_t on (u_t, r_t)
    c, d = m * a, m * b + σ                          # x_t on (u_t, r_t)
    κ_hat = ((a * c * σ_u**2 + b * d * σ_r**2)
             / (c**2 * σ_u**2 + d**2 * σ_r**2))
    return λ, λ - κ_hat * ψ, κ_hat

ψ_grid = np.linspace(1.01, 5.0, 200)
determinate = all(np.sum(np.abs(nk_roots(ψ, ω)) < 1) == 1
                  for ψ in np.r_[ψ_grid, 10.0, 100.0] for ω in [0.25, 1.0])
print("one stable root for every ψ > 1 tried:", determinate)
```

```
one stable root for every ψ > 1 tried: True
```

Here are λ, $\hat{\rho}$, and $\hat{\kappa}$ as functions of nature's Taylor-rule coefficient ϕ.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))
λ_path = [fed_regression(ψ)[0] for ψ in ψ_grid]
axes[0].plot(ψ_grid, λ_path, 'k:', label=r'$\lambda$ (AR(1) root)')
for σ_u in [0.0, 0.1]:
    est = np.array([fed_regression(ψ, σ_u=σ_u)[1:] for ψ in ψ_grid])
    axes[0].plot(ψ_grid, est[:, 0], label=rf'$\hat\rho$, $\sigma_u/\sigma_r = {σ_u}$')
    axes[1].plot(ψ_grid, est[:, 1], label=rf'$\sigma_u/\sigma_r = {σ_u}$')
est = np.array([fed_regression(ψ, σ_u=0.3)[1:] for ψ in ψ_grid])
axes[1].plot(ψ_grid, est[:, 1], label=r'$\sigma_u/\sigma_r = 0.3$')
axes[1].axhline(0, color='k', lw=0.5, ls=':')
axes[0].set_xlabel(r"nature's Taylor-rule coefficient $\phi$")
axes[0].set_ylabel('measured persistence')
axes[1].set_xlabel(r"nature's Taylor-rule coefficient $\phi$")
axes[1].set_ylabel(r'measured slope $\hat\kappa$')
axes[0].legend()
axes[1].legend()
plt.tight_layout()
plt.show()
```

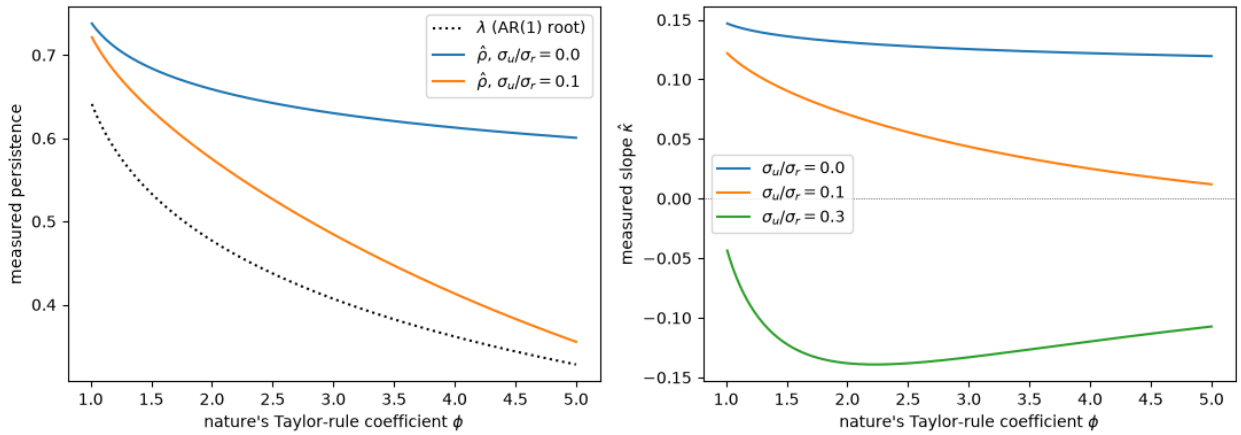


Fig. 109.8: What the Fed's regression measures under different policy rules

The measured persistence $\hat{\rho}$ lies well above the AR(1) root λ, and all three fall as policy grows more aggressive.

With demand shocks only, $\hat{\kappa}$ also falls with ϕ but stays above the structural slope $\kappa = 0.1$, as the formula says.

Even modest cost-push shocks pull it down sharply, and below zero.

The model is a stylized quarterly one, so its numbers are not meant to match the Fed's estimates from four-quarter inflation.

109.8.3 The self-confirming fixed point

Now compose the two maps.

A history of policy ϕ produces measured beliefs $(\hat{\rho}(\phi), \hat{\kappa}(\phi))$.

Fed into the Phelps problem, those beliefs recommend the inflation coefficient

$$T(\phi) = \phi_\pi(\hat{\rho}(\phi), \hat{\kappa}(\phi)).$$

A fixed point $\phi^* = T(\phi^*)$ is a self-confirming equilibrium in the sense of *Self-Confirming Equilibria*.

At ϕ^*, the Fed's rule solves its Phelps problem given its beliefs, and its beliefs satisfy the least-squares orthogonality conditions in the data that its rule generates.

We make two simplifications.

Nature's rule uses only the inflation coefficient that the Phelps problem recommends.

We also hold the Fed's IS beliefs at the illustrative values of the previous section, so the loop runs only through the Phillips curve.

Since $\hat{\rho}$ and $\hat{\kappa}$ fall as ϕ rises, and ϕ_π rises with both (for small κ), T is *decreasing*: the loop is a negative feedback.

```
def T(ψ, **nature):
    "Inflation coefficient recommended by beliefs measured under the rule ψ."
    _, ρ_hat, κ_hat = fed_regression(ψ, **nature)
    return ψ_π(ρ_hat, κ_hat)

def fixed_point(**nature):
    "ψ* = T(ψ*) and the slope T'(ψ*)."
```

(continues on next page)

(continued from previous page)

```
    φ_star = brentq(lambda φ: T(φ, **nature) - φ, 1.01, 5.0)
    h = 1e-5
    slope = (T(φ_star + h, **nature) - T(φ_star - h, **nature)) / (2 * h)
    return φ_star, slope

fig, ax = plt.subplots(figsize=(7, 4.5))
for σ_u in [0.0, 0.1]:
    ax.plot(φ_grid, [T(φ, σ_u=σ_u) for φ in φ_grid],
            label=rf'$T(\phi)$, $\sigma_u/\sigma_r = {σ_u}$')
ax.plot(φ_grid, φ_grid, 'k:', lw=1, label='45° line')
ax.set_xlabel(r'past policy $\phi$')
ax.set_ylabel(r'recommended $\phi_\pi$')
ax.legend()
plt.show()

cases = {'baseline': {}, 'σ_u/σ_r = 0.1': {'σ_u': 0.1},
         'ω = 0.5': {'ω': 0.5}, 'κ = 0.05': {'κ': 0.05}}
for name, nature in cases.items():
    φ_star, slope = fixed_point(**nature)
    _, ρ_hat, κ_hat = fed_regression(φ_star, **nature)
    print(f"{name:>15}: φ* = {φ_star:.3f}, T'(φ*) = {slope:.3f}, "
          f"beliefs (ρ̂, κ̂) = ({ρ_hat:.3f}, {κ_hat:.3f})")
```

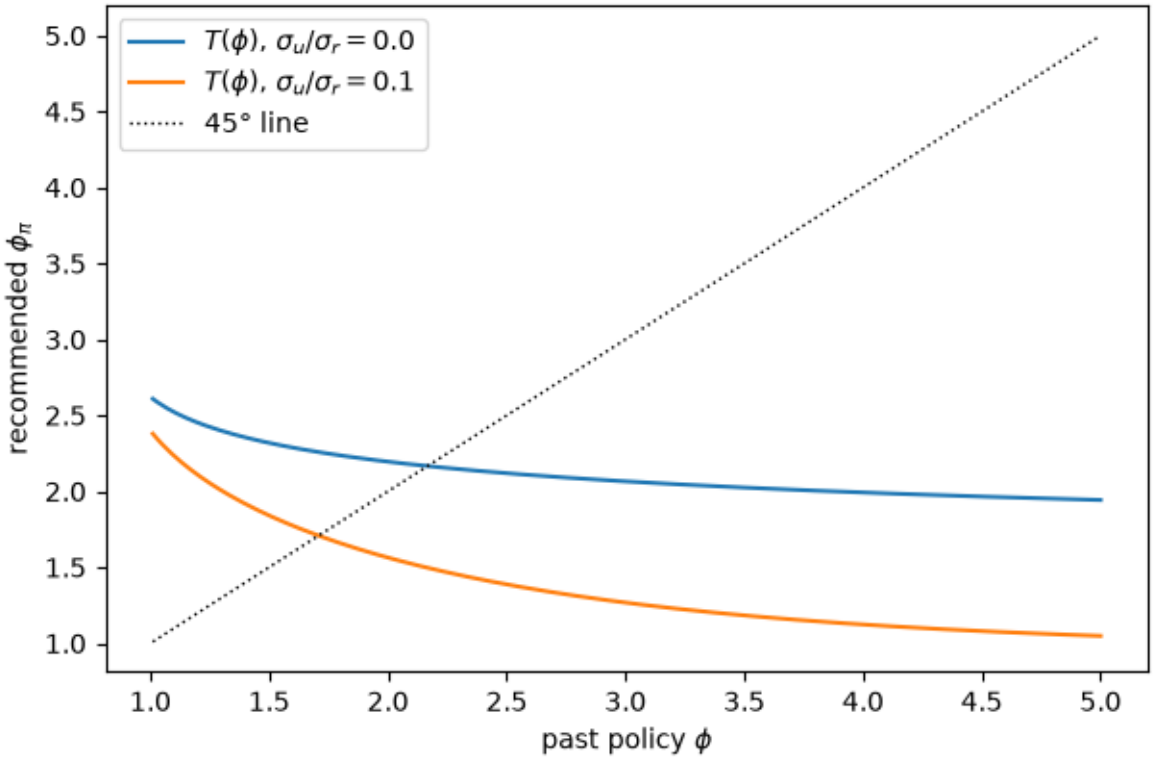


Fig. 109.9: The map from past policy to the policy the Phelps problem recommends

```
      baseline: φ* = 2.166, T'(φ*) = -0.161, beliefs (ρ̂, κ̂) = (0.653, 0.130)
σ_u/σ_r = 0.1: φ* = 1.708, T'(φ*) = -0.568, beliefs (ρ̂, κ̂) = (0.607, 0.082)
       ω = 0.5: φ* = 1.612, T'(φ*) = -0.076, beliefs (ρ̂, κ̂) = (0.434, 0.130)
```

(continues on next page)

(continued from previous page)

```
κ = 0.05: ψ* = 1.871, T'(ψ*) = -0.199, beliefs (ρ̂, κ̂) = (0.707, 0.070)
```

T crosses the 45-degree line once, so the self-confirming equilibrium is unique.

Is it stable under learning?

A Fed that learns with a small gain follows mean dynamics like $d\phi/d\tau = T(\phi) - \phi$, the ordinary differential equation of *One recursion, many learners*.

The rest point is E-stable if $T'(\phi^*) < 1$.

In belief space the same condition appears.

The map $(\rho, \kappa) \mapsto (\hat{\rho}, \hat{\kappa})(\phi_\pi(\rho, \kappa))$ has a rank-one Jacobian whose nonzero eigenvalue is $T'(\phi^*)$.

So the E-stability eigenvalues are -1 and $T'(\phi^*) - 1$, the kind of calculation that *Self-Confirming Equilibria* carries out for the *Conquest*'s model (*Uniqueness and E-stability*).

Because $T' < 0$, the condition holds with room to spare in every case above.

Because $|T'| < 1$ as well, even the gain-one "cobweb" iteration $\phi_{k+1} = T(\phi_k)$ converges.

It describes a Fed that re-optimizes only after its beliefs have fully adjusted to each policy regime:

```
for name in ['baseline', 'σ_u/σ_r = 0.1']:
    ψ_k, path = 3.0, [3.0]                    # start from an aggressive history
    for _ in range(8):
        ψ_k = T(ψ_k, **cases[name])
        path.append(ψ_k)
    print(f"{name:>15}:", np.round(path, 2))
```

```
       baseline: [3.   2.07 2.18 2.16 2.17 2.17 2.17 2.17 2.17]
 σ_u/σ_r = 0.1: [3.   1.27 2.04 1.55 1.81 1.65 1.74 1.69 1.72]
```

An aggressive history is followed by a timid rule, then by a partial correction, in damped oscillations around ϕ^*.

That is a faint echo of "conquest lost and regained", but it dies out.

109.8.4 A trap without an escape route

The fixed point has the structure of the *Conquest*'s self-confirming equilibria.

The anticipated-utility Fed never experiments, by construction (*Adaptive Learning and Escape Dynamics*).

At ϕ^* its beliefs are right on the equilibrium path and wrong off it: under a less aggressive rule, measured persistence and slope would rise.

This is the lack-of-experimentation trap of *Escaping Nash Inflation* (*The experimentation trap*), and it is the sense of the extended abstract's framing.

The difference lies in the dynamics.

In the model of Cho *et al.* [2002], the escape route is self-reinforcing.

A run of shocks that steepens the perceived Phillips curve makes the government cut inflation, and that generates data that steepen the curve further.

Here the feedback runs the other way.

A run of shocks that raises measured persistence makes the Fed more aggressive, and more aggressive policy *lowers* measured persistence.

A constant-gain learner would wander around ϕ^*, but the mean dynamics pull it back from either side: there is no escape route.

Two facts also limit how much this loop can explain about the data.

First, in Fig. 109.1 the fall in ρ_t came after 2008, when the zero lower bound kept policy from being aggressive at all.

Second, the jump in ρ_t in 2022, alongside the inflation peak, is in line with the positive correlation between the level and the persistence of inflation that *Drifts and Volatilities* documents.

Escape-route models predict an *inverse* correlation; Cogley and Sargent [2005] (footnote 28) note that gradualism of the kind described above can produce the positive one.

So the 2020s were not a replay of the *Conquest*'s recurrent escapes.

What carries over is the *Conquest*'s diagnosis.

The Fed treated a reduced form whose coefficients depend on policy as if it were structural, ignoring the Lucas critique in just the way the vindication story of *The Rise and Fall of U.S. Inflation* describes (Sargent and Williams [2025], footnote 6).

But the misspecification now concerns the slope and persistence of the Phillips curve, not the natural-rate hypothesis.

> **Note**
>
> Sargent and Williams [2025] also acknowledge a very different reading in which the accommodation of the 2020s was fiscal in origin, as in the fiscal-theory accounts they cite.
>
> That would be another rationalization of the same policy path.

109.9 Was the conquest lost?

The Rise and Fall of U.S. Inflation ended by asking what the data since 1999 say about its two stories.

Would a credible low-inflation equilibrium keep re-anchoring after each shock, as the triumph of natural-rate theory suggests?

Or could surprises still set a monetary authority's beliefs drifting, as the vindication of econometric policy evaluation suggests?

Following Sargent and Williams [2025], this lecture's answer is that the conquest was lost for a while.

The reason was not a change in the Fed's objectives.

- The artificial Fed's loss function never changes.
- Its beliefs drifted during a long period of low and stable inflation, so that by 2019 its Phelps problem prescribed a long-run response to inflation only a little above one; that made its response to the surge *muted*.
- Real-time measures of slack, which overstated how much room the economy had, made it *late*.
- Once the 2022 data arrived, perceived persistence jumped back toward one, the rule's response to inflation jumped with it, and the conquest was regained.

In the language of *The Rise and Fall of U.S. Inflation*, the episode looks like the vindication story's adaptive policy maker, misled by a reduced form whose coefficients depend on policy.

But it is not a replay of that story's recurrent escapes.

The policy-to-beliefs loop we built has no escape route; drifting beliefs explain why the response was muted, and mismeasured data explain why it was late.

Nor does anything here contradict the triumph story's reading, recorded in *The Rise and Fall of U.S. Inflation*, that long-run inflation expectations stayed anchored: the artificial Fed's beliefs concern the reduced-form dynamics of inflation, not the public's expectations.

The lesson the episode adds to the *Conquest* is the one the extended abstract draws: the lessons of past successful policy can fade during long tranquil periods, and then have to be learned again.

109.10 Exercises

Exercise 109.10.1

The interest-smoothing weight η in the loss (109.3) is, in the paper's words, "one of the most important parameters."

The paper says that a larger η fits the observed rate better but makes the rule harder to interpret as systematic stabilization policy.

Make that tradeoff quantitative.

For $\eta \in \{0.1, 0.5, 2.0\}$:

1. compute the recommended rates and their root mean squared error (RMSE) against the actual funds rate from 1991 on, and compare with the no-change rule $i_t = i_{t-1}$;
2. compute the 2019Q4 rule and its long-run response to inflation, $\phi_\pi/(1-\phi_i)$;
3. plot the recommendations against the actual rate from 2015 on.

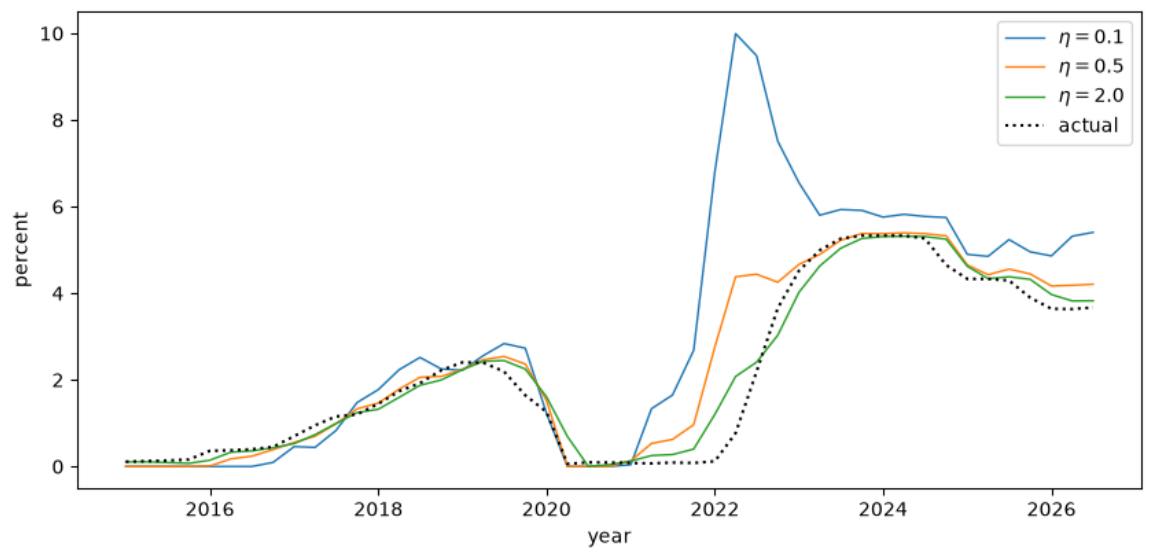
Which η would you choose, and why?

Solution

The beliefs do not depend on η, so we reuse them.

```
fig, ax = plt.subplots(figsize=(10, 4.5))
w = slice('2015', None)
print(f"{'η':>5} {'RMSE':>6} {'φ_π':>7} {'φ_i':>7} {'long run':>9}")
for η_val in [0.1, 0.5, 2.0]:
    rules_η = phelps_rules(beliefs, η=η_val)
    rate = recommended_rate(data, rules_η)
    f = rules_η.loc['2019-10-01']
    print(f"{η_val:>5} {fit(rate)[1]:>6.2f} {f.ψ_π:>7.3f} {f.ψ_i:>7.3f} "
          f"{f.ψ_π / (1 - f.ψ_i):>9.2f}")
    ax.plot(rate[w], lw=1, label=rf'$\eta = {η_val}$')
print(f"no-change rule RMSE: {fit(no_change)[1]:.2f}")
ax.plot(data['i'][w], 'k:', lw=1.5, label='actual')
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()
```

```
    η   RMSE     φ_π     φ_i  long run
  0.1   1.76   0.297   0.862      2.15
  0.5   0.79   0.092   0.929      1.30
  2.0   0.47   0.030   0.966      0.88
no-change rule RMSE: 0.45
```



A small η produces a volatile rate that jumps sharply with inflation and the gap, and fits badly.

Raising η improves the fit a lot, but at $\eta = 2$ the RMSE is close to that of the no-change rule.

Meanwhile the 2019 rule's long-run response to inflation falls below one, so the rule no longer satisfies the Taylor principle.

A large η buys fit by turning the Phelps rule into a no-change rule.

That is the concrete meaning of the paper's "interpretability."

$\eta = 0.5$ keeps a long-run response above one while fitting far better than $\eta = 0.1$.

Exercise 109.10.2

This exercise explores the simplified Phelps problem (109.5).

1. Show that $P_{12} = 0$ and $P_{22} = \lambda_x$ solve its Riccati equation. (Hint: the second column of A is a multiple of B.) Then derive the quadratic equation for P_{11}, and confirm both results with `solve_discrete_are`.
2. For $\rho \in \{0.5, 0.9\}$, plot ϕ_π against $\kappa \in (0, 3]$ and find the slope at which ϕ_π peaks. Explain why ϕ_π eventually falls as κ grows.
3. Add a penalty $\eta_r r_t^2$ on the real rate to the loss. Does the reversal survive for $\eta_r \in \{0.01, 0.1\}$?

Solution

Part 1. With r_t as the control,

$$A = \begin{bmatrix} \rho & \kappa b_1 \\ 0 & b_1 \end{bmatrix}, \qquad B = g \begin{bmatrix} \kappa \\ 1 \end{bmatrix},$$

so the second column of A equals $(b_1/g)B$.

With no control cost, the Riccati equation is

$$P = Q + \beta A^\top (P - PB(B^\top PB)^{-1} B^\top P) A.$$

The matrix in parentheses annihilates B, so the second row and column of $A^\top(\cdot)A$ vanish.

Hence P agrees with $Q = \mathrm{diag}(1, \lambda_x)$ in its second row and column: $P_{12} = 0$ and $P_{22} = \lambda_x$.

Economically, the Fed can offset the predictable part of next quarter's gap completely, so the current gap carries no cost into the future.

Using $B^\top P B = g^2(\kappa^2 P_{11} + \lambda_x)$ and $(PB)_1 = g\kappa P_{11}$, the $(1,1)$ element gives

$$P_{11} = 1 + \beta\rho^2 \frac{\lambda_x P_{11}}{\kappa^2 P_{11} + \lambda_x}, \quad \text{i.e.} \quad \kappa^2 P_{11}^2 + (\lambda_x(1 - \beta\rho^2) - \kappa^2)P_{11} - \lambda_x = 0,$$

whose positive root is the paper's equation (8); $Z = \kappa^2 P_{11}$.

```
def riccati_P(ρ, κ, b1=0.9, g=-0.5, λ_x=0.2, β=0.95, η_r=1e-9):
    A = np.array([[ρ, κ * b1], [0, b1]])
    B = np.array([[κ * g], [g]])
    return solve_discrete_are(np.sqrt(β) * A, np.sqrt(β) * B,
                              np.diag([1, λ_x]), np.array([[η_r]]))

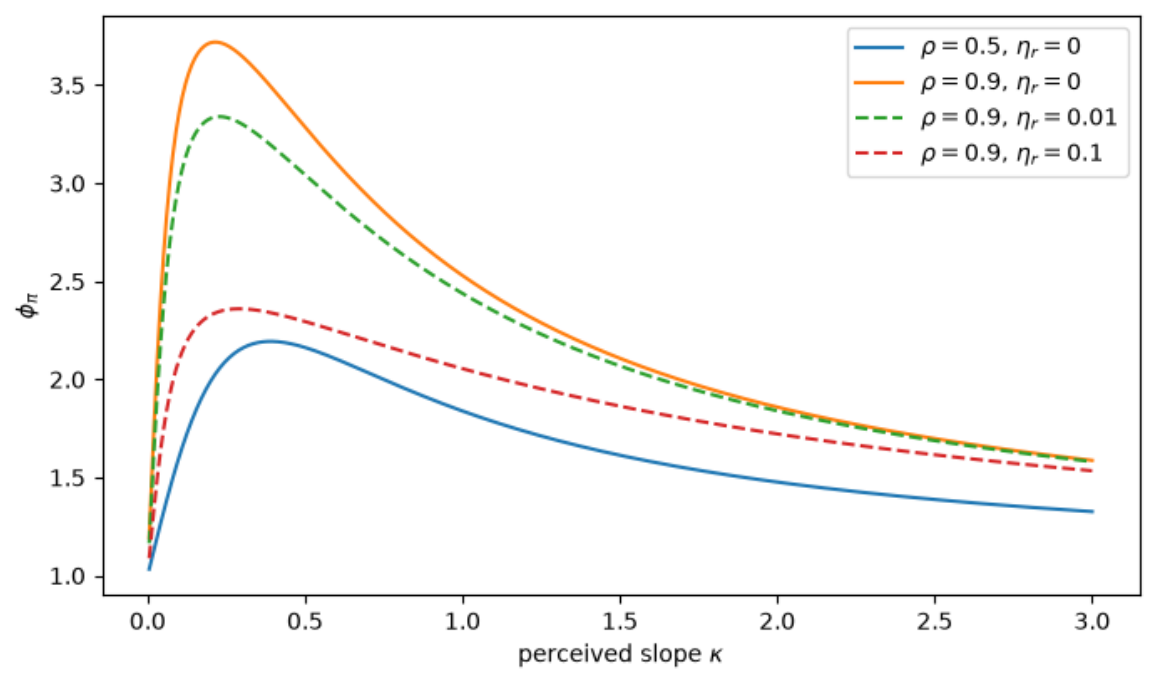
for ρ_, κ_ in [(0.9, 0.05), (0.5, 0.2), (0.95, 1.0)]:
    P = riccati_P(ρ_, κ_)
    Δ = 0.2 * (1 - 0.95 * ρ_**2) - κ_**2
    P11 = (-Δ + np.sqrt(Δ**2 + 4 * 0.2 * κ_**2)) / (2 * κ_**2)
    print(f"ρ = {ρ_}, κ = {κ_}: P12 = {P[0, 1]:.1e}, P22 = {P[1, 1]:.4f}, "
          f"P11 = {P[0, 0]:.4f} (closed form {P11:.4f})")
```

```
ρ = 0.9, κ = 0.05: P12 = 2.9e-09, P22 = 0.2000, P11 = 3.7715 (closed form 3.
↪7715)
ρ = 0.5, κ = 0.2: P12 = 1.8e-09, P22 = 0.2000, P11 = 1.2353 (closed form 1.2353)
ρ = 0.95, κ = 1.0: P12 = 2.9e-09, P22 = 0.2000, P11 = 1.1460 (closed form 1.
↪1460)
```

Parts 2 and 3.

```
κ_wide = np.linspace(0.005, 3.0, 600)
fig, ax = plt.subplots(figsize=(8, 4.5))
for ρ_ in [0.5, 0.9]:
    curve = ψ_π(ρ_, κ_wide)
    ax.plot(κ_wide, curve, label=rf'$\rho = {ρ_}$, $\eta_r = 0$')
    print(f"ρ = {ρ_}, η_r = 0: peak at κ = {κ_wide[np.argmax(curve)]:.3f}")
for η_r in [0.01, 0.1]:
    curve = [ψ_π_riccati(0.9, κ_, η_r=η_r) for κ_ in κ_wide]
    ax.plot(κ_wide, curve, '--', label=rf'$\rho = 0.9$, $\eta_r = {η_r}$')
    print(f"ρ = 0.9, η_r = {η_r}: peak at κ = {κ_wide[np.argmax(curve)]:.3f}, "
          f"monotone in κ: {np.all(np.diff(curve) > 0)}")
ax.set_xlabel(r'perceived slope $\kappa$')
ax.set_ylabel(r'$\phi_\pi$')
ax.legend()
plt.show()
```

```
ρ = 0.5, η_r = 0: peak at κ = 0.390
ρ = 0.9, η_r = 0: peak at κ = 0.215
ρ = 0.9, η_r = 0.01: peak at κ = 0.230, monotone in κ: False
ρ = 0.9, η_r = 0.1: peak at κ = 0.290, monotone in κ: False
```



ϕ_π is hump-shaped in κ.

When the curve is very flat, moving inflation through the gap is expensive, so the Fed does little: $Z/(Z+\lambda_x)$ is small.

When the curve is steep, a small change in the gap moves inflation a lot, so a small rate response suffices: $\phi_\pi \approx 1 + \rho/(|g|\kappa)$ falls toward one.

The peak lies in between, at a smaller κ when ρ is larger.

The paper's footnote 19 suggests that a small control cost makes $\partial\phi_\pi/\partial\kappa > 0$ for all κ.

Numerically the hump survives a penalty on the real rate, so *Proposition 109.7.2* is best read as a statement about small slopes.

The slopes the Fed estimated in the 2010s lie inside that region.

Exercise 109.10.3

The baseline sets the gain to zero in both 2020 and 2021.

Recompute the beliefs and the recommended rates (a) freezing only 2020 and (b) with no freeze at all.

Report ρ_t, κ_t, and the recommended rate for 2021–2022.

Which part of the freeze drives the result, and what does setting the gain to zero in 2021 assume?

Solution

We loop over the three treatments of the pandemic quarters.

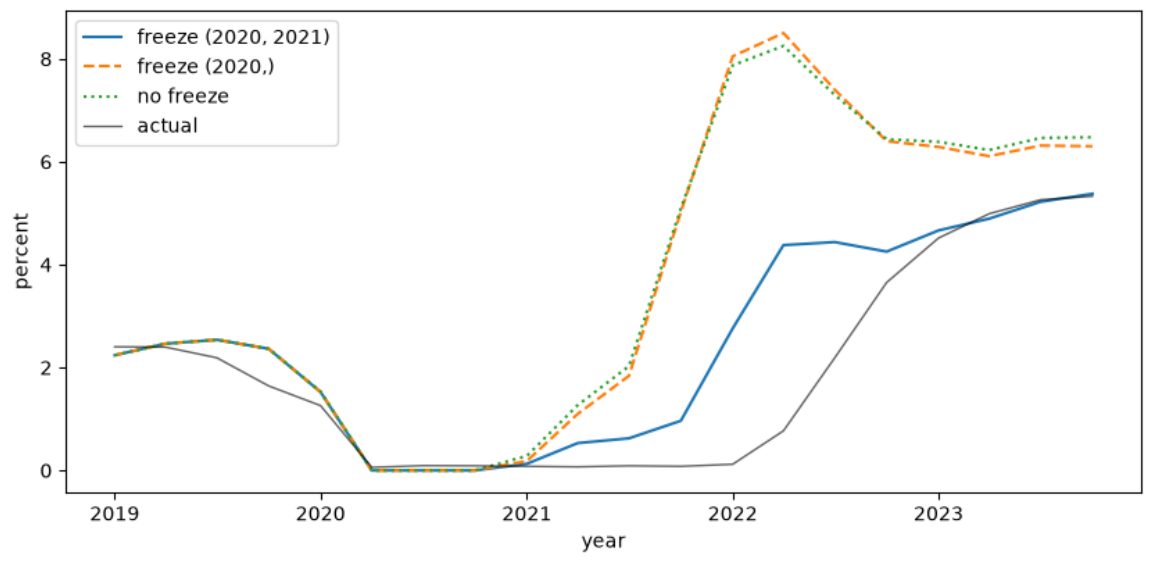
```
fig, ax = plt.subplots(figsize=(10, 4.5))
w = slice('2019', '2023')
for freeze, style in [((2020, 2021), '-'), ((2020,), '--'), ((), ':')]:
```

```
    b = estimate_beliefs(data, freeze=freeze)
    rate = recommended_rate(data, phelps_rules(b))
    label = f"freeze {freeze}" if freeze else 'no freeze'
    print(f"\n{label}")
    print(pd.DataFrame({'ρ': b['rho'], 'κ': b['kappa'], 'rate': rate})
            .loc['2021':'2022'].rename(index=qtr).round(3).T.to_string())
    ax.plot(rate[w], style, label=label)
ax.plot(data['i'][w], 'k', lw=1, alpha=0.6, label='actual')
ax.set_xlabel('year')
ax.set_ylabel('percent')
ax.legend()
plt.show()
```

```
freeze (2020, 2021)
      2021Q1  2021Q2  2021Q3  2021Q4  2022Q1  2022Q2  2022Q3  2022Q4
ρ      0.858   0.858   0.858   0.858   0.951   0.981   0.974   0.947
κ      0.039   0.039   0.039   0.039   0.048   0.047   0.047   0.048
rate   0.129   0.532   0.624   0.965   2.761   4.379   4.438   4.253

freeze (2020,)
      2021Q1  2021Q2  2021Q3  2021Q4  2022Q1  2022Q2  2022Q3  2022Q4
ρ      0.848   0.850   0.885   0.944   0.989   0.995   0.974   0.940
κ      0.042   0.076   0.084   0.106   0.107   0.106   0.109   0.113
rate   0.184   1.097   1.839   5.026   8.048   8.502   7.400   6.400

no freeze
      2021Q1  2021Q2  2021Q3  2021Q4  2022Q1  2022Q2  2022Q3  2022Q4
ρ      0.823   0.831   0.869   0.933   0.980   0.985   0.965   0.931
κ      0.078   0.102   0.107   0.121   0.121   0.120   0.122   0.125
rate   0.283   1.268   2.028   5.090   7.875   8.251   7.297   6.438
```



Freezing only 2020 gives almost the same answer as no freeze at all: the extreme 2020 quarters are not what matter.

What matters is 2021.

A Fed that learned from 2021 would have seen inflation rise with a positive gap, so its perceived slope κ_t would have steepened sharply and its perceived persistence would have risen.

Its recommended rate would have reached several percent by the end of 2021.

So the baseline's zero gain in 2021 builds in the judgment that the 2021 data were uninformative about inflation dynamics, which is the "transitory" view itself.

The freeze is a modeling choice that helps the model rationalize a slow Fed, not a neutral treatment of outliers.

Exercise 109.10.4

Chung [1990] treated the gain of a government's learning algorithm as a free parameter to be estimated (Exercise **??**).

Econometric Policy Evaluation takes up his problem on 1960–1995 data.

Pose Chung's problem for the artificial Fed.

For gains γ between 0.01 and 0.05, recompute the beliefs and recommendations and the RMSE against the actual rate from 1991 on.

Which gain fits best?

What happens to beliefs and recommendations in 2022 at $\gamma = 0.05$?

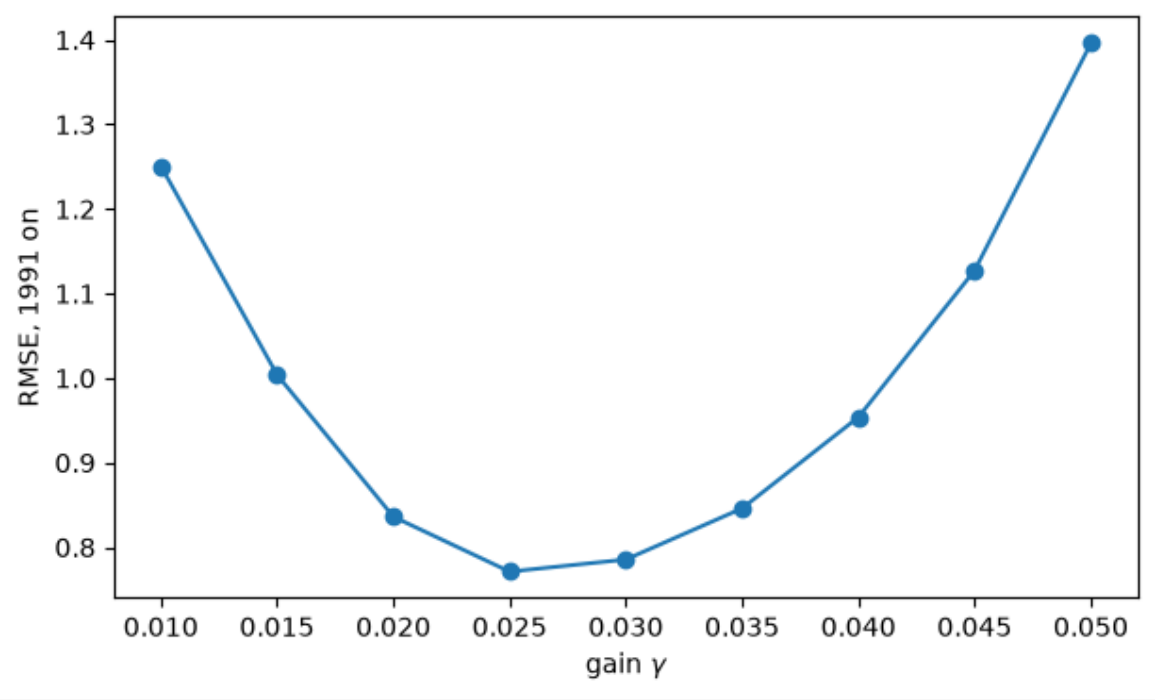
Solution

We keep track of the largest perceived persistence and the largest recommendation in 2022–2023 for each gain.

```
gains = np.arange(0.01, 0.0501, 0.005)
rmse, extremes = [], {}
for γ_val in gains:
    b = estimate_beliefs(data, γ=γ_val)
    rate = recommended_rate(data, phelps_rules(b))
    rmse.append(fit(rate)[1])
    extremes[round(γ_val, 3)] = (b['rho']['2022':'2023'].max(),
                                 rate['2022':'2023'].max())

fig, ax = plt.subplots(figsize=(7, 4))
ax.plot(gains, rmse, marker='o')
ax.set_xlabel(r'gain $\gamma$')
ax.set_ylabel('RMSE, 1991 on')
plt.show()

print(f"best gain on the grid: {gains[np.argmin(rmse)]:.3f}")
for γ_val in [0.03, 0.05]:
    ρ_max, rate_max = extremes[γ_val]
    print(f"γ = {γ_val}: max ρ_t in 2022-23 = {ρ_max:.3f}, "
          f"max recommended rate = {rate_max:.1f}%")
```




```
best gain on the grid: 0.025
γ = 0.03: max ρ_t in 2022-23 = 0.981, max recommended rate = 5.4%
γ = 0.05: max ρ_t in 2022-23 = 1.036, max recommended rate = 11.9%
```

The best fit is close to the paper's $\gamma = 0.03$.

Small gains make beliefs too sluggish to follow the drift.

Large gains make them chase noise.

At $\gamma = 0.05$, the 2022 surprises push perceived persistence above one: the Fed's model becomes explosive, and its Phelps problem calls for an enormous tightening.

The larger the gain, the more a few quarters of surprises can transform the Fed's model.

This criterion is a least-squares fit of the policy rate that treats the loss parameters as known, not Chung's likelihood for the whole system.

Exercise 109.10.5

This exercise reconciles the persistence measure of this lecture with the one in *Drifts and Volatilities* (*Inflation persistence*).

1. Show by simulation that if quarterly inflation were serially uncorrelated, a regression of four-quarter inflation on its one-quarter lag would give a coefficient near 0.75.
2. Rerun `estimate_beliefs` with quarterly annualized inflation $400 \log(P_t / P_{t-1})$ in place of four-quarter inflation, and compare the two paths of ρ_t.
3. When does persistence fall under each measure?

Solution

We start with the overlap, then rerun the beliefs on quarterly data.

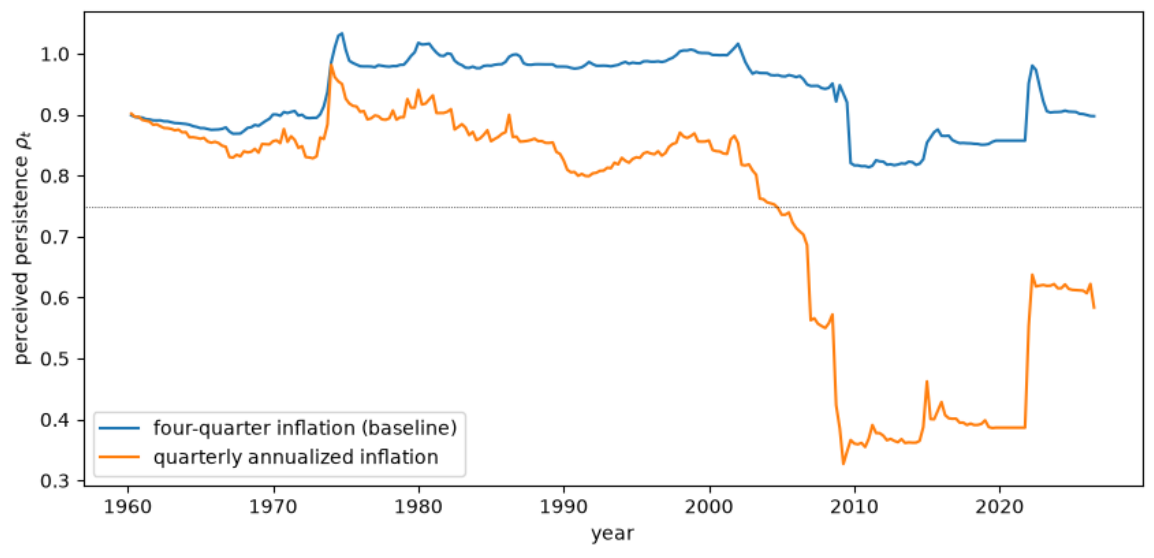
```
rng = np.random.default_rng(0)
quarterly = rng.standard_normal(100_000)
four_quarter = np.convolve(quarterly, np.ones(4) / 4, mode='valid')
coef = np.polyfit(four_quarter[:-1], four_quarter[1:], 1)[0]
print(f"regression coefficient with IID quarterly inflation: {coef:.3f}")

data_q = data.copy()
data_q['pi'] = 400 * np.log(pce / pce.shift(1))[data.index]
beliefs_q = estimate_beliefs(data_q)

fig, ax = plt.subplots(figsize=(10, 4.5))
ax.plot(beliefs['rho'], label='four-quarter inflation (baseline)')
ax.plot(beliefs_q['rho'], label='quarterly annualized inflation')
ax.axhline(0.75, color='k', lw=0.5, ls=':')
ax.set_xlabel('year')
ax.set_ylabel(r'perceived persistence $\rho_t$')
ax.legend()
plt.show()

dates = ['1980-01-01', '1985-01-01', '1995-01-01', '2005-01-01',
         '2015-01-01', '2019-10-01', '2022-04-01']
print(pd.DataFrame({'four-quarter': beliefs['rho'], 'quarterly': beliefs_q['rho']})
        .loc[dates].rename(index=qtr).round(2).T.to_string())
```

```
regression coefficient with IID quarterly inflation: 0.749
              1980Q1  1985Q1  1995Q1  2005Q1  2015Q1  2019Q4  2022Q2
four-quarter    1.02    0.98    0.98    0.96    0.85    0.86    0.98
quarterly       0.94    0.86    0.83    0.74    0.46    0.39    0.64
```



The overlap of four-quarter rates builds in a coefficient of about 0.75 even without any persistence in quarterly inflation, so the baseline ρ_t cannot fall far.

With quarterly inflation, persistence falls modestly after 1980, further by the mid-2000s, and to around 0.4 in the 2010s.

It rises again in 2022, but to a level far below one.

The overlap thus both inflates the level of the baseline measure and hides the early part of its decline.

The quarterly measure is closer to *Drifts and Volatilities*, which finds persistence falling soon after 1980.

Both measures agree on the main fact this lecture uses: by 2019 inflation was much less persistent than it had been.

Part XIV

Optimal Growth

CASS-KOOPMANS MODEL

110.1 Overview

This lecture and *Cass-Koopmans Competitive Equilibrium* describe a model that Tjalling Koopmans [Koopmans, 1965] and David Cass [Cass, 1965] used to analyze optimal growth.

The model extends the model of Robert Solow described in an earlier lecture.

It does so by making saving rate be a decision, instead of a hard-wired constant.

(Solow assumed a constant saving rate determined outside the model.)

We describe two versions of the model, a planning problem in this lecture, and a competitive equilibrium in this lecture *Cass-Koopmans Competitive Equilibrium*.

Together, the two lectures illustrate what is, in fact, a more general connection between a **planned economy** and a decentralized economy organized as a **competitive equilibrium**.

This lecture is devoted to the planned economy version.

In the planned economy, there are

- no prices
- no budget constraints

Instead there is a dictator that tells people

- what to produce
- what to invest in physical capital
- who is to consume what and when

The lecture uses important ideas including

- A min-max problem for solving a planning problem.
- A **shooting algorithm** for solving difference equations subject to initial and terminal conditions.
- A **turnpike** property of optimal paths for long but finite-horizon economies.
- A **stable manifold** and a **phase plane**

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
from numba import jit, float64
from numba.experimental import jitclass
import numpy as np
from quantecon.optimize import brentq
```

110.2 The Model

Time is discrete and takes values $t = 0, 1, \ldots, T$ where T is finite.

(We'll eventually study a limiting case in which $T = +\infty$)

A single good can either be consumed or invested in physical capital.

The consumption good is not durable and depreciates completely if not consumed immediately.

The capital good is durable but depreciates.

We let C_t be the total consumption of a nondurable consumption good at time t.

Let K_t be the stock of physical capital at time t.

Let $\vec{C} = \{C_0, \ldots, C_T\}$ and $\vec{K} = \{K_0, \ldots, K_{T+1}\}$.

110.2.1 Digression: Aggregation Theory

We use a concept of a representative consumer to be thought of as follows.

There is a unit mass of identical consumers indexed by $\omega \in [0, 1]$.

Consumption of consumer ω is $c(\omega)$.

Aggregate consumption is

$$C = \int_0^1 c(\omega) d\omega$$

Consider a welfare problem that chooses an allocation $\{c(\omega)\}$ across consumers to maximize

$$\int_0^1 u(c(\omega)) d\omega$$

where $u(\cdot)$ is a concave utility function with $u' > 0, u'' < 0$ and maximization is subject to

$$C = \int_0^1 c(\omega) d\omega. \tag{110.1}$$

Form a Lagrangian $L = \int_0^1 u(c(\omega)) d\omega + \lambda[C - \int_0^1 c(\omega) d\omega]$.

Differentiate under the integral signs with respect to each ω to obtain the first-order necessary conditions

$$u'(c(\omega)) = \lambda.$$

These conditions imply that $c(\omega)$ equals a constant c that is independent of ω.

To find c, use feasibility constraint (110.1) to conclude that

$$c(\omega) = c = C.$$

This line of argument indicates the special *aggregation theory* that lies beneath outcomes in which a representative consumer consumes amount C.

It appears often in aggregate economics.

We shall use this aggregation theory here and also in this lecture *Cass-Koopmans Competitive Equilibrium*.

An Economy

A representative household is endowed with one unit of labor at each t and likes the consumption good at each t.

The representative household inelastically supplies a single unit of labor N_t at each t, so that $N_t = 1$ for all $t \in \{0, 1, \dots, T\}$.

The representative household has preferences over consumption bundles ordered by the utility functional:

$$U(\vec{C}) = \sum_{t=0}^{T} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \tag{110.2}$$

where $\beta \in (0, 1)$ is a discount factor and $\gamma > 0$ governs the curvature of the one-period utility function.

Larger γ's imply more curvature.

Note that

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma} \tag{110.3}$$

satisfies $u' > 0, u'' < 0$.

$u' > 0$ asserts that the consumer prefers more to less.

$u'' < 0$ asserts that marginal utility declines with increases in C_t.

We assume that $K_0 > 0$ is an exogenous initial capital stock.

There is an economy-wide production function

$$F(K_t, N_t) = AK_t^{\alpha} N_t^{1-\alpha} \tag{110.4}$$

with $0 < \alpha < 1$, $A > 0$.

A feasible allocation $\vec{C}, \vec{K}$ satisfies

$$C_t + K_{t+1} \leq F(K_t, N_t) + (1-\delta)K_t \quad \text{for all } t \in \{0, 1, \dots, T\} \tag{110.5}$$

where $\delta \in (0, 1)$ is a depreciation rate of capital.

110.3 Planning Problem

A planner chooses an allocation $\{\vec{C}, \vec{K}\}$ to maximize (110.2) subject to (110.5).

Let $\vec{\mu} = \{\mu_0, \dots, \mu_T\}$ be a sequence of nonnegative **Lagrange multipliers**.

To find an optimal allocation, form a Lagrangian

$$\mathcal{L}(\vec{C}, \vec{K}, \vec{\mu}) = \sum_{t=0}^{T} \beta^t \left\{ u(C_t) + \mu_t \left(F(K_t, 1) + (1-\delta)K_t - C_t - K_{t+1} \right) \right\} \tag{110.6}$$

and pose the following min-max problem:

$$\min_{\vec{\mu}} \max_{\vec{C},\vec{K}} \mathcal{L}(\vec{C},\vec{K},\vec{\mu}) \tag{110.7}$$

- **Extremization** means maximization with respect to $\vec{C},\vec{K}$ and minimization with respect to $\vec{\mu}$.
- Our problem satisfies conditions that assure that second-order conditions are satisfied at an allocation that satisfies the first-order necessary conditions that we are about to compute.

Before computing first-order conditions, we present some handy formulas.

110.3.1 Useful Properties of Linearly Homogeneous Production Function

The following technicalities will help us.

Notice that

$$F(K_t,N_t) = AK_t^{\alpha}N_t^{1-\alpha} = N_t A\left(\frac{K_t}{N_t}\right)^{\alpha}$$

Define the **output per-capita production function**

$$\frac{F(K_t,N_t)}{N_t} \equiv f\left(\frac{K_t}{N_t}\right) = A\left(\frac{K_t}{N_t}\right)^{\alpha}$$

whose argument is **capital per-capita**.

It is useful to recall the following calculations for the marginal product of capital

$$\begin{aligned}
\frac{\partial F(K_t,N_t)}{\partial K_t} &= \frac{\partial N_t f\left(\frac{K_t}{N_t}\right)}{\partial K_t} \\
&= N_t f'\left(\frac{K_t}{N_t}\right)\frac{1}{N_t} \quad \text{(Chain rule)} \\
&= f'\left(\frac{K_t}{N_t}\right)\Big|_{N_t=1} \\
&= f'(K_t)
\end{aligned} \tag{110.8}$$

and the marginal product of labor

$$\begin{aligned}
\frac{\partial F(K_t,N_t)}{\partial N_t} &= \frac{\partial N_t f\left(\frac{K_t}{N_t}\right)}{\partial N_t} \quad \text{(Product rule)} \\
&= f\left(\frac{K_t}{N_t}\right) + N_t f'\left(\frac{K_t}{N_t}\right)\frac{-K_t}{N_t^2} \quad \text{(Chain rule)} \\
&= f\left(\frac{K_t}{N_t}\right) - \frac{K_t}{N_t}f'\left(\frac{K_t}{N_t}\right)\Big|_{N_t=1} \\
&= f(K_t) - f'(K_t)K_t
\end{aligned}$$

(Here we are using that $N_t = 1$ for all t, so that $K_t = \frac{K_t}{N_t}$.)

110.3.2 First-order necessary conditions

We now compute **first-order necessary conditions** for extremization of Lagrangian (110.6):

$$C_t: \quad u'(C_t) - \mu_t = 0 \quad \text{for all} \quad t = 0, 1, \ldots, T \tag{110.9}$$

$$K_t: \quad \beta\mu_t \left[(1-\delta) + f'(K_t)\right] - \mu_{t-1} = 0 \quad \text{for all} \quad t = 1, 2, \ldots, T \tag{110.10}$$

$$\mu_t: \quad F(K_t, 1) + (1-\delta)K_t - C_t - K_{t+1} = 0 \quad \text{for all} \quad t = 0, 1, \ldots, T \tag{110.11}$$

$$K_{T+1}: \quad -\mu_T \leq 0, \ \leq 0 \text{ if } K_{T+1} = 0; \ = 0 \text{ if } K_{T+1} > 0 \tag{110.12}$$

In computing (110.10) we recognize that K_t appears in both the time t and time $t-1$ feasibility constraints (110.5).

Restrictions (110.12) come from differentiating with respect to K_{T+1} and applying the following **Karush-Kuhn-Tucker condition** (KKT) (see Karush-Kuhn-Tucker conditions):

$$\mu_T K_{T+1} = 0 \tag{110.13}$$

Combining (110.9) and (110.10) gives

$$\beta u'\left(C_t\right)\left[(1-\delta) + f'\left(K_t\right)\right] - u'\left(C_{t-1}\right) = 0 \quad \text{for all } t = 1, 2, \ldots, T+1$$

which can be rearranged to become

$$\beta u'\left(C_{t+1}\right)\left[(1-\delta) + f'\left(K_{t+1}\right)\right] = u'\left(C_t\right) \quad \text{for all } t = 0, 1, \ldots, T \tag{110.14}$$

Applying the inverse marginal utility of consumption function on both sides of the above equation gives

$$C_{t+1} = u'^{-1}\left(\left(\frac{\beta}{u'(C_t)}[f'(K_{t+1}) + (1-\delta)]\right)^{-1}\right)$$

which for our utility function (110.3) becomes the consumption **Euler equation**

$$C_{t+1} = \left(\beta C_t^{\gamma}[f'(K_{t+1}) + (1-\delta)]\right)^{1/\gamma} \tag{110.15}$$

which we can combine with the feasibility constraint (110.5) to get

$$\begin{aligned} C_{t+1} &= C_t\left(\beta[f'(F(K_t, 1) + (1-\delta)K_t - C_t) + (1-\delta)]\right)^{1/\gamma} \\ K_{t+1} &= F(K_t, 1) + (1-\delta)K_t - C_t. \end{aligned} \tag{110.16}$$

This is a pair of non-linear first-order difference equations that map C_t, K_t into C_{t+1}, K_{t+1} and that an optimal sequence $\vec{C}, \vec{K}$ must satisfy.

It must also satisfy the initial condition that K_0 is given and $K_{T+1} = 0$.

Below we define a `jitclass` that stores parameters and functions that define our economy.

```
planning_data = [
    ('γ', float64),     # Coefficient of relative risk aversion
    ('β', float64),     # Discount factor
    ('δ', float64),     # Depreciation rate on capital
    ('α', float64),     # Return to capital per capita
    ('A', float64)      # Technology
]
```

```
@jitclass(planning_data)
class PlanningProblem():

    def __init__(self, γ=2, β=0.95, δ=0.02, α=0.33, A=1):

        self.γ, self.β = γ, β
        self.δ, self.α, self.A = δ, α, A

    def u(self, c):
        '''
        Utility function
        ASIDE: If you have a utility function that is hard to solve by hand
        you can use automatic or symbolic differentiation
        See https://github.com/HIPS/autograd
        '''
        γ = self.γ

        return c ** (1 - γ) / (1 - γ) if γ!= 1 else np.log(c)

    def u_prime(self, c):
        'Derivative of utility'
        γ = self.γ

        return c ** (-γ)

    def u_prime_inv(self, c):
        'Inverse of derivative of utility'
        γ = self.γ

        return c ** (-1 / γ)

    def f(self, k):
        'Production function'
        α, A = self.α, self.A

        return A * k ** α

    def f_prime(self, k):
        'Derivative of production function'
        α, A = self.α, self.A

        return α * A * k ** (α - 1)

    def f_prime_inv(self, k):
        'Inverse of derivative of production function'
        α, A = self.α, self.A

        return (k / (A * α)) ** (1 / (α - 1))

    def next_k_c(self, k, c):
        ''''
        Given the current capital Kt and an arbitrary feasible
        consumption choice Ct, computes Kt+1 by state transition law
        and optimal Ct+1 by Euler equation.
        '''
        β, δ = self.β, self.δ
        u_prime, u_prime_inv = self.u_prime, self.u_prime_inv
```

(continues on next page)

(continued from previous page)

```
        f, f_prime = self.f, self.f_prime

        k_next = f(k) + (1 - δ) * k - c
        c_next = u_prime_inv(u_prime(c) / (β * (f_prime(k_next) + (1 - δ))))

        return k_next, c_next
```

We can construct an economy with the Python code:

```
pp = PlanningProblem()
```

110.4 Shooting Algorithm

We use **shooting** to compute an optimal allocation $\vec{C}, \vec{K}$ and an associated Lagrange multiplier sequence $\vec{\mu}$.

First-order necessary conditions (110.9), (110.10), and (110.11) for the planning problem form a system of **difference equations** with two boundary conditions:

- K_0 is a given **initial condition** for capital
- $K_{T+1} = 0$ is a **terminal condition** for capital that we deduced from the first-order necessary condition for K_{T+1} the KKT condition (110.13)

We have no initial condition for the Lagrange multiplier μ_0.

If we did, our job would be easy:

- Given μ_0 and k_0, we could compute c_0 from equation (110.9) and then k_1 from equation (110.11) and μ_1 from equation (110.10).
- We could continue in this way to compute the remaining elements of $\vec{C}, \vec{K}, \vec{\mu}$.

However, we woujld not be assured that the Kuhn-Tucker condition (110.13) would be satisfied.

Furthermore, we don't have an initial condition for μ_0.

So this won't work.

Indeed, part of our task is to compute the **optimal** value of μ_0.

To compute μ_0 and the other objects we want, a simple modification of the above procedure will work.

It is called the **shooting algorithm**.

It is an instance of a **guess and verify** algorithm that consists of the following steps:

- Guess an initial Lagrange multiplier μ_0.
- Apply the **simple algorithm** described above.
- Compute K_{T+1} and check whether it equals zero.
- If $K_{T+1} = 0$, we have solved the problem.
- If $K_{T+1} > 0$, lower μ_0 and try again.
- If $K_{T+1} < 0$, raise μ_0 and try again.

The following Python code implements the shooting algorithm for the planning problem.

(Actually, we modified the preceding algorithm slightly by starting with a guess for c_0 instead of μ_0 in the following code.)

```
@jit
def shooting(pp, c0, k0, T=10):
    '''
    Given the initial condition of capital k0 and an initial guess
    of consumption c0, computes the whole paths of c and k
    using the state transition law and Euler equation for T periods.
    '''
    if c0 > pp.f(k0) + (1 - pp.δ) * k0:
        print("initial consumption is not feasible")

        return None

    # initialize vectors of c and k
    c_vec = np.empty(T+1)
    k_vec = np.empty(T+2)

    c_vec[0] = c0
    k_vec[0] = k0

    for t in range(T):
        k_vec[t+1], c_vec[t+1] = pp.next_k_c(k_vec[t], c_vec[t])

    k_vec[T+1] = pp.f(k_vec[T]) + (1 - pp.δ) * k_vec[T] - c_vec[T]

    return c_vec, k_vec
```

We'll start with an incorrect guess.

```
paths = shooting(pp, 0.2, 0.3, T=10)
```

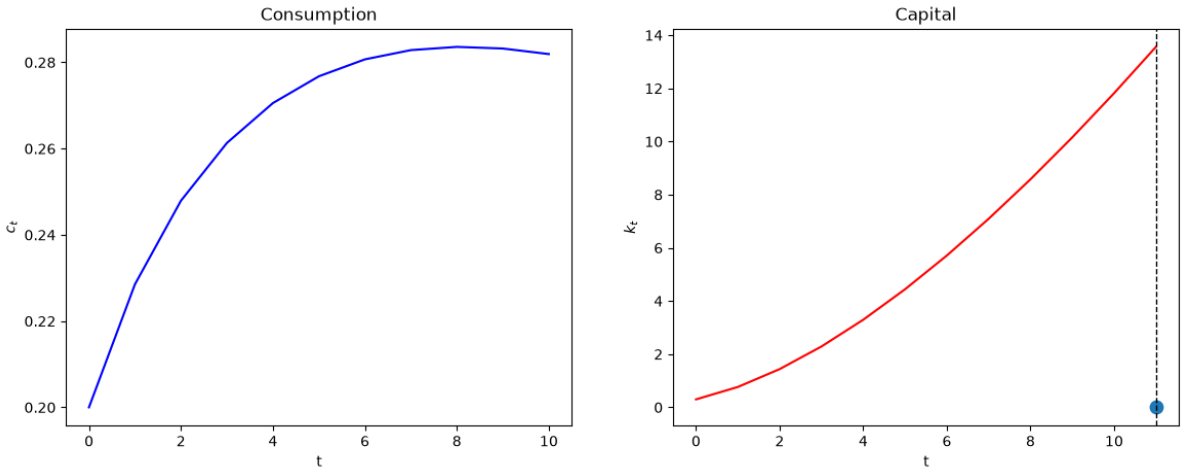
```
fig, axs = plt.subplots(1, 2, figsize=(14, 5))

colors = ['blue', 'red']
titles = ['Consumption', 'Capital']
ylabels = ['$c_t$', '$k_t$']

T = paths[0].size - 1
for i in range(2):
    axs[i].plot(paths[i], c=colors[i])
    axs[i].set(xlabel='t', ylabel=ylabels[i], title=titles[i])

axs[1].scatter(T+1, 0, s=80)
axs[1].axvline(T+1, color='k', ls='--', lw=1)

plt.show()
```



Evidently, our initial guess for μ_0 is too high, so initial consumption too low.

We know this because we miss our $K_{T+1} = 0$ target on the high side.

Now we automate things with a search-for-a-good μ_0 algorithm that stops when we hit the target $K_{t+1} = 0$.

We use a **bisection method**.

We make an initial guess for C_0 (we can eliminate μ_0 because C_0 is an exact function of μ_0).

We know that the lowest C_0 can ever be is 0 and that the largest it can be is initial output $f(K_0)$.

Guess C_0 and shoot forward to $T + 1$.

If $K_{T+1} > 0$, we take it to be our new **lower** bound on C_0.

If $K_{T+1} < 0$, we take it to be our new **upper** bound.

Make a new guess for C_0 that is halfway between our new upper and lower bounds.

Shoot forward again, iterating on these steps until we converge.

When K_{T+1} gets close enough to 0 (i.e., within an error tolerance bounds), we stop.

```
@jit
def bisection(pp, c0, k0, T=10, tol=1e-4, max_iter=500, k_ter=0, verbose=True):

    # initial boundaries for guess c0
    c0_upper = pp.f(k0) + (1 - pp.δ) * k0
    c0_lower = 0

    i = 0
    while True:
        c_vec, k_vec = shooting(pp, c0, k0, T)
        error = k_vec[-1] - k_ter

        # check if the terminal condition is satisfied
        if np.abs(error) < tol:
            if verbose:
                print('Converged successfully on iteration ', i+1)
            return c_vec, k_vec

        i += 1
        if i == max_iter:
            if verbose:
```

(continues on next page)

(continued from previous page)

```
                print('Convergence failed.')
            return c_vec, k_vec

        # if iteration continues, updates boundaries and guess of c0
        if error > 0:
            c0_lower = c0
        else:
            c0_upper = c0

        c0 = (c0_lower + c0_upper) / 2
```

```
def plot_paths(pp, c0, k0, T_arr, k_ter=0, k_ss=None, axs=None):

    if axs is None:
        fix, axs = plt.subplots(1, 3, figsize=(16, 4))
    ylabels = ['$c_t$', '$k_t$', r'$\mu_t$']
    titles = ['Consumption', 'Capital', 'Lagrange Multiplier']

    c_paths = []
    k_paths = []
    for T in T_arr:
        c_vec, k_vec = bisection(pp, c0, k0, T, k_ter=k_ter, verbose=False)
        c_paths.append(c_vec)
        k_paths.append(k_vec)

        μ_vec = pp.u_prime(c_vec)
        paths = [c_vec, k_vec, μ_vec]

        for i in range(3):
            axs[i].plot(paths[i])
            axs[i].set(xlabel='t', ylabel=ylabels[i], title=titles[i])

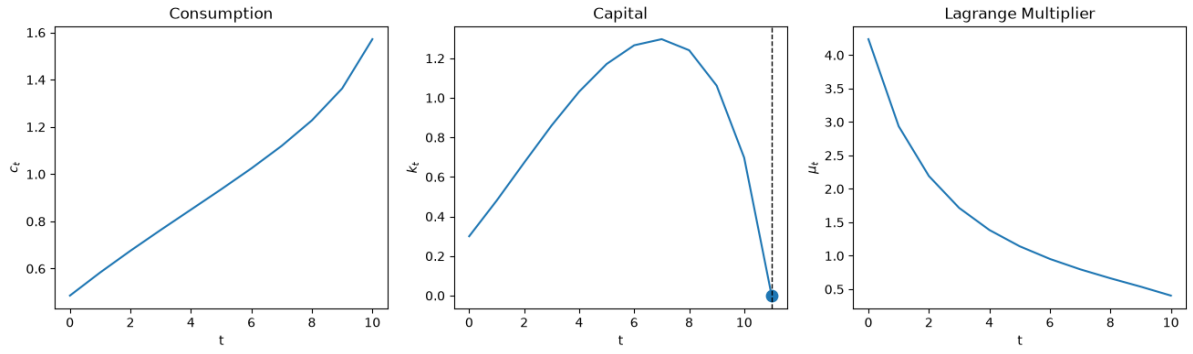
        # Plot steady state value of capital
        if k_ss is not None:
            axs[1].axhline(k_ss, c='k', ls='--', lw=1)

        axs[1].axvline(T+1, c='k', ls='--', lw=1)
        axs[1].scatter(T+1, paths[1][-1], s=80)

    return c_paths, k_paths
```

Now we can solve the model and plot the paths of consumption, capital, and Lagrange multiplier.

```
plot_paths(pp, 0.3, 0.3, [10]);
```



110.5 Setting Initial Capital to Steady State Capital

When $T \to +\infty$, the optimal allocation converges to steady state values of C_t and K_t.

It is instructive to set K_0 equal to the $\lim_{T \to +\infty} K_t$, which we'll call steady state capital.

In a steady state $K_{t+1} = K_t = \bar{K}$ for all very large t.

Evalauating feasibility constraint (110.5) at $\bar{K}$ gives

$$f(\bar{K}) - \delta \bar{K} = \bar{C} \tag{110.17}$$

Substituting $K_t = \bar{K}$ and $C_t = \bar{C}$ for all t into (110.14) gives

$$1 = \beta \frac{u'(\bar{C})}{u'(\bar{C})} [f'(\bar{K}) + (1 - \delta)]$$

Defining $\beta = \frac{1}{1+\rho}$, and cancelling gives

$$1 + \rho = 1[f'(\bar{K}) + (1 - \delta)]$$

Simplifying gives

$$f'(\bar{K}) = \rho + \delta$$

and

$$\bar{K} = f'^{-1}(\rho + \delta)$$

For production function (110.4), this becomes

$$\alpha \bar{K}^{\alpha - 1} = \rho + \delta$$

As an example, after setting $\alpha = .33$, $\rho = 1/\beta - 1 = 1/(19/20) - 1 = 20/19 - 19/19 = 1/19$, $\delta = 1/50$, we get

$$\bar{K} = \left(\frac{\frac{33}{100}}{\frac{1}{50} + \frac{1}{19}} \right)^{\frac{67}{100}} \approx 9.57583$$

Let's verify this with Python and then use this steady state $\bar{K}$ as our initial capital stock K_0.

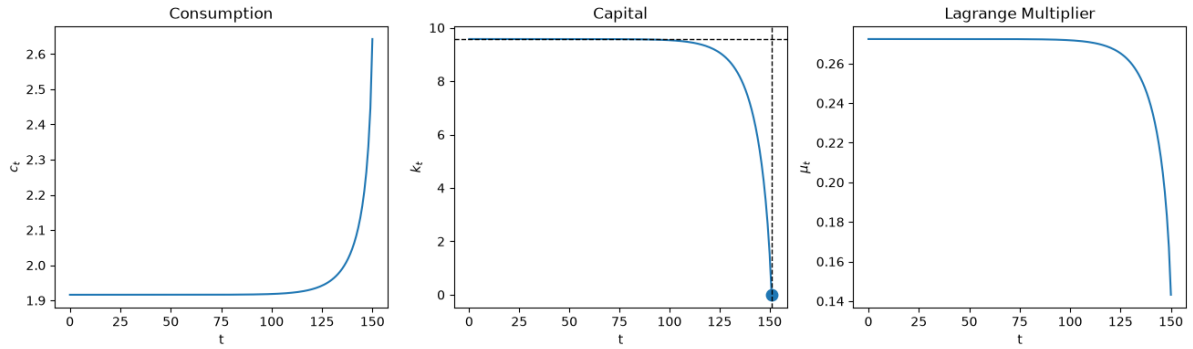
```
ρ = 1 / pp.β - 1
k_ss = pp.f_prime_inv(ρ+pp.δ)

print(f'steady state for capital is: {k_ss}')
```

```
steady state for capital is: 9.57583816331462
```

Now we plot

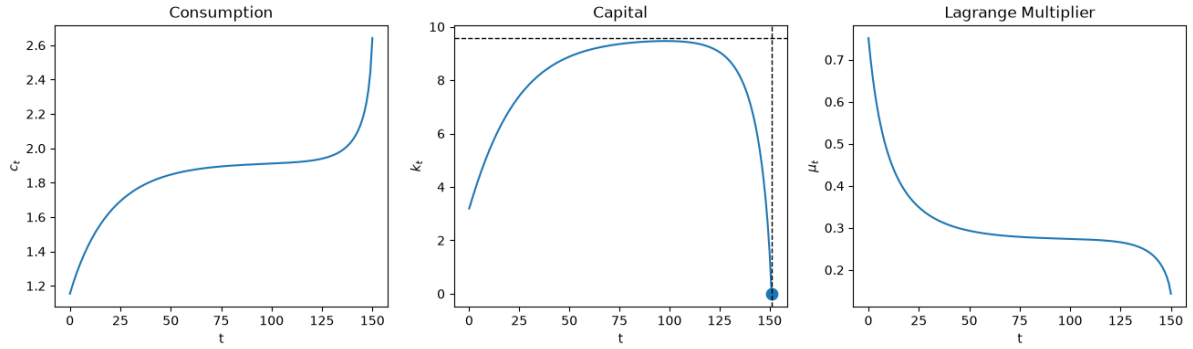
```
plot_paths(pp, 0.3, k_ss, [150], k_ss=k_ss);
```



Evidently, with a large value of T, K_t stays near K_0 until t approaches T closely.

Let's see what the planner does when we set K_0 below $\bar{K}$.

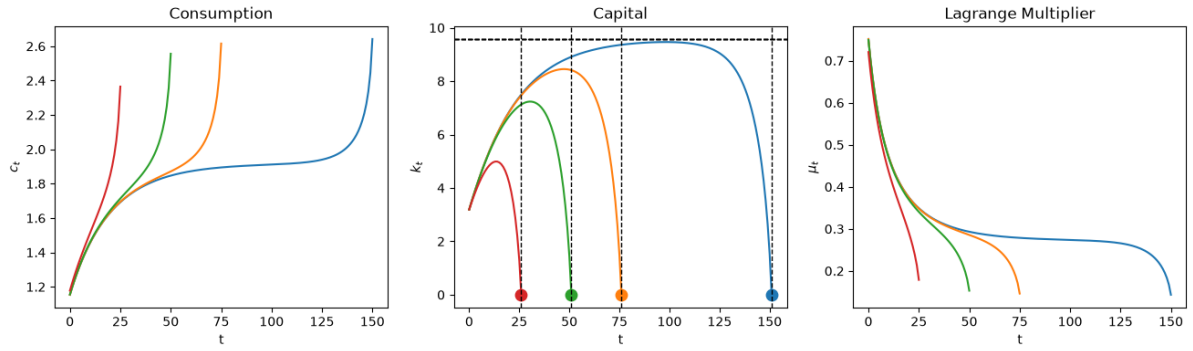
```
plot_paths(pp, 0.3, k_ss/3, [150], k_ss=k_ss);
```



Notice how the planner pushes capital toward the steady state, stays near there for a while, then pushes K_t toward the terminal value $K_{T+1} = 0$ when t closely approaches T.

The following graphs compare optimal outcomes as we vary T.

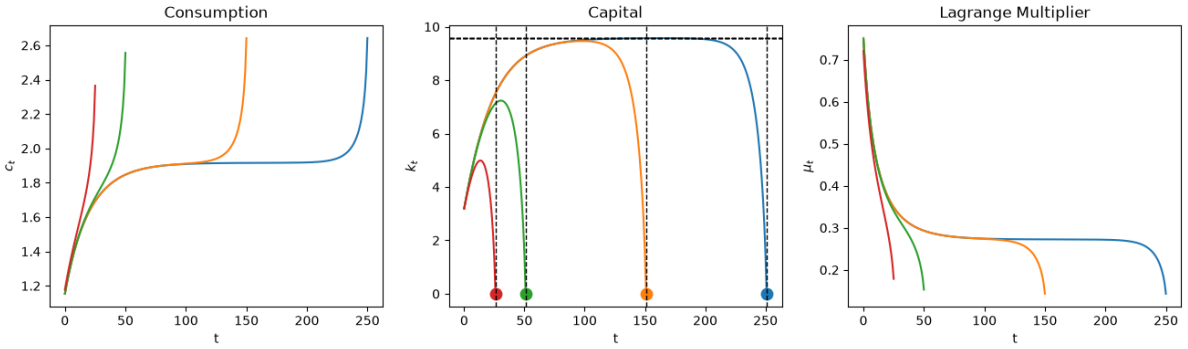
```
plot_paths(pp, 0.3, k_ss/3, [150, 75, 50, 25], k_ss=k_ss);
```



110.6 A Turnpike Property

The following calculation indicates that when T is very large, the optimal capital stock stays close to its steady state value most of the time.

```
plot_paths(pp, 0.3, k_ss/3, [250, 150, 50, 25], k_ss=k_ss);
```



In the above graphs, different colors are associated with different horizons T.

Notice that as the horizon increases, the planner keeps K_t closer to the steady state value $\bar{K}$ for longer.

This pattern reflects a **turnpike** property of the steady state.

A rule of thumb for the planner is

- from K_0, push K_t toward the steady state and stay close to the steady state until time approaches T.

The planner accomplishes this by adjusting the saving rate $\frac{f(K_t)-C_t}{f(K_t)}$ over time.

Exercise 110.6.1

The turnpike property is independent of the initial condition K_0 provided that T is sufficiently large.

Expand the `plot_paths` function so that it plots trajectories for multiple initial points using `k0s = [k_ss*2, k_ss*3, k_ss/3]`.

Solution

Here is one solution

```
def plot_multiple_paths(pp, c0, k0s, T_arr, k_ter=0, k_ss=None, axs=None):
    if axs is None:
        fig, axs = plt.subplots(1, 3, figsize=(16, 4))

    ylabels = ['$c_t$', '$k_t$', r'$\mu_t$']
    titles = ['Consumption', 'Capital', 'Lagrange Multiplier']

    colors = plt.cm.viridis(np.linspace(0, 1, len(k0s)))

    all_c_paths = []
    all_k_paths = []

    for i, k0 in enumerate(k0s):
        k0_c_paths = []
        k0_k_paths = []
```

```
        for T in T_arr:
            c_vec, k_vec = bisection(pp, c0, k0, T, k_ter=k_ter, verbose=False)
            k0_c_paths.append(c_vec)
            k0_k_paths.append(k_vec)

            μ_vec = pp.u_prime(c_vec)
            paths = [c_vec, k_vec, μ_vec]

            for j in range(3):
                axs[j].plot(paths[j], color=colors[i],
                            label=f'$k_0 = {k0:.2f}$' if j == 0 and T == T_arr[0]␣
↪else "", alpha=0.7)
                axs[j].set(xlabel='t', ylabel=ylabels[j], title=titles[j])

            if k_ss is not None and i == 0 and T == T_arr[0]:
                axs[1].axhline(k_ss, c='k', ls='--', lw=1)

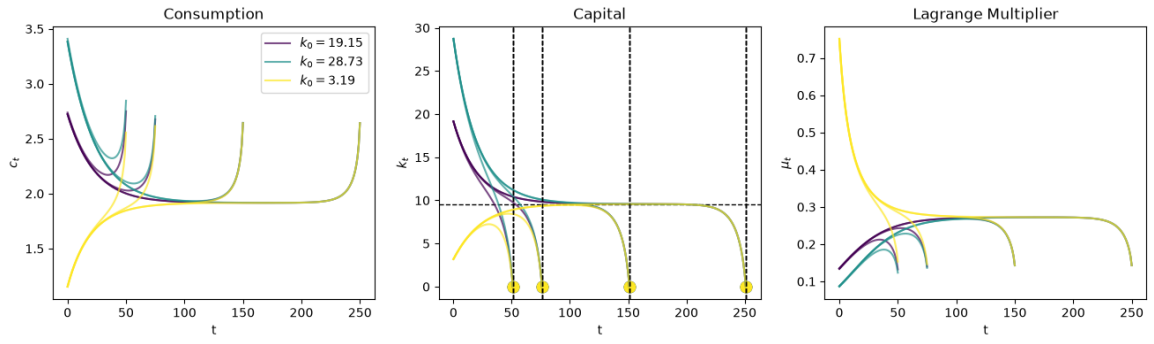
            axs[1].axvline(T+1, c='k', ls='--', lw=1)
            axs[1].scatter(T+1, paths[1][-1], s=80, color=colors[i])

        all_c_paths.append(k0_c_paths)
        all_k_paths.append(k0_k_paths)

    # Add legend if multiple initial points
    if len(k0s) > 1:
        axs[0].legend()

    return all_c_paths, all_k_paths
```

```
_ = plot_multiple_paths(pp, 0.3, [k_ss*2, k_ss*3, k_ss/3], [250, 150, 75, 50], k_
↪ss=k_ss)
```



We see that the turnpike property holds for various initial values of K_0.

Let's calculate and plot the saving rate.

```
@jit
def saving_rate(pp, c_path, k_path):
    'Given paths of c and k, computes the path of saving rate.'
    production = pp.f(k_path[:-1])

    return (production - c_path) / production
```

```
def plot_saving_rate(pp, c0, k0, T_arr, k_ter=0, k_ss=None, s_ss=None):

    fix, axs = plt.subplots(2, 2, figsize=(12, 9))

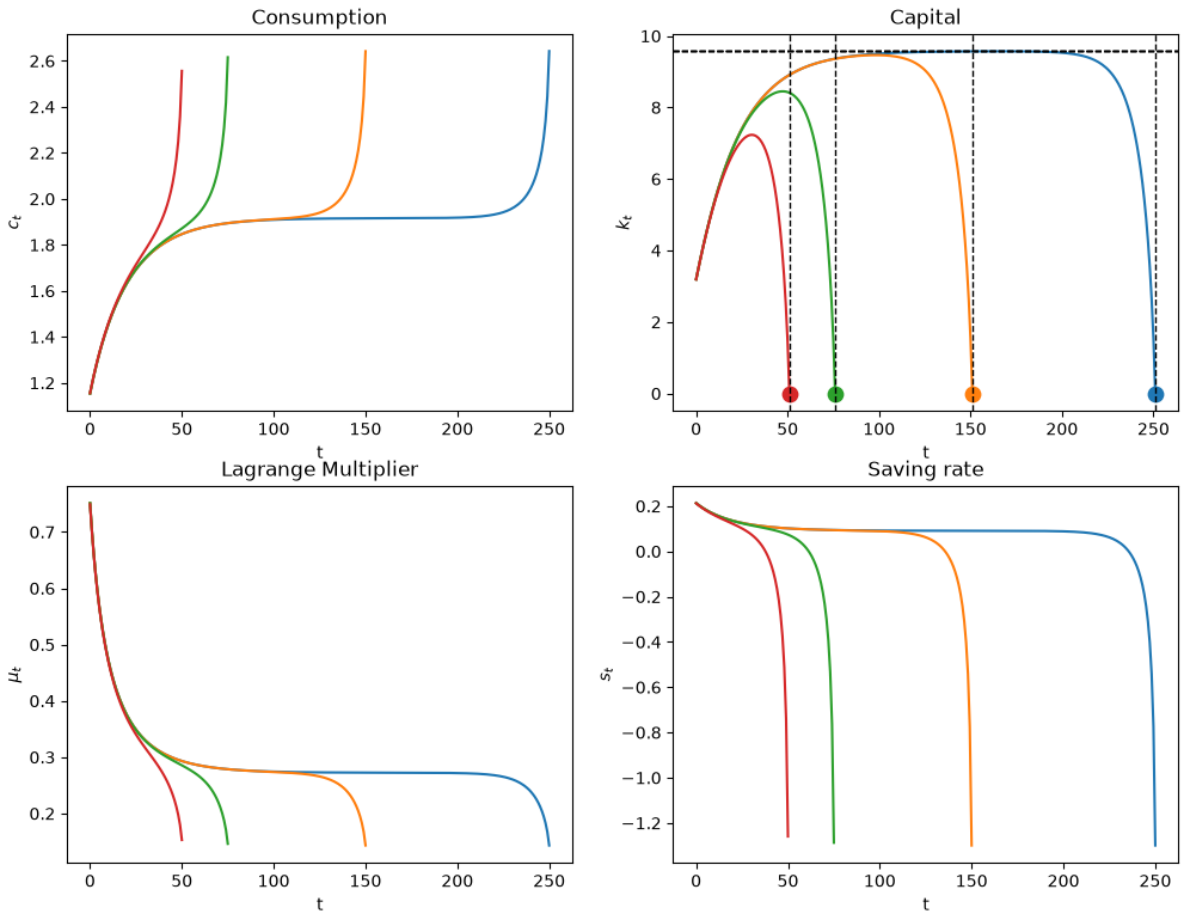
    c_paths, k_paths = plot_paths(pp, c0, k0, T_arr, k_ter=k_ter, k_ss=k_ss, axs=axs.
↪flatten())

    for i, T in enumerate(T_arr):
        s_path = saving_rate(pp, c_paths[i], k_paths[i])
        axs[1, 1].plot(s_path)

    axs[1, 1].set(xlabel='t', ylabel='$s_t$', title='Saving rate')

    if s_ss is not None:
        axs[1, 1].hlines(s_ss, 0, np.max(T_arr), linestyle='--')
```

```
plot_saving_rate(pp, 0.3, k_ss/3, [250, 150, 75, 50], k_ss=k_ss)
```



110.7 A Limiting Infinite Horizon Economy

We want to set $T = +\infty$.

The appropriate thing to do is to replace terminal condition (110.12) with

$$\lim_{T \to +\infty} \beta^T u'(C_T) K_{T+1} = 0,$$

a condition that will be satisfied by a path that converges to an optimal steady state.

We can approximate the optimal path by starting from an arbitrary initial K_0 and shooting towards the optimal steady state K at a large but finite $T + 1$.

In the following code, we do this for a large T and plot consumption, capital, and the saving rate.

We know that in the steady state that the saving rate is constant and that $\bar{s} = \frac{f(\bar{K}) - \bar{C}}{f(\bar{K})}$.

From (110.17) the steady state saving rate equals

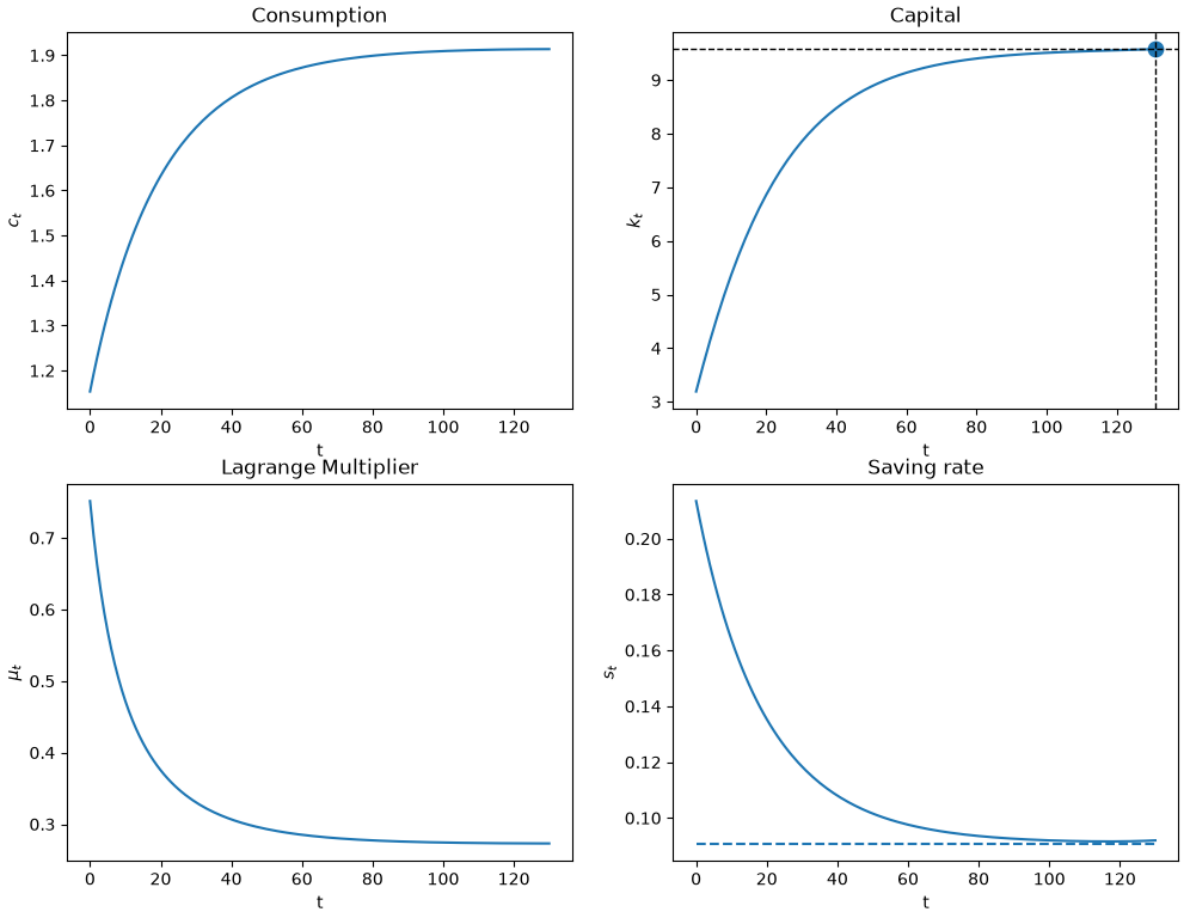
$$\bar{s} = \frac{\delta \bar{K}}{f(\bar{K})}$$

The steady state saving rate $\bar{S} = \bar{s} f(\bar{K})$ is the amount required to offset capital depreciation each period.

We first study optimal capital paths that start below the steady state.

```
# steady state of saving rate
s_ss = pp.δ * k_ss / pp.f(k_ss)

plot_saving_rate(pp, 0.3, k_ss/3, [130], k_ter=k_ss, k_ss=k_ss, s_ss=s_ss)
```



Since $K_0 < \bar{K}$, $f'(K_0) > \rho + \delta$.

The planner chooses a positive saving rate that is higher than the steady state saving rate.

Note that $f''(K) < 0$, so as K rises, $f'(K)$ declines.

The planner slowly lowers the saving rate until reaching a steady state in which $f'(K) = \rho + \delta$.

110.8 Stable Manifold and Phase Diagram

We now describe a classic diagram that describes an optimal (K_{t+1}, C_t) path.

The diagram has K on the ordinate axis and C on the coordinate axis.

Given an arbitrary and fixed K, a fixed point C of the consumption Euler equation (110.15) satisfies

$$C = C\left(\beta\left[f'\left(f\left(K\right) + \left(1-\delta\right)K - C\right) + \left(1-\delta\right)\right]\right)^{1/\gamma}$$

which implies

$$\begin{aligned} C &= f\left(K\right) + \left(1-\delta\right)K - f'^{-1}\left(\frac{1}{\beta} - \left(1-\delta\right)\right) \\ &\equiv \tilde{C}\left(K\right) \end{aligned} \tag{110.18}$$

A positive fixed point $C = \tilde{C}(K)$ exists only if $f\left(K\right) + \left(1-\delta\right)K - f'^{-1}\left(\frac{1}{\beta} - \left(1-\delta\right)\right) > 0$

```
@jit
def C_tilde(K, pp):

    return pp.f(K) + (1 - pp.δ) * K - pp.f_prime_inv(1 / pp.β - 1 + pp.δ)
```

Next note that given a time-invariant arbitrary C, a fixed point K of the feasibility condition (110.5) solves the following equation

$$K = f(K) + (1 - \delta K) - C.$$

A fixed point of the above equation is described by a function

$$K = \tilde{K}(C) \tag{110.19}$$

```
@jit
def K_diff(K, C, pp):
    return pp.f(K) - pp.δ * K - C

@jit
def K_tilde(C, pp):

    res = brentq(K_diff, 1e-6, 100, args=(C, pp))

    return res.root
```

A steady state (K_s, C_s) is a pair (K, C) that satisfies both equations (110.18) and (110.19).

It is thus the intersection of the two curves $\tilde{C}$ and $\tilde{K}$ that we'll plot in Figure Fig. 110.1 below.

We can compute K_s by solving the equation $K_s = \tilde{K}\left(\tilde{C}\left(K_s\right)\right)$

```
@jit
def K_tilde_diff(K, pp):

    K_out = K_tilde(C_tilde(K, pp), pp)

    return K - K_out
```

```
res = brentq(K_tilde_diff, 8, 10, args=(pp,))

Ks = res.root
Cs = C_tilde(Ks, pp)

Ks, Cs
```

```
(9.575838163314447, 1.9160839808123402)
```

We can use the shooting algorithm to compute trajectories that approach (K_s, C_s).

For a given K, let's compute $\vec{C}$ and $\vec{K}$ for a large T , e.g., $= 200$.

We compute C_0 by the bisection algorithm that assures that $K_T = K_s$.

Let's compute two trajectories towards (K_s, C_s) that start from different sides of K_s: $\bar{K}_0 = 1e-3 < K_s < \bar{K}_1 = 15$.

```
c_vec1, k_vec1 = bisection(pp, 5, 15, T=200, k_ter=Ks)
c_vec2, k_vec2 = bisection(pp, 1e-3, 1e-3, T=200, k_ter=Ks)
```

```
Converged successfully on iteration  46
Converged successfully on iteration  51
```

The following code generates Figure Fig. 110.1, which is patterned on a graph that appears on page 411 of [Intriligator, 2002].

Figure Fig. 110.1 is a classic "phase plane" with "state" variable K on the ordinate axis and "co-state" variable C on the coordinate axis.

Figure Fig. 110.1 plots three curves:

- the blue line graphs $C = \tilde{C}(K)$ of fixed points described by equation (110.18).
- the red line graphs $K = \tilde{K}(C)$ of fixed points described by equation (110.19)
- the green line graphs the stable traced out by paths that converge to the steady state starting from an arbitrary K_0 at time 0.
  - for a given K_0, the shooting algorithm sets C_0 to the coordinate on the green line in order to initiate a path that converges to the optimal steady state
  - the arrows on the green line show the direction in which dynamics (110.16) push successive (K_{t+1}, C_t) pairs.

In addition to the three curves, Figure Fig. 110.1 plots arrows that point where the dynamics (110.16) drive the system when, for a given K_0, C_0 is not on the stable manifold depicted in the green line.

- If C_0 is set below the green line for a given K_0, too much capital is accumulated
- If C_0 is set above the green line for a given K_0, too little capital is accumulated

110.9 Concluding Remarks

In *Cass-Koopmans Competitive Equilibrium*, we study a decentralized version of an economy with exactly the same technology and preference structure as deployed here.

In that lecture, we replace the planner of this lecture with Adam Smith's **invisible hand**.

In place of quantity choices made by the planner, there are market prices that are set by a *deus ex machina* from outside the model, a so-called invisible hand.

Equilibrium market prices must reconcile distinct decisions that are made independently by a representative household and a representative firm.

The relationship between a command economy like the one studied in this lecture and a market economy like that studied in *Cass-Koopmans Competitive Equilibrium* is a foundational topic in general equilibrium theory and welfare economics.

110.9.1 Exercise

Exercise 110.9.1

- Plot the optimal consumption, capital, and saving paths when the initial capital level begins at 1.5 times the steady state level as we shoot towards the steady state at $T = 130$.
- Why does the saving rate respond as it does?

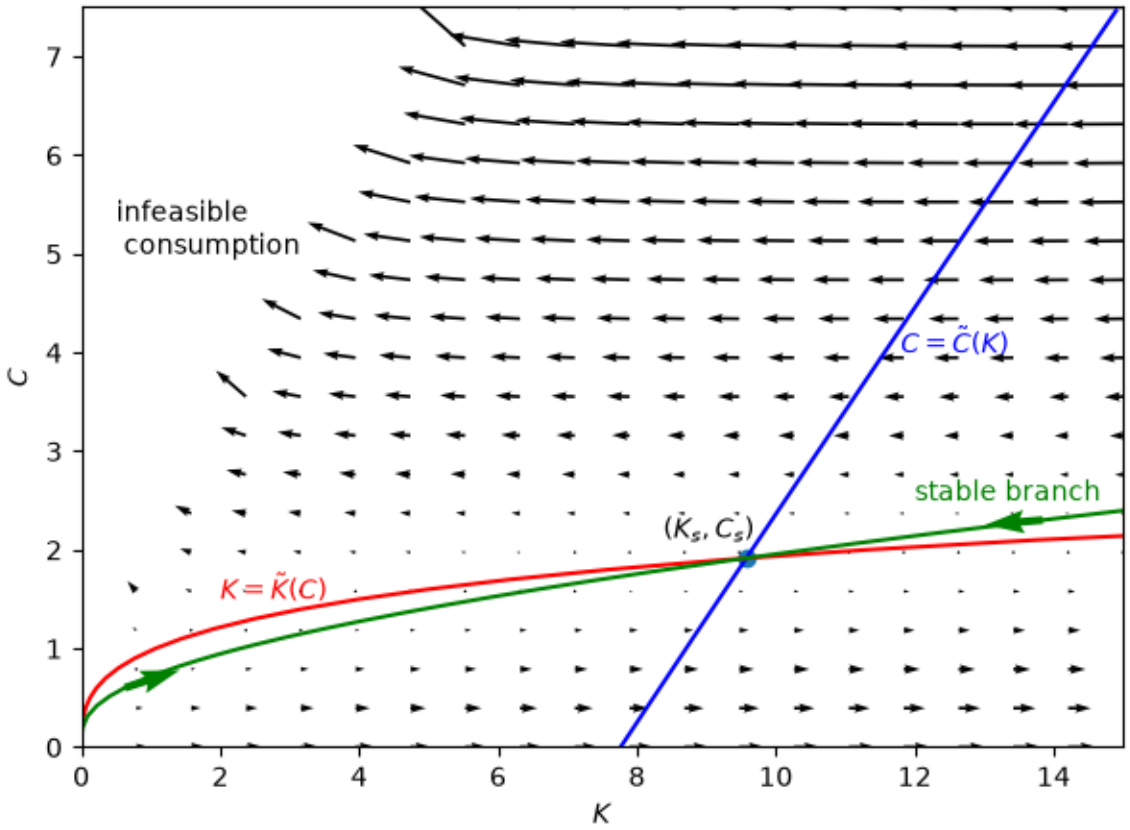
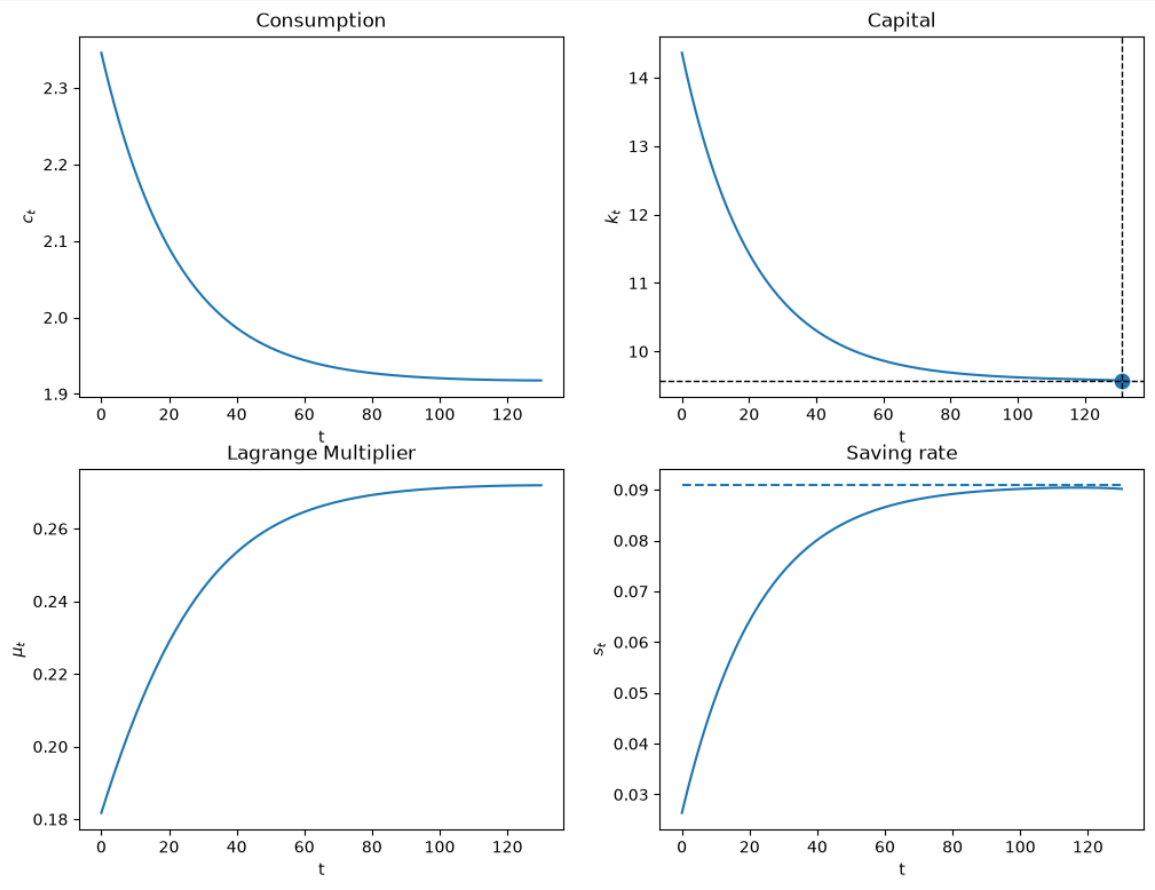


Fig. 110.1: Stable Manifold and Phase Plane

Solution

```
plot_saving_rate(pp, 0.3, k_ss*1.5, [130], k_ter=k_ss, k_ss=k_ss, s_ss=s_ss)
```



CHAPTER

ONE

CASS-KOOPMANS COMPETITIVE EQUILIBRIUM

Contents

- *Cass-Koopmans Competitive Equilibrium*
  - *Overview*
  - *Review of Cass-Koopmans Model*
  - *Competitive Equilibrium*
  - *Market Structure*
  - *Firm Problem*
  - *Household Problem*
  - *Computing a Competitive Equilibrium*
  - *Yield Curves and Hicks-Arrow Prices*

111.1 Overview

This lecture continues our analysis in this lecture *Cass-Koopmans Planning Model* about the model that Tjalling Koopmans [Koopmans, 1965] and David Cass [Cass, 1965] used to study optimal capital accumulation.

This lecture illustrates what is, in fact, a more general connection between a **planned economy** and an economy organized as a competitive equilibrium or a **market economy**.

The earlier lecture *Cass-Koopmans Planning Model* studied a planning problem and used ideas including

- A Lagrangian formulation of the planning problem that leads to a system of difference equations.
- A **shooting algorithm** for solving difference equations subject to initial and terminal conditions.
- A **turnpike** property that describes optimal paths for long-but-finite horizon economies.

The present lecture uses additional ideas including

- Hicks-Arrow prices, named after John R. Hicks and Kenneth Arrow.
- A connection between some Lagrange multipliers from the planning problem and the Hicks-Arrow prices.
- A **Big** K **, little** k trick widely used in macroeconomic dynamics.
  - We shall encounter this trick in this lecture and also in this lecture.

- A non-stochastic version of a theory of the **term structure of interest rates**.
- An intimate connection between two ways to organize an economy, namely:
  - **socialism** in which a central planner commands the allocation of resources, and
  - **competitive markets** in which competitive equilibrium **prices** induce individual consumers and producers to choose a socially optimal allocation as unintended consequences of their selfish decisions

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
from numba import jit, float64
from numba.experimental import jitclass
import numpy as np
```

111.2 Review of Cass-Koopmans Model

The physical setting is identical with that in *Cass-Koopmans Planning Model*.

Time is discrete and takes values $t = 0, 1, \ldots, T$.

Output of a single good can either be consumed or invested in physical capital.

The capital good is durable but partially depreciates each period at a constant rate.

We let C_t be a nondurable consumption good at time t.

Let K_t be the stock of physical capital at time t.

Let $\vec{C} = \{C_0, \ldots, C_T\}$ and $\vec{K} = \{K_0, \ldots, K_{T+1}\}$.

A representative household is endowed with one unit of labor at each t and likes the consumption good at each t.

The representative household inelastically supplies a single unit of labor N_t at each t, so that $N_t = 1$ for all $t \in \{0, 1, \ldots, T\}$.

The representative household has preferences over consumption bundles ordered by the utility functional:

$$U(\vec{C}) = \sum_{t=0}^{T} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}$$

where $\beta \in (0, 1)$ is a discount factor and $\gamma > 0$ governs the curvature of the one-period utility function.

We assume that $K_0 > 0$.

There is an economy-wide production function

$$F(K_t, N_t) = AK_t^{\alpha} N_t^{1-\alpha}$$

with $0 < \alpha < 1$, $A > 0$.

A feasible allocation $\vec{C}, \vec{K}$ satisfies

$$C_t + K_{t+1} \leq F(K_t, N_t) + (1-\delta)K_t \quad \text{for all } t \in \{0, 1, \ldots, T\}$$

where $\delta \in (0, 1)$ is a depreciation rate of capital.

111.2.1 Planning Problem

In this lecture *Cass-Koopmans Planning Model*, we studied a problem in which a planner chooses an allocation $\{\vec{C}, \vec{K}\}$ to maximize (110.2) subject to (110.5).

The allocation that solves the planning problem reappears in a competitive equilibrium, as we shall see below.

111.3 Competitive Equilibrium

We now study a decentralized version of the economy.

It shares the same technology and preference structure as the planned economy studied in this lecture *Cass-Koopmans Planning Model*.

But now there is no planner.

There are (unit masses of) price-taking consumers and firms.

Market prices are set to reconcile distinct decisions that are made separately by a representative consumer and a representative firm.

There is a representative consumer who has the same preferences over consumption plans as did a consumer in the planned economy.

Instead of being told what to consume and save by a planner, a consumer (also known as a *household*) chooses for itself subject to a budget constraint.

- At each time t, the consumer receives wages and rentals of capital from a firm – these comprise its **income** at time t.
- The consumer decides how much income to allocate to consumption or to savings.
- The household can save either by acquiring additional physical capital (it trades one for one with time t consumption) or by acquiring claims on consumption at dates other than t.
- The household owns physical capital and labor and rents them to the firm.
- The household consumes, supplies labor, and invests in physical capital.
- A profit-maximizing representative firm operates the production technology.
- The firm rents labor and capital each period from the representative household and sells its output each period to the household.
- The representative household and the representative firm are both **price takers** who believe that prices are not affected by their choices

> **Note**
>
> Again, we can think of there being unit measures of identical representative consumers and identical representative firms.

111.4 Market Structure

The representative household and the representative firm are both price takers.

The household owns both factors of production, namely, labor and physical capital.

Each period, the firm rents both factors from the household.

There is a **single** grand competitive market in which a household trades date 0 goods for goods at all other dates $t = 1, 2, \ldots, T$.

111.4.1 Prices

There are sequences of prices $\{w_t, \eta_t\}_{t=0}^T = \{\vec{w}, \vec{\eta}\}$ where

- w_t is a wage, i.e., a rental rate, for labor at time t
- η_t is a rental rate for capital at time t

In addition there is a vector $\{q_t^0\}$ of intertemporal prices where

- q_t^0 is the price at time 0 of one unit of the good at date t.

We call $\{q_t^0\}_{t=0}^T$ a vector of **Hicks-Arrow prices**, named after the 1972 economics Nobel prize winners.

Because is a **relative price**. the unit of account in terms of which the prices q_t^0 are stated is; we are free to re-normalize them by multiplying all of them by a positive scalar, say $\lambda > 0$.

Units of q_t^0 could be set so that they are

$$\frac{\text{number of time 0 goods}}{\text{number of time t goods}}$$

In this case, we would be taking the time 0 consumption good to be the **numeraire**.

111.5 Firm Problem

At time t a representative firm hires labor $\tilde{n}_t$ and capital $\tilde{k}_t$.

The firm's profits at time t are

$$F(\tilde{k}_t, \tilde{n}_t) - w_t \tilde{n}_t - \eta_t \tilde{k}_t$$

where w_t is a wage rate at t and η_t is the rental rate on capital at t.

As in the planned economy model

$$F(\tilde{k}_t, \tilde{n}_t) = A \tilde{k}_t^\alpha \tilde{n}_t^{1-\alpha}$$

111.5.1 Zero Profit Conditions

Zero-profits conditions for capital and labor are

$$F_k(\tilde{k}_t, \tilde{n}_t) = \eta_t$$

and

$$F_n(\tilde{k}_t, \tilde{n}_t) = w_t \tag{111.1}$$

These conditions emerge from a no-arbitrage requirement.

To describe this no-arbitrage profits reasoning, we begin by applying a theorem of Euler about linearly homogenous functions.

The theorem applies to the Cobb-Douglas production function because we it displays constant returns to scale:

$$\alpha F(\tilde{k}_t, \tilde{n}_t) = F(\alpha\tilde{k}_t, \alpha\tilde{n}_t)$$

for $\alpha \in (0, 1)$.

Taking partial derivatives $\frac{\partial}{\partial\alpha}$ on both sides of the above equation gives

$$F(\tilde{k}_t, \tilde{n}_t) = \frac{\partial F}{\partial \tilde{k}_t}\tilde{k}_t + \frac{\partial F}{\partial \tilde{n}_t}\tilde{n}_t$$

Rewrite the firm's profits as

$$\frac{\partial F}{\partial \tilde{k}_t}\tilde{k}_t + \frac{\partial F}{\partial \tilde{n}_t}\tilde{n}_t - w_t\tilde{n}_t - \eta_t k_t$$

or

$$\left(\frac{\partial F}{\partial \tilde{k}_t} - \eta_t\right)\tilde{k}_t + \left(\frac{\partial F}{\partial \tilde{n}_t} - w_t\right)\tilde{n}_t$$

Because F is homogeneous of degree 1, it follows that $\frac{\partial F}{\partial \tilde{k}_t}$ and $\frac{\partial F}{\partial \tilde{n}_t}$ are homogeneous of degree 0 and therefore fixed with respect to $\tilde{k}_t$ and $\tilde{n}_t$.

If $\frac{\partial F}{\partial \tilde{k}_t} > \eta_t$, then the firm makes positive profits on each additional unit of $\tilde{k}_t$, so it would want to make $\tilde{k}_t$ arbitrarily large.

But setting $\tilde{k}_t = +\infty$ is not physically feasible, so **equilibrium** prices must take values that present the firm with no such arbitrage opportunity.

A similar argument applies if $\frac{\partial F}{\partial \tilde{n}_t} > w_t$.

If $\frac{\partial \tilde{k}_t}{\partial \tilde{k}_t} < \eta_t$, the firm would want to set $\tilde{k}_t$ to zero, which is not feasible.

It is convenient to define $\vec{w} = \{w_0, \ldots, w_T\}$ and $\vec{\eta} = \{\eta_0, \ldots, \eta_T\}$.

111.6 Household Problem

A representative household lives at $t = 0, 1, \ldots, T$.

At t, the household rents 1 unit of labor and k_t units of capital to a firm and receives income

$$w_t 1 + \eta_t k_t$$

At t the household allocates its income to the following purchases between the following two categories:

- consumption c_t
- net investment $k_{t+1} - (1-\delta)k_t$

Here $(k_{t+1} - (1-\delta)k_t)$ is the household's net investment in physical capital and $\delta \in (0,1)$ is again a depreciation rate of capital.

In period t, the consumer is free to purchase more goods to be consumed and invested in physical capital than its income from supplying capital and labor to the firm, provided that in some other periods its income exceeds its purchases.

A consumer's net excess demand for time t consumption goods is the gap

$$e_t \equiv (c_t + (k_{t+1} - (1-\delta)k_t)) - (w_t 1 + \eta_t k_t)$$

Let $\vec{c} = \{c_0, \ldots, c_T\}$ and let $\vec{k} = \{k_1, \ldots, k_{T+1}\}$.

k_0 is given to the household.

The household faces a **single** budget constraint that requires that the present value of the household's net excess demands must be zero:

$$\sum_{t=0}^{T} q_t^0 e_t \leq 0$$

or

$$\sum_{t=0}^{T} q_t^0 \left(c_t + (k_{t+1} - (1-\delta)k_t)\right) \leq \sum_{t=0}^{T} q_t^0 (w_t 1 + \eta_t k_t)$$

The household faces price system $\{q_t^0, w_t, \eta_t\}$ as a price-taker and chooses an allocation to solve the constrained optimization problem:

$$\max_{\vec{c},\vec{k}} \sum_{t=0}^{T} \beta^t u(c_t)$$

$$\text{subject to } \sum_{t=0}^{T} q_t^0 \left(c_t + (k_{t+1} - (1-\delta)k_t) - (w_t - \eta_t k_t)\right) \leq 0$$

Components of a **price system** have the following units:

- w_t is measured in units of the time t good per unit of time t labor hired
- η_t is measured in units of the time t good per unit of time t capital hired
- q_t^0 is measured in units of a numeraire per unit of the time t good

111.6.1 Definitions

- A **price system** is a sequence $\{q_t^0, \eta_t, w_t\}_{t=0}^T = \{\vec{q}, \vec{\eta}, \vec{w}\}$.
- An **allocation** is a sequence $\{c_t, k_{t+1}, n_t = 1\}_{t=0}^T = \{\vec{c}, \vec{k}, \vec{n}\}$.
- A **competitive equilibrium** is a price system and an allocation with the following properties:
  - Given the price system, the allocation solves the household's problem.
  - Given the price system, the allocation solves the firm's problem.

The vision here is that an equilibrium price system and allocation are determined once and for all.

In effect, we imagine that all trades occur just before time 0.

111.7 Computing a Competitive Equilibrium

We compute a competitive equilibrium by using a **guess and verify** approach.

- We **guess** equilibrium price sequences $\{\vec{q}, \vec{\eta}, \vec{w}\}$.
- We then **verify** that at those prices, the household and the firm choose the same allocation.

111.7.1 Guess for Price System

In this lecture *Cass-Koopmans Planning Model*, we computed an allocation $\{\vec{C}, \vec{K}, \vec{N}\}$ that solves a planning problem.

We use that allocation to construct a guess for the equilibrium price system.

> **Note**
>
> This allocation will constitute the **Big** K to be in the present instance of the **Big** K **, little** k trick that we'll apply to a competitive equilibrium in the spirit of this lecture and this lecture.

In particular, we shall use the following procedure:

- obtain first-order conditions for the representative firm and the representative consumer.
- from these equations, obtain a new set of equations by replacing the firm's choice variables $\tilde{k}, \tilde{n}$ and the consumer's choice variables with the quantities $\vec{C}, \vec{K}$ that solve the planning problem.
- solve the resulting equations for $\{\vec{q}, \vec{\eta}, \vec{w}\}$ as functions of $\vec{C}, \vec{K}$.
- verify that at these prices, $c_t = C_t, k_t = \tilde{k}_t = K_t, \tilde{n}_t = 1$ for $t = 0, 1, \ldots, T$.

Thus, we guess that for $t = 0, \ldots, T$:

$$q_t^0 = \beta^t u'(C_t) \tag{111.2}$$

$$w_t = f(K_t) - K_t f'(K_t) \tag{111.3}$$

$$\eta_t = f'(K_t) \tag{111.4}$$

At these prices, let capital chosen by the household be

$$k_t^*(\vec{q}, \vec{w}, \vec{\eta}), \quad t \geq 0 \tag{111.5}$$

and let the allocation chosen by the firm be

$$\tilde{k}_t^*(\vec{q}, \vec{w}, \vec{\eta}), \quad t \geq 0$$

and so on.

If our guess for the equilibrium price system is correct, then it must occur that

$$k_t^* = \tilde{k}_t^* \tag{111.6}$$

$$1 = \tilde{n}_t^* \tag{111.7}$$

$$c_t^* + k_{t+1}^* - (1-\delta)k_t^* = F(\tilde{k}_t^*, \tilde{n}_t^*)$$

We shall verify that for $t = 0, \ldots, T$ allocations chosen by the household and the firm both equal the allocation that solves the planning problem:

$$k_t^* = \tilde{k}_t^* = K_t, \tilde{n}_t = 1, c_t^* = C_t \tag{111.8}$$

111.7.2 Verification Procedure

Our approach is firsts to stare at first-order necessary conditions for optimization problems of the household and the firm.

At the price system we have guessed, we'll then verify that both sets of first-order conditions are satisfied at the allocation that solves the planning problem.

111.7.3 Household's Lagrangian

To solve the household's problem, we formulate the Lagrangian

$$\mathcal{L}(\vec{c}, \vec{k}, \lambda) = \sum_{t=0}^{T} \beta^t u(c_t) + \lambda \left(\sum_{t=0}^{T} q_t^0 \left(((1-\delta)k_t - w_t) + \eta_t k_t - c_t - k_{t+1} \right) \right)$$

and attack the min-max problem:

$$\min_{\lambda} \max_{\vec{c}, \vec{k}} \mathcal{L}(\vec{c}, \vec{k}, \lambda)$$

First-order conditions are

$$c_t : \quad \beta^t u'(c_t) - \lambda q_t^0 = 0 \quad t = 0, 1, \ldots, T \tag{111.9}$$

$$k_t : \quad -\lambda q_t^0 \left[(1-\delta) + \eta_t \right] + \lambda q_{t-1}^0 = 0 \quad t = 1, 2, \ldots, T+1 \tag{111.10}$$

$$\lambda : \quad \left(\sum_{t=0}^{T} q_t^0 \left(c_t + (k_{t+1} - (1-\delta)k_t) - w_t - \eta_t k_t \right) \right) \leq 0 \tag{111.11}$$

$$k_{T+1} : \quad -\lambda q_0^{T+1} \leq 0, \ \leq 0 \text{ if } k_{T+1} = 0; \ = 0 \text{ if } k_{T+1} > 0 \tag{111.12}$$

Now we plug in our guesses of prices and do some algebra in the hope of recovering all first-order necessary conditions (110.9)-(110.12) for the planning problem from this lecture *Cass-Koopmans Planning Model*.

Combining (111.9) and (111.2), we get:

$$u'(C_t) = \mu_t$$

which is (110.9).

Combining (111.10), (111.2), and (111.4), we get:

$$-\lambda \beta^t \mu_t \left[(1-\delta) + f'(K_t) \right] + \lambda \beta^{t-1} \mu_{t-1} = 0 \tag{111.13}$$

Rewriting (111.13) by dividing by λ on both sides (which is nonzero since u'>0) we get:

$$\beta^t \mu_t [(1 - \delta + f'(K_t)] = \beta^{t-1} \mu_{t-1}$$

or

$$\beta \mu_t [(1 - \delta + f'(K_t)] = \mu_{t-1}$$

which is (110.10).

Combining (111.11), (111.2), (111.3) and (111.4) after multiplying both sides of (111.11) by λ, we get

$$\sum_{t=0}^{T} \beta^t \mu_t \left(C_t + (K_{t+1} - (1-\delta)K_t) - f(K_t) + K_t f'(K_t) - f'(K_t) K_t \right) \leq 0$$

which simplifies to

$$\sum_{t=0}^{T} \beta^t \mu_t \left(C_t + K_{t+1} - (1-\delta) K_t - F(K_t, 1) \right) \leq 0$$

Since $\beta^t \mu_t > 0$ for $t = 0, \ldots, T$, it follows that

$$C_t + K_{t+1} - (1-\delta) K_t - F(K_t, 1) = 0 \quad \text{for all } t \text{ in } \{0, 1, \ldots, T\}$$

which is (110.11).

Combining (111.12) and (111.2), we get:

$$-\beta^{T+1} \mu_{T+1} \leq 0$$

Dividing both sides by β^{T+1} gives

$$-\mu_{T+1} \leq 0$$

which is (110.12) for the planning problem.

Thus, at our guess of the equilibrium price system, the allocation that solves the planning problem also solves the problem faced by a representative household living in a competitive equilibrium.

111.7.4 Representative Firm's Problem

We now turn to the problem faced by a firm in a competitive equilibrium:

If we plug (111.8) into (111.1) for all t, we get

$$\frac{\partial F(K_t, 1)}{\partial K_t} = f'(K_t) = \eta_t$$

which is (111.4).

If we now plug (111.8) into (111.1) for all t, we get:

$$\frac{\partial F(\tilde{K}_t, 1)}{\partial \tilde{L}_t} = f(K_t) - f'(K_t) K_t = w_t$$

which is exactly (111.5).

Thus, at our guess for the equilibrium price system, the allocation that solves the planning problem also solves the problem faced by a firm within a competitive equilibrium.

By (111.6) and (111.7) this allocation is identical to the one that solves the consumer's problem.

Note

Because budget sets are affected only by relative prices, $\{q_t^0\}$ is determined only up to multiplication by a positive constant.

Normalization: We are free to choose a $\{q_t^0\}$ that makes $\lambda = 1$ so that we are measuring q_t^0 in units of the marginal utility of time 0 goods.

We will plot q, w, η below to show these equilibrium prices induce the same aggregate movements that we saw earlier in the planning problem.

To proceed, we bring in Python code that *Cass-Koopmans Planning Model* used to solve the planning problem

First let's define a `jitclass` that stores parameters and functions the characterize an economy.

```
planning_data = [
    ('γ', float64),     # Coefficient of relative risk aversion
    ('β', float64),     # Discount factor
    ('δ', float64),     # Depreciation rate on capital
    ('α', float64),     # Return to capital per capita
    ('A', float64)      # Technology
]
```

```
@jitclass(planning_data)
class PlanningProblem():

    def __init__(self, γ=2, β=0.95, δ=0.02, α=0.33, A=1):

        self.γ, self.β = γ, β
        self.δ, self.α, self.A = δ, α, A

    def u(self, c):
        '''
        Utility function
        ASIDE: If you have a utility function that is hard to solve by hand
        you can use automatic or symbolic differentiation
        See https://github.com/HIPS/autograd
        '''
        γ = self.γ

        return c ** (1 - γ) / (1 - γ) if γ!= 1 else np.log(c)

    def u_prime(self, c):
        'Derivative of utility'
        γ = self.γ

        return c ** (-γ)

    def u_prime_inv(self, c):
        'Inverse of derivative of utility'
        γ = self.γ

        return c ** (-1 / γ)

    def f(self, k):
        'Production function'
        α, A = self.α, self.A

        return A * k ** α

    def f_prime(self, k):
        'Derivative of production function'
        α, A = self.α, self.A

        return α * A * k ** (α - 1)

    def f_prime_inv(self, k):
        'Inverse of derivative of production function'
        α, A = self.α, self.A

        return (k / (A * α)) ** (1 / (α - 1))
```

(continues on next page)

(continued from previous page)

```
    def next_k_c(self, k, c):
        '''
        Given the current capital Kt and an arbitrary feasible
        consumption choice Ct, computes Kt+1 by state transition law
        and optimal Ct+1 by Euler equation.
        '''
        β, δ = self.β, self.δ
        u_prime, u_prime_inv = self.u_prime, self.u_prime_inv
        f, f_prime = self.f, self.f_prime

        k_next = f(k) + (1 - δ) * k - c
        c_next = u_prime_inv(u_prime(c) / (β * (f_prime(k_next) + (1 - δ))))

        return k_next, c_next
```

```
@jit
def shooting(pp, c0, k0, T=10):
    '''
    Given the initial condition of capital k0 and an initial guess
    of consumption c0, computes the whole paths of c and k
    using the state transition law and Euler equation for T periods.
    '''
    if c0 > pp.f(k0):
        print("initial consumption is not feasible")

        return None

    # initialize vectors of c and k
    c_vec = np.empty(T+1)
    k_vec = np.empty(T+2)

    c_vec[0] = c0
    k_vec[0] = k0

    for t in range(T):
        k_vec[t+1], c_vec[t+1] = pp.next_k_c(k_vec[t], c_vec[t])

    k_vec[T+1] = pp.f(k_vec[T]) + (1 - pp.δ) * k_vec[T] - c_vec[T]

    return c_vec, k_vec
```

```
@jit
def bisection(pp, c0, k0, T=10, tol=1e-4, max_iter=500, k_ter=0, verbose=True):

    # initial boundaries for guess c0
    c0_upper = pp.f(k0)
    c0_lower = 0

    i = 0
    while True:
        c_vec, k_vec = shooting(pp, c0, k0, T)
        error = k_vec[-1] - k_ter

        # check if the terminal condition is satisfied
        if np.abs(error) < tol:
            if verbose:
```

(continues on next page)

(continued from previous page)

```
                print('Converged successfully on iteration ', i+1)
            return c_vec, k_vec

        i += 1
        if i == max_iter:
            if verbose:
                print('Convergence failed.')
            return c_vec, k_vec

        # if iteration continues, updates boundaries and guess of c0
        if error > 0:
            c0_lower = c0
        else:
            c0_upper = c0

        c0 = (c0_lower + c0_upper) / 2
```

```
pp = PlanningProblem()

# Steady states
ρ = 1 / pp.β - 1
k_ss = pp.f_prime_inv(ρ+pp.δ)
c_ss = pp.f(k_ss) - pp.δ * k_ss
```

The above code from this lecture *Cass-Koopmans Planning Model* lets us compute an optimal allocation for the planning problem.

- from the preceding analysis, we know that it will also be an allocation associated with a competitive equilibium.

Now we're ready to bring in Python code that we require to compute additional objects that appear in a competitive equilibrium.

```
@jit
def q(pp, c_path):
    # Here we choose numeraire to be u'(c_0) -- this is q^(t_0)_t
    T = len(c_path) - 1
    q_path = np.ones(T+1)
    q_path[0] = 1
    for t in range(1, T+1):
        q_path[t] = pp.β ** t * pp.u_prime(c_path[t])
    return q_path

@jit
def w(pp, k_path):
    w_path = pp.f(k_path) - k_path * pp.f_prime(k_path)
    return w_path

@jit
def η(pp, k_path):
    η_path = pp.f_prime(k_path)
    return η_path
```

Now we calculate and plot for each T

```
T_arr = [250, 150, 75, 50]

fix, axs = plt.subplots(2, 3, figsize=(13, 6))
```

(continues on next page)

(continued from previous page)

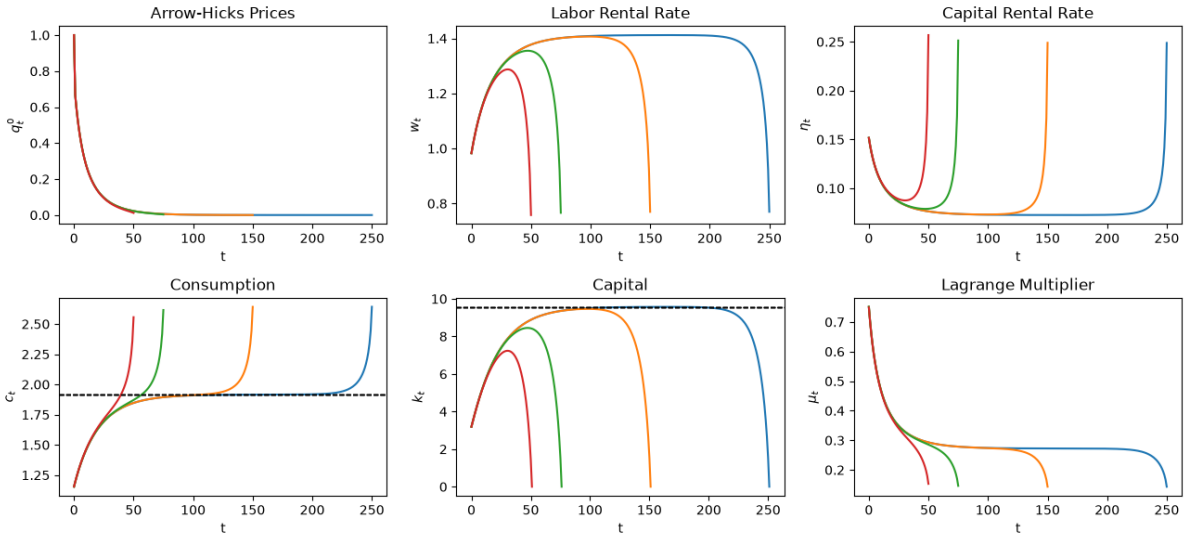
```
titles = ['Arrow-Hicks Prices', 'Labor Rental Rate', 'Capital Rental Rate',
          'Consumption', 'Capital', 'Lagrange Multiplier']
ylabels = ['$q_t^0$', '$w_t$', r'$\eta_t$', '$c_t$', '$k_t$', r'$\mu_t$']

for T in T_arr:
    c_path, k_path = bisection(pp, 0.3, k_ss/3, T, verbose=False)
    μ_path = pp.u_prime(c_path)

    q_path = q(pp, c_path)
    w_path = w(pp, k_path)[:-1]
    η_path = η(pp, k_path)[:-1]
    paths = [q_path, w_path, η_path, c_path, k_path, μ_path]

    for i, ax in enumerate(axs.flatten()):
        ax.plot(paths[i])
        ax.set(title=titles[i], ylabel=ylabels[i], xlabel='t')
        if titles[i] == 'Capital':
            ax.axhline(k_ss, lw=1, ls='--', c='k')
        if titles[i] == 'Consumption':
            ax.axhline(c_ss, lw=1, ls='--', c='k')

plt.tight_layout()
plt.show()
```



Varying Curvature

Now we see how our results change if we keep T constant, but allow the curvature parameter, γ to vary, starting with K_0 below the steady state.

We plot the results for $T = 150$

```
T = 150
γ_arr = [1.1, 4, 6, 8]

fix, axs = plt.subplots(2, 3, figsize=(13, 6))
```

(continues on next page)

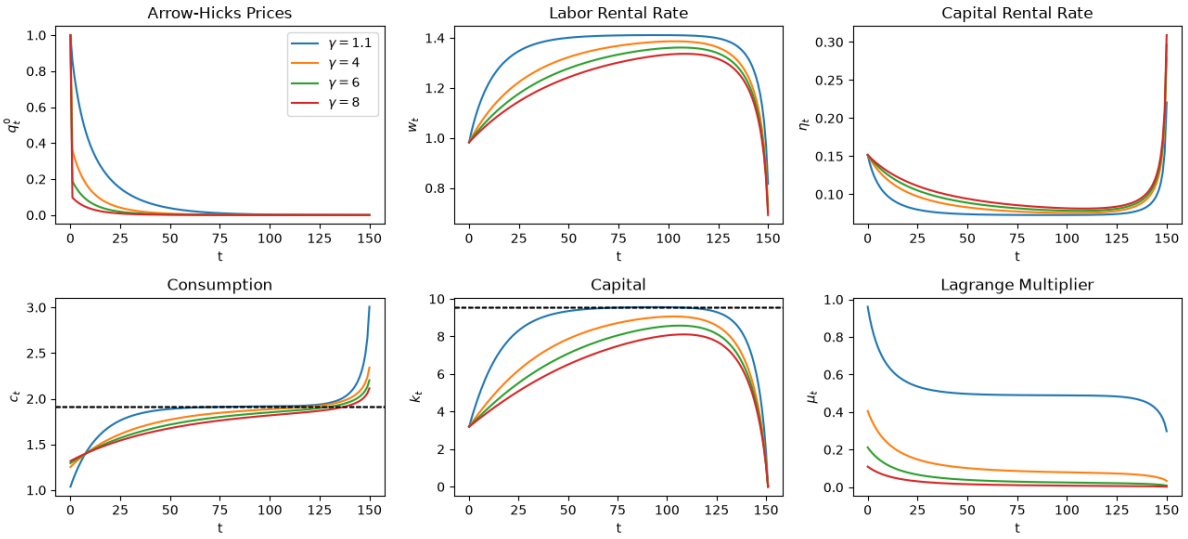
(continued from previous page)

```
for γ in γ_arr:
    pp_γ = PlanningProblem(γ=γ)
    c_path, k_path = bisection(pp_γ, 0.3, k_ss/3, T, verbose=False)
    μ_path = pp_γ.u_prime(c_path)

    q_path = q(pp_γ, c_path)
    w_path = w(pp_γ, k_path)[:-1]
    η_path = η(pp_γ, k_path)[:-1]
    paths = [q_path, w_path, η_path, c_path, k_path, μ_path]

    for i, ax in enumerate(axs.flatten()):
        ax.plot(paths[i], label=fr'$\gamma = {γ}$')
        ax.set(title=titles[i], ylabel=ylabels[i], xlabel='t')
        if titles[i] == 'Capital':
            ax.axhline(k_ss, lw=1, ls='--', c='k')
        if titles[i] == 'Consumption':
            ax.axhline(c_ss, lw=1, ls='--', c='k')

axs[0, 0].legend()
plt.tight_layout()
plt.show()
```



Adjusting γ means adjusting how much individuals prefer to smooth consumption.

Higher γ means individuals prefer to smooth more resulting in slower convergence to a steady state allocation.

Lower γ means individuals prefer to smooth less, resulting in faster convergence to a steady state allocation.

111.8 Yield Curves and Hicks-Arrow Prices

We return to Hicks-Arrow prices and calculate how they are related to **yields** on loans of alternative maturities.

This will let us plot a **yield curve** that graphs yields on bonds of maturities $j = 1, 2, ...$ against $j = 1, 2, ...$

We use the following formulas.

A **yield to maturity** on a loan made at time t_0 that matures at time $t > t_0$

$$r_{t_0,t} = -\frac{\log q_t^{t_0}}{t - t_0}$$

A Hicks-Arrow price system for a base-year $t_0 \leq t$ satisfies

$$q_t^{t_0} = \beta^{t-t_0} \frac{u'(c_t)}{u'(c_{t_0})} = \beta^{t-t_0} \frac{c_t^{-\gamma}}{c_{t_0}^{-\gamma}}$$

We redefine our function for q to allow arbitrary base years, and define a new function for r, then plot both.

We begin by continuing to assume that $t_0 = 0$ and plot things for different maturities $t = T$, with K_0 below the steady state

```
@jit
def q_generic(pp, t0, c_path):
    # simplify notations
    β = pp.β
    u_prime = pp.u_prime

    T = len(c_path) - 1
    q_path = np.zeros(T+1-t0)
    q_path[0] = 1
    for t in range(t0+1, T+1):
        q_path[t-t0] = β ** (t-t0) * u_prime(c_path[t]) / u_prime(c_path[t0])
    return q_path

@jit
def r(pp, t0, q_path):
    '''Yield to maturity'''
    r_path = - np.log(q_path[1:]) / np.arange(1, len(q_path))
    return r_path

def plot_yield_curves(pp, t0, c0, k0, T_arr):

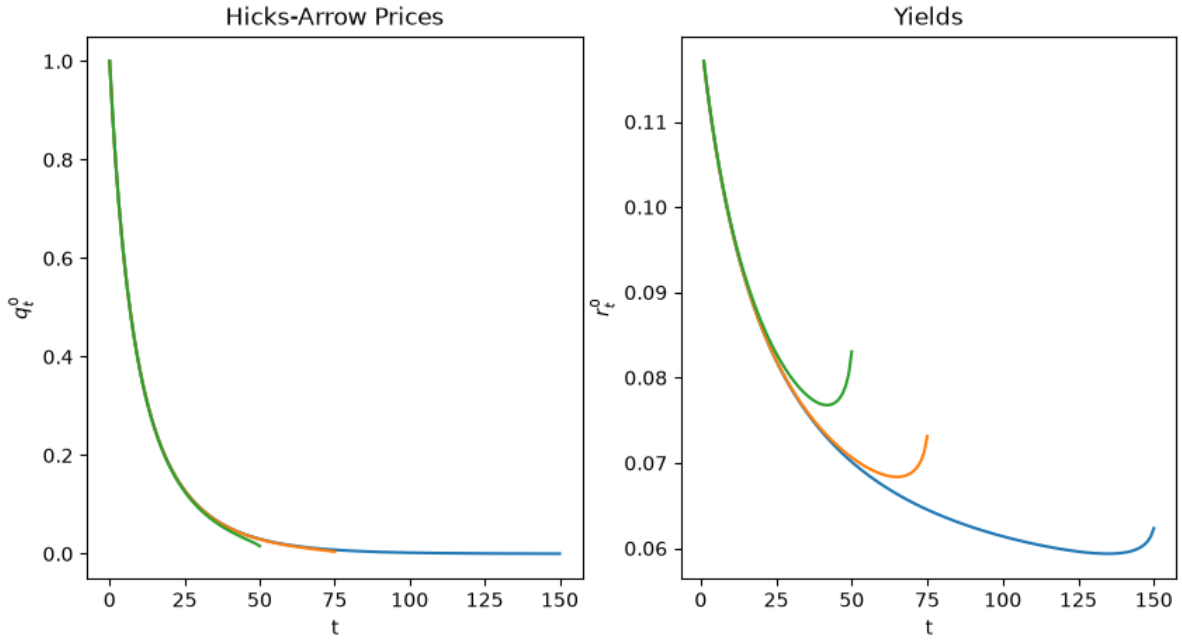
    fig, axs = plt.subplots(1, 2, figsize=(10, 5))

    for T in T_arr:
        c_path, k_path = bisection(pp, c0, k0, T, verbose=False)
        q_path = q_generic(pp, t0, c_path)
        r_path = r(pp, t0, q_path)

        axs[0].plot(range(t0, T+1), q_path)
        axs[0].set(xlabel='t', ylabel='$q_t^0$', title='Hicks-Arrow Prices')

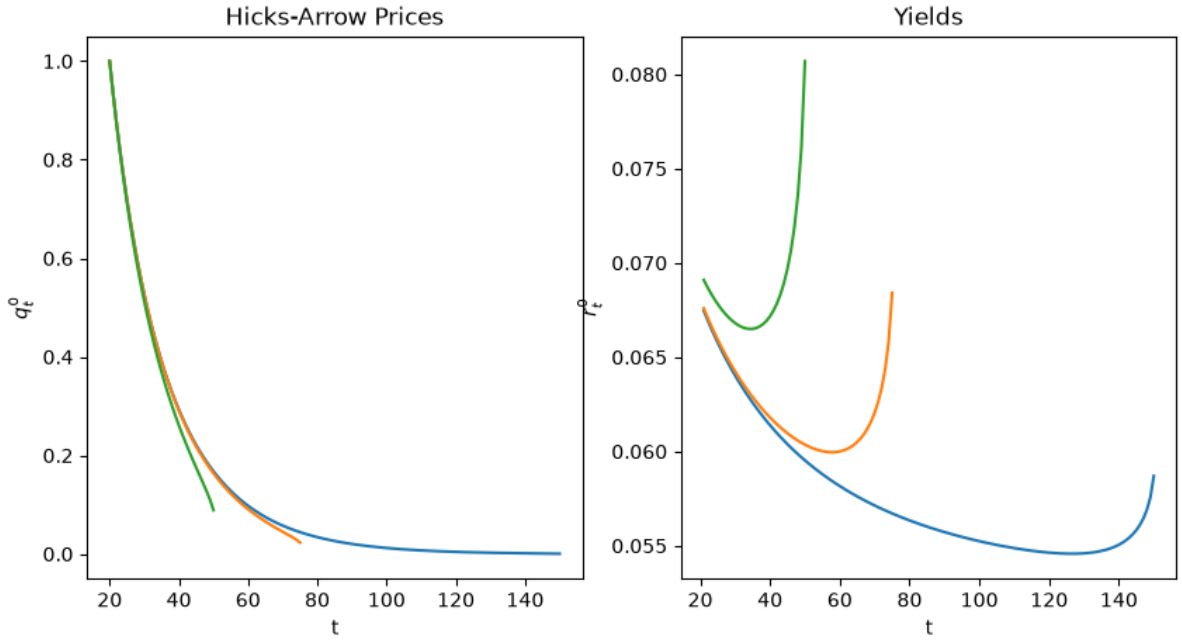
        axs[1].plot(range(t0+1, T+1), r_path)
        axs[1].set(xlabel='t', ylabel='$r_t^0$', title='Yields')
```

```
T_arr = [150, 75, 50]
plot_yield_curves(pp, 0, 0.3, k_ss/3, T_arr)
```



Now we plot when $t_0 = 20$

```
plot_yield_curves(pp, 20, 0.3, k_ss/3, T_arr)
```



CHAPTER

TWO

CASS-KOOPMANS MODEL WITH DISTORTING TAXES

112.1 Overview

This lecture studies effects of foreseen fiscal and technology shocks on competitive equilibrium prices and quantities in a nonstochastic version of a Cass-Koopmans growth model with features described in this QuantEcon lecture *Cass-Koopmans Competitive Equilibrium*.

This model is discussed in more detail in chapter 11 of [Ljungqvist and Sargent, 2018].

We use the model as a laboratory to experiment with numerical techniques for approximating equilibria and to display the structure of dynamic models in which decision makers have perfect foresight about future government decisions.

Following a classic paper by Robert E. Hall [Hall, 1971], we augment a nonstochastic version of the Cass-Koopmans optimal growth model with a government that purchases a stream of goods and that finances its purchases with an sequences of several distorting flat-rate taxes.

Distorting taxes prevent a competitive equilibrium allocation from solving a planning problem.

Therefore, to compute an equilibrium allocation and price system, we solve a system of nonlinear difference equations consisting of the first-order conditions for decision makers and the other equilibrium conditions.

We present two ways to approximate an equilibrium:

- The first is a shooting algorithm like the one that we deployed in *Cass-Koopmans Competitive Equilibrium*.
- The second method is a root-finding algorithm that minimizes residuals from the first-order conditions of the consumer and representative firm.

112.2 The Economy

112.2.1 Technology

Feasible allocations satisfy

$$g_t + c_t + x_t \leq F(k_t, n_t), \tag{112.1}$$

where

- g_t is government purchases of the time t good
- x_t is gross investment, and
- $F(k_t, n_t)$ is a linearly homogeneous production function with positive and decreasing marginal products of capital k_t and labor n_t.

Physical capital evolves according to

$$k_{t+1} = (1-\delta)k_t + x_t,$$

where $\delta \in (0,1)$ is a depreciation rate.

It is sometimes convenient to eliminate x_t from (112.1) and to represent it as as

$$g_t + c_t + k_{t+1} \leq F(k_t, n_t) + (1-\delta)k_t.$$

112.2.2 Components of a competitive equilibrium

All trades occurring at time 0.

The representative household owns capital, makes investment decisions, and rents capital and labor to a representative production firm.

The representative firm uses capital and labor to produce goods with the production function $F(k_t, n_t)$.

A **price system** is a triple of sequences $\{q_t, \eta_t, w_t\}_{t=0}^{\infty}$, where

- q_t is the time 0 pretax price of one unit of investment or consumption at time t (x_t or c_t),
- η_t is the pretax price at time t that the household receives from the firm for renting capital at time t, and
- w_t is the pretax price at time t that the household receives for renting labor to the firm at time t.

The prices w_t and η_t are expressed in terms of time t goods, while q_t is expressed in terms of a numeraire at time 0, as in *Cass-Koopmans Competitive Equilibrium*.

The presence of a government distinguishes this lecture from *Cass-Koopmans Competitive Equilibrium*.

Government purchases of goods at time t are $g_t \geq 0$.

A government expenditure plan is a sequence $g = \{g_t\}_{t=0}^{\infty}$.

A government tax plan is a 4-tuple of sequences $\{\tau_{ct}, \tau_{kt}, \tau_{nt}, \tau_{ht}\}_{t=0}^{\infty}$, where

- τ_{ct} is a tax rate on consumption at time t,
- τ_{kt} is a tax rate on rentals of capital at time t,
- τ_{nt} is a tax rate on wage earnings at time t, and
- τ_{ht} is a lump sum tax on a consumer at time t.

Because lump-sum taxes τ_{ht} are available, the government actually should not use any distorting taxes.

Nevertheless, we include all of these taxes because, like [Hall, 1971], they allow us to analyze how various taxes distort production and consumption decisions.

In the *experiment section*, we shall see how variations in government tax plan affect the transition path and equilibrium.

112.2.3 Representative Household

A representative household has preferences over nonnegative streams of a single consumption good c_t and leisure $1 - n_t$ that are ordered by:

$$\sum_{t=0}^{\infty} \beta^t U(c_t, 1-n_t), \quad \beta \in (0,1), \tag{112.2}$$

where

- U is strictly increasing in c_t, twice continuously differentiable, and strictly concave with $c_t \geq 0$ and $n_t \in [0, 1]$.

The representative hßousehold maximizes (112.2) subject to the single budget constraint:

$$\begin{aligned}\sum_{t=0}^{\infty} q_t \left\{ (1+\tau_{ct})c_t + \underbrace{[k_{t+1} - (1-\delta)k_t]}_{\text{no tax when investing}} \right\} \\ \leq \sum_{t=0}^{\infty} q_t \left\{ \eta_t k_t - \underbrace{\tau_{kt}(\eta_t - \delta)k_t}_{\text{tax on rental return}} + (1-\tau_{nt})w_t n_t - \tau_{ht} \right\}.\end{aligned} \tag{112.3}$$

Here we have assumed that the government gives a depreciation allowance δk_t from the gross rentals on capital $\eta_t k_t$ and so collects taxes $\tau_{kt}(\eta_t - \delta)k_t$ on rentals from capital.

112.2.4 Government

Government plans $\{g_t\}_{t=0}^{\infty}$ for government purchases and taxes $\{\tau_{ct}, \tau_{kt}, \tau_{nt}, \tau_{ht}\}_{t=0}^{\infty}$ must respect the budget constraint

$$\sum_{t=0}^{\infty} q_t g_t \leq \sum_{t=0}^{\infty} q_t \left\{ \tau_{ct} c_t + \tau_{kt}(\eta_t - \delta)k_t + \tau_{nt} w_t n_t + \tau_{ht} \right\}. \tag{112.4}$$

Given a budget-feasible government policy $\{g_t\}_{t=0}^{\infty}$ and $\{\tau_{ct}, \tau_{kt}, \tau_{nt}, \tau_{ht}\}_{t=0}^{\infty}$ subject to (112.4),

- **Household** chooses $\{c_t\}_{t=0}^{\infty}$, $\{n_t\}_{t=0}^{\infty}$, and $\{k_{t+1}\}_{t=0}^{\infty}$ to maximize utility(112.2) subject to budget constraint(112.3), and
- **Firm** chooses sequences of capital $\{k_t\}_{t=0}^{\infty}$ and $\{n_t\}_{t=0}^{\infty}$ to maximize profits

$$\sum_{t=0}^{\infty} q_t [F(k_t, n_t) - \eta_t k_t - w_t n_t] \tag{112.5}$$

- A **feasible allocation** is a sequence $\{c_t, x_t, n_t, k_t\}_{t=0}^{\infty}$ that satisfies feasibility condition (112.1).

112.3 Equilibrium

Definition 112.3.1

A **competitive equilibrium with distorting taxes** is a **budget-feasible government policy**, **a feasible allocation**, and **a price system** for which, given the price system and the government policy, the allocation solves the household's problem and the firm's problem.

112.4 No-arbitrage Condition

A no-arbitrage argument implies a restriction on prices and tax rates across time.

By rearranging (112.3) and group k_t at the same t, we can get

$$\begin{aligned}\sum_{t=0}^{\infty} q_t \left[(1+\tau_{ct})c_t\right] \leq \sum_{t=0}^{\infty} q_t (1-\tau_{nt}) w_t n_t - \sum_{t=0}^{\infty} q_t \tau_{ht} \\ + \sum_{t=1}^{\infty} \left\{ \left[(1-\tau_{kt})(\eta_t - \delta) + 1\right] q_t - q_{t-1} \right\} k_t \\ + \left[(1-\tau_{k0})(\eta_0 - \delta) + 1\right] q_0 k_0 - \lim_{T \to \infty} q_T k_{T+1}\end{aligned} \tag{112.6}$$

The household inherits a given k_0 that it takes as initial condition and is free to choose $\{c_t, n_t, k_{t+1}\}_{t=0}^{\infty}$.

The household's budget constraint (112.3) must be bounded in equilibrium due to finite resources.

This imposes a restriction on price and tax sequences.

Specifically, for $t \geq 1$, the terms multiplying k_t must equal zero.

If they were strictly positive (negative), the household could arbitrarily increase (decrease) the right-hand side of (112.3) by selecting an arbitrarily large positive (negative) k_t, leading to unbounded profit or arbitrage opportunities:

- For strictly positive terms, the household could purchase large capital stocks k_t and profit from their rental services and undepreciated value.
- For strictly negative terms, the household could engage in "short selling" synthetic units of capital. Both cases would make (112.3) unbounded.

Hence, by setting the terms multiplying k_t to 0 we have the non-arbitrage condition:

$$\frac{q_t}{q_{t+1}} = \left[(1 - \tau_{kt+1})(\eta_{t+1} - \delta) + 1\right]. \tag{112.7}$$

Moreover, we have terminal condition:

$$-\lim_{T \to \infty} q_T k_{T+1} = 0. \tag{112.8}$$

Zero-profit conditions for the representative firm impose additional restrictions on equilibrium prices and quantities.

The present value of the firm's profits is

$$\sum_{t=0}^{\infty} q_t \left[F(k_t, n_t) - w_t n_t - \eta_t k_t\right].$$

Applying Euler's theorem on linearly homogeneous functions to $F(k, n)$, the firm's present value is:

$$\sum_{t=0}^{\infty} q_t \left[(F_{kt} - \eta_t)k_t + (F_{nt} - w_t)n_t\right].$$

No-arbitrage (or zero-profit) conditions are:

$$\eta_t = F_{kt}, \quad w_t = F_{nt}. \tag{112.9}$$

112.5 Household's First Order Condition

Household maximize (112.2) under (112.3).

Let $U_1 = \frac{\partial U}{\partial c}, U_2 = \frac{\partial U}{\partial (1-n)} = -\frac{\partial U}{\partial n}$., we can derive FOC from the Lagrangian

$$\mathcal{L} = \sum_{t=0}^{\infty} \beta^t U(c_t, 1 - n_t) + \mu \left(\sum_{t=0}^{\infty} q_t \left[(1 + \tau_{ct})c_t - (1 - \tau_{nt})w_t n_t + ...\right] \right),$$

First-order necessary conditions for the representative household's problem are

$$\frac{\partial \mathcal{L}}{\partial c_t} = \beta^t U_1(c_t, 1 - n_t) - \mu q_t (1 + \tau_{ct}) = 0 \tag{112.10}$$

and

$$\frac{\partial \mathcal{L}}{\partial n_t} = \beta^t \left(-U_{2t}(c_t, 1 - n_t)\right) - \mu q_t (1 - \tau_{nt}) w_t = 0 \tag{112.11}$$

Rearranging (112.10) and (112.11), we have

$$\beta^t U_1(c_t, 1-n_t) = \beta^t U_{1t} = \mu q_t (1+\tau_{ct}), \tag{112.12}$$

$$\beta^t U_2(c_t, 1-n_t) = \beta^t U_{2t} = \mu q_t (1-\tau_{nt}) w_t. \tag{112.13}$$

Plugging (112.12) into (112.8) and replacing q_t, we get terminal condition

$$-\lim_{T\to\infty} \beta^T \frac{U_{1T}}{(1+\tau_{cT})} k_{T+1} = 0. \tag{112.14}$$

112.6 Computing Equilibria

To compute an equilibrium, we seek a price system $\{q_t, \eta_t, w_t\}$, a budget feasible government policy $\{g_t, \tau_t\} \equiv \{g_t, \tau_{ct}, \tau_{nt}, \tau_{kt}, \tau_{ht}\}$, and an allocation $\{c_t, n_t, k_{t+1}\}$ that solve a system of nonlinear difference equations consisting of

- feasibility condition (112.1), no-arbitrage condition for household (112.7) and firms (112.9), household's first order conditions (112.12) and (112.13).
- an initial condition k_0 and a terminal condition (112.14).

112.7 Python Code

We require the following imports

```
import numpy as np
from scipy.optimize import root
import matplotlib.pyplot as plt
from collections import namedtuple
from mpmath import mp, mpf
from warnings import warn

# Set the precision
mp.dps = 40
mp.pretty = True
```

We use the `mpmath` library to perform high-precision arithmetic in the shooting algorithm in cases where the solution diverges due to numerical instability.

Note

In the functions below, we include routines to handle the growth component, which will be discussed further in the section *Exogenous growth*.

We include them here to avoid code duplication.

We set the following parameters

```
# Create a namedtuple to store the model parameters
Model = namedtuple("Model",
                   ["β", "γ", "δ", "α", "A"])
```

(continues on next page)

(continued from previous page)

```
def create_model(β=0.95,  # discount factor
                 γ=2.0,   # relative risk aversion coefficient
                 δ=0.2,   # depreciation rate
                 α=0.33,  # capital share
                 A=1.0    # TFP
                 ):
    """Create a model instance."""
    return Model(β=β, γ=γ, δ=δ, α=α, A=A)

model = create_model()

# Total number of periods
S = 100
```

112.7.1 Inelastic Labor Supply

In this lecture, we consider the special case where $U(c, 1-n) = u(c)$ and $f(k) := F(k, 1)$.

We rewrite (112.1) with $f(k) := F(k, 1)$,

$$k_{t+1} = f(k_t) + (1-\delta)k_t - g_t - c_t. \tag{112.15}$$

```
def next_k(k_t, g_t, c_t, model, μ_t=1):
    """
    Capital next period: k_{t+1} = f(k_t) + (1 - δ) * k_t - c_t - g_t
    with optional growth adjustment: k_{t+1} = (f(k_t) + (1 - δ) * k_t - c_t - g_t) /␣
↪μ_{t+1}
    """
    return (f(k_t, model) + (1 - model.δ) * k_t - g_t - c_t) / μ_t
```

By the properties of a linearly homogeneous production function, we have $F_k(k, n) = f'(k)$ and $F_n(k, 1) = f(k, 1) - f'(k)k$.

Substituting (112.12), (112.9), and (112.15) into (112.7), we obtain the Euler equation

$$\begin{aligned}
&\frac{u'(f(k_t) + (1-\delta)k_t - g_t - k_{t+1})}{(1+\tau_{ct})} \\
&- \beta \frac{u'(f(k_{t+1}) + (1-\delta)k_{t+1} - g_{t+1} - k_{t+2})}{(1+\tau_{ct+1})} \\
&\times [(1-\tau_{kt+1})(f'(k_{t+1}) - \delta) + 1] = 0.
\end{aligned} \tag{112.16}$$

This can be simplified to:

$$u'(c_t) = \beta u'(c_{t+1}) \frac{(1+\tau_{ct})}{(1+\tau_{ct+1})} [(1-\tau_{kt+1})(f'(k_{t+1}) - \delta) + 1]. \tag{112.17}$$

Equation (112.17) will appear prominently in our equilibrium computation algorithms.

112.7.2 Steady state

Tax rates and government expenditures act as **forcing functions** for difference equations (112.15) and (112.17).

Define $z_t = [g_t, \tau_{kt}, \tau_{ct}]'$.

Represent the second-order difference equation as:

$$H(k_t, k_{t+1}, k_{t+2}; z_t, z_{t+1}) = 0. \tag{112.18}$$

We assume that a government policy reaches a steady state such that $\lim_{t\to\infty} z_t = \bar{z}$ and that the steady state prevails for $t > T$.

The terminal steady-state capital stock $\bar{k}$ satisfies:

$$H(\bar{k}, \bar{k}, \bar{k}, \bar{z}, \bar{z}) = 0.$$

From difference equation (112.17), we can infer a restriction on the steady-state:

$$\begin{aligned} u'(\bar{c}) &= \beta u'(\bar{c}) \frac{(1+\bar{\tau}_c)}{(1+\bar{\tau}_c)}[(1-\bar{\tau}_k)(f'(\bar{k}) - \delta) + 1]. \\ \implies 1 &= \beta[(1-\bar{\tau}_k)(f'(\bar{k}) - \delta) + 1]. \end{aligned} \tag{112.19}$$

112.7.3 Other equilibrium quantities and prices

Price:

$$q_t = \frac{\beta^t u'(c_t)}{u'(c_0)} \tag{112.20}$$

```
def compute_q_path(c_path, model, S=100, A_path=None):
    """
    Compute q path: q_t = (β^t * u'(c_t)) / u'(c_0)
    with optional A_path for growth models.
    """
    A = np.ones_like(c_path) if A_path is None else np.asarray(A_path)
    q_path = np.zeros_like(c_path)
    for t in range(S):
        q_path[t] = (model.β ** t *
                     u_prime(c_path[t], model, A[t])) / u_prime(c_path[0], model,␣
↪A[0])
    return q_path
```

Capital rental rate

$$\eta_t = f'(k_t)$$

```
def compute_η_path(k_path, model, S=100, A_path=None):
    """
    Compute η path: η_t = f'(k_t)
    with optional A_path for growth models.
    """
    A = np.ones_like(k_path) if A_path is None else np.asarray(A_path)
    η_path = np.zeros_like(k_path)
    for t in range(S):
        η_path[t] = f_prime(k_path[t], model, A[t])
    return η_path
```

Labor rental rate:

$$w_t = f(k_t) - k_t f'(k_t)$$

```
def compute_w_path(k_path, η_path, model, S=100, A_path=None):
    """
    Compute w path: w_t = f(k_t) - k_t * f'(k_t)
    with optional A_path for growth models.
    """
    A = np.ones_like(k_path) if A_path is None else np.asarray(A_path)
    w_path = np.zeros_like(k_path)
    for t in range(S):
        w_path[t] = f(k_path[t], model, A[t]) - k_path[t] * η_path[t]
    return w_path
```

Gross one-period return on capital:

$$\bar{R}_{t+1} = \frac{(1+\tau_{ct})}{(1+\tau_{ct+1})}\left[(1-\tau_{kt+1})(f'(k_{t+1})-\delta)+1\right] = \frac{(1+\tau_{ct})}{(1+\tau_{ct+1})}R_{t,t+1} \tag{112.21}$$

```
def compute_R_bar(τ_ct, τ_ctp1, τ_ktp1, k_tp1, model):
    """
    Gross one-period return on capital:
    R_bar = [(1 + τ_c_t) / (1 + τ_c_{t+1})]
        * { [1 - τ_k_{t+1}] * [f'(k_{t+1}) - δ] + 1 }
    """
    return ((1 + τ_ct) / (1 + τ_ctp1)) * (
        (1 - τ_ktp1) * (f_prime(k_tp1, model) - model.δ) + 1)
```

```
def compute_R_bar_path(shocks, k_path, model, S=100):
    """
    Compute R_bar path over time.
    """
    R_bar_path = np.zeros(S + 1)
    for t in range(S):
        R_bar_path[t] = compute_R_bar(
            shocks['τ_c'][t], shocks['τ_c'][t + 1], shocks['τ_k'][t + 1],
            k_path[t + 1], model)
    R_bar_path[S] = R_bar_path[S - 1]
    return R_bar_path
```

One-period discount factor:

$$R_{t,t+1}^{-1} = \frac{q_{t+1}}{q_t} = m_{t,t+1} = \beta\frac{u'(c_{t+1})}{u'(c_t)}\frac{(1+\tau_{ct})}{(1+\tau_{ct+1})} \tag{112.22}$$

Net one-period rate of interest:

$$r_{t,t+1} \equiv R_{t,t+1} - 1 = (1-\tau_{k,t+1})(f'(k_{t+1})-\delta) \tag{112.23}$$

By (112.22) and $r_{t,t+1} = -\ln(\frac{q_{t+1}}{q_t})$, we have

$$R_{t,t+s} = e^{s\cdot r_{t,t+s}}.$$

Then by (112.23), we have

$$\frac{q_{t+s}}{q_t} = e^{-s\cdot r_{t,t+s}}.$$

Rearranging the above equation, we have

$$r_{t,t+s} = -\frac{1}{s} \ln \left(\frac{q_{t+s}}{q_t} \right).$$

```
def compute_rts_path(q_path, S, t):
    """
    Compute r path:
    r_t,t+s = - (1/s) * ln(q_{t+s} / q_t)
    """
    s = np.arange(1, S + 1)
    q_path = np.array([float(q) for q in q_path])

    with np.errstate(divide='ignore', invalid='ignore'):
        rts_path = - np.log(q_path[t + s] / q_path[t]) / s
    return rts_path
```

112.8 Some functional forms

We assume that the representative household' period utility has the following CRRA (constant relative risk aversion) form

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}$$

```
def u_prime(c, model, A_t=1):
    """
    Marginal utility: u'(c) = c^{-γ}
    with optional technology adjustment: u'(cA) = (cA)^{-γ}
    """
    return (c * A_t) ** (-model.γ)
```

By substituting (112.21) into (112.17), we obtain

$$c_{t+1} = c_t \left[\beta \frac{(1+\tau_{ct})}{(1+\tau_{ct+1})} \left[(1-\tau_{k,t+1})(f'(k_{t+1}) - \delta) + 1 \right] \right]^{\frac{1}{\gamma}} = c_t \left[\beta \overline{R}_{t+1} \right]^{\frac{1}{\gamma}} \tag{112.24}$$

```
def next_c(c_t, R_bar, model, μ_t=1):
    """
    Consumption next period: c_{t+1} = c_t * (β * R̄)^{1/γ}
    with optional growth adjustment: c_{t+1} = c_t * (β * R_bar)^{1/γ} * μ_{t+1}^{-1}
    """
    return c_t * (model.β * R_bar) ** (1 / model.γ) / μ_t
```

For the production function we assume a Cobb-Douglas form:

$$F(k, 1) = Ak^{\alpha}$$

```
def f(k, model, A=1):
    """
    Production function: f(k) = A * k^{α}
    """
    return A * k ** model.α

def f_prime(k, model, A=1):
```

(continues on next page)

(continued from previous page)

```
    """
    Marginal product of capital: f'(k) = α * A * k^{α - 1}
    """
    return model.α * A * k ** (model.α - 1)
```

112.9 Computation

We describe two ways to compute an equilibrium:

- a shooting algorithm
- a residual-minimization method that focuses on imposing Euler equation (112.17) and the feasibility condition (112.15).

112.9.1 Shooting Algorithm

This algorithm deploys the following steps.

1. Solve the equation (112.19) for the terminal steady-state capital $\bar{k}$ that corresponds to the permanent policy vector $\bar{z}$.
2. Select a large time index $S \gg T$, guess an initial consumption rate c_0, and use the equation (112.15) to solve for k_1.
3. Use the equation (112.24) to determine c_{t+1}. Then, apply the equation (112.15) to compute k_{t+2}.
4. Iterate step 3 to compute candidate values $\hat{k}_t$ for $t = 1, \ldots, S$.
5. Compute the difference $\hat{k}_S - \bar{k}$. If $\left|\hat{k}_S - \bar{k}\right| > \epsilon$ for some small ϵ, adjust c_0 and repeat steps 2-5.
6. Adjust c_0 iteratively using the bisection method to find a value that ensures $\left|\hat{k}_S - \bar{k}\right| < \epsilon$.

The following code implements these steps.

```
# Steady-state calculation
def steady_states(model, g_ss, τ_k_ss=0.0, μ_ss=None):
    """
    Calculate steady state values for capital and
    consumption with optional A_path for growth models.
    """

    β, δ, α, γ = model.β, model.δ, model.α, model.γ

    A = model.A or 1.0

    # growth-adjustment in the numerator: μ^γ or 1
    μ_eff = μ_ss**γ if μ_ss is not None else 1.0

    num = δ + (μ_eff/β - 1) / (1 - τ_k_ss)
    k_ss = (num / (α * A)) ** (1 / (α - 1))

    c_ss = (
        A * k_ss**α - δ * k_ss - g_ss
        if μ_ss is None
        else k_ss**α + (1 - δ - μ_ss) * k_ss - g_ss
```

(continues on next page)

(continued from previous page)

```
    )

    return k_ss, c_ss

def shooting_algorithm(
    c0, k0, shocks, S, model, A_path=None):
    """
    Shooting algorithm for given initial c0 and k0
    with optional A_path for growth models.
    """
    # unpack & mpf-ify shocks, fill μ with ones if missing
    g = np.array(list(map(mpf, shocks['g'])), dtype=object)
    τ_c = np.array(list(map(mpf, shocks['τ_c'])), dtype=object)
    τ_k = np.array(list(map(mpf, shocks['τ_k'])), dtype=object)
    μ = (np.array(list(map(mpf, shocks['μ'])), dtype=object)
             if 'μ' in shocks else np.ones_like(g))
    A = np.ones_like(g) if A_path is None else A_path

    k_path = np.empty(S+1, dtype=object)
    c_path = np.empty(S+1, dtype=object)
    k_path[0], c_path[0] = mpf(k0), mpf(c0)

    for t in range(S):
        k_t, c_t = k_path[t], c_path[t]
        k_tp1 = next_k(k_t, g[t], c_t, model, μ[t+1])
        if k_tp1 < 0:
            return None, None
        k_path[t+1] = k_tp1

        R_bar = compute_R_bar(
            τ_c[t], τ_c[t+1], τ_k[t+1], k_tp1, model
        )
        c_tp1 = next_c(c_t, R_bar, model, μ[t+1])
        if c_tp1 < 0:
            return None, None
        c_path[t+1] = c_tp1

    return k_path, c_path


def bisection_c0(
    c0_guess, k0, shocks, S, model, tol=mpf('1e-6'),
    max_iter=1000, verbose=False, A_path=None):
    """
    Bisection method to find initial c0
    """
    # steady-state uses last shocks (μ=1 if missing)
    g_last    = mpf(shocks['g'][-1])
    τ_k_last  = mpf(shocks['τ_k'][-1])
    μ_last    = mpf(shocks['μ'][-1]) if 'μ' in shocks else mpf('1')
    k_ss_fin, _ = steady_states(model, g_last, τ_k_last, μ_last)

    c0_lo, c0_hi = mpf('0'), f(k_ss_fin, model)
    c0 = mpf(c0_guess)

    for i in range(1, max_iter+1):
        k_path, _ = shooting_algorithm(c0, k0, shocks, S, model, A_path)
```

(continues on next page)

(continued from previous page)

```
        if k_path is None:
            if verbose:
                print(f"[{i}] shoot failed at c0={c0}")
            c0_hi = c0
        else:
            err = k_path[-1] - k_ss_fin
            if verbose and i % 100 == 0:
                print(f"[{i}] c0={c0}, err={err}")
            if abs(err) < tol:
                if verbose:
                    print(f"Converged after {i} iter")
                return c0
            # update bounds in one line
            c0_lo, c0_hi = (c0, c0_hi) if err > 0 else (c0_lo, c0)
        c0 = (c0_lo + c0_hi) / mpf('2')

    warn(f"bisection did not converge after {max_iter} iters; returning c0={c0}")
    return c0


def run_shooting(
    shocks, S, model, A_path=None,
    c0_finder=bisection_c0, shooter=shooting_algorithm):
    """
    Compute initial SS, find c0, and return [k,c] paths
    with optional A_path for growth models.
    """
    # initial SS at t=0 (μ=1 if missing)
    g0    = mpf(shocks['g'][0])
    τ_k0  = mpf(shocks['τ_k'][0])
    μ0    = mpf(shocks['μ'][0]) if 'μ' in shocks else mpf('1')
    k0, c0 = steady_states(model, g0, τ_k0, μ0)

    optimal_c0 = c0_finder(c0, k0, shocks, S, model, A_path=A_path)
    print(f"Model: {model}\nOptimal initial consumption c0 = {mpf(optimal_c0)}")

    k_path, c_path = shooter(optimal_c0, k0, shocks, S, model, A_path)
    return np.column_stack([k_path, c_path])
```

112.9.2 Experiments

Let's run some experiments.

1. A foreseen once-and-for-all increase in g from 0.2 to 0.4 occurring in period 10,
2. A foreseen once-and-for-all increase in τ_c from 0.0 to 0.2 occurring in period 10,
3. A foreseen once-and-for-all increase in τ_k from 0.0 to 0.2 occurring in period 10, and
4. A foreseen one-time increase in g from 0.2 to 0.4 in period 10, after which g reverts to 0.2 permanently.

To start, we prepare sequences that we'll used to initialize our iterative algorithm.

We will start from an initial steady state and apply shocks at an the indicated time.

```
def plot_results(
    solution, k_ss, c_ss, shocks, shock_param, axes, model,
```

(continues on next page)

(continued from previous page)

```
    A_path=None, label='', linestyle='-', T=40):
    """
    Plot simulation results (k, c, R, η, and a policy shock)
    with optional A_path for growth models.
    """
    k_path = solution[:, 0]
    c_path = solution[:, 1]
    T = min(T, k_path.size)

    # handle growth parameters
    μ0 = shocks['μ'][0] if 'μ' in shocks else 1.0
    A0 = A_path[0] if A_path is not None else (model.A or 1.0)

    # steady-state lines
    R_bar_ss = (1 / model.β) * (μ0**model.γ)
    η_ss     = model.α * A0 * k_ss**(model.α - 1)

    # plot k
    axes[0].plot(k_path[:T], linestyle=linestyle, label=label)
    axes[0].axhline(k_ss, linestyle='--', color='black')
    axes[0].set_title('k')

    # plot c
    axes[1].plot(c_path[:T], linestyle=linestyle, label=label)
    axes[1].axhline(c_ss, linestyle='--', color='black')
    axes[1].set_title('c')

    # plot R bar
    S_full    = k_path.size - 1
    R_bar_path = compute_R_bar_path(shocks, k_path, model, S_full)
    axes[2].plot(R_bar_path[:T], linestyle=linestyle, label=label)
    axes[2].axhline(R_bar_ss, linestyle='--', color='black')
    axes[2].set_title(r'$\bar{R}$')

    # plot η
    η_path = compute_η_path(k_path, model, S_full)
    axes[3].plot(η_path[:T], linestyle=linestyle, label=label)
    axes[3].axhline(η_ss, linestyle='--', color='black')
    axes[3].set_title(r'$\eta$')

    # plot shock
    shock_series = np.array(shocks[shock_param], dtype=object)
    axes[4].plot(shock_series[:T], linestyle=linestyle, label=label)
    axes[4].axhline(shock_series[0], linestyle='--', color='black')
    axes[4].set_title(rf'${shock_param}$')

    if label:
        for ax in axes[:5]:
            ax.legend()
```

Experiment 1: Foreseen once-and-for-all increase in g from 0.2 to 0.4 in period 10

The figure below shows consequences of a foreseen permanent increase in g at $t = T = 10$ that is financed by an increase in lump-sum taxes

```
# Define shocks as a dictionary
shocks = {
```

(continues on next page)

(continued from previous page)

```
    'g': np.concatenate(
        (np.repeat(0.2, 10), np.repeat(0.4, S - 9))
    ),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1)
}

k_ss_initial, c_ss_initial = steady_states(model,
                                           shocks['g'][0],
                                           shocks['τ_k'][0])

print(f"Steady-state capital: {k_ss_initial:.4f}")
print(f"Steady-state consumption: {c_ss_initial:.4f}")

solution = run_shooting(shocks, S, model)

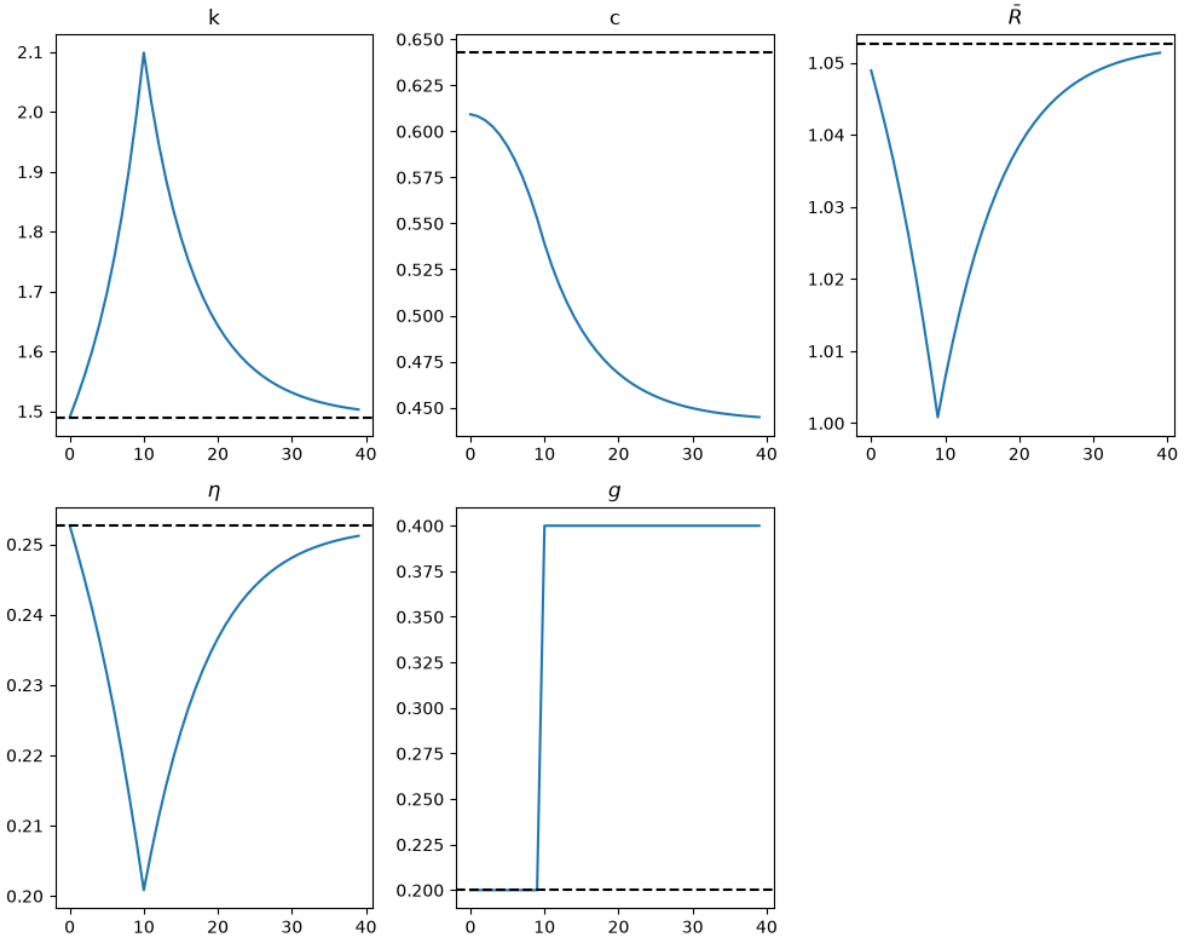
fig, axes = plt.subplots(2, 3, figsize=(10, 8))
axes = axes.flatten()

plot_results(solution, k_ss_initial,
             c_ss_initial, shocks, 'g', axes, model, T=40)

for ax in axes[5:]:
    fig.delaxes(ax)

plt.tight_layout()
plt.show()
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.6092419528879239645312185699727132533517
```



The above figures indicate that an equilibrium **consumption smoothing** mechanism is at work, driven by the representative consumer's preference for smooth consumption paths coming from the curvature of its one-period utility function.

- The steady-state value of the capital stock remains unaffected:
  - This follows from the fact that g disappears from the steady state version of the Euler equation ((112.19)).
- Consumption begins to decline gradually before time T due to increased government consumption:
  - Households reduce consumption to offset government spending, which is financed through increased lump-sum taxes.
  - The competitive economy signals households to consume less through an increase in the stream of lump-sum taxes.
  - Households, caring about the present value rather than the timing of taxes, experience an adverse wealth effect on consumption, leading to an immediate response.
- Capital gradually accumulates between time 0 and T due to increased savings and reduces gradually after time T:
  - This temporal variation in capital stock smooths consumption over time, driven by the representative consumer's consumption-smoothing motive.

Let's collect the procedures used above into a function that runs the solver and draws plots for a given experiment

The following figure compares responses to a foreseen increase in g at $t = 10$ for two economies:

- our original economy with $\gamma = 2$, shown in the solid line, and

- an otherwise identical economy with $\gamma = 0.2$.

This comparison interest us because the utility curvature parameter γ governs the household's willingness to substitute consumption over time, and thus it preferences about smoothness of consumption paths over time.

```
# Solve the model using shooting
solution = run_shooting(shocks, S, model)

# Compute the initial steady states
k_ss_initial, c_ss_initial = steady_states(model,
                                           shocks['g'][0],
                                           shocks['τ_k'][0])

# Plot the solution for γ=2
fig, axes = plt.subplots(2, 3, figsize=(10, 8))
axes = axes.flatten()

label = fr"$\gamma = {model.γ}$"
plot_results(solution, k_ss_initial, c_ss_initial,
             shocks, 'g', axes, model, label=label,
             T=40)

# Solve and plot the result for γ=0.2
model_γ2 = create_model(γ=0.2)
solution = run_shooting(shocks, S, model_γ2)

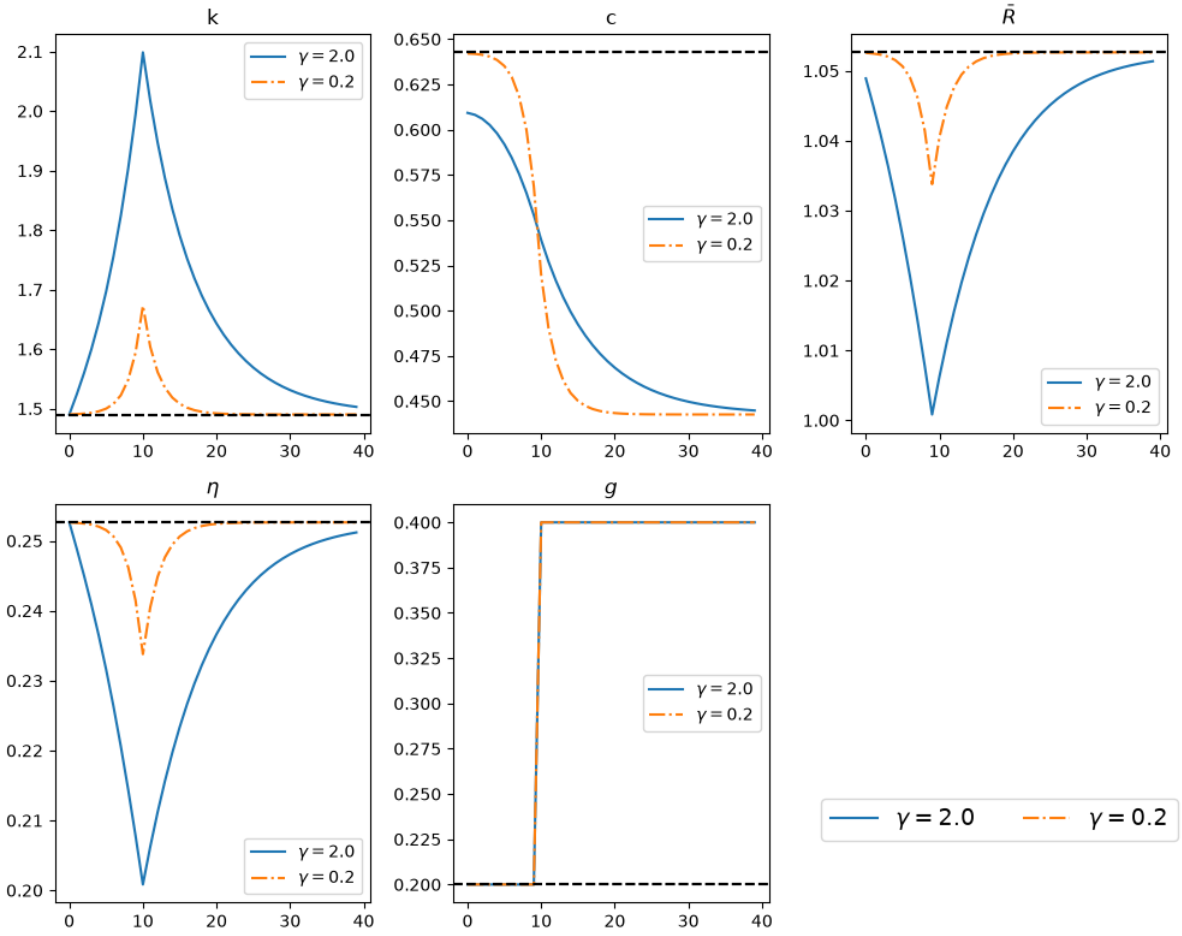
plot_results(solution, k_ss_initial, c_ss_initial,
             shocks, 'g', axes, model_γ2,
             label=fr"$\gamma = {model_γ2.γ}$",
             linestyle='-.', T=40)

handles, labels = axes[0].get_legend_handles_labels()
fig.legend(handles, labels, loc='lower right',
           ncol=3, fontsize=14, bbox_to_anchor=(1, 0.1))

for ax in axes[5:]:
    fig.delaxes(ax)

plt.tight_layout()
plt.show()
```

```
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.60924195288792396453121856997271325335517
Model: Model(β=0.95, γ=0.2, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.64203304129879029264147687246076236817745
```



The outcomes indicate that lowering γ affects both the consumption and capital stock paths because - it increases the representative consumer's willingness to substitute consumption across time:

- Consumption path:
  - When $\gamma = 0.2$, consumption becomes less smooth compared to $\gamma = 2$.
  - For $\gamma = 0.2$, consumption mirrors the government expenditure path more closely, staying higher until $t = 10$.
- Capital stock path:
  - With $\gamma = 0.2$, there are smaller build-ups and drawdowns of capital stock.
  - There are also smaller fluctuations in $\bar{R}$ and η.

Let's write another function that runs the solver and draws plots for these two experiments

Now we plot other equilibrium quantities:

```
def plot_prices(solution, c_ss, shock_param, axes,
                model, label='', linestyle='-', T=40):
    """
    Compares and plots prices
    """
    α, β, δ, γ, A = model.α, model.β, model.δ, model.γ, model.A

    k_path = solution[:, 0]
    c_path = solution[:, 1]
```

(continues on next page)

(continued from previous page)

```
    # Plot for c
    axes[0].plot(c_path[:T], linestyle=linestyle, label=label)
    axes[0].axhline(c_ss, linestyle='--', color='black')
    axes[0].set_title('c')

    # Plot for q
    q_path = compute_q_path(c_path, model, S=S)
    axes[1].plot(q_path[:T], linestyle=linestyle, label=label)
    axes[1].plot(β**np.arange(T), linestyle='--', color='black')
    axes[1].set_title('q')

    # Plot for r_{t,t+1}
    R_bar_path = compute_R_bar_path(shocks, k_path, model, S)
    axes[2].plot(R_bar_path[:T] - 1, linestyle=linestyle, label=label)
    axes[2].axhline(1 / β - 1, linestyle='--', color='black')
    axes[2].set_title('$r_{t,t+1}$')

    # Plot for r_{t,t+s}
    for style, s in zip(['-', '-.', '--'], [0, 10, 60]):
        rts_path = compute_rts_path(q_path, T, s)
        axes[3].plot(rts_path, linestyle=style,
                     color='black' if style == '--' else None,
                     label=f'$t={s}$')
        axes[3].set_xlabel('s')
        axes[3].set_title('$r_{t,t+s}$')

    # Plot for g
    axes[4].plot(shocks[shock_param][:T], linestyle=linestyle, label=label)
    axes[4].axhline(shocks[shock_param][0], linestyle='--', color='black')
    axes[4].set_title(shock_param)
```

For $\gamma = 2$ again, the next figure describes the response of q_t and the term structure of interest rates to a foreseen increase in g_t at $t = 10$

```
solution = run_shooting(shocks, S, model)

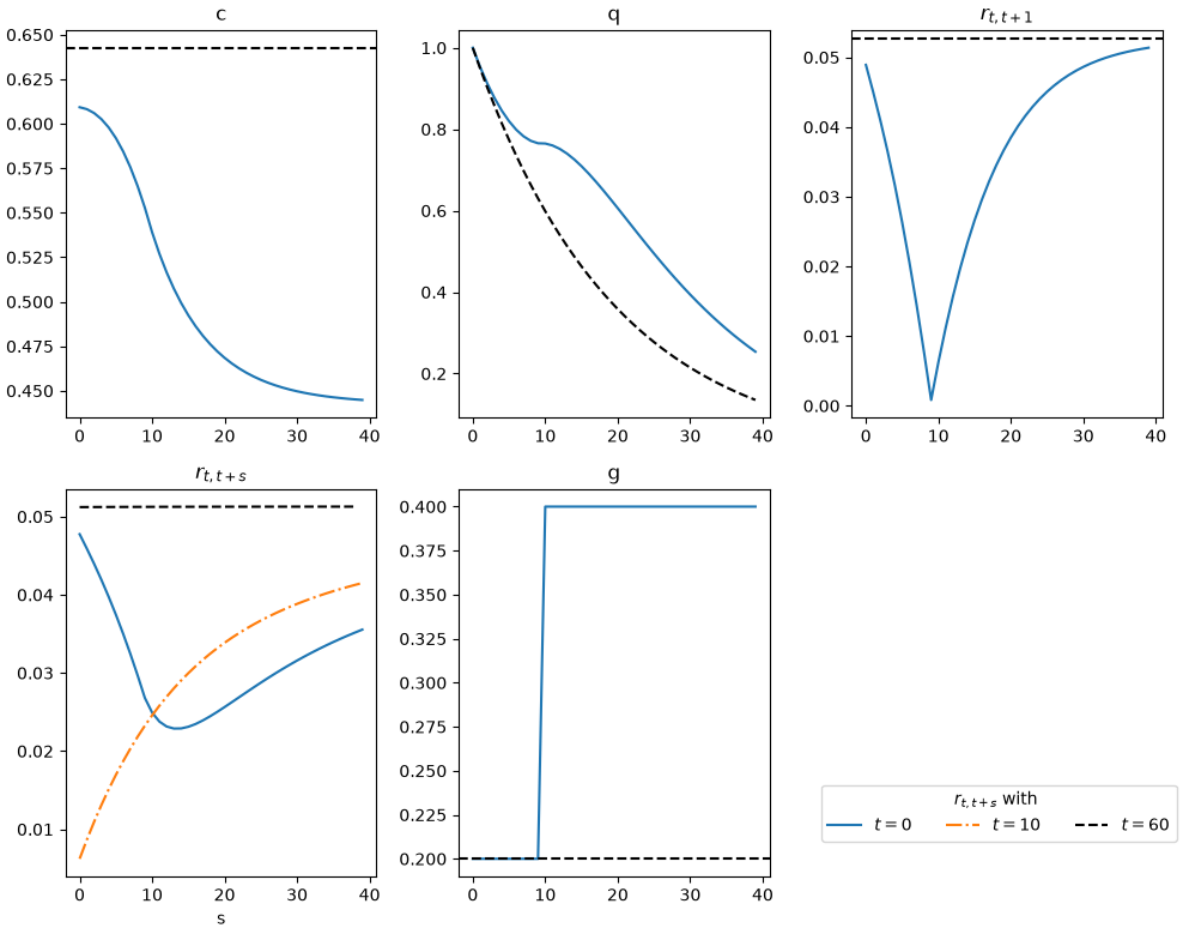
fig, axes = plt.subplots(2, 3, figsize=(10, 8))
axes = axes.flatten()

plot_prices(solution, c_ss_initial, 'g', axes, model, T=40)

for ax in axes[5:]:
    fig.delaxes(ax)

handles, labels = axes[3].get_legend_handles_labels()
fig.legend(handles, labels, title=r"$r_{t,t+s}$ with ", loc='lower right',
           ncol=3, fontsize=10, bbox_to_anchor=(1, 0.1))
plt.tight_layout()
plt.show()
```

```
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.6092419528879239645312185699727132533517
```



The second panel on the top compares q_t for the initial steady state with q_t after the increase in g is foreseen at $t = 0$, while the third panel compares the implied short rate r_t.

The fourth panel shows the term structure of interest rates at $t = 0$, $t = 10$, and $t = 60$.

Notice that, by $t = 60$, the system has converged to the new steady state, and the term structure of interest rates becomes flat.

At $t = 10$, the term structure of interest rates is upward sloping.

This upward slope reflects the expected increase in the rate of growth of consumption over time, as shown in the consumption panel.

At $t = 0$, the term structure of interest rates exhibits a "U-shaped" pattern:

- It declines until maturity at $s = 10$.
- After $s = 10$, it increases for longer maturities.

This pattern aligns with the pattern of consumption growth in the first two figures, which declines at an increasing rate until $t = 10$ and then declines at a decreasing rate afterward.

Experiment 2: Foreseen once-and-for-all increase in τ_c from 0.0 to 0.2 in period 10

With an inelastic labor supply, the Euler equation (112.16) and the other equilibrium conditions show that

- constant consumption taxes do not distort decisions, but
- anticipated changes in them do.

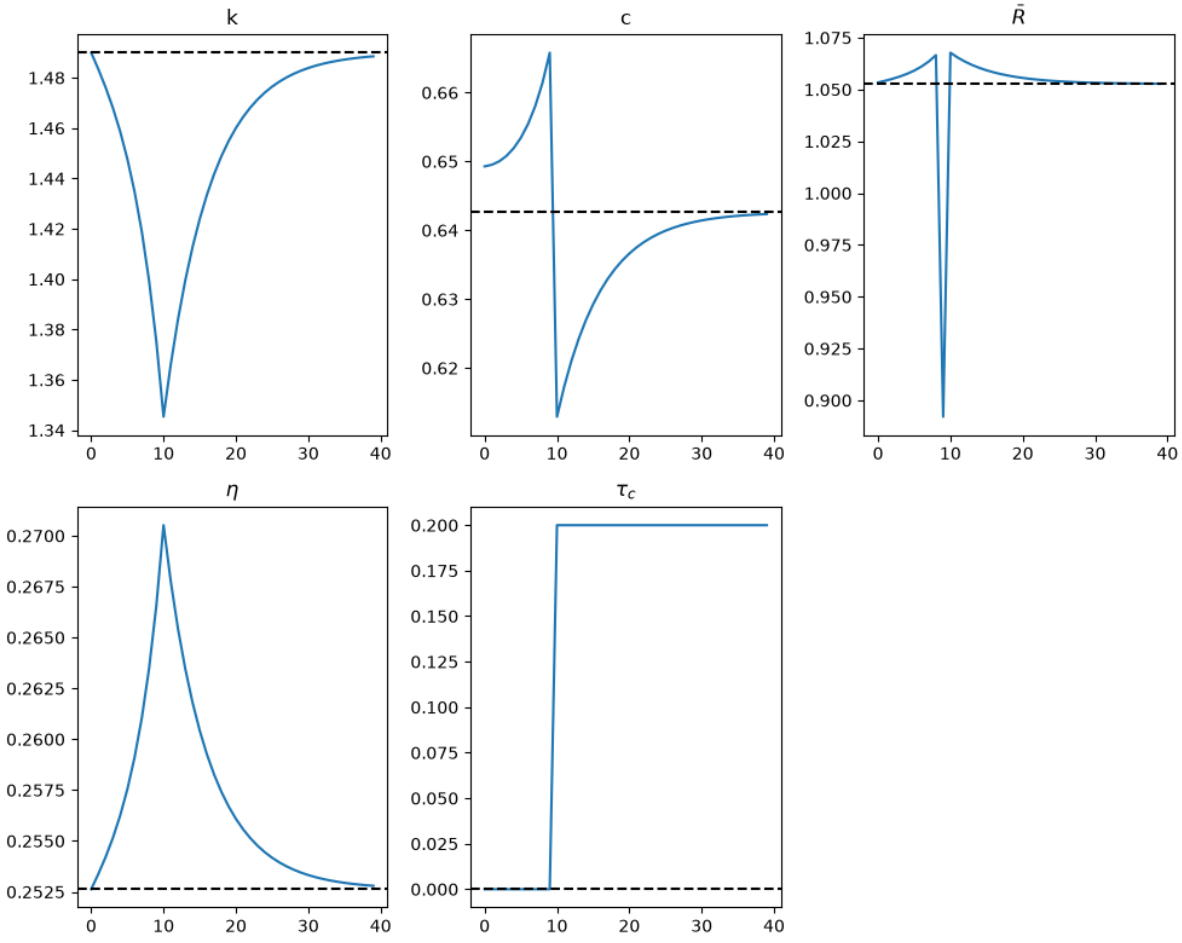
Indeed, (112.16) or (112.17) indicates that a foreseen in- crease in τ_{ct} (i.e., a decrease in $(1+\tau_{ct})$ $(1+\tau_{ct+1})$) operates like an increase in τ_{kt}.

The following figure portrays the response to a foreseen increase in the consumption tax τ_c.

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.concatenate((np.repeat(0.0, 10), np.repeat(0.2, S - 9))),
    'τ_k': np.repeat(0.0, S + 1)
}

experiment_model(shocks, S, model,
                 solver=run_shooting,
                 plot_func=plot_results,
                 policy_shock='τ_c')
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
----------------------------------------------------------------
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.64927956146815433723018647051957883983 96
```



Evidently all variables in the figures above eventually return to their initial steady-state values.

The anticipated increase in τ_{ct} leads to variations in consumption and capital stock across time:

- At $t = 0$:
  - Anticipation of the increase in τ_c causes an *immediate jump in consumption*.
  - This is followed by a *consumption binge* that sends the capital stock downward until $t = T = 10$.
- Between $t = 0$ and $t = T = 10$:
  - The decline in the capital stock raises $\bar{R}$ over time.
  - The equilibrium conditions require the growth rate of consumption to rise until $t = T$.
- At $t = T = 10$:
  - The jump in τ_c depresses $\bar{R}$ below 1, causing a *sharp drop in consumption*.
- After $T = 10$:
  - The effects of anticipated distortion are over, and the economy gradually adjusts to the lower capital stock.
  - Capital must now rise, requiring *austerity* –consumption plummets after $t = T$, indicated by lower levels of consumption.
  - The interest rate gradually declines, and consumption grows at a diminishing rate along the path to the terminal steady-state.

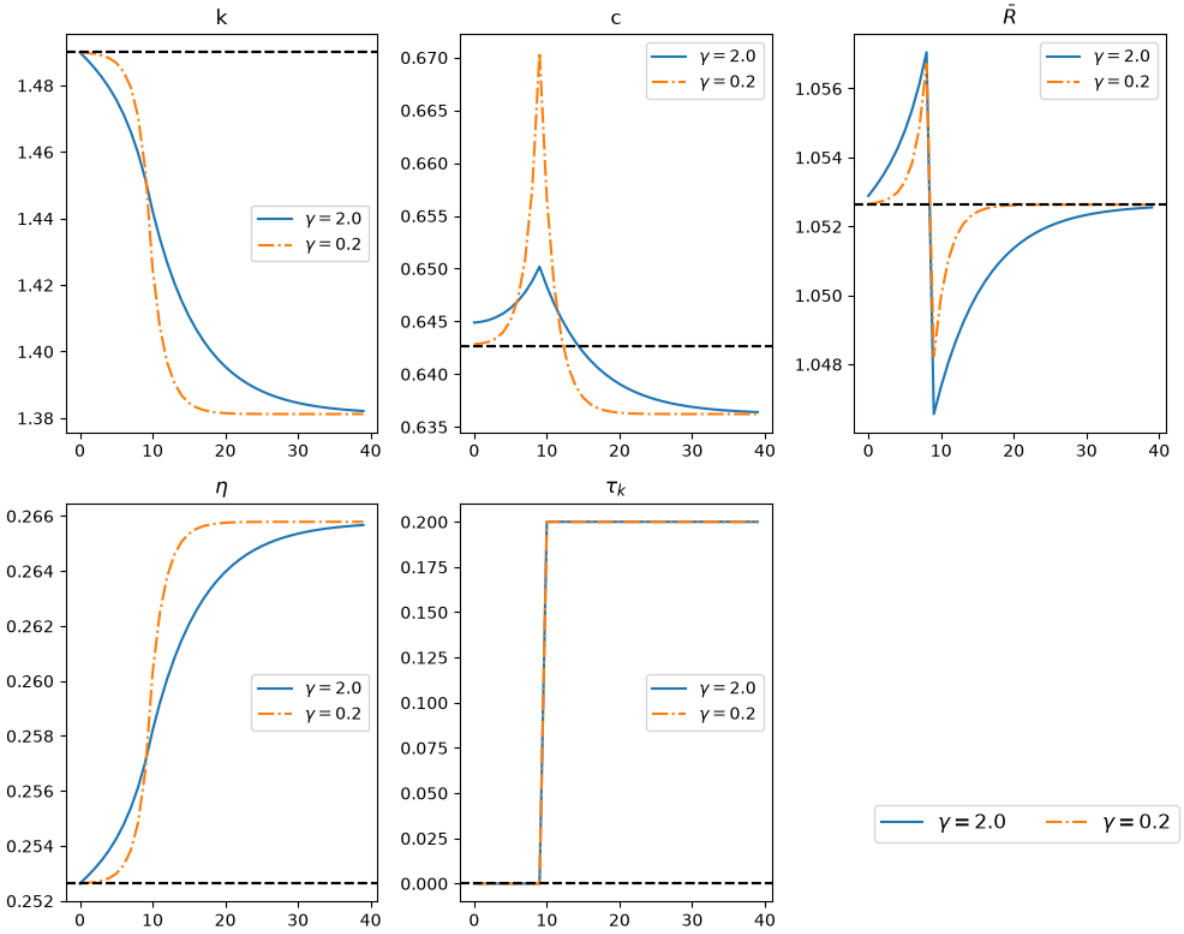
Experiment 3: Foreseen once-and-for-all increase in τ_k from 0.0 to 0.2 in period 10

For the two γ values 2 and 0.2, the next figure shows the response to a foreseen permanent jump in τ_{kt} at $t = T = 10$.

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.concatenate((np.repeat(0.0, 10), np.repeat(0.2, S - 9)))
}

experiment_two_models(shocks, S, model, model_γ2,
                solver=run_shooting,
                 plot_func=plot_results,
                 policy_shock='τ_k')
```

```
Model 1 (γ=2.0): steady state k=1.4900, c=0.6426
Model 2 (γ=0.2): steady state k=1.4900, c=0.6426
----------------------------------------------------------------
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.64488564003186084609963008227070148901 99
Model: Model(β=0.95, γ=0.2, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.64284077722405067274641526959215137674 15
```



The path of government expenditures remains fixed

- the increase in τ_{kt} is offset by a reduction in the present value of lump-sum taxes to keep the budget balanced.

The figure shows that:

- Anticipation of the increase in τ_{kt} leads to immediate decline in capital stock due to increased current consumption and a growing consumption flow.
- $\bar{R}$ starts rising at $t = 0$ and peaks at $t = 9$, and at $t = 10$, $\bar{R}$ drops sharply due to the tax change.
  - Variations in $\bar{R}$ align with the impact of the tax increase at $t = 10$ on consumption across time.
- Transition dynamics push k_t (capital stock) toward a new, lower steady-state level. In the new steady state:
  - Consumption is lower due to reduced output from the lower capital stock.
  - Smoother consumption paths occur when $\gamma = 2$ than when $\gamma = 0.2$.

So far we have explored consequences of foreseen once-and-for-all changes in government policy. Next we describe some experiments in which there is a foreseen one-time change in a policy variable (a "pulse").

Experiment 4: Foreseen one-time increase in g from 0.2 to 0.4 in period 10, after which g returns to 0.2 forever

```
g_path = np.repeat(0.2, S + 1)
g_path[10] = 0.4

shocks = {
```

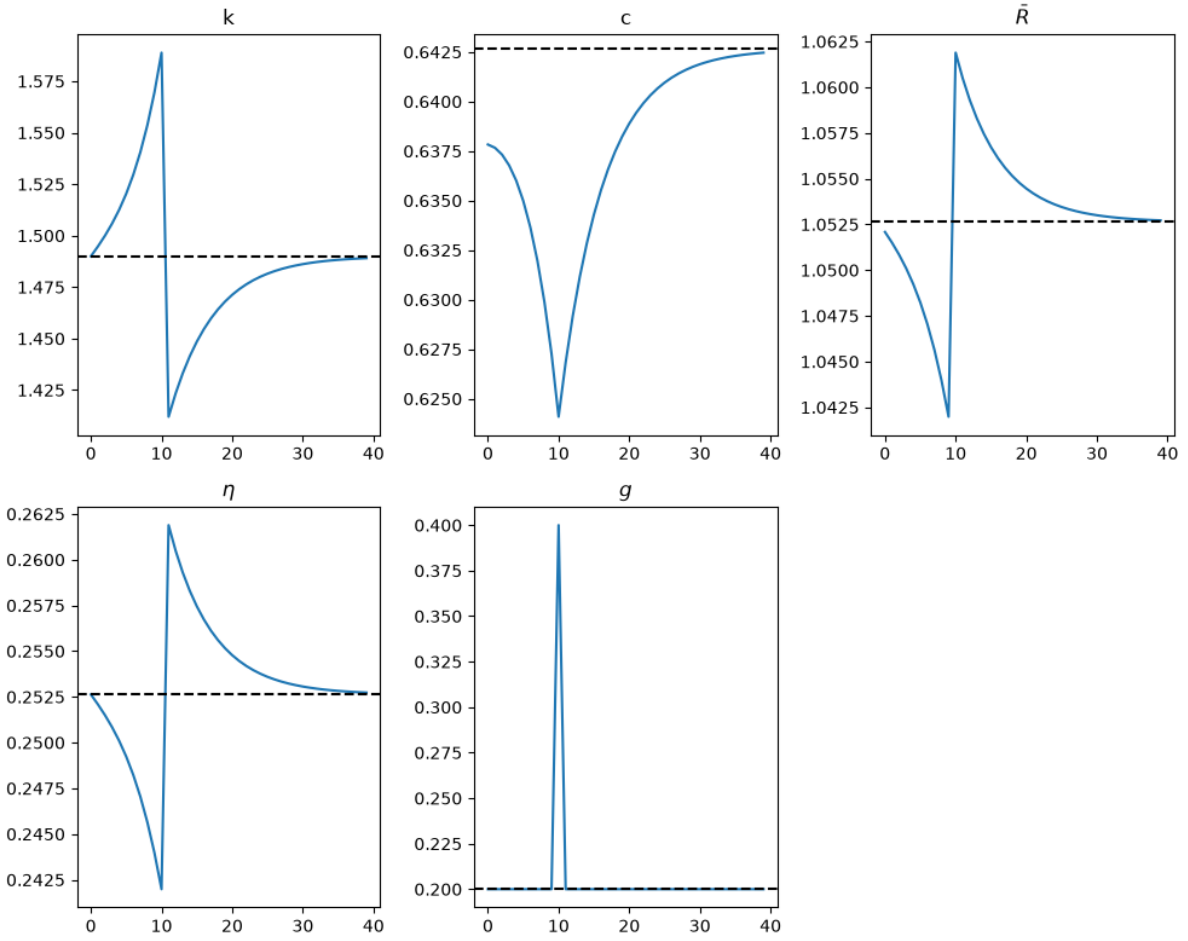
(continues on next page)

(continued from previous page)

```
    'g': g_path,
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1)
}

experiment_model(shocks, S, model,
                 solver=run_shooting,
                 plot_func=plot_results,
                 policy_shock='g')
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
------------------------------------------------------------
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
Optimal initial consumption c0 = 0.6378298012463969247674771825320030214755
```



The figure indicates how:

- Consumption:
  - Drops immediately upon announcement of the policy and continues to decline over time in anticipation of the one-time surge in g.
  - After the shock at $t = 10$, consumption begins to recover, rising at a diminishing rate toward its steady-state value.

- Capital and $\bar{R}$:
  - Before $t = 10$, capital accumulates as interest rate changes induce households to prepare for the anticipated increase in government spending.
  - At $t = 10$, the capital stock sharply decreases as the government consumes part of it.
  - $\bar{R}$ jumps above its steady-state value due to the capital reduction and then gradually declines toward its steady-state level.

112.9.3 Method 2: Residual Minimization

The second method involves minimizing residuals (i.e., deviations from equalities) of the following equations:

- The Euler equation (112.17):

$$1 = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} \frac{(1 + \tau_{ct})}{(1 + \tau_{ct+1})} \left[(1 - \tau_{kt+1})(\alpha A k_{t+1}^{\alpha - 1} - \delta) + 1 \right]$$

- The feasibility condition (112.15):

$$k_{t+1} = A k_t^{\alpha} + (1 - \delta) k_t - g_t - c_t.$$

```
# Euler's equation and feasibility condition
def euler_residual(c_t, c_tp1, τ_c_t, τ_c_tp1, τ_k_tp1, k_tp1, model, μ_tp1=1):
    """
    Computes the residuals for Euler's equation
    with optional growth model parameters μ_tp1.
    """
    R_bar = compute_R_bar(τ_c_t, τ_c_tp1, τ_k_tp1, k_tp1, model)

    c_expected = next_c(c_t, R_bar, model, μ_tp1)

    return c_expected / c_tp1 - 1.0

def feasi_residual(k_t, k_tm1, c_tm1, g_t, model, μ_t=1):
    """
    Computes the residuals for feasibility condition
    with optional growth model parameter μ_t.
    """
    k_t_expected = next_k(k_tm1, g_t, c_tm1, model, μ_t)
    return k_t_expected - k_t
```

The algorithm proceeds follows:

1. Find initial steady state k_0 based on the government plan at $t = 0$.
2. Initialize a sequence of initial guesses $\{\hat{c}_t, \hat{k}_t\}_{t=0}^{S}$.
3. Compute residuals l_a and l_k for $t = 0, \ldots, S$, as well as l_{k_0} for $t = 0$ and l_{k_S} for $t = S$:
   - Compute the Euler equation residual for $t = 0, \ldots, S$ using (112.17):

$$l_{ta} = \beta u'(c_{t+1}) \frac{(1 + \tau_{ct})}{(1 + \tau_{ct+1})} \left[(1 - \tau_{kt+1})(f'(k_{t+1}) - \delta) + 1 \right] - 1$$

   - Compute the feasibility condition residual for $t = 1, \ldots, S - 1$ using (112.15):

$$l_{tk} = k_{t+1} - f(k_t) - (1 - \delta) k_t + g_t + c_t$$

- Compute the residual for the initial condition for k_0 using (112.19) and the initial capital k_0:

$$l_{k_0} = 1 - \beta \left[(1 - \tau_{k0}) \left(f'(k_0) - \delta \right) + 1 \right]$$

- Compute the residual for the terminal condition for $t = S$ using (112.17) under the assumptions $c_t = c_{t+1} = c_S$, $k_t = k_{t+1} = k_S$, $\tau_{ct} = \tau_{ct+1} = \tau_{c_s}$, and $\tau_{kt} = \tau_{kt+1} = \tau_{k_s}$:

$$l_{k_S} = \beta u'(c_S) \frac{(1 + \tau_{c_s})}{(1 + \tau_{c_s})} \left[(1 - \tau_{k_s})(f'(k_S) - \delta) + 1 \right] - 1$$

4. Iteratively adjust guesses for $\{\hat{c}_t, \hat{k}_t\}_{t=0}^{S}$ to minimize residuals l_{k_0}, l_{ta}, l_{tk}, and l_{k_S} for $t = 0, \ldots, S$.

```
def compute_residuals(vars_flat, k_init, S, shock_paths, model):
    """
    Compute the residuals for the Euler equation and feasibility condition.
    """
    g, τ_c, τ_k, μ = (shock_paths[key] for key in ('g','τ_c','τ_k','μ'))
    k, c = vars_flat.reshape((S+1, 2)).T
    res = np.empty(2*S+2, dtype=float)

    # boundary condition on initial capital
    res[0] = k[0] - k_init

    # interior Euler and feasibility
    for t in range(S):
        res[2*t + 1] = euler_residual(
            c[t],    c[t+1],
            τ_c[t],  τ_c[t+1],
            τ_k[t+1],k[t+1],
            model, μ[t+1])
        res[2*t + 2] = feasi_residual(
            k[t+1], k[t], c[t],
            g[t],  model,
            μ[t+1])

    # terminal Euler condition at t=S
    res[-1] = euler_residual(
        c[S],   c[S],
        τ_c[S], τ_c[S],
        τ_k[S], k[S],
        model,
        μ[S])

    return res


def run_min(shocks, S, model, A_path=None):
    """
    Solve for the full (k,c) path by root-finding the residuals.
    """
    shocks['μ'] = shocks['μ'] if 'μ' in shocks else np.ones_like(shocks['g'])

    # compute the steady-state to serve as both initial capital and guess
    k_ss, c_ss = steady_states(
        model,
        shocks['g'][0],
```

(continues on next page)

(continued from previous page)

```
        shocks['τ_k'][0],
        shocks['μ'][0]  # =1 if no growth
    )

    # initial guess: flat at the steady-state
    guess = np.column_stack([
        np.full(S+1, k_ss),
        np.full(S+1, c_ss)
    ]).flatten()

    sol = root(
        compute_residuals,
        guess,
        args=(k_ss, S, shocks, model),
        tol=1e-8
    )

    return sol.x.reshape((S+1, 2))
```

We found that method 2 did not encounter numerical stability issues, so using `mp.mpf` is not necessary.

We leave as exercises replicating some of our experiments by using the second method.

Exercise 112.9.1

Replicate the plots of our four experiments using the second method of residual minimization:

1. A foreseen once-and-for-all increase in g from 0.2 to 0.4 occurring in period 10,
2. A foreseen once-and-for-all increase in τ_c from 0.0 to 0.2 occurring in period 10,
3. A foreseen once-and-for-all increase in τ_k from 0.0 to 0.2 occurring in period 10, and
4. A foreseen one-time increase in g from 0.2 to 0.4 in period 10, after which g reverts to 0.2 permanently,

Solution

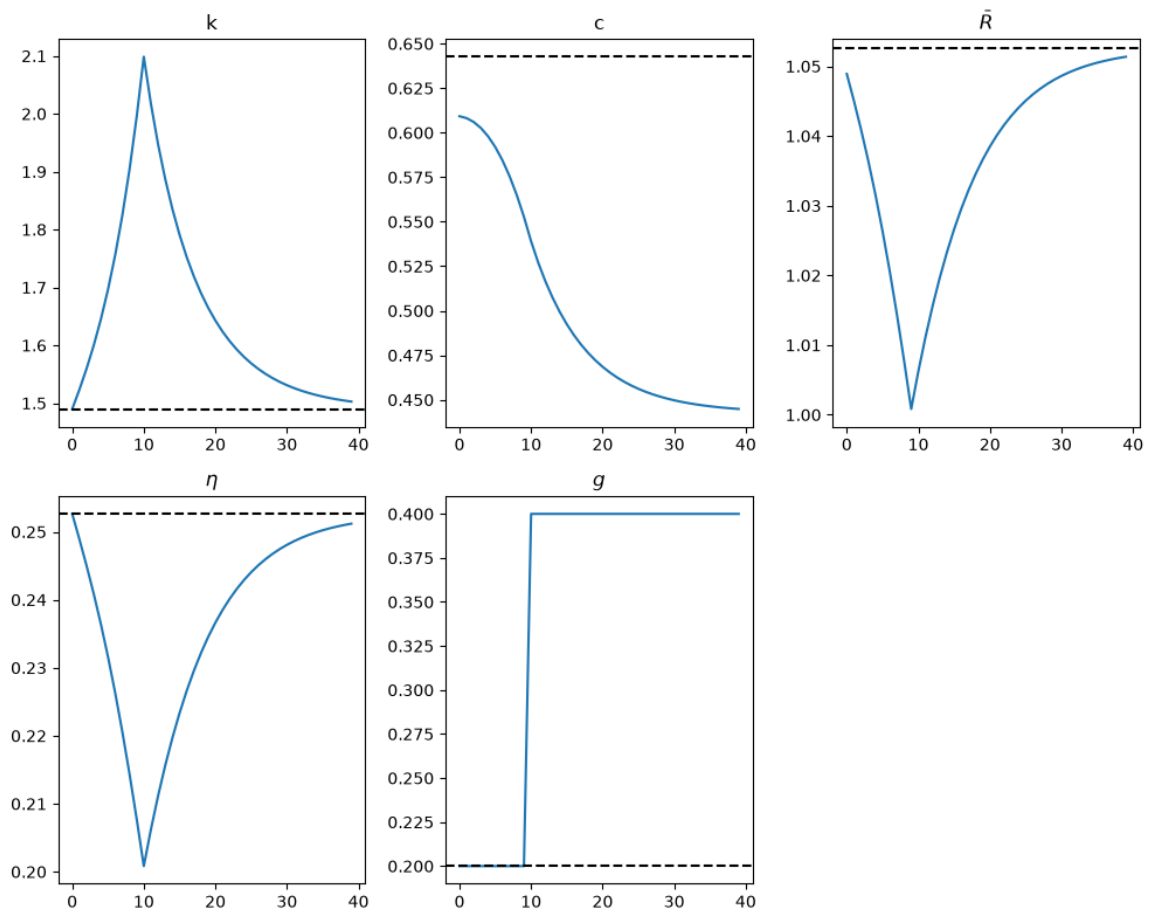
Here is one solution:

Experiment 1: Foreseen once-and-for-all increase in g from 0.2 to 0.4 in period 10

```
shocks = {
    'g': np.concatenate((np.repeat(0.2, 10), np.repeat(0.4, S - 9))),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1)
}

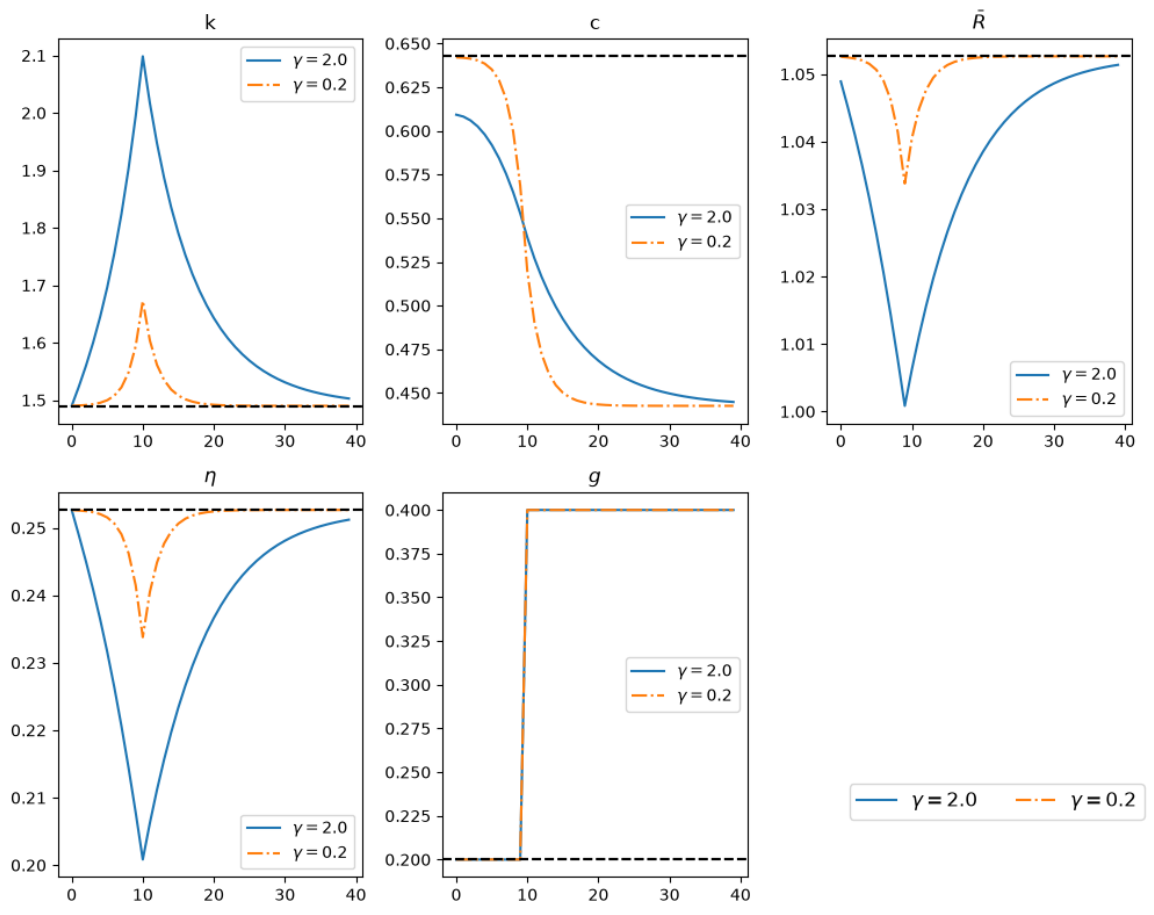
experiment_model(shocks, S, model, solver=run_min,
                 plot_func=plot_results,
                 policy_shock='g')
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
----------------------------------------------------------------
```




```
experiment_two_models(shocks, S, model, model_y2,
                run_min, plot_results, 'g')
```

```
Model 1 (y=2.0): steady state k=1.4900, c=0.6426
Model 2 (y=0.2): steady state k=1.4900, c=0.6426
------------------------------------------------------------
```



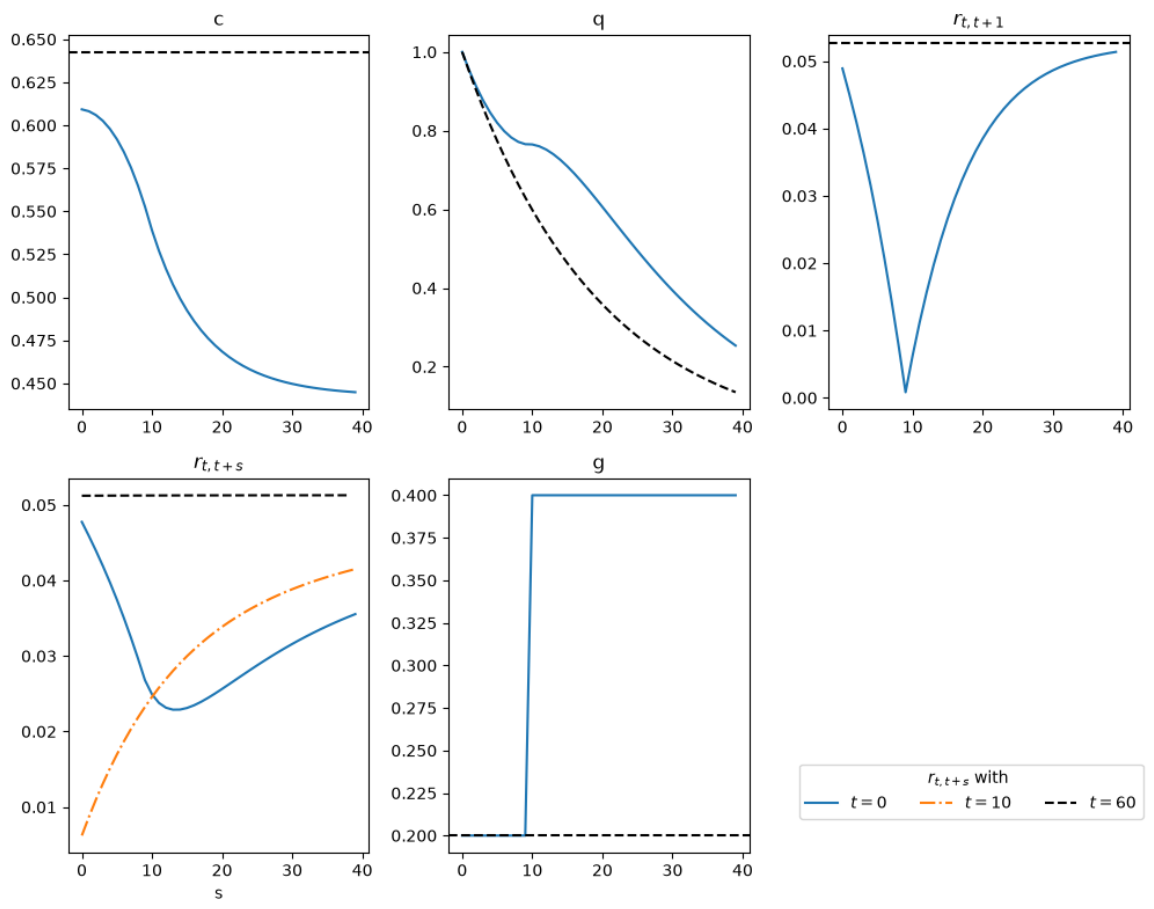

```
solution = run_min(shocks, S, model)

fig, axes = plt.subplots(2, 3, figsize=(10, 8))
axes = axes.flatten()

plot_prices(solution, c_ss_initial, 'g', axes, model, T=40)

for ax in axes[5:]:
    fig.delaxes(ax)

handles, labels = axes[3].get_legend_handles_labels()
fig.legend(handles, labels, title=r"$r_{t,t+s}$ with ", loc='lower right', ncol=3,␣
↪fontsize=10, bbox_to_anchor=(1, 0.1))
plt.tight_layout()
plt.show()
```

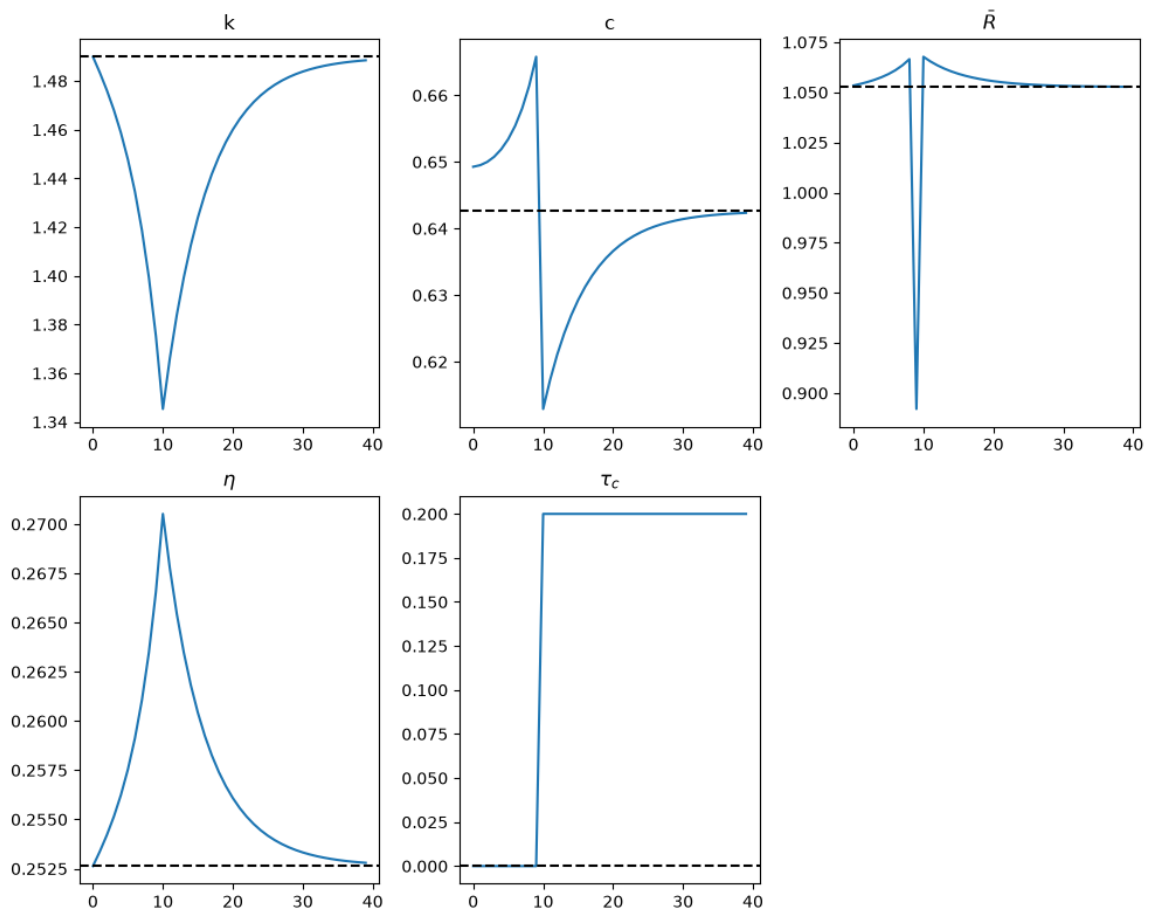


Experiment 2: Foreseen once-and-for-all increase in τ_c from 0.0 to 0.2 in period 10.

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.concatenate((np.repeat(0.0, 10), np.repeat(0.2, S - 9))),
    'τ_k': np.repeat(0.0, S + 1)
}

experiment_model(shocks, S, model, solver=run_min,
                 plot_func=plot_results,
                 policy_shock='τ_c')
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
----------------------------------------------------------------
```

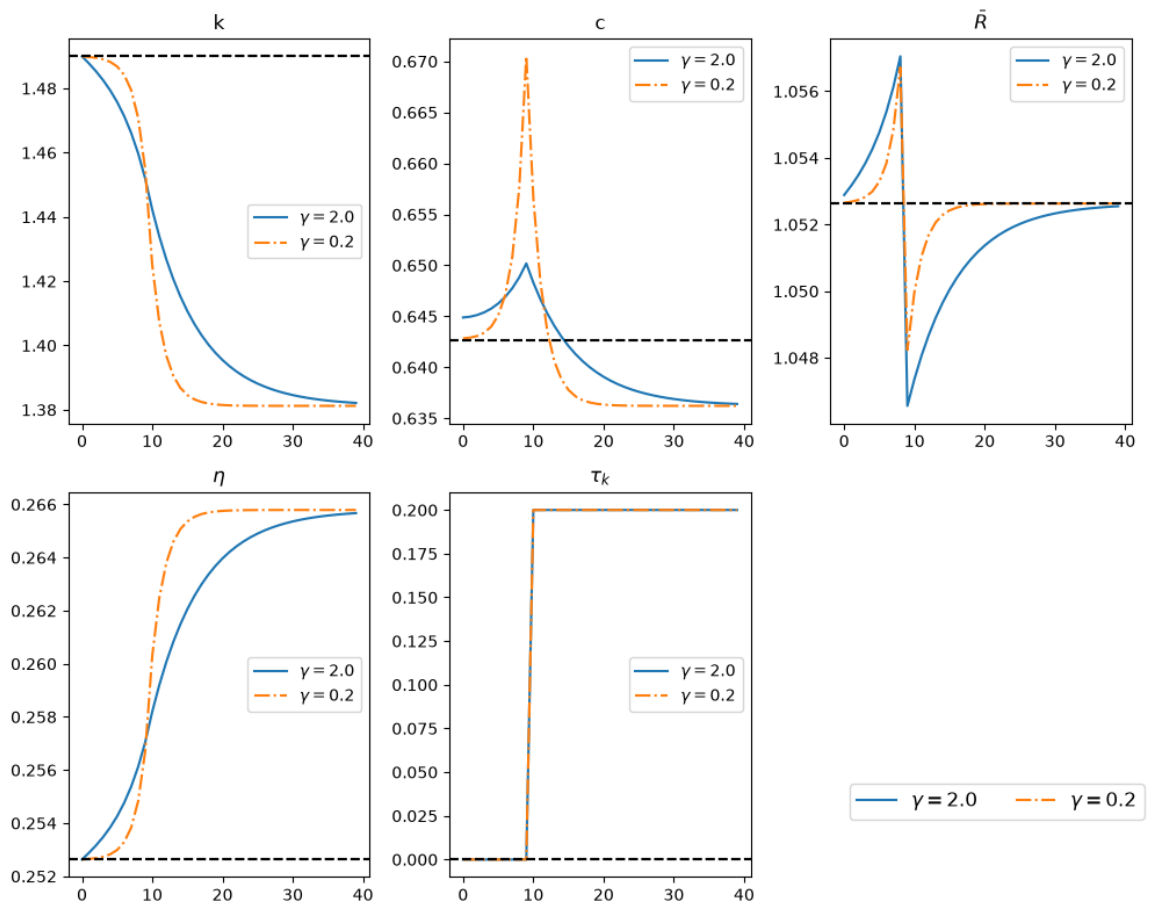


Experiment 3: Foreseen once-and-for-all increase in τ_k from 0.0 to 0.2 in period 10.

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.concatenate((np.repeat(0.0, 10), np.repeat(0.2, S - 9)))
}

experiment_two_models(shocks, S, model, model_γ2,
                solver=run_min,
                 plot_func=plot_results,
                 policy_shock='τ_k')
```

```
Model 1 (γ=2.0): steady state k=1.4900, c=0.6426
Model 2 (γ=0.2): steady state k=1.4900, c=0.6426
----------------------------------------------------------------
```



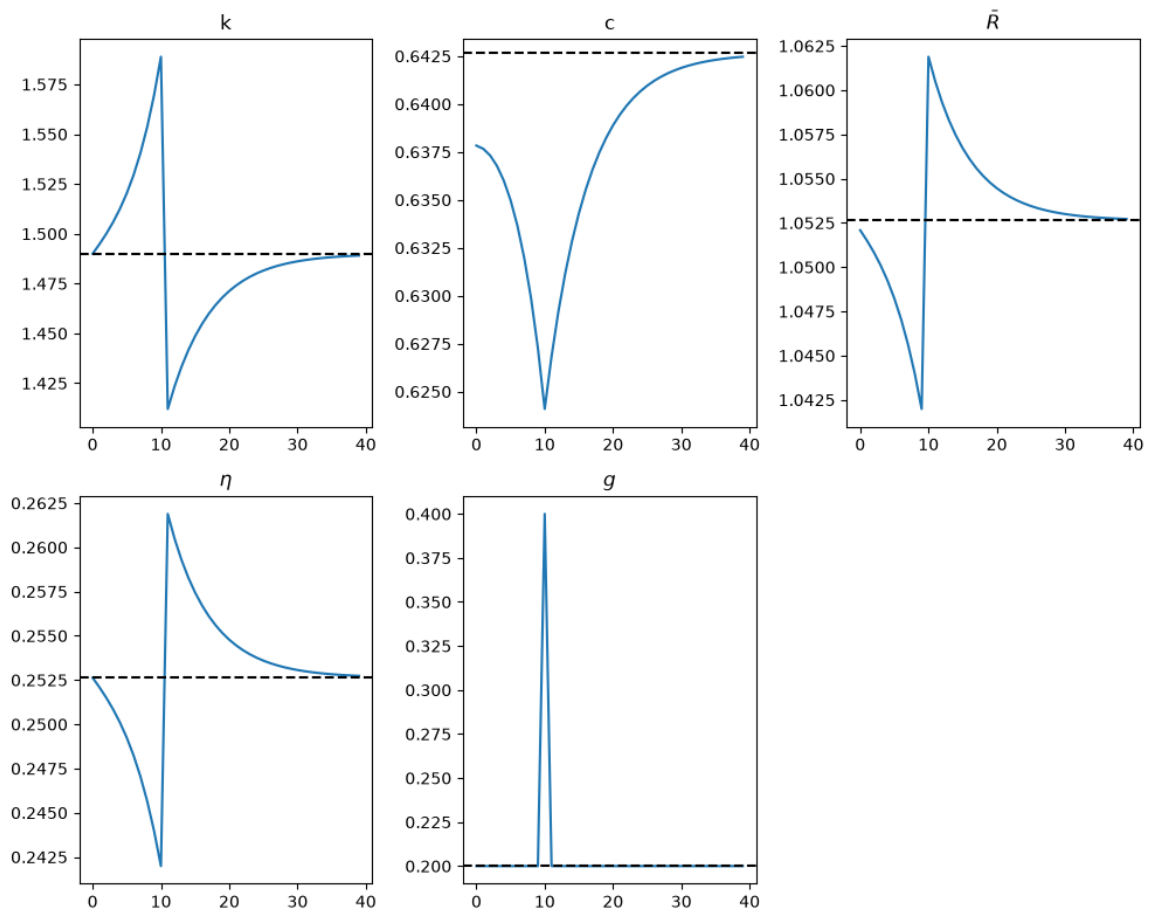
Experiment 4: Foreseen one-time increase in g from 0.2 to 0.4 in period 10, after which g returns to 0.2 forever

```
g_path = np.repeat(0.2, S + 1)
g_path[10] = 0.4

shocks = {
    'g': g_path,
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1)
}

experiment_model(shocks, S, model, solver=run_min,
                 plot_func=plot_results,
                 policy_shock='g')
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
------------------------------------------------------------------
```



Exercise 112.9.2

Design a new experiment where the government expenditure g increases from 0.2 to 0.4 in period 10, and then decreases to 0.1 in period 20 permanently.

Solution

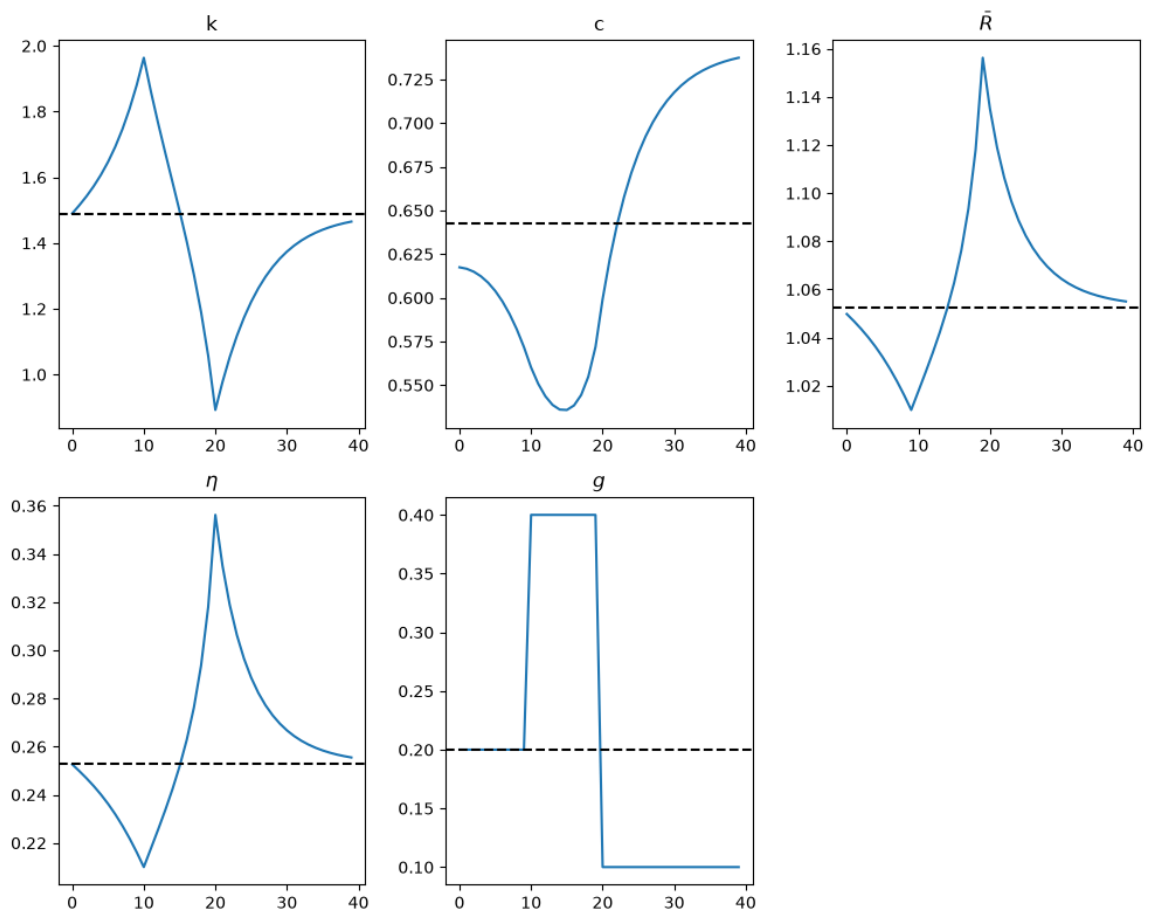
Here is one solution:

```
g_path = np.repeat(0.2, S + 1)
g_path[10:20] = 0.4
g_path[20:] = 0.1

shocks = {
    'g': g_path,
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1)
}

experiment_model(shocks, S, model, solver=run_min,
                 plot_func=plot_results,
                 policy_shock='g')
```

```
Steady-state capital: 1.4900
Steady-state consumption: 0.6426
------------------------------------------------------------------
```



112.10 Exogenous growth

In the previous section, we considered a model without exogenous growth.

We set the term A_t in the production function to a constant by setting $A_t = 1$ for all t.

Now we are ready to consider growth.

To incorporate growth, we modify the production function to be

$$Y_t = F(K_t, A_t n_t)$$

where Y_t is aggregate output, N_t is total employment, A_t is labor-augmenting technical change, and $F(K, AN)$ is the same linearly homogeneous production function as before.

We assume that A_t follows the process

$$A_{t+1} = \mu_{t+1} A_t \tag{112.25}$$

and that $\mu_{t+1} = \bar{\mu} > 1$.

```
# Set the constant A parameter to None
model = create_model(A=None)
```

```
def compute_A_path(A0, shocks, S=100):
    """
    Compute A path over time.
    """
    A_path = np.full(S + 1, A0)
    for t in range(1, S + 1):
        A_path[t] = A_path[t-1] * shocks['μ'][t-1]
    return A_path
```

112.10.1 Inelastic Labor Supply

By linear homogeneity, the production function can be expressed as

$$y_t = f(k_t)$$

where $f(k) = F(k, 1) = k^\alpha$ and $k_t = \frac{K_t}{n_t A_t}$, $y_t = \frac{Y_t}{n_t A_t}$.

k_t and y_t are measured per unit of "effective labor" $A_t n_t$.

We also let $c_t = \frac{C_t}{A_t n_t}$ and $g_t = \frac{G_t}{A_t n_t}$, where C_t and G_t are total consumption and total government expenditures.

We continue to consider the case of inelastic labor supply.

Based on this, feasibility can be summarized by the following modified version of equation (112.15):

$$k_{t+1} = \mu_{t+1}^{-1}[f(k_t) + (1-\delta)k_t - g_t - c_t] \tag{112.26}$$

Again, by the properties of a linearly homogeneous production function, we have

$$\eta_t = F_k(k_t, 1) = f'(k_t), w_t = F_n(k_t, 1) = f(k_t) - f'(k_t)k_t$$

Since per capita consumption is now $c_t A_t$, the counterpart to the Euler equation (112.17) is

$$u'(c_t A_t) = \beta u'(c_{t+1} A_{t+1}) \frac{(1+\tau_{ct})}{(1+\tau_{ct+1})}[(1-\tau_{kt+1})(f'(k_{t+1}) - \delta) + 1]. \tag{112.27}$$

$\bar{R}_{t+1}$ continues to be defined by (112.21), except that now k_t is capital per effective unit of labor.

Thus, substituting (112.21), (112.27) becomes

$$u'(c_t A_t) = \beta u'(c_{t+1} A_{t+1}) \bar{R}_{t+1}$$

Assuming that the household's utility function is the same as before, we have

$$(c_t A_t)^{-\gamma} = \beta (c_{t+1} A_{t+1})^{-\gamma} \bar{R}_{t+1}$$

Thus, the counterpart to (112.24) is

$$c_{t+1} = c_t \left[\beta \bar{R}_{t+1}\right]^{\frac{1}{\gamma}} \mu_{t+1}^{-1} \tag{112.28}$$

112.10.2 Steady State

In a steady state, $c_{t+1} = c_t$. Then (112.27) becomes

$$1 = \mu^{-\gamma}\beta[(1-\tau_k)(f'(k)-\delta)+1] \tag{112.29}$$

from which we can compute that the steady-state level of capital per unit of effective labor satisfies

$$f'(k) = \delta + (\frac{\frac{1}{\beta}\mu^{\gamma}-1}{1-\tau_k}) \tag{112.30}$$

and that

$$\bar{R} = \frac{\mu^{\gamma}}{\beta} \tag{112.31}$$

The steady-state level of consumption per unit of effective labor can be found using (112.26):

$$c = f(k) + (1-\delta-\mu)k - g$$

Since the algorithm and plotting routines are the same as before, we include the steady-state calculations and shooting routine in the section *Python Code*.

112.10.3 Shooting Algorithm

Now we can apply the shooting algorithm to compute equilibrium. We augment the vector of shock variables by including μ_t, then proceed as before.

112.10.4 Experiments

Let's run some experiments:

1. A foreseen once-and-for-all increase in μ from 1.02 to 1.025 in period 10
2. An unforeseen once-and-for-all increase in μ to 1.025 in period 0

Experiment 1: A foreseen increase in μ from 1.02 to 1.025 at t=10

The figures below show the effects of a permanent increase in productivity growth μ from 1.02 to 1.025 at t=10.

They now measure c and k in effective units of labor.

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1),
    'μ': np.concatenate((np.repeat(1.02, 10), np.repeat(1.025, S - 9)))
}

A_path = compute_A_path(1.0, shocks, S)

k_ss_initial, c_ss_initial = steady_states(model,
                                           shocks['g'][0],
                                           shocks['τ_k'][0],
                                           shocks['μ'][0]
```

(continues on next page)

(continued from previous page)

```
                                        )

print(f"Steady-state capital: {k_ss_initial:.4f}")
print(f"Steady-state consumption: {c_ss_initial:.4f}")

# Run the shooting algorithm with the A_path parameter
solution = run_shooting(shocks, S, model, A_path)

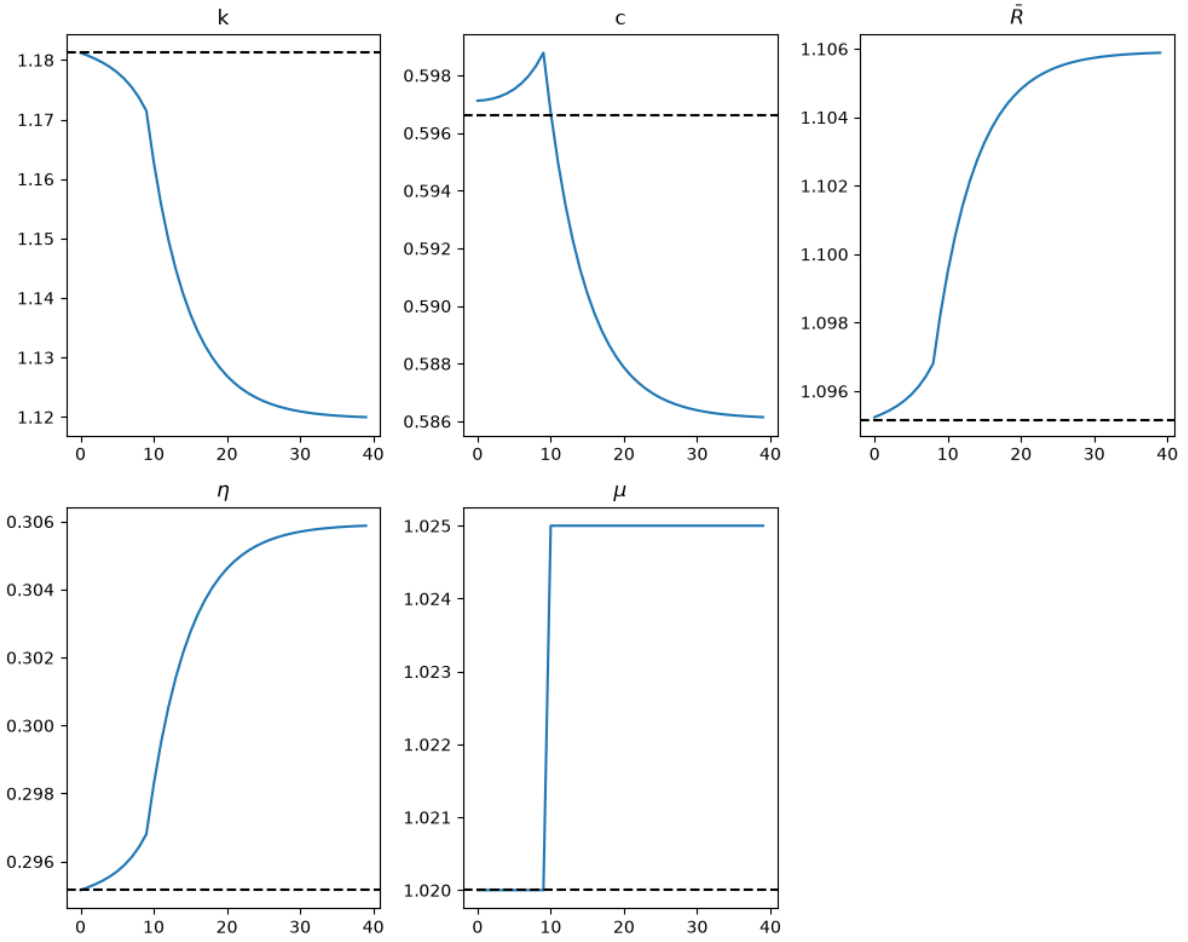
fig, axes = plt.subplots(2, 3, figsize=(10, 8))
axes = axes.flatten()

plot_results(solution, k_ss_initial,
             c_ss_initial, shocks, 'μ', axes, model,
             A_path, T=40)

for ax in axes[5:]:
    fig.delaxes(ax)

plt.tight_layout()
plt.show()
```

```
Steady-state capital: 1.1812
Steady-state consumption: 0.5966
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=None)
Optimal initial consumption c0 = 0.597118474934446239627091833718360733991 9
```



The results in the figures are mainly driven by (112.29) and imply that a permanent increase in μ will lead to a decrease in the steady-state value of capital per unit of effective labor.

The figures indicate the following:

- As capital becomes more efficient, even with less of it, consumption per capita can be raised.
- Consumption smoothing drives an *immediate jump in consumption* in anticipation of the increase in μ.
- The increased productivity of capital leads to an increase in the gross return $\bar{R}$.
- Perfect foresight makes the effects of the increase in the growth of capital precede it, with the effect visible at $t = 0$.

Experiment 2: An unforeseen increase in μ from 1.02 to 1.025 at t=0

The figures below show the effects of an immediate jump in μ to 1.025 at t=0.

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1),
    'μ': np.concatenate((np.repeat(1.02, 1), np.repeat(1.025, S)))
}

A_path = compute_A_path(1.0, shocks, S)
```

(continues on next page)

(continued from previous page)

```
k_ss_initial, c_ss_initial = steady_states(model,
                                           shocks['g'][0],
                                           shocks['τ_k'][0],
                                           shocks['μ'][0]
                                          )

print(f"Steady-state capital: {k_ss_initial:.4f}")
print(f"Steady-state consumption: {c_ss_initial:.4f}")

# Run the shooting algorithm with the A_path parameter
solution = run_shooting(shocks, S, model, A_path)

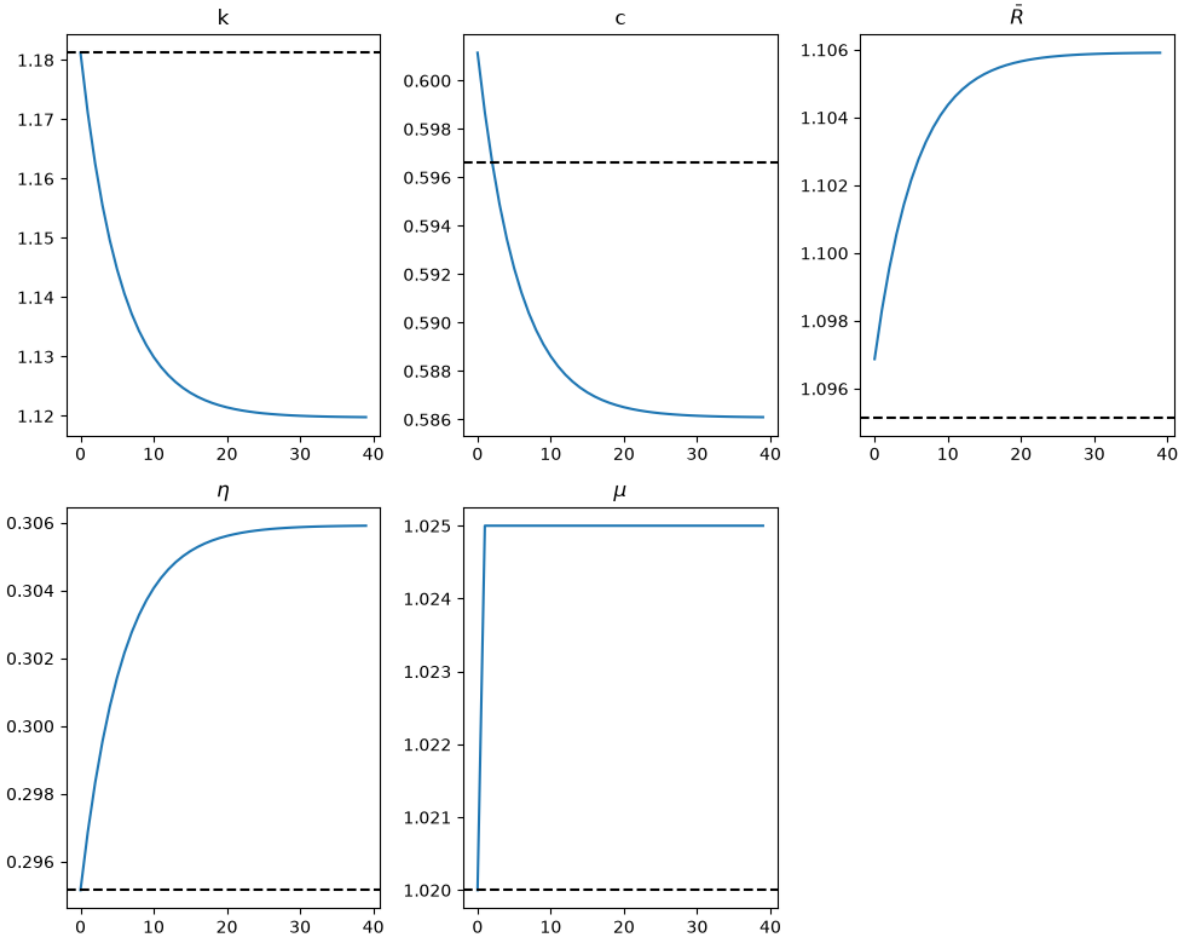
fig, axes = plt.subplots(2, 3, figsize=(10, 8))
axes = axes.flatten()

plot_results(solution, k_ss_initial,
             c_ss_initial, shocks, 'μ', axes, model, A_path, T=40)

for ax in axes[5:]:
    fig.delaxes(ax)

plt.tight_layout()
plt.show()
```

```
Steady-state capital: 1.1812
Steady-state consumption: 0.5966
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=None)
Optimal initial consumption c0 = 0.601149493043064115039588375310958831663 8
```



Again, we can collect the procedures used above into a function that runs the solver and draws plots for a given experiment.

```
def experiment_model(shocks, S, model, A_path, solver, plot_func, policy_shock, T=40):
    """
    Run the shooting algorithm given a model and plot the results.
    """
    k0, c0 = steady_states(model, shocks['g'][0], shocks['τ_k'][0], shocks['μ'][0])

    print(f"Steady-state capital: {k0:.4f}")
    print(f"Steady-state consumption: {c0:.4f}")
    print('-'*64)

    fig, axes = plt.subplots(2, 3, figsize=(10, 8))
    axes = axes.flatten()

    solution = solver(shocks, S, model, A_path)
    plot_func(solution, k0, c0,
              shocks, policy_shock, axes, model, A_path, T=T)

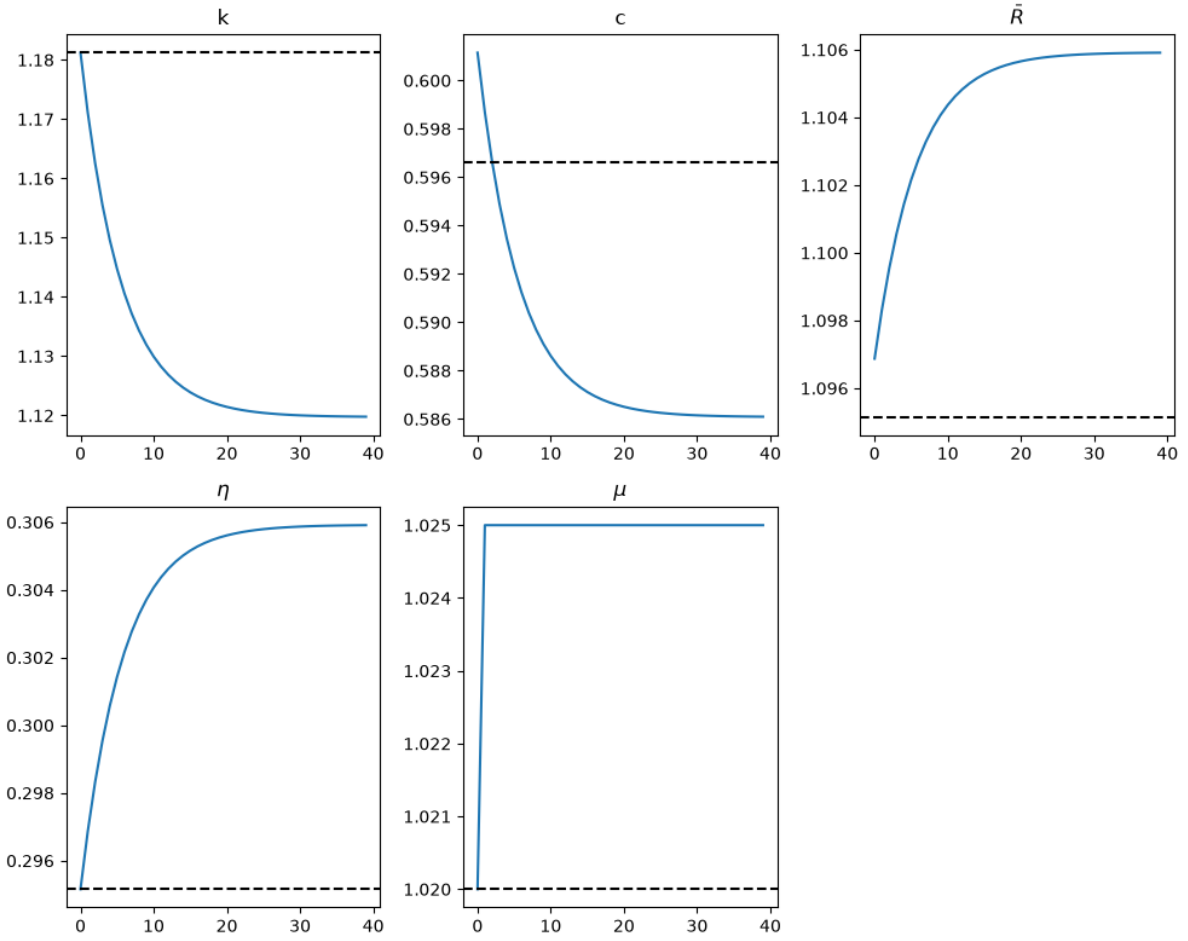
    for ax in axes[5:]:
        fig.delaxes(ax)

    plt.tight_layout()
    plt.show()
```

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1),
    'μ': np.concatenate((np.repeat(1.02, 1), np.repeat(1.025, S)))
}

experiment_model(shocks, S, model, A_path, run_shooting, plot_results, 'μ')
```

```
Steady-state capital: 1.1812
Steady-state consumption: 0.5966
----------------------------------------------------------------------
Model: Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=None)
Optimal initial consumption c0 = 0.601149493043064115039588375310958831663 8
```



The figures show that:

- The paths of all variables are now smooth due to the absence of feedforward effects.
- Capital per effective unit of labor gradually declines to a lower steady-state level.
- Consumption per effective unit of labor jumps immediately and then declines smoothly toward its lower steady-state value.
- The after-tax gross return $\bar{R}$ once again co-moves with the consumption growth rate, verifying the Euler equation (112.29).

Exercise 112.10.1

Replicate the plots of our two experiments using the second method of residual minimization:

1. A foreseen increase in μ from 1.02 to 1.025$ at t=10
2. An unforeseen increase in μ from 1.02 to 1.025$ at t=0

Solution

Here is one solution:

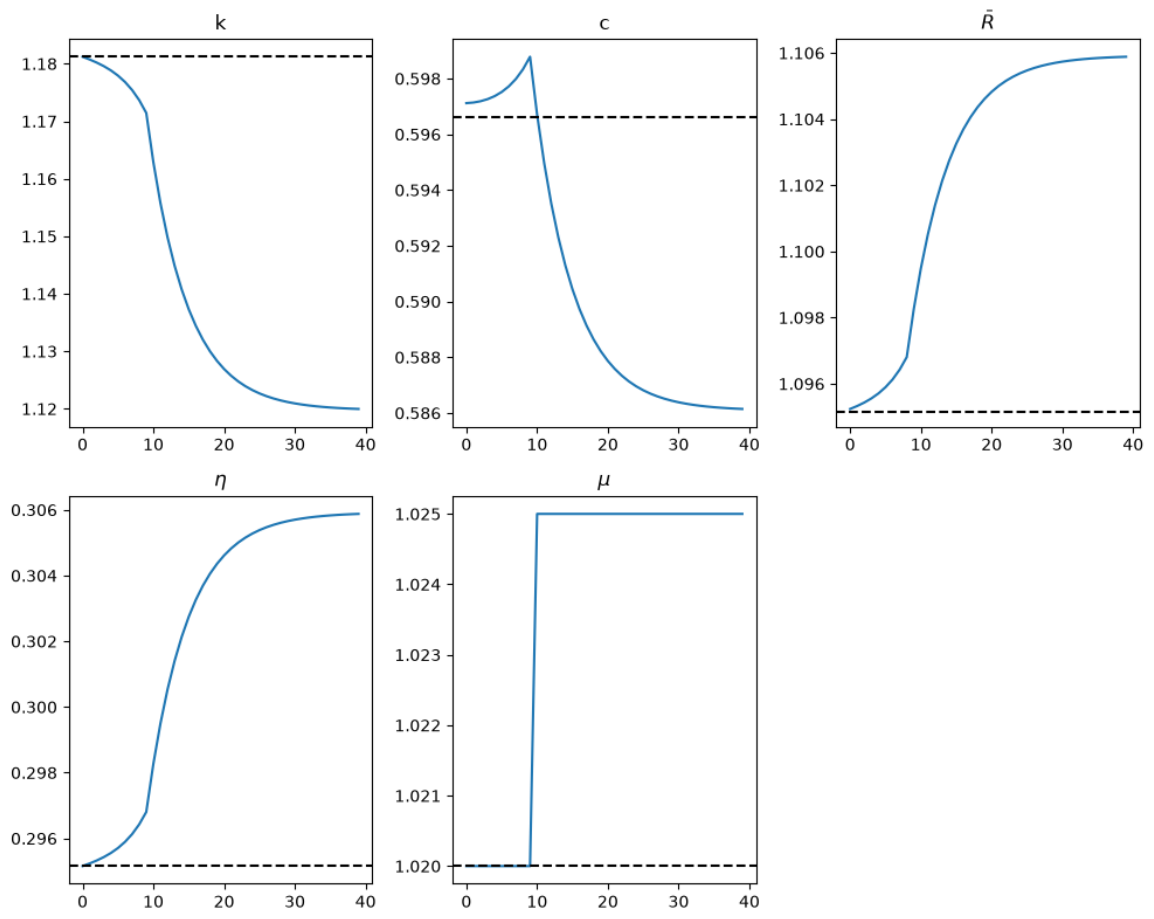
Experiment 1: A foreseen increase in μ from 1.02 to 1.025 at $t = 10$

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1),
    'μ': np.concatenate((np.repeat(1.02, 10), np.repeat(1.025, S - 9)))
}

A_path = compute_A_path(1.0, shocks, S)

experiment_model(shocks, S, model, A_path, run_min, plot_results, 'μ')
```

```
Steady-state capital: 1.1812
Steady-state consumption: 0.5966
----------------------------------------------------------------------
```

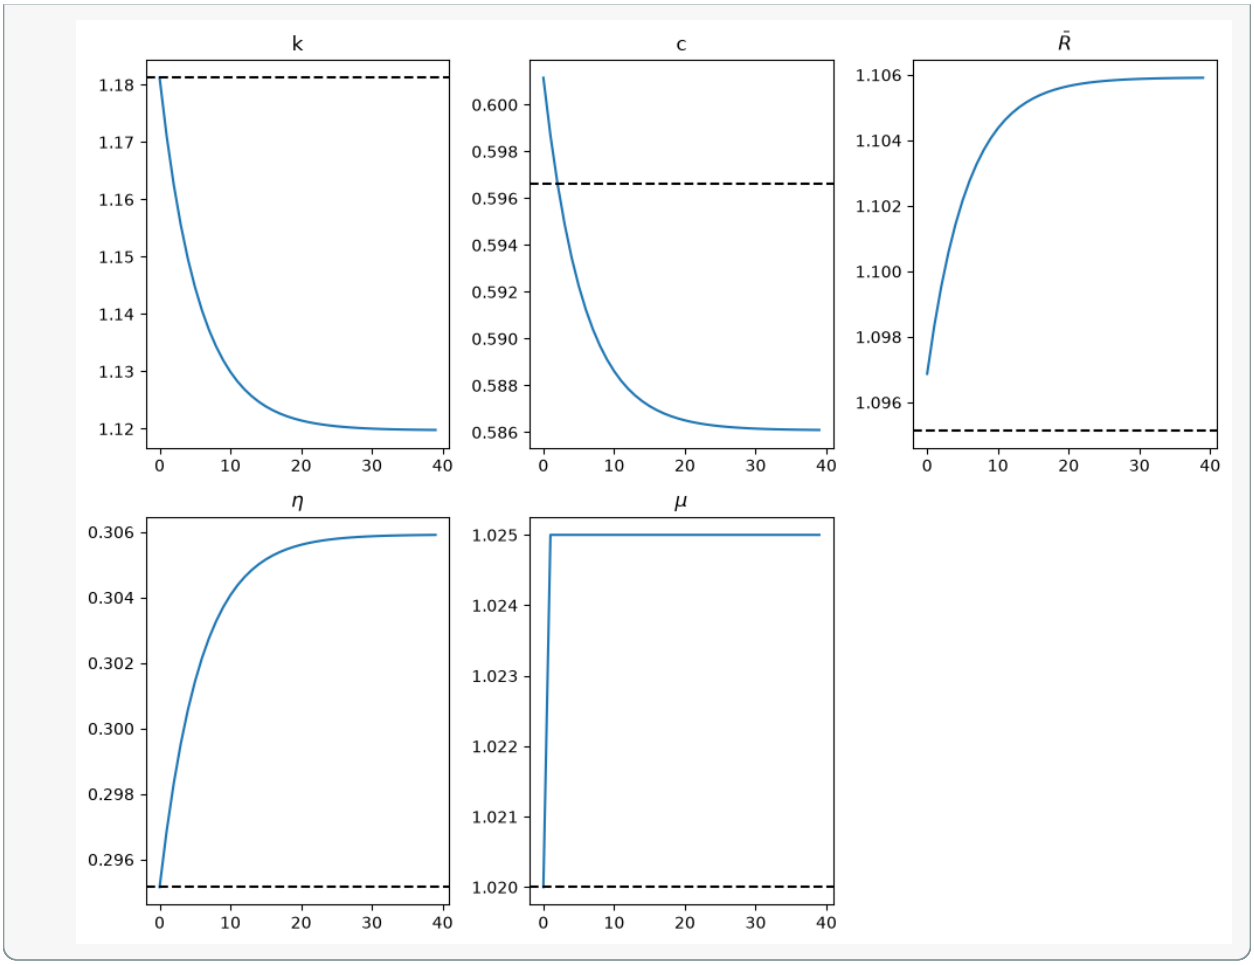


Experiment 2: An unforeseen increase in μ from 1.02 to 1.025 at $t = 0$

```
shocks = {
    'g': np.repeat(0.2, S + 1),
    'τ_c': np.repeat(0.0, S + 1),
    'τ_k': np.repeat(0.0, S + 1),
    'μ': np.concatenate((np.repeat(1.02, 1), np.repeat(1.025, S)))
}

experiment_model(shocks, S, model, A_path, run_min, plot_results, 'μ')
```

```
Steady-state capital: 1.1812
Steady-state consumption: 0.5966
------------------------------------------------------------------
```



In this this sequel *Two-Country Model with Distorting Taxes*, we study a two-country version of our one-country model that is closely related to Mendoza and Tesar [1998].

CHAPTER

THREE

TWO-COUNTRY MODEL WITH DISTORTING TAXES

113.1 Overview

This lecture is a sequel to this QuantEcon lecture *Cass-Koopmans Model with Distorting Taxes* in which we studied consequences of foreseen fiscal and technology shocks on competitive equilibrium prices and quantities in a nonstochastic version of a Cass-Koopmans growth model like the one described in this QuantEcon lecture *Cass-Koopmans Competitive Equilibrium*.

Here we study a two-country version of that model.

We construct it by putting instances of two *Cass-Koopmans Competitive Equilibrium* economies together back to back, and then opening international trade in some commodities, but not in others.

This lets us focus on some of the issues studied by Mendoza and Tesar [1998].

Let's start with some imports:

```
import numpy as np
from scipy.optimize import root
import matplotlib.pyplot as plt
from collections import namedtuple
from mpmath import mp, mpf
from warnings import warn

# Set the precision
mp.dps = 40
mp.pretty = True
```

113.2 A Two-Country Cass-Koopmans Model

This section describes a two-country version of the basic model of *The Economy*.

The model has a structure similar to ones used in the international real business cycle literature and is in the spirit of an analysis of distorting taxes by Mendoza and Tesar [1998].

We allow two countries to trade goods and claims on future goods, but not labor.

Both countries have production technologies, and consumers in each country can hold capital in either country, subject to different tax treatments.

We denote variables in the second country with asterisks (*).

Households in both countries maximize lifetime utility:

$$\sum_{t=0}^{\infty} \beta^t u(c_t) \quad \text{and} \quad \sum_{t=0}^{\infty} \beta^t u(c_t^*),$$

where $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 0$.

There are Cobb-Douglas functions with identical technology parameters in the two countries.

The world resource constraint in this two-country economy is:

$$(c_t + c_t^*) + (g_t + g_t^*) + (k_{t+1} - (1-\delta)k_t) + (k_{t+1}^* - (1-\delta)k_t^*) = f(k_t) + f(k_t^*)$$

which combines the feasibility constraints for the two countries.

Later, we will use this constraint as a global feasibility constraint in our computation.

To connect the two countries, we need to specify how capital flows across borders and how taxes are levied in different jurisdictions.

113.2.1 Capital Mobility and Taxation

A consumer in country one can hold capital in either country but pays taxes on rentals from foreign holdings of capital at the rate set by the foreign country.

Residents in both countries can purchase consumption at date t at a common Arrow-Debreu price q_t. We assume capital markets are complete.

Let B_t^f be the amount of time t goods that the representative domestic consumer raises by issuing a one-period IOU to the representative foreign consumer.

So $B_t^f > 0$ indicates the domestic consumer is borrowing from abroad at t, and $B_t^f < 0$ indicates the domestic consumer is lending abroad at t.

Hence, the budget constraint of a representative consumer in country one is:

$$\sum_{t=0}^{\infty} q_t \left(c_t + (k_{t+1} - (1-\delta)k_t) + (\tilde{k}_{t+1} - (1-\delta)\tilde{k}_t) + R_{t-1,t} B_{t-1}^f \right) \leq$$

$$\sum_{t=0}^{\infty} q_t \left((\eta_t - \tau_{kt}(\eta_t - \delta))k_t + (\eta_t^* - \tau_{kt}^*(\eta_t^* - \delta))\tilde{k}_t + (1-\tau_{nt}) w_t n_t - \tau_{ht} + B_t^f \right).$$

No-arbitrage conditions for k_t and $\tilde{k}_t$ for $t \geq 1$ imply

$$q_{t-1} = [(1-\tau_{kt})(\eta_t - \delta) + 1] q_t,$$
$$q_{t-1} = [(1-\tau_{kt}^*)(\eta_t^* - \delta) + 1] q_t,$$

which together imply that after-tax rental rates on capital are equalized across the two countries:

$$(1-\tau_{kt}^*)(\eta_t^* - \delta) = (1-\tau_{kt})(\eta_t - \delta).$$

The no-arbitrage conditions for B_t^f for $t \geq 0$ are $q_t = q_{t+1} R_{t+1,t}$, which implies that

$$q_{t-1} = q_t R_{t-1,t}$$

for $t \geq 1$.

Since domestic capital, foreign capital, and consumption loans bear the same rates of return, portfolios are indeterminate.

We can set holdings of foreign capital equal to zero in each country if we allow B_t^f to be nonzero.

This way of resolving portfolio indeterminacy is convenient because it reduces the number of initial conditions we need to specify.

Therefore, we set holdings of foreign capital equal to zero in both countries while allowing international lending.

Given an initial level B_{-1}^f of debt from the domestic country to the foreign country, and where $R_{t-1,t} = \frac{q_{t-1}}{q_t}$, international debt dynamics satisfy

$$B_t^f = R_{t-1,t} B_{t-1}^f + c_t + (k_{t+1} - (1-\delta)k_t) + g_t - f(k_t)$$

```
def Bf_path(k, c, g, model):
    """
    Compute B^{f}_t:
      Bf_t = R_{t-1} Bf_{t-1} + c_t + (k_{t+1}-(1-δ)k_t) + g_t - f(k_t)
    with Bf_0 = 0.
    """
    S = len(c) - 1
    R = c[:-1]**(-model.γ) / (model.β * c[1:]**(-model.γ))

    Bf = np.zeros(S + 1)
    for t in range(1, S + 1):
        inv = k[t] - (1 - model.δ) * k[t-1]
        Bf[t] = (
            R[t-1] * Bf[t-1] + c[t] + inv + g[t-1]
            - f(k[t-1], model))
    return Bf

def Bf_ss(c_ss, k_ss, g_ss, model):
    """
    Compute the steady-state B^f
    """
    R_ss   = 1.0 / model.β
    inv_ss = model.δ * k_ss
    num    = c_ss + inv_ss + g_ss - f(k_ss, model)
    den    = 1.0 - R_ss
    return num / den
```

and

$$c_t^* + (k_{t+1}^* - (1-\delta)k_t^*) + g_t^* - R_{t-1,t} B_{t-1}^f = f(k_t^*) - B_t^f.$$

The firms' first-order conditions in the two countries are:

$$\begin{aligned} \eta_t &= f'(k_t), \quad w_t = f(k_t) - k_t f'(k_t) \\ \eta_t^* &= f'(k_t^*), \quad w_t^* = f(k_t^*) - k_t^* f'(k_t^*). \end{aligned}$$

International trade in goods establishes:

$$\frac{q_t}{\beta^t} = \frac{u'(c_t)}{1+\tau_{ct}} = \mu^* \frac{u'(c_t^*)}{1+\tau_{ct}^*},$$

where μ^* is a nonnegative number that is a function of the Lagrange multiplier on the budget constraint for a consumer in country $*$.

We have normalized the Lagrange multiplier on the budget constraint of the domestic country to set the corresponding μ for the domestic country to unity.

```
def compute_rs(c_t, c_tp1, c_s_t, c_s_tp1, τc_t,
               τc_tp1, τc_s_t, τc_s_tp1, model):
    """
    Compute international risk sharing after trade starts.
    """

    return (c_t**(-model.γ)/(1+τc_t)) * ((1+τc_s_t)/c_s_t**(-model.γ)) - (
        c_tp1**(-model.γ)/(1+τc_tp1)) * ((1+τc_s_tp1)/c_s_tp1**(-model.γ))
```

Equilibrium requires that the following two national Euler equations be satisfied for $t \geq 0$:

$$u'(c_t) = \beta u'(c_{t+1}) \left[(1 - \tau_{kt+1})(f'(k_{t+1}) - \delta) + 1\right] \left[\frac{1 + \tau_{ct+1}}{1 + \tau_{ct}}\right],$$

$$u'(c_t^*) = \beta u'(c_{t+1}^*) \left[(1 - \tau_{kt+1}^*)(f'(k_{t+1}^*) - \delta) + 1\right] \left[\frac{1 + \tau_{ct+1}^*}{1 + \tau_{ct}^*}\right].$$

The following code computes both the domestic and foreign Euler equations.

Since they have the same form but use different variables, we can write a single function that handles both cases.

```
def compute_euler(c_t, c_tp1, τc_t,
                  τc_tp1, τk_tp1, k_tp1, model):
    """
    Compute the Euler equation.
    """
    Rbar = (1 - τk_tp1)*(f_prime(k_tp1, model) - model.δ) + 1
    return model.β * (c_tp1/c_t)**(-model.γ) * (1+τc_t)/(1+τc_tp1) * Rbar - 1
```

113.2.2 Initial condition and steady state

For the initial conditions, we choose the pre-trade allocation of capital (k_0, k_0^*) and the initial level B_{-1}^f of international debt owed by the unstarred (domestic) country to the starred (foreign) country.

113.2.3 Equilibrium steady state values

The steady state of the two-country model is characterized by two sets of equations.

First, the following equations determine the steady-state capital-labor ratios $\bar{k}$ and $\bar{k}^*$ in each country:

$$f'(\bar{k}) = \delta + \frac{\rho}{1 - \tau_k} \tag{113.1}$$

$$f'(\bar{k}^*) = \delta + \frac{\rho}{1 - \tau_k^*} \tag{113.2}$$

Given these steady-state capital-labor ratios, the domestic and foreign consumption values $\bar{c}$ and $\bar{c}^*$ are determined by:

$$(\bar{c} + \bar{c}^*) = f(\bar{k}) + f(\bar{k}^*) - \delta(\bar{k} + \bar{k}^*) - (\bar{g} + \bar{g}^*) \tag{113.3}$$

$$\bar{c} = f(\bar{k}) - \delta\bar{k} - \bar{g} - \rho\bar{B}^f \tag{113.4}$$

Equation (113.3) expresses feasibility at the steady state, while equation (113.4) represents the trade balance, including interest payments, at the steady state.

The steady-state level of debt $\bar{B}^f$ from the domestic country to the foreign country influences the consumption allocation between countries but not the total world capital stock.

We assume $\bar{B}^f = 0$ in the steady state, which gives us the following function to compute the steady-state values of capital and consumption

```
def compute_steady_state_global(model, g_ss=0.2):
    """
    Calculate steady state values for capital, consumption, and investment.
    """
    k_ss = ((1/model.β - (1-model.δ)) / (model.A * model.α)) ** (1/(model.α-1))
    c_ss = f(k_ss, model) - model.δ * k_ss - g_ss
    return k_ss, c_ss
```

Now, we can apply the residual minimization method to compute the steady-state values of capital and consumption.

Again, we minimize the residuals of the Euler equation, the global resource constraint, and the no-arbitrage condition.

```
def compute_residuals_global(z, model, shocks, T, k0_ss, k_star, Bf_star):
    """
    Compute residuals for the two-country model.
    """
    k, c, k_s, c_s = z.reshape(T+1, 4).T
    g, gs = shocks['g'], shocks['g_s']
    τc, τk = shocks['τ_c'], shocks['τ_k']
    τc_s, τk_s = shocks['τ_c_s'], shocks['τ_k_s']

    res = [k[0] - k0_ss, k_s[0] - k0_ss]

    for t in range(T):
        e_d = compute_euler(
            c[t], c[t+1],
            τc[t], τc[t+1], τk[t+1],
            k[t+1], model)

        e_f = compute_euler(
            c_s[t], c_s[t+1],
            τc_s[t], τc_s[t+1], τk_s[t+1],
            k_s[t+1], model)

        rs = compute_rs(
            c[t], c[t+1], c_s[t], c_s[t+1],
            τc[t], τc[t+1], τc_s[t], τc_s[t+1],
            model)

        # Global resource constraint
        grc = k[t+1] + k_s[t+1] - (
            f(k[t], model) + f(k_s[t], model) +
            (1-model.δ)*(k[t] + k_s[t]) -
            c[t] - c_s[t] - g[t] - gs[t]
        )

        res.extend([e_d, e_f, rs, grc])

    Bf_term = Bf_path(k, c, shocks['g'], model)[-1]
    res.append(k[T] - k_star)
    res.append(Bf_term - Bf_star)
    return np.array(res)
```

Now we plot the results

```
# Function to plot global two-country model results
def plot_global_results(k, k_s, c, c_s, shocks, model,
                        k0_ss, c0_ss, g_ss, S, T=40, shock='g',
                        # a dictionary storing sequence for lower left panel
                        ll_series='None'):
    """
    Plot results for the two-country model.
    """
    fig, axes = plt.subplots(2, 3, figsize=(10, 8))
    x = np.arange(T)
    τc, τk = shocks['τ_c'], shocks['τ_k']
    Bf = Bf_path(k, c, shocks['g'], model)

    # Compute derived series
    R_ratio = c[:-1]**(-model.γ) / (model.β * c[1:]**(-model.γ)) \
    *(1+τc[:-1])/(1+τc[1:])
    inv = k[1:] - (1-model.δ)*k[:-1]
    inv_s = k_s[1:] - (1-model.δ)*k_s[:-1]

    # Add initial conditions into the series
    R_ratio = np.append(1/model.β, R_ratio)
    c = np.append(c0_ss, c)
    c_s = np.append(c0_ss, c_s)
    k = np.append(k0_ss, k)
    k_s = np.append(k0_ss, k_s)

    # Capital
    axes[0,0].plot(x, k[:T], '-', lw=1.5)
    axes[0,0].plot(x, np.full(T, k0_ss), 'k-.', lw=1.5)
    axes[0,0].plot(x, k_s[:T], '--', lw=1.5)
    axes[0,0].set_title('k')
    axes[0,0].set_xlim(0, T-1)

    # Consumption
    axes[0,1].plot(x, c[:T], '-', lw=1.5)
    axes[0,1].plot(x, np.full(T, c0_ss), 'k-.', lw=1.5)
    axes[0,1].plot(x, c_s[:T], '--', lw=1.5)
    axes[0,1].set_title('c')
    axes[0,1].set_xlim(0, T-1)

    # Interest rate
    axes[0,2].plot(x, R_ratio[:T], '-', lw=1.5)
    axes[0,2].plot(x, np.full(T, 1/model.β), 'k-.', lw=1.5)
    axes[0,2].set_title(r'$\bar{R}$')
    axes[0,2].set_xlim(0, T-1)

    # Investment
    axes[1,0].plot(x, np.full(T, model.δ * k0_ss),
    'k-.', lw=1.5)
    axes[1,0].plot(x, np.append(model.δ*k0_ss, inv[:T-1]),
    '-', lw=1.5)
    axes[1,0].plot(x, np.append(model.δ*k0_ss, inv_s[:T-1]),
    '--', lw=1.5)
    axes[1,0].set_title('x')
    axes[1,0].set_xlim(0, T-1)

    # Shock
```

(continues on next page)

(continued from previous page)

```
    axes[1,1].plot(x, shocks[shock][:T], '-', lw=1.5)
    axes[1,1].plot(x, np.full(T, shocks[shock][0]), 'k-.', lw=1.5)
    axes[1,1].set_title(f'${shock}$')
    axes[1,1].set_ylim(-0.1, 0.5)
    axes[1,1].set_xlim(0, T-1)

    # Capital flow
    axes[1,2].plot(x, np.append(0, Bf[1:T]), lw=1.5)
    axes[1,2].plot(x, np.zeros(T), 'k-.', lw=1.5)
    axes[1,2].set_title(r'$B^{f}$')
    axes[1,2].set_xlim(0, T-1)

    plt.tight_layout()
    return fig, axes
```

As in our in the one-country model in *Cass-Koopmans Model with Distorting Taxes*, we assume a Cobb-Douglas production function:

$$F(k,1) = Ak^{\alpha}$$

```
def f(k, model, A=1):
    """
    Production function: f(k) = A * k^{α}
    """
    return A * k ** model.α

def f_prime(k, model, A=1):
    """
    Marginal product of capital: f'(k) = α * A * k^{α - 1}
    """
    return model.α * A * k ** (model.α - 1)
```

Similarly, we define the capital rental rate

$$\eta_t = f'(k_t)$$

```
def compute_η_path(k_path, model, S=100, A_path=None):
    """
    Compute η path: η_t = f'(k_t)
    with optional A_path for growth models.
    """
    A = np.ones_like(k_path) if A_path is None else np.asarray(A_path)
    η_path = np.zeros_like(k_path)
    for t in range(S):
        η_path[t] = f_prime(k_path[t], model, A[t])
    return η_path
```

Experiment 1: A foreseen increase in g from 0.2 to 0.4 at t=10

The figure below presents transition dynamics after an increase in g in the domestic economy from 0.2 to 0.4 that is announced ten periods in advance.

We start both economies from a steady state with $B_0^f = 0$.

In the figure below, the blue lines represent the domestic economy and orange dotted lines represent the foreign economy.

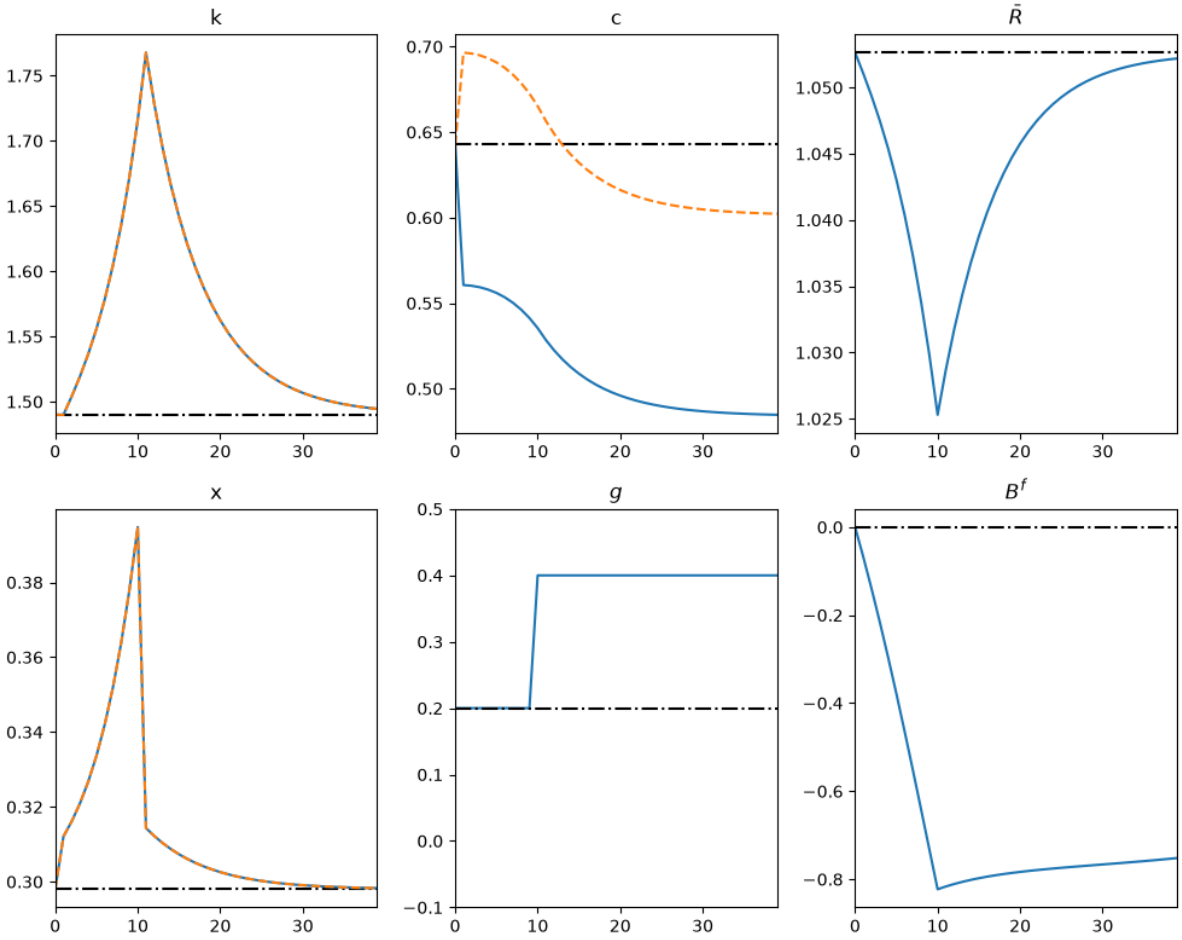
```
Model = namedtuple("Model", ["β", "γ", "δ", "α", "A"])
model = Model(β=0.95, γ=2.0, δ=0.2, α=0.33, A=1.0)
S = 100

shocks_global = {
    'g': np.concatenate((np.full(10, 0.2), np.full(S-9, 0.4))),
    'g_s': np.full(S+1, 0.2),
    'τ_c': np.zeros(S+1),
    'τ_k': np.zeros(S+1),
    'τ_c_s': np.zeros(S+1),
    'τ_k_s': np.zeros(S+1)
}
g_ss = 0.2
k0_ss, c0_ss = compute_steady_state_global(model, g_ss)

k_star = k0_ss
Bf_star = Bf_ss(c0_ss, k_star, g_ss, model)

init_glob = np.tile([k0_ss, c0_ss, k0_ss, c0_ss], S+1)
sol_glob = root(
    lambda z: compute_residuals_global(z, model, shocks_global,
                                       S, k0_ss, k_star, Bf_star),
    init_glob, tol=1e-12
)
k, c, k_s, c_s = sol_glob.x.reshape(S+1, 4).T

# Plot global results via function
plot_global_results(k, k_s, c, c_s,
                        shocks_global, model,
                        k0_ss, c0_ss, g_ss,
                        S)
plt.show()
```



At time 1, the government announces that domestic government purchases g will rise ten periods later, cutting into future private resources.

To smooth consumption, domestic households immediately increase saving, offsetting the anticipated hit to their future wealth.

In a closed economy, they would save solely by accumulating extra domestic capital; with open capital markets, they can also lend to foreigners.

Once the capital flow opens up at time 1, the no-arbitrage conditions connect adjustments of both types of saving: the increase in savings by domestic households will reduce the equilibrium return on bonds and capital in the foreign economy to prevent arbitrage opportunities.

Because no-arbitrage equalizes the ratio of marginal utilities, the resulting paths of consumption and capital are synchronized across the two economies.

Up to the date the higher g takes effect, both countries continue to build their capital stocks.

When government spending finally rises 10 periods later, domestic households begin to draw down part of that capital to cushion consumption.

Again by no-arbitrage conditions, when g actually increases, both countries reduce their investment rates.

The domestic economy, in turn, starts running current-account deficits partially to fund the increase in g.

This means that foreign households begin repaying part of their external debt by reducing their capital stock.

Experiment 2: A foreseen increase in g from 0.2 to 0.4 at t=10

We now explore the impact of an increase in capital taxation in the domestic economy 10 periods after its announcement at $t = 1$.

Because the change is anticipated, households in both countries adjust immediately–even though the tax does not take effect until period $t = 11$.

```
shocks_global = {
    'g': np.full(S+1, g_ss),
    'g_s': np.full(S+1, g_ss),
    'τ_c': np.zeros(S+1),
    'τ_k': np.concatenate((np.zeros(10), np.full(S-9, 0.2))),
    'τ_c_s': np.zeros(S+1),
    'τ_k_s': np.zeros(S+1),
}

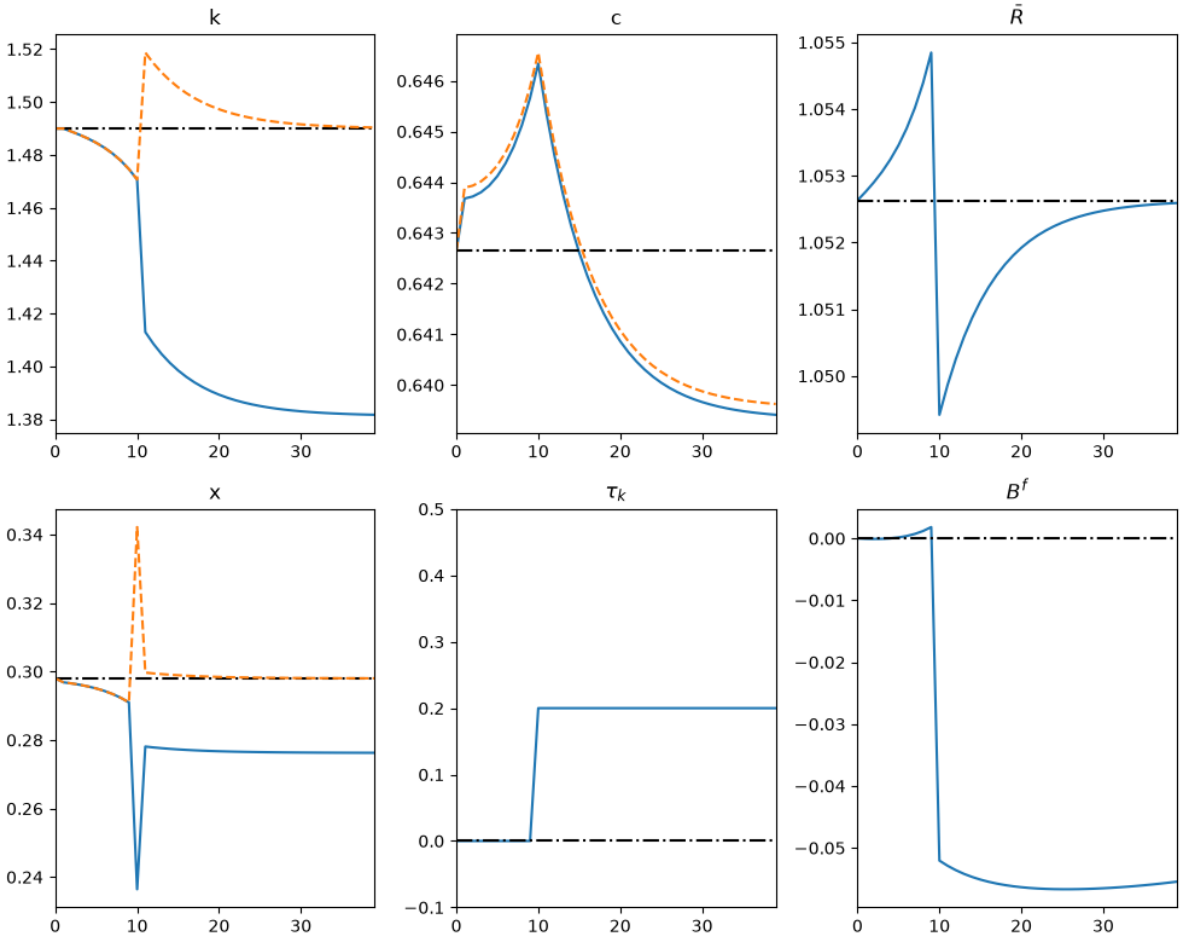
k0_ss, c0_ss = compute_steady_state_global(model, g_ss)
k_star = k0_ss
Bf_star = Bf_ss(c0_ss, k_star, g_ss, model)

init_glob = np.tile([k0_ss, c0_ss, k0_ss, c0_ss], S+1)

sol_glob = root(
    lambda z: compute_residuals_global(z, model,
            shocks_global, S, k0_ss, k_star, Bf_star),
            init_glob, tol=1e-12)

k, c, k_s, c_s = sol_glob.x.reshape(S+1, 4).T

# plot
fig, axes = plot_global_results(k, k_s, c, c_s, shocks_global, model,
                                k0_ss, c0_ss, g_ss, S, shock='τ_k')
plt.tight_layout()
plt.show()
```



After the tax increase is announced, domestic households foresee lower after-tax returns on capital, so they shift toward higher present consumption and allow the domestic capital stock to decline.

This shrinkage of the world capital supply drives the global real interest rate upward, prompting foreign households to raise current consumption as well.

Prior to the actual tax hike, the domestic economy finances part of its consumption by importing capital, generating a current-account deficit.

When τ_k finally rises, international arbitrage leads investors to reallocate capital quickly toward the untaxed foreign market, compressing the yield on bonds everywhere.

The bond-rate drop reflects the lower after-tax return on domestic capital and the higher foreign capital stock, which depresses its marginal product.

Foreign households fund their capital purchases by borrowing abroad, creating a pronounced current-account deficit and a buildup of external debt.

After the policy change, both countries move smoothly toward a new steady state in which:

- Consumption levels in each economy settle below their pre-announcement paths.
- Capital stocks differ just enough to equalize after-tax returns across borders.

Despite carrying positive net liabilities, the foreign country enjoys higher steady-state consumption because its larger capital stock yields greater output.

The episode demonstrates how open capital markets transmit a domestic tax shock internationally: capital flows and

interest-rate movements share the burden, smoothing consumption adjustments in both the taxed and untaxed economies over time.

Exercise 113.2.1

In this exercise, replace the plot for x_t with η_t to replicate the figure in [Ljungqvist and Sargent, 2018].

Compare the figures for k_t and η_t and discuss the economic intuition.

Solution

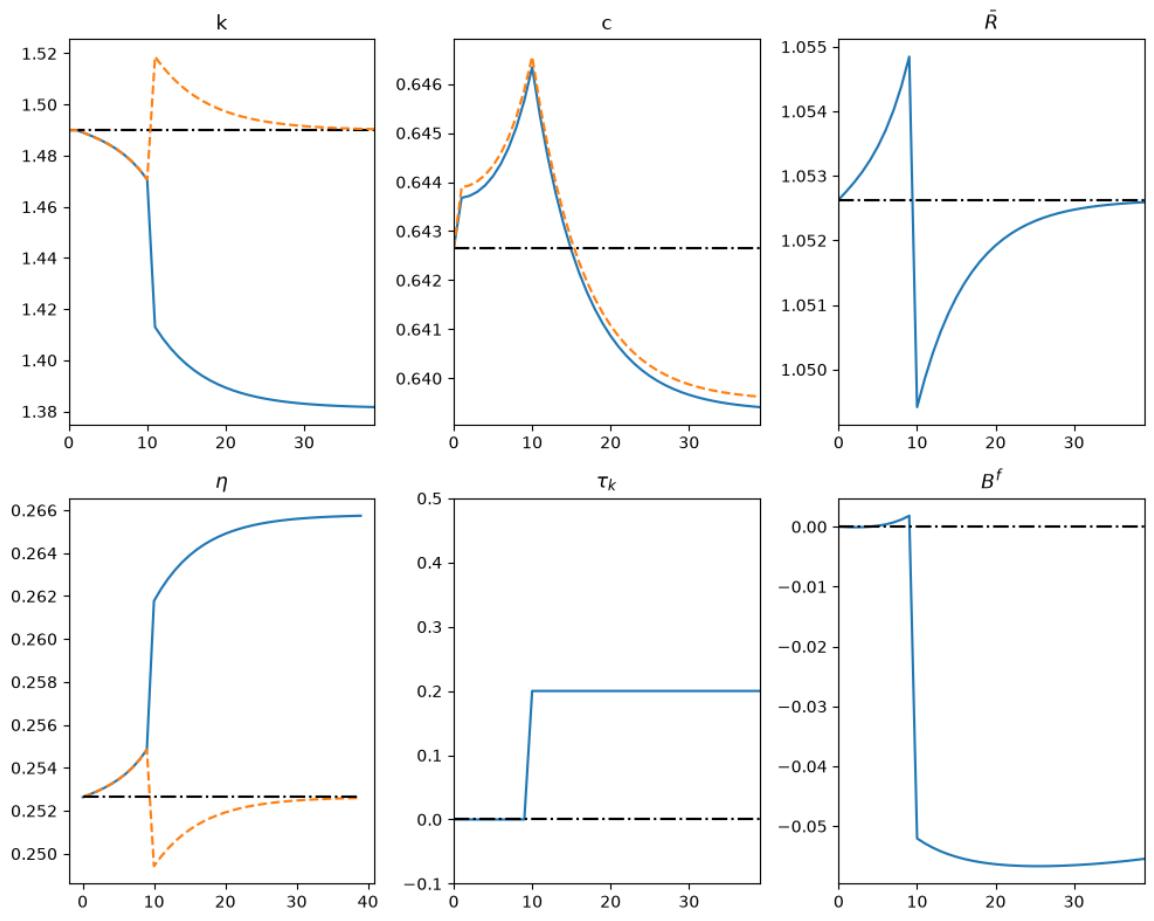
Here is one solution.

```
fig, axes = plot_global_results(k, k_s, c, c_s, shocks_global, model,
                                k0_ss, c0_ss, g_ss, S, shock='τ_k')

# Clear the plot for x_t
axes[1,0].cla()

# Plot η_t
axes[1,0].plot(compute_η_path(k, model)[:40])
axes[1,0].plot(compute_η_path(k_s, model)[:40], '--')
axes[1,0].plot(np.full(40, f_prime(k_s, model)[0]), 'k-.', lw=1.5)
axes[1,0].set_title(r'$\eta$')

plt.tight_layout()
plt.show()
```



When capital k_t decreases in the domestic country after the tax shock, the rental rate η_t increases in that country.

This happens because when capital becomes scarcer, its marginal product rises.

Part XV

Multiple Agent Models

CHAPTER

FOUR

A LAKE MODEL OF EMPLOYMENT AND UNEMPLOYMENT

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *A Lake Model of Employment and Unemployment*
  - *Overview*
  - *The model*
  - *Implementation*
  - *Dynamics of an individual worker*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

114.1 Overview

This lecture describes what has come to be called a *lake model*.

The lake model is a basic tool for modeling unemployment.

It allows us to analyze

- flows between unemployment and employment
- how these flows influence steady state employment and unemployment rates

It is a good model for interpreting monthly labor department reports on gross and net jobs created and destroyed.

The "lakes" in the model are the pools of employed and unemployed.

The "flows" between the lakes are caused by

- firing and hiring
- entry and exit from the labor force

For the first part of this lecture, the parameters governing transitions into and out of unemployment and employment are exogenous.

Later, we'll determine some of these transition rates endogenously using the *McCall search model*.

We'll also use some nifty concepts like ergodicity, which provides a fundamental link between *cross-sectional* and *long run time series* distributions.

These concepts will help us build an equilibrium model of ex-ante homogeneous workers whose different luck generates variations in their ex post experiences.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
from typing import NamedTuple
from quantecon.distributions import BetaBinomial
from functools import partial
import jax.scipy.stats as stats
```

114.1.1 Prerequisites

Before working through what follows, we recommend you read the *lecture on finite Markov chains*.

You will also need some basic *linear algebra* and probability.

114.2 The model

The economy is inhabited by a very large number of ex-ante identical workers.

The workers live forever, spending their lives moving between unemployment and employment.

Their rates of transition between employment and unemployment are governed by the following parameters:

- λ, the job finding rate for currently unemployed workers
- α, the dismissal rate for currently employed workers
- b, the entry rate into the labor force
- d, the exit rate from the labor force

The growth rate of the labor force evidently equals $g = b - d$.

114.2.1 Aggregate variables

We want to derive the dynamics of the following aggregates:

- E_t, the total number of employed workers at date t
- U_t, the total number of unemployed workers at t
- N_t, the number of workers in the labor force at t

114.2.2 Laws of motion for stock variables

We begin by constructing laws of motion for the aggregate variables E_t, U_t, N_t.

Of the mass of workers E_t who are employed at date t,

- $(1-d)E_t$ will remain in the labor force
- of these, $(1-\alpha)(1-d)E_t$ will remain employed

Of the mass of workers U_t workers who are currently unemployed,

- $(1-d)U_t$ will remain in the labor force
- of these, $(1-d)\lambda U_t$ will become employed

Therefore, the number of workers who will be employed at date $t+1$ will be

$$E_{t+1} = (1-d)(1-\alpha)E_t + (1-d)\lambda U_t$$

A similar analysis implies

$$U_{t+1} = (1-d)\alpha E_t + (1-d)(1-\lambda)U_t + b(E_t + U_t)$$

The value $b(E_t + U_t)$ is the mass of new workers entering the labor force unemployed.

The total stock of workers $N_t = E_t + U_t$ evolves as

$$N_{t+1} = (1+b-d)N_t = (1+g)N_t$$

Letting $X_t := \begin{pmatrix} U_t \\ E_t \end{pmatrix}$, the law of motion for X is

$$X_{t+1} = AX_t \quad \text{where} \quad A := \begin{bmatrix} (1-d)(1-\lambda) + b & (1-d)\alpha + b \\ (1-d)\lambda & (1-d)(1-\alpha) \end{bmatrix}$$

This law tells us how total employment and unemployment evolve over time.

114.2.3 Laws of motion for rates

Now let's derive the law of motion for rates.

We want to track the values of the following objects:

- The employment rate $e_t := E_t/N_t$.
- The unemployment rate $u_t := U_t/N_t$.

(Here and below, capital letters represent aggregates and lowercase letters represent rates)

To get these we can divide both sides of $X_{t+1} = AX_t$ by N_{t+1} to get

$$\begin{bmatrix} U_{t+1}/N_{t+1} \\ E_{t+1}/N_{t+1} \end{bmatrix} = \frac{1}{1+g} A \begin{bmatrix} U_t/N_t \\ E_t/N_t \end{bmatrix}$$

Letting

$$x_t := \begin{pmatrix} u_t \\ e_t \end{pmatrix} = \begin{pmatrix} U_t/N_t \\ E_t/N_t \end{pmatrix}$$

we can also write this as

$$x_{t+1} = Rx_t \quad \text{where} \quad R := \frac{1}{1+g} A$$

You can check that $e_t + u_t = 1$ implies that $e_{t+1} + u_{t+1} = 1$.

This follows from the fact that the columns of R sum to 1.

114.3 Implementation

Let's code up these equations.

114.3.1 Model

To begin, we set up a class called `LakeModel` that stores the primitives α, λ, b, d.

```
class LakeModel(NamedTuple):
    """
    Parameters for the lake model
    """
    λ: float
    α: float
    b: float
    d: float
    A: jnp.ndarray
    R: jnp.ndarray
    g: float


def create_lake_model(
        λ: float = 0.283,      # job finding rate
        α: float = 0.013,      # separation rate
        b: float = 0.0124,     # birth rate
        d: float = 0.00822     # death rate
    ) -> LakeModel:
    """
    Create a LakeModel instance with default parameters.

    Computes and stores the transition matrices A and R,
    and the labor force growth rate g.

    """
    # Compute growth rate
```

(continues on next page)

(continued from previous page)

```
    g = b - d

    # Compute transition matrix A
    A = jnp.array([
        [(1-d) * (1-λ) + b,  (1-d) * α + b],
        [(1-d) * λ,          (1-d) * (1-α)]
    ])

    # Compute normalized transition matrix R
    R = A / (1 + g)

    return LakeModel(λ=λ, α=α, b=b, d=d, A=A, R=R, g=g)
```

The default parameter values are:

- $\alpha = 0.013$ and $\lambda = 0.283$ are based on [Davis *et al.*, 2006]
- $b = 0.0124$ and $d = 0.00822$ are set to match monthly birth and death rates, respectively, in the U.S. population

As an experiment, let's create two instances, one with $\alpha = 0.013$ and another with $\alpha = 0.03$

```
model = create_lake_model()
print(f"Default α: {model.α}")
print(f"A matrix:\n{model.A}")
print(f"R matrix:\n{model.R}")
```

```
Default α: 0.013
A matrix:
[[0.7235063  0.02529314]
 [0.28067374 0.97888684]]
R matrix:
[[0.7204946  0.02518786]
 [0.27950543 0.97481215]]
```

```
model_new = create_lake_model(α=0.03)
print(f"New α: {model_new.α}")
print(f"New A matrix:\n{model_new.A}")
print(f"New R matrix:\n{model_new.R}")
```

```
New α: 0.03
New A matrix:
[[0.7235063  0.0421534 ]
 [0.28067374 0.9620266 ]]
New R matrix:
[[0.7204946  0.04197793]
 [0.27950543 0.9580221 ]]
```

114.3.2 Code for dynamics

We will also use a specialized function to generate time series in an efficient JAX-compatible manner.

Iteratively generating time series is somewhat nontrivial in JAX because arrays are immutable.

Here we use `lax.scan`, which allows the function to be jit-compiled.

Readers who prefer to skip the details can safely continue reading after the function definition.

```
@partial(jax.jit, static_argnames=['f', 'num_steps'])
def generate_path(f, initial_state, num_steps, **kwargs):
    """
    Generate a time series by repeatedly applying an update rule.

    Given a map f, initial state x_0, and model parameters, this
    function computes and returns the sequence {x_t}_{t=0}^{T-1} when

        x_{t+1} = f(x_t, **kwargs)

    Args:
        f: Update function mapping (x_t, **kwargs) -> x_{t+1}
        initial_state: Initial state x_0
        num_steps: Number of time steps T to simulate
        **kwargs: Optional extra arguments passed to f

    Returns:
        Array of shape (dim(x), T) containing the time series path
        [x_0, x_1, x_2, ..., x_{T-1}]
    """

    def update_wrapper(state, t):
        """
        Wrapper function that adapts f for use with JAX scan.
        """
        next_state = f(state, **kwargs)
        return next_state, state

    _, path = jax.lax.scan(update_wrapper,
                    initial_state, jnp.arange(num_steps))
    return path.T
```

Here are functions to update X_t and x_t.

```
def stock_update(X: jnp.ndarray, model: LakeModel) -> jnp.ndarray:
    """Apply transition matrix to get next period's stocks."""
    λ, α, b, d, A, R, g = model
    return A @ X

def rate_update(x: jnp.ndarray, model: LakeModel) -> jnp.ndarray:
    """Apply normalized transition matrix for next period's rates."""
    λ, α, b, d, A, R, g = model
    return R @ x
```

114.3.3 Aggregate dynamics

Let's run a simulation under the default parameters starting from $X_0 = (12, 138)$.

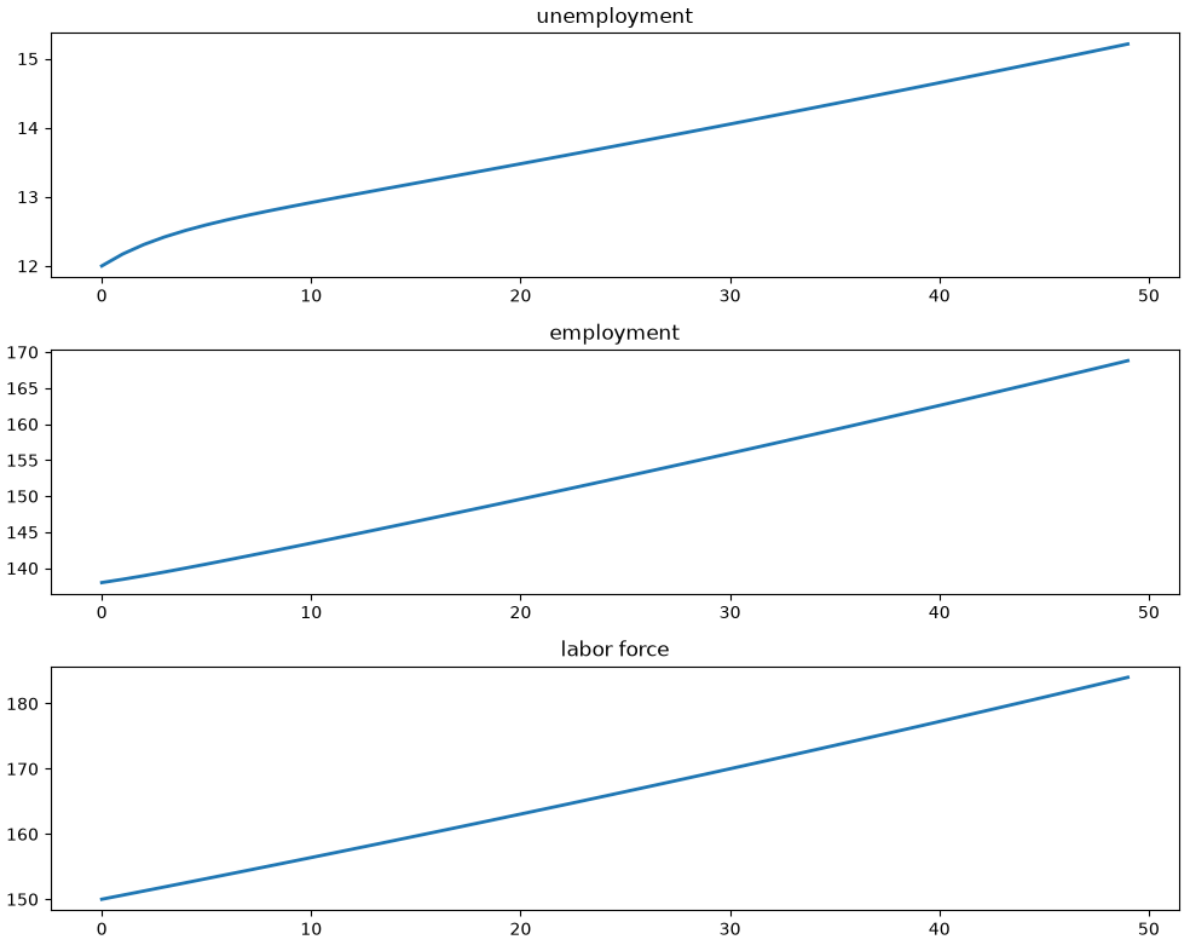
We will plot the sequences $\{E_t\}$, $\{U_t\}$ and $\{N_t\}$.

```
N_0 = 150       # Population
e_0 = 0.92      # Initial employment rate
u_0 = 1 - e_0   # Initial unemployment rate
T = 50          # Simulation length

U_0 = u_0 * N_0
E_0 = e_0 * N_0

# Generate X path
X_0 = jnp.array([U_0, E_0])
X_path = generate_path(stock_update, X_0, T, model=model)

# Plot
fig, axes = plt.subplots(3, 1, figsize=(10, 8))
titles = ['unemployment', 'employment', 'labor force']
data = [X_path[0, :], X_path[1, :], X_path.sum(0)]
for ax, title, series in zip(axes, titles, data):
    ax.plot(series, lw=2)
    ax.set_title(title)
plt.tight_layout()
plt.show()
```



The aggregates E_t and U_t don't converge because their sum $E_t + U_t$ grows at rate g.

114.3.4 Rate dynamics

On the other hand, the vector of employment and unemployment rates x_t can be in a steady state $\bar{x}$ if there exists an $\bar{x}$ such that

- $\bar{x} = R\bar{x}$
- the components satisfy $\bar{e} + \bar{u} = 1$

This equation tells us that a steady state level $\bar{x}$ is an eigenvector of R associated with a unit eigenvalue.

The following function can be used to compute the steady state.

```
@jax.jit
def rate_steady_state(model: LakeModel) -> jnp.ndarray:
    r"""
    Finds the steady state of the system :math:`x_{t+1} = R x_{t}`
    by computing the eigenvector corresponding to the largest eigenvalue.

    By the Perron-Frobenius theorem, since :math:`R` is a non-negative
    matrix with columns summing to 1 (a stochastic matrix), the largest
    eigenvalue equals 1 and the corresponding eigenvector gives the steady state.
    """
```

(continues on next page)

(continued from previous page)

```
    λ, α, b, d, A, R, g = model
    eigenvals, eigenvec = jnp.linalg.eig(R)

    # Find the eigenvector corresponding to the largest eigenvalue
    # (which is 1 for a stochastic matrix by Perron-Frobenius theorem)
    max_idx = jnp.argmax(jnp.abs(eigenvals))

    # Get the corresponding eigenvector
    steady_state = jnp.real(eigenvec[:, max_idx])

    # Normalize to ensure positive values and sum to 1
    steady_state = jnp.abs(steady_state)
    steady_state = steady_state / jnp.sum(steady_state)

    return steady_state
```

We also have $x_t \to \bar{x}$ as $t \to \infty$ provided that the remaining eigenvalue of R has modulus less than 1.

This is the case for our default parameters:

```
model = create_lake_model()
e, f = jnp.linalg.eigvals(model.R)
print(f"Eigenvalue magnitudes: {abs(e):.2f}, {abs(f):.2f}")
```

```
Eigenvalue magnitudes: 0.70, 1.00
```

Let's look at the convergence of the unemployment and employment rates to steady state levels (dashed line)

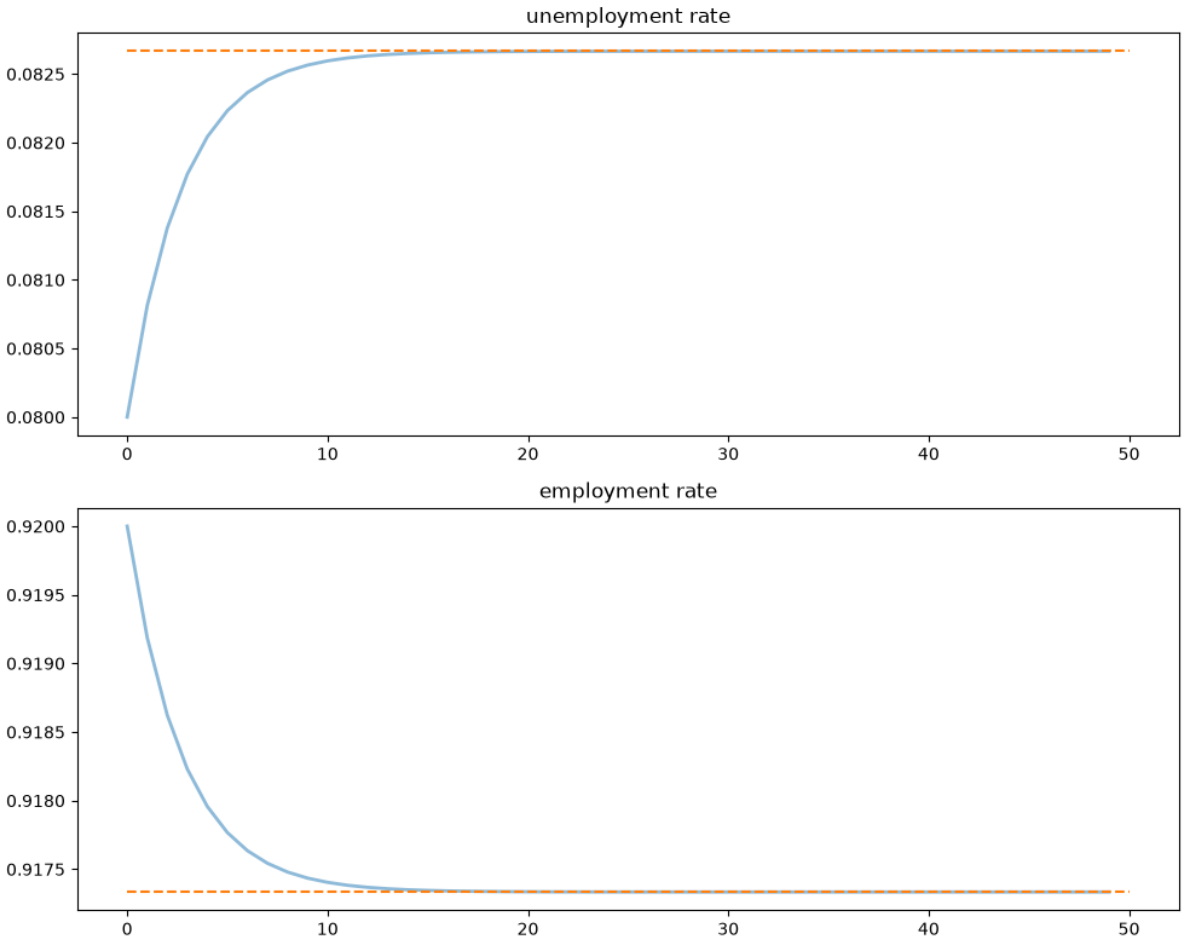
```
xbar = rate_steady_state(model)

fig, axes = plt.subplots(2, 1, figsize=(10, 8))
x_0 = jnp.array([u_0, e_0])
x_path = generate_path(rate_update, x_0, T, model=model)

titles = ['unemployment rate', 'employment rate']

for i, title in enumerate(titles):
    axes[i].plot(x_path[i, :], lw=2, alpha=0.5)
    axes[i].hlines(xbar[i], 0, T, color='C1', linestyle='--')
    axes[i].set_title(title)

plt.tight_layout()
plt.show()
```



Exercise 114.3.1

Use JAX's `vmap` to compute steady-state unemployment rates for a range of job finding rates λ (from 0.1 to 0.5), and plot the relationship.

Solution

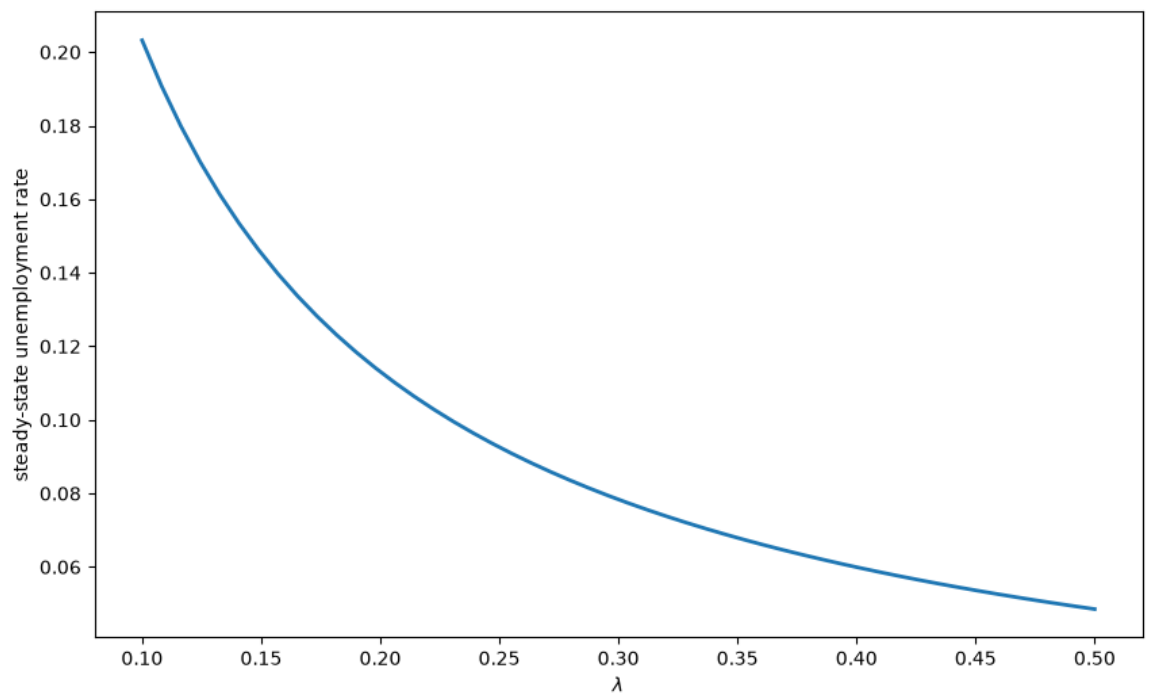
Here is one solution

```
@jax.jit
def compute_unemployment_rate(λ_val):
    """Computes steady-state unemployment for a given λ"""
    model = create_lake_model(λ=λ_val)
    steady_state = rate_steady_state(model)
    return steady_state[0]

# Use vmap to compute for multiple λ values
λ_values = jnp.linspace(0.1, 0.5, 50)
unemployment_rates = jax.vmap(compute_unemployment_rate)(λ_values)

# Plot the results
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(λ_values, unemployment_rates, lw=2)
```

```
ax.set_xlabel(r'$\lambda$')
ax.set_ylabel('steady-state unemployment rate')
plt.show()
```



114.4 Dynamics of an individual worker

An individual worker's employment dynamics are governed by a *finite state Markov process*.

The worker can be in one of two states:

- $s_t = 0$ means unemployed
- $s_t = 1$ means employed

Let's start off under the assumption that $b = d = 0$.

The associated transition matrix is then

$$P = \begin{pmatrix} 1-\lambda & \lambda \\ \alpha & 1-\alpha \end{pmatrix}$$

Let ψ_t denote the *marginal distribution* over employment/unemployment states for the worker at time t.

As usual, we regard it as a row vector.

We know *from an earlier discussion* that ψ_t follows the law of motion

$$\psi_{t+1} = \psi_t P$$

We also know from the *lecture on finite Markov chains* that if $\alpha \in (0, 1)$ and $\lambda \in (0, 1)$, then P has a unique stationary distribution, denoted here by ψ^*.

The unique stationary distribution satisfies

$$\psi^*[0] = \frac{\alpha}{\alpha + \lambda}$$

Not surprisingly, probability mass on the unemployment state increases with the dismissal rate and falls with the job finding rate.

114.4.1 Ergodicity

Let's look at a typical lifetime of employment-unemployment spells.

We want to compute the average amounts of time an infinitely lived worker would spend employed and unemployed.

Let

$$\bar{s}_{u,T} := \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}\{s_t = 0\}$$

and

$$\bar{s}_{e,T} := \frac{1}{T} \sum_{t=1}^{T} \mathbb{1}\{s_t = 1\}$$

(As usual, $\mathbb{1}\{Q\} = 1$ if statement Q is true and 0 otherwise)

These are the fraction of time a worker spends unemployed and employed, respectively, up until period T.

If $\alpha \in (0, 1)$ and $\lambda \in (0, 1)$, then P is *ergodic*, and hence we have

$$\lim_{T \to \infty} \bar{s}_{u,T} = \psi^*[0] \quad \text{and} \quad \lim_{T \to \infty} \bar{s}_{e,T} = \psi^*[1]$$

with probability one.

Inspection tells us that P is exactly the transpose of R under the assumption $b = d = 0$.

Thus, the percentages of time that an infinitely lived worker spends employed and unemployed equal the fractions of workers employed and unemployed in the steady state distribution.

114.4.2 Convergence rate

How long does it take for time series sample averages to converge to cross-sectional averages?

We can investigate this by simulating the Markov chain.

Let's plot the path of the sample averages over 5,000 periods

```
def markov_update(state, P, key):
    """
    Sample next state from transition probabilities.
    """
    probs = P[state]
    state_new = jax.random.choice(key,
                        a=jnp.arange(len(probs)),
                        p=probs)
    return state_new

model_markov = create_lake_model(d=0, b=0)
```

(continues on next page)

(continued from previous page)

```
T = 5000  # Simulation length

α, λ = model_markov.α, model_markov.λ

P = jnp.array([[1 - λ,         λ],
               [    α,     1 - α]])

xbar = rate_steady_state(model_markov)

# Simulate the Markov chain - we need a different approach for random updates
key = jax.random.key(0)

def simulate_markov(P, initial_state, T, key):
    """Simulate Markov chain for T periods"""
    keys = jax.random.split(key, T)

    def scan_fn(state, key):
        next_state = markov_update(state, P, key)
        return next_state, state

    _, path = jax.lax.scan(scan_fn, initial_state, keys)
    return path

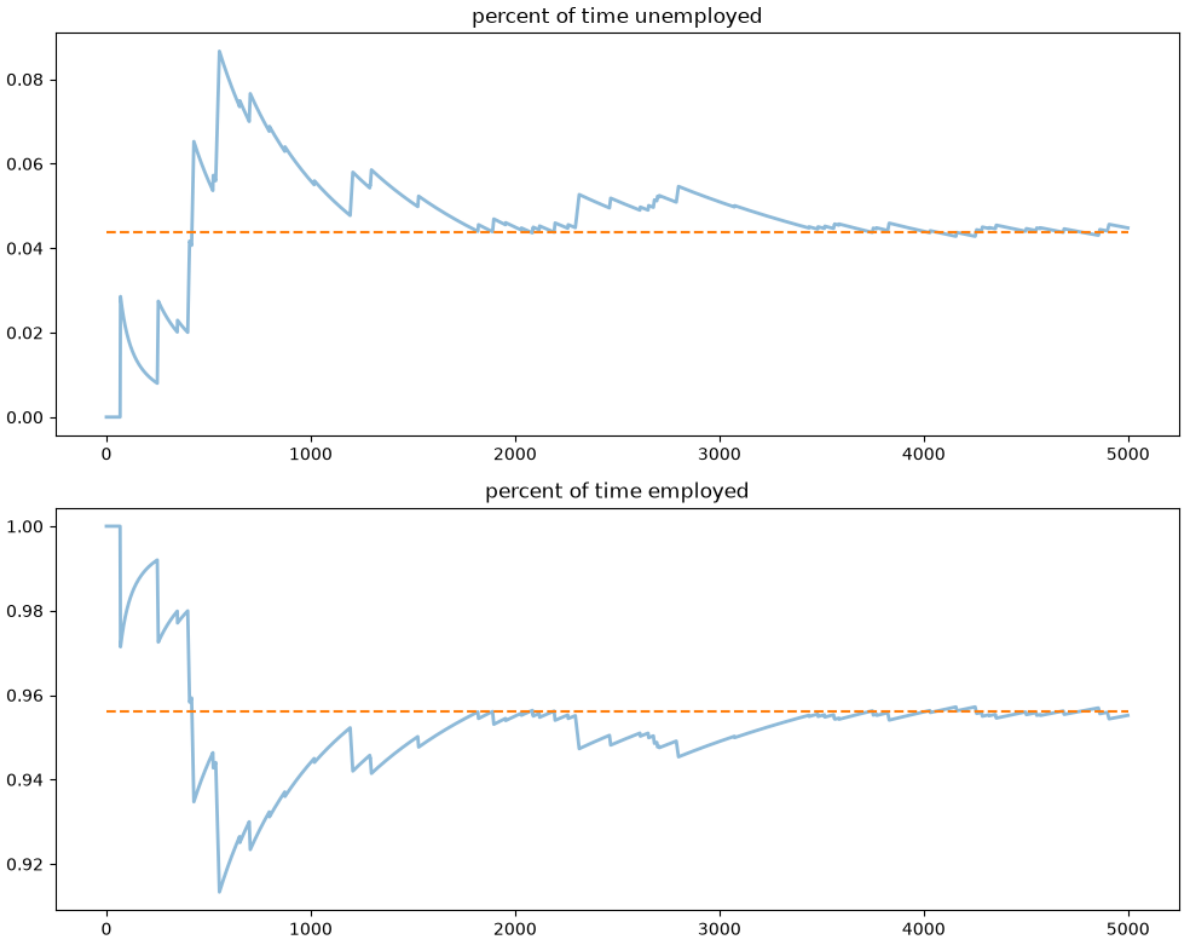
s_path = simulate_markov(P, 1, T, key)

fig, axes = plt.subplots(2, 1, figsize=(10, 8))
s_bar_e = jnp.cumsum(s_path) / jnp.arange(1, T+1)
s_bar_u = 1 - s_bar_e

to_plot = [s_bar_u, s_bar_e]
titles = ['percent of time unemployed', 'percent of time employed']

for i, plot in enumerate(to_plot):
    axes[i].plot(plot, lw=2, alpha=0.5)
    axes[i].hlines(xbar[i], 0, T, color='C1', linestyle='--')
    axes[i].set_title(titles[i])

plt.tight_layout()
plt.show()
```



The stationary probabilities are given by the dashed line.

In this case it takes much of the sample for these two objects to converge.

This is largely due to the high persistence in the Markov chain.

114.5 Exercises

Exercise 114.5.1

Consider an economy with an initial stock of workers $N_0 = 100$ at the steady state level of employment in the baseline parameterization.

Suppose that in response to new legislation the hiring rate reduces to $\lambda = 0.2$.

Plot the transition dynamics of the unemployment and employment stocks for 50 periods.

Plot the transition dynamics for the rates.

How long does the economy take to converge to its new steady state?

What is the new steady state level of employment?

Solution

We begin by constructing the model with default parameters and finding the initial steady state

```
model_initial = create_lake_model()
x0 = rate_steady_state(model_initial)
print(f"Initial Steady State: {x0}")
```

```
Initial Steady State: [0.08266623 0.9173338 ]
```

Initialize the simulation values

```
N0 = 100
T = 50
```

New legislation changes λ to 0.2

```
model_ex2 = create_lake_model(λ=0.2)
xbar = rate_steady_state(model_ex2)  # new steady state

# Simulate paths
X_path = generate_path(stock_update, x0 * N0, T, model=model_ex2)
x_path = generate_path(rate_update, x0, T, model=model_ex2)
print(f"New Steady State: {xbar}")
```

```
New Steady State: [0.11309288 0.8869071 ]
```

Now plot stocks

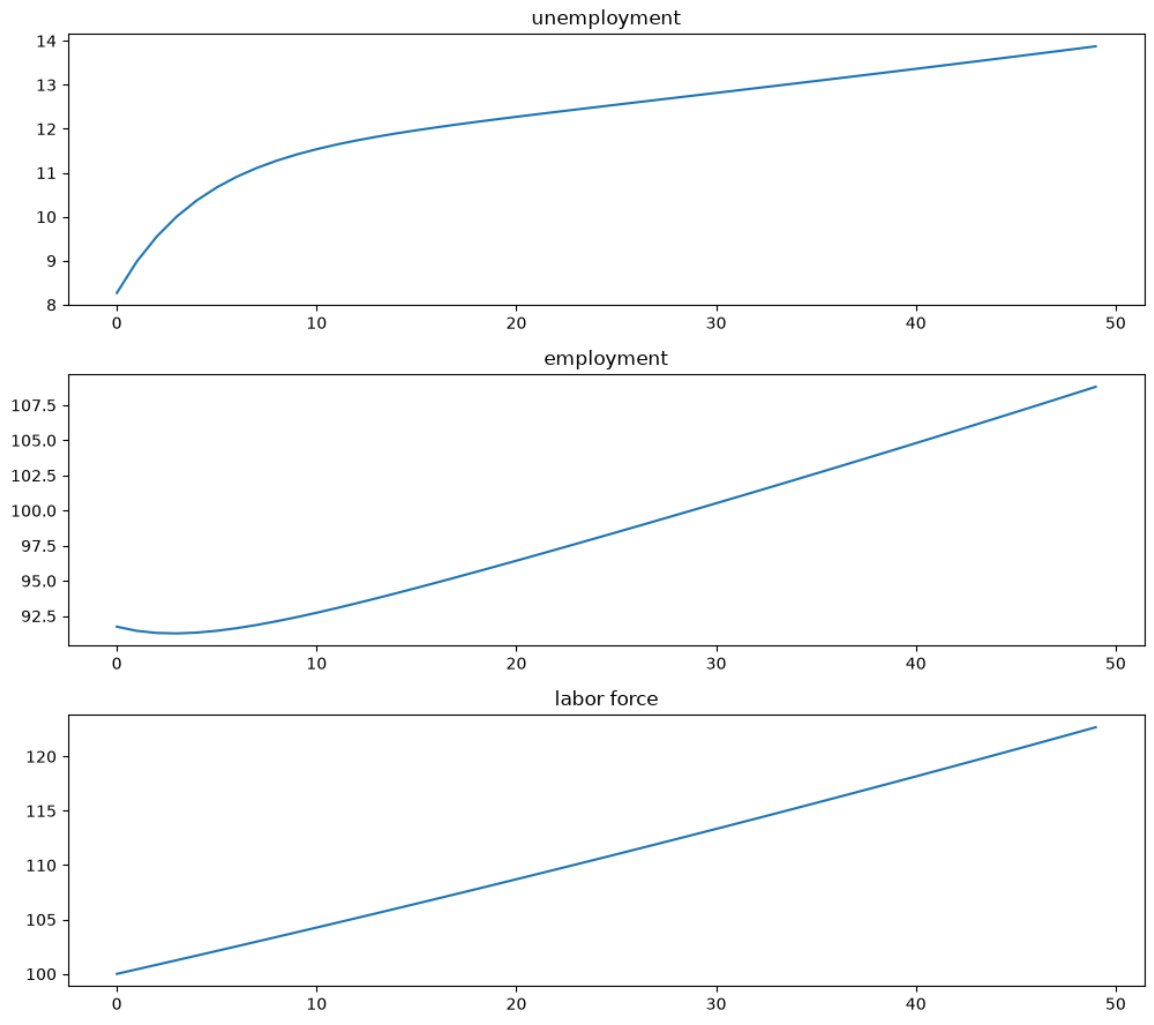
```
fig, axes = plt.subplots(3, 1, figsize=[10, 9])

axes[0].plot(X_path[0, :])
axes[0].set_title('unemployment')

axes[1].plot(X_path[1, :])
axes[1].set_title('employment')

axes[2].plot(X_path.sum(0))
axes[2].set_title('labor force')

plt.tight_layout()
plt.show()
```



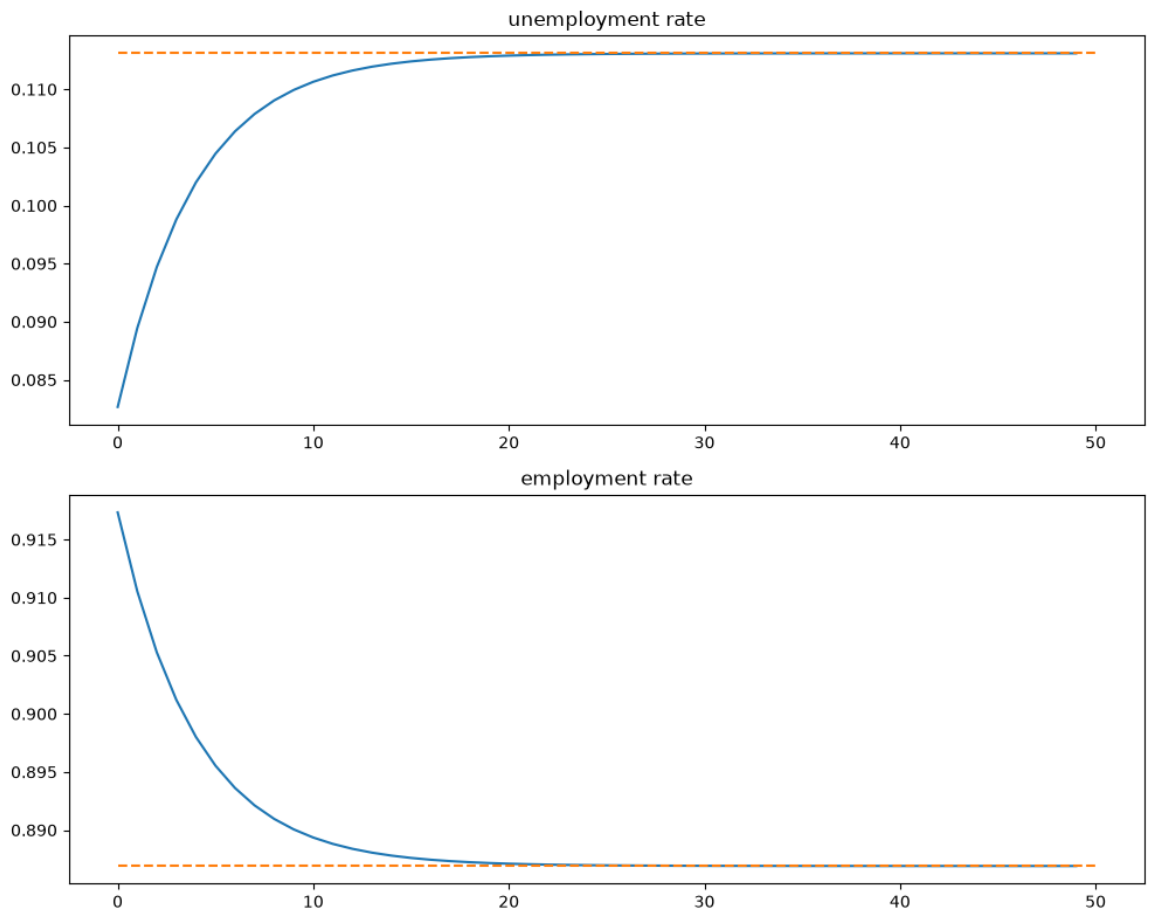
And how the rates evolve

```
fig, axes = plt.subplots(2, 1, figsize=(10, 8))

titles = ['unemployment rate', 'employment rate']

for i, title in enumerate(titles):
    axes[i].plot(x_path[i, :])
    axes[i].hlines(xbar[i], 0, T, color='C1', linestyle='--')
    axes[i].set_title(title)

plt.tight_layout()
plt.show()
```



We see that it takes 20 periods for the economy to converge to its new steady state levels.

Exercise 114.5.2

Consider an economy with an initial stock of workers $N_0 = 100$ at the steady state level of employment in the baseline parameterization.

Suppose that for 20 periods the birth rate was temporarily high ($b = 0.025$) and then returned to its original level.

Plot the transition dynamics of the unemployment and employment stocks for 50 periods.

Plot the transition dynamics for the rates.

How long does the economy take to return to its original steady state?

Solution

This exercise has the economy experiencing a boom in entrances to the labor market and then later returning to the original levels.

For 20 periods the economy has a new entry rate into the labor market.

Let's start off at the baseline parameterization and record the steady state

```
model_baseline = create_lake_model()
x0 = rate_steady_state(model_baseline)
N0 = 100
T = 50
```

Here are the other parameters:

```
b_hat = 0.025
T_hat = 20
```

Let's increase b to the new value and simulate for 20 periods

```
model_high_b = create_lake_model(b=b_hat)

# Simulate stocks and rates for first 20 periods
X_path1 = generate_path(stock_update, x0 * N0, T_hat, model=model_high_b)
x_path1 = generate_path(rate_update, x0, T_hat, model=model_high_b)
```

Now we reset b to the original value and then, using the state after 20 periods for the new initial conditions, we simulate for the additional 30 periods

```
# Use final state from period 20 as initial condition
X_path2 = generate_path(stock_update, X_path1[:, -1], T-T_hat,
                        model=model_baseline)
x_path2 = generate_path(rate_update, x_path1[:, -1], T-T_hat,
                        model=model_baseline)
```

Finally, we combine these two paths and plot

```
# Combine paths
X_path = jnp.hstack([X_path1, X_path2[:, 1:]])
x_path = jnp.hstack([x_path1, x_path2[:, 1:]])

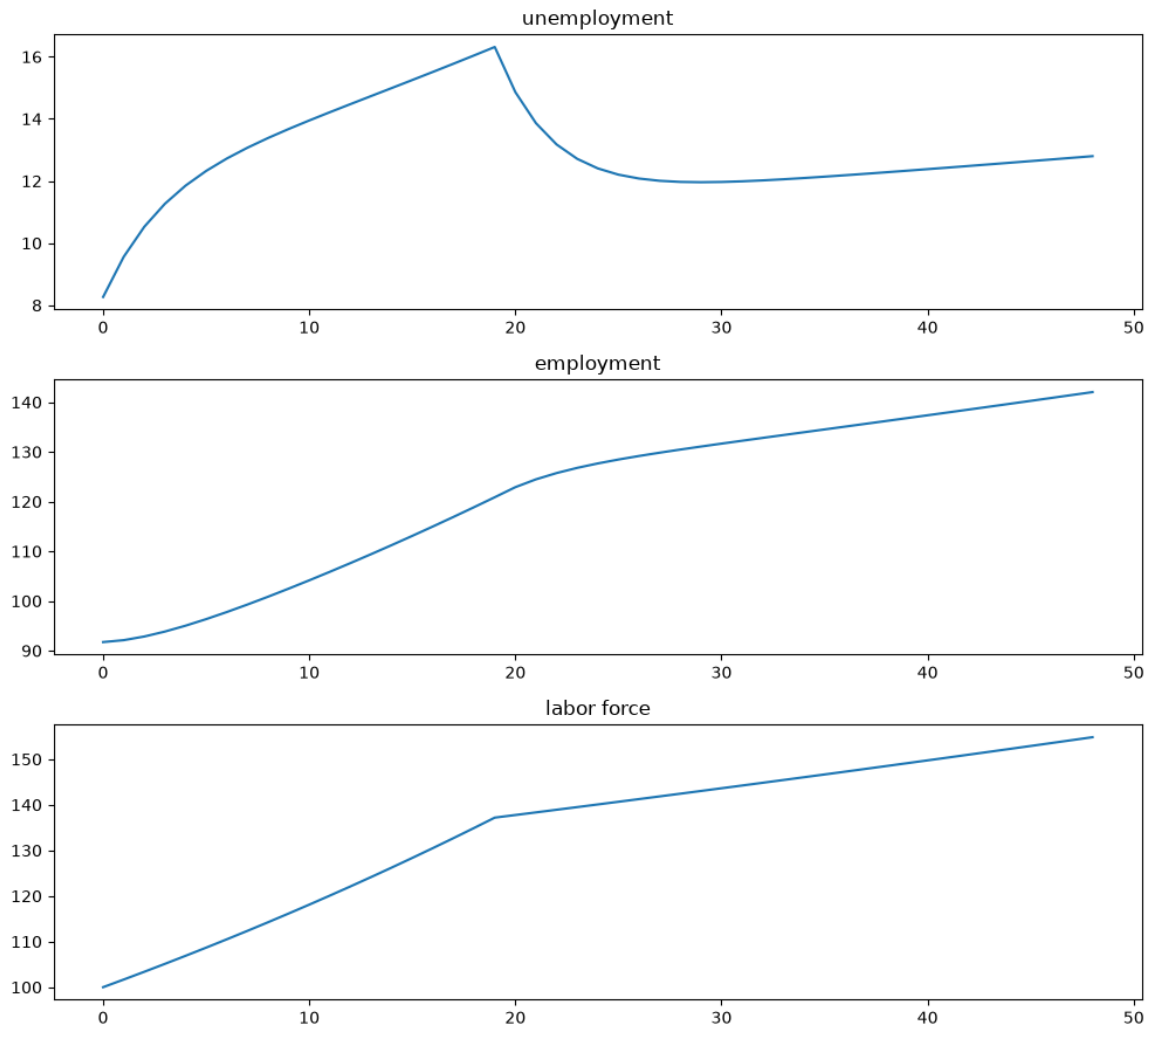
fig, axes = plt.subplots(3, 1, figsize=[10, 9])

axes[0].plot(X_path[0, :])
axes[0].set_title('unemployment')

axes[1].plot(X_path[1, :])
axes[1].set_title('employment')

axes[2].plot(X_path.sum(0))
axes[2].set_title('labor force')

plt.tight_layout()
plt.show()
```



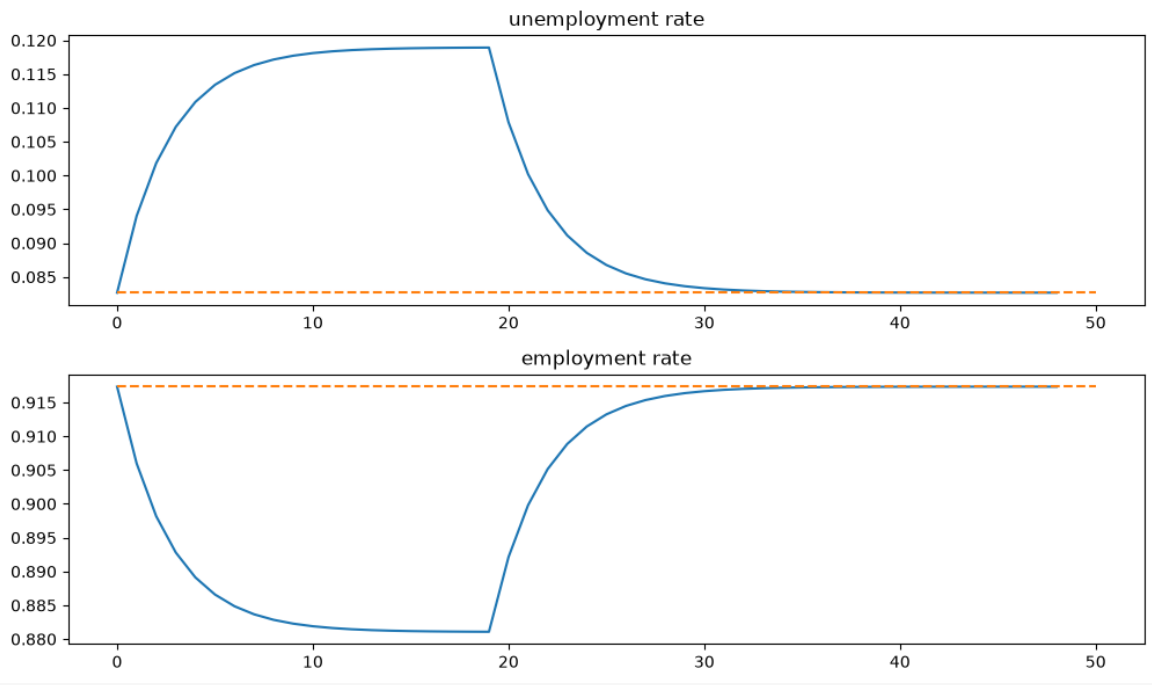
And the rates

```
fig, axes = plt.subplots(2, 1, figsize=[10, 6])

titles = ['unemployment rate', 'employment rate']

for i, title in enumerate(titles):
    axes[i].plot(x_path[i, :])
    axes[i].hlines(x0[i], 0, T, color='C1', linestyle='--')
    axes[i].set_title(title)

plt.tight_layout()
plt.show()
```


unemployment rate
0.120
0.115
0.110
0.105
0.100
0.095
0.090
0.085
0
10
20
30
40
50
employment rate
0.915
0.910
0.905
0.900
0.895
0.890
0.885
0.880
0
10
20
30
40
50

CHAPTER

FIVE

LAKE MODEL WITH AN ENDOGENOUS JOB FINDING RATE

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Lake Model with an Endogenous Job Finding Rate*
  - *Overview*
  - *Set Up*
  - *Fiscal policy*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon jax
```

115.1 Overview

This lecture is a continuation of the *lake model lecture*.

We recommend you read that lecture first before proceeding with this one.

In the previous lecture, we studied a lake model of unemployment and employment where the transition rates between states were exogenous parameters.

In this lecture, we extend the model by making the job finding rate endogenous.

Specifically, the transition rate from unemployment to employment will be determined by the McCall search model [McCall, 1970].

Let's start with some imports:

```
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
from typing import NamedTuple
from quantecon.distributions import BetaBinomial
from functools import partial
import jax.scipy.stats as stats
```

115.2 Set Up

The basic structure of the model will be as discussed in the *lake model lecture*.

The only difference is that the hiring rate is endogenous, determined by the decisions of optimizing agents inhabiting a McCall search model [McCall, 1970] with IID wage offers and job separation at rate α.

115.2.1 Reservation wage

In the model, the optimal policy is characterized by a reservation wage $\bar{w}$

- If the wage offer w in hand is greater than or equal to $\bar{w}$, then the worker accepts.
- Otherwise, the worker rejects.

The reservation wage depends on the wage offer distribution and the parameters

- α, the separation rate
- β, the discount factor
- γ, the offer arrival rate
- c, unemployment compensation

The wage offer distribution will be a discretized version of a lognormal distribution.

We first define a function to create such a discrete distribution.

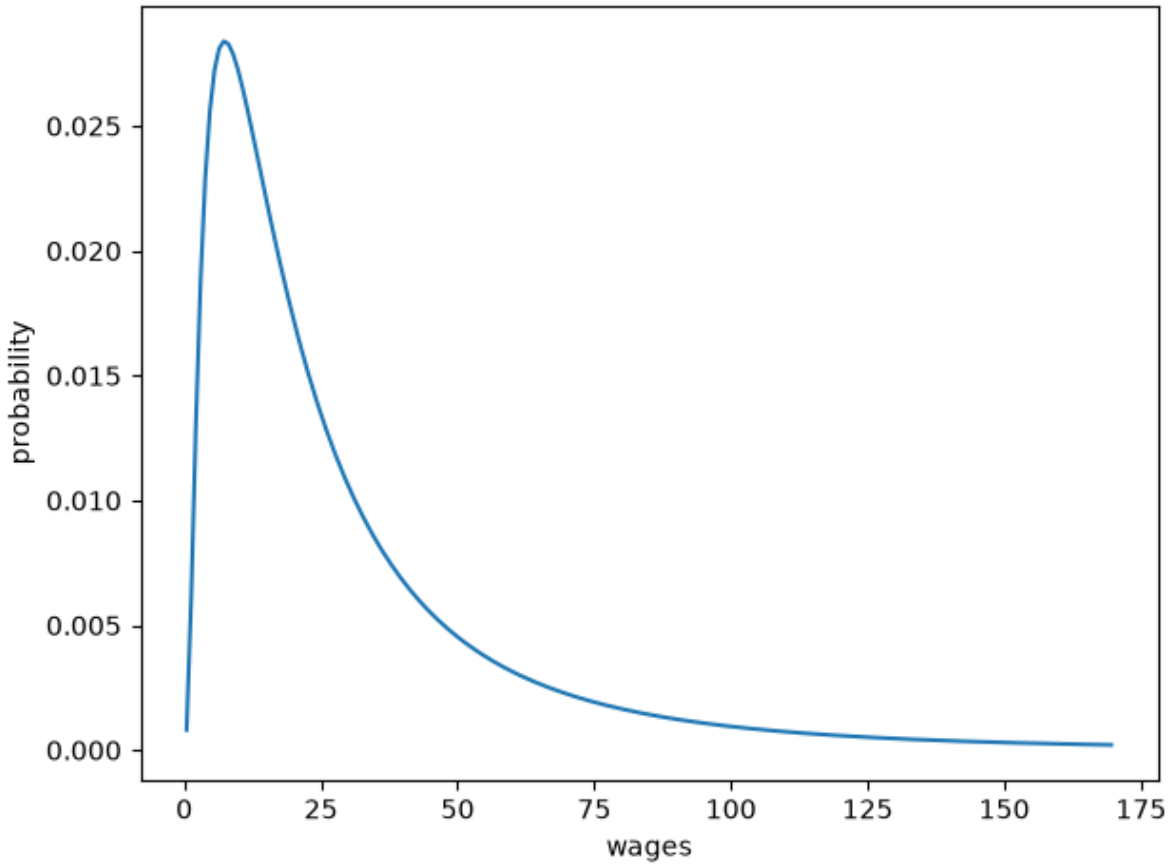
```
def create_wage_distribution(
        max_wage: float,
        wage_grid_size: int,
        log_wage_mean: float
    ):
    """
    Creates a discretized version of a lognormal density LN(log(m),1), where
    m is log_wage_mean.

    """
    w_vec_temp = jnp.linspace(1e-8, max_wage, wage_grid_size + 1)
    cdf = stats.norm.cdf(
        jnp.log(w_vec_temp), loc=jnp.log(log_wage_mean), scale=1
    )
    pdf = cdf[1:] - cdf[:-1]
    p_vec = pdf / pdf.sum()
    w_vec = (w_vec_temp[1:] + w_vec_temp[:-1]) / 2
    return w_vec, p_vec
```

The cell below creates a discretized $LN(\log(20), 1)$ wage distribution and plots it.

```
w_vec, p_vec = create_wage_distribution(170, 200, 20)

fig, ax = plt.subplots()
ax.plot(w_vec, p_vec)
ax.set_xlabel('wages')
ax.set_ylabel('probability')
plt.tight_layout()
plt.show()
```



Now we organize the code for solving the McCall model, given a set of parameters.

For background on the model and our solution method, see the *lecture on the McCall model with separation*

Our first step is to define the utility function and the McCall model data structure.

```
def u(c, σ=2.0):
    return jnp.where(c > 0, (c**(1 - σ) - 1) / (1 - σ), -10e6)


class McCallModel(NamedTuple):
    """
    Stores the parameters for the McCall search model
    """
    α: float            # Job separation rate
    β: float            # Discount rate
    γ: float            # Job offer rate
```

(continues on next page)

(continued from previous page)

```
    c: float            # Unemployment compensation
    σ: float            # Utility parameter
    w_vec: jnp.ndarray  # Possible wage values
    p_vec: jnp.ndarray  # Probabilities over w_vec


def create_mccall_model(
        α=0.2, β=0.98, γ=0.7, c=6.0, σ=2.0,
        w_vec=None,
        p_vec=None
    ) -> McCallModel:
    if w_vec is None:
        n = 60  # Number of possible outcomes for wage
        # Wages between 10 and 20
        w_vec = jnp.linspace(10, 20, n)
        a, b = 600, 400  # Shape parameters
        dist = BetaBinomial(n-1, a, b)
        p_vec = jnp.array(dist.pdf())
    return McCallModel(
        α=α, β=β, γ=γ, c=c, σ=σ, w_vec=w_vec, p_vec=p_vec
    )
```

Next, we implement the Bellman operator

```
def T(mcm: McCallModel, V, U):
    """
    Update the Bellman equations.
    """
    α, β, γ, c, σ = mcm.α, mcm.β, mcm.γ, mcm.c, mcm.σ
    w_vec, p_vec = mcm.w_vec, mcm.p_vec

    V_new = u(w_vec, σ) + β * ((1 - α) * V + α * U)
    U_new = u(c, σ) + β * (1 - γ) * U + β * γ * (jnp.maximum(U, V) @ p_vec)

    return V_new, U_new
```

Now we define the value function iteration solver.

We'll use a compiled while loop for extra speed.

```
@jax.jit
def solve_mccall_model(mcm: McCallModel, tol=1e-5, max_iter=2000):
    """
    Iterates to convergence on the Bellman equations.
    """
    def cond(state):
        V, U, i, error = state
        return jnp.logical_and(error > tol, i < max_iter)

    def update(state):
        V, U, i, error = state
        V_new, U_new = T(mcm, V, U)
        error_1 = jnp.max(jnp.abs(V_new - V))
        error_2 = jnp.abs(U_new - U)
        error_new = jnp.maximum(error_1, error_2)
        return V_new, U_new, i + 1, error_new
```

(continues on next page)

(continued from previous page)

```
    # Initial state
    V_init = jnp.ones(len(mcm.w_vec))
    U_init = 1.0
    i_init = 0
    error_init = tol + 1

    init_state = (V_init, U_init, i_init, error_init)
    V_final, U_final, _, _ = jax.lax.while_loop(
        cond, update, init_state
    )
    return V_final, U_final
```

115.2.2 Lake model code

We also need the lake model functions from the previous lecture to compute steady state unemployment rates:

```
class LakeModel(NamedTuple):
    """
    Parameters for the lake model
    """
    λ: float
    α: float
    b: float
    d: float
    A: jnp.ndarray
    R: jnp.ndarray
    g: float


def create_lake_model(
        λ: float = 0.283,      # job finding rate
        α: float = 0.013,      # separation rate
        b: float = 0.0124,     # birth rate
        d: float = 0.00822     # death rate
    ) -> LakeModel:
    """
    Create a LakeModel instance with default parameters.

    Computes and stores the transition matrices A and R,
    and the labor force growth rate g.

    """
    # Compute growth rate
    g = b - d

    # Compute transition matrix A
    A = jnp.array([
        [(1-d) * (1-λ) + b, (1-d) * α + b],
        [(1-d) * λ,         (1-d) * (1-α)]
    ])

    # Compute normalized transition matrix R
    R = A / (1 + g)

    return LakeModel(λ=λ, α=α, b=b, d=d, A=A, R=R, g=g)
```

(continues on next page)

(continued from previous page)

```
@jax.jit
def rate_steady_state(model: LakeModel) -> jnp.ndarray:
    r"""
    Finds the steady state of the system :math:`x_{t+1} = R x_{t}`
    by computing the eigenvector corresponding to the largest eigenvalue.

    By the Perron-Frobenius theorem, since :math:`R` is a non-negative
    matrix with columns summing to 1 (a stochastic matrix), the largest
    eigenvalue equals 1 and the corresponding eigenvector gives the steady state.
    """
    λ, α, b, d, A, R, g = model
    eigenvals, eigenvec = jnp.linalg.eig(R)

    # Find the eigenvector corresponding to the largest eigenvalue
    # (which is 1 for a stochastic matrix by Perron-Frobenius theorem)
    max_idx = jnp.argmax(jnp.abs(eigenvals))

    # Get the corresponding eigenvector
    steady_state = jnp.real(eigenvec[:, max_idx])

    # Normalize to ensure positive values and sum to 1
    steady_state = jnp.abs(steady_state)
    steady_state = steady_state / jnp.sum(steady_state)

    return steady_state
```

115.2.3 Linking the McCall search model to the lake model

Suppose that all workers inside a lake model behave according to the McCall search model.

The exogenous probability of leaving employment remains α.

But their optimal decision rules determine the probability λ of leaving unemployment.

This is now

$$\lambda = \gamma \mathbb{P}\{w_t \geq \bar{w}\} = \gamma \sum_{w' \geq \bar{w}} p(w') \tag{115.1}$$

Here

- $\bar{w}$ is the reservation wage determined by the parameters and
- p is the wage offer distribution.

Wage offers across the population of workers are independent draws from p.

Here we calculate λ at the default parameters:

```
mcm = create_mccall_model(w_vec=w_vec, p_vec=p_vec)
V, U = solve_mccall_model(mcm)
w_idx = jnp.searchsorted(V - U, 0)
w_bar = jnp.where(w_idx == len(V), jnp.inf, mcm.w_vec[w_idx])
λ = mcm.γ * jnp.sum(p_vec * (w_vec > w_bar))
print(f"Job finding rate at default paramters ={λ}.")
```

```
Job finding rate at default paramters =0.4510621726512909.
```

115.3 Fiscal policy

In this section, we will put the lake model to work, examining outcomes associated with different levels of unemployment compensation.

Our aim is to find an optimal level of unemployment insurance.

We assume that the government sets unemployment compensation c.

The government imposes a lump-sum tax τ sufficient to finance total unemployment payments.

To attain a balanced budget at a steady state, taxes, the steady state unemployment rate u, and the unemployment compensation rate must satisfy

$$\tau = uc$$

The lump-sum tax applies to everyone, including unemployed workers.

- The post-tax income of an employed worker with wage w is $w - \tau$.
- The post-tax income of an unemployed worker is $c - \tau$.

For each specification (c, τ) of government policy, we can solve for the worker's optimal reservation wage.

This determines λ via (115.1) evaluated at post tax wages, which in turn determines a steady state unemployment rate $u(c, \tau)$.

For a given level of unemployment benefit c, we can solve for a tax that balances the budget in the steady state

$$\tau = u(c, \tau)c$$

To evaluate alternative government tax-unemployment compensation pairs, we require a welfare criterion.

We use a steady state welfare criterion

$$W := e\,\mathbb{E}[V \,|\, \text{employed}] + u\,U$$

where the notation V and U is as defined above and the expectation is at the steady state.

115.3.1 Computing optimal unemployment insurance

Now we set up the infrastructure to compute optimal unemployment insurance levels.

First, we define a container for the economy's parameters:

```
class Economy(NamedTuple):
    """Parameters for the economy"""
    α: float
    b: float
    d: float
    β: float
    γ: float
    σ: float
    log_wage_mean: float
    wage_grid_size: int
```

(continues on next page)

(continued from previous page)

```
    max_wage: float

def create_economy(
        α=0.013,
        b=0.0124,
        d=0.00822,
        β=0.98,
        γ=1.0,
        σ=2.0,
        log_wage_mean=20,
        wage_grid_size=200,
        max_wage=170
    ) -> Economy:
    """
    Create an economy with a set of default values"""
    return Economy(α=α, b=b, d=d, β=β, γ=γ, σ=σ,
                           log_wage_mean=log_wage_mean,
                           wage_grid_size=wage_grid_size,
                           max_wage=max_wage)
```

Next, we define a function that computes optimal worker behavior given policy parameters:

```
@jax.jit
def compute_optimal_quantities(
        c: float,
        τ: float,
        economy: Economy,
        w_vec: jnp.array,
        p_vec: jnp.array
    ):
    """
    Compute the reservation wage, job finding rate and value functions
    of the workers given c and τ.

    """
    mcm = create_mccall_model(
        α=economy.α,
        β=economy.β,
        γ=economy.γ,
        c=c-τ,          # Post tax compensation
        σ=economy.σ,
        w_vec=w_vec-τ,  # Post tax wages
        p_vec=p_vec
    )

    # Compute reservation wage under given parameters
    V, U = solve_mccall_model(mcm)
    w_idx = jnp.searchsorted(V - U, 0)
    w_bar = jnp.where(w_idx == len(V), jnp.inf, mcm.w_vec[w_idx])

    # Compute job finding rate
    λ = economy.γ * jnp.sum(p_vec * (w_vec - τ > w_bar))

    return w_bar, λ, V, U
```

This function computes the steady state outcomes given unemployment insurance and tax levels:

```
@jax.jit
def compute_steady_state_quantities(
        c, τ, economy: Economy, w_vec, p_vec
    ):
    """
    Compute the steady state unemployment rate given c and τ using optimal
    quantities from the McCall model and computing corresponding steady
    state quantities

    """

    # Find optimal values and policies by solving the McCall model, as well
    # as the corresponding job finding rate.
    w_bar, λ, V, U = compute_optimal_quantities(c, τ, economy, w_vec, p_vec)

    # Set up a lake model using the given parameters and the job finding rate.
    model = create_lake_model(λ=λ, α=economy.α, b=economy.b, d=economy.d)

    # Compute steady state employment and unemployment rates from this lake
    # model.
    u, e = rate_steady_state(model)

    # Compute expected lifetime value conditional on being employed.
    mask = (w_vec - τ > w_bar)
    w = jnp.sum(V * p_vec * mask) / jnp.sum(p_vec * mask)
    # Compute steady state welfare.
    welfare = e * w + u * U

    return e, u, welfare
```

We need a function to find the tax rate that balances the government budget:

```
def find_balanced_budget_tax(c, economy: Economy, w_vec, p_vec):
    """
    Find the tax rate that will induce a balanced budget given unemployment
    compensation c.

    """

    def steady_state_budget(t):
        """
        For given tax rate t, compute the budget surplus.

        """
        e, u, w = compute_steady_state_quantities(c, t, economy, w_vec, p_vec)
        return t - u * c

    # Use a simple bisection method to find the tax rate that balances the
    # budget (but setting the surplus to zero

    t_low, t_high = 0.0, 0.9 * c
    tol = 1e-6
    max_iter = 100
    for i in range(max_iter):
        t_mid = (t_low + t_high) / 2
        budget = steady_state_budget(t_mid)
        if abs(budget) < tol:
```

(continues on next page)

(continued from previous page)

```
            return t_mid
        elif budget < 0:
            t_low = t_mid
        else:
            t_high = t_mid

    return t_mid
```

Now we compute how employment, unemployment, taxes, and welfare vary with the unemployment compensation rate:

```
# Create economy and wage distribution
economy = create_economy()
w_vec, p_vec = create_wage_distribution(
    economy.max_wage, economy.wage_grid_size, economy.log_wage_mean
)

# Levels of unemployment insurance we wish to study
c_vec = jnp.linspace(5, 140, 40)

tax_vec = []
unempl_vec = []
empl_vec = []
welfare_vec = []

for c in c_vec:
    t = find_balanced_budget_tax(c, economy, w_vec, p_vec)
    e_rate, u_rate, welfare = compute_steady_state_quantities(
        c, t, economy, w_vec, p_vec
    )
    tax_vec.append(t)
    unempl_vec.append(u_rate)
    empl_vec.append(e_rate)
    welfare_vec.append(welfare)
```

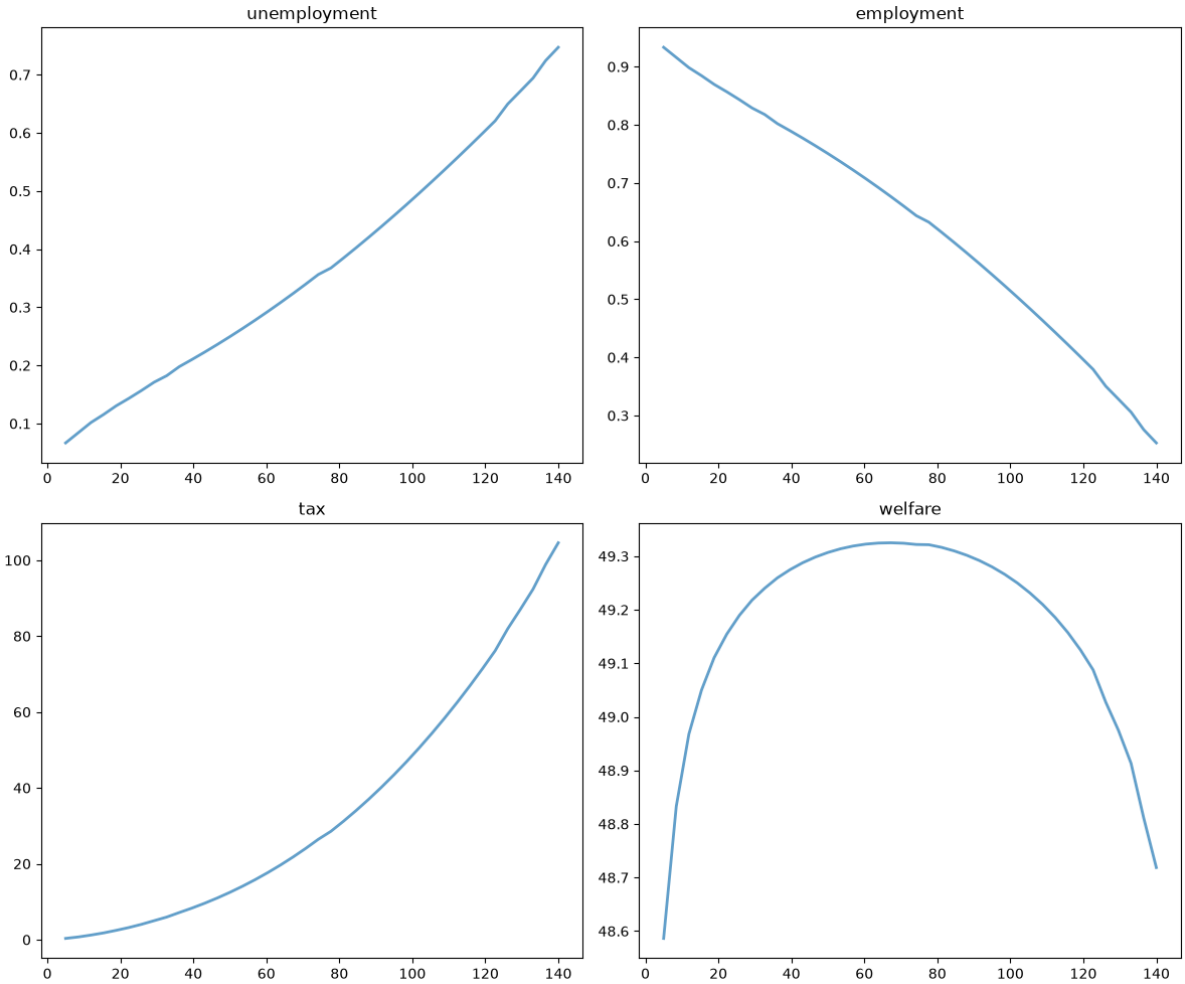
Let's visualize the results:

```
fig, axes = plt.subplots(2, 2, figsize=(12, 10))

plots = [unempl_vec, empl_vec, tax_vec, welfare_vec]
titles = ['unemployment', 'employment', 'tax', 'welfare']

for ax, plot, title in zip(axes.flatten(), plots, titles):
    ax.plot(c_vec, plot, lw=2, alpha=0.7)
    ax.set_title(title)

plt.tight_layout()
plt.show()
```



Welfare first increases and then decreases as unemployment benefits rise.

The level that maximizes steady state welfare is approximately 62.

115.4 Exercises

Exercise 115.4.1

How does the welfare-maximizing level of unemployment compensation c change with the job separation rate α?

Compute and plot the optimal c (the value that maximizes welfare) for a range of separation rates α from 0.01 to 0.04.

For each α value, find the optimal c by computing welfare across the range of c values and selecting the maximum.

Solution

Here is one solution:

```
# Range of separation rates to explore (wider range, fewer points)
α_values = jnp.linspace(0.01, 0.04, 8)

# We'll store the optimal c for each α
optimal_c_values = []

# Use a finer grid for c values to get better resolution
c_vec_fine = jnp.linspace(5, 140, 150)

for α_val in α_values:
    # Create economy parameters with this α
    params_α = create_economy(α=α_val)

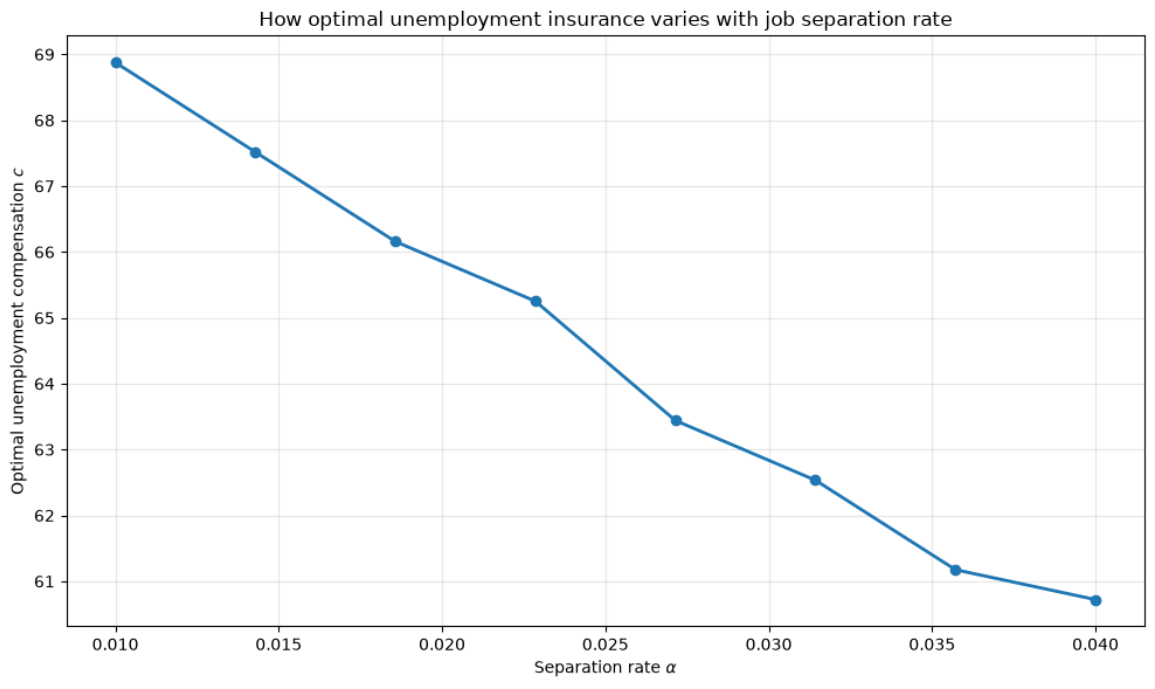
    # Create wage distribution
    w_vec_α, p_vec_α = create_wage_distribution(
        params_α.max_wage, params_α.wage_grid_size, params_α.log_wage_mean
    )

    # Compute welfare for each c value
    welfare_values = []
    for c in c_vec_fine:
        t = find_balanced_budget_tax(c, params_α, w_vec_α, p_vec_α)
        e_rate, u_rate, welfare = compute_steady_state_quantities(
            c, t, params_α, w_vec_α, p_vec_α
        )
        welfare_values.append(welfare)

    # The welfare function is very flat near its maximum.
    # Using argmax on a single point can be unstable due to numerical noise.
    # Instead, we find all c values within 99.9% of maximum welfare and
    # compute their weighted average (centroid). This gives a more stable
    # estimate of the optimal unemployment compensation level.
    welfare_array = jnp.array(welfare_values)
    max_welfare = jnp.max(welfare_array)
    threshold = 0.999 * max_welfare
    near_optimal_mask = welfare_array >= threshold

    # Compute weighted average of c values in the near-optimal region
    optimal_c = jnp.sum(c_vec_fine * near_optimal_mask * welfare_array) / \
                jnp.sum(near_optimal_mask * welfare_array)
    optimal_c_values.append(optimal_c)

# Plot the relationship
fig, ax = plt.subplots(figsize=(10, 6))
ax.plot(α_values, optimal_c_values, lw=2, marker='o')
ax.set_xlabel(r'Separation rate $\alpha$')
ax.set_ylabel('Optimal unemployment compensation $c$')
ax.set_title('How optimal unemployment insurance varies with job separation rate')
ax.grid(True, alpha=0.3)
plt.tight_layout()
plt.show()
```



We see that as the separation rate increases (workers lose their jobs more frequently), the welfare-maximizing level of unemployment compensation decreases.

This occurs because higher separation rates increase steady-state unemployment, which raises the tax burden needed to finance unemployment benefits. The optimal policy balances insurance against distortionary taxation.

CHAPTER

SIX

RATIONAL EXPECTATIONS EQUILIBRIUM

Contents

- *Rational Expectations Equilibrium*
  - *Overview*
  - *Rational expectations equilibrium*
  - *Computing an equilibrium*
  - *Concluding remarks*
  - *Exercises*

"If you're so smart, why aren't you rich?"

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

116.1 Overview

This lecture introduces the concept of a *rational expectations equilibrium*.

To illustrate it, we describe a linear quadratic version of a model due to Lucas and Prescott [Lucas and Prescott, 1971].

That 1971 paper is one of a small number of research articles that ignited a *rational expectations revolution*.

You can regard the present lecture as a "baby" version of that paper, and as an introduction to ideas that two sequels study in more general settings.

Investment Under Uncertainty describes the model that Lucas and Prescott actually built, in which demand is shifted by a Markov process, so that the equilibrium is a stochastic process rather than a deterministic path.

Optimal Growth Under Uncertainty and Tobin's q studies a one-sector optimal growth model with random shocks to production, and uses it to think about Tobin's q.

Both sequels ask a question that cannot arise here, because this lecture has no uncertainty: does the equilibrium have an invariant probability distribution to which it converges from any initial condition?

We follow Lucas and Prescott by employing a setting that is readily "Bellmanized" (i.e., susceptible to being formulated as a dynamic programming problems.

Because we use linear quadratic setups for demand and costs, we can deploy the LQ programming techniques described in *this lecture*.

We will learn about how a representative agent's problem differs from a planner's, and how a planning problem can be used to compute quantities and prices in a rational expectations equilibrium.

We will also learn about how a rational expectations equilibrium can be characterized as a fixed point of a mapping from a *perceived law of motion* to an *actual law of motion*.

Equality between a perceived and an actual law of motion for endogenous market-wide objects captures in a nutshell what the rational expectations equilibrium concept is all about.

Finally, we will learn about the important "Big K, little k" trick, a modeling device widely used in macroeconomics.

Except that for us

- Instead of "Big K" it will be "Big Y".
- Instead of "little k" it will be "little y".

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

We'll also use the LQ class from `QuantEcon.py`.

```
from quantecon import LQ
```

116.1.1 The big Y, little y trick

This widely used method applies in contexts in which a **representative firm** or agent is a "price taker" operating within a competitive equilibrium.

The following setting justifies the concept of a representative firm that stands in for a large number of other firms too.

There is a uniform unit measure of identical firms named $\omega \in \Omega = [0, 1]$.

The output of firm ω is $y(\omega)$.

The output of all firms is $Y = \int_0^1 y(\omega) d\,\omega$.

All firms end up choosing to produce the same output, so that at the end of the day $y(\omega) = y$ and $Y = y = \int_0^1 y(\omega) d\,\omega$.

This setting allows us to speak of a representative firm that chooses to produce y.

We want to impose that

- The representative firm or individual firm takes *aggregate* Y as given when it chooses individual $y(\omega)$, but ....
- At the end of the day, $Y = y(\omega) = y$, so that the representative firm is indeed representative.

The Big Y, little y trick accomplishes these two goals by

- Taking Y as beyond control when posing the choice problem of who chooses y; but ....
- Imposing $Y = y$ *after* having solved the individual's optimization problem.

Please watch for how this strategy is applied as the lecture unfolds.

We begin by applying the Big Y, little y trick in a very simple static context.

A simple static example of the big Y, little y trick

Consider a static model in which a unit measure of firms produce a homogeneous good that is sold in a competitive market.

Each of these firms ends up producing and selling output $y(\omega) = y$.

The price p of the good lies on an inverse demand curve

$$p = a_0 - a_1 Y \tag{116.1}$$

where

- $a_i > 0$ for $i = 0, 1$
- $Y = \int_0^1 y(\omega) d\omega$ is the market-wide level of output

For convenience, we'll often just write y instead of $y(\omega)$ when we are describing the choice problem of an individual firm $\omega \in \Omega$.

Each firm has a total cost function

$$c(y) = c_1 y + 0.5 c_2 y^2, \qquad c_i > 0 \text{ for } i = 1, 2$$

The profits of a representative firm are $py - c(y)$.

Using (116.1), we can express the problem of the representative firm as

$$\max_y \left[(a_0 - a_1 Y) y - c_1 y - 0.5 c_2 y^2 \right] \tag{116.2}$$

In posing problem (116.2), we want the firm to be a *price taker*.

We do that by regarding p and therefore Y as exogenous to the firm.

The essence of the Big Y, little y trick is *not* to set $Y = ny$ *before* taking the first-order condition with respect to y in problem (116.2).

This assures that the firm is a price taker.

The first-order condition for problem (116.2) is

$$a_0 - a_1 Y - c_1 - c_2 y = 0 \tag{116.3}$$

At this point, *but not before*, we substitute $Y = y$ into (116.3) to obtain the following linear equation

$$a_0 - c_1 - (a_1 + c_2) Y = 0 \tag{116.4}$$

to be solved for the competitive equilibrium market-wide output Y.

After solving for Y, we can compute the competitive equilibrium price p from the inverse demand curve (116.1).

116.1.2 Related planning problem

Define **consumer surplus** as the area under the inverse demand curve:

$$S_c(Y) = \int_0^Y (a_0 - a_1 s) ds = a_o Y - \frac{a_1}{2} Y^2.$$

Define the social cost of production as

$$S_p(Y) = c_1 Y + \frac{c_2}{2} Y^2$$

Consider the planning problem

$$\max_Y [S_c(Y) - S_p(Y)]$$

The first-order necessary condition for the planning problem is equation (116.4).

Thus, a Y that solves (116.4) is a competitive equilibrium output as well as an output that solves the planning problem.

This type of outcome provides an intellectual justification for liking a competitive equilibrium.

116.1.3 Further reading

References for this lecture include

- [Lucas and Prescott, 1971]
- [Sargent, 1987], chapter XIV
- [Ljungqvist and Sargent, 2018], chapter 7

The two sequels to this lecture are *Investment Under Uncertainty* and *Optimal Growth Under Uncertainty and Tobin's q*.

116.2 Rational expectations equilibrium

Our first illustration of a rational expectations equilibrium involves a market with a unit measure of identical firms, each of which seeks to maximize the discounted present value of profits in the face of adjustment costs.

The adjustment costs induce the firms to make gradual adjustments, which in turn requires consideration of future prices.

Individual firms understand that, via the inverse demand curve, the price is determined by the amounts supplied by other firms.

Hence each firm wants to forecast future total industry output.

In our context, a forecast is generated by a belief about the law of motion for the aggregate state.

Rational expectations equilibrium prevails when this belief coincides with the actual law of motion generated by production choices induced by this belief.

We formulate a rational expectations equilibrium in terms of a fixed point of an operator that maps beliefs into optimal beliefs.

116.2.1 Competitive equilibrium with adjustment costs

To illustrate, consider a collection of n firms producing a homogeneous good that is sold in a competitive market.

Each firm sell output $y_t(\omega) = y_t$.

The price p_t of the good lies on the inverse demand curve

$$p_t = a_0 - a_1 Y_t \tag{116.5}$$

where

- $a_i > 0$ for $i = 0, 1$
- $Y_t = \int_0^1 y_t(\omega) d\omega = y_t$ is the market-wide level of output

The firm's problem

Each firm is a price taker.

While it faces no uncertainty, it does face adjustment costs

In particular, it chooses a production plan to maximize

$$\sum_{t=0}^{\infty} \beta^t r_t \tag{116.6}$$

where

$$r_t := p_t y_t - \frac{\gamma (y_{t+1} - y_t)^2}{2}, \qquad y_0 \text{ given} \tag{116.7}$$

Regarding the parameters,

- $\beta \in (0, 1)$ is a discount factor
- $\gamma > 0$ measures the cost of adjusting the rate of output

Regarding timing, the firm observes p_t and y_t when it chooses y_{t+1} at time t.

To state the firm's optimization problem completely requires that we specify dynamics for all state variables.

This includes ones that the firm cares about but does not control like p_t.

We turn to this problem now.

Prices and aggregate output

In view of (116.5), the firm's incentive to forecast the market price translates into an incentive to forecast aggregate output Y_t.

Aggregate output depends on the choices of other firms.

The output $y_t(\omega)$ of a single firm ω has a negligible effect on aggregate output $\int_0^1 y_t(\omega) d\omega$.

That justifies firms in regarding their forecasts of aggregate output as being unaffected by their own output decisions.

Representative firm's beliefs

We suppose the firm believes that market-wide output Y_t follows the law of motion

$$Y_{t+1} = H(Y_t) \tag{116.8}$$

where Y_0 is a known initial condition.

The **belief function** H is an equilibrium object, and hence remains to be determined.

Because of this, at this stage Y_{t+1} only means the perceived output in the next period, Y_{t+1}^e.

Optimal behavior given beliefs

For now, let's fix a particular belief H in (116.8) and investigate the firm's response to it.

Let v be the optimal value function for the firm's problem given H.

The value function satisfies the Bellman equation

$$v(y, Y) = \max_{y'} \left\{ a_0 y - a_1 y Y - \frac{\gamma (y' - y)^2}{2} + \beta v(y', H(Y)) \right\} \tag{116.9}$$

Let's denote the firm's optimal policy function by h, so that

$$y_{t+1} = h(y_t, Y_t) \tag{116.10}$$

where

$$h(y, Y) := \text{argmax}_{y'} \left\{ a_0 y - a_1 y Y - \frac{\gamma (y' - y)^2}{2} + \beta v(y', H(Y)) \right\} \tag{116.11}$$

Evidently v and h both depend on H.

Characterization with first-order necessary conditions

In what follows it will be helpful to have a second characterization of h, based on first-order conditions.

The first-order necessary condition for choosing y' is

$$-\gamma (y' - y) + \beta v_y(y', H(Y)) = 0 \tag{116.12}$$

An important useful envelope result of Benveniste-Scheinkman [Benveniste and Scheinkman, 1979] implies that to differentiate v with respect to y we can naively differentiate the right side of (116.9), giving

$$v_y(y, Y) = a_0 - a_1 Y + \gamma (y' - y)$$

and equivalently, $v_y(y', H(Y)) = a_0 - a_1 H(Y) + \gamma (y'' - y')$

Substituting this equation into (116.12) gives the **Euler equation**

$$-\gamma (y_{t+1} - y_t) + \beta [a_0 - a_1 H(Y_t) + \gamma (y_{t+2} - y_{t+1})] = 0 \tag{116.13}$$

The firm optimally sets an output path that satisfies (116.13), taking (116.8) as given, and subject to

- the initial conditions for (y_0, Y_0).
- the terminal condition $\lim_{t \to \infty} \beta^t y_t v_y(y_t, Y_t) = 0$.

This last condition is called the **transversality condition**, and acts as a first-order necessary condition "at infinity".

A representative firm's decision rule solves the difference equation (116.13) subject to the given initial condition y_0 and the transversality condition.

Note that solving the Bellman equation (116.9) for v and then h in (116.11) yields a decision rule that automatically imposes both the Euler equation (116.13) and the transversality condition.

The actual law of motion for output

As we've seen, a given belief translates into a particular decision rule h.

Recalling that in equilbrium $Y_t = y_t$, the **actual law of motion** for market-wide output is then

$$Y_{t+1} = h(Y_t, Y_t) \tag{116.14}$$

Thus, when firms believe that the law of motion for market-wide output is (116.8), their optimizing behavior makes the actual law of motion be (116.14).

116.2.2 Definition of rational expectations equilibrium

Definition

A **rational expectations equilibrium** or **recursive competitive equilibrium** of the model with adjustment costs is a decision rule h and an aggregate law of motion H such that

1. Given belief H, the map h is the firm's optimal policy function.
2. The law of motion H satisfies $H(Y) = h(Y, Y)$ for all Y.

Thus, a rational expectations equilibrium equates the perceived and actual laws of motion (116.8) and (116.14).

Fixed point characterization

As we've seen, the firm's optimum problem induces a mapping Φ from a perceived law of motion H for market-wide output to an actual law of motion $\Phi(H)$.

The mapping Φ is the composition of two mappings, the first of which maps a perceived law of motion into a decision rule via (116.9)–(116.11), the second of which maps a decision rule into an actual law via (116.14).

The H component of a rational expectations equilibrium is a fixed point of Φ.

116.3 Computing an equilibrium

Now let's compute a rational expectations equilibrium.

116.3.1 Failure of contractivity

Readers accustomed to dynamic programming arguments might try to address this problem by choosing some guess H_0 for the aggregate law of motion and then iterating with Φ.

Unfortunately, the mapping Φ is not a contraction.

Indeed, there is no guarantee that direct iterations on Φ converge[1].

There are examples in which these iterations diverge.

[1] A literature that studies whether models populated with agents who learn can converge to rational expectations equilibria features iterations on a modification of the mapping Φ that can be approximated as $\gamma\Phi + (1-\gamma)I$. Here I is the identity operator and $\gamma \in (0, 1)$ is a *relaxation parameter*. See Marcet and Sargent [1989] and Evans and Honkapohja [2001] for statements and applications of this approach to establish conditions under which collections of adaptive agents who use least squares learning to converge to a rational expectations equilibrium.

To see this intuitively, consider Blackwell's sufficient condition and suppose there are two beliefs with $H_a(Y) > H_b(Y)$ for every Y.

By the Euler equation (116.13), the actual law of motion $Y_{t+1} = h(Y_t, Y_t)$ decreases as H increases, so Φ reverses the ordering of beliefs.

Hence the monotonicity required by Blackwell's condition fails.

Fortunately, another method works here.

The method exploits a connection between equilibrium and Pareto optimality expressed in the fundamental theorems of welfare economics (see, e.g, [Mas-Colell *et al.*, 1995]).

Lucas and Prescott [Lucas and Prescott, 1971] used this method to construct a rational expectations equilibrium.

Some details follow.

116.3.2 A planning problem approach

Our plan of attack is to match the Euler equations of the market problem with those for a single-agent choice problem.

As we'll see, this planning problem can be solved by LQ control (*linear regulator*).

Optimal quantities from the planning problem are rational expectations equilibrium quantities.

The rational expectations equilibrium price can be obtained as a shadow price in the planning problem.

We first compute a sum of consumer and producer surplus at time t

$$s(Y_t, Y_{t+1}) := \int_0^{Y_t} (a_0 - a_1 x)\, dx - \frac{\gamma (Y_{t+1} - Y_t)^2}{2} \tag{116.15}$$

The first term is the area under the demand curve, while the second measures the social costs of changing output.

The **planning problem** is to choose a production plan $\{Y_t\}$ to maximize

$$\sum_{t=0}^{\infty} \beta^t s(Y_t, Y_{t+1})$$

subject to an initial condition for Y_0.

116.3.3 Solution of planning problem

Evaluating the integral in (116.15) yields the quadratic form $a_0 Y_t - a_1 Y_t^2/2$.

As a result, the Bellman equation for the planning problem is

$$V(Y) = \max_{Y'} \left\{ a_0 Y - \frac{a_1}{2} Y^2 - \frac{\gamma (Y' - Y)^2}{2} + \beta V(Y') \right\} \tag{116.16}$$

The associated first-order condition is

$$-\gamma (Y' - Y) + \beta V'(Y') = 0 \tag{116.17}$$

Applying the same Benveniste-Scheinkman formula gives

$$V'(Y) = a_0 - a_1 Y + \gamma (Y' - Y)$$

Substituting this into equation (116.17) and rearranging leads to the Euler equation

$$\beta a_0 + \gamma Y_t - [\beta a_1 + \gamma (1 + \beta)] Y_{t+1} + \gamma \beta Y_{t+2} = 0 \tag{116.18}$$

116.3.4 Key insight

Return to equation (116.13) and set $y_t = Y_t$ for all t.

In a rational expectations equilibrium the perceived and actual laws of motion agree, so $H(Y_t) = Y_{t+1}$.

Imposing this alongside $y_t = Y_t$, a small amount of algebra will convince you that equations (116.18) and (116.13) are identical.

Thus, the Euler equation for the planning problem matches the second-order difference equation that we derived by

1. finding the Euler equation of the representative firm and
2. substituting into it the expression $Y_t = y_t$ that "makes the representative firm be representative".

If it is appropriate to apply the same terminal conditions for these two difference equations, which it is, then we have verified that a solution of the planning problem is also a rational expectations equilibrium quantity sequence.

It follows that for this example we can compute equilibrium quantities by forming the optimal linear regulator problem corresponding to the Bellman equation (116.16).

The optimal policy function for the planning problem is the aggregate law of motion H that the representative firm faces within a rational expectations equilibrium.

Structure of the law of motion

As you are asked to show in the exercises, the fact that the planner's problem is an LQ control problem implies an optimal policy — and hence aggregate law of motion — taking the form

$$Y_{t+1} = \kappa_0 + \kappa_1 Y_t \tag{116.19}$$

for some parameter pair κ_0, κ_1.

Now that we know the aggregate law of motion is linear, we can see from the firm's Bellman equation (116.9) that the firm's problem can also be framed as an LQ problem.

As you're asked to show in the exercises, the LQ formulation of the firm's problem implies a law of motion that looks as follows

$$y_{t+1} = h_0 + h_1 y_t + h_2 Y_t \tag{116.20}$$

Hence a rational expectations equilibrium will be defined by the parameters $(\kappa_0, \kappa_1, h_0, h_1, h_2)$ in (116.19)–(116.20).

116.4 Concluding remarks

Three ideas from this lecture recur throughout macroeconomics.

The first is the equilibrium concept itself: a *perceived* law of motion for a market-wide object must coincide with the *actual* law of motion that the resulting decisions generate.

The second is the "Big Y, little y" device that lets a price-taking firm be representative.

The third is the computational strategy: because the mapping from perceived to actual laws of motion is not a contraction, we found an equilibrium by solving a *planning problem* instead, and then read off equilibrium prices as shadow prices.

The two sequels take these ideas into settings with uncertainty.

Investment Under Uncertainty returns to Lucas and Prescott [1971] itself.

There, demand is shifted by a Markov process, firms face a nonlinear technology for converting investment into capacity, and the equilibrium is proved to exist, to be unique, and to solve a planning problem that maximizes discounted consumer surplus.

Because the equilibrium is a Markov process, one can ask whether it converges to an invariant distribution, and whether time averages computed from a single realization converge to population moments.

Affirmative answers are what make such models usable in econometrics.

Optimal Growth Under Uncertainty and Tobin's q pursues the same questions in the one-sector optimal growth model of Brock and Mirman [1972], and then follows Sargent [1980] in adding irreversible investment.

That small change makes the shadow price of installed capital – Tobin's q – diverge from the price of new capital, and it turns the planner's value function into the object on which a theory of investment rests.

The fixed-point idea also survives when beliefs are allowed to be wrong: *Self-Confirming Equilibria* keeps the requirement that perceived and actual laws of motion agree on the outcomes that are actually observed, while letting beliefs be wrong off the equilibrium path.

116.5 Exercises

Exercise 116.5.1

Consider the firm problem *described above*.

Let the firm's belief function H be as given in (116.19).

Formulate the firm's problem as a discounted optimal linear regulator problem, being careful to describe all of the objects needed.

Use the class `LQ` from the QuantEcon.py package to solve the firm's problem for the following parameter values:

$$a_0 = 100, a_1 = 0.05, \beta = 0.95, \gamma = 10, \kappa_0 = 95.5, \kappa_1 = 0.95$$

Express the solution of the firm's problem in the form (116.20) and give the values for each h_j.

If there were a unit measure of identical competitive firms all behaving according to (116.20), what would (116.20) imply for the *actual* law of motion (116.8) for market supply.

Solution

To map a problem into a *discounted optimal linear control problem*, we need to define

- state vector x_t and control vector u_t
- matrices A, B, Q, R that define preferences and the law of motion for the state

For the state and control vectors, we choose

$$x_t = \begin{bmatrix} y_t \\ Y_t \\ 1 \end{bmatrix}, \qquad u_t = y_{t+1} - y_t$$

For B, Q, R we set

$$A = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \kappa_1 & \kappa_0 \\ 0 & 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \quad R = \begin{bmatrix} 0 & a_1/2 & -a_0/2 \\ a_1/2 & 0 & 0 \\ -a_0/2 & 0 & 0 \end{bmatrix}, \quad Q = \gamma/2$$

By multiplying out you can confirm that

- $x_t' R x_t + u_t' Q u_t = -r_t$
- $x_{t+1} = Ax_t + Bu_t$

We'll use the module `lqcontrol.py` to solve the firm's problem at the stated parameter values.

This will return an LQ policy F with the interpretation $u_t = -Fx_t$, or

$$y_{t+1} - y_t = -F_0 y_t - F_1 Y_t - F_2$$

Matching parameters with $y_{t+1} = h_0 + h_1 y_t + h_2 Y_t$ leads to

$$h_0 = -F_2, \quad h_1 = 1 - F_0, \quad h_2 = -F_1$$

Here's our solution

```
# Model parameters

a0 = 100
a1 = 0.05
β = 0.95
γ = 10.0

# Beliefs

κ0 = 95.5
κ1 = 0.95

# Formulate the LQ problem

A = np.array([[1, 0, 0], [0, κ1, κ0], [0, 0, 1]])
B = np.array([1, 0, 0])
B.shape = 3, 1
R = np.array([[0, a1/2, -a0/2], [a1/2, 0, 0], [-a0/2, 0, 0]])
Q = 0.5 * γ

# Solve for the optimal policy

lq = LQ(Q, R, A, B, beta=β)
P, F, d = lq.stationary_values()
F = F.flatten()
out1 = f"F = [{F[0]:.3f}, {F[1]:.3f}, {F[2]:.3f}]"
h0, h1, h2 = -F[2], 1 - F[0], -F[1]
out2 = f"(h0, h1, h2) = ({h0:.3f}, {h1:.3f}, {h2:.3f})"

print(out1)
print(out2)
```

```
F = [-0.000, 0.046, -96.949]
(h0, h1, h2) = (96.949, 1.000, -0.046)
```

The implication is that

$$y_{t+1} = 96.949 + y_t - 0.046\, Y_t$$

For the case $n > 1$, recall that $Y_t = n y_t$, which, combined with the previous equation, yields

$$Y_{t+1} = n\,(96.949 + y_t - 0.046\, Y_t) = n96.949 + (1 - n0.046) Y_t$$

For the case of a unit measure of firms,

$$\int_0^1 y_{t+1}(\omega)\, d\omega = h_0 + h_1 \int_0^1 y_t(\omega)\, d\omega + h_2 Y_t$$
$$Y_{t+1} = h_0 + h_1 Y_t + h_2 Y_t$$
$$Y_{t+1} = 96.949 + (1 - 0.046) Y_t$$

Exercise 116.5.2

Consider the following κ_0, κ_1 pairs as candidates for the aggregate law of motion component of a rational expectations equilibrium (see (116.19)).

Extending the program that you wrote for Exercise 116.5.1, determine which if any satisfy *the definition* of a rational expectations equilibrium

- (94.0886298678, 0.923409232937)
- (93.2119845412, 0.984323478873)
- (95.0818452486, 0.952459076301)

Describe an iterative algorithm that uses the program that you wrote for Exercise 116.5.1 to compute a rational expectations equilibrium.

(You are not being asked actually to use the algorithm you are suggesting)

Solution

To determine whether a κ_0, κ_1 pair forms the aggregate law of motion component of a rational expectations equilibrium, we can proceed as follows:

- Determine the corresponding firm law of motion $y_{t+1} = h_0 + h_1 y_t + h_2 Y_t$.
- Test whether the associated aggregate law :$Y_{t+1} = nh(Y_t/n, Y_t)$ evaluates to $Y_{t+1} = \kappa_0 + \kappa_1 Y_t$.

In the second step, we can use $Y_t = n y_t = y_t$, so that $Y_{t+1} = nh(Y_t/n, Y_t)$ becomes

$$Y_{t+1} = h(Y_t, Y_t) = h_0 + (h_1 + h_2) Y_t$$

Hence to test the second step we can test $\kappa_0 = h_0$ and $\kappa_1 = h_1 + h_2$.

The following code implements this test

```
candidates = ((94.0886298678, 0.923409232937),
              (93.2119845412, 0.984323478873),
              (95.0818452486, 0.952459076301))

for κ0, κ1 in candidates:

    # Form the associated law of motion
    A = np.array([[1, 0, 0], [0, κ1, κ0], [0, 0, 1]])

    # Solve the LQ problem for the firm
    lq = LQ(Q, R, A, B, beta=β)
    P, F, d = lq.stationary_values()
    F = F.flatten()
    h0, h1, h2 = -F[2], 1 - F[0], -F[1]

    # Test the equilibrium condition
    if np.allclose((κ0, κ1), (h0, h1 + h2)):
        print(f'Equilibrium pair = {κ0}, {κ1}')
        print('f(h0, h1, h2) = {h0}, {h1}, {h2}')
        break
```

```
Equilibrium pair = 95.0818452486, 0.952459076301
f(h0, h1, h2) = {h0}, {h1}, {h2}
```

The output tells us that the answer is pair (iii), which implies $(h_0, h_1, h_2) = (95.0819, 1.0000, -.0475)$.

(Notice we use `np.allclose` to test equality of floating-point numbers, since exact equality is too strict).

Regarding the iterative algorithm, one could loop from a given (κ_0, κ_1) pair to the associated firm law and then to a new (κ_0, κ_1) pair.

This amounts to implementing the operator Φ described in the lecture.

(There is in general no guarantee that this iterative process will converge to a rational expectations equilibrium)

Exercise 116.5.3

Recall the planner's problem *described above*

1. Formulate the planner's problem as an LQ problem.
2. Solve it using the same parameter values in exercise 1
   - $a_0 = 100, a_1 = 0.05, \beta = 0.95, \gamma = 10$
3. Represent the solution in the form $Y_{t+1} = \kappa_0 + \kappa_1 Y_t$.
4. Compare your answer with the results from exercise 2.

Solution

We are asked to write the planner problem as an LQ problem.

For the state and control vectors, we choose

$$x_t = \begin{bmatrix} Y_t \\ 1 \end{bmatrix}, \quad u_t = Y_{t+1} - Y_t$$

For the LQ matrices, we set

$$A = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 1 \\ 0 \end{bmatrix}, \quad R = \begin{bmatrix} a_1/2 & -a_0/2 \\ -a_0/2 & 0 \end{bmatrix}, \quad Q = \gamma/2$$

By multiplying out you can confirm that

- $x_t' R x_t + u_t' Q u_t = -s(Y_t, Y_{t+1})$
- $x_{t+1} = A x_t + B u_t$

By obtaining the optimal policy and using $u_t = -F x_t$ or

$$Y_{t+1} - Y_t = -F_0 Y_t - F_1$$

we can obtain the implied aggregate law of motion via $\kappa_0 = -F_1$ and $\kappa_1 = 1 - F_0$.

The Python code to solve this problem is below:

```
# Formulate the planner's LQ problem

A = np.array([[1, 0], [0, 1]])
B = np.array([[1], [0]])
R = np.array([[a1 / 2, -a0 / 2], [-a0 / 2, 0]])
Q = γ / 2

# Solve for the optimal policy

lq = LQ(Q, R, A, B, beta=β)
P, F, d = lq.stationary_values()

# Print the results

F = F.flatten()
κ0, κ1 = -F[1], 1 - F[0]
print(κ0, κ1)
```

```
95.08187459214764 0.952459062703923
```

The output yields the same (κ_0, κ_1) pair obtained as an equilibrium from the previous exercise.

Exercise 116.5.4

A monopolist faces the industry demand curve (116.5) and chooses $\{Y_t\}$ to maximize $\sum_{t=0}^{\infty} \beta^t r_t$ where

$$r_t = p_t Y_t - \frac{\gamma (Y_{t+1} - Y_t)^2}{2}$$

Formulate this problem as an LQ problem.

Compute the optimal policy using the same parameters as Exercise 116.5.2.

In particular, solve for the parameters in

$$Y_{t+1} = m_0 + m_1 Y_t$$

Compare your results with Exercise 116.5.2 – comment.

Solution

The monopolist's LQ problem is almost identical to the planner's problem from the previous exercise, except that

$$R = \begin{bmatrix} a_1 & -a_0/2 \\ -a_0/2 & 0 \end{bmatrix}$$

The problem can be solved as follows

```
A = np.array([[1, 0], [0, 1]])
B = np.array([[1], [0]])
R = np.array([[a1, -a0 / 2], [-a0 / 2, 0]])
Q = γ / 2

lq = LQ(Q, R, A, B, beta=β)
P, F, d = lq.stationary_values()

F = F.flatten()
m0, m1 = -F[1], 1 - F[0]
print(m0, m1)
```

```
73.4729440350284 0.9265270559649701
```

We see that the law of motion for the monopolist is approximately $Y_{t+1} = 73.4729 + 0.9265 Y_t$.

In the rational expectations case, the law of motion was approximately $Y_{t+1} = 95.0818 + 0.9525 Y_t$.

One way to compare these two laws of motion is by their fixed points, which give long-run equilibrium output in each case.

For laws of the form $Y_{t+1} = c_0 + c_1 Y_t$, the fixed point is $c_0/(1 - c_1)$.

If you crunch the numbers, you will see that the monopolist adopts a lower long-run quantity than obtained by the competitive market, implying a higher market price.

This is analogous to the elementary static-case results

CHAPTER

SEVEN

INVESTMENT UNDER UNCERTAINTY

Contents

- *Investment Under Uncertainty*
  - *Overview*
  - *The industry*
  - *The firm and the price of installed capital*
  - *Rational expectations equilibrium*
  - *Equilibrium as a planning problem*
  - *Recursive competitive equilibrium*
  - *A computable version*
  - *Long run behavior with serially independent demand*
  - *Long run behavior with serially correlated demand*
  - *Relation to the rational expectations lecture*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

117.1 Overview

This lecture studies Lucas and Prescott [1971], a paper that helped to ignite a *rational expectations revolution*.

Lucas and Prescott studied a competitive industry in which

- demand shifts randomly each period
- firms face costs of adjusting their capital stocks
- firms must forecast future prices in order to decide how much to invest
- the probability distribution that firms use to forecast prices *equals* the probability distribution that their investment decisions actually generate

That last bullet point is what Muth [1961] called **rational expectations**.

The QuantEcon lecture *Rational Expectations Equilibrium* presents what we might call a "baby" version of the Lucas-Prescott model.

That lecture studies a linear-quadratic industry *without uncertainty*.

The present lecture describes the more ambitious structure that Lucas and Prescott actually built.

Relative to the baby version, Lucas and Prescott

- let demand be shifted by a Markov process $\{u_t\}$, so that the equilibrium is a stochastic process rather than a deterministic path
- allow a nonlinear technology for converting investment into capacity
- prove that a competitive equilibrium *exists* and is *unique*
- show that the equilibrium is the solution of a *planning problem* that maximizes discounted consumer surplus
- show that the equilibrium is a *Markov process* in the state (k_t, u_t)
- provide conditions under which that Markov process has an *invariant probability distribution* to which it converges from any initial condition

The last item is the part of the paper that most influenced later work.

A model whose equilibrium is a Markov process with an invariant distribution is a model that can be taken to time series data.

That observation set the stage for the *rational expectations econometrics* subsequently developed by Lars Peter Hansen and Thomas Sargent [Hansen and Sargent, 1980].

Along the way we'll describe how Prescott and Mehra [1980] later distilled the Lucas-Prescott structure into a general definition of a **recursive competitive equilibrium**.

The sequel *Optimal Growth Under Uncertainty and Tobin's q* studies a paper that asks the same questions about a one-sector optimal growth model, Brock and Mirman [1972], and a paper that uses that model to think about Tobin's q, Sargent [1980].

Let's start with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
import quantecon as qe
from collections import namedtuple
```

117.2 The industry

An industry consists of many small firms.

Each firm produces a single output q_t with a single input, capital k_t.

Production has constant returns to scale, and with a suitable choice of units the production function is

$$0 \leq q_t \leq k_t. \tag{117.1}$$

Because capital is the only input and output is sold at a positive price, every firm produces at capacity, so that $q_t = k_t$.

Let x_t denote gross investment.

Capacity next period is related to capacity this period and investment by

$$k_{t+1} = k_t \, h\left(\frac{x_t}{k_t}\right), \tag{117.2}$$

where h is bounded, continuously differentiable, increasing, and strictly *concave*.

The strict concavity of h is what creates **costs of adjustment**: doubling the rate of investment per unit of capital less than doubles the resulting increment to capacity.

Adjustment costs are why firms change their capital stocks gradually instead of jumping immediately to a long-run target.

Assume that $\delta = h^{-1}(1)$ exists and satisfies $0 < \delta < 1$.

Then $x_t = \delta k_t$ is the investment rate that just maintains capacity, so δ plays the role of a depreciation rate.

Let p_t be the output price and let $\beta = 1/(1+r)$, where $r > 0$ is the cost of capital.

The present value of a firm is

$$V = \sum_{t=0}^{\infty} \beta^t \left[p_t q_t - x_t\right]. \tag{117.3}$$

Because the allocation of a given industry capital stock across firms does not matter, we can use k_t, x_t, q_t interchangeably for firm and industry variables.

Equivalently, we can think of a competitive industry with a single price-taking firm.

Industry demand is subject to random shifts:

$$p_t = D(q_t, u_t), \tag{117.4}$$

where D is continuous and strictly decreasing in q_t and increasing in u_t, so that an increase in u_t shifts the demand curve to the right.

The demand shifter $\{u_t\}$ is a Markov process with transition distribution $p(\cdot, u)$, meaning that the probability that $u_{t+1} \in A$ conditional on $u_t = u$ is $\int_A p(dz, u)$.

117.3 The firm and the price of installed capital

Before studying equilibrium, Lucas and Prescott pause over an instructive question: what does an individual firm actually need to know?

Let w_t be the current market value of a unit of installed capital and let w_t^* be the value per unit expected to prevail next period.

A firm that begins period t with capital k_t and invests x obtains next-period capital worth $\beta k_t h(x/k_t) w_t^*$ at a cost of x, so it solves

$$\max_x \left[-x + \beta k_t h(x/k_t) w_t^*\right].$$

The current value of the firm is

$$w_t k_t = p_t k_t - x + \beta k_t h(x/k_t) w_t^* \tag{117.5}$$

and the first-order condition is

$$0 \geq -1 + \beta h'(x/k_t) w_t^*, \quad \text{with equality if } x > 0. \tag{117.6}$$

Solving (117.5) and (117.6) jointly for x and w_t^* gives an investment function of the form

$$x_t = k_t \, g(w_t - p_t), \qquad g'(\cdot) > 0. \tag{117.7}$$

Lucas and Prescott note that (117.7) is essentially the investment function that Grunfeld had used in empirical work, with the market value of the firm as an explanatory variable.

Their argument is stronger than Grunfeld's, though: a firm does not need to forecast its own future income stream at all.

It needs only to know the value that securities markets place on a unit of installed capital.

Readers will recognize a version of what later became known as Tobin's q theory of investment.

But equation (117.7) is a *consistency requirement*, not yet a theory of capital accumulation, because the path of w_t is still unknown.

To determine w_t we have to study equilibrium.

117.4 Rational expectations equilibrium

Firms must forecast future prices.

Lucas and Prescott describe the usual approach as postulating a forecasting rule – for example "adaptive expectations" – that generates investment behavior, which in conjunction with demand generates an actual price process.

They object that if the underlying disturbance has a regular stochastic character, then, except by coincidence, forecast prices and actual prices will have *different probability distributions*, and the difference will be persistent, costly, and easy to correct.

So they go to the opposite extreme and assume that the actual and anticipated prices have the *same probability distribution*.

To say this precisely, fix an initial state (k_0, u_0).

Because prices depend on the history of demand shocks, an anticipated price process is a sequence $\{p_t\}$ of functions of $(u_1, \ldots, u_t)$.

Similarly, an investment-output plan is a pair of sequences $\{q_t, x_t\}$ of functions of $(u_1, \ldots, u_t)$ – a contingency plan that says in advance what the firm will do after every possible history.

Definition 117.4.1

An **industry equilibrium** for a fixed initial state (k_0, u_0) is a triple of sequences $\{q_t^0, x_t^0, p_t^0\}$ such that

1. the demand curve (117.4) holds for every history, and
2. the plan $\{q_t^0, x_t^0\}$ maximizes expected present value

$$\mathbb{E}\left\{\sum_{t=0}^{\infty} \beta^t \left[p_t^0 q_t - x_t\right]\right\}$$

over all plans $\{q_t, x_t\}$ that satisfy (117.1) and (117.2), *given* the price process $\{p_t^0\}$.

The rational expectations requirement is hiding in plain sight in *Definition 117.4.1*.

The price process $\{p_t^0\}$ that firms take as given when they maximize is the *same* price process that their own decisions generate through the demand curve.

This is exactly the fixed point idea of the lecture *Rational Expectations Equilibrium*, where a **perceived law of motion** H for aggregate output must equal the **actual law of motion** that the resulting decision rule generates.

The difference is that here the fixed point is in a space of sequences of functions of histories rather than in a space of linear decision rules.

Note

Lucas and Prescott are careful about what rationality does and does not assume.

They write that they "surrender, in advance, any hope of shedding light on the process by which firms translate current information into price forecasts."

They also defend the assumption: if the demand shift process really does have a regular, stationary structure, then expectations that are rational in their sense "are surely more plausible than any simple, adaptive scheme"; and if it does not, then adopting some other expectations hypothesis "will certainly not improve matters."

117.5 Equilibrium as a planning problem

How can we compute an object defined by a fixed point in such a large space?

Lucas and Prescott's answer is the device that the lecture *Rational Expectations Equilibrium* also uses: find a **planning problem** whose solution is the equilibrium.

Define **consumer surplus** as the area under the demand curve

$$s(q, u) = \int_0^q D(z, u)\, dz,$$

and define discounted consumer surplus net of investment costs by

$$S = \mathbb{E}\left\{\sum_{t=0}^{\infty} \beta^t \left[s(q_t, u_t) - x_t\right]\right\}. \tag{117.8}$$

Associated with the problem of maximizing S is the functional equation

$$v(k, u) = \sup_{x \geq 0}\left\{s(k, u) - x + \beta \int v\left[kh\left(\frac{x}{k}\right), z\right] p(dz, u)\right\}. \tag{117.9}$$

Theorem 117.5.1

The functional equation (117.9) has a unique bounded solution v, and for each (k, u) the supremum is attained at a unique $x(k, u)$.

In terms of that policy function, the unique industry equilibrium, given (k_0, u_0), is

$$x_t = x(k_t, u_t), \qquad k_{t+1} = k_t h\left(\frac{x(k_t, u_t)}{k_t}\right), \qquad q_t = k_t, \qquad p_t = D(q_t, u_t).$$

The proof has two halves, and both halves matter for later work.

The first half shows that a competitive equilibrium maximizes S, and conversely.

This is an application of the welfare theorems in an infinite-dimensional commodity space, using the valuation equilibria of Debreu and the price systems that Prescott and Lucas developed in a companion paper.

Lucas and Prescott are explicit that they use this connection only as a computational device: "the welfare significance of S is not important. We are interested only in using the connection between the maximization of S and competitive equilibrium in order to determine the properties of the latter."

The second half shows that the planning problem is solved by the functional equation (117.9).

Here they use the operator

$$Tf(k,u) = \sup_{x \geq 0} \left\{ s(k,u) - x + \beta \int f[kh(x/k), z]\, p(dz, u) \right\}$$

and verify that T is monotone and satisfies a discounting property, so that by Blackwell's theorem [Blackwell, 1965] it has a unique fixed point that successive approximations converge to.

They also show that T preserves concavity and monotonicity in k, which delivers a unique and continuous policy function $x(k,u)$.

Note

These arguments are now standard and are treated at length in Stokey *et al.* [1989].

In 1971 they were not standard, which is one reason the paper is hard to read.

Much of its length is devoted to measurability details – Baire functions, Borel sets – that a modern treatment would relegate to an appendix.

Two features of *Theorem 117.5.1* deserve emphasis.

First, the equilibrium is **recursive**: the pair (k_t, u_t) is a Markov process, and equilibrium prices and quantities are time-invariant functions of it.

Second, the equilibrium is computed *without ever iterating on a mapping from beliefs to outcomes*.

The lecture *Rational Expectations Equilibrium* explains why that matters: the mapping Φ from a perceived law of motion to an actual law of motion is *not a contraction*, and iterating on it can diverge.

The planning problem replaces an unreliable fixed point calculation with a dynamic program that is a contraction.

117.6 Recursive competitive equilibrium

Prescott and Mehra [1980] later extracted the general structure that Lucas and Prescott had exploited.

Their goal was to replace a search for equilibrium *sequences of contingency functions*, in the style of Arrow and Debreu, with a search for equilibrium *decision rules*.

Such rules specify current actions as functions of a small number of **state variables** that summarize the effects of past decisions and current information.

As Prescott and Mehra put it, these equilibrium decision rules "must be time invariant in order to apply standard time series methods and this necessitates a recursive structure."

That sentence is the bridge from Lucas and Prescott's theory to econometrics.

In a recursive competitive equilibrium

- the state variables should be of minimal dimension, indexing only the factors that can change over time
- the state is observed, or is an invertible function of observables
- the conditional distribution of next period's state given current decisions and the current state is time invariant

- individual decision rules are optimal given equilibrium pricing functions, and markets clear

Prescott and Mehra note that their structure "subsumes the structure considered in Lucas and Prescott's analysis of equilibrium investment under uncertainty."

Their analysis also establishes optimality of recursive equilibria and supportability of Pareto optima "in a simpler and more direct way" than arguments that pass through equivalence with state-contingent equilibria.

For us, the important point is that *Theorem 117.5.1* produces exactly the objects that rational expectations econometrics needs: time-invariant decision rules, driven by a Markov state, with cross-equation restrictions linking the parameters of the shock process to the parameters of the decision rules.

117.7 A computable version

Let's now compute equilibria of a version of the model.

We take the adjustment technology

$$h(z) = (1 - \delta + z)^\alpha, \qquad 0 < \alpha \leq 1,$$

which satisfies the Lucas-Prescott assumptions: h is increasing and strictly concave for $\alpha < 1$, and $h(\delta) = 1$, so δ is the maintenance investment rate.

When $\alpha = 1$ we recover the familiar linear accumulation equation $k_{t+1} = (1 - \delta)k_t + x_t$.

When $\alpha < 1$ there are adjustment costs.

Note that $h'(\delta) = \alpha$, a fact we will use below.

We take a linear inverse demand curve

$$D(q, u) = a_0 + u - a_1 q,$$

which matches the demand curve in the lecture *Rational Expectations Equilibrium* except that it is shifted by u.

Consumer surplus is then

$$s(k, u) = (a_0 + u)k - \frac{a_1}{2}k^2.$$

The demand shifter follows a Gaussian AR(1) process

$$u_{t+1} = \rho u_t + \sigma \epsilon_{t+1}, \qquad \epsilon_{t+1} \sim N(0, 1),$$

which is an example that Lucas and Prescott themselves offer of a process satisfying their assumptions.

We discretize it with the Tauchen method.

Rather than choosing investment x directly, it is convenient to let the planner choose next period's capital k' on a grid, and to invert (117.2) to find the required investment

$$x = k\left[\left(\frac{k'}{k}\right)^{1/\alpha} - (1 - \delta)\right].$$

```
Model = namedtuple("Model", "r β δ α a0 a1 k u P X feasible s")

def create_model(r=0.05, δ=0.10, α=0.70, a0=1.0, a1=0.01,
                 ρ=0.9, σ=0.02, n_u=9, n_k=400, k_lo=20.0, k_hi=160.0):
    "Discretize the Lucas-Prescott industry."
```

(continues on next page)

(continued from previous page)

```
    β = 1 / (1 + r)
    mc = qe.markov.tauchen(n_u, ρ, σ)
    u, P = mc.state_values, mc.P
    k = np.linspace(k_lo, k_hi, n_k)
    # investment needed to move from k (rows) to k' (columns)
    X = k[:, None] * ((k[None, :] / k[:, None])**(1/α) - (1 - δ))
    feasible = X >= 0
    s = (a0 + u[None, :]) * k[:, None] - a1 * k[:, None]**2 / 2
    return Model(r, β, δ, α, a0, a1, k, u, P, X, feasible, s)
```

We solve the planner's Bellman equation (117.9) by value function iteration, with Howard policy improvement steps to speed convergence.

```
def solve_model(m, tol=1e-8, maxit=1000, howard=30):
    "Solve the planning problem; return value function and policies."
    n_k, n_u = len(m.k), len(m.u)
    R = np.where(m.feasible, -m.X, -1e12)
    rows, cols = np.arange(n_k)[:, None], np.arange(n_u)[None, :]
    v = m.s.copy()

    for it in range(maxit):
        EV = v @ m.P.T                                  # EV[k', u] = E[v(k', u') | u]
        obj = R[:, :, None] + m.β * EV[None, :, :]      # (k, k', u)
        idx = obj.argmax(axis=1)                        # choice of k' given (k, u)
        v_new = m.s + np.take_along_axis(obj, idx[:, None, :], axis=1)[:, 0, :]

        for _ in range(howard):                         # policy evaluation steps
            EV = v_new @ m.P.T
            v_new = m.s + R[rows, idx] + m.β * EV[idx, cols]

        if np.max(np.abs(v_new - v)) < tol:
            v = v_new
            break
        v = v_new

    k_next = m.k[idx]
    x = np.take_along_axis(m.X, idx, axis=1)
    return v, idx, k_next, x

m = create_model()
v, idx, k_next, x = solve_model(m)
print(f"grid: {len(m.k)} capital points, {len(m.u)} demand states")
```

```
grid: 400 capital points, 9 demand states
```

Let's look at the equilibrium investment policy and the law of motion for capital.

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))

for j in [0, len(m.u)//2, len(m.u)-1]:
    axes[0].plot(m.k, x[:, j], label=f'$u = {m.u[j]:.3f}$')
    axes[1].plot(m.k, k_next[:, j], label=f'$u = {m.u[j]:.3f}$')

axes[0].plot(m.k, m.δ * m.k, 'k--', lw=1, label=r'$\delta k$')
axes[0].set_xlabel('$k$'); axes[0].set_ylabel('$x(k, u)$')
axes[0].set_title('investment policy')
```

(continues on next page)

(continued from previous page)

```
axes[1].plot(m.k, m.k, 'k--', lw=1, label='45 degree line')
axes[1].set_xlabel('$k$'); axes[1].set_ylabel("$k'(k, u)$")
axes[1].set_title('law of motion for capital')
for ax in axes:
    ax.legend()
plt.tight_layout()
plt.show()
```

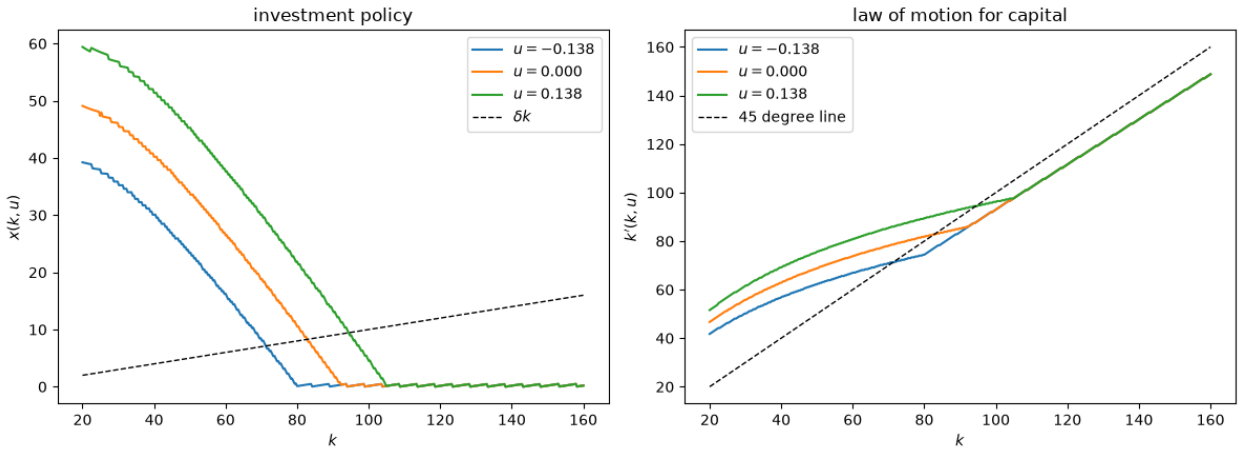


Fig. 117.1: Investment policy and law of motion

Capital rises when $x(k, u)$ lies above the maintenance line δk and falls when it lies below.

Higher demand shifts the investment policy up, so the capital stock that the industry sustains is higher when demand is strong.

117.7.1 The price of installed capital

The planner's problem also delivers the market value w of a unit of installed capital that appeared in (117.5).

Let $z = x/k$ denote the investment rate.

Differentiating the Bellman equation (117.9) and using the envelope condition gives

$$w(k, u) = v_k(k, u) = D(k, u) + \beta \mathbb{E}\left[v_k(k', u') \mid u\right]\left[h(z) - z h'(z)\right],$$

while the first-order condition for x is

$$\beta \mathbb{E}\left[v_k(k', u') \mid u\right] = \frac{1}{h'(z)}. \tag{117.10}$$

Combining them expresses the shadow price in closed form,

$$w(k, u) = D(k, u) + \frac{h(z)}{h'(z)} - z. \tag{117.11}$$

The marginal value of installed capital equals the current price of output plus the value of the capacity that the unit carries into the future.

Notice that (117.10) is the planner's counterpart of the firm's first-order condition (117.6), with $w^* = \mathbb{E}[v_k(k', u') \mid u]$.

That correspondence is the "big K, little k" logic of the lecture *Rational Expectations Equilibrium* in its Lucas-Prescott form.

```
h = lambda z, m: (1 - m.δ + z)**m.α
h_prime = lambda z, m: m.α * (1 - m.δ + z)**(m.α - 1)

z = x / m.k[:, None]
D = m.a0 + m.u[None, :] - m.a1 * m.k[:, None]
w = D + h(z, m) / h_prime(z, m) - z

# check the first-order condition (up to grid error)
Ew = np.take_along_axis(w @ m.P.T, idx, axis=0)
resid = np.abs(m.β * Ew - 1 / h_prime(z, m))
scale = np.median(1 / h_prime(z, m))
near = (m.k > 70) & (m.k < 95)          # capital levels the industry actually visits
print(f"typical size of each side of the FOC: {scale:.3f}")
print(f"median residual, all k:                {np.median(resid):.2e}")
print(f"median residual, 70 < k < 95:          {np.median(resid[near]):.2e}")
```

```
typical size of each side of the FOC: 1.396
median residual, all k:                5.35e-03
median residual, 70 < k < 95:          2.39e-03
```

The residual is a few tenths of one per cent of the magnitude of the terms being compared, which confirms (117.11).

It does not vanish entirely because the planner chooses k' from a finite grid, so the policy function jumps by a whole grid point at a time.

Residuals are much larger at capital levels so extreme that the industry never visits them.

117.8 Long run behavior with serially independent demand

Lucas and Prescott next ask what happens in the long run.

They treat two cases, and the first is the special case in which u_t and u_s are independent for $s \neq t$.

Inspecting the Bellman equation (117.9) when $p(dz, u)$ does not depend on u shows that the optimal investment rate $x(k, u)$ *does not depend on* u.

A demand shift is then a pure windfall: it tells firms nothing about future demand, so it does not change investment.

Consequently the capital stock evolves *deterministically*, according to $k_{t+1} = k_t h(x(k_t)/k_t)$, while output is supplied inelastically and demand shocks move only prices.

Theorem 117.8.1

Under independence, there are two possibilities for the capital stock.

If

$$\int D(0, u)p(du) > \delta + \frac{r}{h'(\delta)}, \tag{117.12}$$

and $k_0 > 0$, then k_t converges monotonically to the unique stationary value k^c given implicitly by

$$\int D(k^c, u)p(du) = \delta + \frac{r}{h'(\delta)}. \tag{117.13}$$

Otherwise capital converges monotonically to zero.

Condition (117.13) has a familiar interpretation.

The left side is expected marginal revenue product of capital, which here is just the expected output price, because the marginal physical product is one.

The right side is a **user cost of capital**: a depreciation term δ plus an interest term $r/h'(\delta)$.

Lucas and Prescott observe that this case corresponds closely to the textbook dichotomy between short-run and long-run supply.

In the short run, capacity is fixed and demand determines price.

In the long run, demand fluctuations play no role at all: capacity is determined entirely by *average* demand.

Let's verify this numerically by setting $\rho = 0$.

```
m_iid = create_model(ρ=0.0)
v_iid, idx_iid, k_next_iid, x_iid = solve_model(m_iid)

# does the policy depend on u?
print("investment policy independent of u:",
      np.allclose(x_iid, x_iid[:, [0]], atol=1e-10))

# stationary capital: where x(k) crosses δ k
def stationary_k(x_col, m):
    "Capital where investment just maintains capacity."
    d = x_col - m.δ * m.k
    i = np.where(np.sign(d[:-1]) != np.sign(d[1:]))[0]
    if len(i) == 0:
        return np.nan
    i = i[0]
    return np.interp(0, [d[i+1], d[i]], [m.k[i+1], m.k[i]])

kc = stationary_k(x_iid[:, 0], m_iid)
user_cost = m_iid.δ + m_iid.r / h_prime(m_iid.δ, m_iid)
print(f"\nstationary capital k^c        = {kc:.3f}")
print(f"expected price at k^c         = {m_iid.a0 - m_iid.a1 * kc:.5f}")
print(f"user cost δ + r / h'(δ)       = {user_cost:.5f}")
```

```
investment policy independent of u: True

stationary capital k^c        = 82.807
expected price at k^c         = 0.17193
user cost δ + r / h'(δ)       = 0.17143
```

The stationary capital stock equates the expected price to the user cost of capital, as (117.13) requires.

Now let's confirm that capital approaches k^c monotonically, and from either direction.

```
def capital_path(m, idx, k0, T=60, u_index=None):
    "Simulate capital, holding the demand state fixed if u_index is given."
    ki = np.abs(m.k - k0).argmin()
    path = np.empty(T)
    j = len(m.u) // 2 if u_index is None else u_index
    for t in range(T):
        path[t] = m.k[ki]
        ki = idx[ki, j]
    return path

fig, ax = plt.subplots(figsize=(8, 4.5))
```

(continues on next page)

(continued from previous page)

```
for k0 in (30.0, 55.0, 110.0, 150.0):
    ax.plot(capital_path(m_iid, idx_iid, k0), lw=2, label=f'$k_0 = {k0:.0f}$')
ax.axhline(kc, color='k', ls='--', lw=1, label='$k^c$')
ax.set_xlabel('$t$'); ax.set_ylabel('$k_t$')
ax.legend()
plt.tight_layout()
plt.show()
```

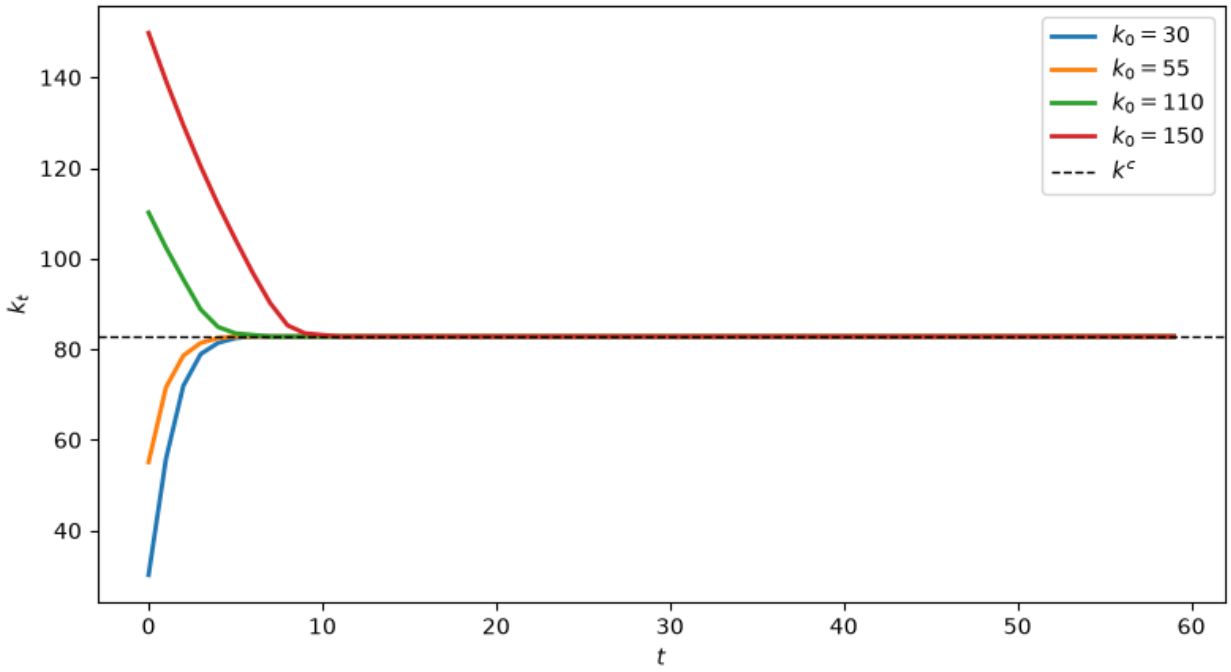


Fig. 117.2: Capital paths under IID demand

Convergence is monotone and the limit does not depend on the initial capital stock.

117.9 Long run behavior with serially correlated demand

The more interesting case allows demand shifts to be positively serially correlated, so that a high demand today signals high demand tomorrow.

Now the current demand state *does* affect investment, and the capital stock is genuinely stochastic.

To characterize the long run, Lucas and Prescott impose additional restrictions on the $\{u_t\}$ process, whose purpose is to guarantee that the distribution of (k_t, u_t) settles down.

In words, they assume that

- from any current u, next period's shock lands in any non-degenerate interval with positive probability
- u_t has a limiting distribution that does not depend on the initial u_0 and that puts positive probability on every non-degenerate interval
- $\mathbb{P}\{u_{t+1} \geq x \mid u_t\}$ is strictly increasing in u_t, so high demand today always signals high demand tomorrow

- consumer surplus $s(k, u)$ converges uniformly as $u \to \pm\infty$

A Gaussian AR(1) process with $0 < \rho < 1$ satisfies these conditions, which is the example they give and the one we simulate.

The analysis then proceeds by bounding the capital stock.

Let $\bar{v}(k)$ and $\underline{v}(k)$ be the limits of the expected value function as $u \to \infty$ and $u \to -\infty$, and let $\bar{x}(k)$ and $\underline{x}(k)$ be the associated investment policies.

Because investment is increasing in u, these bracket the policy for every u.

Let $\bar{k}$ solve $\bar{x}(k) = \delta k$ and let $\underline{k}$ solve $\underline{x}(k) = \delta k$.

These are the capital stocks that would be sustained under permanently maximal and permanently minimal demand.

Lucas and Prescott then prove that

- the sets $(0, \underline{k})$ and $(\bar{k}, \infty)$, paired with any demand state, are **transient**: once the industry leaves them it does not return, and from any starting point it enters $(\underline{k}, \bar{k})$ with probability approaching one
- the set $B = (\underline{k}, \bar{k}) \times E$, where E is the set of possible values of u, is a **single ergodic set**

Theorem 117.9.1

If B is non-empty, then for all (k, u) and every initial state (k_0, u_0),

$$\lim_{t \to \infty} \mathbb{P}\{k_t \leq k, u_t \leq u \mid k_0, u_0\} = P(k, u)$$

exists and *does not depend on* (k_0, u_0).

The function P is a probability distribution that assigns probability zero to the transient sets and positive probability to every subset of B with positive area.

Theorem 117.9.2

If B is non-empty, then for any initial state, with probability one

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} k_t = k^*,$$

where k^* is the mean of k under the invariant distribution P.

The ergodic set is non-empty – and so this long-run behavior applies – if and only if

$$\lim_{u \to \infty} D(0, u) > \delta + \frac{r}{h'(\delta)}, \tag{117.14}$$

which says that demand must sometimes be strong enough to justify holding any capital at all.

Theorem 117.9.1 and *Theorem 117.9.2* are the results that make this model an object that econometricians can use.

The first says that the model implies a well-defined stationary probability distribution for the observable time series.

The second says that sample averages computed from a single long realization converge to the corresponding population moments of that distribution.

Let's compute the bounds $\underline{k}$ and $\bar{k}$ for our parameterization and check that simulations behave as the theorems say.

```
bounds = np.array([stationary_k(x[:, j], m) for j in range(len(m.u))])
k_lo_star, k_hi_star = bounds.min(), bounds.max()
print(f"conditional stationary capital, lowest demand state:  {k_lo_star:.2f}")
print(f"conditional stationary capital, highest demand state: {k_hi_star:.2f}")
print(f"ergodic set for capital: ({k_lo_star:.2f}, {k_hi_star:.2f})")
```

```
conditional stationary capital, lowest demand state:  71.23
conditional stationary capital, highest demand state: 94.39
ergodic set for capital: (71.23, 94.39)
```

Now we simulate the equilibrium Markov process from two very different initial capital stocks, using the *same* sequence of demand shocks.

```
def simulate(m, idx, k0, T=20_000, seed=0):
    "Simulate the equilibrium Markov process for (k, u)."
    mc = qe.MarkovChain(m.P, m.u)
    u_idx = mc.simulate_indices(T, init=len(m.u)//2, random_state=seed)
    ki = np.abs(m.k - k0).argmin()
    k_path = np.empty(T)
    for t in range(T):
        k_path[t] = m.k[ki]
        ki = idx[ki, u_idx[t]]
    p_path = m.a0 + m.u[u_idx] - m.a1 * k_path
    return k_path, p_path

k_low, p_low = simulate(m, idx, k0=30.0, seed=1)
k_high, p_high = simulate(m, idx, k0=150.0, seed=1)

fig, ax = plt.subplots(figsize=(9, 4.5))
ax.plot(k_low[:250], lw=2, label='$k_0 = 30$')
ax.plot(k_high[:250], lw=2, label='$k_0 = 150$')
ax.axhline(k_lo_star, color='k', ls='--', lw=1)
ax.axhline(k_hi_star, color='k', ls='--', lw=1, label='ergodic set')
ax.set_xlabel('$t$'); ax.set_ylabel('$k_t$')
ax.legend()
plt.tight_layout()
plt.show()
```

Both paths are drawn into the ergodic set and then fluctuate inside it forever.

The next figure compares the long-run distributions of capital computed from the two simulations, after discarding a burn-in sample.

```
burn = 2000
fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))

axes[0].hist(k_low[burn:], bins=60, density=True, alpha=0.5, label='from $k_0 = 30$')
axes[0].hist(k_high[burn:], bins=60, density=True, alpha=0.5, label='from $k_0 = 150$
↪')
axes[0].set_xlabel('$k$'); axes[0].set_ylabel('density')
axes[0].set_title('invariant distribution of capital')
axes[0].legend()

axes[1].hist(p_low[burn:], bins=60, density=True, alpha=0.5)
axes[1].set_xlabel('$p$'); axes[1].set_ylabel('density')
```

(continues on next page)

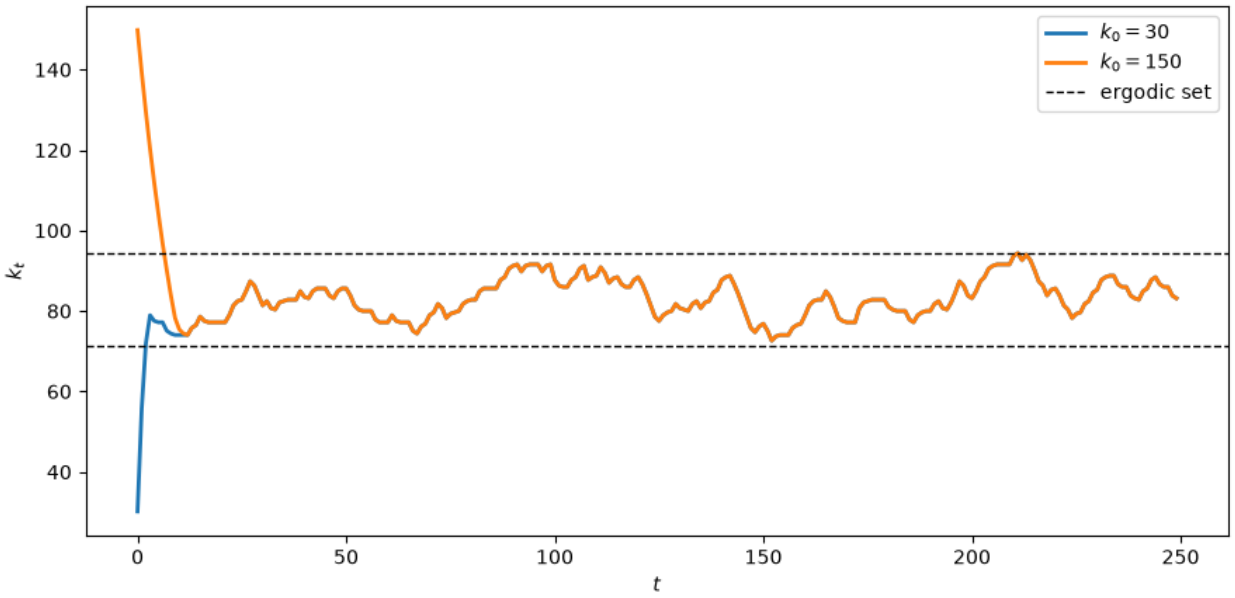


Fig. 117.3: Capital paths entering the ergodic set

(continued from previous page)

```
axes[1].set_title('invariant distribution of price')

plt.tight_layout()
plt.show()

print(f"mean capital from k_0 = 30:  {k_low[burn:].mean():.3f}")
print(f"mean capital from k_0 = 150: {k_high[burn:].mean():.3f}")
print(f"range visited: ({k_low[burn:].min():.2f}, {k_low[burn:].max():.2f})")
```

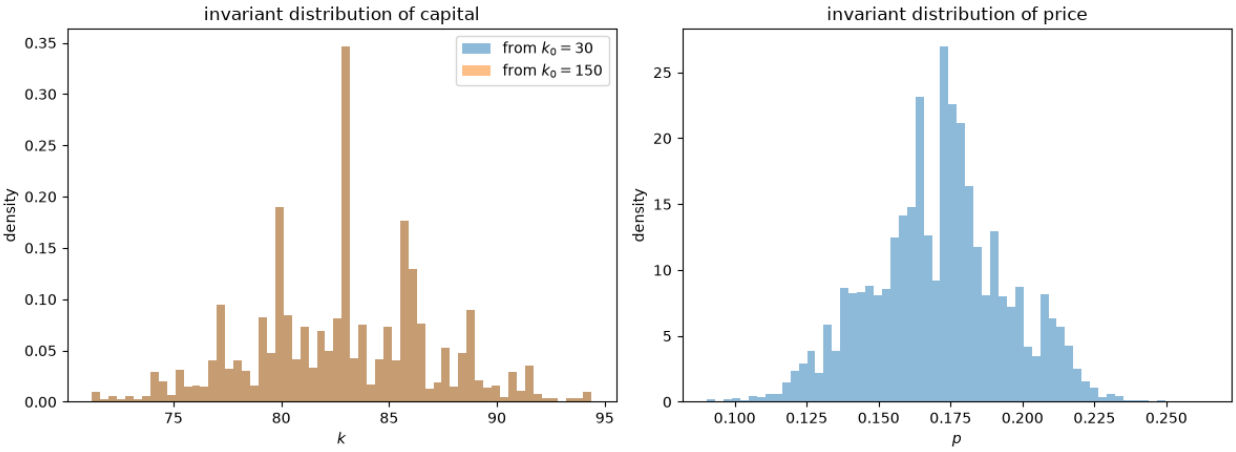


Fig. 117.4: Invariant distributions of capital and price

```
mean capital from k_0 = 30:  82.792
mean capital from k_0 = 150: 82.792
range visited: (71.23, 94.39)
```

The two histograms coincide, as *Theorem 117.9.1* promises, and the capital stock stays inside the ergodic set.

Finally, here is *Theorem 117.9.2* in action: time averages from a single realization converge to the mean of the invariant distribution.

```
running_mean = np.cumsum(k_low) / np.arange(1, len(k_low) + 1)

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(running_mean, lw=2)
ax.axhline(k_low[burn:].mean(), color='k', ls='--', lw=1, label='$k^*$')
ax.set_xlabel('$T$'); ax.set_ylabel(r'$T^{-1}\sum_{t \leq T} k_t$')
ax.set_xscale('log')
ax.legend()
plt.tight_layout()
plt.show()
```

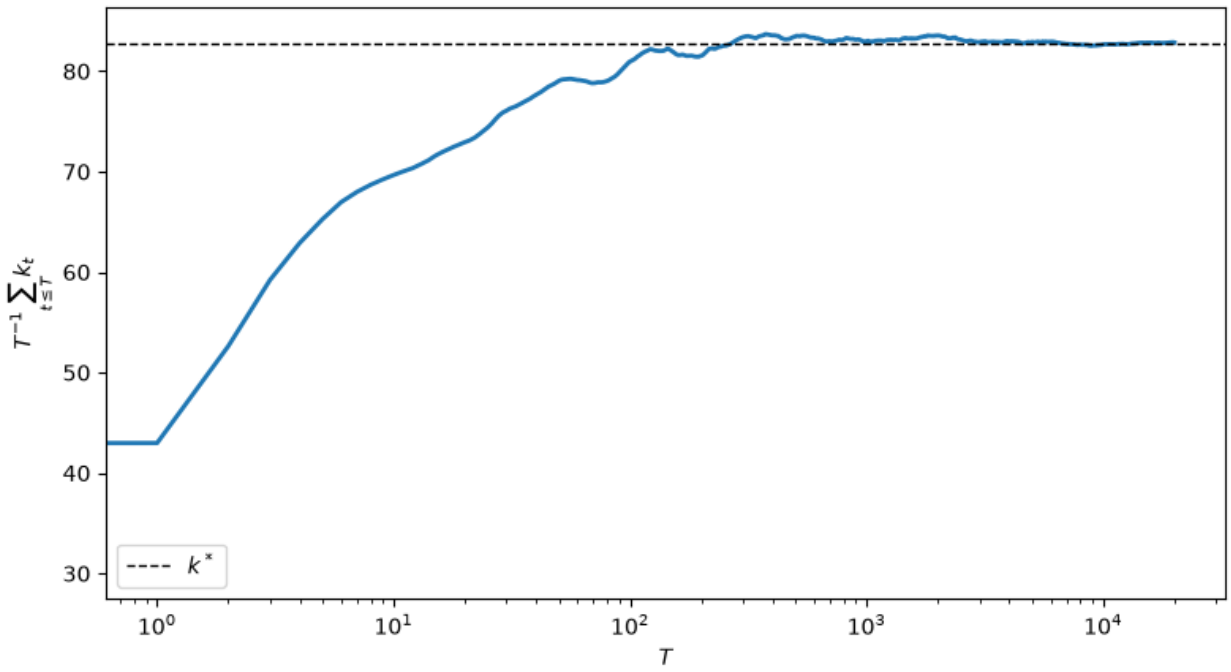


Fig. 117.5: Time average of capital

117.10 Relation to the rational expectations lecture

It is worth collecting the correspondences between this lecture and *Rational Expectations Equilibrium*.

	Rational Expectations Equilibrium	this lecture
uncertainty	none	Markov demand shifter u_t
adjustment costs	quadratic, $\gamma(y' - y)^2/2$	concave technology $k' = kh(x/k)$
equilibrium object	belief H with $Y' = H(Y)$	price process $\{p_t\}$, equivalently policy $x(k, u)$
equilibrium concept	H is a fixed point of Φ	anticipated price distribution equals actual
how it is computed	planning problem, solved as an LQ problem	planning problem, solved by dynamic programming
what the planner maximizes	consumer plus producer surplus	discounted consumer surplus (117.8)
equilibrium dynamics	$Y_{t+1} = \kappa_0 + \kappa_1 Y_t$	Markov process for (k_t, u_t)
long-run behavior	convergence to a steady state	convergence to an invariant distribution

The deepest common element is the strategy for computing an equilibrium.

In both lectures, the direct approach – guess a law of motion, compute the induced best response, and iterate – is unreliable, because that mapping need not be a contraction.

In both lectures, the remedy is to find a planning problem whose Euler equations coincide with the equilibrium conditions, and then to solve the planning problem by dynamic programming.

The lecture *Rational Expectations Equilibrium* verifies this correspondence by matching Euler equations for a particular linear-quadratic example.

Theorem 117.5.1 is the general statement: for this class of economies, the set of competitive equilibria and the set of solutions of the planning problem coincide, and both are singletons.

What the baby version cannot show, because it has no uncertainty, is the payoff that Lucas and Prescott were after: an equilibrium that is a *stationary stochastic process*, with an invariant distribution and ergodic time averages.

That is what makes it possible to confront such a model with data, and what led on to rational expectations econometrics.

The companion lecture *Optimal Growth Under Uncertainty and Tobin's q* pursues exactly this theme in a one-sector growth model.

Brock and Mirman [1972] prove there the counterparts of *Theorem 117.9.1* and *Theorem 117.9.2*: the distribution of capital converges to an invariant distribution that does not depend on initial conditions, and time averages along a single realization converge to population moments.

That lecture also shows what the shadow price of capital in such a planning problem becomes in a competitive equilibrium – namely Tobin's q – and examines a subtle question about the differentiability of the value function that the answer depends on.

117.11 Exercises

Exercise 117.11.1

The user cost of capital in (117.13) depends on the curvature parameter α of the adjustment technology through $h'(\delta) = \alpha$.

1. Explain why a *lower* α – meaning stronger adjustment costs – should reduce the long-run capital stock.
2. For the serially independent case, compute the stationary capital stock k^c for $\alpha \in \{0.4, 0.6, 0.8, 1.0\}$ and verify in each case that the marginal condition (117.13) holds.
3. Confirm that when $\alpha = 1$ the accumulation equation is $k_{t+1} = (1-\delta)k_t + x_t$ and the user cost is the textbook $\delta + r$.

Solution

A lower α makes h more concave, so a unit of investment buys less capacity at the margin.

Since $h'(\delta) = \alpha$, the interest component of the user cost, $r/h'(\delta) = r/\alpha$, rises as α falls.

A higher user cost must be matched by a higher expected price, and since demand slopes down, that means a smaller capital stock.

```
print(f"{'α':>5} {'k^c':>10} {'E[price]':>12} {'user cost':>12}")
for α in (0.4, 0.6, 0.8, 1.0):
    m_α = create_model(ρ=0.0, a=α, k_lo=5.0, k_hi=160.0, n_k=600)
    _, _, _, x_α = solve_model(m_α)
    kc_α = stationary_k(x_α[:, 0], m_α)
    price = m_α.a0 - m_α.a1 * kc_α
    cost = m_α.δ + m_α.r / h_prime(m_α.δ, m_α)
    print(f"{α:>5.1f} {kc_α:>10.3f} {price:>12.5f} {cost:>12.5f}")
```

```
    α        k^c     E[price]    user cost
  0.4     77.454      0.22546      0.22500
  0.6     81.594      0.18406      0.18333
  0.8     83.664      0.16336      0.16250
  1.0     84.958      0.15042      0.15000
```

Stronger adjustment costs (lower α) do indeed lower the long-run capital stock.

With $\alpha = 1$ we have $h(z) = 1 - \delta + z$, so $k' = k(1 - \delta + x/k) = (1-\delta)k + x$, and $h'(\delta) = 1$, so the user cost is $\delta + r$.

Exercise 117.11.2

Theorem 117.5.1 says that the planner's policy *is* the competitive equilibrium.

Verify this numerically, using the "big K, little k" logic of *Rational Expectations Equilibrium*.

Solve the problem of an individual price-taking firm that

- owns capital k_i and chooses k_i' subject to the same accumulation technology
- takes as given the aggregate capital stock K, which evolves according to the planner's policy computed above
- takes as given the price $p = a_0 + u - a_1 K$, which depends on aggregate, not own, capital

Then check that when the firm's own capital equals aggregate capital, $k_i = K$, the firm chooses exactly what the planner chooses.

Use a coarser grid for the firm's own capital to keep the computation small.

Solution

The firm's Bellman equation is

$$v_i(k_i, K, u) = \max_{k_i'} \{p(K,u)k_i - x(k_i, k_i') + \beta \mathbb{E}\left[v_i(k_i', K', u') \mid u\right]\}$$

where K' follows the planner's law of motion.

Note that the firm's own capital affects its revenue but not the price.

```
def firm_problem(m, idx_agg, n_i=80, tol=1e-8, maxit=1000, howard=20):
    "Solve an individual firm's problem taking the aggregate law of motion as␣
↪given."
    sub = np.linspace(0, len(m.k) - 1, n_i).astype(int)    # firm grid ⊂ aggregate␣
↪grid
    ki = m.k[sub]
    n_K, n_u = len(m.k), len(m.u)

    Xi = ki[:, None] * ((ki[None, :] / ki[:, None])**(1/m.α) - (1 - m.δ))
    Ri = np.where(Xi >= 0, -Xi, -1e12)                     # (k_i, k_i')
    price = m.a0 + m.u[None, :] - m.a1 * m.k[:, None]      # (K, u)
    revenue = ki[:, None, None] * price[None, :, :]        # (k_i, K, u)

    v_i = np.zeros((n_i, n_K, n_u))
    u_cols = np.arange(n_u)[None, :]
    for it in range(maxit):
        EV = np.tensordot(v_i, m.P, axes=([2], [1]))       # E[v_i(k_i', K', u') |␣
↪u]
        cont = EV[:, idx_agg, u_cols]                      # impose K' = planner's␣
↪choice
        obj = Ri[:, :, None, None] + m.β * cont[None, :, :, :]
        pol = obj.argmax(axis=1)
        v_new = revenue + np.take_along_axis(obj, pol[:, None, :, :], axis=1)[:, 0,
↪ :, :]

        for _ in range(howard):
            EV = np.tensordot(v_new, m.P, axes=([2], [1]))
            cont = EV[:, idx_agg, u_cols]
            v_new = (revenue + Ri[np.arange(n_i)[:, None, None], pol]
                     + m.β * np.take_along_axis(cont, pol, axis=0))

        if np.max(np.abs(v_new - v_i)) < tol:
            v_i = v_new
            break
        v_i = v_new

    return ki, sub, pol

ki, sub, pol_firm = firm_problem(m, idx)

# compare the firm's choice with the planner's, evaluated at k_i = K
gaps = []
```

```
for a, K_i in enumerate(sub):
    for j in range(len(m.u)):
        gaps.append(abs(ki[pol_firm[a, K_i, j]] - m.k[idx[K_i, j]]))
gaps = np.array(gaps)

print(f"firm grid spacing:                    {np.diff(ki).mean():.3f}")
print(f"mean |firm choice - planner choice|: {gaps.mean():.3f}")
print(f"max  |firm choice - planner choice|: {gaps.max():.3f}")
```

```
firm grid spacing:                    1.772
mean |firm choice - planner choice|: 0.590
max  |firm choice - planner choice|: 1.404
```

The discrepancies are smaller than the spacing of the firm's own capital grid.

So the price-taking firm, responding optimally to the price process that the planner's allocation generates, chooses to do exactly what the planner does.

That is the content of *Theorem 117.5.1*, and it is the Lucas-Prescott counterpart of the fixed point condition $H(Y) = h(Y, Y)$ in *Rational Expectations Equilibrium*.

Exercise 117.11.3

Theorem 117.9.1 says that the invariant distribution does not depend on initial conditions, but it says nothing about how *wide* that distribution is.

Investigate how serial correlation in demand affects the ergodic set.

1. For $\rho \in \{0.0, 0.5, 0.9, 0.98\}$, compute the ergodic bounds $\underline{k}$ and $\bar{k}$.
2. Simulate each economy and compare the invariant distributions of capital.
3. Explain the pattern. Why does the $\rho = 0$ case produce a degenerate distribution for capital?

Solution

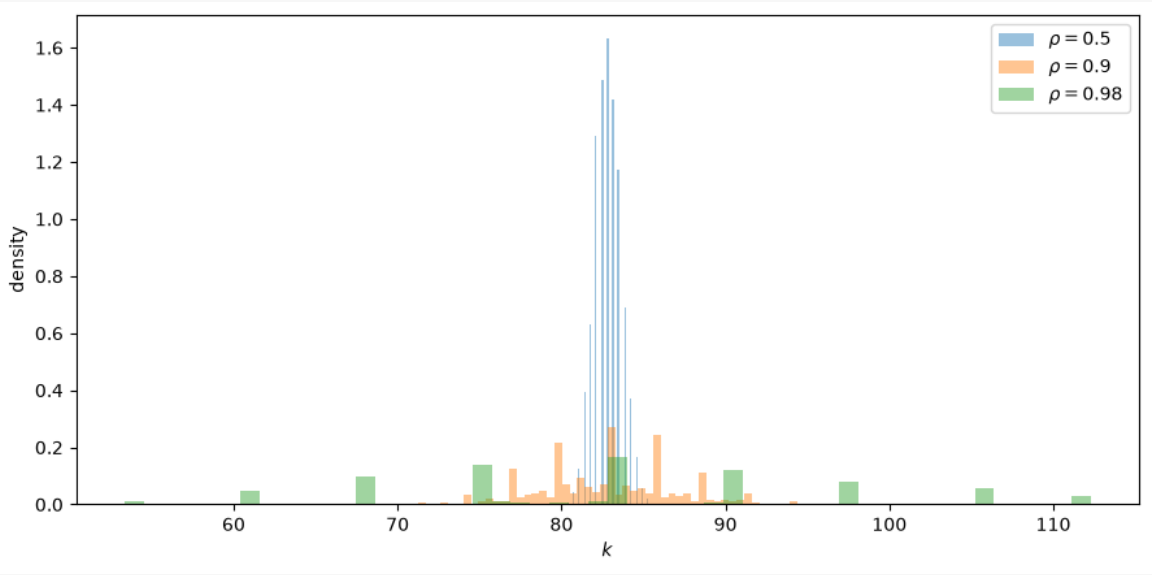
Here is one solution.

```
fig, ax = plt.subplots(figsize=(9, 4.5))
print(f"{'ρ':>6} {'k_lo':>9} {'k_hi':>9} {'width':>9} {'std(k)':>9}")

for ρ in (0.0, 0.5, 0.9, 0.98):
    m_ρ = create_model(ρ=ρ)
    _, idx_ρ, _, x_ρ = solve_model(m_ρ)
    b = np.array([stationary_k(x_ρ[:, j], m_ρ) for j in range(len(m_ρ.u))])
    k_ρ, _ = simulate(m_ρ, idx_ρ, k0=80.0, seed=3)
    print(f"{ρ:>6.2f} {b.min():>9.2f} {b.max():>9.2f} "
          f"{b.max()-b.min():>9.2f} {k_ρ[2000:].std():>9.3f}")
    if ρ > 0:      # the ρ = 0 distribution is a spike at k^c, so we omit it here
        ax.hist(k_ρ[2000:], bins=50, density=True, alpha=0.45, label=f'$\\rho = {ρ}
↪$')

ax.set_xlabel('$k$'); ax.set_ylabel('density')
ax.legend()
plt.tight_layout()
plt.show()
```

```
     ρ      k_lo      k_hi     width    std(k)
0.00     82.81     82.81      0.00     0.000
0.50     80.00     85.61      5.61     0.797
0.90     71.23     94.39     23.16     4.156
0.98     53.33    112.28     58.95    13.552
```



The more persistent is demand, the wider the ergodic set and the more dispersed the invariant distribution of capital.

(The figure omits $\rho = 0$, whose distribution is a spike at k^c that would dwarf the others.)

The reason is the one Lucas and Prescott emphasize.

Investment responds to *news about future demand*, not to current demand as such.

When $\rho = 0$, a demand shift conveys no information about the future, so investment does not respond at all, and capital converges to the single deterministic value k^c of *Theorem 117.8.1*: the invariant distribution of capital is degenerate even though prices keep fluctuating.

As ρ rises, a high demand state signals a sustained period of high prices, so firms invest more, and the capital stock inherits the persistence of demand.

CHAPTER

EIGHT

EQUILIBRIUM SEARCH AND UNEMPLOYMENT

Contents

- *Equilibrium Search and Unemployment*
  - *Overview*
  - *The economy*
  - *A worker's problem*
  - *Employment, wages, and two thresholds*
  - *A computable version*
  - *Economy-wide equilibrium*
  - *Relation to other lectures*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

118.1 Overview

This lecture studies Lucas and Prescott [1974], the labor-market companion to the paper on investment that we studied in *Investment Under Uncertainty*.

The two papers were written by the same authors at the same place within a few years of each other, and they use the same tools: a Bellman equation solved by successive approximations, Blackwell's theorem, and a Markov process whose invariant distribution describes the long run.

What changes is the question.

The question is one that the search literature of the 1960s had raised but could not answer.

Job Search I: The McCall Search Model describes that literature's central problem: a worker draws wage offers from a known distribution and decides when to stop searching.

It is a model of an individual, and it takes the wage distribution as given.

But where does the distribution come from?

Lucas and Prescott set out to answer that, and discovered that the answer requires a general equilibrium model:

> *optimal labor supply behavior and the wage distribution on which it is based must be simultaneously determined within a model of market (as opposed to individual) behavior.*

They build an economy of many spatially separate markets – "islands" – each with its own randomly shifting demand for labor.

Workers in a market can work at the wage that clears it, or leave.

Leaving means earning nothing this period and arriving somewhere else next period, so **unemployment is the act of moving between markets**.

Everything is competitive, everyone optimizes, and expectations are rational in the sense of Muth [1961]: the distribution that workers use to evaluate their prospects is the one that their own decisions generate.

The paper's central object is a single number, λ, the expected present value of search.

An individual market takes λ as given, in the way a competitive firm takes a price as given.

Economy-wide, λ adjusts until the workforce that markets want to hold equals the workforce the economy has.

In equilibrium some markets always have more workers than they want to employ, so

Note

Unemployment is positive in every period.

Lucas and Prescott argue that this unemployment need not signal an inefficiency, and the argument is worth stating in their terms, although the paper proves no welfare theorem, leaves open whether actual labor markets have important external effects, and this lecture proves nothing either.

There is no congestion externality in the model and no wage rigidity.

The searching worker who crowds his fellows does so in the same way that any seller of a good crowds other sellers, and λ, like any other equilibrium price, signals what the paper calls "the correct social return" to an additional searcher.

Let's start with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
import quantecon as qe
from collections import namedtuple
```

118.2 The economy

There are many markets, indexed by the state of demand in them.

Demand shifts are *independent across markets*, so that aggregate demand is constant, and *persistent within* a market, so that today's demand tells a worker something about tomorrow's.

In each market a competitive industry produces a single good with labor as the only input.

If $p = D(s, Q)$ is the inverse demand curve in demand state s and $q(n)$ is the production function of each of the m identical firms, then profit maximization sets the wage equal to the marginal value product of labor, and the function

$$R(s, n) = D\big(s, m\, q(n)\big)\, q'(n) \tag{118.1}$$

summarizes everything about the demand side of this labor market.

Lucas and Prescott discard D and q and place assumptions directly on R: it is positive, bounded and differentiable, with

$$R_s(s,n) > 0, \qquad R_n(s,n) < 0.$$

The demand state s follows a Markov process with transition density $f(s',s)$ possessing a unique stationary distribution.

118.2.1 Timing

A market begins a period with a **workforce** y, which is the most it can employ.

Then, within the period:

1. demand s is realized and the labor market clears, employing $n \leq y$ workers at the wage $w = R(s,n)$
2. the $y - n$ workers who are not employed earn nothing and join the pool of searchers
3. employed workers remain in the market into next period; searchers are allocated across markets and arrive next period

So a worker is *unemployed for exactly one period* each time he moves, and the unemployed are workers in transit.

118.2.2 The value of search

Because demand shifts are independent across markets and there are many of them, the economy-wide distribution of workers across markets is constant even though each market's workforce fluctuates.

A searcher's prospects therefore do not depend on the market he is leaving, and the expected present value of search is a single number λ.

Searchers are allocated to markets in a way that *equalizes* the expected value of arriving, so no market offers an arriving worker more than λ.

We take λ as given for now and determine it in *Economy-wide equilibrium*.

118.3 A worker's problem

Let $v(s,y,\lambda)$ be the expected present value of the earnings of a worker who finds himself in a market with demand s and workforce y at the start of a period, and let $\beta \in (0,1)$ be the discount factor.

Lucas and Prescott arrive at a functional equation for v by considering three cases.

Case A: some workers leave. Departing workers get λ. The workers who stay can have no more, since the leavers could have stayed, and no less, since they could have left. So $v = \lambda$.

Case B1: all workers stay and no new workers arrive. Current employment is y, the wage is $R(s,y)$, and next period's workforce is still y:

$$v(s,y,\lambda) = R(s,y) + \beta \int v(s',y,\lambda) f(s',s)\, ds'.$$

Case B2: all workers stay and new workers arrive. Arriving workers are searchers, so the continuation value of being in this market next period must be exactly λ:

$$v(s,y,\lambda) = R(s,y) + \lambda.$$

Cases B1 and B2 differ according to whether the continuation value at the unchanged workforce exceeds λ, and in case A the worker takes λ.

All three collapse into one equation:

$$v(s, y, \lambda) = \max\left\{\lambda,\ R(s, y) + \min\left[\lambda,\ \beta \int v(s', y, \lambda) f(s', s)\, ds'\right]\right\}. \tag{118.2}$$

Equation (118.2) repays a careful look, because its two operators do different jobs.

The *max* is the worker's own option: he can always leave and collect λ.

The *min* is not the worker's choice at all.

It is a restriction that competition among searchers imposes: if the continuation value in this market exceeded λ, searchers would flow in until it did not.

Proposition 118.3.1

For each λ, equation (118.2) has a unique bounded continuous solution v, and successive approximations converge to it from any starting point.

The solution is nondecreasing in s and in λ and nonincreasing in y.

The proof is the one used throughout *Investment Under Uncertainty*: the operator defined by the right side of (118.2) is monotone and satisfies a discounting condition, so Blackwell's theorem [Blackwell, 1965] applies.

That v is nonincreasing in y is worth a word: a worker is worse off the more crowded his market, because a larger workforce means a lower marginal product and a lower wage.

118.4 Employment, wages, and two thresholds

With v in hand, the rest of a market's behavior follows from two threshold rules.

118.4.1 How many work

Let $\bar{n}(s, \lambda)$ be the employment at which the marginal worker is just indifferent between staying and searching:

$$R(s, \bar{n}) + \beta \int v(s', \bar{n}, \lambda) f(s', s)\, ds' = \lambda. \tag{118.3}$$

The left side is what a worker gets by staying when employment and next period's workforce are both $\bar{n}$: this period's wage plus the discounted continuation.

It is non-increasing in $\bar{n}$.

Once a market is so crowded that the marginal product has reached zero, the value of being there has fallen to λ and the left side is flat at $\beta\lambda < \lambda$, so (118.3) cannot hold there.

On the integer workforce grid used below, (118.3) rarely holds with equality, and employment is the largest $\bar{n}$ at which the left side is at least λ.

Employment is then

$$n(s, y, \lambda) = \min\{\bar{n}(s, \lambda),\ y\}, \qquad w = R(s, n), \tag{118.4}$$

and unemployment in this market is $y - n$, which is positive exactly when the workforce exceeds $\bar{n}$.

118.4.2 How many arrive

A market attracts searchers only when it employs its whole workforce and the continuation value still exceeds λ.

Let $\bar{y}(s, \lambda)$ be the workforce at which that stops being true:

$$\beta \int v(s', \bar{y}, \lambda) f(s', s)\, ds' = \lambda. \tag{118.5}$$

Arrivals fill a market up to $\bar{y}(s, \lambda)$ and no further.

Note that $\bar{y} \leq \bar{n}$, since $R \geq 0$ means that any workforce whose continuation alone covers λ also covers it once the wage is added.

The inequality is strict wherever the marginal product is positive: a market just worth entering for a searcher is then strictly inside the range where everyone works.

In the worst demand states of the calibration below both thresholds are zero, and the market employs nobody and attracts nobody.

118.4.3 Three regions, one transition law

Together, (118.4) and (118.5) partition the (s, y) plane into the three regions of the paper's Figure 3, and they give the law of motion for the workforce a very simple form:

$$y' = \min\left\{ \max\{y,\ \bar{y}(s, \lambda)\},\ \bar{n}(s, \lambda)\right\}. \tag{118.6}$$

A market that is too crowded sheds workers down to $\bar{n}(s, \lambda)$; a market that is too empty is filled up to $\bar{y}(s, \lambda)$; in between, nobody moves.

It is worth pausing on an apparent inconsistency between (118.6) and (118.2), because it is the place where a reader is most likely to suspect an error.

The Bellman equation evaluates the continuation at the *unchanged* workforce y, while the transition law moves the workforce somewhere else.

Both are right, and the reason is that the max and the min have already dealt with the two cases in which the workforce moves: where arrivals come, the min replaces the continuation by λ, which is exactly what a market filled to $\bar{y}$ is worth to anyone in it, and where workers leave, the max replaces the whole expression by λ.

The continuation at the unchanged y is therefore used only in the region where the workforce really is unchanged.

Note

The paper lets arrivals be random: a market short of workers receives $a(s, y, \lambda)\, x$ of them, where x is a positive random variable with mean one, so that search "eliminates rents on average" rather than exactly.

Like the numerical examples in the paper's appendix, we set $x \equiv 1$, which gives (118.6).

118.5 A computable version

We follow the appendix of Lucas and Prescott [1974]: a finite number of demand states, an integer workforce, and value function iteration on (118.2).

For the demand process we use the example in the paper's own footnote: $\log s$ follows an AR(1), discretized with the method of Tauchen.

For R we take a marginal revenue schedule that falls linearly in employment and reaches zero at a capacity N,

$$R(s, n) = s\,(1 - n/N)^{+},$$

which is the shape of the schedules drawn in the top panels of the paper's Figure 5.

The capacity matters for more than convenience.

Lucas and Prescott require that employment stay bounded as demand grows without bound, and a schedule that never reaches zero violates that: markets would absorb unlimited numbers of workers in good states, and the answers would depend on where we truncated the grid rather than on the economics.

We will see below that this assumption is not a technicality.

The shape of R does something else that is easy to miss and that Exercise 118.8.5 takes up.

It sets how far employment moves when demand moves, and therefore how much there is to gain from sending a worker from one market to another.

```
Model = namedtuple('Model', ['P', 's', 'y', 'R', 'β'])

def create_model(n_s=7, ρ=0.9, sd=0.26, β=0.9, N=200, y_max=260):
    """An island economy.

    ρ and sd are the persistence and the stationary standard deviation of
    log demand, so that changing ρ alone does not change how variable
    demand is.  N is a market's capacity: marginal revenue reaches zero there.
    """
    mc = qe.markov.tauchen(n_s, ρ, sd*np.sqrt(1 - ρ**2))
    s = np.exp(mc.state_values)
    y = np.arange(y_max + 1)
    R = np.maximum(s[:, None]*(1 - y[None, :]/N), 0.0)
    R[:, 0] = 0.0
    return Model(P=mc.P, s=s, y=y, R=R, β=β)

model = create_model()
print("demand states:", model.s.round(3))
```

```
demand states: [0.458 0.595 0.771 1.    1.297 1.682 2.181]
```

The Bellman operator is (118.2) with the integral replaced by a sum.

We need some value of search to solve it, so we use $\lambda = 7.1$ here, a number that *Economy-wide equilibrium* will turn out to have chosen for a reason.

```
def worker_value(model, λ, tol=1e-11, max_iter=10_000):
    "Solve the worker's functional equation by successive approximations."
    P, R, β = model.P, model.R, model.β
    v = np.full(R.shape, λ)
    for i in range(max_iter):
        v_new = np.maximum(λ, R + np.minimum(λ, β*(P @ v)))
```

(continues on next page)

(continued from previous page)

```
        v_new[:, 0] = λ                    # an empty market holds no workers
        if np.max(np.abs(v_new - v)) < tol:
            return v_new, i
        v = v_new
    return v, max_iter

v, iters = worker_value(model, 7.1)
print(f"converged in {iters} iterations")
```

```
converged in 107 iterations
```

The two thresholds are the largest employment, and the largest workforce, at which the relevant value still covers λ.

```
def thresholds(model, v, λ):
    "The employment threshold n̄(s) and the arrival threshold ȳ(s)."
    P, R, β = model.P, model.R, model.β
    EV = β*(P @ v)                 # the continuation value at an unchanged workforce
    H = R + EV                     # the value of staying on as the marginal worker
    last = lambda A: np.array([np.max(np.where(A[i] >= λ)[0]) if (A[i, 1:] >= λ).any()
                               else 0 for i in range(len(model.s))])
    n_bar = last(H)
    y_bar = last(EV)               # ≤ n_bar automatically, because R ≥ 0
    return n_bar, y_bar

n_bar, y_bar = thresholds(model, v, 7.1)
print("n̄(s) =", n_bar)
print("ȳ(s) =", y_bar)
```

```
n̄(s) = [  0   5  41  71  97 118 134]
ȳ(s) = [  0   0   0  32  63  88 107]
```

Employment is capped by $\bar{n}$, arrivals fill up to $\bar{y}$, and (118.6) moves the workforce.

```
def stationary_distribution(model, n_bar, y_bar, tol=1e-14, max_iter=100_000):
    "The invariant distribution of the Markov chain on (s, y)."
    P, y = model.P, model.y
    n_s, n_y = len(model.s), len(y)
    ψ = np.ones((n_s, n_y))/(n_s*n_y)
    y_next = np.clip(y[None, :], y_bar[:, None], n_bar[:, None])
    for _ in range(max_iter):
        moved = np.array([np.bincount(y_next[i], weights=ψ[i], minlength=n_y)
                          for i in range(n_s)])
        ψ_new = P.T @ moved
        if np.max(np.abs(ψ_new - ψ)) < tol:
            return ψ_new
        ψ = ψ_new
    return ψ

def market_outcomes(model, λ):
    "Everything about a market at a given value of search."
    v, _ = worker_value(model, λ)
    n_bar, y_bar = thresholds(model, v, λ)
    ψ = stationary_distribution(model, n_bar, y_bar)
    y = model.y
    employment = np.minimum(y[None, :], n_bar[:, None])
    unemployment = y[None, :] - employment
```

(continues on next page)

(continued from previous page)

```
    wage = np.take_along_axis(model.R, np.maximum(employment, 1), axis=1)
    total = (ψ*employment).sum()
    return dict(v=v, n_bar=n_bar, y_bar=y_bar, ψ=ψ,
                workforce=(ψ*y).sum(),
                employment=total,
                unemployment=(ψ*unemployment).sum(),
                # the mean wage per employed worker; empty markets get no weight
                wage=(ψ*employment*wage).sum()/total if total > 0 else np.nan)

out = market_outcomes(model, 7.1)
print(f"workforce {out['workforce']:.2f}, employment {out['employment']:.2f}, "
      f"unemployment {out['unemployment']:.2f}")
```

```
workforce 50.49, employment 49.16, unemployment 1.33
```

118.5.1 What a market looks like

The thresholds describe the whole of a market's behavior, so it is worth seeing them.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))

axes[0].plot(model.s, out['n_bar'], 'o-', lw=2, label=r'$\bar n(s)$: employment')
axes[0].plot(model.s, out['y_bar'], 's--', lw=2, label=r'$\bar y(s)$: arrivals fill to
↪')
axes[0].set_xlabel('demand state $s$')
axes[0].set_ylabel('workers')
axes[0].legend()

employs = out['n_bar'] > 0                # a state with n̄(s) = 0 employs nobody
wage_at_nbar = model.R[np.arange(len(model.s)), out['n_bar']]
axes[1].plot(out['n_bar'][employs], wage_at_nbar[employs], 'o-', lw=2)
axes[1].set_xlabel(r'employment $\bar n(s)$')
axes[1].set_ylabel('wage')
plt.tight_layout()
plt.show()
```

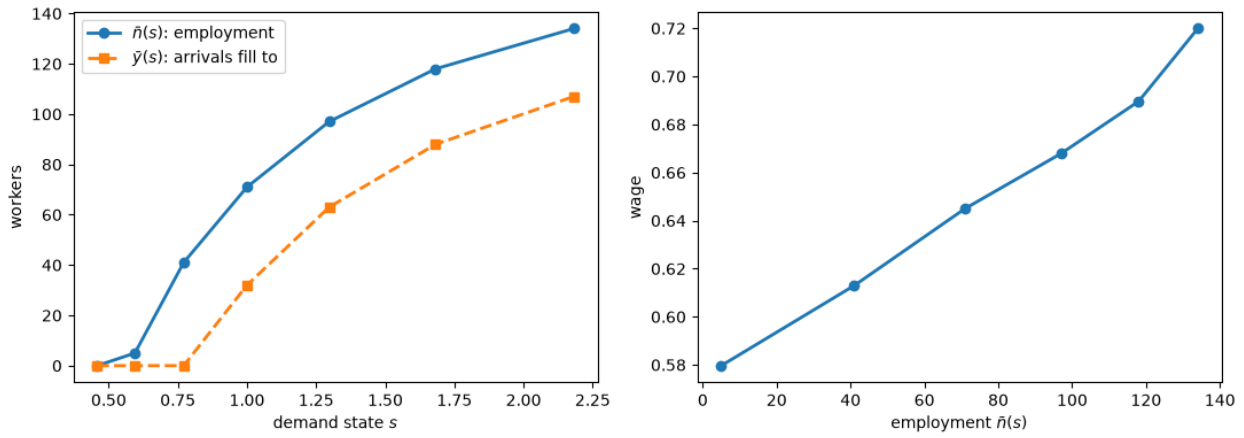


Fig. 118.1: The two thresholds, and the locus of employment and wages as demand shifts

The right panel is the curve that Lucas and Prescott call SS, the pairs of employment and wages that a market would trace out as demand shifted if its workforce never constrained it.

It leaves out the worst demand states, where $\bar{n}(s) = 0$, because a market there employs nobody and pays no wage.

The curve is *not flat*.

A flat SS would mean that workers had a single reservation wage, accepting work above it and searching below it, which is what the partial-equilibrium search models of *Job Search I: The McCall Search Model* *derive*: there a reservation wage is the conclusion, and here it fails.

The reason is instructive: a change in demand moves the current wage, but it also carries information about future wages, and it changes the workforce that the market will carry into the future.

Those effects pull in opposite directions, so the locus can slope either way; in this calibration it happens to slope upward throughout.

118.5.2 The workforce distribution and a sample path

Because arrivals are deterministic, the workforce moves between a small number of values, each one a threshold of some demand state.

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4))

marginal = out['ψ'].sum(axis=0)
support = marginal > 1e-10
axes[0].stem(model.y[support], marginal[support], basefmt=' ')
axes[0].set_xlabel('workforce $y$')
axes[0].set_ylabel('probability')

mc = qe.MarkovChain(model.P)
idx = mc.simulate_indices(60, init=3, random_state=0)
path_y = np.empty(60, dtype=int)
path_y[0] = out['n_bar'][idx[0]]
for t in range(59):
    path_y[t+1] = np.clip(path_y[t], out['y_bar'][idx[t]], out['n_bar'][idx[t]])
path_n = np.minimum(path_y, out['n_bar'][idx])

axes[1].plot(path_y, lw=2, label='workforce')
axes[1].plot(path_n, lw=2, ls='--', label='employment')
axes[1].fill_between(range(60), path_n, path_y, alpha=0.3, label='unemployed')
axes[1].set_xlabel('period')
axes[1].set_ylabel('workers')
axes[1].legend()
plt.tight_layout()
plt.show()
```

The right panel shows the mechanism of the model in one picture.

A market is at rest whenever demand does not move.

When demand falls, the gap between the two lines opens for exactly one period: that period's leavers are the unemployed, and by the next period they are somewhere else.

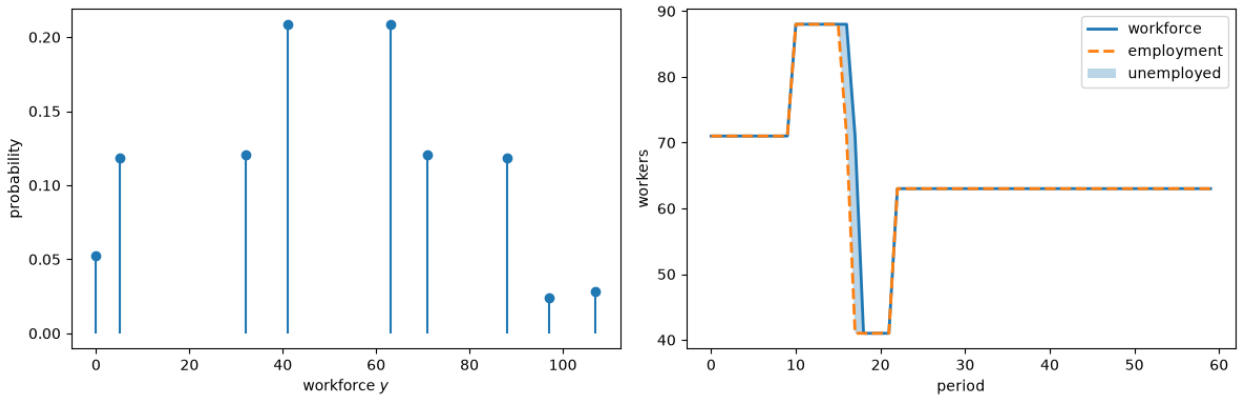


Fig. 118.2: The invariant distribution of the workforce, and one market over sixty periods

118.6 Economy-wide equilibrium

So far λ has been a parameter.

Now close the model.

The economy has a fixed workforce, μ per market.

For any λ, the calculation above delivers the workforce that a market holds on average, and summing across markets gives the economy's **demand for workforce**,

$$E(y;\lambda) = \int y\,\Psi(ds, dy; \lambda). \tag{118.7}$$

A higher λ makes every market shed workers sooner and fill up less, so $E(y;\lambda)$ is decreasing.

Equilibrium is the λ^* at which markets want to hold exactly the workers the economy has:

$$E(y;\lambda^*) = \mu. \tag{118.8}$$

Because the schedule is decreasing, bisection finds λ^*.

```
def workforce_demand(model, λ):
    "The average workforce per market at the value of search λ."
    return market_outcomes(model, λ)['workforce']

def equilibrium_λ(model, μ, lo=0.3, hi=40.0, tol=1e-4):
    "Bisect on the decreasing demand schedule E(y; λ) = μ."
    while hi - lo > tol:
        mid = 0.5*(lo + hi)
        if workforce_demand(model, mid) > μ:
            lo = mid
        else:
            hi = mid
    return 0.5*(lo + hi)

μ = 50                          # workers per market, against a capacity of 200
λ_star = equilibrium_λ(model, μ)
eq = market_outcomes(model, λ_star)

print(f"equilibrium value of search  λ* = {λ_star:.3f}")
```

(continues on next page)

(continued from previous page)

```
print(f"workforce per market            = {eq['workforce']:.2f}   (target {μ})")
print(f"employment per market           = {eq['employment']:.2f}")
print(f"unemployment per market         = {eq['unemployment']:.2f}")
print(f"unemployment rate               = {100*eq['unemployment']/eq['workforce']:.2f}
 ↪%")
print(f"separation rate                 = "
      f"{100*eq['unemployment']/eq['employment']:.2f}% of the employed per period")
print(f"mean wage                       = {eq['wage']:.3f}")
```

```
equilibrium value of search  λ* = 7.121
workforce per market            = 50.00   (target 50)
employment per market           = 48.66
unemployment per market         = 1.34
unemployment rate               = 2.68%
separation rate                 = 2.75% of the employed per period
mean wage                       = 0.780
```

Lucas and Prescott draw this as a demand curve meeting a vertical supply curve, which is exactly what it is.

```
grid = np.linspace(0.6, 15.0, 45)
demand = np.array([market_outcomes(model, λ) for λ in grid])
E_y = np.array([d['workforce'] for d in demand])
E_u = np.array([d['unemployment'] for d in demand])

fig, axes = plt.subplots(1, 2, figsize=(11, 4))

axes[0].plot(E_y, grid, lw=2, label=r'demand $E(y;\lambda)$')
axes[0].axvline(μ, color='C1', lw=2, ls='--', label=r'supply $\mu$')
axes[0].plot(eq['workforce'], λ_star, 'ko', ms=8)
axes[0].set_xlabel('workforce per market')
axes[0].set_ylabel(r'value of search $\lambda$')
axes[0].legend()

axes[1].plot(E_y, E_u, lw=2)
axes[1].plot(eq['workforce'], eq['unemployment'], 'ko', ms=8)
axes[1].set_xlabel('workforce per market')
axes[1].set_ylabel('unemployment per market')
plt.tight_layout()
plt.show()
```

The left panel is the paper's Figure 4.

The right panel corresponds to the bottom panels of the paper's Figure 5, which plot average unemployment against λ for two numerical examples; since $E(y;\lambda)$ is decreasing, plotting it against the workforce instead shows the same shape.

The right panel is the more interesting of the two, because unemployment is *not monotone* in the size of the workforce, a feature that Lucas and Prescott report from their own examples.

Where workers are scarce, λ is high, every market employs all it has, and almost nobody is ever shed.

Where workers are abundant, λ is low, every market is pressed against its capacity, and the two thresholds are squeezed together there, so even a large shift in demand moves few workers.

Unemployment is largest in between, and this is where the bounded capacity earns its keep: with a marginal revenue schedule that never reached zero, the thresholds would keep rising as λ fell, unemployment would rise without bound, and the right panel would have no peak at all.

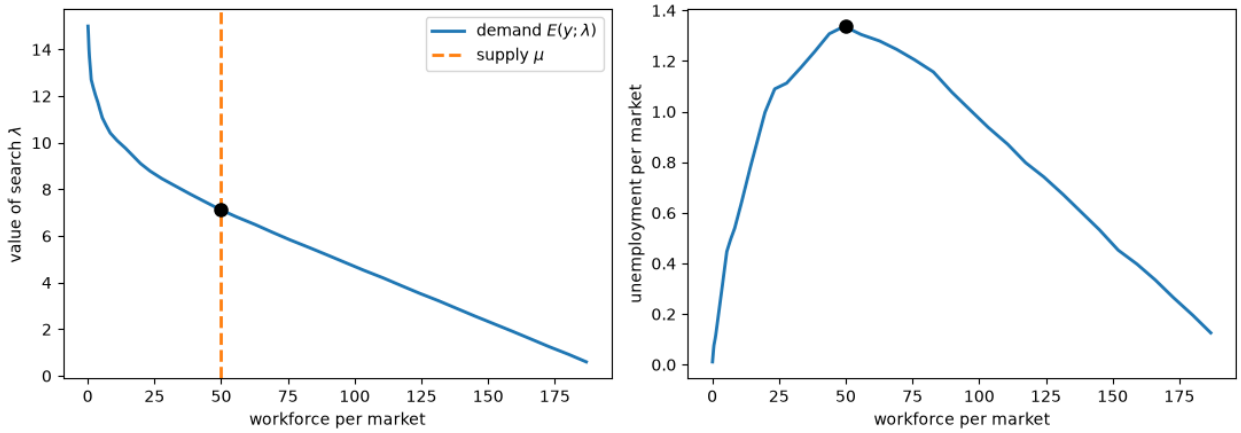


Fig. 118.3: The market for workers, and unemployment along the demand curve

118.6.1 Would anyone find it?

Lucas and Prescott devote a section of the paper to a question that the equilibrium itself does not answer: we have assumed that workers know λ, and said nothing about how they would come to know it.

Their answer is an adaptive story.

Suppose workers believe the value of search is higher than it is.

Then more of them leave their markets than can be placed at that return, so searchers and stayers alike are disappointed in the wages they expected, and they revise the number down.

Writing g for a decreasing function that vanishes at zero,

$$\frac{d\lambda}{dt} = g(\lambda - \lambda^*),$$

and since the demand schedule slopes down, the equilibrium is stable under this revision.

Note

That story and the algorithm we just ran rest on the same fact.

`equilibrium_λ` bisects on a decreasing schedule: when the candidate λ is too high, markets want fewer workers than the economy has, and the candidate comes down.

Bisection brackets the root rather than following the paper's adjustment $d\lambda/dt = g(\lambda - \lambda^*)$, but both work for the same reason: $E(y;\lambda)$ is decreasing in λ.

Lucas and Prescott are careful to add that this kind of stability "can have no relevance to the dynamic response to regularly recurring shocks", which is the distinction between learning an equilibrium and responding within one.

118.6.2 Where markets sit

The three regions of (118.6) are a partition of the state space, so the invariant distribution says how much of the economy is in each.

```
y2 = model.y[None, :]
sheds = y2 > eq['n_bar'][:, None]                    # region I: workers leave
fills = y2 < eq['y_bar'][:, None]                    # region III: workers arrive
still = ~sheds & ~fills                              # region II: nobody moves

for name, region in (('shedding workers', sheds), ('receiving workers', fills),
                     ('neither', still)):
    print(f"{name:>20}: {100*eq['ψ'][region].sum():5.1f}% of markets")
print(f"{'with no workers':>20}: {100*eq['ψ'][:, 0].sum():5.1f}% of markets")
```

```
    shedding workers:   5.6% of markets
   receiving workers:   5.2% of markets
             neither:  89.2% of markets
      with no workers:   5.2% of markets
```

Nearly nine markets in ten are completely still in any given period.

All of the economy's unemployment comes from the one in twenty that is contracting, and an equal number of markets is taking those workers in.

The last line is the extreme case of a contracting market, and it is visible in the thresholds printed earlier as the leading zeros of $\bar{n}(s)$ and $\bar{y}(s)$: in the worst demand states a market employs nobody and attracts nobody.

Those markets close, and they stay closed until they draw a state good enough to be worth entering again.

118.6.3 The wage distribution

We can now answer the question the Overview began with.

The search literature hands its worker a distribution of wage offers and asks what he should do with it.

Here that distribution is an output: every market pays the marginal product of labor in it, and what a market pays depends on the demand it drew and on how many workers chose to be there.

```
employed = np.minimum(model.y[None, :], eq['n_bar'][:, None])
wage = np.take_along_axis(model.R, np.maximum(employed, 1), axis=1)
weight = eq['ψ']*employed                    # weight a market by the workers in it
weight = weight/weight.sum()

w, p = wage.ravel(), weight.ravel()
order = np.argsort(w)
w, p = w[order], p[order]
cdf = np.cumsum(p)
quantile = lambda a: w[np.searchsorted(cdf, a)]
mean = (w*p).sum()

fig, ax = plt.subplots(figsize=(7.5, 4))
ax.hist(w, bins=30, weights=p, color='C0', alpha=0.85)
ax.axvline(mean, color='k', lw=1.5, label=f'mean {mean:.3f}')
ax.axvline(λ_star*(1 - model.β), color='C1', lw=1.5, ls='--',
           label=f'flow value of search {λ_star*(1-model.β):.3f}')
ax.set_xlabel('wage')
```

(continues on next page)

(continued from previous page)

```
ax.set_ylabel('share of workers')
ax.set_xlim(0.5, 1.45)            # a thin tail runs past this; see the print below
ax.legend()
plt.tight_layout()
plt.show()

print(f"mean {mean:.3f}, "
      f"p10 {quantile(0.1):.3f}, median {quantile(0.5):.3f}, p90 {quantile(0.9):.3f}")
print(f"p90/p10 = {quantile(0.9)/quantile(0.1):.2f},  "
      f"highest wage paid {w[p > 1e-12].max():.3f}")
print(f"share paid less than the flow value of search: "
      f"{100*p[w < λ_star*(1 - model.β)].sum():.0f}%")
```

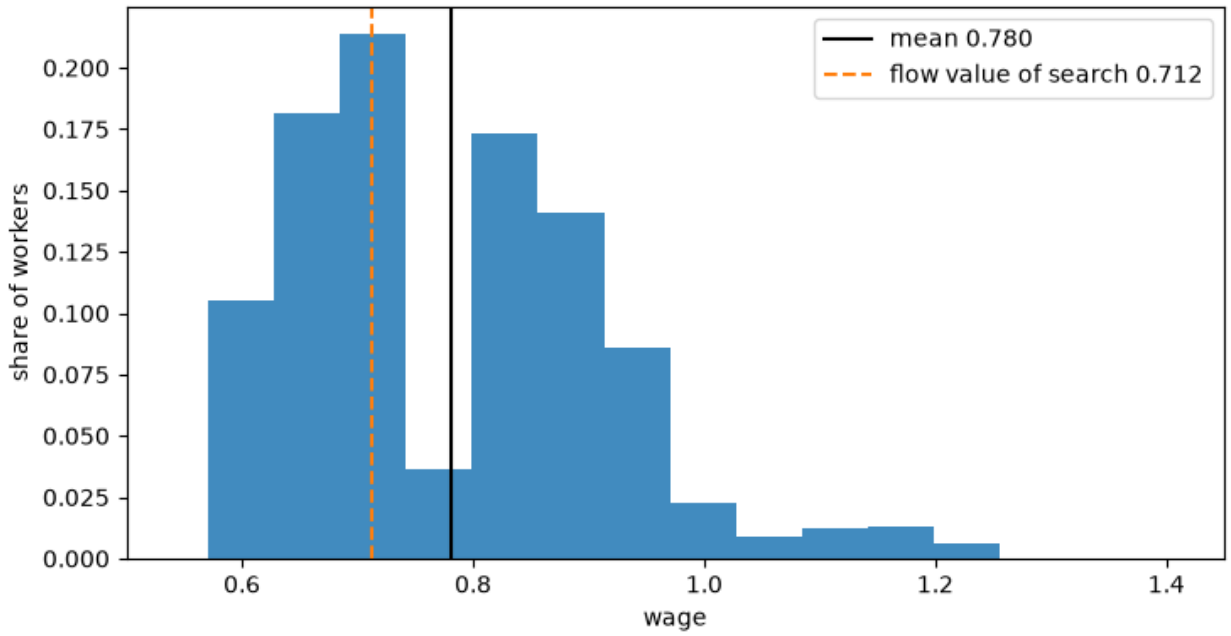


Fig. 118.4: Wages across identical workers, at the equilibrium

```
mean 0.780, p10 0.617, median 0.733, p90 0.950
p90/p10 = 1.54,  highest wage paid 1.832
share paid less than the flow value of search: 40%
```

Workers in this economy are identical.

They have the same skills, the same preferences and the same information, and they are free to leave at any time.

They are nevertheless paid differently, by half as much again at the ninetieth percentile as at the tenth, and the dispersion is not a friction or a mistake: it is what equalizes the *value* of being in one market or another when demand has moved markets apart and it takes a period to cross between them.

The dashed line is the flow value of search, $(1-\beta)\lambda^*$, which is what a searcher's position is worth per period.

Two fifths of workers are paid less than that, and stay.

There is no puzzle in it, but there is something worth seeing: what a worker holds is not his current wage.

It is a position in a particular market, which carries a claim on that market's future and the option to walk away later, and a market can be worth being in today precisely because it is depressed today and will not be for long.

Note

Notice what this model does *not* contain.

There is no matching function, no bargaining, and no wage rigidity.

Unemployment is here because moving takes a period, and in equilibrium somebody is always moving.

118.7 Relation to other lectures

Baley *et al.* [2023] open by observing that there is no standard model of frictional unemployment, but three workable classes of them: matching models in the Diamond-Mortensen-Pissarides tradition, equilibrium versions of McCall [1970], and search-island models in the tradition of Lucas and Prescott [1974].

Calibrated versions of all three fit data on labor market flows and deliver plausible responses of unemployment to policies like unemployment insurance and layoff taxes.

This lecture is the third class, and the McCall lectures in this series are the second.

It is worth being explicit about what this model supplies that the search lectures elsewhere in this series take as given.

Job Search I: The McCall Search Model and its sequels study a worker who draws from a known wage offer distribution and chooses a reservation wage.

That worker's problem is a piece of this one: here too a worker compares what a market offers him with the value of looking elsewhere.

What is added is everything on the other side of the offer.

The wage a worker is offered is the marginal product of labor in a market whose workforce was determined by where other workers decided to go, and the value of looking elsewhere, λ, is the number that clears the economy-wide market for workers.

Two consequences are worth noting.

The first is that there is no reservation wage, and the reason runs deeper than a comparison of numbers.

At the equilibrium computed above, workers stay on, and nobody leaves, in markets paying as little as 0.583, while elsewhere workers walk away from markets paying 0.674, so no cutoff between the two sorts the markets a worker accepts from those he leaves.

But the leavers are not turning down an offer of 0.674 in favor of something better, and nobody turns them away from it.

At $\bar{n}(s)$ every worker in the market, stayer or leaver, gets exactly λ.

The leavers could stay, but then employment would be the whole workforce y, the wage would fall to $R(s, y)$, and staying would be worth less than λ.

So no offer is rejected against a cutoff wage, and nobody is denied a job at the going wage.

The margin is the market's level of employment, and leaving is a choice made at a point where staying and leaving are worth the same.

The second is about the flow rates between employment and unemployment, and it is worth setting out as a ladder.

A Lake Model of Employment and Unemployment takes both rates as parameters: a job finding rate of 0.283 and a separation rate of 0.013 go in, and the unemployment rate comes out.

Lake Model with an Endogenous Job Finding Rate climbs one rung, making the job finding rate the outcome of a McCall problem, though still against a wage distribution that is given.

This lecture is the next rung: both rates are equilibrium objects.

The separation rate printed above, around 2.7 percent of the employed per period, is nothing but the frequency with which demand shifts carry markets below their employment thresholds.

Warning

The job finding rate in this model is *one*.

A worker who leaves is placed in some market at the start of the next period, always, so every spell of unemployment lasts exactly one period and the model has nothing to say about unemployment *duration*.

That is a direct consequence of directed search, and it is the first thing Alvarez and Veracierto [1999] give up: when searchers arrive uniformly, a worker can land in a market that is itself shedding workers and have to move again.

Duration becomes an equilibrium object there, and averages 2.4 periods in their calibration.

The connection to *Investment Under Uncertainty* is methodological, and close.

In the 1971 paper a competitive industry's capital stock is a Markov process whose invariant distribution describes the long run, and the equilibrium is computed by solving a fixed point problem in a function.

Here a market's workforce is a Markov process whose invariant distribution describes the long run, and the equilibrium is computed by solving a fixed point problem in a function and then one more, in the single number λ.

The same authors wrote both papers within three years, and the second is the first one's method turned on a question about labor.

118.7.1 What the island model became

Alvarez and Veracierto [1999] took this model and made it a tool for quantitative policy analysis, in their words by "modifying Lucas and Prescott's (1974) island model to incorporate undirected search and out-of-the-labor-force participation".

Four things change.

Search becomes **undirected**: searchers "arrive uniformly across all islands", so an island's workforce next period is its own employment plus the same inflow every other island gets, whether it wants the workers or not.

The nonemployed gain a second option, home production, so that labor force participation is a margin and the model has three states rather than two.

Islands produce with a decreasing returns technology subject to persistent productivity shocks rather than facing a shifting demand curve.

And agents trade a complete set of state-contingent securities, which the authors justify by their own earlier finding that the insurance role of the policies they study is quantitatively small.

With that machinery they evaluate minimum wages, two models of unions, firing taxes and unemployment insurance, at a model period of six weeks.

Note

The first of those changes is worth dwelling on, because it shows what directed search is doing here.

In this lecture λ is a **price**: raise it and every market sheds workers sooner and fills up less, so the workforce that markets want to hold falls, and (118.8) pins λ down exactly as a market-clearing condition pins down any other price.

Behind that is the min in (118.2): because searchers go only where it pays them to go, every market that receives workers is worth exactly λ to them, and λ needs no condition of its own.

Under undirected search that link is broken.

A searcher lands in a market drawn at random, so λ must equal the discounted value of landing in a randomly drawn market, and the common inflow U becomes a second unknown.

With a fixed population there are still two conditions for these two unknowns, the population and the consistency of λ, so the level of the workforce is not left free.

What is lost is the simple structure: the consistency condition need not be monotone in λ, and nothing guarantees a unique equilibrium.

Exercise 118.8.4 works through both points.

Alvarez and Veracierto add the participation margin for a different reason.

The policies they study act on whether people look for work at all, as well as on firms' hiring and firing and on workers' search, and with home production the inflow of searchers is set by the indifference between searching and staying home, so that the size of the labor force becomes an outcome.

118.7.2 Returns to labor mobility

Baley *et al.* [2025] and Baley *et al.* [2023] argue that models of this family live or die by one quantity, which they call the **returns to labor mobility**: how much more a worker produces where he is going than where he is.

It is the quantity that decides what a layoff tax does to unemployment, and whether skill losses at the time of a quit can reverse the effect of turbulence on unemployment.

And it is not a property of the shock alone.

It is a property of the shock *and* of the technology that intermediates it, which is why the two papers insist on calibrating productivity processes against evidence on firm size dynamics rather than on per-worker productivity statistics alone.

The sharpest way to see this is to compare our two-state economy of Exercise 118.8.3 with Alvarez and Veracierto [2001], whose calibration Baley *et al.* [2025] take from establishment data on job creation and destruction.

The two economies face almost the same shock: ours doubles demand between the low and the high state, theirs has a productivity ratio of 2.12.

What the shock does to employment could hardly be more different.

In our two-state economy mean employment is about **2.6 times** larger in the high state than in the low one, as Exercise 118.8.5 computes.

In theirs, Baley *et al.* [2025] compute from Alvarez and Veracierto's parameters that a low-productivity establishment employs only 3.81 percent as many workers as a high-productivity one, a ratio of about **26**, and that high-productivity establishments hold more than 96 percent of all employment.

Most of the gap is the technology.

In Alvarez and Veracierto an establishment produces $xk^{\xi}n^{\psi}$ with $\xi + \psi < 1$, and equating marginal products across establishments gives

$$\frac{n_{\text{high}}}{n_{\text{low}}} = \left(\frac{x_{\text{high}}}{x_{\text{low}}}\right)^{1/(1-\xi-\psi)}, \tag{118.9}$$

so a shock is amplified by an exponent that explodes as returns to scale approach one.

With their parameters the exponent is about 4.3, which turns a productivity ratio of 2.12 into an employment ratio of about 26.

That is a frictionless benchmark: it equates marginal products across establishments, as instantaneous and costless reallocation would.

Our marginal revenue schedule, which falls linearly to zero at a capacity, amplifies a doubling of demand about as much as a decreasing returns technology whose labor exponent is a little above one half.

Exercise 118.8.5 puts (118.9) into this model and shows what it does to unemployment.

118.8 Exercises

Exercise 118.8.1

Lucas and Prescott [1974] report two comparative statics from their own numerical examples:

> *Overall, we found for low and high persistence in demand that unemployment rates were low. In the former case, there was little gained by reallocating workers, while in the latter reallocation occurred infrequently. As expected, the greater the variability of demand, holding the degree of persistence fixed, the greater the level of unemployment.*

Check both in this calibration.

1. Compute the equilibrium unemployment rate for persistence $\rho \in \{0, 0.3, 0.6, 0.8, 0.9, 0.95, 0.99\}$, holding the *stationary* standard deviation of log s fixed at 0.26 and μ at 50. Is the relationship monotone?
2. Compute it for stationary standard deviations $\{0.10, 0.18, 0.26, 0.34, 0.45\}$, holding $\rho = 0.9$.

Explain the shape you find in part 1.

The two parameters you are varying are the same two that *Job Search VI: Persistent and Transitory Wage Shocks* varies, there as the persistence and volatility of a worker's wage offers rather than of his market's demand.

The difference is what the experiment delivers: a reservation wage there, an unemployment rate here.

Solution

Holding the stationary standard deviation fixed as ρ varies is what the argument `sd` of `create_model` already does, so each experiment is one call.

```
print(f"{'ρ':>8}{'λ*':>9}{'workforce':>12}{'u rate':>10}")
u_by_ρ = []
for ρ in (0.0, 0.3, 0.6, 0.8, 0.9, 0.95, 0.99):
    m_ρ = create_model(ρ=ρ)
    λ_ρ = equilibrium_λ(m_ρ, μ)
    r = market_outcomes(m_ρ, λ_ρ)
    u_by_ρ.append(100*r['unemployment']/r['workforce'])
    print(f"{ρ:>8}{λ_ρ:>9.3f}{r['workforce']:>12.1f}{u_by_ρ[-1]:>9.2f}%")
```

```
       ρ       λ*   workforce    u rate
     0.0    7.247        49.7     0.00%
     0.3    7.153        50.4     0.02%
     0.6    7.230        50.2     1.21%
     0.8    7.186        50.1     2.54%
     0.9    7.121        50.0     2.68%
    0.95    7.108        50.1     1.60%
    0.99    7.065        50.0     0.01%
```

The second experiment varies how far demand moves, holding how long it stays.

```
print(f"{'sd':>8}{'λ*':>9}{'workforce':>12}{'u rate':>10}")
for sd in (0.10, 0.18, 0.26, 0.34, 0.45):
    m_sd = create_model(sd=sd)
    λ_sd = equilibrium_λ(m_sd, μ)
    r = market_outcomes(m_sd, λ_sd)
    print(f"{sd:>8}{λ_sd:>9.3f}{r['workforce']:>12.1f}"
          f"{100*r['unemployment']/r['workforce']:>9.2f}%")
```

```
      sd       λ*   workforce    u rate
     0.1    7.092        49.9     0.37%
    0.18    7.099        50.1     1.45%
    0.26    7.121        50.0     2.68%
    0.34    7.292        49.9     3.40%
    0.45    7.537        50.0     4.97%
```

Variability does what one expects: demand that moves further moves more workers, and the unemployment rate rises steadily with it.

Persistence does not.

The unemployment rate is near zero at both ends and reaches its maximum in between, which is the shape Lucas and Prescott describe.

The two ends are near zero for opposite reasons.

When demand is nearly independent across periods, a market that is unlucky today expects to be average tomorrow, so its workers see little point in moving: the thresholds $\bar{n}$ and $\bar{y}$ are close together and almost nobody is ever shed.

When demand is nearly permanent, the thresholds are far apart, but a market almost never crosses from one to the other, so the reallocation that would make workers move hardly ever happens.

Unemployment needs demand shifts that are *both* worth responding to and frequent, which is a combination that only intermediate persistence delivers.

Exercise 118.8.2

In their concluding section, Lucas and Prescott consider a lump-sum cost c of leaving a market to search, so that the right side of (118.3) becomes $\lambda - c$ rather than λ: a worker who leaves collects the value of search net of what it costs him to get there.

The arrival condition (118.5) is unchanged, since it prices the market for workers who are already searching.

They report that the cost lowers unemployment, and remark that this

> *shows that lower average unemployment is not, in general, associated with higher welfare for workers.*

Modify the code to allow a cost c, and for $c \in \{0, 0.2, 0.5, 1, 2\}$ report the equilibrium λ^*, the unemployment rate, the expected value of a worker under the invariant distribution, which weights each market by the workers in it,

$$\frac{\sum_{s,y} \psi(s,y)\, y\, v(s,y)}{\sum_{s,y} \psi(s,y)\, y},$$

and output per market,

$$\sum_{s,y} \psi(s,y) \int_0^{n(s,y)} R(s,m)\, dm,$$

which is the area under the marginal revenue schedule up to employment.

Does lower unemployment mean better off?

Solution

Only two lines change: the outside option inside the max of (118.2), and the threshold that defines $\bar{n}$.

```
def worker_value_cost(model, λ, c, tol=1e-11, max_iter=10_000):
    "The worker's problem when leaving costs c."
    P, R, β = model.P, model.R, model.β
    v = np.full(R.shape, λ - c)
    for _ in range(max_iter):
        v_new = np.maximum(λ - c, R + np.minimum(λ, β*(P @ v)))
        v_new[:, 0] = λ - c
        if np.max(np.abs(v_new - v)) < tol:
            return v_new
        v = v_new
    return v

def outcomes_cost(model, λ, c):
    v = worker_value_cost(model, λ, c)
    P, R, β, y = model.P, model.R, model.β, model.y
    EV = β*(P @ v)
    H = R + EV
    last = lambda A, bar: np.array([np.max(np.where(A[i] >= bar)[0])
                                    if (A[i, 1:] >= bar).any() else 0
                                    for i in range(len(model.s))])
    n_bar = last(H, λ - c)       # staying is worth λ - c to the leaver
    y_bar = last(EV, λ)          # arrivals are priced at λ; still ≤ n_bar
    ψ = stationary_distribution(model, n_bar, y_bar)
    employment = np.minimum(y[None, :], n_bar[:, None])
    output = np.take_along_axis(np.cumsum(R, axis=1), employment, axis=1)
    return dict(workforce=(ψ*y).sum(),
                unemployment=(ψ*(y[None, :] - employment)).sum(),
                welfare=(ψ*y*v).sum()/(ψ*y).sum(),    # per worker, not per market
                output=(ψ*output).sum())

def equilibrium_λ_cost(model, μ, c, lo=0.3, hi=40.0, tol=1e-4):
    while hi - lo > tol:
        mid = 0.5*(lo + hi)
        if outcomes_cost(model, mid, c)['workforce'] > μ:
            lo = mid
        else:
            hi = mid
    return 0.5*(lo + hi)
```

Now sweep the cost.

```
print(f"{'c':>6}{'λ*':>9}{'λ*-c':>9}{'u rate':>9}{'E[v]':>9}{'output':>9}")
for c in (0.0, 0.2, 0.5, 1.0, 2.0):
    λ_c = equilibrium_λ_cost(model, μ, c)
    r = outcomes_cost(model, λ_c, c)
    print(f"{c:>6}{λ_c:>9.3f}{λ_c - c:>9.3f}"
          f"{100*r['unemployment']/r['workforce']:>8.2f}%{r['welfare']:>9.3f}"
          f"{r['output']:>9.3f}")
```

```
     c       λ*     λ*-c   u rate     E[v]   output
  0.0    7.121    7.121    2.68%    7.575   48.100
  0.2    7.227    7.027    2.24%    7.582   48.055
  0.5    7.403    6.903    1.58%    7.594   47.758
  1.0    7.681    6.681    1.14%    7.632   47.363
  2.0    8.290    6.290    0.50%    7.698   46.592
```

A cost of two cuts measured unemployment by more than four fifths, from 2.7 percent to 0.5 percent.

It also cuts output, by about three percent.

That is the paper's point, and it is worth stating plainly: the unemployment that the cost eliminated was *productive*.

Those workers were not idle in any useful sense; they were in transit from markets where their marginal product had fallen to markets where it was high, and the cost keeps them where they are.

The rest of the table shows why "welfare" needs saying carefully.

The value of the option a worker holds, $\lambda^* - c$, falls at every step: the outside option is worth less, which is what makes workers stay.

Yet the expected present value of a worker's own earnings *rises*.

A worker's income is the wage when he is employed and nothing when he is searching, and the cost buys him fewer periods of nothing.

What he gains that way, the economy loses more than twice over in output, because the match between workers and markets is worse.

So unemployment is not a welfare index, which is the lesson Lucas and Prescott draw, and in this calibration even the two welfare measures move in opposite directions.

It would be fair to suspect that so stark a result is an artifact of a model this stripped down.

It is not.

Alvarez and Veracierto [1999] run the same experiment as a firing tax inside a calibrated version of this economy, with taxes worth between three and twelve months of average wages.

At twelve months the tax takes the unemployment rate *down* from 5.3 to 3.7 percent and pushes the average duration of unemployment *up* from 2.4 to 5.1 periods.

When home and market goods are perfect substitutes it also cuts employment to 86 percent of its laissez-faire level, because so many fewer people search; with less substitutability employment falls much less, and with little enough it rises.

The number you have just computed has a use beyond the comparison.

Baley *et al.* [2023] make a diagnostic out of a closely related number: the minimum layoff tax that shuts down voluntary separations altogether.

Here even a cost of 2, which is about 2.6 times a period's average wage, does not do that: it takes unemployment to a fifth of its laissez-faire rate, not to zero.

Their point is that this number has to be *large* to be credible, because extensive reallocation is observed in economies whose layoff costs differ enormously; a model in which a modest cost stops workers moving is a model whose returns to labor mobility are too low.

Exercise 118.8.5 asks what sets them.

Exercise 118.8.3

The numerical examples in the appendix of Lucas and Prescott [1974] use just two demand states, with

$$P = \begin{bmatrix} 0.9 & 0.1 \\ 0.1 & 0.9 \end{bmatrix}, \qquad s \in \{1, 2\}, \qquad \beta = 0.9.$$

Build that economy, keeping $R(s, n) = s\,(1 - n/N)^+$ from the lecture, with capacity $N = 200$.

1. For λ on a grid, plot workforce demand and average unemployment against λ, as in the bottom panels of the paper's Figure 5.
2. Show that the invariant distribution puts all of its mass on *two* values of the workforce, and say which two.
3. Use that fact to derive a formula for average unemployment per market in terms of $\bar{n}$, $\bar{y}$ and the persistence ρ, and check it against the computed value.

Solution

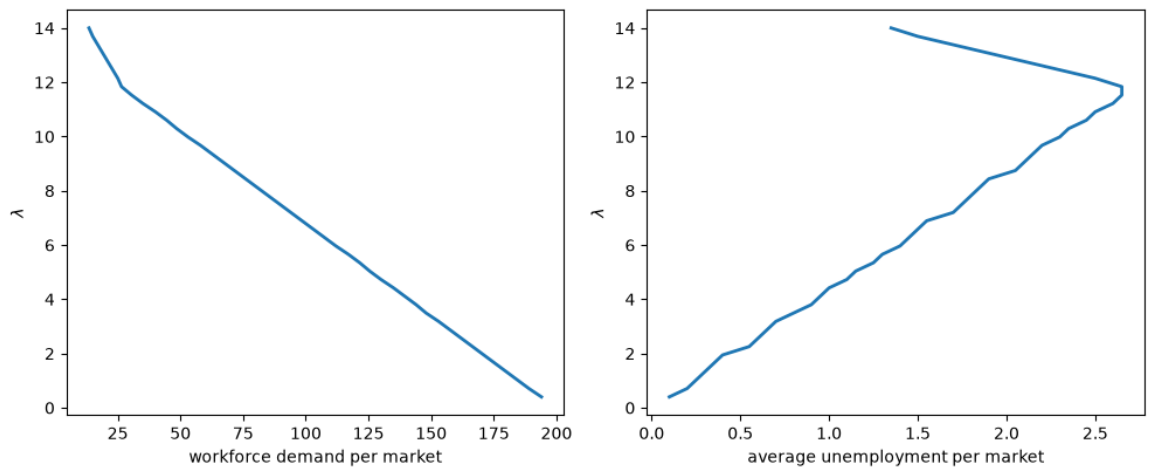
The two-state chain needs its own constructor, but everything else is reused.

```
def two_state_model(ρ=0.9, s_lo=1.0, s_hi=2.0, β=0.9, N=200, y_max=260):
    P = np.array([[ρ, 1-ρ], [1-ρ, ρ]])
    s = np.array([s_lo, s_hi])
    y = np.arange(y_max + 1)
    R = np.maximum(s[:, None]*(1 - y[None, :]/N), 0.0)
    R[:, 0] = 0.0
    return Model(P=P, s=s, y=y, R=R, β=β)

two = two_state_model()
grid2 = np.linspace(0.4, 14.0, 45)
res2 = [market_outcomes(two, λ) for λ in grid2]
Ey2 = np.array([r['workforce'] for r in res2])
Eu2 = np.array([r['unemployment'] for r in res2])
```

The paper puts λ on the vertical axis.

```
fig, axes = plt.subplots(1, 2, figsize=(10, 4.2))
axes[0].plot(Ey2, grid2, lw=2)
axes[0].set_xlabel('workforce demand per market')
axes[1].plot(Eu2, grid2, lw=2)
axes[1].set_xlabel('average unemployment per market')
for ax in axes:
    ax.set_ylabel(r'$\lambda$')
plt.tight_layout()
plt.show()
```



The right panel has the shape of the bottom panels of the paper's Figure 5: average unemployment is largest at intermediate values of λ.

For part 2, look at the support of the invariant distribution.

```
r2 = market_outcomes(two, 6.0)
support = np.where(r2['ψ'].sum(axis=0) > 1e-12)[0]
print("workforce values with positive probability:", support)
print("n̄(s) =", r2['n_bar'], "   ȳ(s) =", r2['y_bar'])
```

```
workforce values with positive probability: [ 98 125]
n̄(s) = [ 98 140]    ȳ(s) = [ 78 125]
```

The two values are $\bar{n}(s_{\text{low}})$ and $\bar{y}(s_{\text{high}})$.

To see why, follow (118.6).

In the low state the workforce is clipped into $[\bar{y}(s_{\text{low}}), \bar{n}(s_{\text{low}})]$ and in the high state into $[\bar{y}(s_{\text{high}}), \bar{n}(s_{\text{high}})]$, and with these parameters the second interval lies entirely above the first.

So one period in the low state carries any workforce to $\bar{n}(s_{\text{low}})$, one period in the high state carries it to $\bar{y}(s_{\text{high}})$, and the market shuttles between the two.

For part 3, unemployment is positive only in the first period of a downturn: the market enters with $\bar{y}(s_{\text{high}})$ workers, demand has just fallen, and employment is $\bar{n}(s_{\text{low}})$.

That event is "the state was high last period and is low now", which has stationary probability $\frac{1}{2}(1-\rho)$, so

$$E[u] = \tfrac{1}{2}(1-\rho)\,[\bar{y}(s_{\text{high}}) - \bar{n}(s_{\text{low}})].$$

```
predicted = 0.5*(1 - 0.9)*(r2['y_bar'][1] - r2['n_bar'][0])
print(f"formula:  {predicted:.4f}")
print(f"computed: {r2['unemployment']:.4f}")
```

```
formula:  1.3500
computed: 1.3500
```

The formula also explains the hump.

Average unemployment is the gap between two thresholds, multiplied by the chance of crossing it.

As λ falls, both thresholds rise, but not in step, and the gap between them widens and then narrows again.

Exercise 118.8.4

This exercise asks what the model's directed search is doing.

Alvarez and Veracierto [1999] replace it with **undirected** search: a searcher cannot choose where to go, and every market receives the same inflow U, which is the economy's unemployment per market.

1. Show that under undirected search the min disappears from (118.2), and say which economic condition it was enforcing.
2. Write down the transition law for a market's workforce, and show that at any stationary distribution the inflow U equals the average outflow. Does the same identity hold under directed search?
3. Under directed search, (118.8) pins down λ. Under undirected search, what condition must λ satisfy instead? Count the unknowns and the conditions of an equilibrium with a fixed workforce μ per market.
4. Take the lecture's seven-state calibration and an inflow of $U = 1$, and then of $U = 3$. For λ on a grid from 3 to 10, compute the gap between the value of search implied by λ and λ itself, and find the values of λ at which it vanishes. What does this say about equilibrium under undirected search?

Solution

For part 1, the min in (118.2) is not a choice that anyone makes.

It is the restriction that the continuation value in a market cannot exceed λ, which competition among searchers imposes: if some market promised more, searchers would go there until it no longer did.

Take away their ability to choose where to go and nothing enforces it.

A worker who stays simply collects the wage and the continuation value his own market happens to deliver,

$$v(s, y) = \max\left\{\lambda,\ R(s, y) + \beta \int v(s', y'(s, y)) f(s', s)\, ds'\right\},$$

and the continuation can be worth more or less than λ, because nobody can act on the difference.

For part 2, employment is still $n(s, y) = \min\{\bar{n}(s), y\}$, those not employed leave, and the inflow is the same everywhere, so

$$y' = n(s, y) + U.$$

At a stationary distribution $E[y'] = E[y]$, so averaging across markets gives

$$U = E[y] - E[n(s, y)] = E[\, y - n(s, y)\,] :$$

the inflow equals the average outflow.

That is an accounting identity, and it holds under directed search too.

At the stationary distribution of the lecture's own transition law (118.6), the workers arriving in markets that fill up equal the workers leaving markets that shed, at every value of λ and not only at λ^*.

```
y = model.y
print(f"{'λ':>8}{'E[y]':>9}{'outflow':>10}{'inflow':>9}")
for λ in (4.0, λ_star, 10.0):
    r = market_outcomes(model, λ)
    y_next = np.clip(y[None, :], r['y_bar'][:, None], r['n_bar'][:, None])
    outflow = (r['ψ']*np.maximum(y[None, :] - y_next, 0)).sum()
    inflow = (r['ψ']*np.maximum(y_next - y[None, :], 0)).sum()
    print(f"{λ:>8.3f}{r['workforce']:>9.2f}{outflow:>10.4f}{inflow:>9.4f}")
```

```
       λ      E[y]   outflow   inflow
   4.000   114.76    0.8259   0.8259
   7.121    50.00    1.3377   1.3377
  10.000    12.00    0.6825   0.6825
```

So the identity cannot be what distinguishes the two models.

What distinguishes them is how λ is determined.

For part 3, a searcher under undirected search lands in a market drawn at random from the invariant distribution, so the value of search must satisfy

$$\lambda = \beta \sum_{s,y} \psi(s,y)\, v(s,y),$$

the discounted value of landing in a randomly drawn market.

An equilibrium with a fixed population therefore has two unknowns, λ and U, and two conditions, $\mu = \sum_{s,y} \psi(s,y)\, y$ and the condition above.

The level of the workforce is not left free: a given μ selects U, and the consistency condition selects λ.

The difference from directed search lies in how λ is pinned.

There the min makes every market that receives workers worth exactly λ, so λ needs no separate consistency condition, and μ pins it down alone, through a schedule that slopes down.

Here λ needs a condition of its own, and nothing makes that condition monotone in λ.

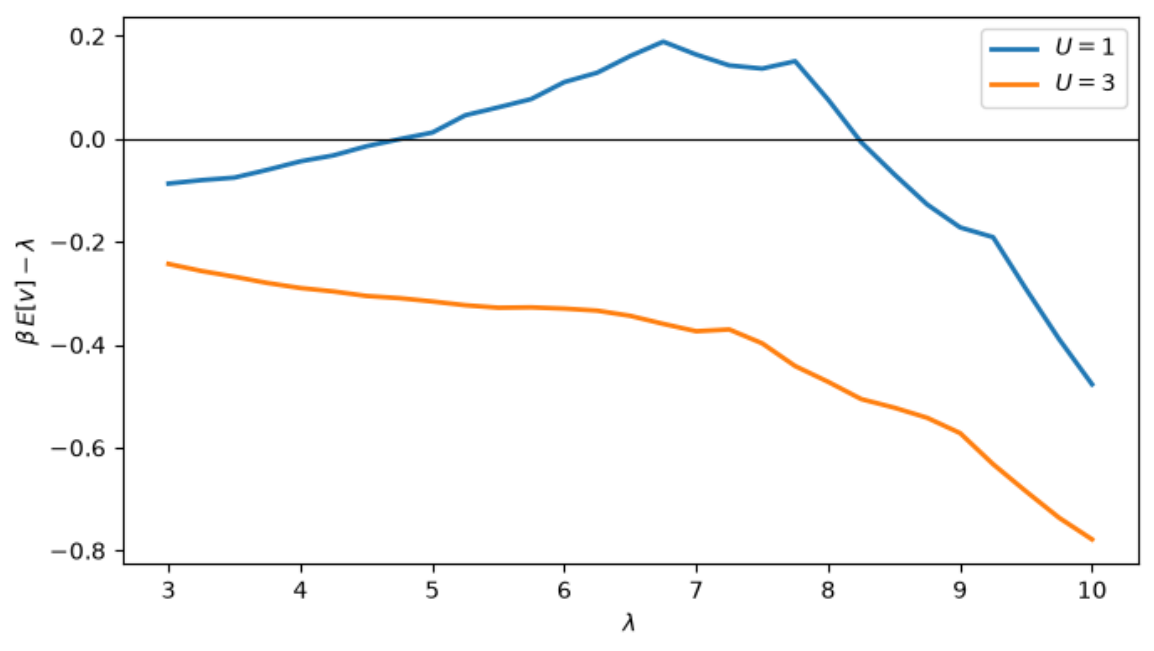
For part 4, a market under undirected search keeps the Bellman equation of part 1, with next period's workforce $n(s,y) + U$, and its employment threshold is the largest workforce at which staying is still worth λ.

```
def undirected_market(model, λ, U, tol=1e-10, max_iter=20_000):
    "Every market receives U searchers each period, whether it wants them or not."
    P, R, β, y = model.P, model.R, model.β, model.y
    shift = np.minimum(y + U, y[-1])      # next workforce of a market that keeps␣
↪everyone
    v = np.full(R.shape, λ)
    for _ in range(max_iter):
        v_new = np.maximum(λ, R + β*(P @ v)[:, shift])
        if np.max(np.abs(v_new - v)) < tol:
            break
        v = v_new
    v = v_new
    stay = R + β*(P @ v)[:, shift]
    n_bar = np.array([np.max(np.where(stay[i] >= λ)[0]) if (stay[i, 1:] >= λ).any()
                      else 0 for i in range(len(model.s))])
    # the invariant distribution of y' = min(y, n̄(s)) + U
    y_next = np.minimum(np.minimum(y[None, :], n_bar[:, None]) + U, y[-1])
    ψ = np.ones(R.shape)/R.size
    for _ in range(100_000):
        moved = np.array([np.bincount(y_next[i], weights=ψ[i], minlength=len(y))
                          for i in range(len(model.s))])
        ψ_new = P.T @ moved
        if np.max(np.abs(ψ_new - ψ)) < 1e-13:
            break
        ψ = ψ_new
    ψ = ψ_new
    return β*(ψ*v).sum(), (ψ*y).sum()      # the implied value of search, mean␣
↪workforce
```

Now sweep λ for each inflow and look for the values at which the implied value of search equals λ.

```
λ_grid = np.arange(3.0, 10.01, 0.25)
fig, ax = plt.subplots(figsize=(7, 4))
for U in (1, 3):
    implied, mean_y = np.array([undirected_market(model, λ, U) for λ in λ_grid]).T
    gap = implied - λ_grid
    ax.plot(λ_grid, gap, lw=2, label=f'$U = {U}$')
    crossings = np.where(np.diff(np.sign(gap)) != 0)[0]
    print(f"U = {U}: {len(crossings)} solution(s)")
    for i in crossings:
        w = gap[i]/(gap[i] - gap[i+1])                    # linear interpolation
        print(f"    λ ≈ {λ_grid[i] + w*(λ_grid[i+1] - λ_grid[i]):.2f},"
              f"  mean workforce ≈ {mean_y[i] + w*(mean_y[i+1] - mean_y[i]):.0f}")
ax.axhline(0, color='k', lw=0.8)
ax.set_xlabel(r'$\lambda$')
ax.set_ylabel(r'$\beta\, E[v] - \lambda$')
ax.legend()
plt.tight_layout()
plt.show()
```

```
U = 1: 2 solution(s)
    λ ≈ 4.76,  mean workforce ≈ 97
    λ ≈ 8.23,  mean workforce ≈ 26
U = 3: 0 solution(s)
```



With $U = 1$ the gap is negative at both ends of the grid and positive in between, so the consistency condition has two solutions, one with a crowded economy and a low value of search and one with a sparse economy and a high one.

With $U = 3$ the gap is negative everywhere on the grid, and there is no solution at all.

So under undirected search the level of the workforce is still pinned down, but neither existence nor uniqueness comes for free, as they do under directed search, where Proposition 5 of Lucas and Prescott [1974] delivers a unique λ^* because $E(y; \lambda)$ is decreasing.

The shape of the gap has a reading.

A searcher's prospects come from markets whose demand has just improved and that a trickle of U arrivals per period has not yet filled.

When λ is low, the employment thresholds of all demand states are crowded together near capacity, so a change in demand opens little room; when λ is high, many markets close, and a searcher who lands in one has to move again.

The rents that make searching worthwhile are largest in between.

What Alvarez and Veracierto add is the **participation margin**: the nonemployed choose between searching and home production.

That replaces the fixed population with an indifference condition, the inflow of searchers adjusts until searching and staying home are equally good, and the size of the labor force becomes an outcome.

Their own reason for adding it is the policies they study, which act on participation as well as on firms' employment decisions and on workers' search.

Exercise 118.8.5

How far employment moves when demand moves is not a property of the demand process alone.

It also depends on the technology through which the shock reaches the labor market, and Baley *et al.* [2025] argue that this is the quantity a calibration most needs to get right.

Replace the marginal revenue schedule of this lecture with the one implied by a decreasing returns production function.

If a market produces $s\,n^a$ with $0 < a < 1$, the marginal product of labor is

$$R(s, n) = s\,a\,(n + n_0)^{a-1},$$

where the shift n_0 keeps it bounded at $n = 0$.

1. Build the two-state economy of Exercise 118.8.3 with this schedule, for $a \in \{0.4, 0.6, 0.75\}$, and solve each for its own equilibrium. Report mean employment in each demand state, the ratio of the two, and the unemployment rate.
2. Compare with (118.9). Equating marginal products $s\,a\,(n + n_0)^{a-1}$ across the two states gives $(n_{\text{high}} + n_0)/(n_{\text{low}} + n_0) = 2^{1/(1-a)}$, so for these economies the formula applies to $n + n_0$ rather than to n.
3. What happens to unemployment as the returns to labor mobility rise, and why?

Solution

Only the construction of `R` changes, and we add a function that averages employment within each demand state.

```
def two_state_dr(a, ρ=0.9, s_lo=1.0, s_hi=2.0, β=0.9, n0=20.0, y_max=3000):
    "Two demand states, output s n^a, so R = s a (n + n0)^(a-1)."
    P = np.array([[ρ, 1-ρ], [1-ρ, ρ]])
    s = np.array([s_lo, s_hi])
    y = np.arange(y_max + 1)
    R = s[:, None]*a*(y[None, :] + n0)**(a - 1)
    R[:, 0] = 0.0
    return Model(P=P, s=s, y=y, R=R, β=β)

def employment_by_state(m, r):
    "Mean employment in each demand state under the invariant distribution."
    emp = np.minimum(m.y[None, :], r['n_bar'][:, None])
```

```
    return (r['ψ']*emp).sum(axis=1) / r['ψ'].sum(axis=1)

def report(name, m, r, a=None, n0=None):
    "One row of the table; a and n0 are given only for the decreasing returns␣
 ↪schedules."
    En = employment_by_state(m, r)
    shifted = '--' if n0 is None else f"{(En[1] + n0)/(En[0] + n0):.2f}"
    formula = '--' if a is None else f"{2**(1/(1 - a)):.2f}"
    print(f"{name:>18}{En[0]:>9.1f}{En[1]:>9.1f}{En[1]/En[0]:>8.2f}{shifted:>10}"
          f"{formula:>10}{100*r['unemployment']/r['workforce']:>8.2f}%")

print(f"{'technology':>18}{'E[n|lo]':>9}{'E[n|hi]':>9}{'ratio':>8}{'+n0 ratio':>10}
 ↪"
      f"{'formula':>10}{'u rate':>9}")
report('this lecture', two, market_outcomes(two, equilibrium_λ(two, μ)))
for a in (0.4, 0.6, 0.75):
    m_a = two_state_dr(a)
    r_a = market_outcomes(m_a, equilibrium_λ(m_a, μ, lo=0.01, hi=80.0))
    report(f'output s n^{a}', m_a, r_a, a=a, n0=20.0)
```

```
        technology  E[n|lo]  E[n|hi]   ratio +n0 ratio    formula   u rate
      this lecture     26.0     68.3    2.63        --         --    4.75%
    output s n^0.4     32.0     64.4    2.01      1.62       3.17    3.60%
    output s n^0.6     23.0     71.6    3.11      2.13       5.66    5.40%
   output s n^0.75     10.0     81.1    8.11      3.37      16.00    7.98%
```

The same doubling of demand moves mean employment by a factor of about 2.6 under this lecture's schedule and by a factor of about 8 when output is $s\,n^{0.75}$.

Mean employment is not the employment threshold $\bar{n}$.

In this two-state economy a market in the high state never reaches $\bar{n}(s_{\text{high}})$: arrivals fill it only to $\bar{y}(s_{\text{high}})$, and in the first period of an upturn it still holds the $\bar{n}(s_{\text{low}})$ workers it carried out of the low state, so mean employment in the high state is $\rho\,\bar{y}(s_{\text{high}}) + (1-\rho)\,\bar{n}(s_{\text{low}})$.

For part 2, the formula is a frictionless benchmark, and the model responds less than it in every case, by a margin that grows with a.

That is what a one-period moving friction should do.

A frictionless economy would move workers until marginal products were equal across markets; here workers move only until arriving is worth λ, arrivals stop at $\bar{y}(s_{\text{high}})$ because demand may fall back, and an upturn takes a period to be met at all.

Ratios of n alone are a poor guide in these economies, because when low-state employment is small the shift n_0 is large relative to it.

```
a = 0.75
for n0 in (10.0, 20.0, 40.0):
    m_a = two_state_dr(a, n0=n0)
    En = employment_by_state(m_a, market_outcomes(m_a, equilibrium_λ(m_a, μ, lo=0.
 ↪01, hi=80.0)))
    ratio = En[1]/En[0] if En[0] > 0 else np.inf
    print(f"n0 = {n0:4.0f}:  E[n|lo] = {En[0]:5.1f},  ratio of n = {ratio:6.2f},"
          f"  ratio of n + n0 = {(En[1] + n0)/(En[0] + n0):.2f}")
```

```
n0 =   10:  E[n|lo] =  16.0,  ratio of n =   4.77,  ratio of n + n0 = 3.32
n0 =   20:  E[n|lo] =  10.0,  ratio of n =   8.11,  ratio of n + n0 = 3.37
n0 =   40:  E[n|lo] =   0.0,  ratio of n =    inf,  ratio of n + n0 = 3.25
```

The ratio of n swings from finite to infinite as n_0 rises, while the ratio of $n + n_0$, which is what the technology governs, barely moves.

The ratio of about 26 that Baley *et al.* [2025] compute for Alvarez and Veracierto's parameters comes from the frictionless formula with an exponent of about 4.3, larger than any exponent here, and a slightly bigger shock: it measures what their technology would do to a shock of that size if reallocation were instantaneous.

As for part 3, unemployment rises with the returns to labor mobility, from 3.6 to 8.0 percent across these economies, and it does so for other values of the shift n_0 as well.

That is not a coincidence and it is not a cost.

Unemployment in this model is the time workers spend moving, and the more there is to gain by moving, the more of them move.

This is the sense in which the quantity Baley *et al.* [2025] are after governs everything else: an economy with high returns to labor mobility will show a lot of reallocation, will keep reallocating even when a layoff tax makes it expensive, and will record the unemployment that goes with it.

CHAPTER

NINE

OPTIMAL GROWTH UNDER UNCERTAINTY AND TOBIN'S Q

Contents

- *Optimal Growth Under Uncertainty and Tobin's q*
  - *Overview*
  - *The Brock-Mirman planning problem*
  - *Long-run behavior*
  - *Irreversible investment and Tobin's q*
  - *Differentiability of the value function*
  - *Computing the model*
  - *Exercises*

119.1 Overview

This lecture studies two papers.

The first is Brock and Mirman [1972], a classic analysis of optimal one-sector growth when the production function is hit by random shocks.

The second is Sargent [1980], which adds one small ingredient – **irreversible investment** – and uses the result to think about James Tobin's q theory of investment.

The lecture *Investment Under Uncertainty* studied a competitive industry whose equilibrium is the solution of a planning problem, and whose equilibrium is a Markov process with an invariant distribution.

Brock and Mirman ask the same questions about a one-sector growth model:

- does the planning problem have a well-behaved solution, with continuous and monotone policy functions?
- the optimal policy makes capital a Markov process; does that process have an *invariant distribution*?
- does the distribution of capital converge to it from *any* initial condition?
- do *time averages* computed from a single realization converge to population moments?

The last two questions are the reason both papers matter for econometrics.

A model that answers them affirmatively delivers a stationary, ergodic stochastic process for observable time series.

Sample moments computed from one long realization then estimate population moments, which is exactly what the rational expectations econometrics of Hansen and Sargent [1980] requires.

Sargent [1980] puts the Brock-Mirman apparatus to work on a question in macroeconomics: when, and in what sense, is there an investment demand schedule relating investment to Tobin's q?

Along the way, this lecture pauses over a technical issue that turns out to be central:

- the derivative of the planner's value function with respect to capital *is* the competitive price of used capital, i.e. it is essentially Tobin's q
- so the analysis stands or falls on whether that derivative exists
- irreversible investment creates a **corner**, and at a corner the standard argument for differentiability breaks down

We'll describe the difficulty, note a gap in the published argument, and give a corrected proof.

Let's start with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from collections import namedtuple
```

119.2 The Brock-Mirman planning problem

A planner chooses consumption to maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t), \qquad 0 < \beta < 1, \tag{119.1}$$

subject to

$$c_t + i_t \leq f(k_t, r_t), \qquad c_t \geq 0, \quad i_t \geq 0,$$

where k_t is capital per worker, i_t is investment, and $\{r_t\}$ is a sequence of independent and identically distributed shocks to production.

Brock and Mirman assume that for each realization of the shock, $f(\cdot, r)$ is increasing, strictly concave, and satisfies the Inada conditions

$$f(0, r) = 0, \qquad f'(0, r) = +\infty, \qquad f'(\infty, r) = 0,$$

and that f is increasing in r, so that a higher shock means more output.

The period utility function u is increasing, strictly concave, and differentiable.

In their formulation capital depreciates fully in one period, so that $k_{t+1} = i_t$ and current resources are $s_t = f(k_{t-1}, r_{t-1})$.

The Bellman equation is

$$v(s) = \max_{0 \leq c \leq s} \left\{ u(c) + \beta \int v(f(s-c, r))\, \nu(dr) \right\}, \tag{119.2}$$

where ν is the distribution of the shock.

Standard arguments – a contraction mapping, plus preservation of monotonicity and concavity – give a unique bounded solution v, attained by a unique consumption policy $c = g(s)$ and investment policy $x = h(s) = s - g(s)$.

Proposition 119.2.1

The optimal policies g and h are increasing and continuous, with $g(0) = h(0) = 0$.

Brock and Mirman prove monotonicity from the first-order condition

$$u'(g(s)) = \beta \int u'(g(f(s - g(s), r))) \, f'(s - g(s), r) \, \nu(dr), \tag{119.3}$$

which is the stochastic analogue of the Euler equation.

If g did not increase with s, the left side of (119.3) would stay constant while the right side would fall, because f increases and u' and f' decrease.

119.3 Long-run behavior

The optimal policy makes capital a Markov process,

$$k_{t+1} = h(f(k_t, r_t)), \tag{119.4}$$

and Brock and Mirman devote most of their paper to the long-run behavior of this process.

The strategy resembles the one used in *Investment Under Uncertainty*.

Consider the deterministic difference equations obtained by fixing the shock at its lowest possible value $\underline{r}$ and at its highest possible value $\bar{r}$:

$$k_{t+1} = h(f(k_t, \underline{r})), \qquad k_{t+1} = h(f(k_t, \bar{r})).$$

Let $\underline{k}$ be a positive fixed point of the first and $\bar{k}$ a positive fixed point of the second.

Because f is increasing in the shock and h is increasing, $\underline{k} \leq \bar{k}$.

Brock and Mirman say that the model has a **stable configuration of fixed points** when $\underline{k}$ and $\bar{k}$ are positive and $\underline{k} < \bar{k}$.

They then show that

- the sets of capital stocks below $\underline{k}$ and above $\bar{k}$ are **transient**: the process leaves them and does not return
- the interval $[\underline{k}, \bar{k}]$ is where the process eventually lives

Theorem 119.3.1

There is a distribution function F such that the distribution F_t of k_t converges to F uniformly, and F does not depend on the initial capital stock.

Theorem 119.3.1 is the counterpart of the invariant-distribution theorem of *Investment Under Uncertainty*.

An important part of the proof is an argument about the **inverse optimal process**, which runs the Markov process backwards in time; Brock and Mirman use it to show that a stationary distribution exists and is unique without appealing to heavier machinery from probability theory.

Convergence of distributions is not by itself enough for econometrics.

What an econometrician needs in addition is a **mean-ergodic theorem**: a guarantee that averages computed along a single realization converge to moments of F,

$$\frac{1}{T} \sum_{t=1}^{T} \phi(k_t) \to \int \phi \, dF \qquad \text{with probability one.} \tag{119.5}$$

For Markov processes with a unique invariant distribution and a stable configuration of fixed points, such laws of large numbers are available; Sargent [1980] invokes the version in Doob's book to justify computing population moments of his model and comparing them with sample moments.

This is the bridge from Brock and Mirman's theory to the rational expectations econometrics developed in the late 1970s and early 1980s.

Without *Theorem 119.3.1* and (119.5), a likelihood function computed from a single time series would have no firm justification.

Note

Brock and Mirman [1972] study the case in which the shocks are independent and identically distributed.

The lecture *Investment Under Uncertainty* describes a parallel set of results for a model in which the shock is serially correlated.

In both cases the economics is the same: the state must eventually wander inside a compact set on which the process is well behaved, and the boundaries of that set are the stationary points associated with the most and least favorable shocks.

119.4 Irreversible investment and Tobin's q

Sargent [1980] modifies the model in one respect: capital, once installed, cannot be eaten.

Output can be consumed or added to the capital stock, but the conversion does not run backwards,

$$K_{t+1} = (1-\delta)K_t + I_t, \qquad I_t \geq 0, \tag{119.6}$$

where $0 < \delta < 1$ is the depreciation rate.

Production is $y_t = f(K_t)\theta_t$, and the one-period utility function $u(c_t, e_t)$ is hit by a preference shock e_t.

The shocks (θ_t, e_t) are independent and identically distributed over time.

Why does this small change matter?

In the model with reversible investment, the price of a unit of installed capital always equals the price of a unit of newly produced output, so Tobin's q is identically one and firms have no investment demand schedule at all.

The irreversibility constraint is a **friction** that lets the price of installed capital fall below the price of new capital.

Sargent describes a competitive economy in which households own capital, rent it to firms, and trade claims to installed capital at a relative price p_{Kt}, which is Tobin's q.

Following Lucas and Prescott [1971] – exactly the strategy of *Investment Under Uncertainty* – he studies the competitive equilibrium indirectly, through the planning problem that generates it.

The planner solves

$$v(K,\theta,e) = \max_{K' \geq (1-\delta)K} \left\{ u(f(K)\theta + (1-\delta)K - K', e) + \beta \mathbb{E}\left[v(K', \theta', e') \right] \right\}. \tag{119.7}$$

Let $\lambda \geq 0$ be the multiplier on the irreversibility constraint.

The first-order condition is

$$u_c(c,e) = \beta \mathbb{E}\left[v_K(K', \theta', e') \right] + \lambda, \qquad \lambda \geq 0, \quad \lambda I = 0, \tag{119.8}$$

and the price of installed capital is

$$q = \frac{\beta \mathbb{E}[v_K(K', \theta', e')]}{u_c(c, e)} = 1 - \frac{\lambda}{u_c(c, e)}. \tag{119.9}$$

So this model implies

- $q \leq 1$ always
- $q = 1$ exactly when investment is positive
- $q < 1$ exactly when the irreversibility constraint binds, which is when investment is zero

Notice what (119.9) says: *Tobin's q is the derivative of the planner's value function*, normalized by marginal utility.

Everything therefore depends on whether v_K exists.

119.5 Differentiability of the value function

Here is the difficulty.

The standard tool for differentiability of a value function is the theorem of Benveniste and Scheinkman [1979].

Theorem 119.5.1 (Benveniste-Scheinkman)

Let v be concave on an open set D and let $K_0 \in D$.

Suppose W is concave, differentiable at K_0, and satisfies $W(K) \leq v(K)$ on a neighborhood of K_0, with $W(K_0) = v(K_0)$.

Then v is differentiable at K_0 and $v'(K_0) = W'(K_0)$.

The usual way to build such a W is to take the optimal plan at K_0, *freeze* next period's capital at its optimal value K_0', and let consumption absorb the change in K:

$$W_1(K) = u(f(K)\theta + (1-\delta)K - K_0', e) + \beta \mathbb{E}[v(K_0', \theta', e')]. \tag{119.10}$$

This is a feasible plan, so $W_1 \leq v$, with equality at K_0, and it is differentiable.

But feasibility requires $K_0' \geq (1-\delta)K$, that is $K \leq K_0'/(1-\delta)$.

When investment is positive, $K_0' > (1-\delta)K_0$ and the plan is feasible on a full neighborhood of K_0.

When investment is *zero*, $K_0' = (1-\delta)K_0$ exactly, and the plan is infeasible for every $K > K_0$.

At a corner, the standard support function is available only from the left.

This is precisely where the published argument becomes delicate, and in our reading it has gaps.

Note

Two specific difficulties.

First, Sargent [1980] states the derivative of the value function in the form

$$v_K(K, \theta, e) = u_c(c, e)[f'(K)\theta + (1-\delta)],$$

which is obtained by substituting the first-order condition (119.8) *with equality* into the envelope condition.

That substitution is legitimate only where investment is positive.

Where investment is zero, $\lambda > 0$ and the formula overstates v_K; the paper's own subsequent inequality records this, so the two statements are not consistent with each other.

Second, the argument establishes differentiability of the value function by induction along the iterates $v^j = T^j v^0$, and then passes to the limit.

That route requires knowing how the set of capital stocks at which the constraint just binds is structured.

The published proof handles three configurations, drawn in its Figures 3-5, and appeals to an appendix that assumes the iterates are twice differentiable almost everywhere.

The appendix in turn justifies this by differentiating the first-order condition and observing that "since the right-hand side exists almost everywhere, so does the left," which does not follow as stated, and it asserts without proof that the set of binding points has Lebesgue measure zero.

Fortunately the conclusion is true, and there is a route to it that avoids the induction on iterates entirely.

The key observation is that at a corner a *different* feasible plan supplies the needed support function: investing nothing is feasible at *every* capital stock, so it works on both sides.

Proposition 119.5.1

Assume u and f are increasing, strictly concave and continuously differentiable, with $u_c(0, e) = \infty$ and $f'(0) = \infty$.

Then for each (θ, e) the value function $v(\cdot, \theta, e)$ defined by (119.7) is continuously differentiable on $(0, \infty)$, and

$$v_K(K, \theta, e) = u_c(c, e) f'(K)\theta + \beta(1-\delta)\mathbb{E}\big[v_K(K', \theta', e')\big], \tag{119.11}$$

where c and K' are optimal at (K, θ, e).

Where investment is positive, (119.11) reduces to $v_K = u_c(c, e)[f'(K)\theta + (1-\delta)]$.

Where investment is zero, $v_K < u_c(c, e)[f'(K)\theta + (1-\delta)]$.

Proof. **Step 0.** The operator associated with (119.7) maps bounded continuous concave functions into bounded continuous strictly concave functions, so $v(\cdot, \theta, e)$ is concave and continuous, and the optimal policy is single valued and continuous.

Step 1 (investment positive). Suppose $I(K_0, \theta, e) > 0$, and let K_0' be the optimal choice.

Then $K_0' > (1-\delta)K_0$, so by continuity the plan that freezes next period's capital at K_0' is feasible for all K in a neighborhood of K_0.

The function W_1 of (119.10) is therefore concave, continuously differentiable, dominated by v, and equal to v at K_0.

Theorem 119.5.1 gives differentiability at K_0 with

$$v_K(K_0, \theta, e) = u_c(c, e)\big[f'(K_0)\theta + (1-\delta)\big].$$

Because the first-order condition holds with equality here, $\beta\mathbb{E}[v_K(K_0', \cdot)] = u_c(c, e)$, and substituting gives (119.11).

Step 2 (investment zero, conditional on a smaller capital stock). Suppose $I(K_0, \theta, e) = 0$, so that $K_0' = (1-\delta)K_0$ and investing nothing is optimal at K_0.

Consider instead the plan that invests nothing this period and behaves optimally thereafter:

$$W_2(K) = u(f(K)\theta, e) + \beta \mathbb{E}[v((1-\delta)K, \theta', e')].$$

Investing nothing is feasible at *every* capital stock, so $W_2 \leq v$ everywhere, with $W_2(K_0) = v(K_0)$, and W_2 is concave.

Suppose that $v(\cdot, \theta', e')$ is differentiable at $(1-\delta)K_0$ for every (θ', e').

Then W_2 is differentiable at K_0; differentiation under the expectation is legitimate because the functions $v(\cdot, \theta', e')$ are concave, hence locally Lipschitz with a common bound on a compact neighborhood.

Theorem 119.5.1 then gives differentiability of v at K_0, with

$$v_K(K_0, \theta, e) = u_c(c, e) f'(K_0)\theta + \beta(1-\delta)\mathbb{E}[v_K((1-\delta)K_0, \theta', e')],$$

which is (119.11).

Step 3 (a region where the corner cannot bind). Because $f'(0) = \infty$, there is an $\eta > 0$ such that $I(K, \theta, e) > 0$ for every $K \leq \eta$ and every (θ, e).

The intuition is that the marginal product of capital becomes arbitrarily large as capital approaches zero, so that giving up a little consumption today buys a great deal of consumption tomorrow.

Making this precise requires comparing the rates at which the two sides of (119.8) diverge as $K \downarrow 0$, since with, say, logarithmic utility both sides become unbounded.

Sargent [1980] proves the statement in his Appendix A (his Proposition A2), under an explicit restriction on u and f that he imposes for this purpose.

That result is the one ingredient of his appendix that the argument below borrows, and it is independent of the differentiability question at issue here.

By Step 1, $v(\cdot, \theta, e)$ is therefore differentiable on $(0, \eta]$ for every (θ, e).

Step 4 (bootstrap). Let

$$D = \{K > 0 : v(\cdot, \theta, e) \text{ is differentiable at } K \text{ for every } (\theta, e)\}.$$

We claim that if $(0, A) \subseteq D$ then $(0, A/(1-\delta)) \subseteq D$.

Take $K < A/(1-\delta)$ and any (θ, e).

If $I(K, \theta, e) > 0$, Step 1 applies directly.

If $I(K, \theta, e) = 0$, then $(1-\delta)K < A$, so $v(\cdot, \theta', e')$ is differentiable at $(1-\delta)K$ for every (θ', e'), and Step 2 applies.

By Step 3 we may start the induction at $A = \eta$.

Since $1/(1-\delta) > 1$, iterating the claim gives $(0, \eta(1-\delta)^{-n}) \subseteq D$ for every n, and therefore $D = (0, \infty)$.

Step 5 (continuity). A concave function that is differentiable on an open interval has a continuous derivative there, so $v(\cdot, \theta, e)$ is continuously differentiable.

Finally, where investment is zero the first-order condition holds with $\lambda > 0$, so $\beta\mathbb{E}[v_K(K', \cdot)] < u_c$, and (119.11) gives $v_K < u_c[f'(K)\theta + (1-\delta)]$.

The corrected formula has a natural reading.

Iterating (119.11) forward gives

$$v_K(K_t, \theta_t, e_t) = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j (1-\delta)^j \, u_c(c_{t+j}, e_{t+j}) \, f'(K_{t+j}) \theta_{t+j}. \tag{119.12}$$

A marginal unit of capital installed today yields marginal products for as long as it survives, and each is valued at the marginal utility of consumption on the date it arrives.

When investment is positive, the market prices that whole stream at the cost of a unit of new output, and (119.12) collapses to $u_c[f'(K)\theta + (1-\delta)]$.

When investment is zero, it does not, and the difference is exactly what pushes q below one.

119.6 Computing the model

We now solve the model of (119.7) with

$$u(c, e) = e \ln c, \qquad f(K) = K^a.$$

The shocks θ and e each take two values with equal probability and are independent of each other and over time.

One numerical detail matters a great deal.

The irreversibility constraint says $K' \geq (1-\delta)K$, so we want $(1-\delta)K$ to be a point of the capital grid whenever K is.

Following Sargent [1980], we use a *geometric* grid with ratio $(1-\delta)^{1/z}$ for an integer z, so that $(1-\delta)K_i = K_{i-z}$ exactly.

```
Model = namedtuple("Model", "β δ a z K θ e W nθ ne")

def create_model(β=0.95, δ=0.05, a=0.25, z=10, n_K=500, K_hi=20.0,
                 θ_vals=(0.85, 1.15), e_vals=(0.6, 1.4)):
    "Geometric capital grid so that (1-δ)K is itself a grid point."
    ratio = (1 - δ)**(1/z)
    K = K_hi * ratio**np.arange(n_K - 1, -1, -1)
    θ, e = np.array(θ_vals), np.array(e_vals)
    W = np.outer(np.ones(len(θ))/len(θ), np.ones(len(e))/len(e))
    return Model(β, δ, a, z, K, θ, e, W, len(θ), len(e))

f  = lambda m, K: K**m.a
f_prime = lambda m, K: m.a * K**(m.a - 1)
```

The planner chooses next period's capital from the grid, subject to positive consumption and to irreversibility.

```
def solve_model(m, irreversible=True, tol=1e-10, maxit=5000, howard=50):
    "Value function iteration with Howard policy improvement steps."
    nK = len(m.K)
    C = f(m, m.K)[:, None] * m.θ[None, :]                       # output (K, θ)
    C = C[:, None, :] + (1 - m.δ)*m.K[:, None, None] - m.K[None, :, None]
    ok = C > 1e-12                                              # (K, K', θ)

    if irreversible:                                            # K' ≥ (1-δ)K
        irr = np.zeros((nK, nK), bool)
        for i in range(nK):
            irr[i, max(i - m.z, 0):] = True
        ok = ok & irr[:, :, None]
```

(continues on next page)

(continued from previous page)

```
    U = np.where(ok[:, :, :, None],
                 m.e[None, None, None, :] * np.log(np.where(ok, C, 1.0))[:, :, :,␣
↪None],
                 -1e12)

    v = np.zeros((nK, m.nθ, m.ne))
    for it in range(maxit):
        EV = np.tensordot(v, m.W, axes=([1, 2], [0, 1]))        # E v(K') over (θ', e
↪')
        obj = U + m.β * EV[None, :, None, None]
        idx = obj.argmax(axis=1)
        v_new = np.take_along_axis(obj, idx[:, None, :, :], axis=1)[:, 0, :, :]

        for _ in range(howard):
            EV = np.tensordot(v_new, m.W, axes=([1, 2], [0, 1]))
            v_new = (np.take_along_axis(U, idx[:, None, :, :], axis=1)[:, 0, :, :]
                     + m.β * EV[idx])

        if np.max(np.abs(v_new - v)) < tol:
            v = v_new
            break
        v = v_new

    K_next = m.K[idx]
    I = K_next - (1 - m.δ) * m.K[:, None, None]
    C_pol = (f(m, m.K)[:, None, None] * m.θ[None, :, None]
             + (1 - m.δ) * m.K[:, None, None] - K_next)
    return v, idx, K_next, I, C_pol
```

To compute q we need v_K, and *Proposition 119.5.1* tells us how to get it: equation (119.11) is a linear fixed point problem in v_K given the optimal policy, and it is a contraction with modulus $\beta(1-\delta)$.

```
def marginal_value(m, idx, C_pol, tol=1e-13, maxit=50_000):
    "Solve the envelope equation v_K = u_c f'(K)θ + β(1-δ) E v_K(K')."
    vK = np.zeros_like(C_pol)
    direct = (m.e[None, None, :]/C_pol) * f_prime(m, m.K)[:, None, None] * m.θ[None,␣
↪:, None]
    for _ in range(maxit):
        EvK = np.tensordot(vK, m.W, axes=([1, 2], [0, 1]))
        new = direct + m.β * (1 - m.δ) * EvK[idx]
        if np.max(np.abs(new - vK)) < tol:
            return new
        vK = new
    return vK

m = create_model()
v, idx, K_next, I, C_pol = solve_model(m)
vK = marginal_value(m, idx, C_pol)

u_c = m.e[None, None, :] / C_pol
E_vK = np.tensordot(vK, m.W, axes=([1, 2], [0, 1]))
q = m.β * E_vK[idx] / u_c
corner = I <= 1e-12

print(f"share of states with zero investment: {corner.mean():.3f}")
```

```
share of states with zero investment: 0.567
```

Let's check *Proposition 119.5.1* numerically, by comparing the envelope formula with a finite-difference derivative of the computed value function.

```
vK_fd = np.gradient(v, m.K, axis=0)
naive = u_c * (f_prime(m, m.K)[:, None, None] * m.θ[None, :, None] + (1 - m.δ))
mid = ((m.K > 2.2) & (m.K < 7.0))[:, None, None] & np.ones_like(corner)

rel = lambda x: np.median((np.abs(x - vK_fd)/np.abs(vK_fd))[mid & sel])

sel = ~corner
print(f"investment positive: envelope error {rel(vK):.4f}, "
      f"naive formula error {rel(naive):.4f}")
sel = corner
print(f"investment zero:     envelope error {rel(vK):.4f}, "
      f"naive formula error {rel(naive):.4f}")
```

```
investment positive: envelope error 0.0005, naive formula error 0.0066
investment zero:     envelope error 0.0001, naive formula error 0.3856
```

Where investment is positive, both formulas agree with the numerical derivative.

Where investment is zero, the envelope formula (119.11) remains accurate while the formula $u_c[f'(K)\theta + (1-\delta)]$ is off by tens of percent.

This is the practical content of the correction: at corners the two expressions are genuinely different objects, and it is the envelope formula that gives the price of installed capital.

We can also verify Step 3 of the proof, which asserts a region of small capital stocks where the corner never binds.

```
m_wide = create_model(n_K=700)        # a grid reaching lower capital stocks
_, idx_w, _, I_w, _ = solve_model(m_wide)
frac_corner = (I_w <= 1e-12).reshape(len(m_wide.K), -1).mean(axis=1)
first = np.argmax(frac_corner > 0)
print(f"grid runs from K = {m_wide.K.min():.3f}")
print(f"investment is positive for every shock when K < {m_wide.K[first]:.3f}")
```

```
grid runs from K = 0.555
investment is positive for every shock when K < 1.058
```

119.6.1 Policies and the price of installed capital

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))
show = (m.K > 1.5) & (m.K < 9)

for j in range(m.nθ):
    for l in range(m.ne):
        lab = rf'$\theta={m.θ[j]}, e={m.e[l]}$'
        axes[0].plot(m.K[show], K_next[show, j, l], label=lab)
        axes[1].plot(m.K[show], q[show, j, l], label=lab)

axes[0].plot(m.K[show], m.K[show], 'k--', lw=1, label='45 degree line')
axes[0].plot(m.K[show], (1-m.δ)*m.K[show], 'k:', lw=1, label=r'$(1-\delta)K$')
axes[0].set_xlabel('$K$'); axes[0].set_ylabel("$K'$"); axes[0].set_title('law of␣
```

(continues on next page)

(continued from previous page)

```
↪motion')
axes[1].axhline(1.0, color='k', ls='--', lw=1)
axes[1].set_xlabel('$K$'); axes[1].set_ylabel('$q$'); axes[1].set_title("Tobin's $q$")
axes[0].legend(fontsize=8); axes[1].legend(fontsize=8)
plt.tight_layout()
plt.show()
```

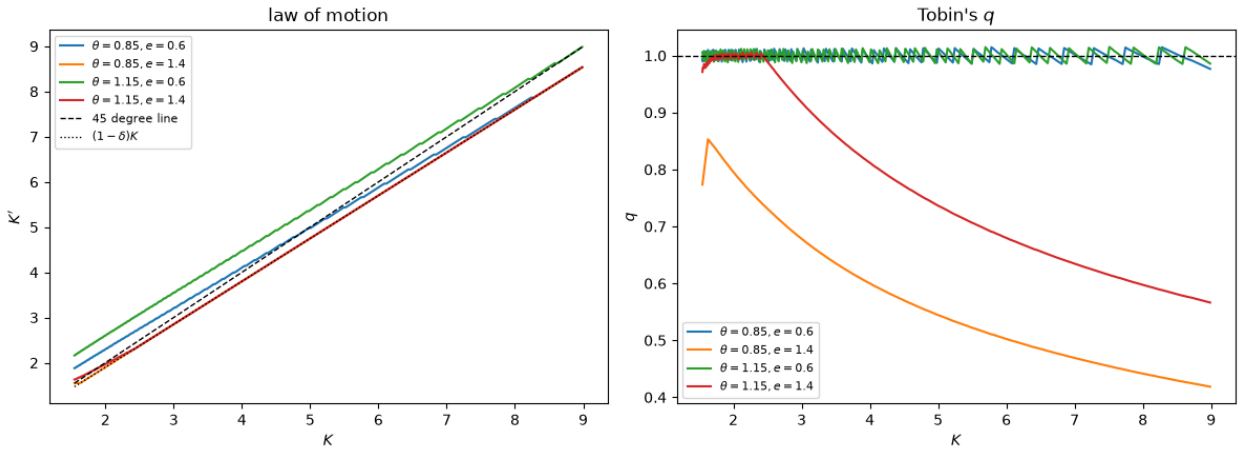


Fig. 119.1: Law of motion and Tobin's q

The law of motion runs along the lower dotted line $(1-\delta)K$ whenever the irreversibility constraint binds.

Exactly on that region, q falls below one.

119.6.2 The invariant distribution

As in *Investment Under Uncertainty*, we check that the long-run distribution does not depend on where the economy starts.

```
def simulate(m, idx, K0, T=50_000, seed=0, burn=500):
    "Simulate the equilibrium Markov process."
    rng = np.random.default_rng(seed)
    j_idx = rng.integers(0, m.nθ, T)
    l_idx = rng.integers(0, m.ne, T)
    k = np.abs(m.K - K0).argmin()
    path = np.empty(T, int)
    for t in range(T):
        path[t] = k
        k = idx[k, j_idx[t], l_idx[t]]
    return path[burn:], j_idx[burn:], l_idx[burn:]

p_lo, j_lo, l_lo = simulate(m, idx, K0=2.0, seed=7)
p_hi, j_hi, l_hi = simulate(m, idx, K0=15.0, seed=7)

print(f"mean capital starting from K0 = 2:  {m.K[p_lo].mean():.4f}")
print(f"mean capital starting from K0 = 15: {m.K[p_hi].mean():.4f}")
print(f"capital visited: [{m.K[p_lo].min():.2f}, {m.K[p_lo].max():.2f}]")
```

```
mean capital starting from K0 = 2:  4.5395
mean capital starting from K0 = 15: 4.5395
capital visited: [2.06, 7.24]
```

```
fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))

axes[0].hist(m.K[p_lo], bins=60, density=True, alpha=0.5, label='from $K_0 = 2$')
axes[0].hist(m.K[p_hi], bins=60, density=True, alpha=0.5, label='from $K_0 = 15$')
axes[0].set_xlabel('$K$'); axes[0].set_ylabel('density')
axes[0].set_title('invariant distribution of capital')
axes[0].legend()

q_sim, I_sim = q[p_lo, j_lo, l_lo], I[p_lo, j_lo, l_lo]
axes[1].scatter(q_sim, I_sim, s=4, alpha=0.2)
axes[1].set_xlabel('$q$'); axes[1].set_ylabel('$I$')
axes[1].set_title("investment and Tobin's $q$")

plt.tight_layout()
plt.show()
```

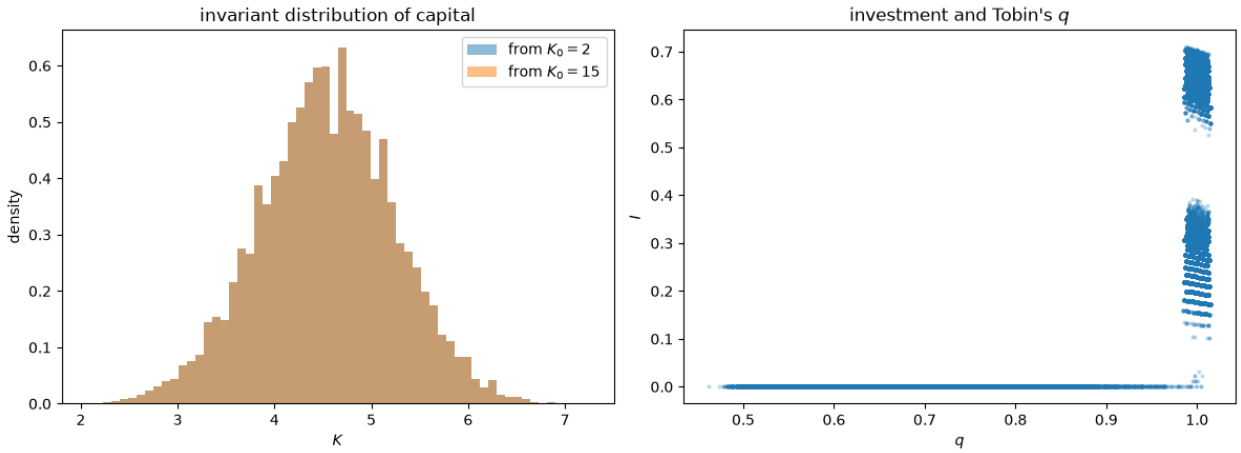


Fig. 119.2: Invariant distribution and investment against q

The two histograms coincide, as *Theorem 119.3.1* leads us to expect.

The scatter plot on the right reproduces the central picture of Sargent [1980].

Investment is positive only when q equals one; when q is below one, investment is exactly zero.

```
print(f"q when I > 0: [{q_sim[I_sim > 1e-12].min():.4f}, {q_sim[I_sim > 1e-12].max():.
↪4f}]")
print(f"q when I = 0: [{q_sim[I_sim <= 1e-12].min():.4f}, {q_sim[I_sim <= 1e-12].
↪max():.4f}]")
```

```
q when I > 0: [0.9851, 1.0148]
q when I = 0: [0.4619, 1.0040]
```

The small departures from exactly one reflect the discreteness of the capital grid.

Sargent's point is now visible.

There *is* a positive relationship between investment and q in these data, and an econometrician could certainly run that regression.

But the relationship is not an investment demand schedule.

It is a **mongrel relation** that mixes together preferences, technology and the distribution of the shocks, and it will shift whenever any of them changes.

Exercise 119.7.3 asks you to demonstrate this.

119.7 Exercises

Exercise 119.7.1

Brock and Mirman [1972] is famous partly for a special case that can be solved by hand.

Let $u(c) = \ln c$, let $f(k, r) = rk^\alpha$ with $0 < \alpha < 1$, and let capital depreciate fully each period, so that $k_{t+1} = f(k_t, r_t) - c_t$.

1. Verify that the optimal policy is $k_{t+1} = \alpha\beta\, r_t k_t^\alpha$.
2. Show that $\ln k_t$ follows an AR(1) process, and derive the mean and variance of its invariant distribution when $\ln r_t \sim N(\mu, \sigma^2)$.
3. Simulate the model and check both the invariant distribution and the mean-ergodic property (119.5).

Solution

Guess that a constant fraction of resources is saved, $k_{t+1} = s\, r_t k_t^\alpha$.

The Euler equation for this problem is

$$\frac{1}{c_t} = \beta \mathbb{E}_t \left[\frac{1}{c_{t+1}} \alpha r_{t+1} k_{t+1}^{\alpha-1} \right].$$

With $c_t = (1-s) r_t k_t^\alpha$ and $k_{t+1} = s r_t k_t^\alpha$, the right side becomes

$$\beta \mathbb{E}_t \left[\frac{\alpha r_{t+1} k_{t+1}^{\alpha-1}}{(1-s) r_{t+1} k_{t+1}^\alpha} \right] = \frac{\beta\alpha}{(1-s) k_{t+1}} = \frac{\beta\alpha}{(1-s) s r_t k_t^\alpha},$$

while the left side is $1/[(1-s) r_t k_t^\alpha]$.

Equating the two gives $s = \alpha\beta$.

Taking logs of the policy,

$$\ln k_{t+1} = \ln(\alpha\beta) + \alpha \ln k_t + \ln r_t,$$

an AR(1) with coefficient α.

With $\ln r_t \sim N(\mu, \sigma^2)$, the invariant distribution of $\ln k$ is normal with

$$\text{mean} = \frac{\ln(\alpha\beta) + \mu}{1-\alpha}, \qquad \text{variance} = \frac{\sigma^2}{1-\alpha^2}.$$

```
α, β_d, μ, σ = 0.4, 0.95, 0.0, 0.1
T = 200_000
rng = np.random.default_rng(0)
ln_r = μ + σ * rng.normal(size=T)

k = np.empty(T + 1)
k[0] = 1.0
for t in range(T):
    k[t+1] = α * β_d * np.exp(ln_r[t]) * k[t]**α

ln_k = np.log(k[1000:])
print(f"mean of ln k: simulated {ln_k.mean():.4f}, "
      f"theory {(np.log(α*β_d) + μ)/(1-α):.4f}")
print(f"var  of ln k: simulated {ln_k.var():.4f}, "
      f"theory {σ**2/(1-α**2):.4f}")
```

```
mean of ln k: simulated -1.6126, theory -1.6126
var  of ln k: simulated 0.0119, theory 0.0119
```

The time averages match the population moments of the invariant distribution, which is the mean-ergodic property (119.5) in action.

Exercise 119.7.2

This exercise examines the difference between the two candidate formulas for v_K discussed above.

For the baseline model, compute at every state

- the envelope formula (119.11)
- the formula $u_c(c,e)[f'(K)\theta + (1-\delta)]$
- a finite-difference derivative of the value function

Then

1. Confirm that all three agree where investment is positive.
2. Confirm that the second disagrees with the other two where investment is zero, and report by how much.
3. Explain, using (119.8), why the second formula is an upper bound for v_K.
4. What would go wrong with the computed q if you used the second formula everywhere?

Solution

The second formula is what you get by substituting the first-order condition (119.8) into the envelope condition *as though* $\lambda = 0$.

Since $\lambda \geq 0$, dropping it can only raise the expression, so

$$v_K = u_c f'(K)\theta + (1-\delta)[u_c - \lambda] \leq u_c[f'(K)\theta + (1-\delta)],$$

with equality exactly when $\lambda = 0$, that is, when investment is positive.

```
gap = np.abs(naive - vK_fd)/np.abs(vK_fd)
env = np.abs(vK - vK_fd)/np.abs(vK_fd)
```

```
for name, sel in (("investment positive", ~corner), ("investment zero", corner)):
    s = mid & sel
    print(f"{name:20}: envelope {np.median(env[s]):.4f}   "
          f"naive {np.median(gap[s]):.4f}   max naive {gap[s].max():.4f}")

q_naive = m.β * np.tensordot(naive, m.W, axes=([1, 2], [0, 1]))[idx] / u_c
print(f"\nusing the naive formula, max q = {q_naive.max():.3f} "
      f"(theory says q ≤ 1)")
print(f"fraction of states with q > 1: {(q_naive > 1 + 1e-8).mean():.3f}")
```

```
investment positive : envelope 0.0005   naive 0.0066   max naive 0.0134
investment zero     : envelope 0.0001   naive 0.3856   max naive 1.0175

using the naive formula, max q = 1.863 (theory says q ≤ 1)
fraction of states with q > 1: 0.578
```

Using the second formula everywhere produces values of q that exceed one, which the theory rules out.

The reason is economic, not numerical: at a corner the market does not price installed capital at the cost of new capital, precisely because the household would like to sell capital and cannot.

Exercise 119.7.3

Show that the regression of investment on q is not structural.

Solve and simulate the model for several economies that differ in

- the amount of preference risk (the spread of e)
- the amount of technology risk (the spread of θ)
- the depreciation rate δ

For each economy, report the slope and R^2 of a regression of I on q, together with the frequency with which investment is zero.

Comment on what this implies for an econometrician who estimates an "investment demand schedule" relating investment to q.

Solution

Here is one solution.

```
def q_regression(**kwargs):
    "Solve an economy, simulate it, and regress I on q."
    mm = create_model(**kwargs)
    vv, ii, _, II, CC = solve_model(mm)
    vvK = marginal_value(mm, ii, CC)
    uu_c = mm.e[None, None, :]/CC
    qq = mm.β * np.tensordot(vvK, mm.W, axes=([1, 2], [0, 1]))[ii] / uu_c
    p, j, l = simulate(mm, ii, K0=3.5)
    q_s, I_s = qq[p, j, l], II[p, j, l]
    slope = np.polyfit(q_s, I_s, 1)[0]
    corr = np.corrcoef(I_s, q_s)[0, 1]
    return slope, corr**2, (I_s <= 1e-12).mean()

cases = [("baseline",              {}),
```

```
        ("less preference risk",  dict(e_vals=(0.8, 1.2))),
        ("more preference risk",  dict(e_vals=(0.4, 1.6))),
        ("less technology risk",  dict(θ_vals=(0.95, 1.05))),
        ("more technology risk",  dict(θ_vals=(0.7, 1.3))),
        ("faster depreciation",   dict(δ=0.10))]

print(f"{'economy':24}{'slope':>9}{'R^2':>8}{'zero I':>9}")
for name, kw in cases:
    slope, r2, zero = q_regression(**kw)
    print(f"{name:24}{slope:9.3f}{r2:8.3f}{zero:9.3f}")
```

```
economy                  slope     R^2   zero I
baseline                 1.127   0.594    0.500
less preference risk     1.094   0.283    0.279
more preference risk     1.131   0.780    0.501
less technology risk     1.308   0.904    0.501
more technology risk     0.780   0.234    0.529
faster depreciation      1.271   0.436    0.456
```

The slope and the fit both move around substantially across economies that share the same preferences and technology parameters but differ in the distribution of shocks or in the depreciation rate.

Nothing in the regression is invariant, so it cannot be used to predict what investment would do under a policy that changed any of these features.

This is the message of Sargent [1980], and it is a concrete instance of the Lucas critique: the very friction that makes q interesting – the occasionally binding irreversibility constraint – also makes the relationship between I and q a reduced-form artifact rather than a decision rule.

STABILITY IN LINEAR RATIONAL EXPECTATIONS MODELS

Contents

- *Stability in Linear Rational Expectations Models*
  - *Overview*
  - *Linear Difference Equations*
  - *Illustration: Cagan's Model*
  - *Some Python Code*
  - *Alternative Code*
  - *Another Perspective*
  - *Log money Supply Feeds Back on Log Price Level*
  - *Big P, Little p Interpretation*
  - *Fun with SymPy*

In addition to what's in Anaconda, this lecture deploys the following libraries:

```
!pip install quantecon
```

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
from sympy import init_printing, symbols, Matrix
init_printing()
```

120.1 Overview

This lecture studies stability in the context of an elementary rational expectations model.

We study a rational expectations version of Philip Cagan's model [Cagan, 1956] linking the price level to the money supply.

Cagan did not use a rational expectations version of his model, but Sargent [Sargent, 1977] did.

We study a rational expectations version of this model because it is intrinsically interesting and because it has a mathematical structure that appears in virtually all linear rational expectations model, namely, that a key endogenous variable equals a mathematical expectation of a geometric sum of future values of another variable.

The model determines the price level or rate of inflation as a function of the money supply or the rate of change in the money supply.

In this lecture, we'll encounter:

- a convenient formula for the expectation of geometric sum of future values of a variable
- a way of solving an expectational difference equation by mapping it into a vector first-order difference equation and appropriately manipulating an eigen decomposition of the transition matrix in order to impose stability
- a way to use a Big K, little k argument to allow apparent feedback from endogenous to exogenous variables within a rational expectations equilibrium
- a use of eigenvector decompositions of matrices that allowed Blanchard and Khan (1981) [Blanchard and Kahn, 1980] and Whiteman (1983) [Whiteman, 1983] to solve a class of linear rational expectations models
- how to use **SymPy** to get analytical formulas for some key objects comprising a rational expectations equilibrium

Matrix decompositions employed here are described in more depth in this lecture *Lagrangian formulations*.

We formulate a version of Cagan's model under rational expectations as an **expectational difference equation** whose solution is a rational expectations equilibrium.

We'll start this lecture with a quick review of deterministic (i.e., non-random) first-order and second-order linear difference equations.

120.2 Linear Difference Equations

We'll use the **backward shift** or **lag** operator L.

The lag operator L maps a sequence $\{x_t\}_{t=0}^{\infty}$ into the sequence $\{x_{t-1}\}_{t=0}^{\infty}$

We'll deploy L by using the equality $Lx_t \equiv x_{t-1}$ in algebraic expressions.

Further, the inverse L^{-1} of the lag operator is the **forward shift** operator.

We'll often use the equality $L^{-1}x_t \equiv x_{t+1}$ below.

The algebra of lag and forward shift operators can simplify representing and solving linear difference equations.

120.2.1 First Order

We want to solve a linear first-order scalar difference equation.

Let $|\lambda| < 1$ and let $\{u_t\}_{t=-\infty}^{\infty}$ be a bounded sequence of scalar real numbers.

Let L be the lag operator defined by $Lx_t \equiv x_{t-1}$ and let L^{-1} be the forward shift operator defined by $L^{-1}x_t \equiv x_{t+1}$.

Then

$$(1 - \lambda L)y_t = u_t, \forall t \tag{120.1}$$

has solutions

$$y_t = (1 - \lambda L)^{-1}u_t + k\lambda^t \tag{120.2}$$

or

$$y_t = \sum_{j=0}^{\infty} \lambda^j u_{t-j} + k\lambda^t$$

for any real number k.

You can verify this fact by applying $(1 - \lambda L)$ to both sides of equation (120.2) and noting that $(1 - \lambda L)\lambda^t = 0$.

To pin down k we need one condition imposed from outside (e.g., an initial or terminal condition) on the path of y.

Now let $|\lambda| > 1$.

Rewrite equation (120.1) as

$$y_{t-1} = \lambda^{-1} y_t - \lambda^{-1} u_t, \forall t \tag{120.3}$$

or

$$(1 - \lambda^{-1} L^{-1}) y_t = -\lambda^{-1} u_{t+1}. \tag{120.4}$$

A solution is

$$y_t = -\lambda^{-1} \left(\frac{1}{1 - \lambda^{-1} L^{-1}} \right) u_{t+1} + k\lambda^t \tag{120.5}$$

for any k.

To verify that this is a solution, check the consequences of operating on both sides of equation (120.5) by $(1 - \lambda L)$ and compare to equation (120.1).

For any bounded $\{u_t\}$ sequence, solution (120.2) exists for $|\lambda| < 1$ because the **distributed lag** in u converges.

Solution (120.5) exists when $|\lambda| > 1$ because the **distributed lead** in u converges.

When $|\lambda| > 1$, the distributed lag in u in (120.2) may diverge, in which case a solution of this form does not exist.

The distributed lead in u in (120.5) need not converge when $|\lambda| < 1$.

120.2.2 Second Order

Now consider the second order difference equation

$$(1 - \lambda_1 L)(1 - \lambda_2 L) y_{t+1} = u_t \tag{120.6}$$

where $\{u_t\}$ is a bounded sequence, y_0 is an initial condition, $|\lambda_1| < 1$ and $|\lambda_2| > 1$.

We seek a bounded sequence $\{y_t\}_{t=0}^{\infty}$ that satisfies (120.6). Using insights from our analysis of the first-order equation, operate on both sides of (120.6) by the forward inverse of $(1 - \lambda_2 L)$ to rewrite equation (120.6) as

$$(1 - \lambda_1 L) y_{t+1} = -\frac{\lambda_2^{-1}}{1 - \lambda_2^{-1} L^{-1}} u_{t+1}$$

or

$$y_{t+1} = \lambda_1 y_t - \lambda_2^{-1} \sum_{j=0}^{\infty} \lambda_2^{-j} u_{t+j+1}. \tag{120.7}$$

Thus, we obtained equation (120.7) by solving a stable root (in this case λ_1) **backward**, and an unstable root (in this case λ_2) **forward**.

Equation (120.7) has a form that we shall encounter often.

- $\lambda_1 y_t$ is called the **feedback part**
- $-\frac{\lambda_2^{-1}}{1 - \lambda_2^{-1} L^{-1}} u_{t+1}$ is called the **feedforward part**

120.3 Illustration: Cagan's Model

Now let's use linear difference equations to represent and solve Sargent's [Sargent, 1977] rational expectations version of Cagan's model [Cagan, 1956] that connects the price level to the public's anticipations of future money supplies.

Cagan did not use a rational expectations version of his model, but Sargent [Sargent, 1977]

Let

- m_t^d be the log of the demand for money
- m_t be the log of the supply of money
- p_t be the log of the price level

It follows that $p_{t+1} - p_t$ is the rate of inflation.

The logarithm of the demand for real money balances $m_t^d - p_t$ is an inverse function of the expected rate of inflation $p_{t+1} - p_t$ for $t \geq 0$:

$$m_t^d - p_t = -\beta(p_{t+1} - p_t), \quad \beta > 0$$

Equate the demand for log money m_t^d to the supply of log money m_t in the above equation and rearrange to deduce that the logarithm of the price level p_t is related to the logarithm of the money supply m_t by

$$p_t = (1 - \lambda)m_t + \lambda p_{t+1} \tag{120.8}$$

where $\lambda \equiv \frac{\beta}{1+\beta} \in (0, 1)$.

(We note that the characteristic polynomial if $1 - \lambda^{-1}z^{-1} = 0$ so that the zero of the characteristic polynomial in this case is $\lambda \in (0, 1)$ which here is **inside** the unit circle.)

Solving the first order difference equation (120.8) forward gives

$$p_t = (1 - \lambda) \sum_{j=0}^{\infty} \lambda^j m_{t+j}, \tag{120.9}$$

which is the unique **stable** solution of difference equation (120.8) among a class of more general solutions

$$p_t = (1 - \lambda) \sum_{j=0}^{\infty} \lambda^j m_{t+j} + c\lambda^{-t} \tag{120.10}$$

that is indexed by the real number $c \in \mathbf{R}$.

Because we want to focus on stable solutions, we set $c = 0$.

Equation (120.10) attributes **perfect foresight** about the money supply sequence to the holders of real balances.

We begin by assuming that the log of the money supply is **exogenous** in the sense that it is an autonomous process that does not feed back on the log of the price level.

In particular, we assume that the log of the money supply is described by the linear state space system

$$\begin{aligned} m_t &= Gx_t \\ x_{t+1} &= Ax_t \end{aligned} \tag{120.11}$$

where x_t is an $n \times 1$ vector that does not include p_t or lags of p_t, A is an $n \times n$ matrix with eigenvalues that are less than λ^{-1} in absolute values, and G is a $1 \times n$ selector matrix.

Variables appearing in the vector x_t contain information that might help predict future values of the money supply.

We'll start with an example in which x_t includes only m_t, possibly lagged values of m, and a constant.

An example of such an $\{m_t\}$ process that fits info state space system (120.11) is one that satisfies the second order linear difference equation

$$m_{t+1} = \alpha + \rho_1 m_t + \rho_2 m_{t-1}$$

where the zeros of the characteristic polynomial $(1 - \rho_1 z - \rho_2 z^2)$ are strictly greater than 1 in modulus.

(Please see *this* QuantEcon lecture for more about characteristic polynomials and their role in solving linear difference equations.)

We seek a stable or non-explosive solution of the difference equation (120.8) that obeys the system comprised of (120.8)-(120.11).

By stable or non-explosive, we mean that neither m_t nor p_t diverges as $t \to +\infty$.

This requires that we shut down the term $c\lambda^{-t}$ in equation (120.10) above by setting $c = 0$

The solution we are after is

$$p_t = F x_t \tag{120.12}$$

where

$$F = (1 - \lambda) G (I - \lambda A)^{-1} \tag{120.13}$$

Note

As mentioned above, an **explosive solution** of difference equation (120.8) can be constructed by adding to the right hand of (120.12) a sequence $c\lambda^{-t}$ where c is an arbitrary positive constant.

120.4 Some Python Code

We'll construct examples that illustrate (120.11).

Our first example takes as the law of motion for the log money supply the second order difference equation

$$m_{t+1} = \alpha + \rho_1 m_t + \rho_2 m_{t-1} \tag{120.14}$$

that is parameterized by ρ_1, ρ_2, α

To capture this parameterization with system (120.9) we set

$$x_t = \begin{bmatrix} 1 \\ m_t \\ m_{t-1} \end{bmatrix}, \quad A = \begin{bmatrix} 1 & 0 & 0 \\ \alpha & \rho_1 & \rho_2 \\ 0 & 1 & 0 \end{bmatrix}, \quad G = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}$$

Here is Python code

```
λ = .9

α = 0
ρ1 = .9
ρ2 = .05

A = np.array([[1,  0,  0],
              [α, ρ1, ρ2],
              [0,  1,  0]])
G = np.array([[0, 1, 0]])
```

The matrix A has one eigenvalue equal to unity.

It is associated with the A_{11} component that captures a constant component of the state x_t.

We can verify that the two eigenvalues of A not associated with the constant in the state x_t are strictly less than unity in modulus.

```
eigvals = np.linalg.eigvals(A)
print(eigvals)
```

```
[-0.05249378  0.95249378  1.        ]
```

```
(abs(eigvals) <= 1).all()
```

```
np.True_
```

Now let's compute F in formulas (120.12) and (120.13).

```
# compute the solution, i.e. forumula (3)
F = (1 - λ) * G @ np.linalg.inv(np.eye(A.shape[0]) - λ * A)
print("F= ",F)
```

```
F=  [[0.         0.66889632 0.03010033]]
```

Now let's simulate paths of m_t and p_t starting from an initial value x_0.

```
# set the initial state
x0 = np.array([1, 1, 0])

T = 100 # length of simulation

m_seq = np.empty(T+1)
p_seq = np.empty(T+1)

[m_seq[0]] = G @ x0
[p_seq[0]] = F @ x0

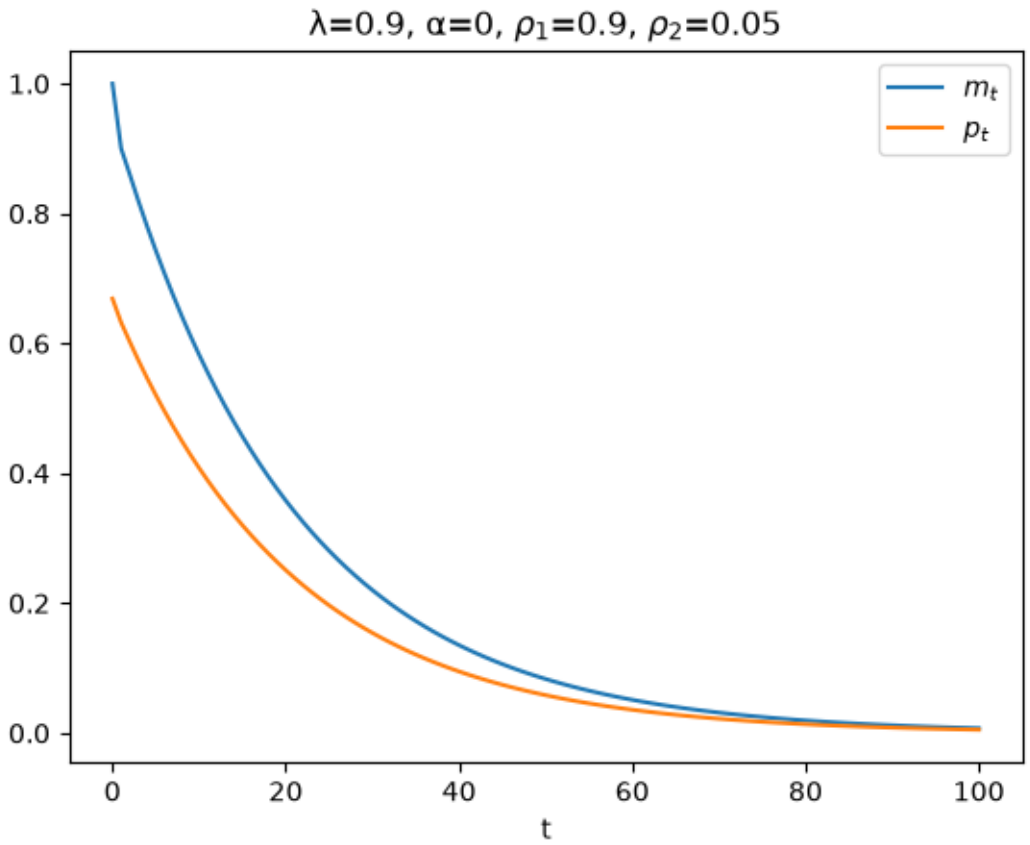
# simulate for T periods
x_old = x0
for t in range(T):

    x = A @ x_old

    [m_seq[t+1]] = G @ x
    [p_seq[t+1]] = F @ x

    x_old = x
```

```
plt.figure()
plt.plot(range(T+1), m_seq, label=r'$m_t$')
plt.plot(range(T+1), p_seq, label=r'$p_t$')
plt.xlabel('t')
plt.title(rf'λ={λ}, α={α}, $ρ_1$={ρ1}, $ρ_2$={ρ2}')
plt.legend()
plt.show()
```



In the above graph, why is the log of the price level always less than the log of the money supply?

Because

- according to equation (120.9), p_t is a geometric weighted average of current and future values of m_t, and
- it happens that in this example future m's are always less than the current m

120.5 Alternative Code

We could also have run the simulation using the quantecon **LinearStateSpace** code.

The following code block performs the calculation with that code.

```
# construct a LinearStateSpace instance

# stack G and F
G_ext = np.vstack([G, F])

C = np.zeros((A.shape[0], 1))

ss = qe.LinearStateSpace(A, C, G_ext, mu_0=x0)
```

```
T = 100

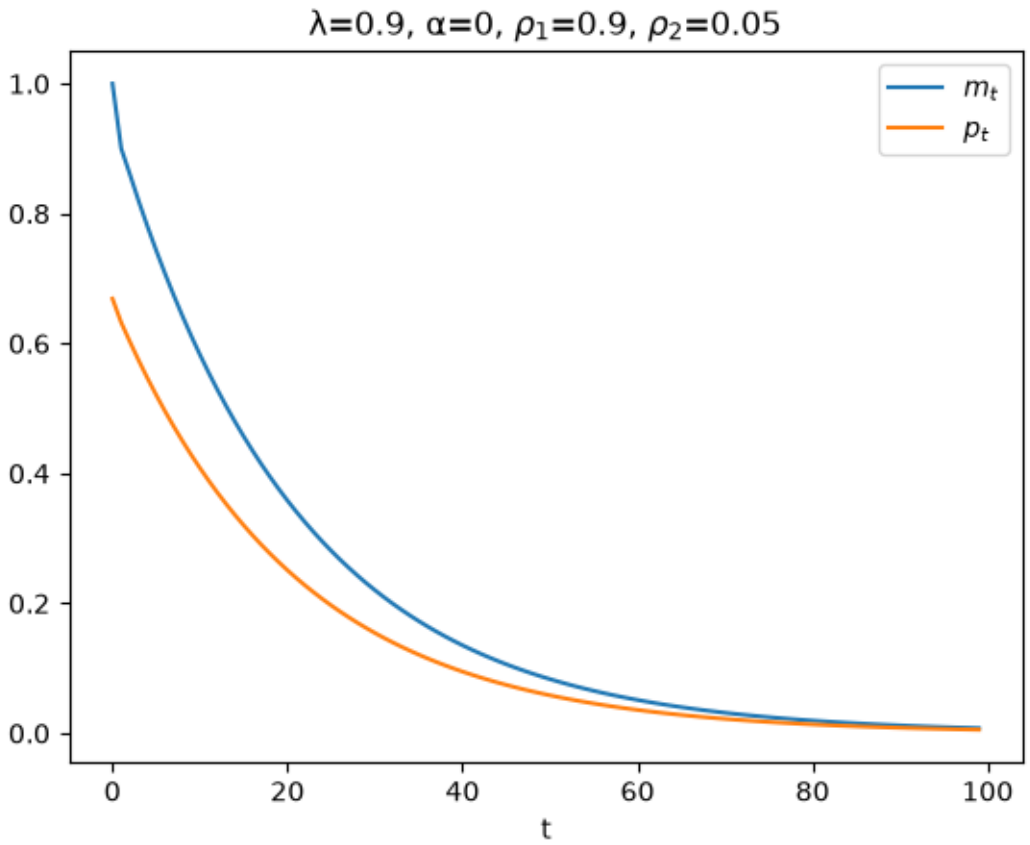
# simulate using LinearStateSpace
```

(continues on next page)

(continued from previous page)

```
x, y = ss.simulate(ts_length=T)

# plot
plt.figure()
plt.plot(range(T), y[0,:], label='$m_t$')
plt.plot(range(T), y[1,:], label='$p_t$')
plt.xlabel('t')
plt.title(f'λ={λ}, α={α}, $ρ_1$={ρ1}, $ρ_2$={ρ2}')
plt.legend()
plt.show()
```



120.5.1 Special Case

To simplify our presentation in ways that will let focus on an important idea, in the above second-order difference equation (120.14) that governs m_t, we now set $\alpha = 0$, $\rho_1 = \rho \in (-1, 1)$, and $\rho_2 = 0$ so that the law of motion for m_t becomes

$$m_{t+1} = \rho m_t \tag{120.15}$$

and the state x_t becomes

$$x_t = m_t.$$

Consequently, we can set $G = 1, A = \rho$ making our formula (120.13) for F become

$$F = (1 - \lambda)(1 - \lambda\rho)^{-1}.$$

so that the log the log price level satisfies

$$p_t = Fm_t.$$

Please keep these formulas in mind as we investigate an alternative route to and interpretation of our formula for F.

120.6 Another Perspective

Above, we imposed stability or non-explosiveness on the solution of the key difference equation (120.8) in Cagan's model by solving the unstable root of the characteristic polynomial forward.

To shed light on the mechanics involved in imposing stability on a solution of a potentially unstable system of linear difference equations and to prepare the way for generalizations of our model in which the money supply is allowed to feed back on the price level itself, we stack equations (120.8) and (120.15) to form the system

$$\begin{bmatrix} m_{t+1} \\ p_{t+1} \end{bmatrix} = \begin{bmatrix} \rho & 0 \\ -(1-\lambda)/\lambda & \lambda^{-1} \end{bmatrix} \begin{bmatrix} m_t \\ p_t \end{bmatrix} \tag{120.16}$$

or

$$y_{t+1} = Hy_t, \quad t \geq 0 \tag{120.17}$$

where

$$H = \begin{bmatrix} \rho & 0 \\ -(1-\lambda)/\lambda & \lambda^{-1} \end{bmatrix}. \tag{120.18}$$

Transition matrix H has eigenvalues $\rho \in (0,1)$ and $\lambda^{-1} > 1$.

Because an eigenvalue of H exceeds unity, if we iterate on equation (120.17) starting from an arbitrary initial vector $y_0 = \begin{bmatrix} m_0 \\ p_0 \end{bmatrix}$ with $m_0 > 0, p_0 > 0$, we discover that in general absolute values of both components of y_t diverge toward $+\infty$ as $t \to +\infty$.

To substantiate this claim, we can use the eigenvector matrix decomposition of H that is available to us because the eigenvalues of H are distinct

$$H = Q\Lambda Q^{-1}.$$

Here Λ is a diagonal matrix of eigenvalues of H and Q is a matrix whose columns are eigenvectors associated with the corresponding eigenvalues.

Note that

$$H^t = Q\Lambda^t Q^{-1}$$

so that

$$y_t = Q\Lambda^t Q^{-1} y_0$$

For almost all initial vectors y_0, the presence of the eigenvalue $\lambda^{-1} > 1$ causes both components of y_t to diverge in absolute value to $+\infty$.

To explore this outcome in more detail, we can use the following transformation

$$y_t^* = Q^{-1} y_t$$

that allows us to represent the dynamics in a way that isolates the source of the propensity of paths to diverge:

$$y^*_{t+1} = \Lambda^t y^*_t$$

Staring at this equation indicates that unless

$$y^*_0 = \begin{bmatrix} y^*_{1,0} \\ 0 \end{bmatrix} \tag{120.19}$$

the path of y^*_t and therefore the paths of both components of $y_t = Q y^*_t$ will diverge in absolute value as $t \to +\infty$. (We say that the paths *explode*)

Equation (120.19) also leads us to conclude that there is a unique setting for the initial vector y_0 for which both components of y_t do not diverge.

The required setting of y_0 must evidently have the property that

$$Q y_0 = y^*_0 = \begin{bmatrix} y^*_{1,0} \\ 0 \end{bmatrix}.$$

But note that since $y_0 = \begin{bmatrix} m_0 \\ p_0 \end{bmatrix}$ and m_0 is given to us an initial condition, p_0 has to do all the adjusting to satisfy this equation.

Sometimes this situation is described by saying that while m_0 is truly a **state** variable, p_0 is a **jump** variable that must adjust at $t = 0$ in order to satisfy the equation.

Thus, in a nutshell the unique value of the vector y_0 for which the paths of y_t do not diverge must have second component p_0 that verifies equality (120.19) by setting the second component of y^*_0 equal to zero.

The component p_0 of the initial vector $y_0 = \begin{bmatrix} m_0 \\ p_0 \end{bmatrix}$ must evidently satisfy

$$Q^{\{2\}} y_0 = 0$$

where $Q^{\{2\}}$ denotes the second row of Q^{-1}, a restriction that is equivalent to

$$Q^{21} m_0 + Q^{22} p_0 = 0 \tag{120.20}$$

where Q^{ij} denotes the (i, j) component of Q^{-1}.

Solving this equation for p_0, we find

$$p_0 = -(Q^{22})^{-1} Q^{21} m_0. \tag{120.21}$$

This is the unique **stabilizing value** of p_0 expressed as a function of m_0.

120.6.1 Refining the Formula

We can get an even more convenient formula for p_0 that is cast in terms of components of Q instead of components of Q^{-1}.

To get this formula, first note that because $(Q^{21} \; Q^{22})$ is the second row of the inverse of Q and because $Q^{-1}Q = I$, it follows that

$$\begin{bmatrix} Q^{21} & Q^{22} \end{bmatrix} \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix} = 0$$

which implies that

$$Q^{21}Q_{11} + Q^{22}Q_{21} = 0.$$

Therefore,

$$-(Q^{22})^{-1}Q^{21} = Q_{21}Q_{11}^{-1}.$$

So we can write

$$p_0 = Q_{21}Q_{11}^{-1}m_0. \tag{120.22}$$

It can be verified that this formula replicates itself over time in the sense that

$$p_t = Q_{21}Q_{11}^{-1}m_t. \tag{120.23}$$

To implement formula (120.23), we want to compute Q_1 the eigenvector of Q associated with the stable eigenvalue ρ of Q.

By hand it can be verified that the eigenvector associated with the stable eigenvalue ρ is proportional to

$$Q_1 = \begin{bmatrix} 1 - \lambda\rho \\ 1 - \lambda \end{bmatrix}.$$

Notice that if we set $A = \rho$ and $G = 1$ in our earlier formula for p_t we get

$$p_t = G(I - \lambda A)^{-1}m_t = (1 - \lambda)(1 - \lambda\rho)^{-1}m_t,$$

a formula that is equivalent with

$$p_t = Q_{21}Q_{11}^{-1}m_t,$$

where

$$Q_1 = \begin{bmatrix} Q_{11} \\ Q_{21} \end{bmatrix}.$$

120.6.2 Remarks about Feedback

We have expressed (120.16) in what superficially appears to be a form in which y_{t+1} feeds back on y_t, even though what we actually want to represent is that the component p_t feeds **forward** on p_{t+1}, and through it, on future m_{t+j}, $j = 0, 1, 2, \ldots$.

A tell-tale sign that we should look beyond its superficial "feedback" form is that $\lambda^{-1} > 1$ so that the matrix H in (120.16) is **unstable**

- it has one eigenvalue ρ that is less than one in modulus that does not imperil stability, but ...
- it has a second eigenvalue λ^{-1} that exceeds one in modulus and that makes H an unstable matrix

We'll keep these observations in mind as we turn now to a case in which the log money supply actually does feed back on the log of the price level.

120.7 Log money Supply Feeds Back on Log Price Level

An arrangement of eigenvalues that split around unity, with one being below unity and another being greater than unity, sometimes prevails when there is *feedback* from the log price level to the log money supply.

Let the feedback rule be

$$m_{t+1} = \rho m_t + \delta p_t \tag{120.24}$$

where $\rho \in (0, 1)$ and where we shall now allow $\delta \neq 0$.

Warning: If things are to fit together as we wish to deliver a stable system for some initial value p_0 that we want to determine uniquely, δ cannot be too large.

The forward-looking equation (120.8) continues to describe equality between the demand and supply of money.

We assume that equations (120.8) and (120.24) govern $y_t \equiv \begin{bmatrix} m_t \\ p_t \end{bmatrix}$ for $t \geq 0$.

The transition matrix H in the law of motion

$$y_{t+1} = H y_t$$

now becomes

$$H = \begin{bmatrix} \rho & \delta \\ -(1-\lambda)/\lambda & \lambda^{-1} \end{bmatrix}.$$

We take m_0 as a given initial condition and as before seek an initial value p_0 that stabilizes the system in the sense that y_t converges as $t \to +\infty$.

Our approach is identical with the one followed above and is based on an eigenvalue decomposition in which, cross our fingers, one eigenvalue exceeds unity and the other is less than unity in absolute value.

When $\delta \neq 0$ as we now assume, the eigenvalues of H will no longer be $\rho \in (0, 1)$ and $\lambda^{-1} > 1$

We'll just calculate them and apply the same algorithm that we used above.

That algorithm remains valid so long as the eigenvalues split around unity as before.

Again we assume that m_0 is an initial condition, but that p_0 is not given but to be solved for.

Let's write and execute some Python code that will let us explore how outcomes depend on δ.

```
def construct_H(ρ, λ, δ):
    "contruct matrix H given parameters."

    H = np.empty((2, 2))
    H[0, :] = ρ,δ
    H[1, :] = - (1 - λ) / λ, 1 / λ

    return H

def H_eigvals(ρ=.9, λ=.5, δ=0):
    "compute the eigenvalues of matrix H given parameters."

    # construct H matrix
    H = construct_H(ρ, λ, δ)

    # compute eigenvalues
    eigvals = np.linalg.eigvals(H)

    return eigvals
```

```
H_eigvals()
```

```
array([2. , 0.9])
```

Notice that a negative δ will not imperil the stability of the matrix H, even if it has a big absolute value.

```
# small negative δ
H_eigvals(δ=-0.05)
```

```
array([0.8562829, 2.0437171])
```

```
# large negative δ
H_eigvals(δ=-1.5)
```

```
array([0.10742784, 2.79257216])
```

A sufficiently small positive δ also causes no problem.

```
# sufficiently small positive δ
H_eigvals(δ=0.05)
```

```
array([0.94750622, 1.95249378])
```

But a large enough positive δ makes both eigenvalues of H strictly greater than unity in modulus.

For example,

```
H_eigvals(δ=0.2)
```

```
array([1.12984379, 1.77015621])
```

We want to study systems in which one eigenvalue exceeds unity in modulus while the other is less than unity in modulus, so we avoid values of δ that are too.

That is, we want to avoid too much positive feedback from p_t to m_{t+1}.

```
def magic_p0(m0, ρ=.9, λ=.5, δ=0):
    """
    Use the magic formula (8) to compute the level of p0
    that makes the system stable.
    """

    H = construct_H(ρ, λ, δ)
    eigvals, Q = np.linalg.eig(H)

    # find the index of the smaller eigenvalue
    ind = 0 if eigvals[0] < eigvals[1] else 1

    # verify that the eigenvalue is less than unity
    if eigvals[ind] > 1:

        print("both eigenvalues exceed unity in modulus")

        return None
```

(continues on next page)

(continued from previous page)

```
    p0 = Q[1, ind] / Q[0, ind] * m0

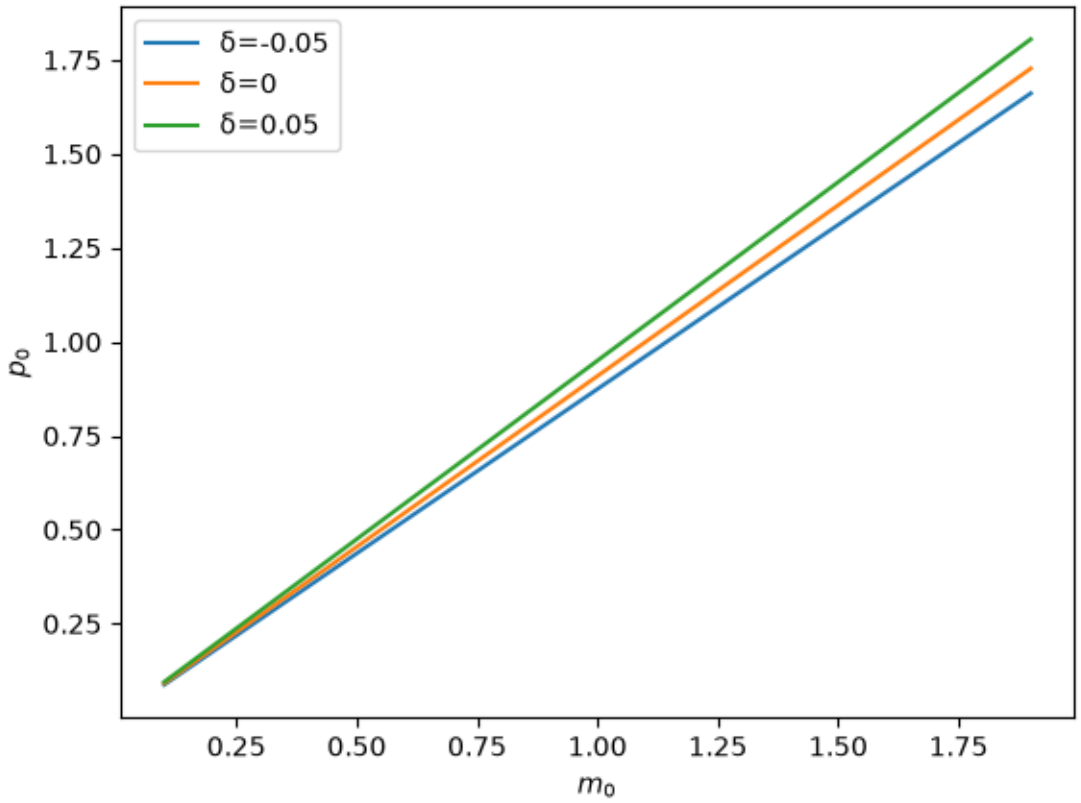
    return p0
```

Let's plot how the solution p_0 changes as m_0 changes for different settings of δ.

```
m_range = np.arange(0.1, 2., 0.1)

for δ in [-0.05, 0, 0.05]:
    plt.plot(m_range, [magic_p0(m0, δ=δ) for m0 in m_range], label=f"δ={δ}")
plt.legend()

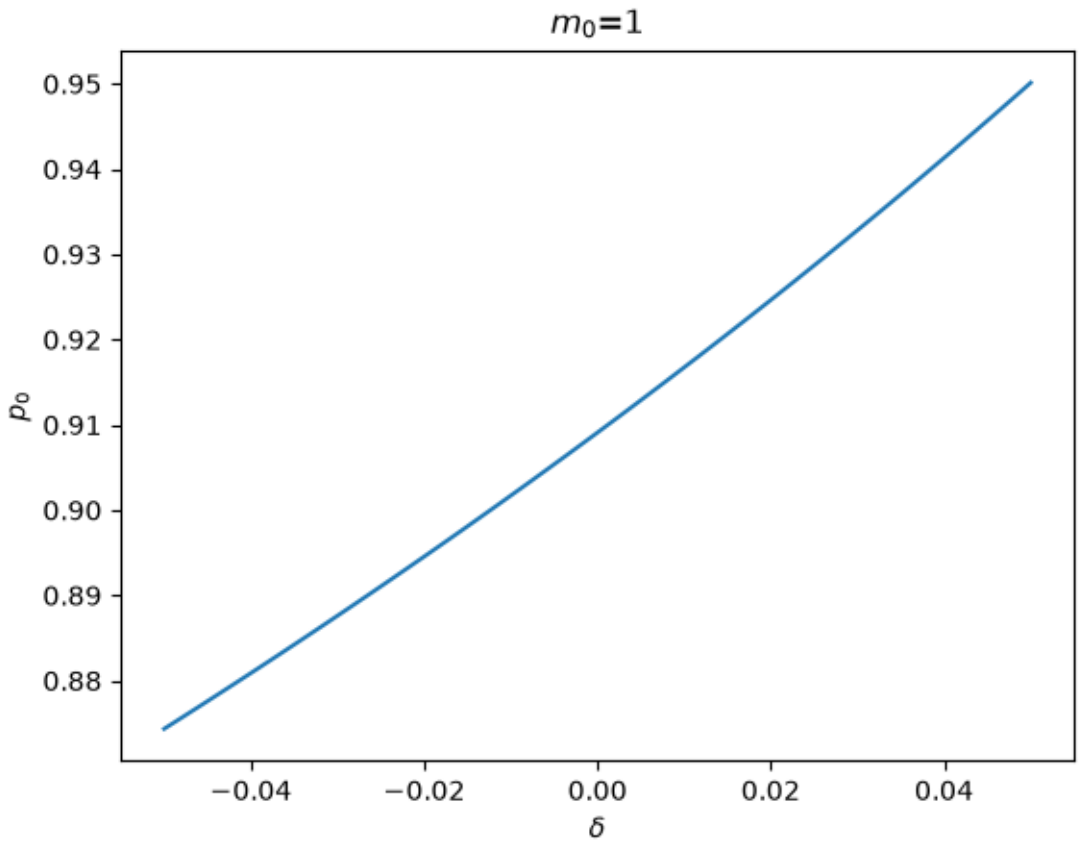
plt.xlabel(r"$m_0$")
plt.ylabel(r"$p_0$")
plt.show()
```



To look at things from a different angle, we can fix the initial value m_0 and see how p_0 changes as δ changes.

```
m0 = 1

δ_range = np.linspace(-0.05, 0.05, 100)
plt.plot(δ_range, [magic_p0(m0, δ=δ) for δ in δ_range])
plt.xlabel(r'$\delta$')
plt.ylabel(r'$p_0$')
plt.title(rf'$m_0$={m0}')
plt.show()
```



Notice that when δ is large enough, both eigenvalues exceed unity in modulus, causing a stabilizing value of p_0 not to exist.

```
magic_p0(1, δ=0.2)
```

```
both eigenvalues exceed unity in modulus
```

120.8 Big P, Little p Interpretation

It is helpful to view our solutions of difference equations having feedback from the price level or inflation to money or the rate of money creation in terms of the Big K, little k idea discussed in *Rational Expectations Models*.

This will help us sort out what is taken as given by the decision makers who use the difference equation (120.9) to determine p_t as a function of their forecasts of future values of m_t.

Let's write the stabilizing solution that we have computed using the eigenvector decomposition of H as $P_t = F^* m_t$, where

$$F^* = Q_{21} Q_{11}^{-1}.$$

Then from $P_{t+1} = F^* m_{t+1}$ and $m_{t+1} = \rho m_t + \delta P_t$ we can deduce the recursion $P_{t+1} = F^* \rho m_t + F^* \delta P_t$ and create the stacked system

$$\begin{bmatrix} m_{t+1} \\ P_{t+1} \end{bmatrix} = \begin{bmatrix} \rho & \delta \\ F^* \rho & F^* \delta \end{bmatrix} \begin{bmatrix} m_t \\ P_t \end{bmatrix}$$

or

$$x_{t+1} = Ax_t$$

where $x_t = \begin{bmatrix} m_t \\ P_t \end{bmatrix}$.

Apply formula (120.13) for F to deduce that

$$p_t = F \begin{bmatrix} m_t \\ P_t \end{bmatrix} = F \begin{bmatrix} m_t \\ F^* m_t \end{bmatrix}$$

which implies that

$$p_t = \begin{bmatrix} F_1 & F_2 \end{bmatrix} \begin{bmatrix} m_t \\ F^* m_t \end{bmatrix} = F_1 m_t + F_2 F^* m_t$$

so that we can anticipate that

$$F^* = F_1 + F_2 F^*$$

We shall verify this equality in the next block of Python code that implements the following computations.

1. For the system with $\delta \neq 0$ so that there is feedback, we compute the stabilizing solution for p_t in the form $p_t = F^* m_t$ where $F^* = Q_{21} Q_{11}^{-1}$ as above.
2. Recalling the system (120.11), (120.12), and (120.13) above, we define $x_t = \begin{bmatrix} m_t \\ P_t \end{bmatrix}$ and notice that it is Big P_t and not little p_t here. Then we form A and G as $A = \begin{bmatrix} \rho & \delta \\ F^* \rho & F^* \delta \end{bmatrix}$ and $G = \begin{bmatrix} 1 & 0 \end{bmatrix}$ and we compute $\begin{bmatrix} F_1 & F_2 \end{bmatrix} \equiv F$ from equation (120.13) above.
3. We compute $F_1 + F_2 F^*$ and compare it with F^* and check for the anticipated equality.

```
# set parameters
ρ = .9
λ = .5
δ = .05
```

```
# solve for F_star
H = construct_H(ρ, λ, δ)
eigvals, Q = np.linalg.eig(H)

ind = 0 if eigvals[0] < eigvals[1] else 1
F_star = Q[1, ind] / Q[0, ind]
F_star
```

```
0.950124378879 11
```

```
# solve for F_check
A = np.empty((2, 2))
A[0, :] = ρ, δ
A[1, :] = F_star * A[0, :]

G = np.array([1, 0])

F_check= (1 - λ) * G @ np.linalg.inv(np.eye(2) - λ * A)
F_check
```

```
array([0.92755597, 0.02375311])
```

Compare F^* with $F_1 + F_2 F^*$

```
F_check[0] + F_check[1] * F_star, F_star
```

$$(0.950124378879\,11,\ 0.950124378879\,11)$$

120.9 Fun with SymPy

This section is a gift for readers who have made it this far.

It puts SymPy to work on our model.

Thus, we use Sympy to compute some key objects comprising the eigenvector decomposition of H.

We start by generating an H with nonzero δ.

```
λ, δ, ρ = symbols('λ, δ, ρ')
```

```
H1 = Matrix([[ρ,δ], [- (1 - λ) / λ, λ ** -1]])
```

```
H1
```

$$\begin{bmatrix} \rho & \delta \\ \frac{\lambda-1}{\lambda} & \frac{1}{\lambda} \end{bmatrix}$$

```
H1.eigenvals()
```

$$\left\{(\lambda\rho + 1)/2\lambda - \sqrt{4\delta\lambda^2 - 4\delta\lambda + \lambda^2\rho^2 - 2\lambda\rho + 1}/2\lambda : 1,\ (\lambda\rho + 1)/2\lambda + \sqrt{4\delta\lambda^2 - 4\delta\lambda + \lambda^2\rho^2 - 2\lambda\rho + 1}/2\lambda : 1\right\}$$

```
H1.eigenvects()
```

$$\left[\left(\frac{\lambda\rho+1}{2\lambda} - \frac{\sqrt{4\delta\lambda^2 - 4\delta\lambda + \lambda^2\rho^2 - 2\lambda\rho + 1}}{2\lambda},\ 1,\ \left[\begin{bmatrix} \frac{\lambda\left(\frac{\lambda\rho+1}{2\lambda} - \frac{\sqrt{4\delta\lambda^2-4\delta\lambda+\lambda^2\rho^2-2\lambda\rho+1}}{2\lambda}\right)}{\lambda-1} - \frac{1}{\lambda-1} \\ 1 \end{bmatrix}\right]\right),\ \left(\frac{\lambda\rho+1}{2\lambda} + \frac{\sqrt{4\delta\lambda^2 - 4\delta\lambda}\right.$$

Now let's compute H when δ is zero.

```
H2 = Matrix([[ρ,0], [- (1 - λ) / λ, λ ** -1]])
```

```
H2
```

$$\begin{bmatrix} \rho & 0 \\ \frac{\lambda-1}{\lambda} & \frac{1}{\lambda} \end{bmatrix}$$

```
H2.eigenvals()
```

$$\{1/\lambda : 1,\ \rho : 1\}$$

```
H2.eigenvects()
```

$$\left[\left(\frac{1}{\lambda},\ 1,\ \left[\begin{bmatrix}0\\1\end{bmatrix}\right]\right),\ \left(\rho,\ 1,\ \left[\begin{bmatrix}\frac{\lambda\rho-1}{\lambda-1}\\1\end{bmatrix}\right]\right)\right]$$

Below we do induce SymPy to do the following fun things for us analytically:

1. We compute the matrix Q whose first column is the eigenvector associated with ρ. and whose second column is the eigenvector associated with λ^{-1}.
2. We use SymPy to compute the inverse Q^{-1} of Q (both in symbols).
3. We use SymPy to compute $Q_{21}Q_{11}^{-1}$ (in symbols).
4. Where Q^{ij} denotes the (i,j) component of Q^{-1}, we use SymPy to compute $-(Q^{22})^{-1}Q^{21}$ (again in symbols)

```
# construct Q
vec = []
for i, (eigval, _, eigvec) in enumerate(H2.eigenvects()):

    vec.append(eigvec[0])

    if eigval == ρ:
        ind = i

Q = vec[ind].col_insert(1, vec[1-ind])
```

```
Q
```

$$\begin{bmatrix}\frac{\lambda\rho-1}{\lambda-1} & 0\\ 1 & 1\end{bmatrix}$$

Q^{-1}

```
Q_inv = Q ** (-1)
Q_inv
```

$$\begin{bmatrix}\frac{\lambda-1}{\lambda\rho-1} & 0\\ \frac{1-\lambda}{\lambda\rho-1} & 1\end{bmatrix}$$

$Q_{21}Q_{11}^{-1}$

```
Q[1, 0] / Q[0, 0]
```

$$\frac{\lambda-1}{\lambda\rho-1}$$

$-(Q^{22})^{-1}Q^{21}$

```
- Q_inv[1, 0] / Q_inv[1, 1]
```

$$-\frac{1-\lambda}{\lambda\rho-1}$$

CHAPTER

ONE

MARKOV PERFECT EQUILIBRIUM

Contents

- *Markov Perfect Equilibrium*
  - *Overview*
  - *Background*
  - *Linear Markov Perfect Equilibria*
  - *Application*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

121.1 Overview

This lecture describes the concept of Markov perfect equilibrium.

Markov perfect equilibrium is a key notion for analyzing economic problems involving dynamic strategic interaction, and a cornerstone of applied game theory.

In this lecture, we teach Markov perfect equilibrium by example.

We will focus on settings with

- two players
- quadratic payoff functions
- linear transition rules for the state

Other references include chapter 7 of [Ljungqvist and Sargent, 2018].

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
```

121.2 Background

Markov perfect equilibrium is a refinement of the concept of Nash equilibrium.

It is used to study settings where multiple decision-makers interact non-cooperatively over time, each pursuing its own objective.

The agents in the model face a common state vector, the time path of which is influenced by – and influences – their decisions.

In particular, the transition law for the state that confronts each agent is affected by decision rules of other agents.

Individual payoff maximization requires that each agent solve a dynamic programming problem that includes this transition law.

Markov perfect equilibrium prevails when no agent wishes to revise its policy, taking as given the policies of all other agents.

Well known examples include

- Choice of price, output, location or capacity for firms in an industry (e.g., [Ericson and Pakes, 1995], [Ryan, 2012], [Doraszelski and Satterthwaite, 2010]).
- Rate of extraction from a shared natural resource, such as a fishery (e.g., [Levhari and Mirman, 1980], [Van Long, 2011]).

Let's examine a model of the first type.

121.2.1 Example: A Duopoly Model

Two firms are the only producers of a good, the demand for which is governed by a linear inverse demand function

$$p = a_0 - a_1(q_1 + q_2) \tag{121.1}$$

Here $p = p_t$ is the price of the good, $q_i = q_{it}$ is the output of firm $i = 1, 2$ at time t and $a_0 > 0, a_1 > 0$.

In (121.1) and what follows,

- the time subscript is suppressed when possible to simplify notation
- $\hat{x}$ denotes a next period value of variable x

Each firm recognizes that its output affects total output and therefore the market price.

The one-period payoff function of firm i is price times quantity minus adjustment costs:

$$\pi_i = pq_i - \gamma(\hat{q}_i - q_i)^2, \quad \gamma > 0, \tag{121.2}$$

Substituting the inverse demand curve (121.1) into (121.2) lets us express the one-period payoff as

$$\pi_i(q_i, q_{-i}, \hat{q}_i) = a_0 q_i - a_1 q_i^2 - a_1 q_i q_{-i} - \gamma(\hat{q}_i - q_i)^2, \tag{121.3}$$

where q_{-i} denotes the output of the firm other than i.

The objective of the firm is to maximize $\sum_{t=0}^{\infty} \beta^t \pi_{it}$.

Firm i chooses a decision rule that sets next period quantity $\hat{q}_i$ as a function f_i of the current state (q_i, q_{-i}).

An essential aspect of a Markov perfect equilibrium is that each firm takes the decision rule of the other firm as known and given.

Given f_{-i}, the Bellman equation of firm i is

$$v_i(q_i, q_{-i}) = \max_{\hat{q}_i} \{\pi_i(q_i, q_{-i}, \hat{q}_i) + \beta v_i(\hat{q}_i, f_{-i}(q_{-i}, q_i))\} \tag{121.4}$$

Definition 121.2.1 (Markov Perfect Equilibrium)

A **Markov perfect equilibrium** of the duopoly model is a pair of value functions (v_1, v_2) and a pair of policy functions (f_1, f_2) such that, for each $i \in \{1, 2\}$ and each possible state,

- The value function v_i satisfies Bellman equation (121.4).
- The maximizer on the right side of (121.4) equals $f_i(q_i, q_{-i})$.

The adjective "Markov" denotes that the equilibrium decision rules depend only on the current values of the state variables, not other parts of their histories.

"Perfect" means complete, in the sense that the equilibrium is constructed by backward induction and hence builds in optimizing behavior for each firm at all possible future states.

- These include many states that will not be reached when we iterate forward on the pair of equilibrium strategies f_i starting from a given initial state.

121.2.2 Computation

One strategy for computing a Markov perfect equilibrium is iterating to convergence on pairs of Bellman equations and decision rules.

In particular, let v_i^j, f_i^j be the value function and policy function for firm i at the j-th iteration.

Imagine constructing the iterates

$$v_i^{j+1}(q_i, q_{-i}) = \max_{\hat{q}_i} \left\{\pi_i(q_i, q_{-i}, \hat{q}_i) + \beta v_i^j(\hat{q}_i, f_{-i}(q_{-i}, q_i))\right\} \tag{121.5}$$

These iterations can be challenging to implement computationally.

However, they simplify for the case in which one-period payoff functions are quadratic and transition laws are linear — which takes us to our next topic.

121.3 Linear Markov Perfect Equilibria

As we saw in the duopoly example, the study of Markov perfect equilibria in games with two players leads us to an interrelated pair of Bellman equations.

In linear-quadratic dynamic games, these "stacked Bellman equations" become "stacked Riccati equations" with a tractable mathematical structure.

We'll lay out that structure in a general setup and then apply it to some simple problems.

121.3.1 Coupled Linear Regulator Problems

We consider a general linear-quadratic regulator game with two players.

For convenience, we'll start with a finite horizon formulation, where t_0 is the initial date and t_1 is the common terminal date.

Player i takes $\{u_{-it}\}$ as given and minimizes

$$\sum_{t=t_0}^{t_1-1} \beta^{t-t_0} \{x_t' R_i x_t + u_{it}' Q_i u_{it} + u_{-it}' S_i u_{-it} + 2x_t' W_i u_{it} + 2u_{-it}' M_i u_{it}\} \tag{121.6}$$

while the state evolves according to

$$x_{t+1} = Ax_t + B_1 u_{1t} + B_2 u_{2t} \tag{121.7}$$

Here

- x_t is an $n \times 1$ state vector and u_{it} is a $k_i \times 1$ vector of controls for player i
- R_i is $n \times n$
- S_i is $k_{-i} \times k_{-i}$
- Q_i is $k_i \times k_i$
- W_i is $n \times k_i$
- M_i is $k_{-i} \times k_i$
- A is $n \times n$
- B_i is $n \times k_i$

121.3.2 Computing Equilibrium

We formulate a linear Markov perfect equilibrium as follows.

Player i employs linear decision rules $u_{it} = -F_{it} x_t$, where F_{it} is a $k_i \times n$ matrix.

A Markov perfect equilibrium is a pair of sequences $\{F_{1t}, F_{2t}\}$ over $t = t_0, \ldots, t_1 - 1$ such that

- $\{F_{1t}\}$ solves player 1's problem, taking $\{F_{2t}\}$ as given, and
- $\{F_{2t}\}$ solves player 2's problem, taking $\{F_{1t}\}$ as given

If we take $u_{2t} = -F_{2t} x_t$ and substitute it into (121.6) and (121.7), then player 1's problem becomes minimization of

$$\sum_{t=t_0}^{t_1-1} \beta^{t-t_0} \{x_t' \Pi_{1t} x_t + u_{1t}' Q_1 u_{1t} + 2u_{1t}' \Gamma_{1t} x_t\} \tag{121.8}$$

subject to

$$x_{t+1} = \Lambda_{1t} x_t + B_1 u_{1t}, \tag{121.9}$$

where

- $\Lambda_{it} := A - B_{-i} F_{-it}$
- $\Pi_{it} := R_i + F_{-it}' S_i F_{-it}$
- $\Gamma_{it} := W_i' - M_i' F_{-it}$

This is an LQ dynamic programming problem that can be solved by working backwards.

Decision rules that solve this problem are

$$F_{1t} = (Q_1 + \beta B_1' P_{1t+1} B_1)^{-1} (\beta B_1' P_{1t+1} \Lambda_{1t} + \Gamma_{1t}) \tag{121.10}$$

where P_{1t} solves the matrix Riccati difference equation

$$P_{1t} = \Pi_{1t} - (\beta B_1' P_{1t+1} \Lambda_{1t} + \Gamma_{1t})' (Q_1 + \beta B_1' P_{1t+1} B_1)^{-1} (\beta B_1' P_{1t+1} \Lambda_{1t} + \Gamma_{1t}) + \beta \Lambda_{1t}' P_{1t+1} \Lambda_{1t} \tag{121.11}$$

Similarly, decision rules that solve player 2's problem are

$$F_{2t} = (Q_2 + \beta B_2' P_{2t+1} B_2)^{-1} (\beta B_2' P_{2t+1} \Lambda_{2t} + \Gamma_{2t}) \tag{121.12}$$

where P_{2t} solves

$$P_{2t} = \Pi_{2t} - (\beta B_2' P_{2t+1} \Lambda_{2t} + \Gamma_{2t})' (Q_2 + \beta B_2' P_{2t+1} B_2)^{-1} (\beta B_2' P_{2t+1} \Lambda_{2t} + \Gamma_{2t}) + \beta \Lambda_{2t}' P_{2t+1} \Lambda_{2t} \tag{121.13}$$

Here, in all cases $t = t_0, \ldots, t_1 - 1$ and the terminal conditions are $P_{it_1} = 0$.

The solution procedure is to use equations (121.10), (121.11), (121.12), and (121.13), and "work backwards" from time $t_1 - 1$.

Since we're working backward, P_{1t+1} and P_{2t+1} are taken as given at each stage.

Moreover, since

- some terms on the right-hand side of (121.10) contain F_{2t}
- some terms on the right-hand side of (121.12) contain F_{1t}

we need to solve these $k_1 + k_2$ equations simultaneously.

Key Insight

A key insight is that equations (121.10) and (121.12) are linear in F_{1t} and F_{2t}.

After these equations have been solved, we can take F_{it} and solve for P_{it} in (121.11) and (121.13).

Infinite Horizon

We often want to compute the solutions of such games for infinite horizons, in the hope that the decision rules F_{it} settle down to be time-invariant as $t_1 \to +\infty$.

In practice, we usually fix t_1 and compute the equilibrium of an infinite horizon game by driving $t_0 \to -\infty$.

This is the approach we adopt in the next section.

121.3.3 Implementation

We use the function nnash from QuantEcon.py that computes a Markov perfect equilibrium of the infinite horizon linear-quadratic dynamic game in the manner described above.

121.4 Application

Let's use these procedures to treat some applications, starting with the duopoly model.

121.4.1 A Duopoly Model

To map the duopoly model into coupled linear-quadratic dynamic programming problems, define the state and controls as

$$x_t := \begin{bmatrix} 1 \\ q_{1t} \\ q_{2t} \end{bmatrix} \quad \text{and} \quad u_{it} := q_{i,t+1} - q_{it}, \quad i = 1, 2$$

If we write

$$x_t' R_i x_t + u_{it}' Q_i u_{it}$$

where $Q_1 = Q_2 = \gamma$,

$$R_1 := \begin{bmatrix} 0 & -\frac{a_0}{2} & 0 \\ -\frac{a_0}{2} & a_1 & \frac{a_1}{2} \\ 0 & \frac{a_1}{2} & 0 \end{bmatrix} \quad \text{and} \quad R_2 := \begin{bmatrix} 0 & 0 & -\frac{a_0}{2} \\ 0 & 0 & \frac{a_1}{2} \\ -\frac{a_0}{2} & \frac{a_1}{2} & a_1 \end{bmatrix}$$

then we recover the one-period payoffs in expression (121.3).

The law of motion for the state x_t is $x_{t+1} = Ax_t + B_1 u_{1t} + B_2 u_{2t}$ where

$$A := \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad B_1 := \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \quad B_2 := \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

The optimal decision rule of firm i will take the form $u_{it} = -F_i x_t$, inducing the following closed-loop system for the evolution of x in the Markov perfect equilibrium:

$$x_{t+1} = (A - B_1 F_1 - B_1 F_2) x_t \tag{121.14}$$

121.4.2 Parameters and Solution

Consider the previously presented duopoly model with parameter values of:

- $a_0 = 10$
- $a_1 = 2$
- $\beta = 0.96$
- $\gamma = 12$

From these, we compute the infinite horizon MPE using the preceding code

```
import numpy as np
import quantecon as qe

# Parameters
a0 = 10.0
a1 = 2.0
```

(continues on next page)

(continued from previous page)

```
β = 0.96
γ = 12.0

# In LQ form
A = np.eye(3)
B1 = np.array([[0.], [1.], [0.]])
B2 = np.array([[0.], [0.], [1.]])


R1 = [[      0.,     -a0 / 2,           0.],
      [-a0 / 2.,           a1,     a1 / 2.],
      [       0,     a1 / 2.,           0.]]

R2 = [[     0.,           0.,      -a0 / 2],
      [     0.,           0.,      a1 / 2.],
      [-a0 / 2,      a1 / 2.,           a1]]

Q1 = Q2 = γ
S1 = S2 = W1 = W2 = M1 = M2 = 0.0

# Solve using QE's nnash function
F1, F2, P1, P2 = qe.nnash(A, B1, B2, R1, R2, Q1,
                          Q2, S1, S2, W1, W2, M1,
                          M2, beta=β)

# Display policies
print("Computed policies for firm 1 and firm 2:\n")
print(f"F1 = {F1}")
print(f"F2 = {F2}")
print("\n")
```

```
Computed policies for firm 1 and firm 2:

F1 = [[-0.66846615  0.29512482  0.07584666]]
F2 = [[-0.66846615  0.07584666  0.29512482]]
```

Running the code produces the following output.

One way to see that F_i is indeed optimal for firm i taking F_2 as given is to use QuantEcon.py's LQ class.

In particular, let's take F2 as computed above, plug it into (121.8) and (121.9) to get firm 1's problem and solve it using LQ.

We hope that the resulting policy will agree with F1 as computed above

```
beta = 0.96
Λ1 = A - B2 @ F2
lq1 = qe.LQ(Q1, R1, Λ1, B1, beta=beta)
P1_ih, F1_ih, d = lq1.stationary_values()
F1_ih
```

```
array([[-0.66846613,  0.29512482,  0.07584666]])
```

This is close enough for rock and roll, as they say in the trade.

Indeed, np.allclose agrees with our assessment

```
np.allclose(F1, F1_ih)
```

```
True
```

121.4.3 Dynamics

Let's now investigate the dynamics of price and output in this simple duopoly model under the MPE policies.

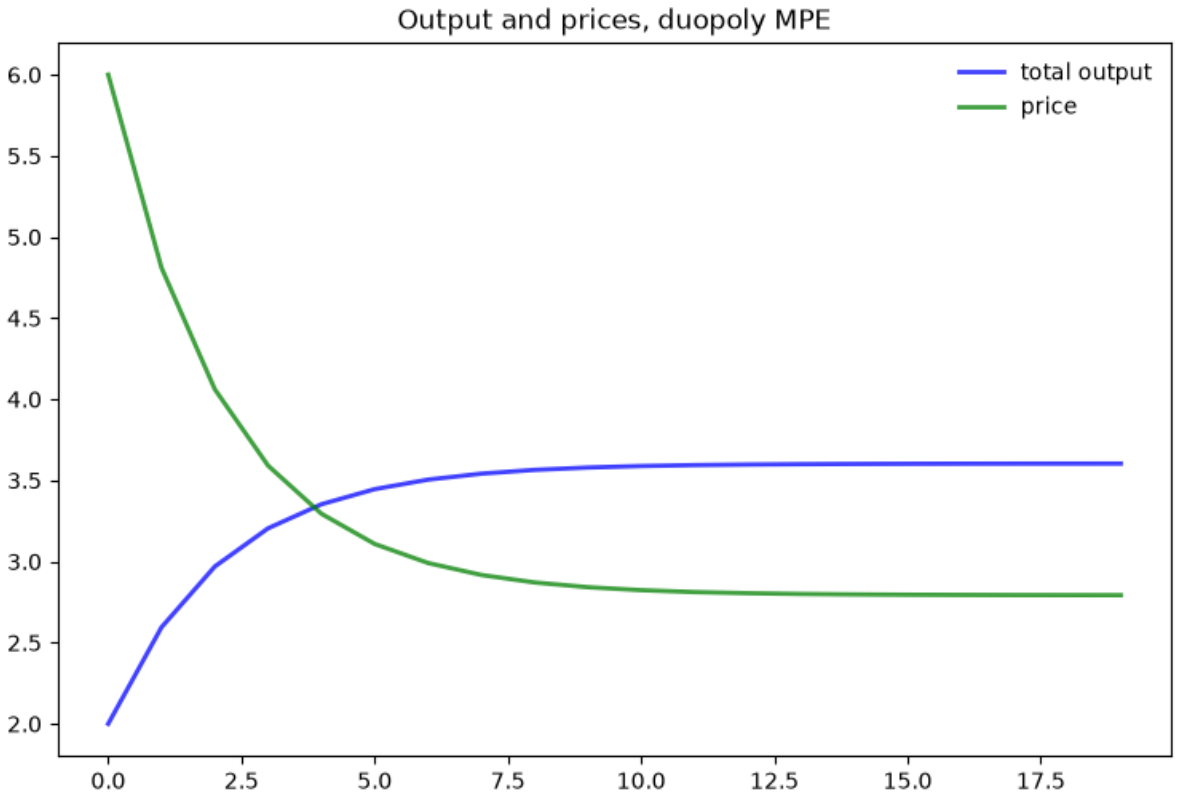
Given our optimal policies $F1$ and $F2$, the state evolves according to (121.14).

The following program

- imports $F1$ and $F2$ from the previous program along with all parameters.
- computes the evolution of x_t using (121.14).
- extracts and plots industry output $q_t = q_{1t} + q_{2t}$ and price $p_t = a_0 - a_1 q_t$.

```
AF = A - B1 @ F1 - B2 @ F2
n = 20
x = np.empty((3, n))
x[:, 0] = 1, 1, 1
for t in range(n-1):
    x[:, t+1] = AF @ x[:, t]
q1 = x[1, :]
q2 = x[2, :]
q = q1 + q2        # Total output, MPE
p = a0 - a1 * q    # Price, MPE

fig, ax = plt.subplots(figsize=(9, 5.8))
ax.plot(q, 'b-', lw=2, alpha=0.75, label='total output')
ax.plot(p, 'g-', lw=2, alpha=0.75, label='price')
ax.set_title('Output and prices, duopoly MPE')
ax.legend(frameon=False)
plt.show()
```



Note that the initial condition has been set to $q_{10} = q_{20} = 1.0$.

To gain some perspective we can compare this to what happens in the monopoly case.

The first panel in the next figure compares output of the monopolist and industry output under the MPE, as a function of time.

The second panel shows analogous curves for price.

Here parameters are the same as above for both the MPE and monopoly solutions.

The monopolist initial condition is $q_0 = 2.0$ to mimic the industry initial condition $q_{10} = q_{20} = 1.0$ in the MPE case.

As expected, output is higher and prices are lower under duopoly than monopoly.

121.5 Exercises

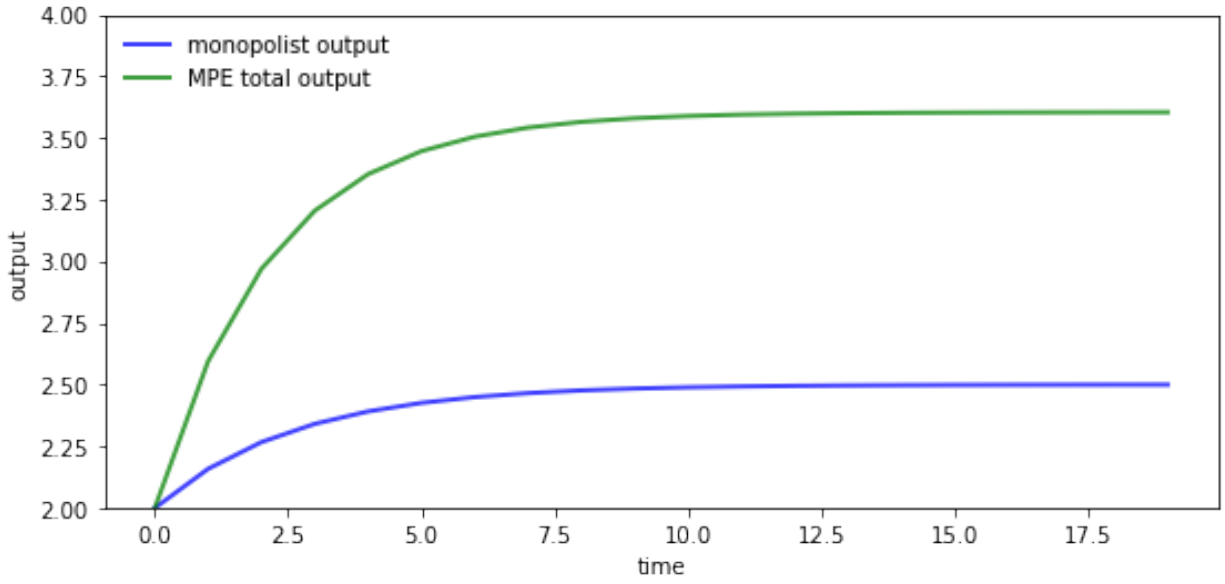
Exercise 121.5.1

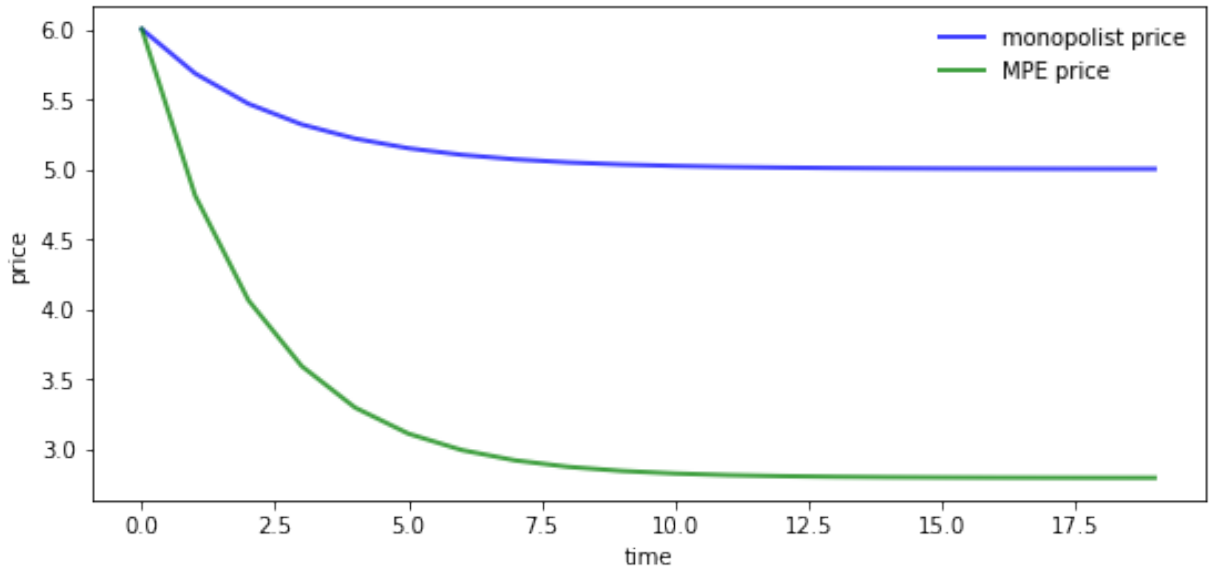
Replicate the *pair of figures* showing the comparison of output and prices for the monopolist and duopoly under MPE.

Parameters are as in duopoly_mpe.py and you can use that code to compute MPE policies under duopoly.

The optimal policy in the monopolist case can be computed using QuantEcon.py's LQ class.

Solution


monopolist output
MPE total output
output
time
4.00
3.75
3.50
3.25
3.00
2.75
2.50
2.25
2.00
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5


monopolist price
MPE price
price
time
6.0
5.5
5.0
4.5
4.0
3.5
3.0
0.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5

First, let's compute the duopoly MPE under the stated parameters

```
# == Parameters == #
a0 = 10.0
a1 = 2.0
β = 0.96
γ = 12.0

# == In LQ form == #
A  = np.eye(3)
B1 = np.array([[0.], [1.], [0.]])
B2 = np.array([[0.], [0.], [1.]])
R1 = [[      0.,      -a0/2,          0.],
      [-a0 / 2.,         a1,     a1 / 2.],
      [       0,    a1 / 2.,          0.]]

R2 = [[     0.,          0.,     -a0 / 2],
      [     0.,          0.,     a1 / 2.],
      [-a0 / 2,     a1 / 2.,          a1]]

Q1 = Q2 = γ
S1 = S2 = W1 = W2 = M1 = M2 = 0.0

# == Solve using QE's nnash function == #
F1, F2, P1, P2 = qe.nnash(A, B1, B2, R1, R2, Q1,
                          Q2, S1, S2, W1, W2, M1,
                          M2, beta=β)
```

Now we evaluate the time path of industry output and prices given initial condition $q_{10} = q_{20} = 1$.

```
AF = A - B1 @ F1 - B2 @ F2
n = 20
x = np.empty((3, n))
x[:, 0] = 1, 1, 1
for t in range(n-1):
    x[:, t+1] = AF @ x[:, t]
q1 = x[1, :]
q2 = x[2, :]
q = q1 + q2        # Total output, MPE
p = a0 - a1 * q    # Price, MPE
```

Next, let's have a look at the monopoly solution.

For the state and control, we take

$$x_t = q_t - \bar{q} \quad \text{and} \quad u_t = q_{t+1} - q_t$$

To convert to an LQ problem we set

$$R = a_1 \quad \text{and} \quad Q = \gamma$$

in the payoff function $x_t' R x_t + u_t' Q u_t$ and

$$A = B = 1$$

in the law of motion $x_{t+1} = A x_t + B u_t$.

We solve for the optimal policy $u_t = -F x_t$ and track the resulting dynamics of $\{q_t\}$, starting at $q_0 = 2.0$.

```
R = a1
Q = γ
```

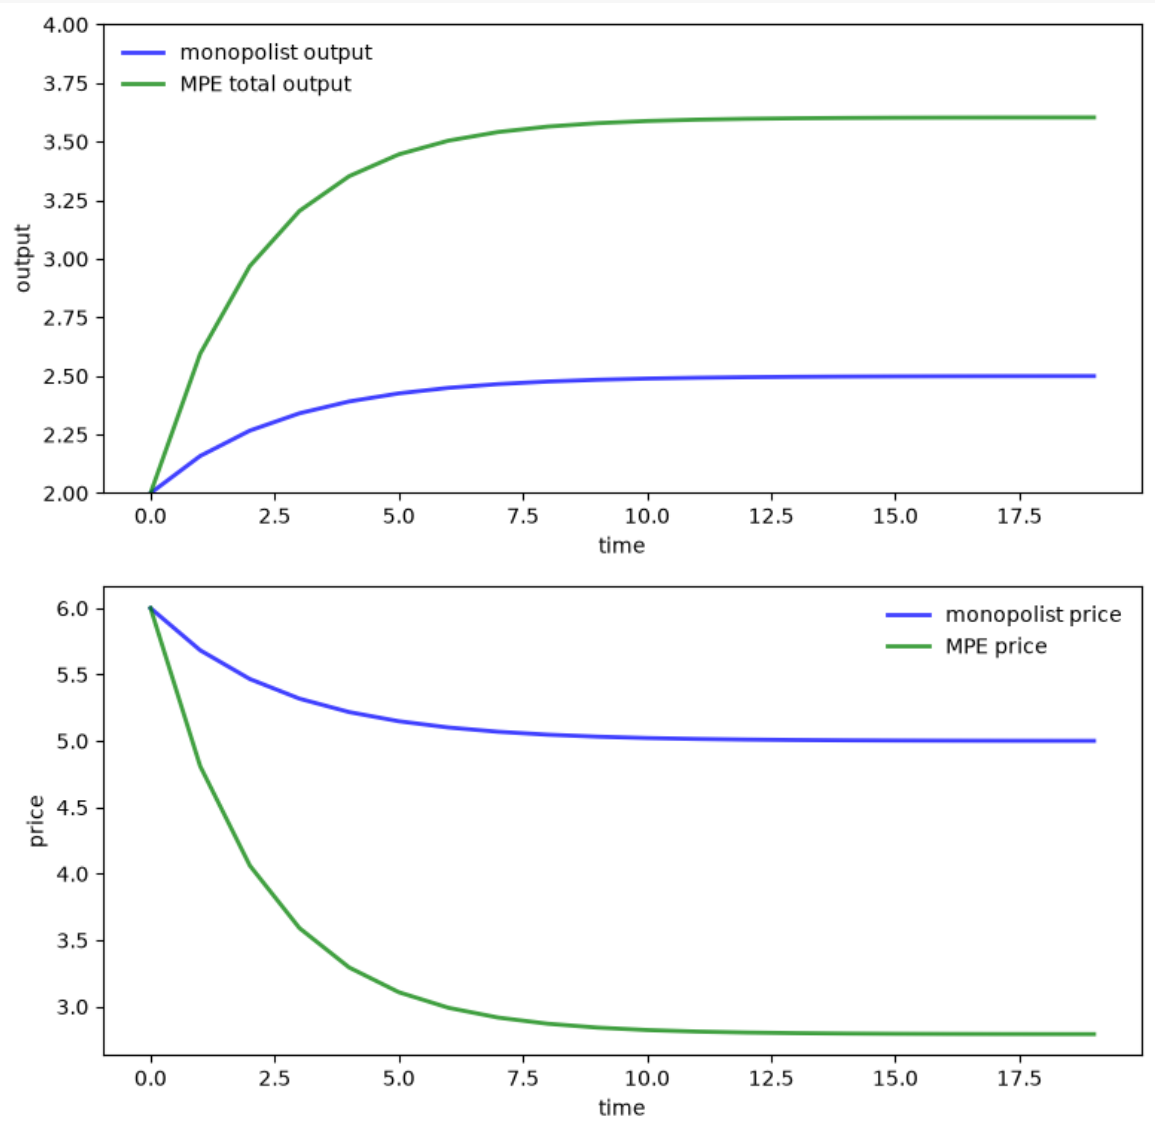
```
A = B = 1
lq_alt = qe.LQ(Q, R, A, B, beta=β)
P, F, d = lq_alt.stationary_values()
q_bar = a0 / (2.0 * a1)
qm = np.empty(n)
qm[0] = 2
x0 = qm[0] - q_bar
x = x0
for i in range(1, n):
    x = A * x - B * F * x
    qm[i] = float(x.item()) + q_bar
pm = a0 - a1 * qm
```

Let's have a look at the different time paths

```
fig, axes = plt.subplots(2, 1, figsize=(9, 9))

ax = axes[0]
ax.plot(qm, 'b-', lw=2, alpha=0.75, label='monopolist output')
ax.plot(q, 'g-', lw=2, alpha=0.75, label='MPE total output')
ax.set(ylabel="output", xlabel="time", ylim=(2, 4))
ax.legend(loc='upper left', frameon=0)

ax = axes[1]
ax.plot(pm, 'b-', lw=2, alpha=0.75, label='monopolist price')
ax.plot(p, 'g-', lw=2, alpha=0.75, label='MPE price')
ax.set(ylabel="price", xlabel="time")
ax.legend(loc='upper right', frameon=0)
plt.show()
```



Exercise 121.5.2

In this exercise, we consider a slightly more sophisticated duopoly problem.

It takes the form of infinite horizon linear-quadratic game proposed by Judd [Judd, 1990].

Two firms set prices and quantities of two goods interrelated through their demand curves.

Relevant variables are defined as follows:

- I_{it} = inventories of firm i at beginning of t
- q_{it} = production of firm i during period t
- p_{it} = price charged by firm i during period t
- S_{it} = sales made by firm i during period t
- E_{it} = costs of production of firm i during period t

- C_{it} = costs of carrying inventories for firm i during t

The firms' cost functions are

- $C_{it} = c_{i1} + c_{i2} I_{it} + 0.5 c_{i3} I_{it}^2$
- $E_{it} = e_{i1} + e_{i2} q_{it} + 0.5 e_{i3} q_{it}^2$ where e_{ij}, c_{ij} are positive scalars

Inventories obey the laws of motion

$$I_{i,t+1} = (1 - \delta) I_{it} + q_{it} - S_{it}$$

Demand is governed by the linear schedule

$$S_t = D p_{it} + b$$

where

- $S_t = \begin{bmatrix} S_{1t} & S_{2t} \end{bmatrix}'$
- D is a 2×2 negative definite matrix and
- b is a vector of constants

Firm i maximizes the undiscounted sum

$$\lim_{T \to \infty} \frac{1}{T} \sum_{t=0}^{T} (p_{it} S_{it} - E_{it} - C_{it})$$

We can convert this to a linear-quadratic problem by taking

$$u_{it} = \begin{bmatrix} p_{it} \\ q_{it} \end{bmatrix} \quad \text{and} \quad x_t = \begin{bmatrix} I_{1t} \\ I_{2t} \\ 1 \end{bmatrix}$$

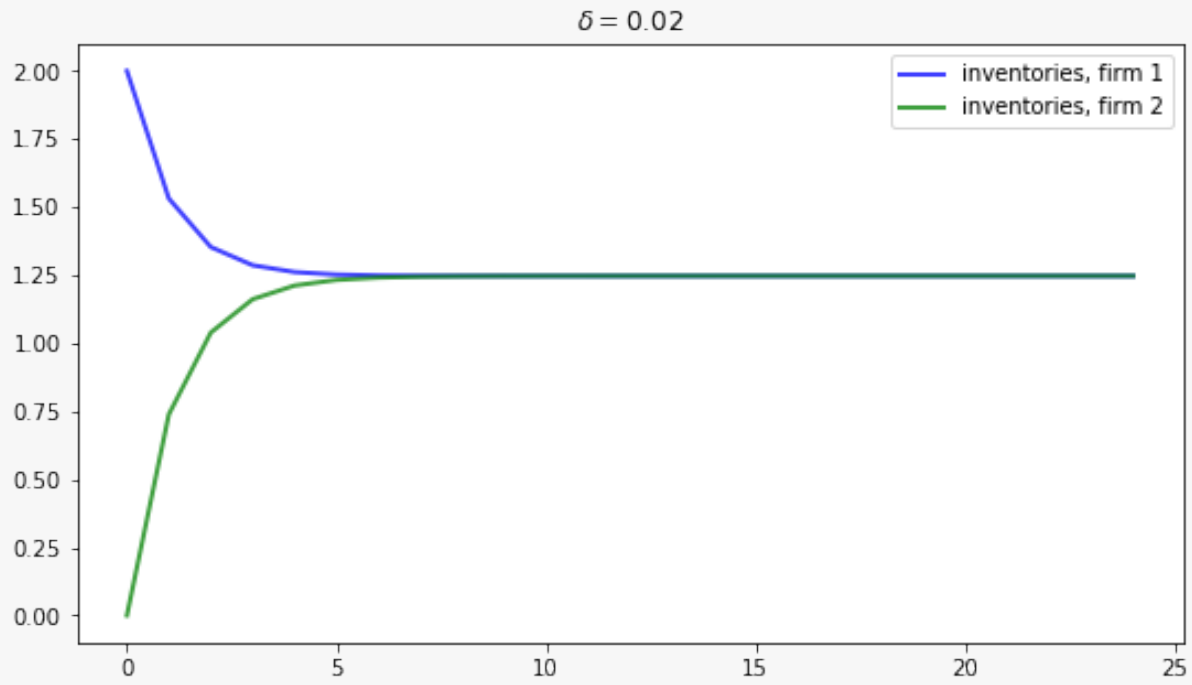
Decision rules for price and quantity take the form $u_{it} = -F_i x_t$.

The Markov perfect equilibrium of Judd's model can be computed by filling in the matrices appropriately.

The exercise is to calculate these matrices and compute the following figures.

The first figure shows the dynamics of inventories for each firm when the parameters are

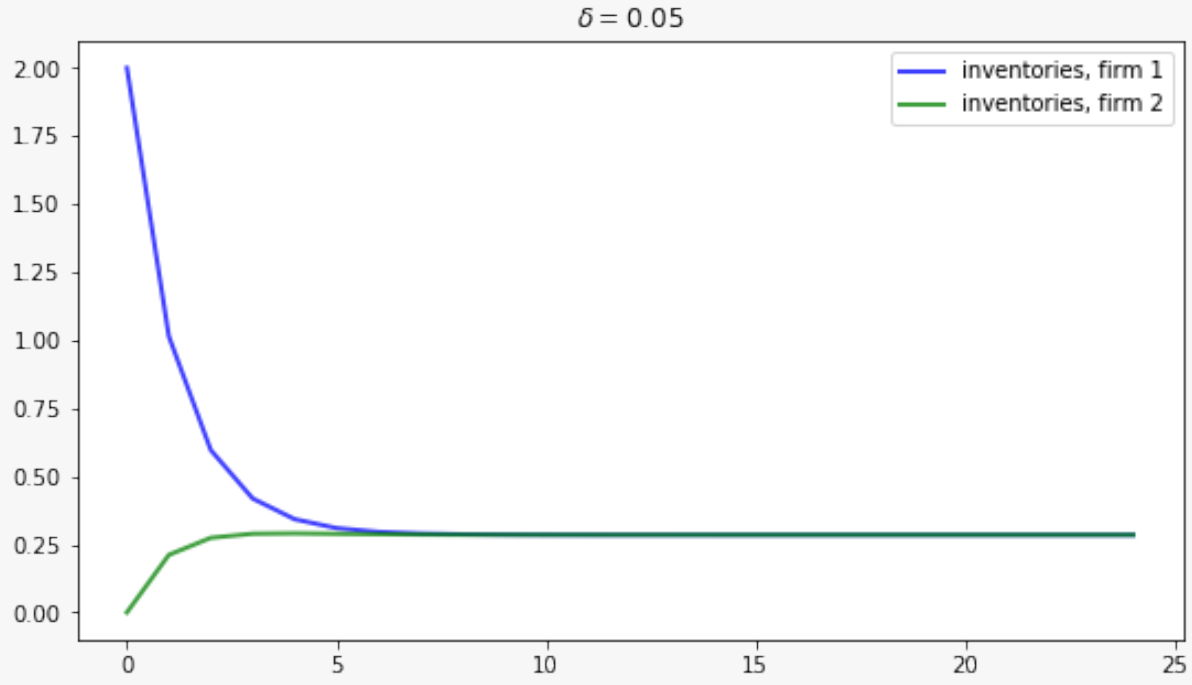
```
δ = 0.02
D = np.array([[-1, 0.5], [0.5, -1]])
b = np.array([25, 25])
c1 = c2 = np.array([1, -2, 1])
e1 = e2 = np.array([10, 10, 3])
```



Inventories trend to a common steady state.

If we increase the depreciation rate to $\delta = 0.05$, then we expect steady state inventories to fall.

This is indeed the case, as the next figure shows



In this exercise, reproduce the figure when $\delta = 0.02$.

Solution

We treat the case $\delta = 0.02$

```
δ = 0.02
D = np.array([[-1, 0.5], [0.5, -1]])
b = np.array([25, 25])
c1 = c2 = np.array([1, -2, 1])
e1 = e2 = np.array([10, 10, 3])

δ_1 = 1 - δ
```

Recalling that the control and state are

$$u_{it} = \begin{bmatrix} p_{it} \\ q_{it} \end{bmatrix} \quad \text{and} \quad x_t = \begin{bmatrix} I_{1t} \\ I_{2t} \\ 1 \end{bmatrix}$$

we set up the matrices as follows:

```
# ==  Create matrices needed to compute the Nash feedback equilibrium == #

A = np.array([[δ_1,      0,     -δ_1 * b[0]],
              [  0,    δ_1,     -δ_1 * b[1]],
              [  0,      0,               1]])

B1 = δ_1 * np.array([[1, -D[0, 0]],
                     [0, -D[1, 0]],
                     [0,       0]])
B2 = δ_1 * np.array([[0, -D[0, 1]],
                     [1, -D[1, 1]],
                     [0,       0]])

R1 = -np.array([[0.5 * c1[2],      0,     0.5 * c1[1]],
                [          0,      0,               0],
                [0.5 * c1[1],      0,          c1[0]]])
R2 = -np.array([[0,                0,               0],
                [0,      0.5 * c2[2],     0.5 * c2[1]],
                [0,      0.5 * c2[1],          c2[0]]])

Q1 = np.array([[-0.5 * e1[2], 0], [0, D[0, 0]]])
Q2 = np.array([[-0.5 * e2[2], 0], [0, D[1, 1]]])

S1 = np.zeros((2, 2))
S2 = np.copy(S1)

W1 = np.array([[           0,              0],
               [           0,              0],
               [-0.5 * e1[1],     b[0] / 2.]])
W2 = np.array([[           0,              0],
               [           0,              0],
               [-0.5 * e2[1],     b[1] / 2.]])

M1 = np.array([[0, 0], [0, D[0, 1] / 2.]])
M2 = np.copy(M1)
```

We can now compute the equilibrium using `qe.nnash`

```
F1, F2, P1, P2 = qe.nnash(A, B1, B2, R1,
                          R2, Q1, Q2, S1,
                          S2, W1, W2, M1, M2)

print("\nFirm 1's feedback rule:\n")
print(F1)

print("\nFirm 2's feedback rule:\n")
print(F2)
```

```
Firm 1's feedback rule:

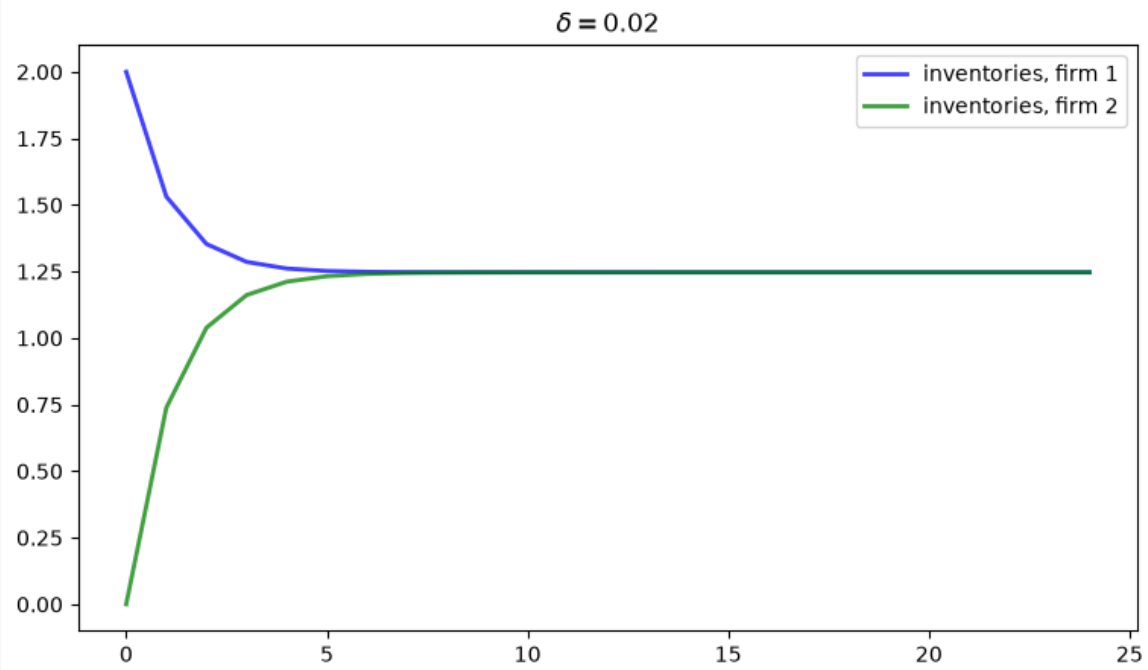
[[ 2.43666582e-01  2.72360627e-02 -6.82788293e+00]
 [ 3.92370734e-01  1.39696451e-01 -3.77341073e+01]]

Firm 2's feedback rule:

[[ 2.72360627e-02  2.43666582e-01 -6.82788293e+00]
 [ 1.39696451e-01  3.92370734e-01 -3.77341073e+01]]
```

Now let's look at the dynamics of inventories, and reproduce the graph corresponding to $\delta = 0.02$

```
AF = A - B1 @ F1 - B2 @ F2
n = 25
x = np.empty((3, n))
x[:, 0] = 2, 0, 1
for t in range(n-1):
    x[:, t+1] = AF @ x[:, t]
I1 = x[0, :]
I2 = x[1, :]
fig, ax = plt.subplots(figsize=(9, 5))
ax.plot(I1, 'b-', lw=2, alpha=0.75, label='inventories, firm 1')
ax.plot(I2, 'g-', lw=2, alpha=0.75, label='inventories, firm 2')
ax.set_title(rf'$\delta = {δ}$')
ax.legend()
plt.show()
```


$\delta = 0.02$
inventories, firm 1
inventories, firm 2
2.00
1.75
1.50
1.25
1.00
0.75
0.50
0.25
0.00
0
5
10
15
20
25

CHAPTER

TWO

UNCERTAINTY TRAPS

Contents

- *Uncertainty Traps*
  - *Overview*
  - *The Model*
  - *Implementation*
  - *Results*
  - *Exercises*

122.1 Overview

In this lecture, we study a simplified version of an uncertainty traps model of Fajgelbaum, Schaal and Taschereau-Dumouchel [Fajgelbaum *et al.*, 2015].

The model features self-reinforcing uncertainty that has big impacts on economic activity.

In the model,

- Fundamentals vary stochastically and are not fully observable.
- At any moment there are both active and inactive entrepreneurs; only active entrepreneurs produce.
- Agents – active and inactive entrepreneurs – have beliefs about the fundamentals expressed as probability distributions.
- Greater uncertainty means greater dispersions of these distributions.
- Entrepreneurs are risk-averse and hence less inclined to be active when uncertainty is high.
- The output of active entrepreneurs is observable, supplying a noisy signal that helps everyone inside the model infer fundamentals.
- Entrepreneurs update their beliefs about fundamentals using Bayes' Law, implemented via *Kalman filtering*.

Uncertainty traps emerge because:

- High uncertainty discourages entrepreneurs from becoming active.
- A low level of participation – i.e., a smaller number of active entrepreneurs – diminishes the flow of information about fundamentals.

- Less information translates to higher uncertainty, further discouraging entrepreneurs from choosing to be active, and so on.

Uncertainty traps stem from a positive externality: high aggregate economic activity levels generates valuable information.

Let's start with some standard imports:

```
import matplotlib.pyplot as plt
import numpy as np
```

122.2 The Model

The original model described in [Fajgelbaum *et al.*, 2015] has many interesting moving parts.

Here we examine a simplified version that nonetheless captures many of the key ideas.

122.2.1 Fundamentals

The evolution of the fundamental process $\{\theta_t\}$ is given by

$$\theta_{t+1} = \rho\theta_t + \sigma_\theta w_{t+1}$$

where

- $\sigma_\theta > 0$ and $0 < \rho < 1$
- $\{w_t\}$ is IID and standard normal

The random variable θ_t is not observable at any time.

122.2.2 Output

There is a total $\bar{M}$ of risk-averse entrepreneurs.

Output of the m-th entrepreneur, conditional on being active in the market at time t, is equal to

$$x_m = \theta + \epsilon_m \quad \text{where} \quad \epsilon_m \sim N\left(0, \gamma_x^{-1}\right) \tag{122.1}$$

Here the time subscript has been dropped to simplify notation.

The inverse of the shock variance, γ_x, is called the shock's **precision**.

The higher is the precision, the more informative x_m is about the fundamental.

Output shocks are independent across time and firms.

122.2.3 Information and Beliefs

All entrepreneurs start with identical beliefs about θ_0.

Signals are publicly observable and hence all agents have identical beliefs always.

Dropping time subscripts, beliefs for current θ are represented by the normal distribution $N(\mu, \gamma^{-1})$.

Here γ is the precision of beliefs; its inverse is the degree of uncertainty.

These parameters are updated by Kalman filtering.

Let

- $\mathbb{M} \subset \{1, \ldots, \bar{M}\}$ denote the set of currently active firms.
- $M := |\mathbb{M}|$ denote the number of currently active firms.
- X be the average output $\frac{1}{M}\sum_{m \in \mathbb{M}} x_m$ of the active firms.

With this notation and primes for next period values, we can write the updating of the mean and precision via

$$\mu' = \rho \frac{\gamma\mu + M\gamma_x X}{\gamma + M\gamma_x} \tag{122.2}$$

$$\gamma' = \left(\frac{\rho^2}{\gamma + M\gamma_x} + \sigma_\theta^2 \right)^{-1} \tag{122.3}$$

These are standard Kalman filtering results applied to the current setting.

Exercise 1 provides more details on how (122.2) and (122.3) are derived and then asks you to fill in remaining steps.

The next figure plots the law of motion for the precision in (122.3) as a 45 degree diagram, with one curve for each $M \in \{0, \ldots, 6\}$.

The other parameter values are $\rho = 0.99, \gamma_x = 0.5, \sigma_\theta = 0.5$

Points where the curves hit the 45 degree lines are long-run steady states for precision for different values of M.

Thus, if one of these values for M remains fixed, a corresponding steady state is the equilibrium level of precision

- high values of M correspond to greater information about the fundamental, and hence more precision in steady state
- low values of M correspond to less information and more uncertainty in steady state

In practice, as we'll see, the number of active firms fluctuates stochastically.

122.2.4 Participation

Omitting time subscripts once more, entrepreneurs enter the market in the current period if

$$\mathbb{E}[u(x_m - F_m)] > c \tag{122.4}$$

Here

- the mathematical expectation of x_m is based on (122.1) and beliefs $N(\mu, \gamma^{-1})$ for θ
- F_m is a stochastic but pre-visible fixed cost, independent across time and firms
- c is a constant reflecting opportunity costs

The statement that F_m is pre-visible means that it is realized at the start of the period and treated as a constant in (122.4).

The utility function has the constant absolute risk aversion form

$$u(x) = \frac{1}{a}\left(1 - \exp(-ax)\right) \tag{122.5}$$

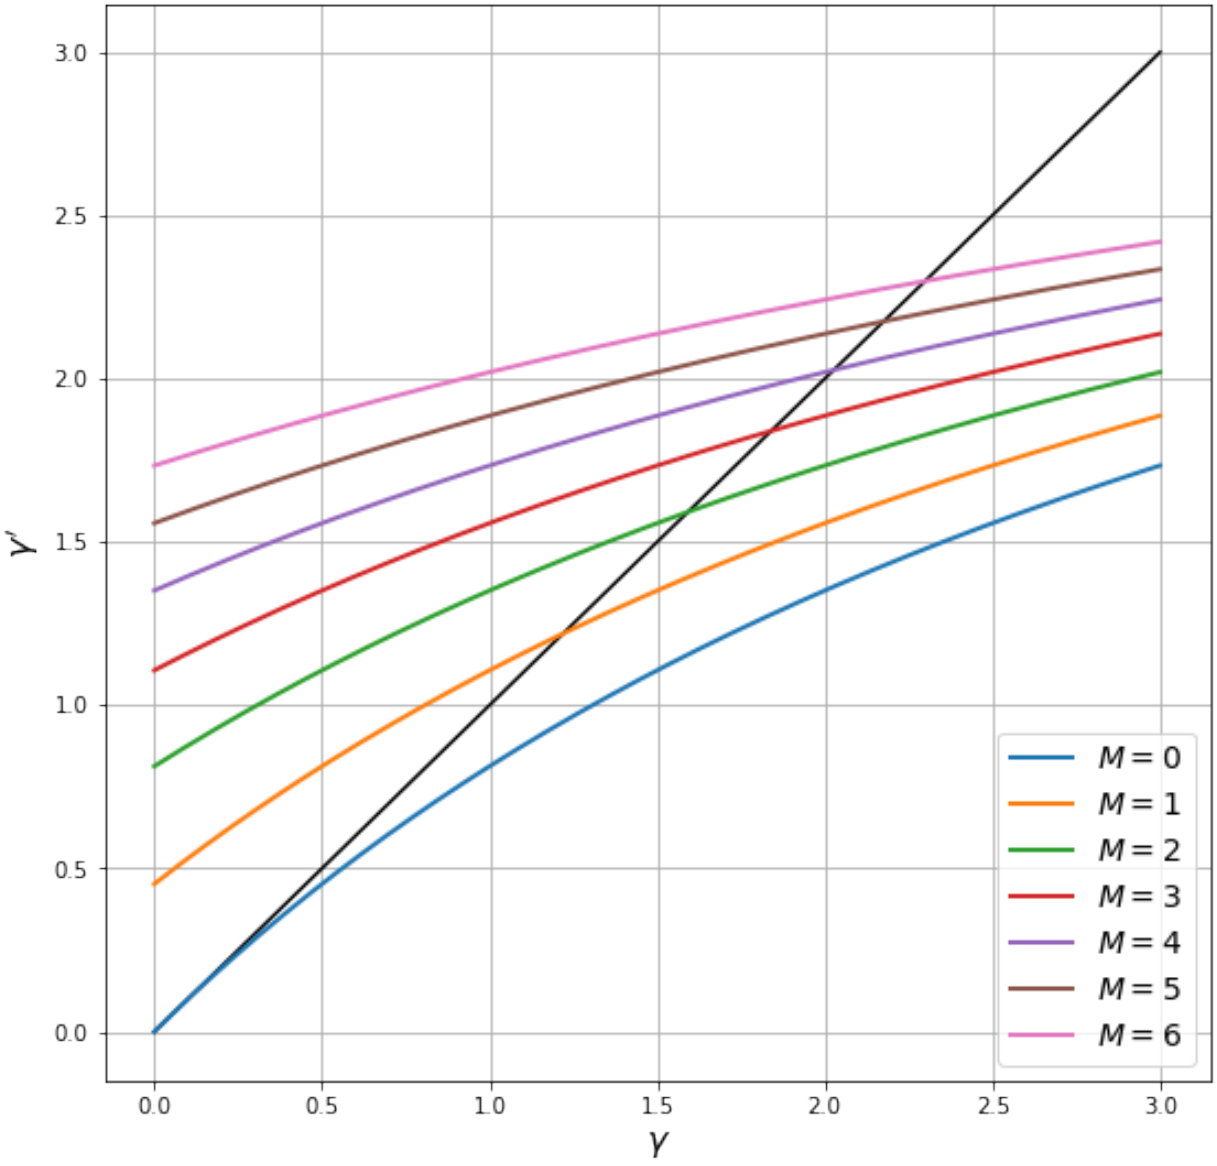
where a is a positive parameter.

Combining (122.4) and (122.5), entrepreneur m participates in the market (or is said to be active) when

$$\frac{1}{a}\left\{1 - \mathbb{E}[\exp\left(-a(\theta + \epsilon_m - F_m)\right)]\right\} > c$$

Using standard formulas for expectations of lognormal random variables, this is equivalent to the condition

$$\psi(\mu, \gamma, F_m) := \frac{1}{a}\left(1 - \exp\left(-a\mu + aF_m + \frac{a^2\left(\frac{1}{\gamma} + \frac{1}{\gamma_x}\right)}{2}\right)\right) - c > 0 \tag{122.6}$$


3.0
2.5
2.0
1.5
1.0
0.5
0.0
y'
0.0
0.5
1.0
1.5
2.0
2.5
3.0
y
$M = 0$
$M = 1$
$M = 2$
$M = 3$
$M = 4$
$M = 5$
$M = 6$

122.3 Implementation

We want to simulate this economy.

As a first step, let's put together a class that bundles

- the parameters, the current value of θ and the current values of the two belief parameters μ and γ
- methods to update θ, μ and γ, as well as to determine the number of active firms and their outputs

The updating methods follow the laws of motion for θ, μ and γ given above.

The method to evaluate the number of active firms generates $F_1, \ldots, F_{\bar{M}}$ and tests condition (122.6) for each firm.

The **init** method encodes as default values the parameters we'll use in the simulations below

```
class UncertaintyTrapEcon:

    def __init__(self,
                 a=1.5,           # Risk aversion
                 γ_x=0.5,         # Production shock precision
                 ρ=0.99,          # Correlation coefficient for θ
                 σ_θ=0.5,         # Standard dev of θ shock
                 num_firms=100,   # Number of firms
                 σ_F=1.5,         # Standard dev of fixed costs
                 c=-420,          # External opportunity cost
                 μ_init=0,        # Initial value for μ
                 γ_init=4,        # Initial value for γ
                 θ_init=0):       # Initial value for θ

        # == Record values == #
        self.a, self.γ_x, self.ρ, self.σ_θ = a, γ_x, ρ, σ_θ
        self.num_firms, self.σ_F, self.c, = num_firms, σ_F, c
        self.σ_x = np.sqrt(1/γ_x)

        # == Initialize states == #
        self.γ, self.μ, self.θ = γ_init, μ_init, θ_init

    def ψ(self, F):
        temp1 = -self.a * (self.μ - F)
        temp2 = self.a**2 * (1/self.γ + 1/self.γ_x) / 2
        return (1 / self.a) * (1 - np.exp(temp1 + temp2)) - self.c

    def update_beliefs(self, X, M):
        """
        Update beliefs (μ, γ) based on aggregates X and M.
        """
        # Simplify names
        γ_x, ρ, σ_θ = self.γ_x, self.ρ, self.σ_θ
        # Update μ
        temp1 = ρ * (self.γ * self.μ + M * γ_x * X)
        temp2 = self.γ + M * γ_x
        self.μ = temp1 / temp2
        # Update γ
        self.γ = 1 / (ρ**2 / (self.γ + M * γ_x) + σ_θ**2)

    def update_θ(self, w):
        """
        Update the fundamental state θ given shock w.
```

(continues on next page)

(continued from previous page)

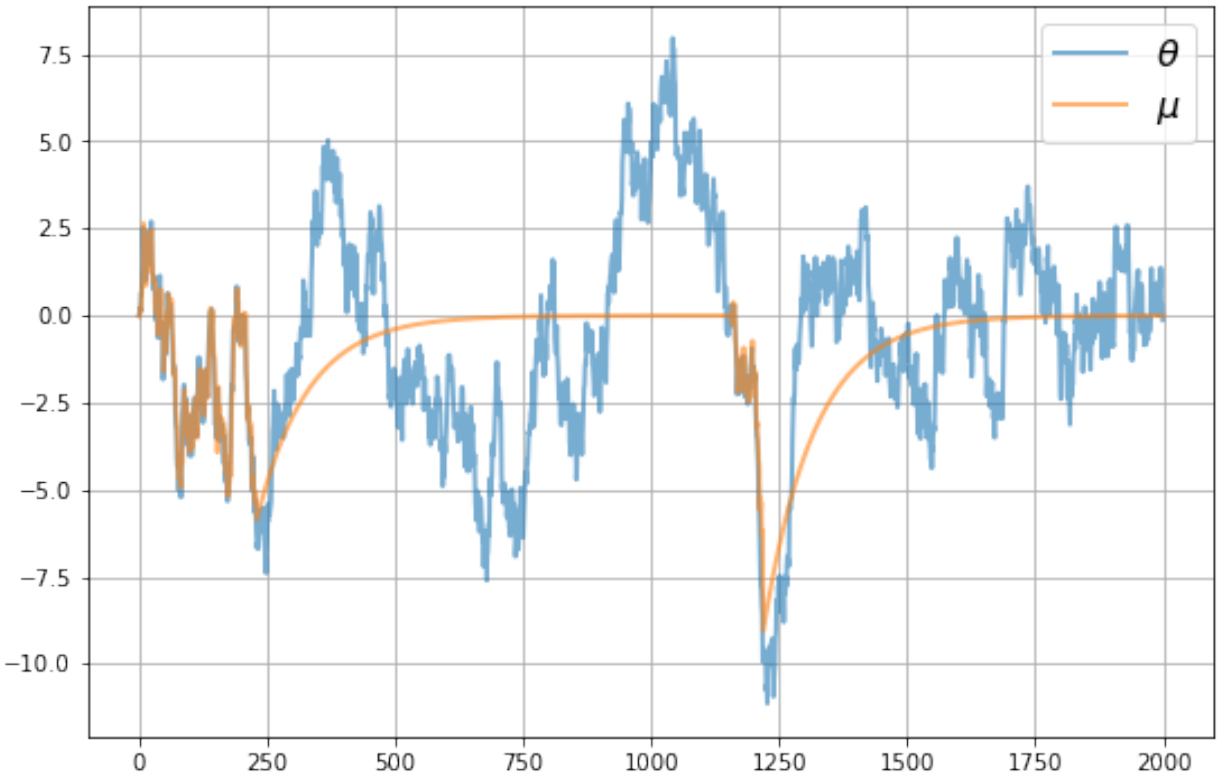
```
        """
        self.θ = self.ρ * self.θ + self.σ_θ * w

    def gen_aggregates(self, rng):
        """
        Generate aggregates based on current beliefs (μ, γ). This
        is a simulation step that depends on the draws for F.
        """
        F_vals = self.σ_F * rng.standard_normal(self.num_firms)
        M = np.sum(self.ψ(F_vals) > 0)  # Counts number of active firms
        if M > 0:
            x_vals = self.θ + self.σ_x * rng.standard_normal(M)
            X = x_vals.mean()
        else:
            X = 0
        return X, M
```

In the results below we use this code to simulate time series for the major variables.

122.4 Results

Let's look first at the dynamics of μ, which the agents use to track θ



We see that μ tracks θ well when there are sufficient firms in the market.

However, there are times when μ tracks θ poorly due to insufficient information.

These are episodes where the uncertainty traps take hold.

During these episodes

- precision is low and uncertainty is high
- few firms are in the market

To get a clearer idea of the dynamics, let's look at all the main time series at once, for a given set of shocks

Notice how the traps only take hold after a sequence of bad draws for the fundamental.

Thus, the model gives us a **propagation mechanism** that maps bad random draws into long downturns in economic activity.

122.5 Exercises

Exercise 122.5.1

Fill in the details behind (122.2) and (122.3) based on the following standard result (see, e.g., p. 24 of [Young and Smith, 2005]).

Fact Let $\mathbf{x} = (x_1, \ldots, x_M)$ be a vector of IID draws from common distribution $N(\theta, 1/\gamma_x)$ and let $\bar{x}$ be the sample mean. If γ_x is known and the prior for θ is $N(\mu, 1/\gamma)$, then the posterior distribution of θ given $\mathbf{x}$ is

$$\pi(\theta \,|\, \mathbf{x}) = N(\mu_0, 1/\gamma_0)$$

where

$$\mu_0 = \frac{\mu\gamma + M\bar{x}\gamma_x}{\gamma + M\gamma_x} \quad \text{and} \quad \gamma_0 = \gamma + M\gamma_x$$

Solution

This exercise asked you to validate the laws of motion for γ and μ given in the lecture, based on the stated result about Bayesian updating in a scalar Gaussian setting. The stated result tells us that after observing average output X of the M firms, our posterior beliefs will be

$$N(\mu_0, 1/\gamma_0)$$

where

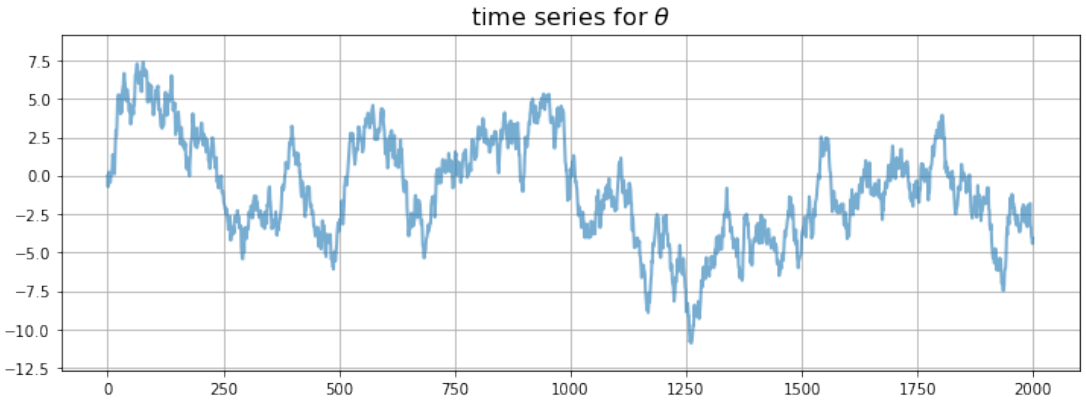
$$\mu_0 = \frac{\mu\gamma + MX\gamma_x}{\gamma + M\gamma_x} \quad \text{and} \quad \gamma_0 = \gamma + M\gamma_x$$

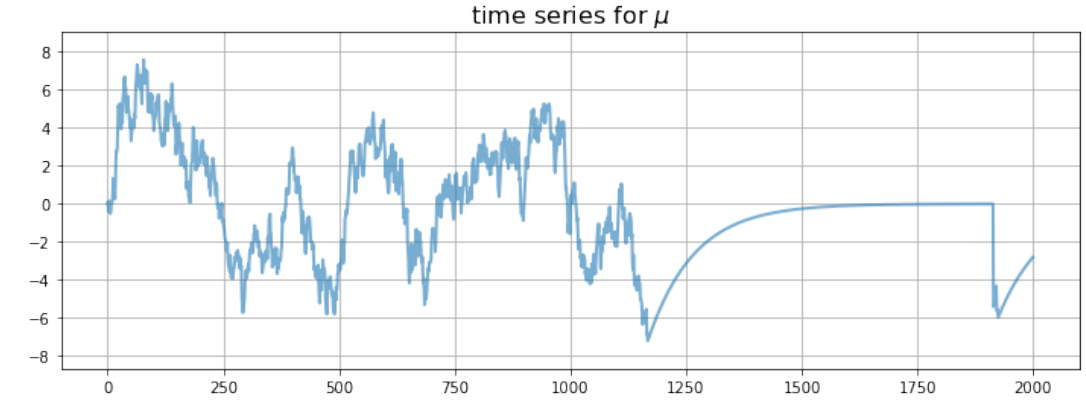
If we take a random variable θ with this distribution and then evaluate the distribution of $\rho\theta + \sigma_\theta w$ where w is independent and standard normal, we get the expressions for μ' and γ' given in the lecture.

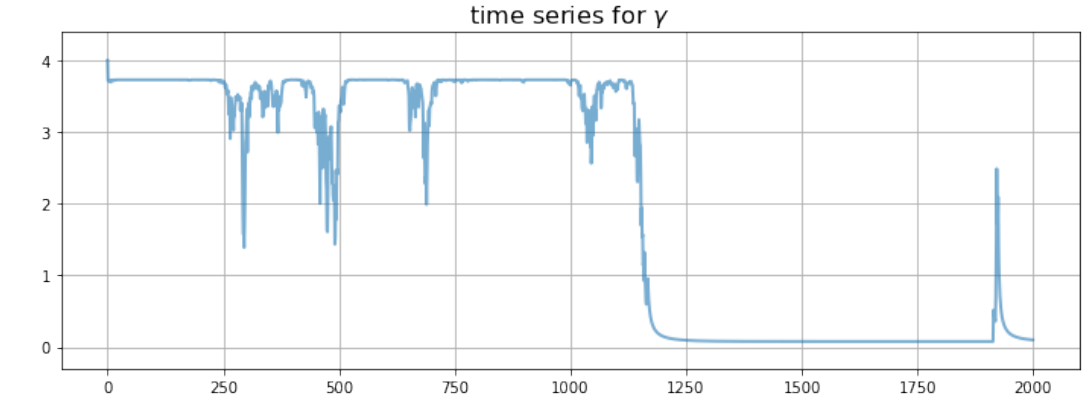
Exercise 122.5.2

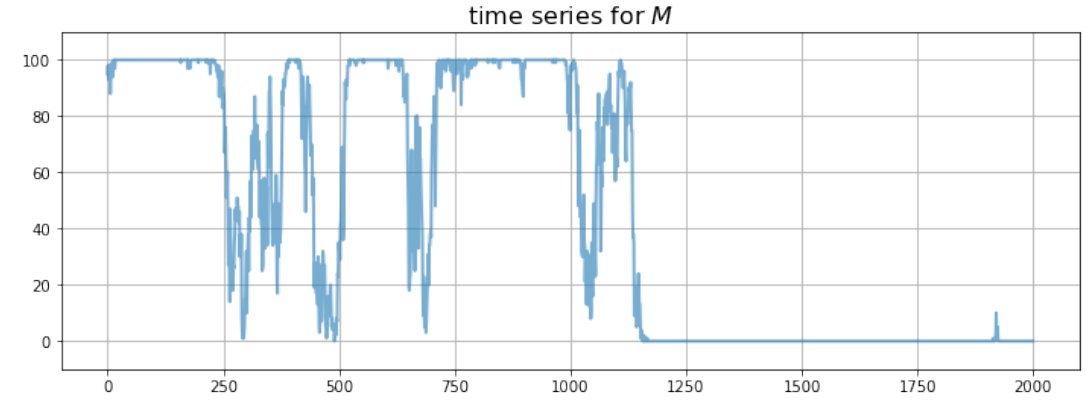
Modulo randomness, replicate the simulation figures shown above.

- Use the parameter values listed as defaults in the **init** method of the UncertaintyTrapEcon class.


time series for θ
7.5
5.0
2.5
0.0
−2.5
−5.0
−7.5
−10.0
−12.5
0
250
500
750
1000
1250
1500
1750
2000


time series for μ
8
6
4
2
0
−2
−4
−6
−8
0
250
500
750
1000
1250
1500
1750
2000


time series for γ
4
3
2
1
0
0
250
500
750
1000
1250
1500
1750
2000


time series for M
100
80
60
40
20
0
0
250
500
750
1000
1250
1500
1750
2000

Solution

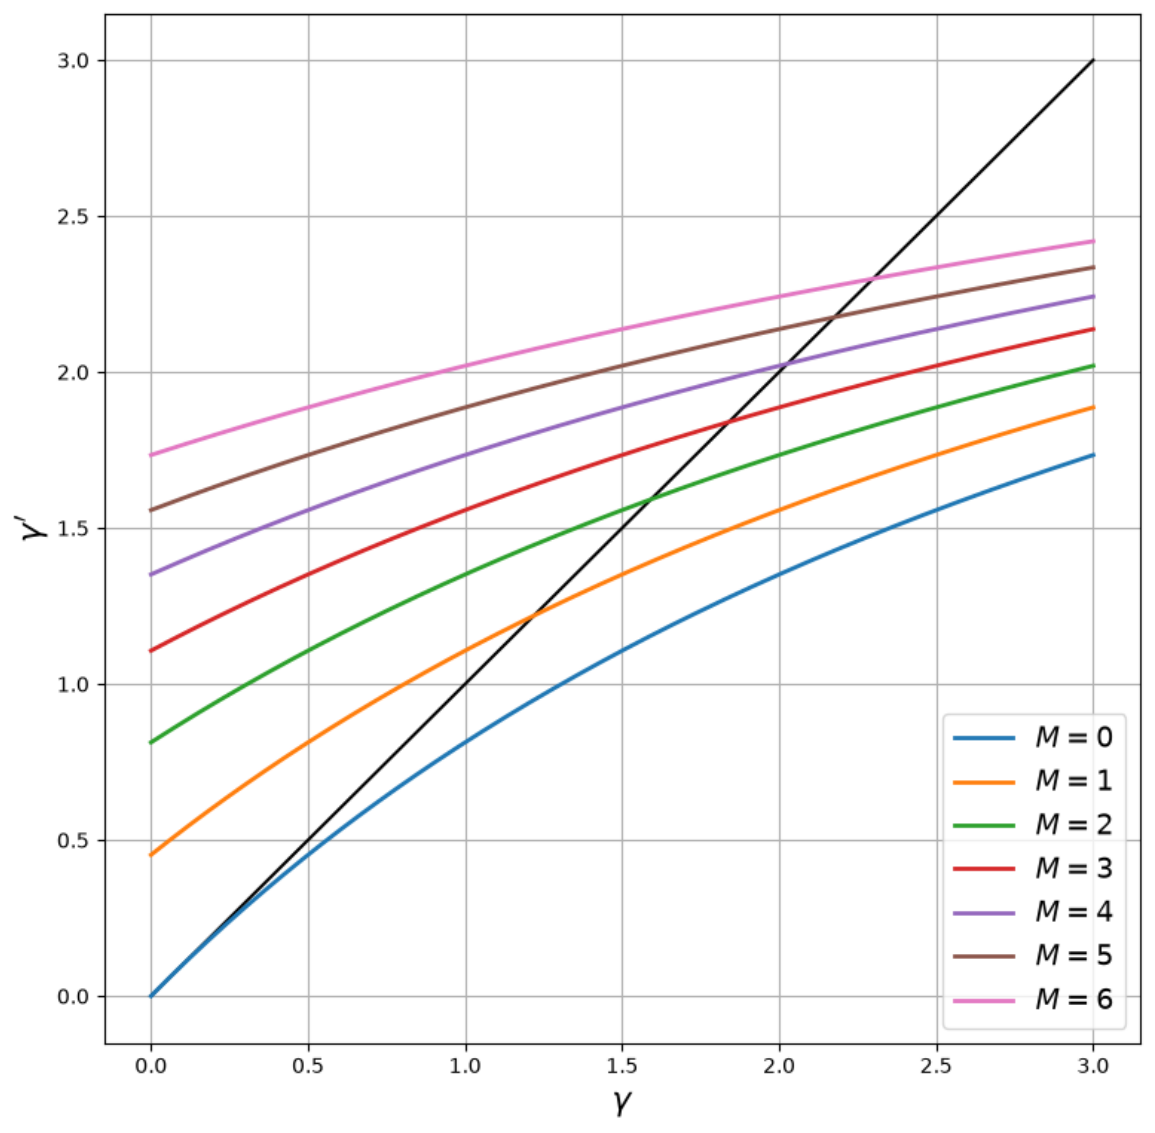
First, let's replicate the plot that illustrates the law of motion for precision, which is

$$\gamma_{t+1} = \left(\frac{\rho^2}{\gamma_t + M\gamma_x} + \sigma_\theta^2 \right)^{-1}$$

Here M is the number of active firms. The next figure plots γ_{t+1} against γ_t on a 45 degree diagram for different values of M

```
econ = UncertaintyTrapEcon()
ρ, σ_θ, γ_x = econ.ρ, econ.σ_θ, econ.γ_x    # Simplify names
γ = np.linspace(1e-10, 3, 200)              # γ grid
fig, ax = plt.subplots(figsize=(9, 9))
ax.plot(γ, γ, 'k-')                         # 45 degree line

for M in range(7):
    γ_next = 1 / (ρ**2 / (γ + M * γ_x) + σ_θ**2)
    label_string = f"$M = {M}$"
    ax.plot(γ, γ_next, lw=2, label=label_string)
ax.legend(loc='lower right', fontsize=14)
ax.set_xlabel(r'$\gamma$', fontsize=16)
ax.set_ylabel(r"$\gamma'$", fontsize=16)
ax.grid()
plt.show()
```



The points where the curves hit the 45 degree lines are the long-run steady states corresponding to each M, if that value of M was to remain fixed. As the number of firms falls, so does the long-run steady state of precision.

Next let's generate time series for beliefs and the aggregates – that is, the number of active firms and average output

```
sim_length=2000

μ_vec = np.empty(sim_length)
θ_vec = np.empty(sim_length)
γ_vec = np.empty(sim_length)
X_vec = np.empty(sim_length)
M_vec = np.empty(sim_length)

μ_vec[0] = econ.μ
γ_vec[0] = econ.γ
θ_vec[0] = 0

rng = np.random.default_rng()
w_shocks = rng.standard_normal(sim_length)

for t in range(sim_length-1):
    X, M = econ.gen_aggregates(rng)
    X_vec[t] = X
    M_vec[t] = M

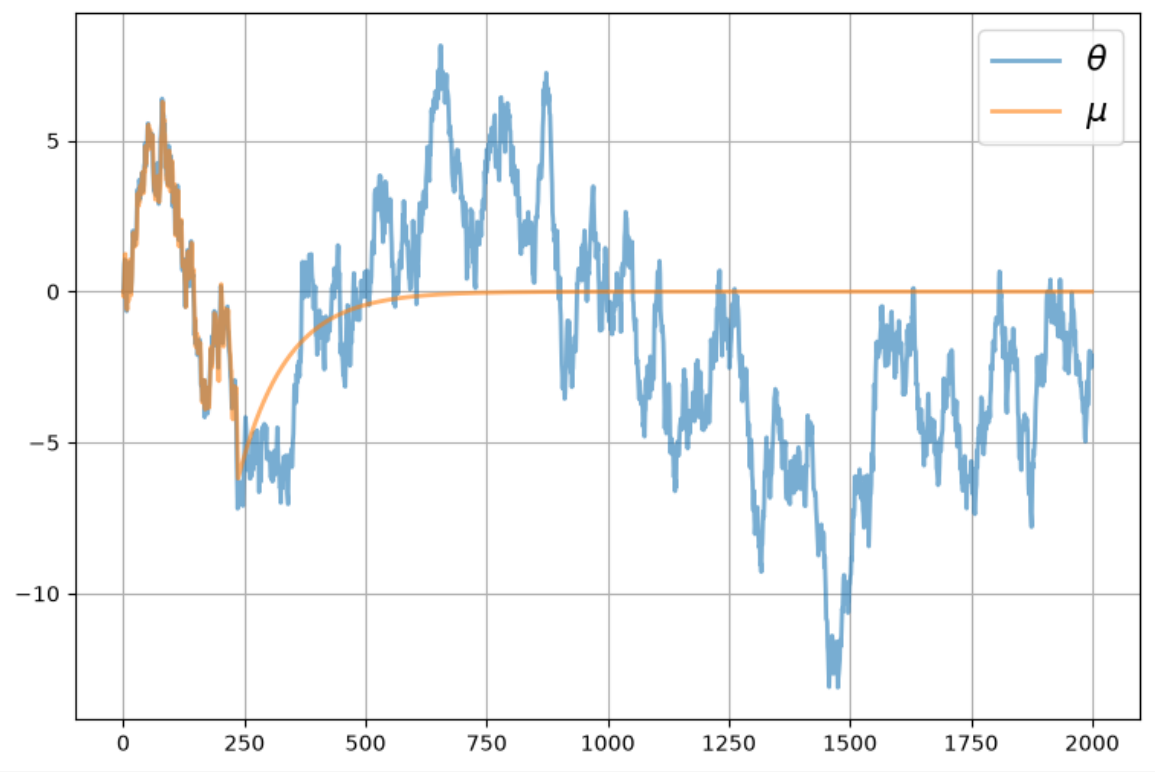
    econ.update_beliefs(X, M)
    econ.update_θ(w_shocks[t])

    μ_vec[t+1] = econ.μ
    γ_vec[t+1] = econ.γ
    θ_vec[t+1] = econ.θ

# Record final values of aggregates
X, M = econ.gen_aggregates(rng)
X_vec[-1] = X
M_vec[-1] = M
```

First, let's see how well μ tracks θ in these simulations

```
fig, ax = plt.subplots(figsize=(9, 6))
ax.plot(range(sim_length), θ_vec, alpha=0.6, lw=2, label=r"$\theta$")
ax.plot(range(sim_length), μ_vec, alpha=0.6, lw=2, label=r"$\mu$")
ax.legend(fontsize=16)
ax.grid()
plt.show()
```



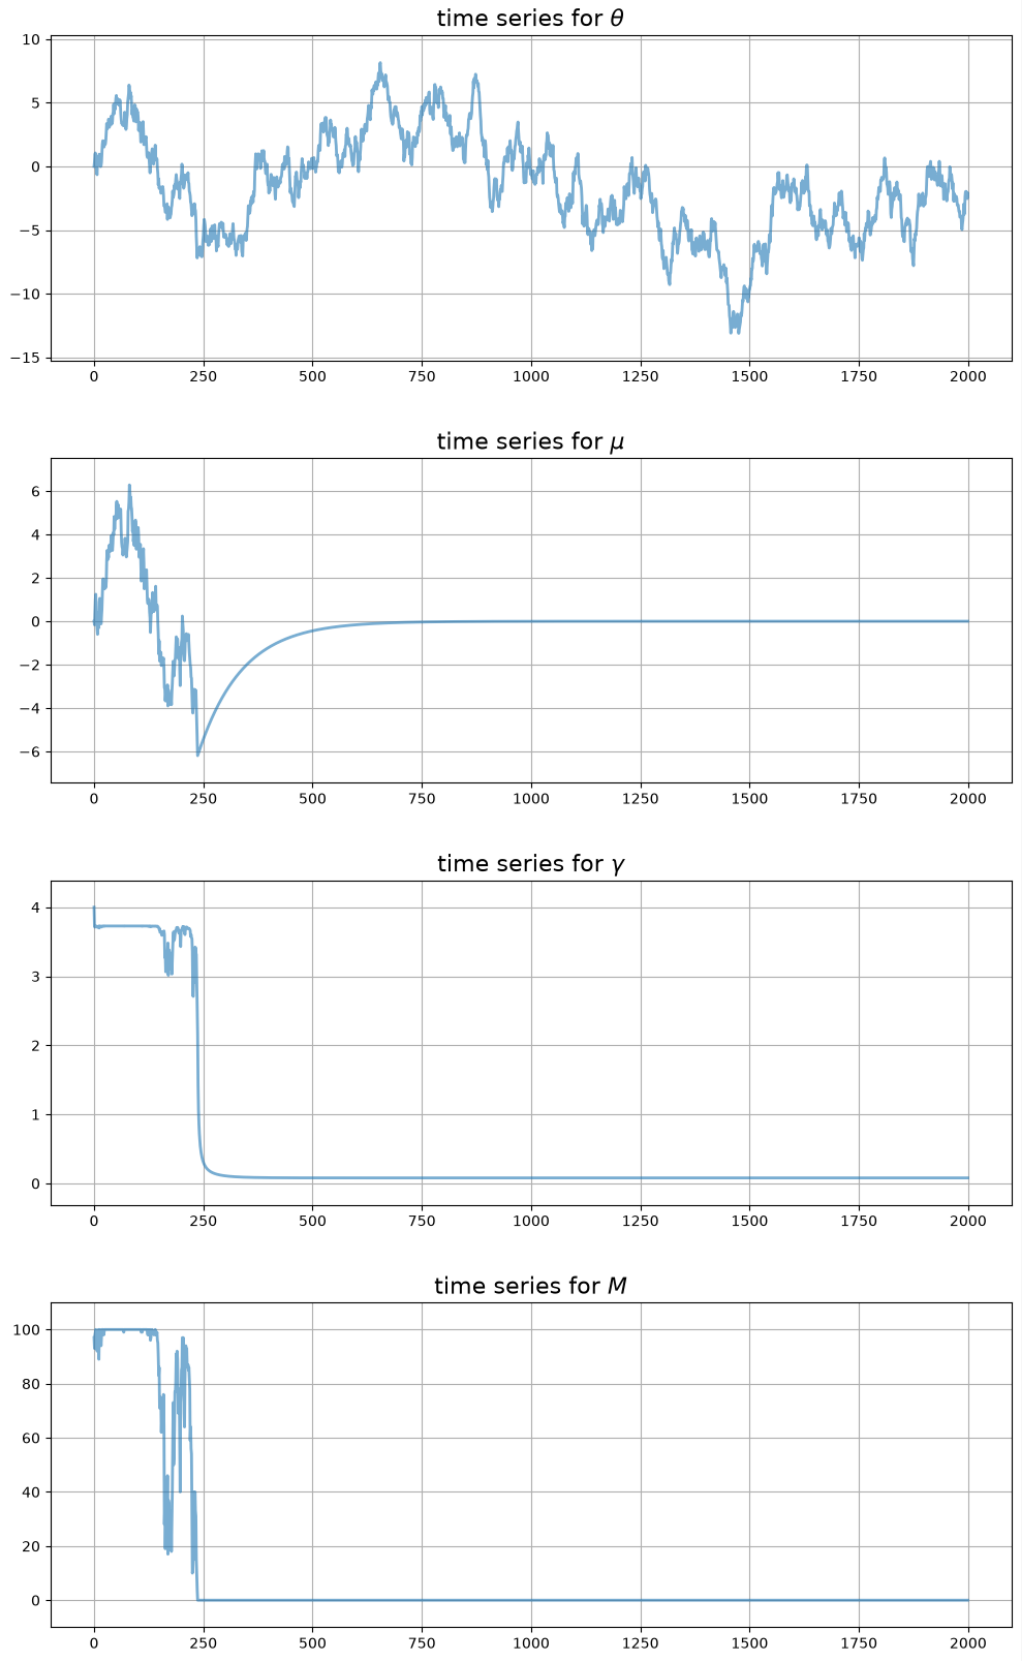
Now let's plot the whole thing together

```
fig, axes = plt.subplots(4, 1, figsize=(12, 20))
# Add some spacing
fig.subplots_adjust(hspace=0.3)

series = (θ_vec, μ_vec, γ_vec, M_vec)
names = r'$\theta$', r'$\mu$', r'$\gamma$', r'$M$'

for ax, vals, name in zip(axes, series, names):
    # Determine suitable y limits
    s_max, s_min = max(vals), min(vals)
    s_range = s_max - s_min
    y_max = s_max + s_range * 0.1
    y_min = s_min - s_range * 0.1
    ax.set_ylim(y_min, y_max)
    # Plot series
    ax.plot(range(sim_length), vals, alpha=0.6, lw=2)
    ax.set_title(f"time series for {name}", fontsize=16)
    ax.grid()

plt.show()
```


time series for θ
time series for μ
time series for γ
time series for M

If you run the code above you'll get different plots, of course.

Try experimenting with different parameters to see the effects on the time series.

(It would also be interesting to experiment with non-Gaussian distributions for the shocks, but this is a big exercise since it takes us outside the world of the standard Kalman filter)

CHAPTER

THREE

THE AIYAGARI MODEL

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *The Aiyagari Model*
  - *Overview*
  - *The Economy*
  - *Implementation*
  - *Exercises*

In addition to what's included in base Anaconda, we need to install JAX

```
!pip install quantecon jax
```

123.1 Overview

In this lecture, we describe the structure of a class of models that build on work by Truman Bewley [Bewley, 1977].

We begin by discussing an example of a Bewley model due to Rao Aiyagari [Aiyagari, 1994].

The model features

- heterogeneous agents
- a single exogenous vehicle for borrowing and lending
- limits on amounts individual agents may borrow

The Aiyagari model has been used to investigate many topics, including

- precautionary savings and the effect of liquidity constraints [Aiyagari, 1994]
- risk sharing and asset pricing [Heaton and Lucas, 1996]
- the shape of the wealth distribution [Benhabib *et al.*, 2015]
- etc., etc., etc.

123.1.1 Preliminaries

We use the following imports:

```
import quantecon as qe
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
from typing import NamedTuple
from scipy.optimize import bisect
```

We will use 64-bit floats with JAX in order to increase precision.

```
jax.config.update("jax_enable_x64", True)
```

We will use the following function to compute stationary distributions of stochastic matrices (for a reference to the algorithm, see p. 88 of Economic Dynamics).

```
@jax.jit
def compute_stationary(P):
    n = P.shape[0]
    I = jnp.identity(n)
    O = jnp.ones((n, n))
    A = I - jnp.transpose(P) + O
    return jnp.linalg.solve(A, jnp.ones(n))
```

123.1.2 References

The primary reference for this lecture is [Aiyagari, 1994].

A textbook treatment is available in chapter 18 of [Ljungqvist and Sargent, 2018].

A continuous time version of the model by SeHyoun Ahn and Benjamin Moll can be found here.

123.2 The Economy

123.2.1 Households

Infinitely lived households / consumers face idiosyncratic income shocks.

A unit interval of *ex-ante* identical households face a common borrowing constraint.

The savings problem faced by a typical household is

$$\max \mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$a_{t+1} + c_t \leq wz_t + (1+r)a_t \quad c_t \geq 0, \quad \text{and} \quad a_t \geq -B$$

where

- c_t is current consumption
- a_t is assets
- z_t is an exogenous component of labor income capturing stochastic unemployment risk, etc.
- w is a wage rate
- r is a net interest rate
- B is the maximum amount that the agent is allowed to borrow

The exogenous process $\{z_t\}$ follows a finite state Markov chain with given stochastic matrix P.

The wage and interest rate are fixed over time.

In this simple version of the model, households supply labor inelastically because they do not value leisure.

123.2.2 Firms

Firms produce output by hiring capital and labor.

Firms act competitively and face constant returns to scale.

Since returns to scale are constant, the number of firms does not matter.

Hence we can consider a single (but nonetheless competitive) representative firm.

The firm's output is

$$Y = AK^\alpha N^{1-\alpha}$$

where

- A and α are parameters with $A > 0$ and $\alpha \in (0, 1)$
- K is aggregate capital
- N is total labor supply (which is constant in this simple version of the model)

The firm's problem is

$$\max_{K,N} \left\{ AK^\alpha N^{1-\alpha} - (r+\delta)K - wN \right\}$$

The parameter δ is the depreciation rate.

These parameters are stored in the following namedtuple:

```
class Firm(NamedTuple):
    A: float = 1.0    # Total factor productivity
    N: float = 1.0    # Total labor supply
    α: float = 0.33   # Capital share
    δ: float = 0.05   # Depreciation rate
```

From the first-order condition with respect to capital, the firm's inverse demand for capital is

$$r = A\alpha \left(\frac{N}{K} \right)^{1-\alpha} - \delta \tag{123.1}$$

```
def r_given_k(K, firm):
    """
    Inverse demand curve for capital. The interest rate associated with a
    given demand for capital K.
    """
    A, N, α, δ = firm
    return A * α * (N / K)**(1 - α) - δ
```

Using this expression and the firm's first-order condition for labor, we can pin down the equilibrium wage rate as a function of r as

$$w(r) = A(1-\alpha)(A\alpha/(r+\delta))^{\alpha/(1-\alpha)} \tag{123.2}$$

```
def r_to_w(r, firm):
    """
    Equilibrium wages associated with a given interest rate r.
    """
    A, N, α, δ = firm
    return A * (1 - α) * (A * α / (r + δ))**(α / (1 - α))
```

123.2.3 Equilibrium

We construct a **stationary rational expectations equilibrium (SREE)**.

In such an equilibrium

- prices induce behavior that generates aggregate quantities consistent with the prices
- aggregate quantities and prices are constant over time

In more detail, an SREE lists a set of prices, savings and production policies such that

- households want to choose the specified savings policies taking the prices as given
- firms maximize profits taking the same prices as given
- the resulting aggregate quantities are consistent with the prices; in particular, the demand for capital equals the supply
- aggregate quantities (defined as cross-sectional averages) are constant

123.3 Implementation

Let's look at how we might compute such an equilibrium in practice.

Below we provide code to solve the household problem, taking r and w as fixed.

123.3.1 Primitives and operators

We will solve the household problem using value function iteration.

First we set up a `NamedTuple` to store the parameters that define a household asset accumulation problem, as well as the grids used to solve it

```
class Household(NamedTuple):
    β: float            # Discount factor
    a_grid: jnp.ndarray # Asset grid
    z_grid: jnp.ndarray # Exogenous states
    Π: jnp.ndarray      # Transition matrix

def create_household(β=0.96,                       # Discount factor
                     Π=[[0.9, 0.1], [0.1, 0.9]],  # Markov chain
                     z_grid=[0.1, 1.0],           # Exogenous states
                     a_min=1e-10, a_max=12.5,     # Asset grid
                     a_size=100):
    """
    Create a Household namedtuple with custom grids.
    """
    a_grid = jnp.linspace(a_min, a_max, a_size)
    z_grid, Π = map(jnp.array, (z_grid, Π))
    return Household(β=β, a_grid=a_grid, z_grid=z_grid, Π=Π)
```

For now we assume that $u(c) = \log(c)$

```
u = jnp.log
```

Here's a namedtuple that stores the wage rate and interest rate with default values

```
class Prices(NamedTuple):
    r: float = 0.01  # Interest rate
    w: float = 1.0   # Wages
```

Now we set up a vectorized version of the right-hand side of the Bellman equation (before maximization), which is a 3D array representing

$$B(a, z, a') = u(wz + (1 + r)a - a') + \beta \sum_{z'} v(a', z')\Pi(z, z')$$

for all (a, z, a').

```
def B(v, household, prices):
    # Unpack
    β, a_grid, z_grid, Π = household
    a_size, z_size = len(a_grid), len(z_grid)
    r, w = prices

    # Compute current consumption as array c[i, j, ip]
    a  = jnp.reshape(a_grid, (a_size, 1, 1))    # a[i]   ->  a[i, j, ip]
    z  = jnp.reshape(z_grid, (1, z_size, 1))    # z[j]   ->  z[i, j, ip]
    ap = jnp.reshape(a_grid, (1, 1, a_size))    # ap[ip] -> ap[i, j, ip]
    c = w * z + (1 + r) * a - ap

    # Calculate continuation rewards at all combinations of (a, z, ap)
    v = jnp.reshape(v, (1, 1, a_size, z_size)) # v[ip, jp] -> v[i, j, ip, jp]
    Π = jnp.reshape(Π, (1, z_size, 1, z_size)) # Π[j, jp]  -> Π[i, j, ip, jp]
```

(continues on next page)

(continued from previous page)

```
    EV = jnp.sum(v * Π, axis=-1)               # sum over last index jp

    # Compute the right-hand side of the Bellman equation
    return jnp.where(c > 0, u(c) + β * EV, -jnp.inf)
```

The next function computes greedy policies

```
def get_greedy(v, household, prices):
    """
    Computes a v-greedy policy σ, returned as a set of indices. If
    σ[i, j] equals ip, then a_grid[ip] is the maximizer at i, j.
    """
    # argmax over ap
    return jnp.argmax(B(v, household, prices), axis=-1)
```

We define the Bellman operator T, which takes a value function v and returns Tv as given in the Bellman equation

```
def T(v, household, prices):
    """
    The Bellman operator. Takes a value function v and returns Tv.
    """
    return jnp.max(B(v, household, prices), axis=-1)
```

Here's value function iteration, which repeatedly applies the Bellman operator until convergence

```
@jax.jit
def value_function_iteration(household, prices, tol=1e-4, max_iter=10_000):
    """
    Implements value function iteration using a compiled JAX loop.
    """
    β, a_grid, z_grid, Π = household
    a_size, z_size = len(a_grid), len(z_grid)

    def condition_function(loop_state):
        i, v, error = loop_state
        return jnp.logical_and(error > tol, i < max_iter)

    def update(loop_state):
        i, v, error = loop_state
        v_new = T(v, household, prices)
        error = jnp.max(jnp.abs(v_new - v))
        return i + 1, v_new, error

    # Initial loop state
    v_init = jnp.zeros((a_size, z_size))
    loop_state_init = (0, v_init, tol + 1)

    # Run the fixed point iteration
    i, v, error = jax.lax.while_loop(condition_function, update, loop_state_init)

    return get_greedy(v, household, prices)
```

As a first example of what we can do, let's compute and plot an optimal accumulation policy at fixed prices

```
# Create an instance of Household
household = create_household()
prices = Prices()
```

(continues on next page)

(continued from previous page)

```
r, w = prices
print(f"Interest rate: {r}, Wage: {w}")
```

```
Interest rate: 0.01, Wage: 1.0
```

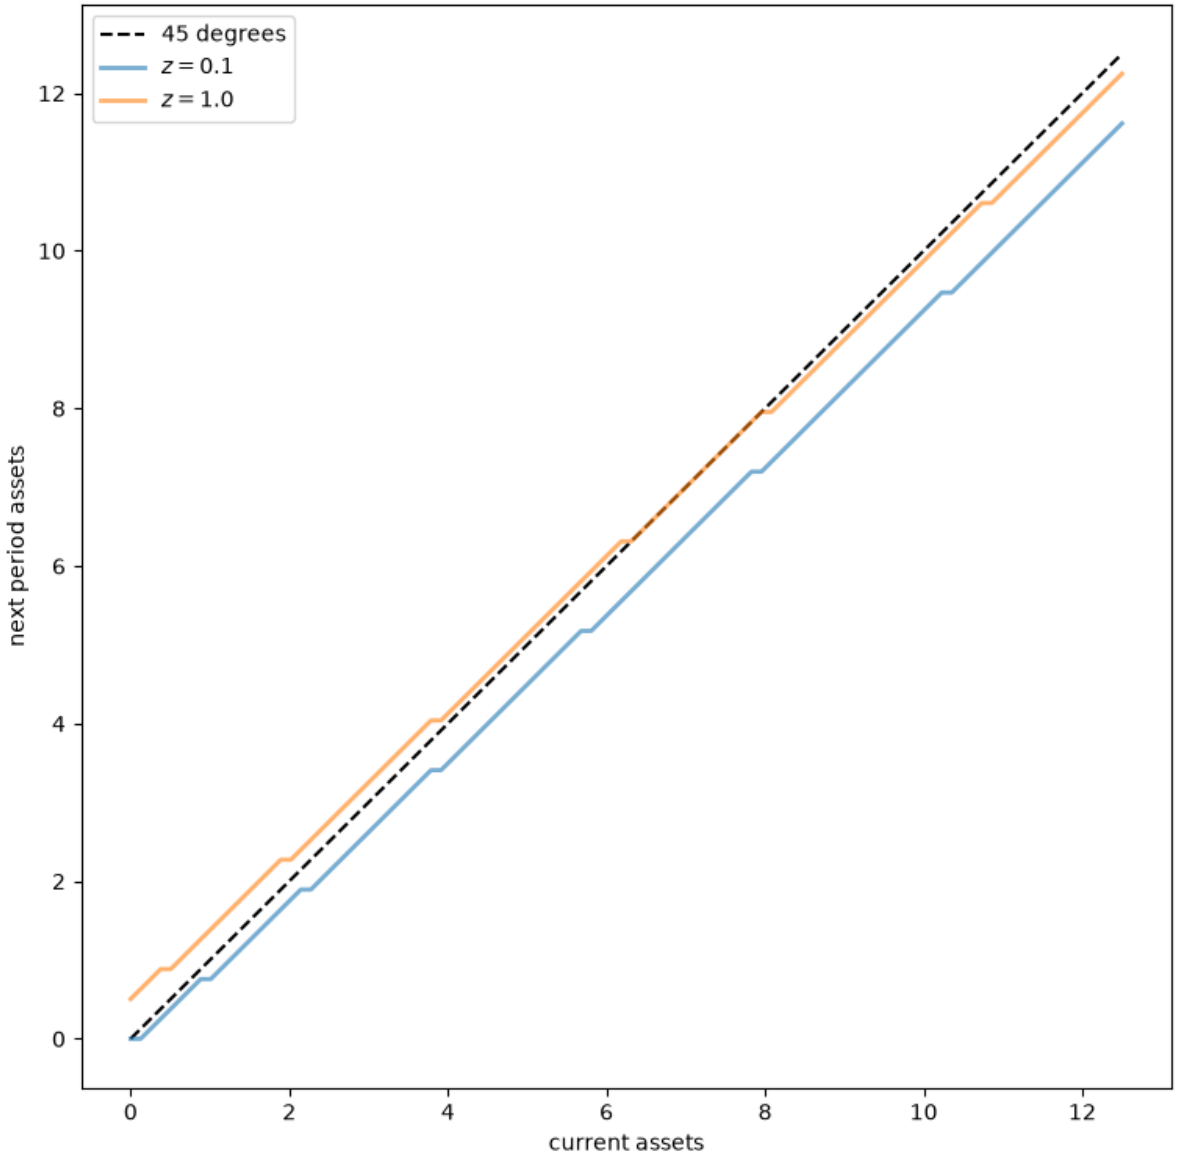
```
with qe.Timer():
    σ_star = value_function_iteration(household, prices).block_until_ready()
```

```
0.2403 seconds elapsed
```

The next plot shows asset accumulation policies at different values of the exogenous state

```
β, a_grid, z_grid, Π = household

fig, ax = plt.subplots(figsize=(9, 9))
ax.plot(a_grid, a_grid, 'k--', label="45 degrees")
for j, z in enumerate(z_grid):
    lb = f'$z = {z:.2}$'
    policy_vals = a_grid[σ_star[:, j]]
    ax.plot(a_grid, policy_vals, lw=2, alpha=0.6, label=lb)
    ax.set_xlabel('current assets')
    ax.set_ylabel('next period assets')
ax.legend(loc='upper left')
plt.show()
```



The plot shows asset accumulation policies at different values of the exogenous state.

123.3.2 Capital supply

To start thinking about equilibrium, we need to know how much capital households supply at a given interest rate r.

This quantity can be calculated by taking the stationary distribution of assets under the optimal policy and computing the mean.

The next function computes the stationary distribution for a given policy σ via the following steps:

- Compute the stationary distribution $\psi = (\psi(a, z))$ of P_σ, which defines the Markov chain of the state (a_t, z_t) under policy σ.
- Sum out z_t to get the marginal distribution for a_t.

```
@jax.jit
def compute_asset_stationary(σ, household):
    # Unpack
    β, a_grid, z_grid, Π = household
    a_size, z_size = len(a_grid), len(z_grid)

    # Construct P_σ as an array of the form P_σ[i, j, ip, jp]
    ap_idx = jnp.arange(a_size)
    ap_idx = jnp.reshape(ap_idx, (1, 1, a_size, 1))
    σ = jnp.reshape(σ, (a_size, z_size, 1, 1))
    A = jnp.where(σ == ap_idx, 1, 0)
    Π = jnp.reshape(Π, (1, z_size, 1, z_size))
    P_σ = A * Π

    # Reshape P_σ into a matrix
    n = a_size * z_size
    P_σ = jnp.reshape(P_σ, (n, n))

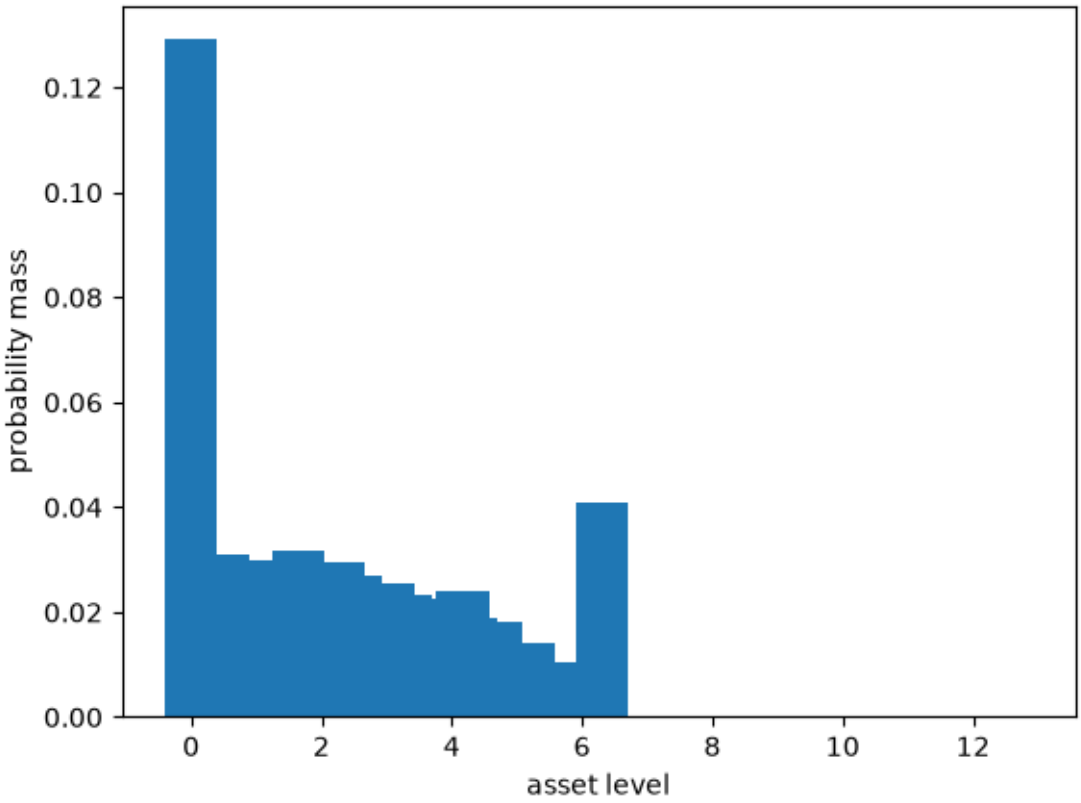
    # Get stationary distribution and reshape back onto [i, j] grid
    ψ = compute_stationary(P_σ)
    ψ = jnp.reshape(ψ, (a_size, z_size))

    # Sum along the rows to get the marginal distribution of assets
    ψ_a = jnp.sum(ψ, axis=1)
    return ψ_a
```

Let's give this a test run.

```
ψ_a = compute_asset_stationary(σ_star, household)

fig, ax = plt.subplots()
ax.bar(household.a_grid, ψ_a)
ax.set_xlabel("asset level")
ax.set_ylabel("probability mass")
plt.show()
```



The distribution should sum to one:

```
ψ_a.sum()
```

```
Array(1., dtype=float64)
```

The next function computes aggregate capital supply by households under policy σ, given wages and interest rates

```
def capital_supply(σ, household):
    """
    Induced level of capital stock under the policy, taking r and w as given.
    """
    β, a_grid, z_grid, Π = household
    ψ_a = compute_asset_stationary(σ, household)
    return float(jnp.sum(ψ_a * a_grid))
```

123.3.3 Equilibrium

We compute a SREE as follows:

1. Set $n = 0$ and start with an initial guess K_0 for aggregate capital.
2. Determine prices r, w from the firm decision problem, given K_n.
3. Compute the optimal savings policy of households given these prices.
4. Compute aggregate capital K_{n+1} as the mean of steady-state capital given this savings policy.
5. If $K_{n+1} \approx K_n$, stop; otherwise, go to step 2.

We can write the sequence of operations in steps 2-4 as

$$K_{n+1} = G(K_n)$$

If K_{n+1} agrees with K_n then we have a SREE.

In other words, our problem is to find the fixed point of the one-dimensional map G.

Here's G expressed as a Python function

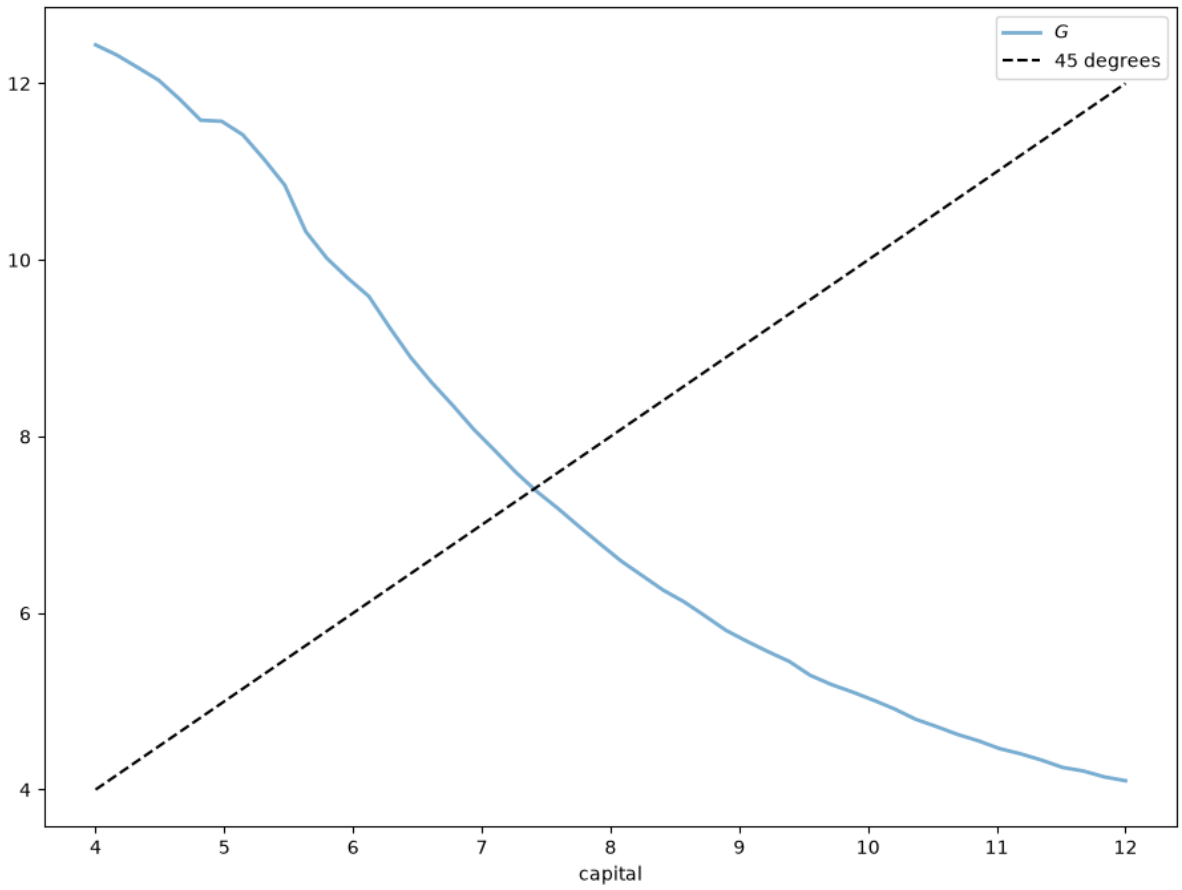
```
def G(K, firm, household):
    # Get prices r, w associated with K
    r = r_given_k(K, firm)
    w = r_to_w(r, firm)

    # Generate a household object with these prices, compute
    # aggregate capital.
    prices = Prices(r=r, w=w)
    σ_star = value_function_iteration(household, prices)
    return capital_supply(σ_star, household)
```

Let's inspect visually as a first pass

```
num_points = 50
firm = Firm()
household = create_household()
k_vals = jnp.linspace(4, 12, num_points)
out = [G(k, firm, household) for k in k_vals]

fig, ax = plt.subplots(figsize=(11, 8))
ax.plot(k_vals, out, lw=2, alpha=0.6, label='$G$')
ax.plot(k_vals, k_vals, 'k--', label="45 degrees")
ax.set_xlabel('capital')
ax.legend()
plt.show()
```



Now let's compute the equilibrium.

Looking at the figure above, we see that a simple iteration scheme $K_{n+1} = G(K_n)$ will cycle from high to low values, leading to slow convergence.

As a result, we use a damped iteration scheme of the form

$$K_{n+1} = \alpha K_n + (1 - \alpha) G(K_n)$$

```
def compute_equilibrium(firm, household,
                        K0=6, α=0.99, max_iter=1_000, tol=1e-4,
                        print_skip=10, verbose=False):
    n = 0
    K = K0
    error = tol + 1
    while error > tol and n < max_iter:
        new_K = α * K + (1 - α) * G(K, firm, household)
        error = abs(new_K - K)
        K = new_K
        n += 1
        if verbose and n % print_skip == 0:
            print(f"At iteration {n} with error {error}")
    return K, n
```

```
firm = Firm()
household = create_household()
```

(continues on next page)

(continued from previous page)

```
print("\nComputing equilibrium capital stock")
with qe.Timer():
    K_star, n = compute_equilibrium(firm, household, K0=6.0)
print(f"Computed equilibrium {K_star:.5} in {n} iterations")
```

```
Computing equilibrium capital stock
2.3051 seconds elapsed
Computed equilibrium 7.4095 in 259 iterations
```

This convergence is not very fast, given how quickly we can solve the household problem.

You can try varying α, but usually this parameter is hard to set a priori.

In the exercises below you will be asked to use bisection instead, which generally performs better.

123.3.4 Supply and demand curves

We can visualize the equilibrium using supply and demand curves.

The following code draws the aggregate supply and demand curves.

The intersection gives the equilibrium interest rate and capital

```
def prices_to_capital_stock(household, r, firm):
    """
    Map prices to the induced level of capital stock.
    """
    w = r_to_w(r, firm)
    prices = Prices(r=r, w=w)

    # Compute the optimal policy
    σ_star = value_function_iteration(household, prices)

    # Compute capital supply
    return capital_supply(σ_star, household)

# Create a grid of r values to compute demand and supply of capital
num_points = 20
r_vals = jnp.linspace(0.005, 0.04, num_points)

# Compute supply of capital
k_vals = []
for r in r_vals:
    k_vals.append(prices_to_capital_stock(household, r, firm))

# Plot against demand for capital by firms
fig, ax = plt.subplots(figsize=(11, 8))
ax.plot(k_vals, r_vals, lw=2, alpha=0.6,
        label='supply of capital')
ax.plot(k_vals, r_given_k(
        jnp.array(k_vals), firm), lw=2, alpha=0.6,
        label='demand for capital')

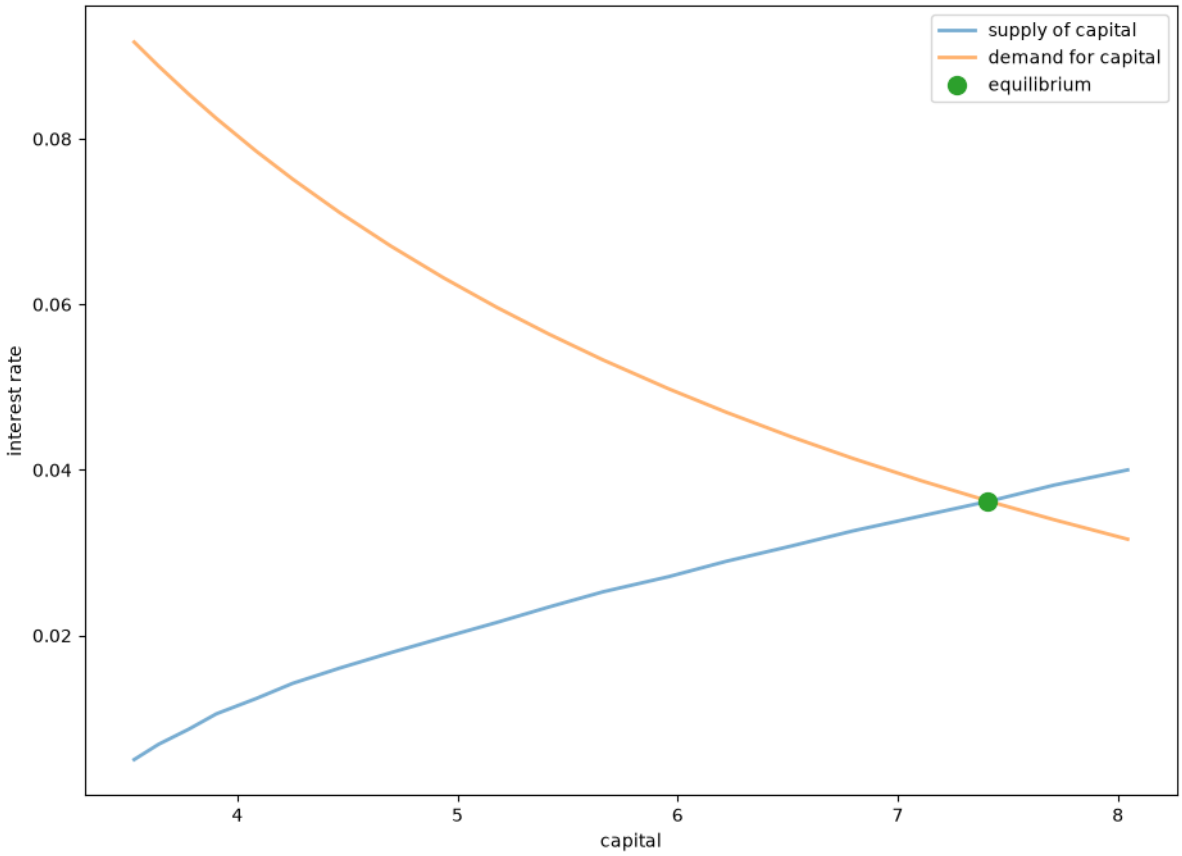
# Add marker at equilibrium
r_star = r_given_k(K_star, firm)
ax.plot(K_star, r_star, 'o', markersize=10, label='equilibrium')
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel('capital')
ax.set_ylabel('interest rate')
ax.legend(loc='upper right')

plt.show()
```



123.4 Exercises

Exercise 123.4.1

Write a new version of `compute_equilibrium` that uses `bisect` from `scipy.optimize` instead of damped iteration.

See if you can make it faster than the previous version.

In `bisect`,

- you should set `xtol=1e-4` to have the same error tolerance as the previous version.
- for the lower and upper bounds of the bisection routine try `a = 1.0` and `b = 20.0`.

Solution

We use bisection to find the zero of the function $h(k) = k - G(k)$

```
def compute_equilibrium_bisect(firm, household, a=1.0, b=20.0):
    K = bisect(lambda k: k - G(k, firm, household), a, b, xtol=1e-4)
    return K

firm = Firm()
household = create_household()
print("\nComputing equilibrium capital stock using bisection")
with qe.Timer():
    K_star = compute_equilibrium_bisect(firm, household)
print(f"Computed equilibrium capital stock {K_star:.5}")
```

```
Computing equilibrium capital stock using bisection
0.1810 seconds elapsed
Computed equilibrium capital stock 7.4178
```

The bisection method is faster than the damped iteration scheme.

Exercise 123.4.2

Show how equilibrium capital stock changes with β.

Use the following values of β and plot the relationship you find.

```
β_vals = jnp.linspace(0.94, 0.98, 20)
```

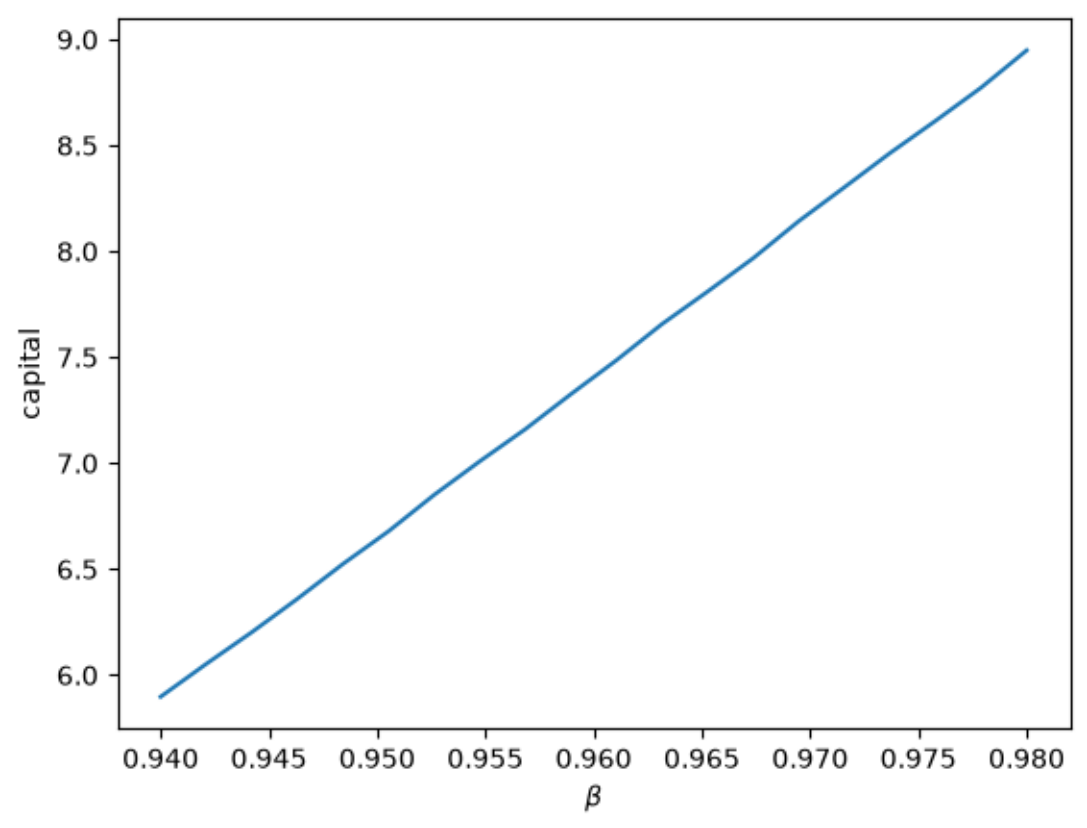
Solution

```
K_vals = []
K = 6.0  # initial guess

for β in β_vals:
    household = create_household(β=β)
    K = compute_equilibrium_bisect(firm, household, 0.5 * K, 1.5 * K)
    print(f"Computed equilibrium {K:.4} at β = {β}")
    K_vals.append(K)

fig, ax = plt.subplots()
ax.plot(β_vals, K_vals, ms=2)
ax.set_xlabel(r'$\beta$')
ax.set_ylabel('capital')
plt.show()
```

```
Computed equilibrium 5.897 at β = 0.94
Computed equilibrium 6.052 at β = 0.9421052631578948
Computed equilibrium 6.202 at β = 0.9442105263157894
Computed equilibrium 6.36 at β = 0.9463157894736841
Computed equilibrium 6.524 at β = 0.9484210526315789
Computed equilibrium 6.679 at β = 0.9505263157894737
Computed equilibrium 6.85 at β = 0.9526315789473683
Computed equilibrium 7.008 at β = 0.9547368421052631
Computed equilibrium 7.16 at β = 0.9568421052631579
Computed equilibrium 7.325 at β = 0.9589473684210527
Computed equilibrium 7.486 at β = 0.9610526315789474
Computed equilibrium 7.657 at β = 0.963157894736842
Computed equilibrium 7.81 at β = 0.9652631578947368
Computed equilibrium 7.967 at β = 0.9673684210526317
Computed equilibrium 8.142 at β = 0.9694736842105263
Computed equilibrium 8.302 at β = 0.971578947368421
Computed equilibrium 8.463 at β = 0.9736842105263158
Computed equilibrium 8.616 at β = 0.9757894736842105
Computed equilibrium 8.773 at β = 0.9778947368421053
Computed equilibrium 8.948 at β = 0.98
```



Exercise 123.4.3

In this lecture, we used value function iteration to solve the household problem.

An alternative is Howard policy iteration (HPI), which is discussed in detail in Dynamic Programming.

HPI can be faster than VFI for some problems because it uses fewer but more computationally intensive iterations.

Your task is to implement Howard policy iteration and compare the results with value function iteration.

Key concepts you'll need:

Howard policy iteration requires computing the value v_σ of a policy σ, defined as:

$$v_\sigma = (I - \beta P_\sigma)^{-1} r_\sigma$$

where r_σ is the reward vector under policy σ, and P_σ is the transition matrix induced by σ.

To solve this, you'll need to:

1. Compute current rewards $r_\sigma(a, z) = u((1 + r)a + wz - \sigma(a, z))$
2. Set up the linear operator R_σ where $(R_\sigma v)(a, z) = v(a, z) - \beta \sum_{z'} v(\sigma(a, z), z')\Pi(z, z')$
3. Solve $v_\sigma = R_\sigma^{-1} r_\sigma$ using `jax.scipy.sparse.linalg.bicgstab`

You can use the `get_greedy` function that's already defined in this lecture.

Implement the following Howard policy iteration routine:

```
def howard_policy_iteration(household, prices,
                            tol=1e-4, max_iter=10_000, verbose=False):
    """
    Howard policy iteration routine.
    """
    # Your code here
    pass
```

Once implemented, compute the equilibrium capital stock using HPI and verify that it produces approximately the same result as VFI at the default parameter values.

Solution

First, we need to implement the helper functions for Howard policy iteration.

The following function computes the array r_σ which gives current rewards given policy σ:

```
def compute_r_σ(σ, household, prices):
    """
    Compute current rewards at each i, j under policy σ. In particular,

        r_σ[i, j] = u((1 + r)a[i] + wz[j] - a'[ip])

    when ip = σ[i, j].
    """
    # Unpack
    β, a_grid, z_grid, Π = household
    a_size, z_size = len(a_grid), len(z_grid)
    r, w = prices

    # Compute r_σ[i, j]
    a = jnp.reshape(a_grid, (a_size, 1))
    z = jnp.reshape(z_grid, (1, z_size))
    ap = a_grid[σ]
    c = (1 + r) * a + w * z - ap
    r_σ = u(c)

    return r_σ
```

The linear operator R_σ is defined as:

```
def R_σ(v, σ, household):
    # Unpack
    β, a_grid, z_grid, Π = household
    a_size, z_size = len(a_grid), len(z_grid)

    # Set up the array v[σ[i, j], jp]
    zp_idx = jnp.arange(z_size)
    zp_idx = jnp.reshape(zp_idx, (1, 1, z_size))
    σ = jnp.reshape(σ, (a_size, z_size, 1))
    V = v[σ, zp_idx]

    # Expand Π[j, jp] to Π[i, j, jp]
    Π = jnp.reshape(Π, (1, z_size, z_size))

    # Compute and return v[i, j] - β Σ_jp v[σ[i, j], jp] * Π[j, jp]
    return v - β * jnp.sum(V * Π, axis=-1)
```

The next function computes the lifetime value of a given policy:

```
def get_value(σ, household, prices):
    """
    Get the lifetime value of policy σ by computing

        v_σ = R_σ^{-1} r_σ
    """
    r_σ = compute_r_σ(σ, household, prices)

    # Reduce R_σ to a function in v
    _R_σ = lambda v: R_σ(v, σ, household)

    # Compute v_σ = R_σ^{-1} r_σ using an iterative routine.
    return jax.scipy.sparse.linalg.bicgstab(_R_σ, r_σ)[0]
```

Now we can implement Howard policy iteration:

```
@jax.jit
def howard_policy_iteration(household, prices, tol=1e-4, max_iter=10_000):
    """
    Howard policy iteration routine using a compiled JAX loop.
    """
    β, a_grid, z_grid, Π = household
    a_size, z_size = len(a_grid), len(z_grid)

    def condition_function(loop_state):
        i, σ, v_σ, error = loop_state
        return jnp.logical_and(error > tol, i < max_iter)

    def update(loop_state):
        i, σ, v_σ, error = loop_state
        σ_new = get_greedy(v_σ, household, prices)
        v_σ_new = get_value(σ_new, household, prices)
        error = jnp.max(jnp.abs(v_σ_new - v_σ))
        return i + 1, σ_new, v_σ_new, error

    # Initial loop state
    σ_init = jnp.zeros((a_size, z_size), dtype=int)
    v_σ_init = get_value(σ_init, household, prices)
    loop_state_init = (0, σ_init, v_σ_init, tol + 1)
```

```
    # Run the fixed point iteration
    i, σ, v_σ, error = jax.lax.while_loop(condition_function, update, loop_state_
↪init)

    return σ
```

Now let's create a modified version of the G function that uses HPI:

```
def G_hpi(K, firm, household):
    # Get prices r, w associated with K
    r = r_given_k(K, firm)
    w = r_to_w(r, firm)

    # Generate prices and compute aggregate capital using HPI.
    prices = Prices(r=r, w=w)
    σ_star = howard_policy_iteration(household, prices)
    return capital_supply(σ_star, household)
```

And compute the equilibrium using HPI:

```
def compute_equilibrium_bisect_hpi(firm, household, a=1.0, b=20.0):
    K = bisect(lambda k: k - G_hpi(k, firm, household), a, b, xtol=1e-4)
    return K

firm = Firm()
household = create_household()
print("\nComputing equilibrium capital stock using HPI")
with qe.Timer():
    K_star_hpi = compute_equilibrium_bisect_hpi(firm, household)
print(f"Computed equilibrium capital stock with HPI: {K_star_hpi:.5}")
print(f"Previous equilibrium capital stock with VFI: {K_star:.5}")
print(f"Difference: {abs(K_star_hpi - K_star):.6}")
```

```
Computing equilibrium capital stock using HPI
0.8993 seconds elapsed
Computed equilibrium capital stock with HPI: 7.4172
Previous equilibrium capital stock with VFI: 7.4178
Difference: 0.000579834
```

The results show that both methods produce approximately the same equilibrium, confirming that HPI is a valid alternative to VFI.

CHAPTER

FOUR

THE AIYAGARI MODEL WITH ENDOGENOUS GRID METHOD

Contents

- *The Aiyagari Model with Endogenous Grid Method*
  - *Overview*
  - *The Economy*
  - *Implementation with EGM*
  - *Computing Aggregate Capital by Simulation*
  - *Computing Equilibrium*
  - *Wealth Distribution*
  - *Summary and Comparison*
  - *Exercises*

In addition to what's included in base Anaconda, we need to install QuantEcon's Python library and JAX.

```
!pip install quantecon jax
```

124.1 Overview

This lecture combines two important computational methods in macroeconomics:

1. The **Aiyagari model** [Aiyagari, 1994] – a heterogeneous agent model with incomplete markets
2. The **endogenous grid method** (EGM) [Carroll, 2006] – an efficient algorithm for solving dynamic programming problems

In the *standard Aiyagari lecture*, we solved the household problem using discretization and value function iteration.

We then computed aggregate capital at a given set of prices using the stationary distribution of the finite Markov chain.

In this lecture, we take a different approach:

1. We use the **endogenous grid method** to solve the household problem via the Euler equation and linear interpolation.
2. We compute aggregate capital by **simulation** rather than an algebraic technique (which only works for the finite case).

These modifications make the solution method faster and more flexible, especially when dealing with more complex models.

We use JAX throughout, so that the EGM operator, the solver and the simulation are all JIT-compiled and vectorized.

124.1.1 References

The primary references for this lecture are:

- our *previous Aiyagari lecture* for the key ideas
- [Aiyagari, 1994] for the economic model
- [Carroll, 2006] for the endogenous grid method
- Chapter 18 of [Ljungqvist and Sargent, 2018] for textbook treatment

124.1.2 Preliminaries

We use the following imports:

```
import quantecon as qe
import matplotlib.pyplot as plt
import jax
import jax.numpy as jnp
import numpy as np
from typing import NamedTuple
from functools import partial
from scipy.optimize import bisect
```

We will use 64-bit floats with JAX in order to increase precision.

```
jax.config.update("jax_enable_x64", True)
```

124.2 The Economy

The economy consists of households and a representative firm.

124.2.1 Households

Infinitely lived households face idiosyncratic income shocks and a borrowing constraint.

The savings problem faced by a typical household is

$$\max \mathbb{E} \sum_{t=0}^{\infty} \beta^t u(c_t)$$

subject to

$$a_{t+1} + c_t \leq w z_t + (1+r) a_t \quad c_t \geq 0, \quad \text{and} \quad a_t \geq -B$$

where

- c_t is current consumption

- a_t is assets
- z_t is an exogenous component of labor income (stochastic employment status)
- w is the wage rate
- r is the interest rate
- B is the maximum amount that the agent is allowed to borrow

The exogenous process $\{z_t\}$ follows a finite state Markov chain with stochastic matrix Π.

Optimal interior consumption choices satisfy the Euler equation

$$u'(c) = \beta \mathbb{E}_z[(1+r)u'(c')]$$

(We use $'$ symbols for both derivatives and future values, which is not ideal but convenient and common.)

In terms of assets, this is

$$u'(wz + (1+r)a - a') = \beta(1+r)\sum_{z'} u'(wz' + (1+r)a' - s(a', z'))\Pi(z, z')$$

where s is the optimal savings policy function.

124.2.2 Firms

Firms produce output by hiring capital and labor under constant returns to scale.

The representative firm's output is

$$Y = AK^\alpha N^{1-\alpha}$$

where

- A and α are parameters with $A > 0$ and $\alpha \in (0, 1)$
- K is aggregate capital
- N is total labor supply (normalized to 1)

These parameters are stored in the following namedtuple:

```
class Firm(NamedTuple):
    A: float = 1.0    # Total factor productivity
    N: float = 1.0    # Total labor supply
    α: float = 0.33   # Capital share
    δ: float = 0.05   # Depreciation rate
```

From the firm's first-order condition, the inverse demand for capital is

$$r = A\alpha \left(\frac{N}{K}\right)^{1-\alpha} - \delta \tag{124.1}$$

```
def r_given_k(K, firm):
    """
    Inverse demand curve for capital.
    """
    A, N, α, δ = firm
    return A * α * (N / K)**(1 - α) - δ
```

The equilibrium wage rate as a function of r is

$$w(r) = A(1-\alpha)(A\alpha/(r+\delta))^{\alpha/(1-\alpha)} \tag{124.2}$$

```
def r_to_w(r, firm):
    """
    Equilibrium wages associated with a given interest rate r.
    """
    A, N, α, δ = firm
    return A * (1 - α) * (A * α / (r + δ))**(α / (1 - α))
```

124.2.3 Equilibrium

A **stationary rational expectations equilibrium (SREE)** consists of prices and policies such that:

- Households optimize given prices
- Firms maximize profits given prices
- Markets clear: aggregate capital supply equals aggregate capital demand
- Aggregate quantities are constant over time

124.3 Implementation with EGM

124.3.1 Household primitives

First we set up the household parameters and grids:

```
class Household(NamedTuple):
    β: float            # Discount factor
    a_grid: jnp.ndarray # Asset grid
    z_grid: jnp.ndarray # Exogenous states
    Π: jnp.ndarray      # Transition matrix

def create_household(β=0.96,                       # Discount factor
                     Π=[[0.9, 0.1], [0.1, 0.9]],   # Markov chain
                     z_grid=[0.1, 1.0],            # Exogenous states
                     a_min=1e-10, a_max=50.0,      # Asset grid
                     a_size=200):
    """
    Create a Household namedtuple with custom grids.
    """
    a_grid = jnp.linspace(a_min, a_max, a_size)
    z_grid, Π = map(jnp.array, (z_grid, Π))
    return Household(β=β, a_grid=a_grid, z_grid=z_grid, Π=Π)
```

For utility, we assume $u(c) = \log(c)$, which gives us $u'(c) = 1/c$ and $(u')^{-1}(x) = 1/x$.

```
@jax.jit
def u_prime(c):
    return 1 / c

@jax.jit
```

(continues on next page)

(continued from previous page)

```
def u_prime_inv(x):
    return 1 / x
```

Here's a namedtuple for prices:

```
class Prices(NamedTuple):
    r: float = 0.01  # Interest rate
    w: float = 1.0   # Wages
```

124.3.2 The EGM operator

The key insight of EGM is to avoid root-finding by choosing the asset grid exogenously and computing the consumption values directly from the Euler equation.

The Coleman-Reffett operator using EGM works as follows:

1. Start with a consumption policy σ represented on an exogenous grid of (next-period) assets $\{a_i\}$.
2. For each asset level a_i and current employment state z_j:
   - Compute the right-hand side of the Euler equation: \$RHS $= \beta(1+r)\sum_{z'} \Pi(z_j, z')\, u'(\sigma(a_i, z'))$\$
   - Use the inverse marginal utility to get current consumption: \$$c_{ij} = (u')^{-1}$(RHS)\$
   - Recover the implied current asset level from the budget constraint: \$$a_{ij} = \frac{c_{ij}+a_i-wz_j}{1+r}$\$
3. Reconstruct the new policy $K\sigma$ on the original asset grid by interpolating (a_{ij}, c_{ij}), handling the borrowing constraint where it binds.

The whole operation vectorizes cleanly, so we write it as a single JIT-compiled function and use `vmap` over the employment states:

```
@jax.jit
def K_egm(σ, household, prices):
    """
    The Coleman-Reffett operator using EGM for the Aiyagari model.

    Here σ[i, j] is consumption when next-period assets are a_grid[i]
    and the current employment state is z_grid[j].
    """
    β, a_grid, z_grid, Π = household
    r, w = prices
    z_size = len(z_grid)

    # Expectation E[u'(c(a', z')) | z] over the next-period shock
    Eu_prime = (Π @ u_prime(σ).T).T                          # (a_size, z_size)

    # Euler equation -> consumption on the endogenous grid
    c_endo = u_prime_inv(β * (1 + r) * Eu_prime)             # (a_size, z_size)

    # Implied current assets: a = (c + a' - w z) / (1 + r)
    a_endo = (c_endo + a_grid[:, None] - w * z_grid[None, :]) / (1 + r)

    # Interpolate back onto the exogenous grid for each employment state
    def interp_policy(j):
        # Where today's assets fall below the endogenous grid the borrowing
        # constraint binds, so the household saves a_grid[0] and consumes
```

(continues on next page)

(continued from previous page)

```
        # the rest of current income.
        return jnp.where(
            a_grid < a_endo[0, j],
            w * z_grid[j] + (1 + r) * a_grid - a_grid[0],
            jnp.interp(a_grid, a_endo[:, j], c_endo[:, j])
        )

    σ_new = jax.vmap(interp_policy)(jnp.arange(z_size))     # (z_size, a_size)
    return σ_new.T                                          # (a_size, z_size)
```

124.3.3 Solving the household problem

We solve for the optimal policy by iterating the EGM operator to convergence.

The solver is fully JIT-compiled and uses `jax.lax.while_loop` for the iteration:

```
@jax.jit
def solve_household(household, prices, tol=1e-6, max_iter=10_000):
    """
    Solve the household problem by iterating the EGM operator.

    Returns the optimal consumption policy σ[i, j], where i indexes
    next-period assets and j indexes employment states.
    """
    β, a_grid, z_grid, Π = household
    r, w = prices

    # Initial guess: consume half of current income
    income = w * z_grid[None, :] + (1 + r) * a_grid[:, None]
    σ_init = 0.5 * income

    def condition(state):
        i, σ, error = state
        return (error > tol) & (i < max_iter)

    def body(state):
        i, σ, error = state
        σ_new = K_egm(σ, household, prices)
        error = jnp.max(jnp.abs(σ_new - σ))
        return i + 1, σ_new, error

    i, σ, error = jax.lax.while_loop(condition, body, (0, σ_init, tol + 1))
    return σ
```

Let's test this on an example:

```
household = create_household()
prices = Prices(r=0.01, w=1.0)

with qe.Timer():
    σ_star = solve_household(household, prices)
    jax.block_until_ready(σ_star)
```

```
0.5037 seconds elapsed
```

We can check that the policy is a fixed point of the EGM operator by measuring the residual:

```
residual = jnp.max(jnp.abs(K_egm(σ_star, household, prices) - σ_star))
print(f"Final Euler residual: {residual:.2e}")
```

```
Final Euler residual: 8.34e-07
```

Let's plot the resulting policy functions:

```
β, a_grid, z_grid, Π = household
r, w = prices

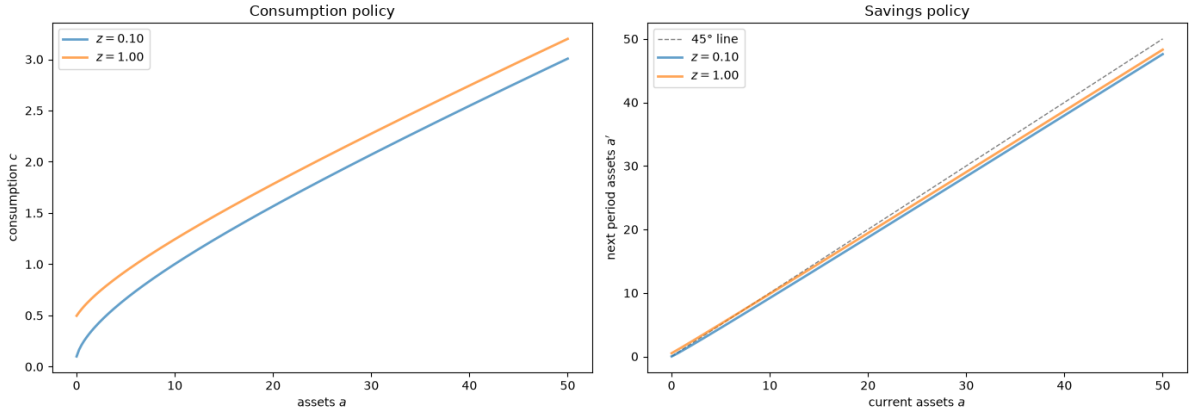
# Convert consumption policy to savings policy
income = w * z_grid[None, :] + (1 + r) * a_grid[:, None]
savings = income - σ_star

fig, axes = plt.subplots(1, 2, figsize=(14, 5))

# Plot consumption policy
ax = axes[0]
for j, z in enumerate(z_grid):
    ax.plot(a_grid, σ_star[:, j], label=f'$z={z:.2f}$', lw=2, alpha=0.7)
ax.set_xlabel('assets $a$')
ax.set_ylabel('consumption $c$')
ax.set_title('Consumption policy')
ax.legend()

# Plot savings policy
ax = axes[1]
ax.plot(a_grid, a_grid, 'k--', lw=1, alpha=0.5, label='45° line')
for j, z in enumerate(z_grid):
    ax.plot(a_grid, savings[:, j], label=f'$z={z:.2f}$', lw=2, alpha=0.7)
ax.set_xlabel('current assets $a$')
ax.set_ylabel("next period assets $a'$")
ax.set_title('Savings policy')
ax.legend()

plt.tight_layout()
plt.show()
```



124.4 Computing Aggregate Capital by Simulation

Instead of computing the stationary distribution of the Markov chain analytically, we compute aggregate capital by simulating a large cross-section of households.

This approach:

- is more flexible (works with continuous shocks, non-linear policies, etc.)
- avoids storing and manipulating large transition matrices
- is conceptually simpler

The simulation is fully vectorized: we advance all households simultaneously with `jax.lax.fori_loop`, drawing the employment transitions by inverse-CDF sampling and looking up consumption with a `vmap`-ed interpolation.

```
@partial(jax.jit, static_argnames=('num_households', 'num_periods'))
def simulate_cross_section(σ, household, prices, key,
                           num_households=50_000, num_periods=1_000):
    """
    Simulate a panel of households forward and return the terminal
    cross-section of assets and employment states.
    """
    β, a_grid, z_grid, Π = household
    r, w = prices

    # CDF of each row of Π, used for inverse-CDF sampling of transitions
    Π_cdf = jnp.cumsum(Π, axis=1)

    # Vectorized consumption lookup: interpolate σ along assets for each z
    @jax.vmap
    def consume(a, j):
        return jnp.interp(a, a_grid, σ[:, j])

    # Initial conditions: everyone at the middle of the grid, in state 0
    assets = jnp.full(num_households, a_grid[len(a_grid) // 2])
    z_idx = jnp.zeros(num_households, dtype=jnp.int32)

    def step(t, state):
        assets, z_idx, key = state
        key, subkey = jax.random.split(key)
        unif = jax.random.uniform(subkey, (num_households,))
        # Markov transition via inverse CDF
        z_idx = (unif[:, None] > Π_cdf[z_idx]).sum(axis=1).astype(jnp.int32)
        # Budget constraint: consume, then carry assets to next period
        income = w * z_grid[z_idx] + (1 + r) * assets
        assets = income - consume(assets, z_idx)
        # Enforce the asset grid bounds
        assets = jnp.clip(assets, a_grid[0], a_grid[-1])
        return assets, z_idx, key

    assets, z_idx, key = jax.lax.fori_loop(
        0, num_periods, step, (assets, z_idx, key)
    )
    return assets, z_idx
```

Now we can compute capital supply for given prices by solving the household problem and averaging assets across the simulated cross-section:

```
def capital_supply(household, prices, key,
                   num_households=50_000, num_periods=1_000):
    """
    Compute aggregate capital supply by simulation.
    """
    σ = solve_household(household, prices)
    assets, _ = simulate_cross_section(
        σ, household, prices, key,
        num_households=num_households, num_periods=num_periods
    )
    return float(jnp.mean(assets))
```

Let's test it:

```
household = create_household()
prices = Prices(r=0.01, w=1.0)
key = jax.random.key(42)

with qe.Timer():
    K_supply = capital_supply(household, prices, key)

print(f"Capital supply: {K_supply:.4f}")
```

```
0.8977 seconds elapsed
Capital supply: 2.5863
```

124.5 Computing Equilibrium

Now we can compute the equilibrium by finding the capital stock at which capital supply equals capital demand.

Given K, the equilibrium mapping G computes:

1. prices (r, w) from the firm's first-order conditions,
2. the household's optimal policy given those prices,
3. aggregate capital supply via simulation.

```
def G(K, firm, household, key,
      num_households=50_000, num_periods=1_000):
    """
    The equilibrium mapping K -> capital supply.
    """
    r = r_given_k(K, firm)
    w = r_to_w(r, firm)
    prices = Prices(r=r, w=w)
    return capital_supply(household, prices, key,
                          num_households=num_households,
                          num_periods=num_periods)
```

We compute the equilibrium by applying bisection to the excess demand $K - G(K)$.

We pass a fixed random key to every evaluation so that the excess demand function is deterministic, as required by the root finder.

```
def compute_equilibrium(firm, household, key,
                        K_min=4.0, K_max=12.0,
                        num_households=50_000, num_periods=1_000,
                        xtol=1e-2):
    """
    Compute the equilibrium capital stock using bisection.
    """
    def excess_demand(K):
        return K - G(K, firm, household, key,
                     num_households=num_households,
                     num_periods=num_periods)

    return bisect(excess_demand, K_min, K_max, xtol=xtol)
```

Let's compute the equilibrium:

```
firm = Firm()
household = create_household()
key = jax.random.key(42)

with qe.Timer():
    K_star = compute_equilibrium(firm, household, key)

r_star = r_given_k(K_star, firm)
w_star = r_to_w(r_star, firm)

print(f"\nEquilibrium capital:       {K_star:.4f}")
print(f"Equilibrium interest rate: {r_star:.4f}")
print(f"Equilibrium wage:          {w_star:.4f}")
```

```
1.9688 seconds elapsed

Equilibrium capital:       8.1484
Equilibrium interest rate: 0.0309
Equilibrium wage:          1.3388
```

124.5.1 Visualizing equilibrium

Let's plot the supply and demand curves:

```
# Supply curve: capital supplied by households as a function of r
r_vals = np.linspace(0.005, 0.04, 10)
K_supply_vals = []

for r in r_vals:
    w = r_to_w(r, firm)
    prices = Prices(r=r, w=w)
    K_supply_vals.append(capital_supply(household, prices, key))

# Demand curve: capital demanded by firms as a function of r
K_vals = np.linspace(4, 12, 50)
r_demand_vals = r_given_k(K_vals, firm)

fig, ax = plt.subplots(figsize=(10, 6))
```

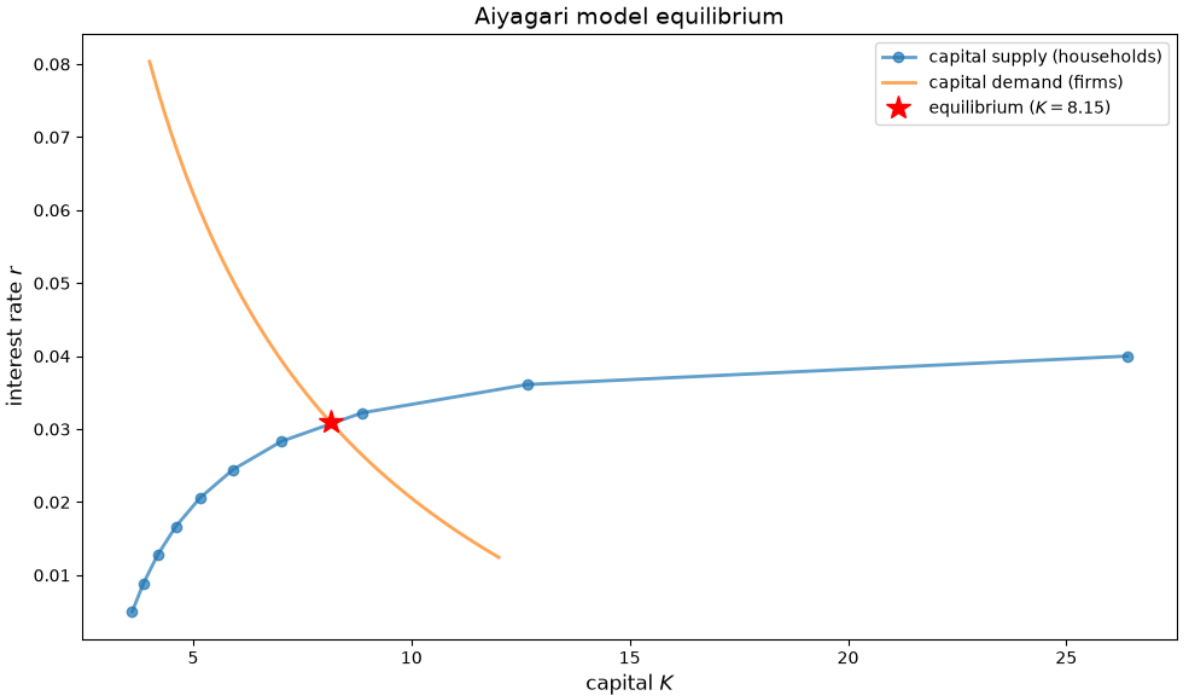
(continues on next page)

(continued from previous page)

```
ax.plot(K_supply_vals, r_vals, 'o-', lw=2, alpha=0.7,
        label='capital supply (households)', markersize=6)
ax.plot(K_vals, r_demand_vals, lw=2, alpha=0.7,
        label='capital demand (firms)')
ax.plot(K_star, r_star, 'r*', markersize=15, zorder=5,
        label=f'equilibrium ($K={K_star:.2f}$)')

ax.set_xlabel('capital $K$', fontsize=12)
ax.set_ylabel('interest rate $r$', fontsize=12)
ax.set_title('Aiyagari model equilibrium', fontsize=14)
ax.legend(fontsize=10)

plt.tight_layout()
plt.show()
```



124.6 Wealth Distribution

One advantage of the simulation approach is that we can easily examine the wealth distribution.

We reuse the cross-section simulated at equilibrium prices:

```
prices_star = Prices(r=r_star, w=w_star)
σ_star = solve_household(household, prices_star)
assets_dist, z_dist = simulate_cross_section(σ_star, household, prices_star, key)

assets_dist = np.asarray(assets_dist)
```

We use QuantEcon's `lorenz_curve` and `gini_coefficient` to summarize inequality:

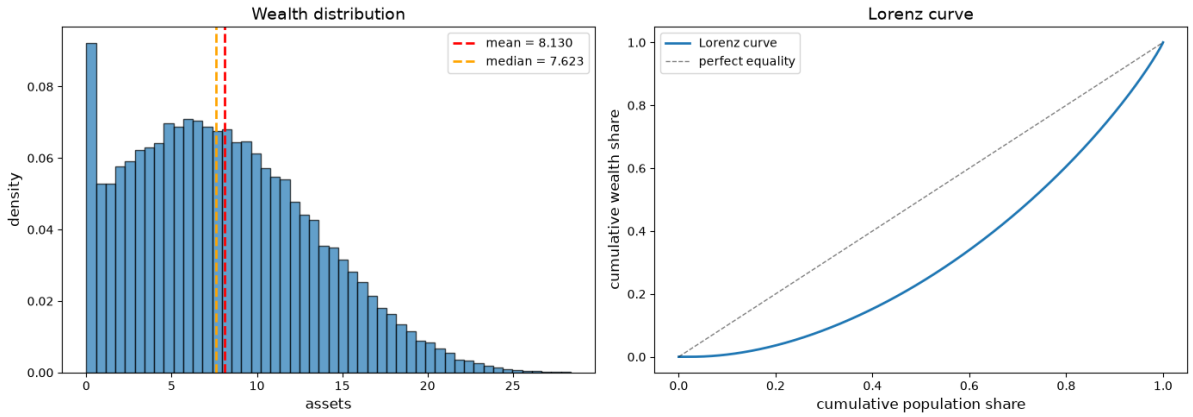
```
fig, axes = plt.subplots(1, 2, figsize=(14, 5))

# Histogram
ax = axes[0]
ax.hist(assets_dist, bins=50, density=True, alpha=0.7, edgecolor='black')
ax.axvline(np.mean(assets_dist), color='red', linestyle='--', linewidth=2,
           label=f'mean = {np.mean(assets_dist):.3f}')
ax.axvline(np.median(assets_dist), color='orange', linestyle='--', linewidth=2,
           label=f'median = {np.median(assets_dist):.3f}')
ax.set_xlabel('assets', fontsize=12)
ax.set_ylabel('density', fontsize=12)
ax.set_title('Wealth distribution', fontsize=14)
ax.legend()

# Lorenz curve
ax = axes[1]
cum_pop, cum_wealth = qe.lorenz_curve(assets_dist)
ax.plot(cum_pop, cum_wealth, lw=2, label='Lorenz curve')
ax.plot([0, 1], [0, 1], 'k--', lw=1, alpha=0.5, label='perfect equality')
ax.set_xlabel('cumulative population share', fontsize=12)
ax.set_ylabel('cumulative wealth share', fontsize=12)
ax.set_title('Lorenz curve', fontsize=14)
ax.legend()

plt.tight_layout()
plt.show()

gini = qe.gini_coefficient(assets_dist)
print(f"\nGini coefficient: {gini:.4f}")
```




```
Gini coefficient: 0.3645
```

124.7 Summary and Comparison

This lecture demonstrated how to solve the Aiyagari model using:

1. **Endogenous Grid Method (EGM)** for the household problem
    - avoids costly root-finding by working backwards from the Euler equation
    - directly computes consumption from marginal utility
    - more efficient than value function iteration
2. **Simulation** for computing aggregate capital
    - simulates a large cross-section of households
    - more flexible than analytical stationary distributions
    - allows easy computation of wealth inequality measures

124.7.1 Comparison with standard approach

Compared to the *standard Aiyagari lecture*:

Advantages:

- EGM avoids the root-finding required by value function iteration
- simulation is more flexible (works with continuous shocks, non-linear policies)
- it is easy to compute distributional statistics (Gini, percentiles, etc.)
- the approach is simpler to extend to more complex models

Disadvantages:

- simulation requires a large number of households for accuracy
- equilibrium computation is subject to Monte Carlo noise
- it is less precise than the analytical stationary distribution

124.7.2 Extensions

This framework can be easily extended to:

- continuous income shocks (e.g., lognormal)
- more complex preference specifications
- aggregate shocks and heterogeneous agent New Keynesian (HANK) models
- life-cycle models with age-dependent policies

124.8 Exercises

Exercise 124.8.1

Compare the speed and accuracy of EGM against the value function iteration approach used in the *standard Aiyagari lecture*.

1. Solve the household problem with both methods at the same prices.
2. Time both methods and compare the resulting policies.
3. Which method is faster? Are the policies close?

Exercise 124.8.2

Study how the wealth distribution changes with the discount factor β.

1. Compute equilibria for $\beta \in \{0.94, 0.95, 0.96, 0.97\}$.
2. For each β, compute and plot the wealth distribution.
3. How does the Gini coefficient change with β?
4. Explain the economic intuition.

Exercise 124.8.3

The simulation method introduced in this lecture uses a fixed number of periods. Investigate the impact of this choice:

1. Vary `num_periods` from 200 to 2000.
2. For each value, compute the mean assets multiple times with different random keys.
3. Plot the standard deviation of the capital estimate as a function of `num_periods`.
4. What is the trade-off between accuracy and computational cost?

Exercise 124.8.4

Extend the model to include a third employment state (e.g., unemployed, part-time, full-time):

1. Set up a 3-state Markov chain with appropriate transition probabilities.
2. Define income levels for each state.
3. Re-compute the equilibrium.
4. How does the additional heterogeneity affect the wealth distribution?

CHAPTER

FIVE

TRANSITIONS IN AN OVERLAPPING GENERATIONS MODEL

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install --upgrade quantecon
```

125.1 Overview

This lecture presents a life-cycle model consisting of overlapping generations of two-period lived people proposed by Peter Diamond [Diamond, 1965].

We'll present the version that was analyzed in chapter 2 of Auerbach and Kotlikoff (1987) [Auerbach and Kotlikoff, 1987].

Auerbach and Kotlikoff (1987) used their two period model as a warm-up for their analysis of overlapping generation models of long-lived people that is the main topic of their book.

Their model of two-period lived overlapping generations is a useful starting point because

- it sets forth the structure of interactions between generations of different agents who are alive at a given date
- it activates forces and tradeoffs confronting the government and successive generations of people
- it is a good laboratory for studying connections between government tax and subsidy programs and for policies for issuing and servicing government debt
- some interesting experiments involving transitions from one steady state to another can be computed by hand
- it is a good setting for illustrating a **shooting method** for solving a system of non-linear difference equations with initial and terminal condition

Note

Auerbach and Kotlikoff use computer code to calculate transition paths for their models with long-lived people.

We take the liberty of extending Auerbach and Kotlikoff's chapter 2 model to study some arrangements for redistributing resources across generations

- these take the form of a sequence of age-specific lump sum taxes and transfers

We study how these arrangements affect capital accumulation and government debt

125.2 Setting

Time is discrete and is indexed by $t = 0, 1, 2, ...$

The economy lives forever, but the people inside it do not.

At each time $t \geq 0$ a representative old person and a representative young person are alive.

At time t a representative old person coexists with a representative young person who will become an old person at time $t + 1$.

We assume that the population size is constant over time.

A young person works, saves, and consumes.

An old person dissaves and consumes, but does not work.

A government lives forever, i.e., at $t = 0, 1, 2, ...$

Each period $t \geq 0$, the government taxes, spends, transfers, and borrows.

Initial conditions set outside the model at time $t = 0$ are

- K_0 – initial capital stock brought into time $t = 0$ by a representative initial old person
- D_0 – government debt falling due at $t = 0$ and owned by a representative old person at time $t = 0$

K_0 and D_0 are both measured in units of time 0 goods.

A government **policy** consists of five sequences $\{G_t, D_t, \tau_t, \delta_{ot}, \delta_{yt}\}_{t=0}^{\infty}$ whose components are

- τ_t – flat rate tax at time t on wages and earnings from capital and government bonds
- D_t – one-period government bond principal due at time t, per capita
- G_t – government purchases of goods at time t, per capita
- δ_{yt} – a lump sum tax on each young person at time t
- δ_{ot} – a lump sum tax on each old person at time t

An **allocation** is a collection of sequences $\{C_{yt}, C_{ot}, K_{t+1}, L_t, Y_t, G_t\}_{t=0}^{\infty}$; constituents of the sequences include

- K_t – physical capital per capita
- L_t – labor per capita
- Y_t – output per capita

and also

- C_{yt} – consumption of young person at time $t \geq 0$
- C_{ot} – consumption of old person at time $t \geq 0$
- $K_{t+1} - K_t \equiv I_t$ – investment in physical capital at time $t \geq 0$
- G_t – government purchases

National income and product accounts consist of a sequence of equalities

- $Y_t = C_{yt} + C_{ot} + (K_{t+1} - K_t) + G_t, \quad t \geq 0$

A **price system** is a pair of sequences $\{W_t, r_t\}_{t=0}^{\infty}$; constituents of a price sequence include rental rates for the factors of production

- W_t – rental rate for labor at time $t \geq 0$
- r_t – rental rate for capital at time $t \geq 0$

125.3 Production

There are two factors of production, physical capital K_t and labor L_t.

Capital does not depreciate, so r_t below is a net rate of return.

Because each young person supplies one unit of labor inelastically and the population is constant, labor supply satisfies $L_t = 1$ for all t.

We use that fact to simplify formulas below.

The initial capital stock K_0 is owned by the representative initial old person, who rents it to the firm at time 0.

Net investment rate I_t at time t is

$$I_t = K_{t+1} - K_t$$

The capital stock at time t emerges from cumulating past rates of investment:

$$K_t = K_0 + \sum_{s=0}^{t-1} I_s$$

A Cobb-Douglas technology converts physical capital K_t and labor services L_t into output Y_t

$$Y_t = K_t^\alpha L_t^{1-\alpha}, \quad \alpha \in (0,1) \tag{125.1}$$

125.4 Government

At time $t-1$, the government issues one-period risk-free debt that promises to pay D_t time t goods per capita at time t.

Young people at time t purchase government debt D_{t+1} that matures at time $t+1$.

Government debt issued at t bears a before-tax net rate of interest r_{t+1} that is paid at time $t+1$.

The government budget constraint at time $t \geq 0$ is

$$D_{t+1} - D_t = r_t D_t + G_t - T_t$$

or

$$D_{t+1} = (1 + r_t) D_t + G_t - T_t. \tag{125.2}$$

Total tax collections net of transfers equal T_t and satisfy

$$T_t = \tau_t W_t L_t + \tau_t r_t (D_t + K_t) + \delta_{yt} + \delta_{ot}$$

Because one young person and one old person are alive at each date, the two lump sum taxes enter T_t with unit weights.

125.5 Activities in factor markets

Old people: At each $t \geq 0$, a representative old person

- brings K_t and D_t into the period,
- rents capital to a representative firm for $r_t K_t$,

- pays taxes $\tau_t r_t (K_t + D_t)$ on its rental and interest earnings,
- pays a lump sum tax δ_{ot} to the government,
- sells K_t to a young person.

Young people: At each $t \geq 0$, a representative young person

- sells one unit of labor services to a representative firm for W_t in wages,
- pays taxes $\tau_t W_t$ on its labor earnings
- pays a lump sum tax δ_{yt} to the government,
- spends C_{yt} on consumption,
- acquires non-negative assets A_{t+1} consisting of a sum of physical capital K_{t+1} and one-period government bonds D_{t+1} that mature at $t+1$.

Note

If a lump-sum tax is negative, it means that the government pays the person a subsidy.

125.6 Representative firm's problem

The representative firm hires labor services from young people at competitive wage rate W_t and hires capital from old people at competitive rental rate r_t.

The rental rate on capital r_t equals the interest rate on government one-period bonds.

Units of the rental rates are:

- for W_t, output at time t per unit of labor at time t
- for r_t, output at time t per unit of capital at time t

We take output at time t as **numeraire**, so the price of output at time t is one.

The firm's profits at time t are

$$K_t^\alpha L_t^{1-\alpha} - r_t K_t - W_t L_t.$$

To maximize profits a firm equates marginal products to rental rates:

$$\begin{aligned} W_t &= (1-\alpha) K_t^\alpha L_t^{-\alpha} \\ r_t &= \alpha K_t^{\alpha-1} L_t^{1-\alpha} \end{aligned} \tag{125.3}$$

Output can be consumed either by old people or young people; or sold to young people who use it to augment the capital stock; or sold to the government for uses that do not generate utility for the people in the model (i.e., "it is thrown into the ocean").

The firm thus sells output to old people, young people, and the government.

125.7 Individuals' problems

125.7.1 Initial old person

At time $t = 0$, a representative initial old person is endowed with $(1+r_0(1-\tau_0))A_0$ in initial assets, where $A_0 = K_0+D_0$.

It must pay a lump sum tax to (if positive) or receive a subsidy from (if negative) δ_{o0} the government.

An old person's budget constraint is

$$C_{o0} = (1 + r_0(1 - \tau_0))A_0 - \delta_{o0}. \tag{125.4}$$

An initial old person's utility function is C_{o0}, so the person's optimal consumption plan is provided by equation (125.4).

125.7.2 Young person

At each $t \geq 0$, a young person inelastically supplies one unit of labor and in return receives pre-tax labor earnings of W_t units of output.

A young person's post-tax-and-transfer earnings are $W_t(1 - \tau_t) - \delta_{yt}$.

At each $t \geq 0$, a young person chooses a consumption plan C_{yt}, C_{ot+1} to maximize the Cobb-Douglas utility function

$$U_t = C_{yt}^{\beta} C_{o,t+1}^{1-\beta}, \quad \beta \in (0, 1) \tag{125.5}$$

subject to the following budget constraints at times t and $t + 1$:

$$\begin{aligned} C_{yt} + A_{t+1} &= W_t(1 - \tau_t) - \delta_{yt} \\ C_{ot+1} &= (1 + r_{t+1}(1 - \tau_{t+1}))A_{t+1} - \delta_{o,t+1} \end{aligned} \tag{125.6}$$

Solving the second equation of (125.6) for savings A_{t+1} and substituting it into the first equation implies the present value budget constraint

$$C_{yt} + \frac{C_{ot+1}}{1 + r_{t+1}(1 - \tau_{t+1})} = W_t(1 - \tau_t) - \delta_{yt} - \frac{\delta_{o,t+1}}{1 + r_{t+1}(1 - \tau_{t+1})} \tag{125.7}$$

To solve the young person's choice problem, form a Lagrangian

$$\begin{aligned} \mathcal{L} &= C_{yt}^{\beta} C_{o,t+1}^{1-\beta} \\ &\quad + \lambda\Big[W_t(1 - \tau_t) - \delta_{yt} - \frac{\delta_{o,t+1}}{1 + r_{t+1}(1 - \tau_{t+1})} - C_{yt} - \frac{C_{ot+1}}{1 + r_{t+1}(1 - \tau_{t+1})}\Big], \end{aligned} \tag{125.8}$$

where $\lambda \geq 0$ is a Lagrange multiplier on the intertemporal budget constraint (125.7).

After several lines of algebra, the intertemporal budget constraint (125.7) and the first-order conditions for maximizing $\mathcal{L}$ with respect to C_{yt}, C_{ot+1} imply that an optimal consumption plan satisfies

$$\begin{aligned} C_{yt} &= \beta\Big[W_t(1 - \tau_t) - \delta_{yt} - \frac{\delta_{o,t+1}}{1 + r_{t+1}(1 - \tau_{t+1})}\Big] \\ \frac{C_{o,t+1}}{1 + r_{t+1}(1 - \tau_{t+1})} &= (1 - \beta)\Big[W_t(1 - \tau_t) - \delta_{yt} - \frac{\delta_{o,t+1}}{1 + r_{t+1}(1 - \tau_{t+1})}\Big] \end{aligned} \tag{125.9}$$

The first-order condition for the Lagrangian (125.8) with respect to the Lagrange multiplier λ recovers the budget constraint (125.7), which, using (125.9) gives the optimal savings plan

$$A_{t+1} = (1 - \beta)[(1 - \tau_t)W_t - \delta_{yt}] + \beta\frac{\delta_{o,t+1}}{1 + r_{t+1}(1 - \tau_{t+1})} \tag{125.10}$$

125.8 Equilibrium

Definition: An equilibrium is an allocation, a government policy, and a price system with the properties that

- given the price system and the government policy, the allocation solves
  - representative firms' problems for $t \geq 0$
  - individual persons' problems for $t \geq 0$
- given the price system and the allocation, the government budget constraint is satisfied for all $t \geq 0$.

125.9 Next steps

To begin our analysis of equilibrium outcomes, we'll study the special case of the model with which Auerbach and Kotlikoff (1987) [Auerbach and Kotlikoff, 1987] began their analysis in chapter 2.

It can be solved by hand.

We shall do that next.

After we derive a closed form solution, we'll pretend that we don't know and will compute equilibrium outcome paths.

We'll do that by first formulating an equilibrium as a fixed point of a mapping from sequences of factor prices and tax rates to sequences of factor prices and tax rates.

We'll compute an equilibrium by iterating to convergence on that mapping.

125.10 Closed form solution

To get the special chapter 2 case of Auerbach and Kotlikoff (1987) [Auerbach and Kotlikoff, 1987], we set both δ_{ot} and δ_{yt} to zero.

As our special case of (125.9), we compute the following consumption-savings plan for a representative young person:

$$\begin{aligned} C_{yt} &= \beta(1-\tau_t)W_t \\ A_{t+1} &= (1-\beta)(1-\tau_t)W_t \end{aligned}$$

Using (125.3) and $A_t = K_t + D_t$, we obtain the following closed form transition law for capital:

$$K_{t+1} = K_t^{\alpha}\left(1-\tau_t\right)\left(1-\alpha\right)\left(1-\beta\right) - D_{t+1} \tag{125.11}$$

The debt term is dated $t+1$ because the assets A_{t+1} that a young person carries into time $t+1$ consist of physical capital K_{t+1} *and* the government bonds D_{t+1} that mature at $t+1$.

125.10.1 Steady states

From (125.11) and the government budget constraint (125.2), we compute **time-invariant** or **steady state values** $\hat{K}, \hat{D}, \hat{T}$:

$$\begin{aligned} \hat{K} &= \hat{K}\left(1-\hat{\tau}\right)\left(1-\alpha\right)\left(1-\beta\right) - \hat{D} \\ \hat{D} &= (1+\hat{r})\hat{D} + \hat{G} - \hat{T} \\ \hat{T} &= \hat{\tau}\hat{Y} + \hat{\tau}\hat{r}\hat{D}. \end{aligned} \tag{125.12}$$

These imply

$$\hat{\tau} = \frac{\hat{G} + \hat{r}\hat{D}}{\hat{Y} + \hat{r}\hat{D}}$$

When $\hat{D} = 0$, the first equation of (125.12) can be solved in closed form for

$$\hat{K} = \left[(1 - \hat{\tau})(1 - \alpha)(1 - \beta)\right]^{\frac{1}{1-\alpha}}$$

When $\hat{D} \neq 0$, that equation must be solved numerically, a task that Exercise 125.12.1 takes up.

Let's take an example in which

1. there is no initial government debt, $D_t = 0$,
2. government consumption G_t equals 15% of output Y_t

Our formulas for steady-state values tell us that

$$\begin{aligned}\hat{D} &= 0 \\ \hat{G} &= 0.15\hat{Y} \\ \hat{\tau} &= 0.15\end{aligned}$$

125.10.2 Implementation

```
import numpy as np
import matplotlib.pyplot as plt
from numba import jit
from quantecon.optimize import brent_max
```

For parameters $\alpha = 0.3$ and $\beta = 0.5$, let's compute $\hat{K}$:

```
# parameters
α = 0.3
β = 0.5

# steady states of τ and D
τ_hat = 0.15
D_hat = 0.

# solve for steady state of K
K_hat = ((1 - τ_hat) * (1 - α) * (1 - β)) ** (1 / (1 - α))
K_hat
```

```
0.17694509514972878
```

Knowing $\hat{K}$, we can calculate other equilibrium objects.

Let's first define some Python helper functions.

```
@jit
def K_to_Y(K, α):

    return K ** α

@jit
```

(continues on next page)

(continued from previous page)

```
def K_to_r(K, α):

    return α * K ** (α - 1)

@jit
def K_to_W(K, α):

    return (1 - α) * K ** α

@jit
def K_to_C(K, D, τ, r, α, β):

    # optimal consumption for the old when δ=0
    A = K + D
    Co = A * (1 + r * (1 - τ))

    # optimal consumption for the young when δ=0
    W = K_to_W(K, α)
    Cy = β * W * (1 - τ)

    return Cy, Co
```

We can use these helper functions to obtain steady state values $\hat{Y}$, $\hat{r}$, and $\hat{W}$ associated with steady state values $\hat{K}$ and $\hat{r}$.

```
Y_hat, r_hat, W_hat = K_to_Y(K_hat, α), K_to_r(K_hat, α), K_to_W(K_hat, α)
Y_hat, r_hat, W_hat
```

```
(0.5947734290747186, 1.0084033613445376, 0.41634140035230305)
```

The steady state net rate of return $\hat{r}$ is close to one, i.e., about 100 percent.

That is not a mistake: a period in this model is half of a lifetime, so $\hat{r}$ cumulates returns over a generation.

Since steady state government debt $\hat{D}$ is 0, all taxes are used to pay for government expenditures

```
G_hat = τ_hat * Y_hat
G_hat
```

```
0.0892160143612078
```

We use the optimal consumption plans to find steady state consumptions for young and old

```
Cy_hat, Co_hat = K_to_C(K_hat, D_hat, τ_hat, r_hat, α, β)
Cy_hat, Co_hat
```

```
(0.17694509514972878, 0.32861231956378195)
```

Let's store the steady state quantities and prices using an array called `init_ss`

```
init_ss = np.array([K_hat, Y_hat, Cy_hat, Co_hat,      # quantities
                    W_hat, r_hat,                      # prices
                    τ_hat, D_hat, G_hat                # policies
                    ])
```

125.10.3 Transitions

We have computed a steady state in which the government policy sequences are each constant over time.

We'll use this steady state as an initial condition at time $t = 0$ for another economy in which government policy sequences are with time-varying sequences.

To make sense of our calculation, we'll treat $t = 0$ as time when a huge unanticipated shock occurs in the form of

- a time-varying government policy sequences that disrupts an original steady state
- new government policy sequences are eventually time-invariant in the sense that after some date $T > 0$, each sequence is constant over time.
- sudden revelation of a new government policy in the form of sequences starting at time $t = 0$

We assume that everyone, including old people at time $t = 0$, knows the new government policy sequence and chooses accordingly.

As the capital stock and other aggregates adjust to the fiscal policy change over time, the economy will approach a new steady state.

We can find a transition path from an old steady state to a new steady state by employing a fixed-point algorithm in a space of sequences.

But in our special case with its closed form solution, we have available a simpler and faster approach.

Here we define a Python class `ClosedFormTrans` that computes length T transition path in response to a particular fiscal policy change.

We choose T large enough so that we have gotten very close to a new steady state after T periods.

The class takes three keyword arguments, `τ_pol`, `D_pol`, and `G_pol`.

These are sequences of tax rate, government debt level, and government purchases, respectively.

In each policy experiment below, we will pass two out of three as inputs required to depict a fiscal policy.

We'll then compute the single remaining undetermined policy variable from the government budget constraint.

When we simulate transition paths, it is useful to distinguish **state variables** at time t such as K_t, Y_t, D_t, W_t, r_t from **control variables** that include $C_{yt}, C_{ot}, \tau_t, G_t$.

```
class ClosedFormTrans:
    """
    This class simulates length T transitional path of a economy
    in response to a fiscal policy change given its initial steady
    state. The simulation is based on the closed form solution when
    the lump sum taxations are absent.

    """

    def __init__(self, α, β):

        self.α, self.β = α, β

    def simulate(self,
                 T,           # length of transitional path to simulate
                 init_ss,     # initial steady state
                 τ_pol=None,  # sequence of tax rates
                 D_pol=None,  # sequence of government debt levels
                 G_pol=None): # sequence of government purchases
```

(continues on next page)

(continued from previous page)

```
        α, β = self.α, self.β

        # unpack the steady state variables
        K_hat, Y_hat, Cy_hat, Co_hat = init_ss[:4]
        W_hat, r_hat = init_ss[4:6]
        τ_hat, D_hat, G_hat = init_ss[6:9]

        # initialize array containers
        # K, Y, Cy, Co
        quant_seq = np.empty((T+1, 4))

        # W, r
        price_seq = np.empty((T+1, 2))

        # τ, D, G
        policy_seq = np.empty((T+2, 3))

        # t=0, starting from steady state
        K0, Y0 = K_hat, Y_hat
        W0, r0 = W_hat, r_hat
        D0 = D_hat

        # fiscal policy
        if τ_pol is None:
            D1 = D_pol[1]
            G0 = G_pol[0]
            τ0 = (G0 + (1 + r0) * D0 - D1) / (Y0 + r0 * D0)
        elif D_pol is None:
            τ0 = τ_pol[0]
            G0 = G_pol[0]
            D1 = (1 + r0) * D0 + G0 - τ0 * (Y0 + r0 * D0)
        elif G_pol is None:
            D1 = D_pol[1]
            τ0 = τ_pol[0]
            G0 = τ0 * (Y0 + r0 * D0) + D1 - (1 + r0) * D0

        # optimal consumption plans
        Cy0, Co0 = K_to_C(K0, D0, τ0, r0, α, β)

        # t=0 economy
        quant_seq[0, :] = K0, Y0, Cy0, Co0
        price_seq[0, :] = W0, r0
        policy_seq[0, :] = τ0, D0, G0
        policy_seq[1, 1] = D1

        # starting from t=1 to T
        for t in range(1, T+1):

            # transition of K
            K_old, τ_old = quant_seq[t-1, 0], policy_seq[t-1, 0]
            D = policy_seq[t, 1]
            K = K_old ** α * (1 - τ_old) * (1 - α) * (1 - β) - D

            # output, capital return, wage
            Y, r, W = K_to_Y(K, α), K_to_r(K, α), K_to_W(K, α)

            # to satisfy the government budget constraint
```

(continues on next page)

(continued from previous page)

```
            if τ_pol is None:
                D = D_pol[t]
                D_next = D_pol[t+1]
                G = G_pol[t]
                τ = (G + (1 + r) * D - D_next) / (Y + r * D)
            elif D_pol is None:
                τ = τ_pol[t]
                G = G_pol[t]
                D = policy_seq[t, 1]
                D_next = (1 + r) * D + G - τ * (Y + r * D)
            elif G_pol is None:
                D = D_pol[t]
                D_next = D_pol[t+1]
                τ = τ_pol[t]
                G = τ * (Y + r * D) + D_next - (1 + r) * D

            # optimal consumption plans
            Cy, Co = K_to_C(K, D, τ, r, α, β)

            # store time t economy aggregates
            quant_seq[t, :] = K, Y, Cy, Co
            price_seq[t, :] = W, r
            policy_seq[t, 0] = τ
            policy_seq[t+1, 1] = D_next
            policy_seq[t, 2] = G

        self.T, self.init_ss = T, init_ss
        self.quant_seq = quant_seq
        self.price_seq = price_seq
        self.policy_seq = policy_seq

        return quant_seq, price_seq, policy_seq

    def plot(self):

        T, init_ss = self.T, self.init_ss
        quant_seq = self.quant_seq
        price_seq = self.price_seq
        policy_seq = self.policy_seq

        fig, axs = plt.subplots(3, 3, figsize=(14, 10))

        # quantities
        for i, name in enumerate(['K', 'Y', 'Cy', 'Co']):
            ax = axs[i//3, i%3]
            ax.plot(range(T+1), quant_seq[:T+1, i], label=name)
            ax.hlines(init_ss[i], 0, T+1, color='r', linestyle='--')
            ax.legend()
            ax.set_xlabel('t')

        # prices
        for i, name in enumerate(['W', 'r']):
            ax = axs[(i+4)//3, (i+4)%3]
            ax.plot(range(T+1), price_seq[:T+1, i], label=name)
            ax.hlines(init_ss[i+4], 0, T+1, color='r', linestyle='--')
            ax.legend()
            ax.set_xlabel('t')
```

(continues on next page)

(continued from previous page)

```
        # policies
        for i, name in enumerate(['τ', 'D', 'G']):
            ax = axs[(i+6)//3, (i+6)%3]
            ax.plot(range(T+1), policy_seq[:T+1, i], label=name)
            ax.hlines(init_ss[i+6], 0, T+1, color='r', linestyle='--')
            ax.legend()
            ax.set_xlabel('t')
```

We can create an instance `closed` for model parameters $\{\alpha, \beta\}$ and use it for various fiscal policy experiments.

```
closed = ClosedFormTrans(α, β)
```

125.10.4 Experiment 1: a tax cut

To illustrate the power of `ClosedFormTrans`, let's first experiment with the following fiscal policy change:

1. at $t = 0$, the government unexpectedly announces a one-period tax cut, $\tau_0 = (1 - \frac{1}{3})\hat{\tau}$, by issuing government debt $\bar{D}$
2. from $t = 1$, the government will keep $D_t = \bar{D}$ and adjust τ_t to collect taxation to pay for the government consumption and interest payments on the debt
3. government consumption G_t will be fixed at $0.15\hat{Y}$

The following equations completely characterize the equilibrium transition path originating from the initial steady state

$$
\begin{aligned}
K_{t+1} &= K_t^{\alpha}\left(1-\tau_t\right)\left(1-\alpha\right)\left(1-\beta\right) - \bar{D} \\
\tau_0 &= (1 - \frac{1}{3})\hat{\tau} \\
\bar{D} &= \hat{G} - \tau_0 \hat{Y} \\
\tau_t &= \frac{\hat{G} + r_t \bar{D}}{\hat{Y} + r_t \bar{D}}
\end{aligned}
$$

We can simulate the transition for 20 periods, after which the economy will be close to a new steady state.

The first step is to prepare sequences of policy variables that describe fiscal policy.

We must define sequences of government expenditure $\{G_t\}_{t=0}^{T}$ and debt level $\{D_t\}_{t=0}^{T+1}$ in advance, then pass them to the solver.

```
T = 20

# tax cut
τ0 = τ_hat * (1 - 1/3)

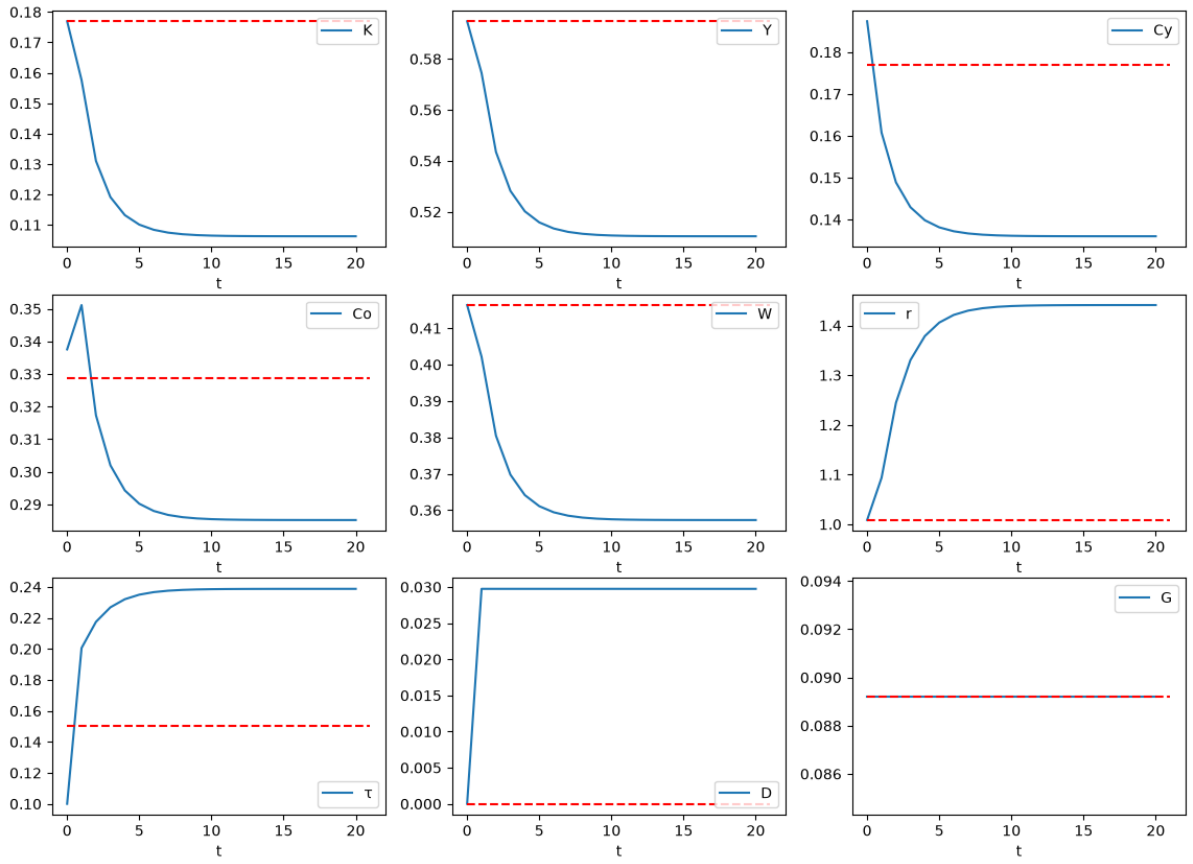
# sequence of government purchase
G_seq = τ_hat * Y_hat * np.ones(T+1)

# sequence of government debt
D_bar = G_hat - τ0 * Y_hat
D_seq = np.ones(T+2) * D_bar
D_seq[0] = D_hat
```

Let's use the `simulate` method of `closed` to compute dynamic transitions.

Note that we leave `τ_pol` as `None`, since the tax rates need to be determined to satisfy the government budget constraint.

```
quant_seq1, price_seq1, policy_seq1 = closed.simulate(T, init_ss,
                                                      D_pol=D_seq,
                                                      G_pol=G_seq)
closed.plot()
```



We can also experiment with a lower tax cut rate, such as 0.2.

```
# lower tax cut rate
τ0 = 0.15 * (1 - 0.2)

# the corresponding debt sequence
D_bar = G_hat - τ0 * Y_hat
D_seq = np.ones(T+2) * D_bar
D_seq[0] = D_hat

quant_seq2, price_seq2, policy_seq2 = closed.simulate(T, init_ss,
                                                      D_pol=D_seq,
                                                      G_pol=G_seq)
```

```
fig, axs = plt.subplots(3, 3, figsize=(14, 10))

# quantities
for i, name in enumerate(['K', 'Y', 'Cy', 'Co']):
    ax = axs[i//3, i%3]
    ax.plot(range(T+1), quant_seq1[:T+1, i], label=f'{name}, 1/3')
    ax.plot(range(T+1), quant_seq2[:T+1, i], label=f'{name}, 0.2')
```

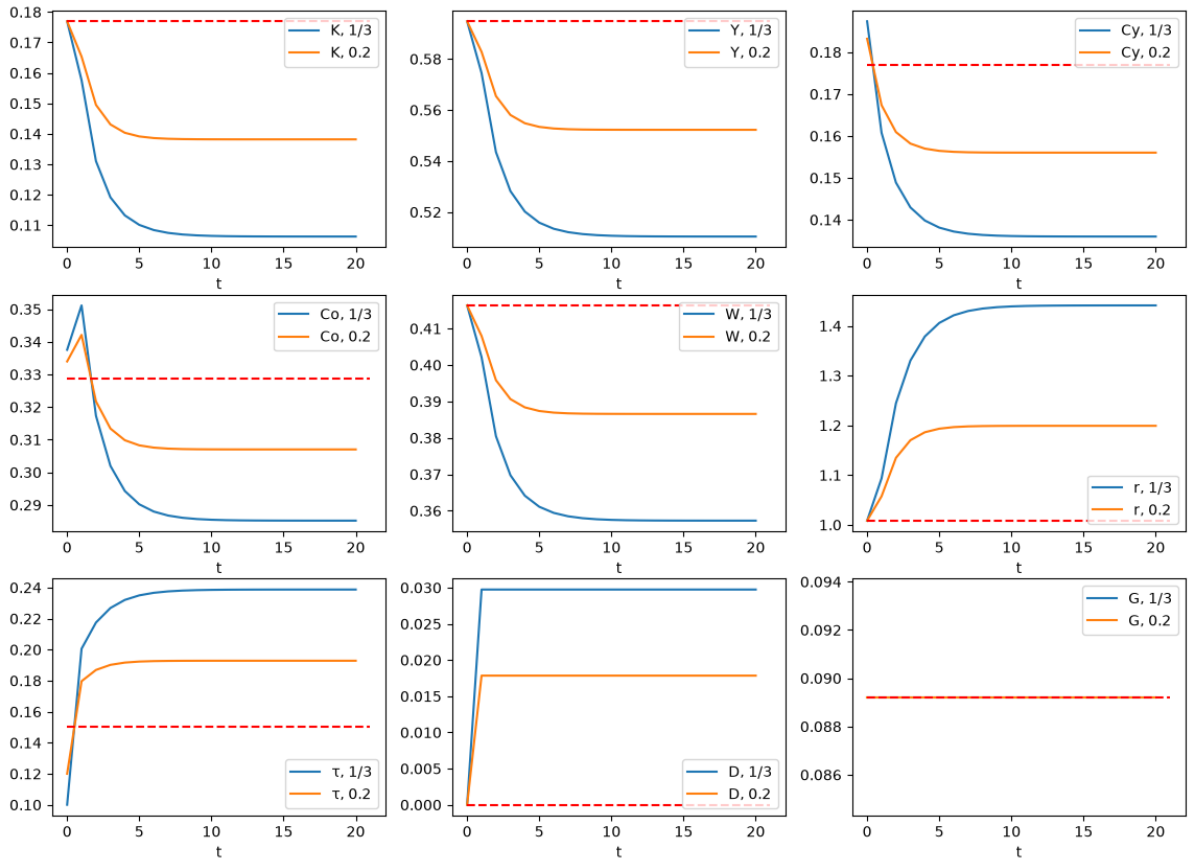
(continues on next page)

(continued from previous page)

```
    ax.hlines(init_ss[i], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')

# prices
for i, name in enumerate(['W', 'r']):
    ax = axs[(i+4)//3, (i+4)%3]
    ax.plot(range(T+1), price_seq1[:T+1, i], label=f'{name}, 1/3')
    ax.plot(range(T+1), price_seq2[:T+1, i], label=f'{name}, 0.2')
    ax.hlines(init_ss[i+4], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')

# policies
for i, name in enumerate(['τ', 'D', 'G']):
    ax = axs[(i+6)//3, (i+6)%3]
    ax.plot(range(T+1), policy_seq1[:T+1, i], label=f'{name}, 1/3')
    ax.plot(range(T+1), policy_seq2[:T+1, i], label=f'{name}, 0.2')
    ax.hlines(init_ss[i+6], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')
```



The economy with lower tax cut rate at $t = 0$ has the same transitional pattern, but is less distorted, and it converges to a new steady state with higher physical capital stock.

125.10.5 Experiment 2: government asset accumulation

Assume that the economy is initially in the same steady state.

Now the government promises to cut its spending on services and goods by half $\forall t \geq 0$.

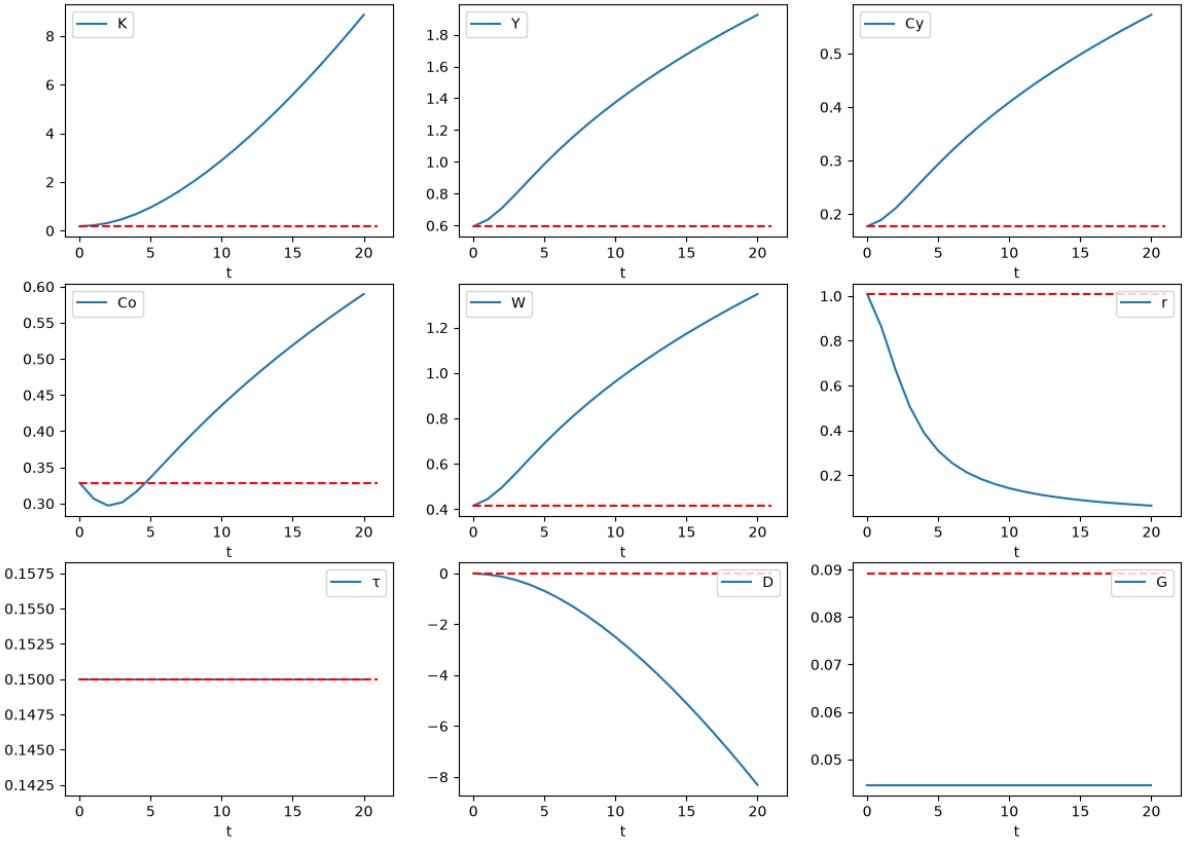
The government targets the same tax rate $\tau_t = \hat{\tau}$ and to accumulate assets $-D_t$ over time.

To conduct this experiment, we pass `τ_seq` and `G_seq` as inputs and let `D_pol` be determined along the path by satisfying the government budget constraint.

```
# government expenditure cut by a half
G_seq = τ_hat * 0.5 * Y_hat * np.ones(T+1)

# targeted tax rate
τ_seq = τ_hat * np.ones(T+1)

closed.simulate(T, init_ss, τ_pol=τ_seq, G_pol=G_seq);
closed.plot()
```



As the government accumulates assets, it supplies funds that would otherwise have been supplied by private savers, so the capital stock rises and the rental rate on capital falls.

As a result, the ratio $-\frac{D_t}{K_t}$ of the government asset to physical capital used in production will increase over time

```
plt.plot(range(T+1), -closed.policy_seq[:-1, 1] / closed.quant_seq[:, 0])
plt.xlabel('t')
plt.ylabel('-D/K')
plt.show()
```

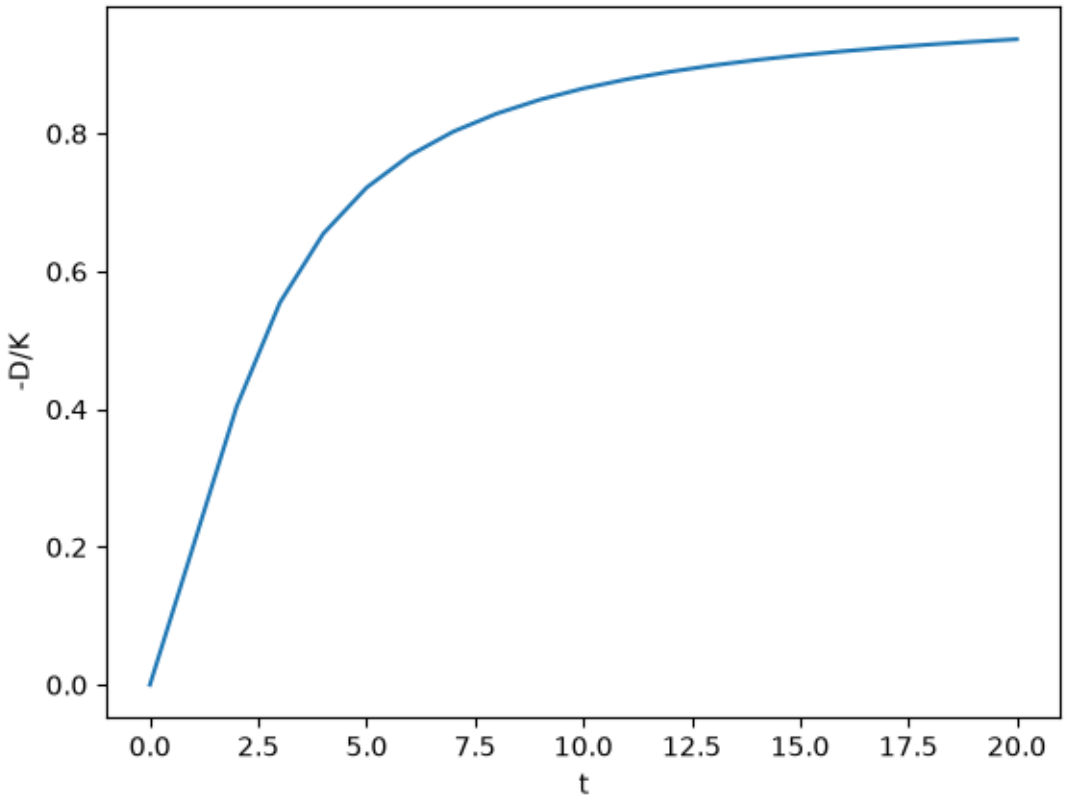


Fig. 125.1: Ratio of government assets to capital

We want to know how this policy experiment affects individuals.

In the long run, future cohorts will enjoy higher consumption throughout their lives because they will earn higher labor income when they work.

In the short run, early cohorts lose.

The representative old person at $t = 0$ is unaffected, because K_0, r_0, and τ_0 are all unchanged at $t = 0$.

But the cohorts that are young at $t = 0$ and $t = 1$ save at wages that the policy has not yet raised, and they collect their returns at the much lower interest rates that the larger capital stock brings.

Exercise 125.12.2 computes these gains and losses cohort by cohort.

Such distinct long run and short run effects motivate us to study transition paths.

Note

Although the consumptions in the new steady state are strictly higher, it is at a cost of fewer public services and goods.

125.10.6 Experiment 3: temporary expenditure cut

Let's now investigate a scenario in which the government again accumulates assets by cutting its spending.

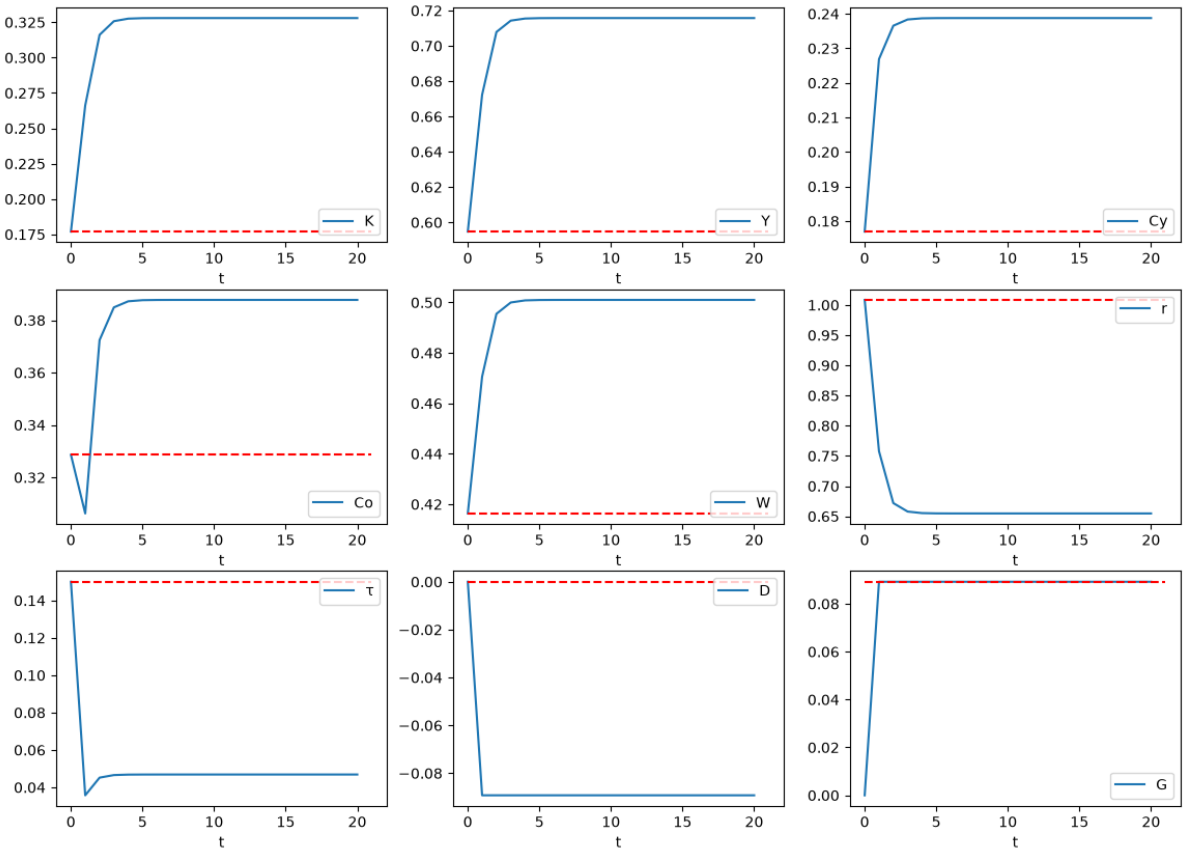
But now the cut lasts for one period only: the government sets $G_0 = 0$ and returns its purchases to $\hat{G}$ from $t = 1$ onwards.

From $t \geq 1$, government expenditures return to $\hat{G}$ and τ_t adjusts to maintain the asset level $-D_t = -D_1$.

```
# sequence of government purchase
G_seq = τ_hat * Y_hat * np.ones(T+1)
G_seq[0] = 0

# sequence of government debt
D_bar = G_seq[0] - τ_hat * Y_hat
D_seq = D_bar * np.ones(T+2)
D_seq[0] = D_hat

closed.simulate(T, init_ss, D_pol=D_seq, G_pol=G_seq);
closed.plot()
```



The economy quickly converges to a new steady state with higher physical capital stock, lower interest rate, higher wage rate, and higher consumptions for both the young and the old.

Even though government expenditure G_t returns to its high initial level from $t \geq 1$, the government can balance the budget at a lower tax rate because it gathers additional revenue $-r_t D_t$ from the asset accumulated during the temporary cut in the spendings.

As in *Experiment 2: government asset accumulation*, the cohorts that are young early in the transition lose from this policy shock, because they save before wages rise and earn the lower interest rates that follow.

125.11 A computational strategy

With the preceding calculations, we studied dynamic transitions instigated by alternative fiscal policies.

In all these experiments, we maintained the assumption that lump sum taxes were absent so that $\delta_{yt} = 0, \delta_{ot} = 0$.

In this section, we investigate the transition dynamics when the lump sum taxes are present.

The government will use lump sum taxes and transfers to redistribute resources across successive generations.

Including lump sum taxes disrupts closed form solution because of how they make optimal consumption and saving plans depend on future prices and tax rates.

Therefore, we compute equilibrium transitional paths by finding a fixed point of a mapping from sequences to sequences.

- that fixed point pins down an equilibrium

To set the stage for the entry of the mapping whose fixed point we seek, we return to concepts introduced in section *Equilibrium*.

Definition: Given parameters $\{\alpha, \beta\}$, a competitive equilibrium consists of

- sequences of optimal consumptions $\{C_{yt}, C_{ot}\}$
- sequences of prices $\{W_t, r_t\}$
- sequences of capital stock and output $\{K_t, Y_t\}$
- sequences of tax rates, government assets (debt), government purchases, and lump sum taxes $\{\tau_t, D_t, G_t, \delta_{yt}, \delta_{ot}\}$

with the properties that

- given the price system and government fiscal policy, consumption plans are optimal
- the government budget constraints are satisfied for all t

An equilibrium transition path can be computed by "guessing and verifying" some endogenous sequences.

In our *Experiment 1: a tax cut* example, sequences $\{D_t\}_{t=0}^T$ and $\{G_t\}_{t=0}^T$ are exogenous.

In addition, we assume that the lump sum taxes $\{\delta_{yt}, \delta_{ot}\}_{t=0}^T$ are given and known to everybody inside the model.

We can solve for sequences of other equilibrium sequences following the steps below

1. guess prices $\{W_t, r_t\}_{t=0}^T$ and tax rates $\{\tau_t\}_{t=0}^T$
2. solve for optimal consumption and saving plans $\{C_{yt}, C_{ot}\}_{t=0}^T$, treating the guesses of future prices and taxes as true
3. solve for transition of the capital stock $\{K_t\}_{t=0}^T$
4. update the guesses for prices and tax rates with the values implied by the equilibrium conditions
5. iterate until convergence

Let's implement this "guess and verify" approach

We start by defining the Cobb-Douglas utility function

```
@jit
def U(Cy, Co, β):

    return (Cy ** β) * (Co ** (1-β))
```

We use `Cy_val` to compute the lifetime value of an arbitrary consumption plan, C_y, given the intertemporal budget constraint.

Note that it requires knowing future prices r_{t+1} and tax rate τ_{t+1}.

```
@jit
def Cy_val(Cy, W, r_next, τ, τ_next, δy, δo_next, β):

    # Co given by the budget constraint
    Co = (W * (1 - τ) - δy - Cy) * (1 + r_next * (1 - τ_next)) - δo_next

    return U(Cy, Co, β)
```

An optimal consumption plan C_y^* can be found by maximizing `Cy_val`.

Here is an example that computes optimal consumption $C_y^* = \hat{C}_y$ in the steady state with $\delta_{yt} = \delta_{ot} = 0$, like one that we studied earlier

```
W, r_next, τ, τ_next = W_hat, r_hat, τ_hat, τ_hat
δy, δo_next = 0, 0

Cy_opt, U_opt, _ = brent_max(Cy_val,             # maximand
                             1e-6,               # lower bound
                             W*(1-τ)-δy-1e-6,    # upper bound
                             args=(W, r_next, τ, τ_next, δy, δo_next, β))

Cy_opt, U_opt
```

```
(0.17694509514972878, 0.241135518231111)
```

Let's define a Python class `AK2` that computes the transition paths with the fixed-point algorithm.

It can handle nonzero lump sum taxes

```
class AK2():
    """
    This class simulates length T transitional path of a economy
    in response to a fiscal policy change given its initial steady
    state. The transitional path is found by employing a fixed point
    algorithm to satisfy the equilibrium conditions.

    """

    def __init__(self, α, β):

        self.α, self.β = α, β

    def simulate(self,
                 T,           # length of transitional path to simulate
                 init_ss,     # initial steady state
                 δy_seq,      # sequence of lump sum tax for the young
                 δo_seq,      # sequence of lump sum tax for the old
                 τ_pol=None,  # sequence of tax rates
                 D_pol=None,  # sequence of government debt levels
                 G_pol=None,  # sequence of government purchases
                 verbose=False,
                 max_iter=500,
                 tol=1e-5):

        α, β = self.α, self.β

        # unpack the steady state variables
        K_hat, Y_hat, Cy_hat, Co_hat = init_ss[:4]
        W_hat, r_hat = init_ss[4:6]
        τ_hat, D_hat, G_hat = init_ss[6:9]

        # K, Y, Cy, Co
        quant_seq = np.empty((T+2, 4))

        # W, r
        price_seq = np.empty((T+2, 2))

        # τ, D, G
        policy_seq = np.empty((T+2, 3))
        policy_seq[:, 1] = D_pol
```

(continues on next page)

(continued from previous page)

```
        policy_seq[:, 2] = G_pol

        # initial guesses of prices
        price_seq[:, 0] = np.ones(T+2) * W_hat
        price_seq[:, 1] = np.ones(T+2) * r_hat

        # initial guesses of policies
        policy_seq[:, 0] = np.ones(T+2) * τ_hat

        # t=0, starting from steady state
        quant_seq[0, :2] = K_hat, Y_hat

        if verbose:
            # prepare to plot iterations until convergence
            fig, axs = plt.subplots(1, 3, figsize=(14, 4))

        # containers for checking convergence
        price_seq_old = np.empty_like(price_seq)
        policy_seq_old = np.empty_like(policy_seq)

        # start iteration
        i_iter = 0
        while True:

            if verbose:
                # plot current prices at ith iteration
                for i, name in enumerate(['W', 'r']):
                    axs[i].plot(range(T+1), price_seq[:T+1, i])
                    axs[i].set_title(name)
                    axs[i].set_xlabel('t')
                axs[2].plot(range(T+1), policy_seq[:T+1, 0],
                            label=f'{i_iter}th iteration')
                axs[2].legend(bbox_to_anchor=(1.05, 1), loc='upper left')
                axs[2].set_title('τ')
                axs[2].set_xlabel('t')

            # store old prices from last iteration
            price_seq_old[:] = price_seq
            policy_seq_old[:] = policy_seq

            # start updating quantities and prices
            for t in range(T+1):
                K, Y = quant_seq[t, :2]
                W, r = price_seq[t, :]
                r_next = price_seq[t+1, 1]
                τ, D, G = policy_seq[t, :]
                τ_next, D_next, G_next = policy_seq[t+1, :]
                δy, δo = δy_seq[t], δo_seq[t]
                δy_next, δo_next = δy_seq[t+1], δo_seq[t+1]

                # consumption for the old
                Co = (1 + r * (1 - τ)) * (K + D) - δo

                # optimal consumption for the young
                out = brent_max(Cy_val, 1e-6, W*(1-τ)-δy-1e-6,
                                args=(W, r_next, τ, τ_next,
                                      δy, δo_next, β))
```

(continues on next page)

(continued from previous page)

```
                Cy = out[0]

                quant_seq[t, 2:] = Cy, Co
                τ_num = ((1 + r) * D + G - D_next - δy - δo)
                τ_denom = (Y + r * D)
                policy_seq[t, 0] = τ_num / τ_denom

                # saving of the young
                A_next = W * (1 - τ) - δy - Cy

                # transition of K
                K_next = A_next - D_next
                Y_next = K_to_Y(K_next, α)

                # implied next-period prices
                W_imp, r_imp = K_to_W(K_next, α), K_to_r(K_next, α)

                quant_seq[t+1, :2] = K_next, Y_next
                price_seq[t+1, :] = W_imp, r_imp

            i_iter += 1

            if (np.max(np.abs(price_seq_old - price_seq)) < tol) & \
               (np.max(np.abs(policy_seq_old - policy_seq)) < tol):
                if verbose:
                    print(f"Converge using {i_iter} iterations")
                break

            if i_iter > max_iter:
                if verbose:
                    print(f"Fail to converge using {i_iter} iterations")
                break

        self.T, self.init_ss = T, init_ss
        self.quant_seq = quant_seq
        self.price_seq = price_seq
        self.policy_seq = policy_seq

        return quant_seq, price_seq, policy_seq

    def plot(self):

        T, init_ss = self.T, self.init_ss
        quant_seq = self.quant_seq
        price_seq = self.price_seq
        policy_seq = self.policy_seq

        fig, axs = plt.subplots(3, 3, figsize=(14, 10))

        # quantities
        for i, name in enumerate(['K', 'Y', 'Cy', 'Co']):
            ax = axs[i//3, i%3]
            ax.plot(range(T+1), quant_seq[:T+1, i], label=name)
            ax.hlines(init_ss[i], 0, T+1, color='r', linestyle='--')
            ax.legend()
            ax.set_xlabel('t')
```

(continues on next page)

(continued from previous page)

```
        # prices
        for i, name in enumerate(['W', 'r']):
            ax = axs[(i+4)//3, (i+4)%3]
            ax.plot(range(T+1), price_seq[:T+1, i], label=name)
            ax.hlines(init_ss[i+4], 0, T+1, color='r', linestyle='--')
            ax.legend()
            ax.set_xlabel('t')

        # policies
        for i, name in enumerate(['τ', 'D', 'G']):
            ax = axs[(i+6)//3, (i+6)%3]
            ax.plot(range(T+1), policy_seq[:T+1, i], label=name)
            ax.hlines(init_ss[i+6], 0, T+1, color='r', linestyle='--')
            ax.legend()
            ax.set_xlabel('t')
```

We can initialize an instance of class `AK2` with model parameters $\{\alpha, \beta\}$ and then use it to conduct fiscal policy experiments.

```
ak2 = AK2(α, β)
```

We first examine that the "guess and verify" method leads to the same numerical results as we obtain with the closed form solution when lump sum taxes are muted

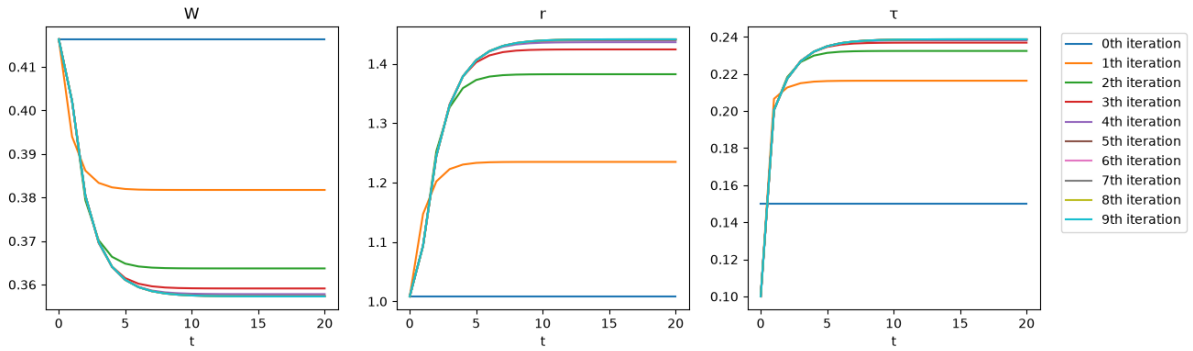
```
δy_seq = np.ones(T+2) * 0.
δo_seq = np.ones(T+2) * 0.

D_pol = np.zeros(T+2)
G_pol = np.ones(T+2) * G_hat

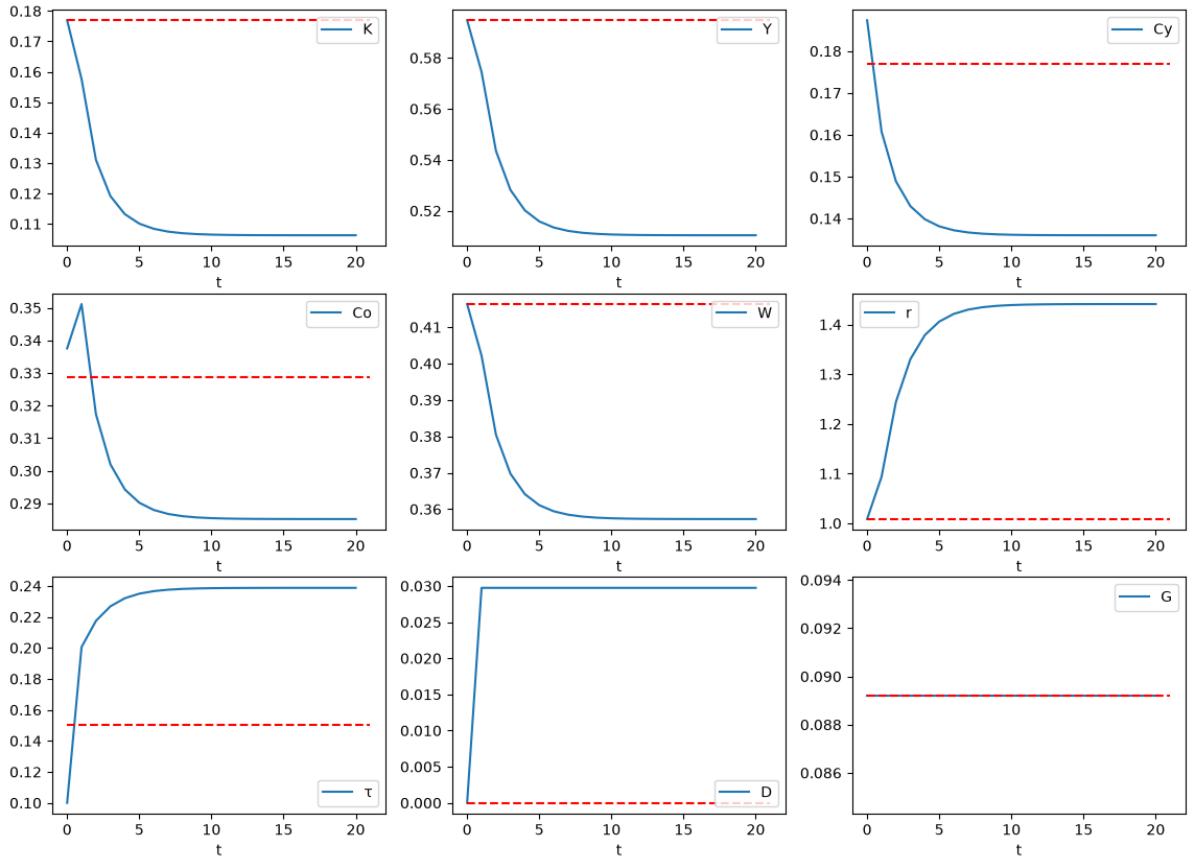
# tax cut
τ0 = τ_hat * (1 - 1/3)
D1 = D_hat * (1 + r_hat * (1 - τ0)) + G_hat - τ0 * Y_hat - δy_seq[0] - δo_seq[0]
D_pol[0] = D_hat
D_pol[1:] = D1
```

```
quant_seq3, price_seq3, policy_seq3 = ak2.simulate(T, init_ss,
                                                   δy_seq, δo_seq,
                                                   D_pol=D_pol, G_pol=G_pol,
                                                   verbose=True)
```

```
Converge using 10 iterations
```



```
ak2.plot()
```



Next, we activate lump sum taxes.

Let's alter our *Experiment 1: a tax cut* fiscal policy experiment by assuming that the government also increases lump sum taxes for both young and old people $\delta_{yt} = \delta_{ot} = 0.005, t \geq 0$.

```
δy_seq = np.ones(T+2) * 0.005
δo_seq = np.ones(T+2) * 0.005

D1 = D_hat * (1 + r_hat * (1 - τ0)) + G_hat - τ0 * Y_hat - δy_seq[0] - δo_seq[0]
D_pol[1:] = D1

quant_seq4, price_seq4, policy_seq4 = ak2.simulate(T, init_ss,
                                                   δy_seq, δo_seq,
                                                   D_pol=D_pol, G_pol=G_pol)
```

Note how "crowding out" has been mitigated.

```
fig, axs = plt.subplots(3, 3, figsize=(14, 10))

# quantities
for i, name in enumerate(['K', 'Y', 'Cy', 'Co']):
    ax = axs[i//3, i%3]
    ax.plot(range(T+1), quant_seq3[:T+1, i], label=rf'{name}, $\delta$s=0')
    ax.plot(range(T+1), quant_seq4[:T+1, i], label=rf'{name}, $\delta$s=0.005')
    ax.hlines(init_ss[i], 0, T+1, color='r', linestyle='--')
```

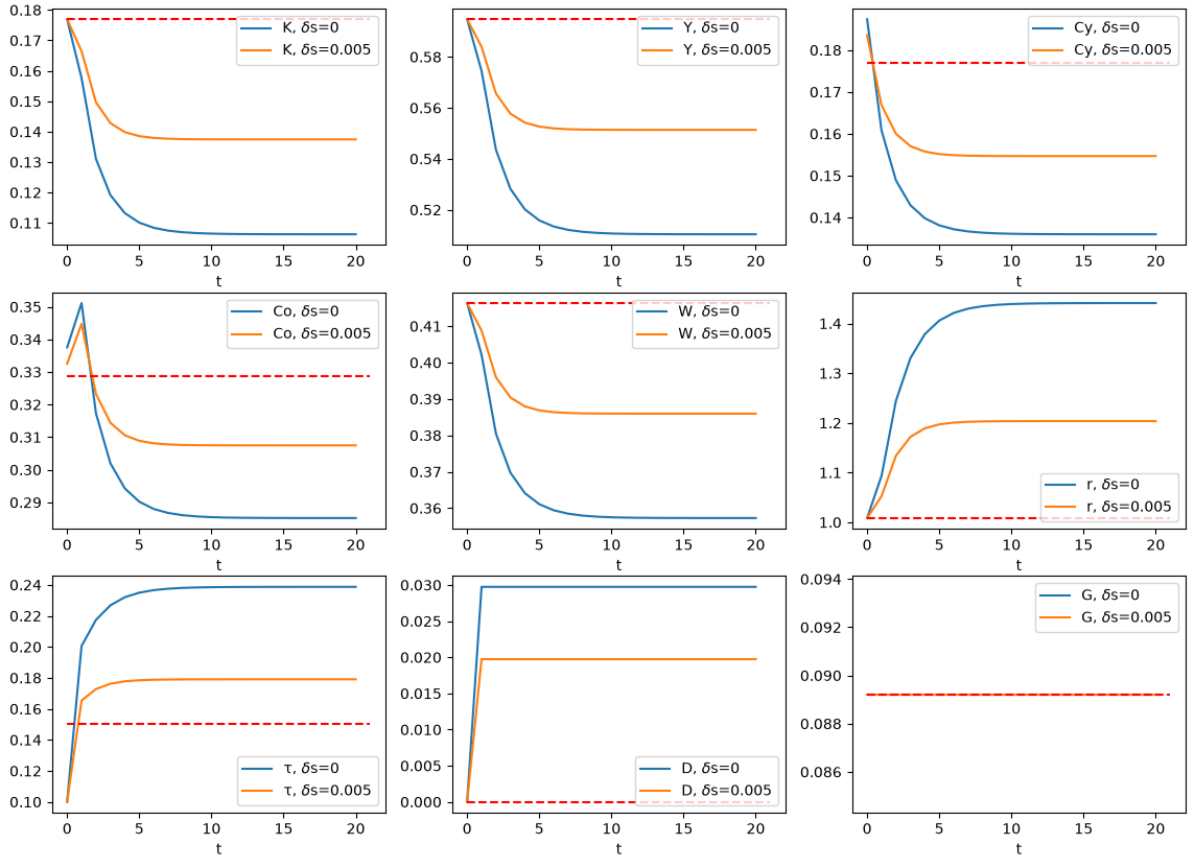
(continues on next page)

(continued from previous page)

```
    ax.legend()
    ax.set_xlabel('t')

# prices
for i, name in enumerate(['W', 'r']):
    ax = axs[(i+4)//3, (i+4)%3]
    ax.plot(range(T+1), price_seq3[:T+1, i], label=rf'{name}, $\delta$s=0')
    ax.plot(range(T+1), price_seq4[:T+1, i], label=rf'{name}, $\delta$s=0.005')
    ax.hlines(init_ss[i+4], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')

# policies
for i, name in enumerate(['τ', 'D', 'G']):
    ax = axs[(i+6)//3, (i+6)%3]
    ax.plot(range(T+1), policy_seq3[:T+1, i], label=rf'{name}, $\delta$s=0')
    ax.plot(range(T+1), policy_seq4[:T+1, i], label=rf'{name}, $\delta$s=0.005')
    ax.hlines(init_ss[i+6], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')
```



Comparing to *Experiment 1: a tax cut*, the government raises lump-sum taxes to finance the increasing debt interest payment, which is less distortionary comparing to raising the capital income tax rate.

125.11.1 Experiment 4: an unfunded social security system

In this experiment, lump-sum taxes are of equal magnitudes for old and the young, but of opposite signs.

A negative lump-sum tax is a subsidy.

Thus, in this experiment we tax the young and subsidize the old.

We start the economy at the same initial steady state that we assumed in several earlier experiments.

The government sets the lump sum taxes $\delta_{y,t} = -\delta_{o,t} = 10\% \hat{C}_y$ starting from $t = 0$.

It keeps debt levels and expenditures at their steady state levels $\hat{D}$ and $\hat{G}$.

In effect, this experiment amounts to launching an unfunded social security system.

We can use our code to compute the transition ignited by launching this system.

Let's compare the results to the *Experiment 1: a tax cut*.

```
δy_seq = np.ones(T+2) * Cy_hat * 0.1
δo_seq = np.ones(T+2) * -Cy_hat * 0.1

D_pol[:] = D_hat

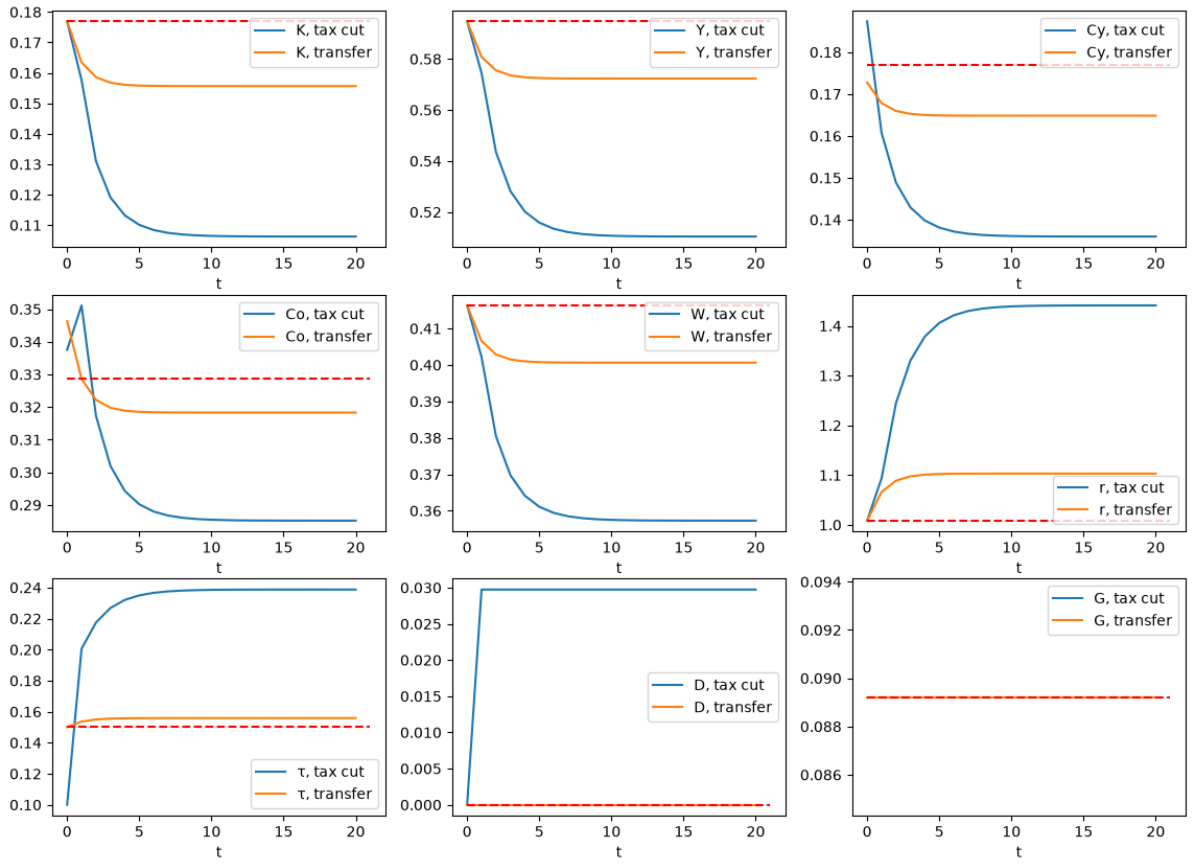
quant_seq5, price_seq5, policy_seq5 = ak2.simulate(T, init_ss,
                                                   δy_seq, δo_seq,
                                                   D_pol=D_pol, G_pol=G_pol)
```

```
fig, axs = plt.subplots(3, 3, figsize=(14, 10))

# quantities
for i, name in enumerate(['K', 'Y', 'Cy', 'Co']):
    ax = axs[i//3, i%3]
    ax.plot(range(T+1), quant_seq3[:T+1, i], label=f'{name}, tax cut')
    ax.plot(range(T+1), quant_seq5[:T+1, i], label=f'{name}, transfer')
    ax.hlines(init_ss[i], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')

# prices
for i, name in enumerate(['W', 'r']):
    ax = axs[(i+4)//3, (i+4)%3]
    ax.plot(range(T+1), price_seq3[:T+1, i], label=f'{name}, tax cut')
    ax.plot(range(T+1), price_seq5[:T+1, i], label=f'{name}, transfer')
    ax.hlines(init_ss[i+4], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')

# policies
for i, name in enumerate(['τ', 'D', 'G']):
    ax = axs[(i+6)//3, (i+6)%3]
    ax.plot(range(T+1), policy_seq3[:T+1, i], label=f'{name}, tax cut')
    ax.plot(range(T+1), policy_seq5[:T+1, i], label=f'{name}, transfer')
    ax.hlines(init_ss[i+6], 0, T+1, color='r', linestyle='--')
    ax.legend()
    ax.set_xlabel('t')
```



An initial old person benefits especially when the social security system is launched because they receive a transfer but pay nothing for it.

But in the long run, consumption of both young and old people decreases because the social security system reduces incentives to save.

That lowers the stock of physical capital and consequently lowers output.

The government must then raise tax rate in order to pay for its expenditures.

The higher rate on capital income further distorts incentives to save.

125.12 Exercises

Exercise 125.12.1

This exercise studies how government debt affects the steady state capital stock.

1. Explain why the law of motion (125.11) contains D_{t+1} rather than D_t.
2. Show that a steady state with a constant debt level $\hat{D}$ satisfies

$$\hat{K} = c\,\hat{K}^{\alpha} - \hat{D}, \qquad c \equiv (1-\hat{\tau})\,(1-\alpha)\,(1-\beta)$$

3. Holding $\hat{\tau} = 0.15$, compute $\hat{K}$ for $\hat{D} \in \{0, 0.02, 0.05, 0.07\}$.
4. Find the largest $\hat{D}$ for which a steady state exists, and count the number of steady states when $0 < \hat{D} < \hat{D}_{\max}$.

Solution

Part 1. A young person at t saves $A_{t+1} = (1-\beta)(1-\tau_t)W_t$.

Those savings are held as physical capital K_{t+1} plus the government bonds D_{t+1} that mature at $t+1$, so $K_{t+1} = A_{t+1} - D_{t+1}$.

Part 2. Set $K_{t+1} = K_t = \hat{K}$, $\tau_t = \hat{\tau}$, and $D_{t+1} = \hat{D}$ in (125.11), and use $W_t = (1-\alpha)\hat{K}^\alpha$.

Parts 3 and 4. The function $cK^\alpha - K$ is strictly concave in K, equals zero at $K = 0$, and attains its maximum at $K^\star = (\alpha c)^{1/(1-\alpha)}$.

So the steady state equation $cK^\alpha - K = \hat{D}$ has two solutions when $0 < \hat{D} < \hat{D}_{\max} \equiv c(K^\star)^\alpha - K^\star$, one solution when $\hat{D} = \hat{D}_{\max}$, and none when $\hat{D} > \hat{D}_{\max}$.

```
from scipy.optimize import brentq

c = (1 - τ_hat) * (1 - α) * (1 - β)
K_star = (α * c) ** (1 / (1 - α))
D_max = c * K_star ** α - K_star

def K_ss(D, high=True):
    "Steady state K solving c K^α - D - K = 0."
    f = lambda K: c * K ** α - D - K
    return brentq(f, K_star, 10) if high else brentq(f, 1e-12, K_star)

print(f"D_max = {D_max:.5f} at K = {K_star:.5f}\n")
print(f"{'D':>6}  {'K (high)':>9}  {'K (low)':>9}  {'r (high)':>9}")
for D in [0.0, 0.02, 0.05, 0.07]:
    K_high = K_ss(D)
    K_low = K_ss(D, high=False) if D > 0 else 0.0
    print(f"{D:>6.2f}  {K_high:>9.5f}  {K_low:>9.5f}  {K_to_r(K_high, α):>9.4f}")
```

```
D_max = 0.07393 at K = 0.03169

     D   K (high)    K (low)   r (high)
  0.00    0.17695    0.00000     1.0084
  0.02    0.14756    0.00013     1.1451
  0.05    0.09837    0.00323     1.5210
  0.07    0.05380    0.01593     2.3204
```

Government debt crowds out capital: at $\hat{D} = 0.05$ the high-capital steady state holds only a little more than half of the capital that it holds when $\hat{D} = 0$.

Beyond $\hat{D}_{\max} \approx 0.074$, the young simply cannot save enough to absorb the debt, and no steady state exists.

The debt level $\bar{D}$ created by the tax cut in *Experiment 1: a tax cut* is well inside this limit.

Exercise 125.12.2

In *Experiment 2: government asset accumulation*, the government permanently halves its purchases and accumulates assets while holding $\tau_t = \hat{\tau}$.

Compute the consumption of the representative old person at $t = 0$ and the lifetime utility

$$U_t = C_{yt}^{\beta} C_{o,t+1}^{1-\beta}$$

of the cohorts that are young at $t = 0, 1, \ldots, 7$, each measured relative to its initial steady state value.

Which generations gain, which lose, and why?

Solution

```
G_seq_ex = τ_hat * 0.5 * Y_hat * np.ones(T+1)
τ_seq_ex = τ_hat * np.ones(T+1)
quant_ex, price_ex, policy_ex = closed.simulate(T, init_ss,
                                                τ_pol=τ_seq_ex,
                                                G_pol=G_seq_ex)

U = lambda Cy, Co: Cy ** β * Co ** (1 - β)
U_hat = U(Cy_hat, Co_hat)

print(f"initial old: Co_0 = {quant_ex[0, 3]:.5f}, "
      f"steady state Co_hat = {Co_hat:.5f}\n")
print(f"{'cohort young at t':>18}  {'Cy':>7}  {'Co next':>8}  {'r next':>7}  {'U/U_
↪hat - 1':>12}")
for t in range(8):
    gain = U(quant_ex[t, 2], quant_ex[t+1, 3]) / U_hat - 1
    print(f"{t:>18}  {quant_ex[t, 2]:>7.4f}  {quant_ex[t+1, 3]:>8.4f}"
          f"  {price_ex[t+1, 1]:>7.4f}  {gain:>11.2%}")
```

```
initial old: Co_0 = 0.32861, steady state Co_hat = 0.32861

 cohort young at t       Cy   Co next   r next   U/U_hat - 1
                 0   0.1769    0.3065   0.8616       -3.42%
                 1   0.1893    0.2971   0.6699       -1.66%
                 2   0.2108    0.3016   0.5062        4.57%
                 3   0.2377    0.3164   0.3894       13.75%
                 4   0.2660    0.3360   0.3096       23.99%
                 5   0.2935    0.3569   0.2540       34.22%
                 6   0.3195    0.3776   0.2139       44.04%
                 7   0.3439    0.3977   0.1839       53.37%
```

The representative old person at $t = 0$ is unaffected to the last digit.

That person's consumption is $(1 + r_0(1 - \tau_0))(K_0 + D_0)$, and the policy changes none of K_0, r_0, τ_0, or D_0.

The cohorts that are young at $t = 0$ and $t = 1$ lose about 3.4 and 1.7 percent of lifetime utility.

They save out of wages that the policy has barely raised, and they earn the much lower interest rates that the rising capital stock produces, with r falling from 1.01 to 0.86 and then to 0.67.

From $t = 2$ onwards every cohort gains, by 4.6 percent and then by steadily more, because by then the wage increases dominate the interest rate declines.

This is the sense in which the policy redistributes from early to later generations.

Exercise 125.12.3

This exercise uses the lump sum taxes to study an unfunded social security system in the **steady state**, complementing the transition computed in *Experiment 4: an unfunded social security system*.

Set $\delta_{yt} = -\delta_{ot} = d$ for all t, with $D_t = 0$ and $G_t = \hat{G}$.

1. Use (125.10) to show that

$$A_{t+1} = (1-\beta)\left[(1-\tau_t)W_t - d\right] - \beta\frac{d}{1+r_{t+1}(1-\tau_{t+1})}$$

2. Use the class `AK2` to compute the new steady state for $d \in \{0,\ 0.05\hat{C}_y,\ 0.1\hat{C}_y,\ 0.2\hat{C}_y\}$, reporting K, r, τ, C_y, and C_o at the end of the transition.
3. The transfer scheme is balanced period by period. Why does it nevertheless lower both consumptions in the steady state?

Solution

Part 1. Substitute $\delta_{yt} = d$ and $\delta_{o,t+1} = -d$ into (125.10).

Part 2.

```
D_pol_ex = np.zeros(T+2)
G_pol_ex = np.ones(T+2) * G_hat

print(f"{'d / Cy_hat':>10}  {'K':>8}  {'r':>7}  {'τ':>7}  {'Cy':>7}  {'Co':>7}")
for frac in [0.0, 0.05, 0.1, 0.2]:
    d = frac * Cy_hat
    δy_ex, δo_ex = np.ones(T+2) * d, -np.ones(T+2) * d
    q_ex, p_ex, pol_ex = ak2.simulate(T, init_ss, δy_ex, δo_ex,
                                      D_pol=D_pol_ex, G_pol=G_pol_ex)
    print(f"{frac:>10.2f}  {q_ex[T, 0]:>8.5f}  {p_ex[T, 1]:>7.4f}"
          f"  {pol_ex[T, 0]:>7.5f}  {q_ex[T, 2]:>7.4f}  {q_ex[T, 3]:>7.4f}")
```

```
d / Cy_hat         K        r         τ       Cy       Co
      0.00   0.17695   1.0084  0.15000   0.1769   0.3286
      0.05   0.16637   1.0529  0.15280   0.1710   0.3236
      0.10   0.15567   1.1030  0.15588   0.1648   0.3183
      0.20   0.13376   1.2266  0.16313   0.1512   0.3064
```

Part 3. The system pays the old out of taxes on the young, so it promises old-age income that a young person does not have to save to obtain.

Both terms in the savings rule of part 1 fall with d: the young have less after-tax income, and the prospect of a transfer when old reduces the saving that they want to do with what remains.

Lower savings mean a lower capital stock, which lowers the wage and raises the interest rate.

Lower output then requires a higher flat rate tax τ to pay for the unchanged purchases $\hat{G}$, which depresses saving further.

CHAPTER

SIX

A LONG-LIVED, HETEROGENEOUS AGENTS, OVERLAPPING GENERATIONS MODEL

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's in Anaconda, this lecture will need the following library

```
!pip install jax
```

126.1 Overview

This lecture describes an overlapping generations model with these features:

- A competitive equilibrium with incomplete markets determines prices and quantities
- Agents live many periods as in [Auerbach and Kotlikoff, 1987]
- Agents receive idiosyncratic labor productivity shocks that cannot be fully insured as in [Aiyagari, 1994]
- Government fiscal policy instruments include tax rates, debt, and transfers as in chapter 2 of [Auerbach and Kotlikoff, 1987] and *Transitions in an Overlapping Generations Model*
- Among other equilibrium objects, a competitive equilibrium determines a sequence of cross-section densities of heterogeneous agents' consumptions, labor incomes, and savings

We use the model to study:

- How fiscal policies affect different generations
- How market incompleteness promotes precautionary savings
- How life-cycle savings and buffer-stock savings motives interact
- How fiscal policies redistribute resources across and within generations

As prerequisites for this lecture, we recommend two quantecon lectures:

1. Discrete State Dynamic Programming
2. *Transitions in an Overlapping Generations Model*

as well as the optional reading *The Aiyagari Model*

As usual, let's start by importing some Python modules

```
from collections import namedtuple
import numpy as np
import matplotlib.pyplot as plt
import jax.numpy as jnp
import jax.scipy as jsp
import jax
```

126.2 Environment

We start by introducing the economic environment we are operating in.

126.2.1 Demographics and time

We work in discrete time indexed by $t = 0, 1, 2, ...$

Each agent lives for $J = 50$ periods and faces no mortality risk.

We index age by $j = 0, 1, ..., 49$.

Each cohort has mass $1/J$ and total population is one, so a cross-section average equals $\frac{1}{J}\sum_j$ of cohort averages.

Throughout, we normalize each cohort's density $\mu_{j,t}$ to integrate to one, so that population weights $1/J$ appear explicitly in aggregates.

126.2.2 Individuals' state variables

Each agent i of age j at time t is characterized by two state variables: asset holdings $a_{i,j,t}$ and idiosyncratic labor productivity $\gamma_{i,j,t}$.

The idiosyncratic labor productivity process follows a two-state Markov chain that takes values γ_l and γ_h with transition matrix Π.

Newborn agents begin with an initial distribution $\pi = [0.5, 0.5]$ over these productivity states.

126.2.3 Labor supply

An agent with productivity $\gamma_{i,j,t}$ supplies $l(j)\gamma_{i,j,t}$ efficiency units of labor.

$l(j)$ is a deterministic age-specific labor efficiency units profile.

An agent's effective labor supply depends on a life-cycle efficiency profile and an idiosyncratic stochastic process.

126.2.4 Initial conditions

Newborns start with zero assets $a_{i,0,t} = 0$.

Initial idiosyncratic productivities are drawn from distribution π.

Agents leave no bequests and have terminal value function $V_J(a) = 0$.

126.3 Production

A representative firm operates a constant returns to scale Cobb-Douglas production:

$$Y_t = Z_t K_t^\alpha L_t^{1-\alpha}$$

where:

- K_t is aggregate capital
- L_t is aggregate efficiency units of labor
- Z_t is total factor productivity
- α is the capital share

Capital does not depreciate, so the rental rate r_t below is a net rate of return.

126.4 Government

The government follows a fiscal policy that includes debt, taxes, transfers, and government spending.

The government issues one-period debt D_t to finance its operations and collects revenues through a flat-rate tax τ_t on both labor and capital income.

The government also implements age-specific lump-sum taxes or transfers $\delta_{j,t}$ that can redistribute resources across different age groups.

Additionally, it makes government purchases G_t for public goods and services.

The government budget constraint at time t is

$$D_{t+1} - D_t = r_t D_t + G_t - T_t$$

where total tax revenues T_t satisfy

$$T_t = \tau_t w_t L_t + \tau_t r_t (D_t + K_t) + \frac{1}{J} \sum_j \delta_{j,t}$$

All quantities here are per capita, which is why the age-specific lump sum taxes enter with the population weights $1/J$.

126.5 Activities in factor markets

At each time $t \geq 0$, agents supply labor and capital.

126.5.1 Age-specific labor supplies

Agents of age $j \in \{0, 1, ..., J-1\}$ supply labor according to:

- Their deterministic age-efficiency profile $l(j)$
- Their current idiosyncratic productivity shock $\gamma_{i,j,t}$

Each agent supplies $l(j)\gamma_{i,j,t}$ effective units of labor and earns a competitive wage w_t per effective unit, subject to a flat tax rate τ_t on labor earnings.

126.5.2 Asset market participation

Summarizing activities in the asset market, all agents, regardless of age $j \in \{0, 1, ..., J-1\}$, can:

- Hold assets $a_{i,j,t}$ (subject to borrowing constraints)
- Earn a risk-free one-period return r_t on savings
- Pay capital income taxes at flat rate τ_t
- Receive or pay age-specific transfers $\delta_{j,t}$

126.5.3 Key features

Lifecycle patterns shape economic behavior across ages:

- Labor productivity varies systematically with age according to the profile $l(j)$, while asset holdings follow a lifecycle pattern of accumulation during the first part of life and decumulation late in life.
- There is no retirement in this model, since $l(j) > 0$ at every age, so agents draw down assets late in life because they approach the end of life and not because they stop working.
- Age-specific fiscal transfers $\delta_{j,t}$ redistribute resources across generations.

Within-cohort heterogeneity creates dispersion among agents of the same age:

- Agents of the same age differ in their asset holdings $a_{i,j,t}$ due to different histories of idiosyncratic productivity shocks, their current productivities $\gamma_{i,j,t}$, and consequently their labor incomes and financial wealth.

Cross-cohort interactions determine equilibrium outcomes through market aggregation:

- All cohorts participate together in factor markets, with asset supplies from all cohorts determining aggregate capital and effective labor supplies from all cohorts determining aggregate labor.
- Equilibrium prices reflect both lifecycle and redistributive forces.

126.6 Representative firm's problem

A representative firm chooses capital and effective labor to maximize profits

$$\max_{K,L} Z_t K_t^{\alpha} L_t^{1-\alpha} - r_t K_t - w_t L_t$$

First-order necessary conditions imply that

$$w_t = (1-\alpha) Z_t (K_t / L_t)^{\alpha}$$

and

$$r_t = \alpha Z_t (K_t / L_t)^{\alpha - 1}$$

126.7 Households' problems

A household's value function satisfies a Bellman equation

$$V_{j,t}(a, \gamma) = \max_{c, a'} \{u(c) + \beta \mathbb{E}[V_{j+1,t+1}(a', \gamma')]\}$$

where maximization is subject to

$$c + a' = (1 + r_t(1 - \tau_t))a + (1 - \tau_t) w_t l(j) \gamma - \delta_{j,t}$$

$$c \geq 0, \qquad a' \geq 0$$

and a terminal condition $V_{J,t}(a, \gamma) = 0$

The constraint $a' \geq 0$ rules out borrowing, so an agent can self-insure only by accumulating assets.

That is the sense in which markets are incomplete here.

126.8 Population dynamics

The joint probability density function $\mu_{j,t}(a, \gamma)$ of asset holdings and idiosyncratic labor productivity evolves according to

- For newborns $(j = 0)$:

$$\mu_{0,t+1}(a', \gamma') = \begin{cases} \pi(\gamma') & \text{if } a' = 0, \\ 0, & \text{otherwise} \end{cases}$$

Each cohort density integrates to one, so this density carries the cohort weight $1/J$ in the aggregates below.

- For other cohorts:

$$\mu_{j+1,t+1}(a', \gamma') = \int \mathbf{1}_{\sigma_{j,t}(a,\gamma)=a'} \Pi(\gamma, \gamma') \mu_{j,t}(a, \gamma) d(a, \gamma)$$

where $\sigma_{j,t}(a, \gamma)$ is the optimal saving policy function.

126.9 Equilibrium

An equilibrium consists of:

- Value functions $V_{j,t}$
- Policy functions $\sigma_{j,t}$
- Joint probability distributions $\mu_{j,t}$
- Prices r_t, w_t
- Government policies $\tau_t, D_t, \delta_{j,t}, G_t$

that satisfy the following conditions

- Given prices and government policies, value and policy functions solve households' problems
- Given prices, the representative firm maximizes profits
- Government budget constraints are satisfied
- Markets clear:
  - Asset market: $K_t = \frac{1}{J}\sum_j \int a\mu_{j,t}(a,\gamma)d(a,\gamma) - D_t$
  - Labor market: $L_t = \frac{1}{J}\sum_j \int l(j)\gamma\mu_{j,t}(a,\gamma)d(a,\gamma)$

Relative to the model presented in *Transitions in an Overlapping Generations Model*, the present model adds

- Heterogeneity within generations due to productivity shocks
- A precautionary savings motive
- More re-distributional effects
- More complicated transition dynamics

126.10 Implementation

Using tools in Discrete State Dynamic Programming, we solve our model by combining value function iteration with equilibrium price determination.

A sensible approach is to nest a discrete DP solver inside an outer loop that searches for market-clearing prices.

For a candidate sequence of prices interest rates r_t and wages w_t, we can solve individual households' dynamic programming problems using either value function iteration or policy iteration to obtain optimal policy functions.

We then deduce associated stationary joint probability distributions of asset holdings and idiosyncratic labor efficiency units for each age cohort.

This will give us an aggregate capital supply (from household savings) and a labor supply (from the age-efficiency profile and productivity shocks).

We can then compare these with capital and labor demand from firms, compute deviations between factor market supplies and demands, then update price guesses until we find market-clearing prices.

To construct transition dynamics, we can compute sequences of time-varying prices by using *backward induction* to compute value and policy functions, and *forward iteration* for the distributions of agents across states:

1. Outer loop (market clearing)
   - Guess initial prices (r_t, w_t)

- Iterate until asset and labor markets clear
- Use firms' first-order necessary conditions to update prices

2. Inner loop (individual dynamic programming)
   - For each age cohort:
     - Discretize asset and productivity state space
     - Use value function iteration or policy iteration
     - Solve for optimal savings policies
     - Compute stationary distributions
3. Aggregation
   - Sum across individual states within each cohort
   - Sum across cohorts both
     - Aggregate capital supply, and
     - Aggregate effective labor supply
   - Take into account population weights $1/J$
4. Transition dynamics
   - Backward induction:
     - Start from final steady state
     - Solve sequence of value functions
   - Forward iteration:
     - Start from initial distribution
     - Track cohort distributions over time
   - Market clearing in each period:
     - Solve for price sequences
     - Update until all markets clear in all periods

We start coding by defining helper functions that describe preferences, firms, and government budget constraints.

```
ϕ, k_bar = 0., 0.

@jax.jit
def V_bar(a):
    "Terminal value function depending on the asset holding."

    return - ϕ * (a - k_bar) ** 2
```

```
ν = 0.5

@jax.jit
def u(c):
    "Utility from consumption."

    return c ** (1 - ν) / (1 - ν)
```

(continues on next page)

(continued from previous page)

```
l1, l2, l3 = 0.5, 0.05, -0.0008

@jax.jit
def l(j):
    "Age-specific wage profile."

    return l1 + l2 * j + l3 * j ** 2
```

Let's define a `Firm` namedtuple that contains parameters governing the production technology.

```
Firm = namedtuple("Firm", ("α", "Z"))

def create_firm(α=0.3, Z=1):

    return Firm(α=α, Z=Z)
```

```
firm = create_firm()
```

The following helper functions link aggregates (K, L) and prices (w, r) that emerge from the representative firm's first-order necessary conditions.

```
@jax.jit
def KL_to_r(K, L, firm):

    α, Z = firm

    return Z * α * (K / L) ** (α - 1)

@jax.jit
def KL_to_w(K, L, firm):

    α, Z = firm

    return Z * (1 - α) * (K / L) ** α
```

We use a function `find_τ` to find flat tax rates that balance the government budget constraint given other policy variables that include debt levels, government spending, and transfers.

```
@jax.jit
def find_τ(policy, price, aggs):

    D, D_next, G, δ = policy
    r, w = price
    K, L = aggs

    # each cohort has mass 1/J, so per capita transfers are δ.sum()/J
    J = δ.shape[-1]
    num = r * D + G - D_next + D - δ.sum(axis=-1) / J
    denom = w * L + r * (D + K)

    return num / denom
```

We use a namedtuple `Household` to store parameters that characterize households' problems.

```
Household = namedtuple("Household", ("j_grid", "a_grid", "y_grid",
                                     "Π", "β", "init_μ", "VJ"))
```

(continues on next page)

(continued from previous page)

```
def create_household(
        a_min=0., a_max=40, a_size=200,
        Π=[[0.9, 0.1], [0.1, 0.9]],
        y_grid=[0.5, 1.5],
        β=0.96, J=50
    ):

    j_grid = jnp.arange(J)

    a_grid = jnp.linspace(a_min, a_max, a_size)

    y_grid, Π = map(jnp.array, (y_grid, Π))
    y_size = len(y_grid)

    # Population distribution of new borns
    init_μ = jnp.zeros((a_size * y_size))

    # Newborns are endowed with zero asset
    # and equal probability of y
    init_μ = init_μ.at[:y_size].set(1 / y_size)

    # Terminal value V_bar(a)
    VJ = jnp.empty(a_size * y_size)
    for a_i in range(a_size):
        a = a_grid[a_i]
        VJ = VJ.at[a_i*y_size:(a_i+1)*y_size].set(V_bar(a))

    return Household(j_grid=j_grid, a_grid=a_grid, y_grid=y_grid,
                     Π=Π, β=β, init_μ=init_μ, VJ=VJ)
```

```
hh = create_household()
```

We apply discrete state dynamic programming tools.

Initial steps involve preparing rewards and transition matrices R and Q for our discretized Bellman equations.

```
@jax.jit
def populate_Q(household):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household

    num_state = a_grid.size * y_grid.size
    num_action = a_grid.size

    Q = jsp.linalg.block_diag(*[Π]*a_grid.size)
    Q = Q.reshape((num_state, num_action, y_grid.size))
    Q = jnp.tile(Q, a_grid.size).T

    return Q

@jax.jit
def populate_R(j, r, w, τ, δ, household):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household

    num_state = a_grid.size * y_grid.size
```

(continues on next page)

(continued from previous page)

```
    num_action = a_grid.size

    a = jnp.reshape(a_grid, (a_grid.size, 1, 1))
    y = jnp.reshape(y_grid, (1, y_grid.size, 1))
    ap = jnp.reshape(a_grid, (1, 1, a_grid.size))
    c = (1 + r*(1-τ)) * a + (1-τ) * w * l(j) * y - δ[j] - ap

    return jnp.reshape(jnp.where(c > 0, u(c), -jnp.inf),
                       (num_state, num_action))
```

126.11 Computing a steady state

We first compute a steady state.

Given guesses of prices and taxes, we can use backwards induction to solve for value functions and optimal consumption and saving policies at all ages.

The function `backwards_opt` solves for optimal values by applying the discretized bellman operator backwards.

We use `jax.lax.scan` to facilitate sequential and recurrent computations efficiently.

```
@jax.jit
def backwards_opt(prices, taxes, household, Q):

    r, w = prices
    τ, δ = taxes

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household
    J = j_grid.size

    num_state = a_grid.size * y_grid.size
    num_action = a_grid.size

    def bellman_operator_j(V_next, j):
        "Solve household optimization problem at age j given Vj+1"

        Rj = populate_R(j, r, w, τ, δ, household)
        vals = Rj + β * Q.dot(V_next)
        σ_j = jnp.argmax(vals, axis=1)
        V_j = vals[jnp.arange(num_state), σ_j]

        return V_j, (V_j, σ_j)

    js = jnp.arange(J-1, -1, -1)
    init_V = VJ

    # Iterate from age J to 1
    _, outputs = jax.lax.scan(bellman_operator_j, init_V, js)
    V, σ = outputs
    V = V[::-1]
    σ = σ[::-1]

    return V, σ
```

```
r, w = 0.05, 1
τ, δ = 0.15, np.zeros(hh.j_grid.size)

Q = populate_Q(hh)
```

```
V, σ = backwards_opt([r, w], [τ, δ], hh, Q)
```

Let's time the computation with `block_until_ready()` to ensure that all JAX operations are complete

```
%time backwards_opt([r, w], [τ, δ], hh, Q)[0].block_until_ready();
```

```
CPU times: user 460 µs, sys: 1.09 ms, total: 1.55 ms
Wall time: 50.9 ms
```

From the optimal consumption and saving choices by each cohort, we can compute a joint probability distribution of asset levels and idiosyncratic productivity levels in a steady state.

```
@jax.jit
def popu_dist(σ, household, Q):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household

    J = household.j_grid.size
    num_state = household.a_grid.size * household.y_grid.size

    def update_popu_j(μ_j, j):
        "Update population distribution from age j to j+1"

        Qσ = Q[jnp.arange(num_state), σ[j]]
        μ_next = μ_j @ Qσ

        return μ_next, μ_next

    js = jnp.arange(J-1)

    # iterate from age 1 to J
    _, μ = jax.lax.scan(update_popu_j, init_μ, js)
    μ = jnp.concatenate([init_μ[jnp.newaxis], μ], axis=0)

    return μ
```

```
μ = popu_dist(σ, hh, Q)
```

Let's time the computation

```
%time popu_dist(σ, hh, Q)[0].block_until_ready();
```

```
CPU times: user 45 ms, sys: 5.95 ms, total: 50.9 ms
Wall time: 49.9 ms
```

Below we plot the marginal distribution of savings for each age group.

```
for j in [0, 5, 20, 45, 49]:
    plt.plot(hh.a_grid, jnp.sum(μ[j].reshape((hh.a_grid.size, hh.y_grid.size)),␣
 ↪axis=1), label=f'j={j}')
```

(continues on next page)

(continued from previous page)

```
plt.legend()
plt.xlabel('a')
plt.ylabel(r'$\sum_\gamma \mu_j(a, \gamma)$')

plt.show()
```

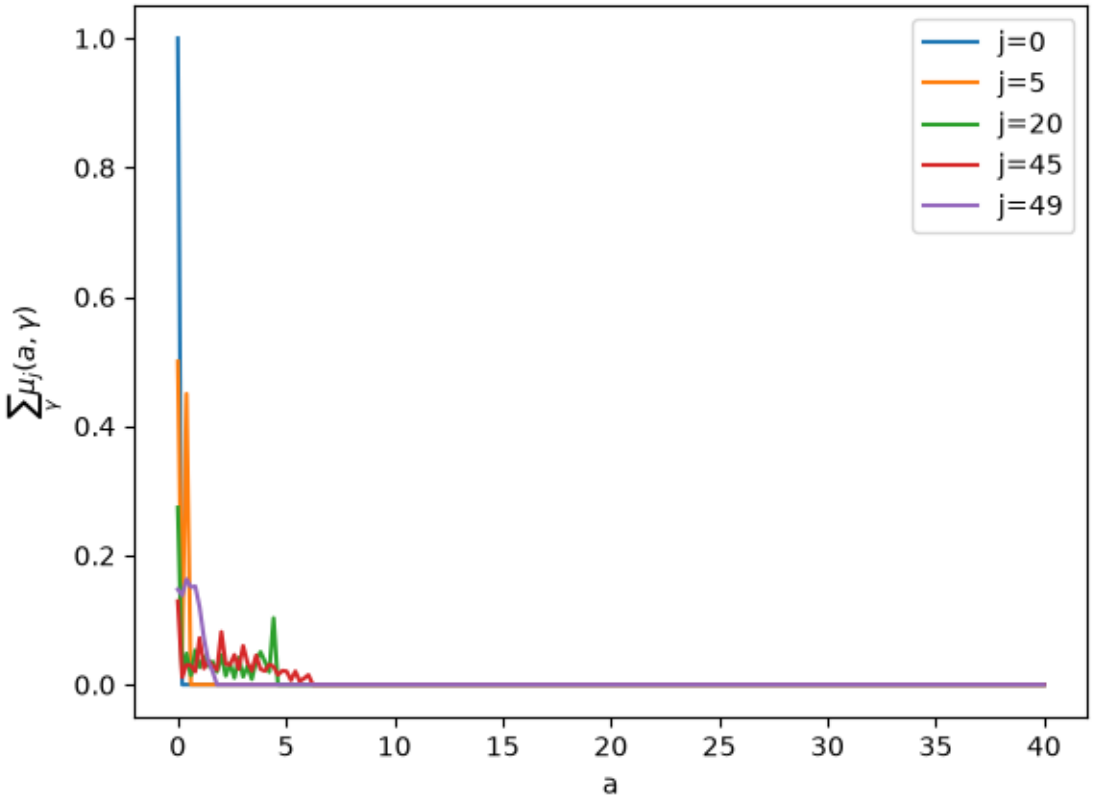


Fig. 126.1: Marginal distribution of assets by age

A grid that is too narrow would distort these distributions by piling probability mass on the highest asset level.

Let's verify that this does not happen here.

```
top_mass = μ.reshape((hh.j_grid.size, hh.a_grid.size,
                      hh.γ_grid.size))[:, -1, :].sum() / hh.j_grid.size
print(f"population share at a_max = {hh.a_grid[-1]:.0f}: {top_mass:.3%}")
```

```
population share at a_max = 40: 0.000%
```

Exercise 126.15.1 asks what happens when the upper bound binds.

These marginal distributions confirm that new agents enter the economy with no asset holdings.

- the blue $j = 0$ distribution has mass only at $a = 0$.

As agents age, at first they gradually accumulate assets.

- the orange $j = 5$ distribution puts positive mass on positive but low asset levels
- the green $j = 20$ distribution puts positive mass on a much wider range of asset levels.
- the red $j = 45$ distribution is even wider

At a later age, they gradually deplete their asset holdings.

- the purple $j = 49$ distribution illustrates this

At the end of life, they will have drawn down all of their assets.

Let's now look at age-specific optimal saving policies that generate the preceding marginal distributions of assets at different ages.

We'll plot some saving functions with the following Python code.

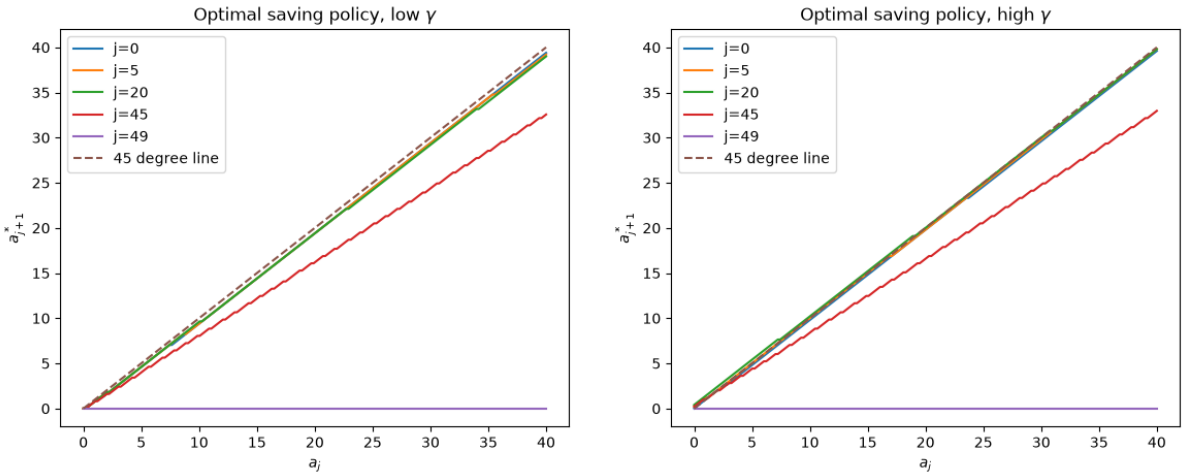
```
σ_reshaped = σ.reshape(hh.j_grid.size, hh.a_grid.size, hh.y_grid.size)
j_labels = [f'j={j}' for j in [0, 5, 20, 45, 49]]

fig, axs = plt.subplots(1, 2, figsize=(14, 5))

axs[0].plot(hh.a_grid, hh.a_grid[σ_reshaped[[0, 5, 20, 45, 49], :, 0].T])
axs[0].plot(hh.a_grid, hh.a_grid, '--')
axs[0].set_xlabel("$a_{j}$")
axs[0].set_ylabel("$a^*_{j+1}$")
axs[0].legend(j_labels+['45 degree line'])
axs[0].set_title(r"Optimal saving policy, low $\gamma$")

axs[1].plot(hh.a_grid, hh.a_grid[σ_reshaped[[0, 5, 20, 45, 49], :, 1].T])
axs[1].plot(hh.a_grid, hh.a_grid, '--')
axs[1].set_xlabel("$a_{j}$")
axs[1].set_ylabel("$a^*_{j+1}$")
axs[1].legend(j_labels+['45 degree line'])
axs[1].set_title(r"Optimal saving policy, high $\gamma$")

plt.show()
```



From an implied stationary population distribution, we can compute the aggregate labor supply L and private savings A.

```
@jax.jit
def compute_aggregates(μ, household):
```

(continues on next page)

(continued from previous page)

```
    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household

    J, a_size, y_size = j_grid.size, a_grid.size, y_grid.size

    μ = μ.reshape((J, household.a_grid.size, household.y_grid.size))

    # Compute private savings
    a = a_grid.reshape((1, a_size, 1))
    A = (a * μ).sum() / J

    y = y_grid.reshape((1, 1, y_size))
    lj = l(j_grid).reshape((J, 1, 1))
    L = (lj * y * μ).sum() / J

    return A, L
```

```
A, L = compute_aggregates(μ, hh)
A, L
```

```
(Array(1.8474663, dtype=float32), Array(1.0781993, dtype=float32))
```

The capital stock in this economy equals $A - D$.

```
D = 0
K = A - D
```

The firm's optimality conditions imply interest rate r and wage rate w.

```
KL_to_r(K, L, firm), KL_to_w(K, L, firm)
```

```
(Array(0.20578183, dtype=float32), Array(0.82273746, dtype=float32))
```

The implied prices (r, w) differ from our guesses, so we must update our guesses and iterate until we find a fixed point.

This is our outer loop.

```
@jax.jit
def find_ss(household, firm, pol_target, Q, tol=1e-6, max_iter=200):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household
    J = j_grid.size
    num_state = a_grid.size * y_grid.size

    D, G, δ = pol_target

    # Initial guesses of prices
    r, w = 0.05, 1.

    # Initial guess of τ
    τ = 0.15

    def cond_fn(state):
        "The convergence criteria."

        V, σ, μ, K, L, r, w, τ, D, G, δ, r_old, w_old, i = state
```

(continues on next page)

(continued from previous page)

```
        error = (r - r_old) ** 2 + (w - w_old) ** 2

        return (error > tol) & (i < max_iter)

    def body_fn(state):
        "The main body of iteration."

        V, σ, μ, K, L, r, w, τ, D, G, δ, r_old, w_old, i = state
        r_old, w_old = r, w

        # Household optimal decisions and values
        V, σ = backwards_opt([r, w], [τ, δ], household, Q)

        # Compute the stationary distribution
        μ = popu_dist(σ, household, Q)

        # Compute aggregates
        A, L = compute_aggregates(μ, household)
        K = A - D

        # Update prices
        r, w = KL_to_r(K, L, firm), KL_to_w(K, L, firm)

        # Find τ
        D_next = D
        τ = find_τ([D, D_next, G, δ],
                   [r, w],
                   [K, L])

        r = (r + r_old) / 2
        w = (w + w_old) / 2

        return V, σ, μ, K, L, r, w, τ, D, G, δ, r_old, w_old, i + 1

    # Initial state
    V = jnp.empty((J, num_state), dtype=float)
    σ = jnp.empty((J, num_state), dtype=int)
    μ = jnp.empty((J, num_state), dtype=float)

    K, L = 1., 1.
    initial_state = (V, σ, μ, K, L, r, w, τ, D, G, δ, r-1, w-1, 0)
    V, σ, μ, K, L, r, w, τ, D, G, δ, _, _, i = jax.lax.while_loop(
                                    cond_fn, body_fn, initial_state)

    # i equals max_iter if the loop stopped before converging
    return V, σ, μ, K, L, r, w, τ, D, G, δ, i
```

```
ss1 = find_ss(hh, firm, [0, 0.1, np.zeros(hh.j_grid.size)], Q)

print(f"iterations used: {ss1[-1]}")
```

```
iterations used: 10
```

Let's time the computation

```
%time find_ss(hh, firm, [0, 0.1, np.zeros(hh.j_grid.size)], Q)[0].block_until_ready();
```

```
CPU times: user 278 ms, sys: 0 ns, total: 278 ms
Wall time: 276 ms
```

```
hh_out_ss1 = ss1[:3]
quant_ss1 = ss1[3:5]
price_ss1 = ss1[5:7]
policy_ss1 = ss1[7:11]
```

```
# V, σ, μ
V_ss1, σ_ss1, μ_ss1 = hh_out_ss1
```

```
# K, L
K_ss1, L_ss1 = quant_ss1

K_ss1, L_ss1
```

```
(Array(9.524267, dtype=float32), Array(1.0781993, dtype=float32))
```

```
# r, w
r_ss1, w_ss1 = price_ss1

r_ss1, w_ss1
```

```
(Array(0.06527545, dtype=float32), Array(1.3466222, dtype=float32))
```

```
# τ, D, G, δ
τ_ss1, D_ss1, G_ss1, δ_ss1 = policy_ss1

τ_ss1, D_ss1, G_ss1, δ_ss1
```

```
(Array(0.04824599, dtype=float32),
 Array(0, dtype=int32, weak_type=True),
 Array(0.1, dtype=float32, weak_type=True),
 Array([0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0.,
        0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0.,
        0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0., 0.],      ␣
↪dtype=float32))
```

126.12 Transition dynamics

We compute transition dynamics using a function `path_iteration`.

In an outer loop, we iterate over guesses of prices and taxes.

In an inner loop, we compute the optimal consumption and saving choices by each cohort j in each time t, then find the implied evolution of the joint distribution of assets and productivities.

We then update our guesses of prices and taxes given the aggregate labor supply and capital stock in the economy.

We use `solve_backwards` to solve for optimal saving choices given price and tax sequences and `simulate_forward` to compute the evolution of the joint distributions.

We require two steady states as inputs: the initial steady state to provide the initial condition for `simulate_forward`, and the final steady state to provide continuation values for `solve_backwards`.

```
@jax.jit
def bellman_operator(prices, taxes, V_next, household, Q):

    r, w = prices
    τ, δ = taxes

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household
    J = j_grid.size

    num_state = a_grid.size * y_grid.size
    num_action = a_grid.size

    def bellman_operator_j(j):
        Rj = populate_R(j, r, w, τ, δ, household)
        vals = Rj + β * Q.dot(V_next[j+1])
        σ_j = jnp.argmax(vals, axis=1)
        V_j = vals[jnp.arange(num_state), σ_j]

        return V_j, σ_j

    V, σ = jax.vmap(bellman_operator_j, (0,))(jnp.arange(J-1))

    # The last life stage
    j = J-1
    Rj = populate_R(j, r, w, τ, δ, household)
    vals = Rj + β * Q.dot(VJ)
    σ = jnp.concatenate([σ, jnp.argmax(vals, axis=1)[jnp.newaxis]])
    V = jnp.concatenate([V, vals[jnp.arange(num_state), σ[j]][jnp.newaxis]])

    return V, σ
```

```
@jax.jit
def solve_backwards(V_ss2, σ_ss2, household, firm, price_seq, pol_seq, Q):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household
    J = j_grid.size
    num_state = a_grid.size * y_grid.size

    τ_seq, D_seq, G_seq, δ_seq = pol_seq
    r_seq, w_seq = price_seq

    T = r_seq.size

    def solve_backwards_t(V_next, t):

        prices = (r_seq[t], w_seq[t])
        taxes = (τ_seq[t], δ_seq[t])
        V, σ = bellman_operator(prices, taxes, V_next, household, Q)

        return V, (V,σ)

    ts = jnp.arange(T-2, -1, -1)
    init_V = V_ss2

    _, outputs = jax.lax.scan(solve_backwards_t, init_V, ts)
```

(continues on next page)

(continued from previous page)

```
    V_seq, σ_seq = outputs
    V_seq = V_seq[::-1]
    σ_seq = σ_seq[::-1]

    V_seq = jnp.concatenate([V_seq, V_ss2[jnp.newaxis]])
    σ_seq = jnp.concatenate([σ_seq, σ_ss2[jnp.newaxis]])

    return V_seq, σ_seq
```

```
@jax.jit
def population_evolution(σt, μt, household, Q):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household

    J = household.j_grid.size
    num_state = household.a_grid.size * household.y_grid.size

    def population_evolution_j(j):

        Qσ = Q[jnp.arange(num_state), σt[j]]
        μ_next = μt[j] @ Qσ

        return μ_next

    μ_next = jax.vmap(population_evolution_j, (0,))(jnp.arange(J-1))
    μ_next = jnp.concatenate([init_μ[jnp.newaxis], μ_next])

    return μ_next
```

```
@jax.jit
def simulate_forwards(σ_seq, D_seq, μ_ss1, K_ss1, L_ss1, household, Q):

    j_grid, a_grid, y_grid, Π, β, init_μ, VJ = household

    J, num_state = μ_ss1.shape

    T = σ_seq.shape[0]

    def simulate_forwards_t(μ, t):

        μ_next = population_evolution(σ_seq[t], μ, household, Q)

        A, L = compute_aggregates(μ_next, household)
        K = A - D_seq[t+1]

        return μ_next, (μ_next, K, L)

    ts = jnp.arange(T-1)
    init_μ = μ_ss1

    _, outputs = jax.lax.scan(simulate_forwards_t, init_μ, ts)
    μ_seq, K_seq, L_seq = outputs

    μ_seq = jnp.concatenate([μ_ss1[jnp.newaxis], μ_seq])
    K_seq = jnp.concatenate([K_ss1[jnp.newaxis], K_seq])
    L_seq = jnp.concatenate([L_ss1[jnp.newaxis], L_seq])
```

(continues on next page)

(continued from previous page)

```
    return μ_seq, K_seq, L_seq
```

The following algorithm describes the path iteration procedure:

Algorithm 126.12.1 (AK-Aiyagari transition path algorithm)

Inputs Given initial steady state ss_1, final steady state ss_2, time horizon T, and policy sequences (D, G, δ)

Output Compute equilibrium transition paths for value functions V, policy functions σ, distributions μ, and prices (r, w, τ)

1. Initialize from steady states:
   - $(V_1, \sigma_1, \mu_1) \leftarrow ss_1$ *(Initial steady state)*
   - $(V_2, \sigma_2, \mu_2) \leftarrow ss_2$ *(Final steady state)*
   - $(r, w, \tau) \leftarrow initialize_prices(T)$ *(Linear interpolation)*
   - $error \leftarrow \infty, i \leftarrow 0$
2. **While** $error > \varepsilon$ and $i \leq max_iter$:
   1. $i \leftarrow i + 1$
   2. $(r_{\text{old}}, w_{\text{old}}, \tau_{\text{old}}) \leftarrow (r, w, \tau)$
   3. **Backward induction:** For $t \in [T, 1]$:
      - For $j \in [0, J-1]$ *(age groups)*:
        - $V[t, j] \leftarrow \max_{a'}\{u(c) + \beta\mathbb{E}[V[t+1, j+1]]\}$
        - $\sigma[t, j] \leftarrow \arg\max_{a'}\{u(c) + \beta\mathbb{E}[V[t+1, j+1]]\}$
   4. **Forward simulation:** For $t \in [1, T]$:
      - $\mu[t] \leftarrow \Gamma(\sigma[t], \mu[t-1])$ *(Distribution evolution)*
      - $K[t] \leftarrow \int a \, d\mu[t] - D[t]$ *(Aggregate capital)*
      - $L[t] \leftarrow \int l(j)\gamma \, d\mu[t]$ *(Aggregate labor)*
      - $r[t] \leftarrow \alpha Z(K[t]/L[t])^{\alpha-1}$ *(Interest rate)*
      - $w[t] \leftarrow (1-\alpha)Z(K[t]/L[t])^{\alpha}$ *(Wage rate)*
      - $\tau[t] \leftarrow solve_budget(r[t], w[t], K[t], L[t], D[t], G[t])$
   5. Compute convergence metric:
      - $error \leftarrow \|r - r_{\text{old}}\| + \|w - w_{\text{old}}\| + \|\tau - \tau_{\text{old}}\|$
   6. Update prices with dampening:
      - $r \leftarrow \lambda r + (1-\lambda) r_{\text{old}}$
      - $w \leftarrow \lambda w + (1-\lambda) w_{\text{old}}$
      - $\tau \leftarrow \lambda \tau + (1-\lambda) \tau_{\text{old}}$
3. **Return** $(V, \sigma, \mu, r, w, \tau)$

```
def path_iteration(ss1, ss2, pol_target, household, firm, Q, tol=1e-4,
                   max_iter=100, verbose=False):

    # Starting point: initial steady state
    V_ss1, σ_ss1, μ_ss1 = ss1[:3]
    K_ss1, L_ss1 = ss1[3:5]
    r_ss1, w_ss1 = ss1[5:7]
    τ_ss1, D_ss1, G_ss1, δ_ss1 = ss1[7:11]

    # Ending point: converging new steady state
    V_ss2, σ_ss2, μ_ss2 = ss2[:3]
    K_ss2, L_ss2 = ss2[3:5]
    r_ss2, w_ss2 = ss2[5:7]
    τ_ss2, D_ss2, G_ss2, δ_ss2 = ss2[7:11]

    # The given policies: D, G, δ
    D_seq, G_seq, δ_seq = pol_target
    T = G_seq.shape[0]

    # Initial guesses of prices
    r_seq = jnp.linspace(0, 1, T) * (r_ss2 - r_ss1) + r_ss1
    w_seq = jnp.linspace(0, 1, T) * (w_ss2 - w_ss1) + w_ss1

    # Initial guess of policy
    τ_seq = jnp.linspace(0, 1, T) * (τ_ss2 - τ_ss1) + τ_ss1

    error = 1
    num_iter = 0

    if verbose:
        fig, axs = plt.subplots(1, 3, figsize=(14, 3))
        axs[0].plot(jnp.arange(T), r_seq)
        axs[1].plot(jnp.arange(T), w_seq)
        axs[2].plot(jnp.arange(T), τ_seq, label=f'iter {num_iter}')

    while (error > tol) and (num_iter < max_iter):
        # Repeat until finding the fixed point, or until max_iter

        r_old, w_old, τ_old = r_seq, w_seq, τ_seq

        pol_seq = (τ_seq, D_seq, G_seq, δ_seq)
        price_seq = (r_seq, w_seq)

        # Solve optimal policies backwards
        V_seq, σ_seq = solve_backwards(
            V_ss2, σ_ss2, household, firm, price_seq, pol_seq, Q)

        # Compute population evolution forwards
        μ_seq, K_seq, L_seq = simulate_forwards(
            σ_seq, D_seq, μ_ss1, K_ss1, L_ss1, household, Q)

        # Update prices by aggregate capital and labor supply
        r_seq = KL_to_r(K_seq, L_seq, firm)
        w_seq = KL_to_w(K_seq, L_seq, firm)

        # Find taxes that balance the government budget constraint
        τ_seq = find_τ([D_seq[:-1], D_seq[1:], G_seq, δ_seq],
```

(continues on next page)

(continued from previous page)

```
                    [r_seq, w_seq],
                    [K_seq, L_seq])

        # Distance between new and old guesses
        error = jnp.sum((r_old - r_seq) ** 2) + \
                jnp.sum((w_old - w_seq) ** 2) + \
                jnp.sum((τ_old - τ_seq) ** 2)

        num_iter += 1
        if verbose:
            print(f"Iteration {num_iter:3d}: error = {error:.6e}")
            axs[0].plot(jnp.arange(T), r_seq)
            axs[1].plot(jnp.arange(T), w_seq)
            axs[2].plot(jnp.arange(T), τ_seq, label=f'iter {num_iter}')

        r_seq = (r_seq + r_old) / 2
        w_seq = (w_seq + w_old) / 2
        τ_seq = (τ_seq + τ_old) / 2

    if error > tol:
        print(f"Warning: stopped at {num_iter} iterations with error {error:.2e}")

    if verbose:
        axs[0].set_xlabel('t')
        axs[1].set_xlabel('t')
        axs[2].set_xlabel('t')

        axs[0].set_title('r')
        axs[1].set_title('w')
        axs[2].set_title('τ')

        axs[2].legend(loc='center left', bbox_to_anchor=(1, 0.5))

    return V_seq, σ_seq, μ_seq, K_seq, L_seq, r_seq, w_seq, \
           τ_seq, D_seq, G_seq, δ_seq
```

We can now compute equilibrium transitions that are ignited by fiscal policy reforms.

126.13 Experiment 1: an immediate tax cut

At $t = 0$, the government unexpectedly announces that it will issue debt.

From $t = 0$ to 19, debt D_{t+1} increases linearly for 20 periods, reaching the new target level $D_{20} = D_0 + 1 = \bar{D} + 1$.

Government spending $\bar{G}$ and transfers $\bar{\delta}_j$ remain constant.

The debt path is the policy; the flat rate tax τ_t is then whatever balances the government budget at each date.

At the announcement the residual tax rate drops well below its initial value, which is why we call this an immediate tax cut.

It then climbs steadily as debt accumulates, and once debt stops growing it settles permanently above its initial value, because the government must service the larger debt forever.

We want to compute the equilibrium transition path.

Our first step is to prepare appropriate policy variable arrays `D_seq`, `G_seq`, `δ_seq`

We'll compute a `τ_seq` that balances government budgets.

```
T = 150

D_seq = jnp.ones(T+1) * D_ss1
D_seq = D_seq.at[:21].set(D_ss1 + jnp.linspace(0, 1, 21))
D_seq = D_seq.at[21:].set(D_seq[20])

G_seq = jnp.ones(T) * G_ss1

δ_seq = jnp.repeat(δ_ss1, T).reshape((T, δ_ss1.size))
```

In order to iterate the path, we need to first find its destination, which is the new steady state under the new fiscal policy.

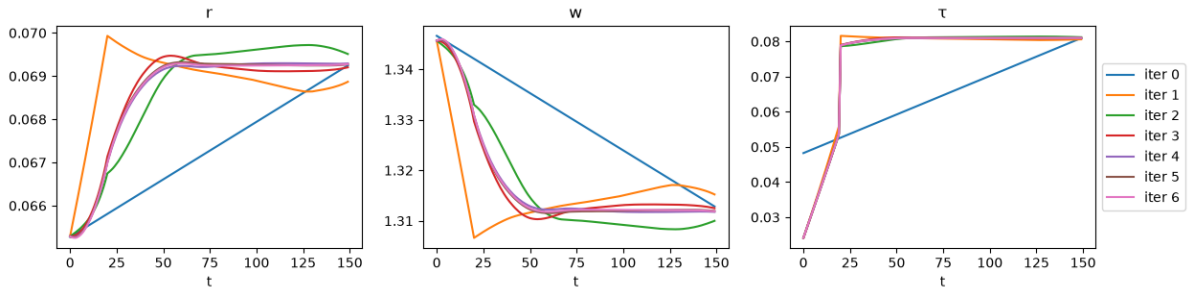
```
ss2 = find_ss(hh, firm, [D_seq[-1], G_seq[-1], δ_seq[-1]], Q)
```

We can use `path_iteration` to find equilibrium transition dynamics.

Setting the key argument `verbose=True` tells the function `path_iteration` to display convergence information.

```
paths = path_iteration(ss1, ss2, [D_seq, G_seq, δ_seq], hh, firm, Q, verbose=True)
```

```
Iteration    1: error = 8.831175e-02
Iteration    2: error = 2.078297e-02
Iteration    3: error = 5.341858e-03
Iteration    4: error = 9.897855e-04
Iteration    5: error = 3.227252e-04
Iteration    6: error = 5.921904e-05
```



Having successfully computed transition dynamics, let's study them.

```
V_seq, σ_seq, μ_seq = paths[:3]
K_seq, L_seq = paths[3:5]
r_seq, w_seq = paths[5:7]
τ_seq, D_seq, G_seq, δ_seq = paths[7:11]
```

```
ap = hh.a_grid[σ_seq[0]]
```

```
j = jnp.reshape(hh.j_grid, (hh.j_grid.size, 1, 1))
lj = l(j)
a = jnp.reshape(hh.a_grid, (1, hh.a_grid.size, 1))
y = jnp.reshape(hh.y_grid, (1, 1, hh.y_grid.size))
```

```
t = 0

ap = hh.a_grid[σ_seq[t]]
```

(continues on next page)

(continued from previous page)

```
δ = δ_seq[t].reshape((hh.j_grid.size, 1, 1))

inc = (1 + r_seq[t]*(1-τ_seq[t])) * a \
        + (1-τ_seq[t]) * w_seq[t] * lj * y - δ
inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

c = inc - ap

c_mean0 = (c * μ_seq[t]).sum(axis=1)
```

We care about how the policy change affects consumption across cohorts and across time.

We can study age-specific average consumption levels.

```
for t in [1, 10, 20, 50, 149]:

    ap = hh.a_grid[σ_seq[t]]
    δ = δ_seq[t].reshape((hh.j_grid.size, 1, 1))

    inc = (1 + r_seq[t]*(1-τ_seq[t])) * a + (1-τ_seq[t]) * w_seq[t] * lj * y - δ
    inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

    c = inc - ap

    c_mean = (c * μ_seq[t]).sum(axis=1)

    plt.plot(range(hh.j_grid.size), c_mean-c_mean0, label=f't={t}')

plt.legend()
plt.xlabel(r'j')
plt.ylabel(r'$\Delta$ mean $C(j)$')
plt.show()
```

To summarize the transition, we can plot paths as we did in *Transitions in an Overlapping Generations Model*.

But unlike the setup in that two-period lived overlapping generations model, we no longer have representative old and young agents.

- now we have 50 cohorts of different ages at each time

To proceed, we construct two age groups of equal size – young and old.

- at age 25, someone moves from being young to becoming old

```
ap = hh.a_grid[σ_ss1]
J = hh.j_grid.size
δ = δ_ss1.reshape((hh.j_grid.size, 1, 1))

inc = (1 + r_ss1*(1-τ_ss1)) * a + (1-τ_ss1) * w_ss1 * lj * y - δ
inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

c = inc - ap

Cy_ss1 = (c[:J//2] * μ_ss1[:J//2]).sum() / (J // 2)
Co_ss1 = (c[J//2:] * μ_ss1[J//2:]).sum() / (J // 2)
```

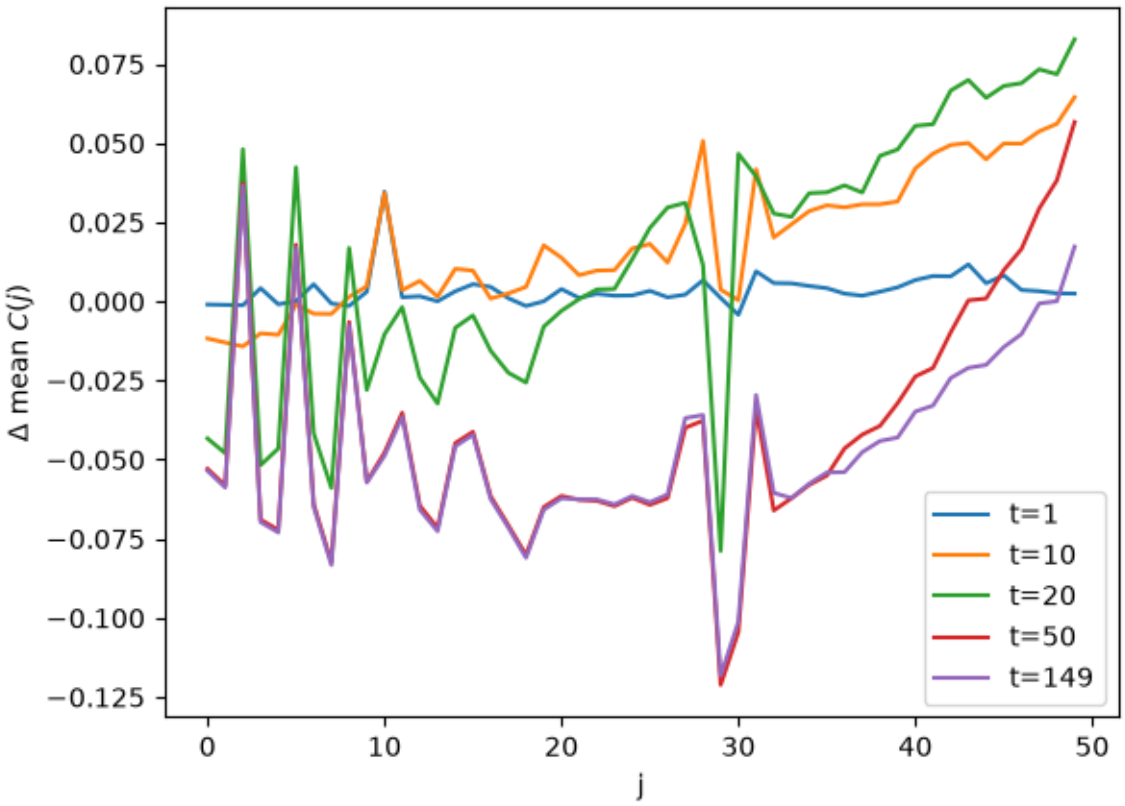


Fig. 126.2: Change in mean consumption by age

```
T = σ_seq.shape[0]
J = σ_seq.shape[1]

Cy_seq = np.empty(T)
Co_seq = np.empty(T)

for t in range(T):
    ap = hh.a_grid[σ_seq[t]]
    δ = δ_seq[t].reshape((hh.j_grid.size, 1, 1))

    inc = (1 + r_seq[t]*(1-τ_seq[t])) * a + (1-τ_seq[t]) * w_seq[t] * lj * y - δ
    inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

    c = inc - ap

    Cy_seq[t] = (c[:J//2] * μ_seq[t, :J//2]).sum() / (J // 2)
    Co_seq[t] = (c[J//2:] * μ_seq[t, J//2:]).sum() / (J // 2)
```

```
fig, axs = plt.subplots(3, 3, figsize=(14, 10))

# Cy (j=0-24)
axs[0, 0].plot(Cy_seq)
axs[0, 0].hlines(Cy_ss1, 0, T, color='r', linestyle='--')
axs[0, 0].set_title('Cy (j < 25)')

# Cy (j=25-49)
axs[0, 1].plot(Co_seq)
axs[0, 1].hlines(Co_ss1, 0, T, color='r', linestyle='--')
axs[0, 1].set_title(r'Co (j $\geq$ 25)')

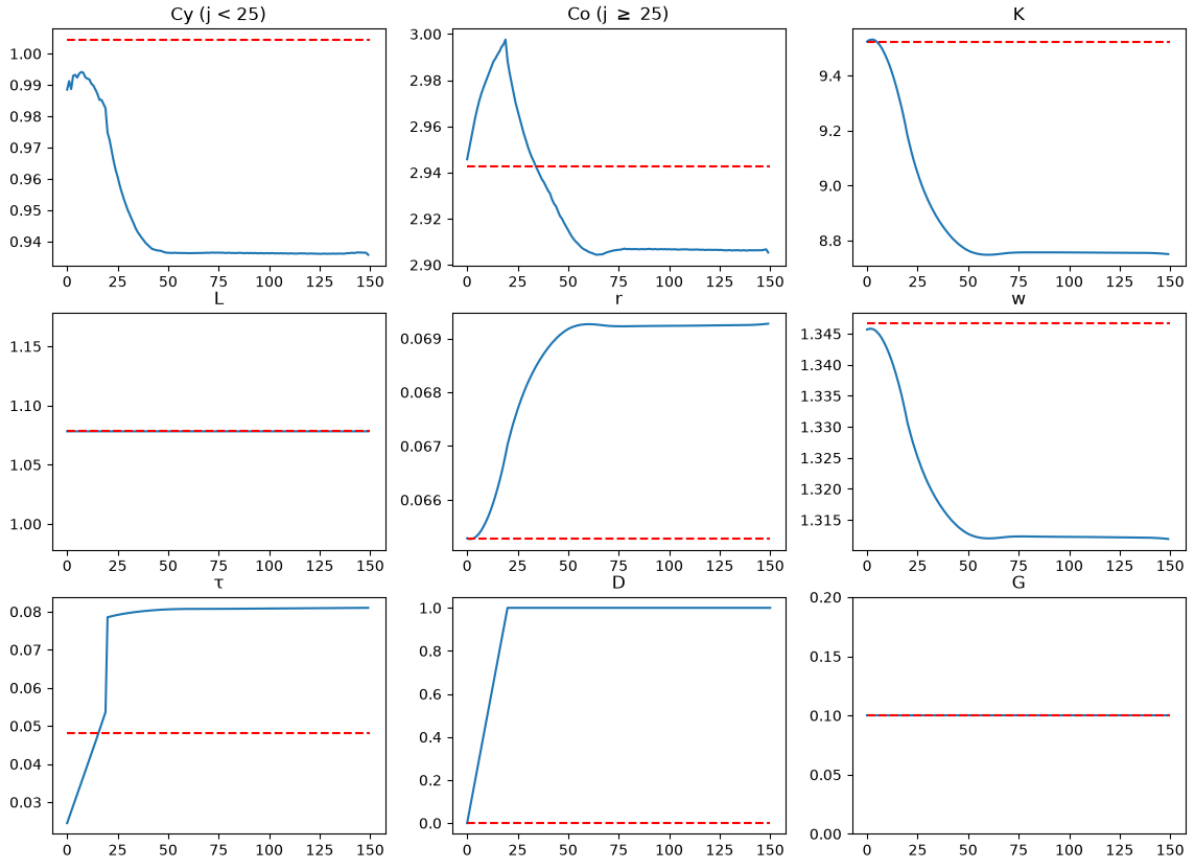
names = ['K', 'L', 'r', 'w', 'τ', 'D', 'G']
for i in range(len(names)):
    i_var = i + 3
    i_axes = i + 2

    row_i = i_axes // 3
    col_i = i_axes % 3

    axs[row_i, col_i].plot(paths[i_var])
    axs[row_i, col_i].hlines(ss1[i_var], 0, T, color='r', linestyle='--')
    axs[row_i, col_i].set_title(names[i])

# ylims
axs[1, 0].set_ylim([ss1[4]-0.1, ss1[4]+0.1])
axs[2, 2].set_ylim([ss1[9]-0.1, ss1[9]+0.1])

plt.show()
```



Now let's compute the mean and variance of consumption conditional on age at each time t.

```
Cmean_seq = np.empty((T, J))
Cvar_seq = np.empty((T, J))

for t in range(T):
    ap = hh.a_grid[σ_seq[t]]
    δ = δ_seq[t].reshape((hh.j_grid.size, 1, 1))

    inc = (1 + r_seq[t]*(1-τ_seq[t])) * a + (1-τ_seq[t]) * w_seq[t] * lj * y - δ
    inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

    c = inc - ap

    Cmean_seq[t] = (c * μ_seq[t]).sum(axis=1)
    Cvar_seq[t] = ((c - Cmean_seq[t].reshape((J, 1))) ** 2 * μ_seq[t]).sum(axis=1)
```

```
J_seq, T_range = np.meshgrid(np.arange(J), np.arange(T))

fig = plt.figure(figsize=[20, 20])

# Plot the consumption mean over age and time
ax1 = fig.add_subplot(121, projection='3d')
ax1.plot_surface(T_range, J_seq, Cmean_seq, rstride=1, cstride=1,
                 cmap='viridis', edgecolor='none')
ax1.set_title(r"Mean of consumption")
ax1.set_xlabel(r"t")
```

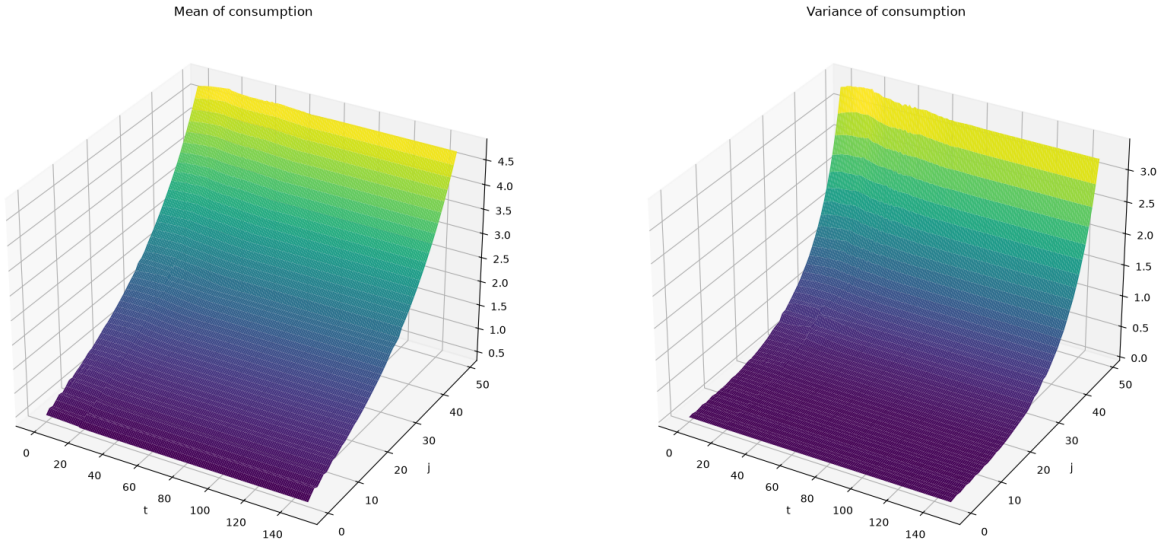
(continues on next page)

(continued from previous page)

```
ax1.set_ylabel(r"j")

# plot the consumption variance over age and time
ax2 = fig.add_subplot(122, projection='3d')
ax2.plot_surface(T_range, J_seq, Cvar_seq, rstride=1, cstride=1,
                 cmap='viridis', edgecolor='none')
ax2.set_title(r"Variance of consumption")
ax2.set_xlabel(r"t")
ax2.set_ylabel(r"j")

plt.show()
```



126.14 Experiment 2: a preannounced tax cut

Now the government announces at time 0 that it will issue the same amount of debt, but will start doing so only after 20 periods.

As in Experiment 1, the tax rate is residual: it falls while the debt is being issued, from $t = 20$ to $t = 40$, and settles permanently above its initial level thereafter.

We will use the same key toolkit `path_iteration`.

We must specify `D_seq` appropriately.

```
T = 150

D_t = 20
D_seq = jnp.ones(T+1) * D_ss1
D_seq = D_seq.at[D_t:D_t+21].set(D_ss1 + jnp.linspace(0, 1, 21))
D_seq = D_seq.at[D_t+21:].set(D_seq[D_t+20])

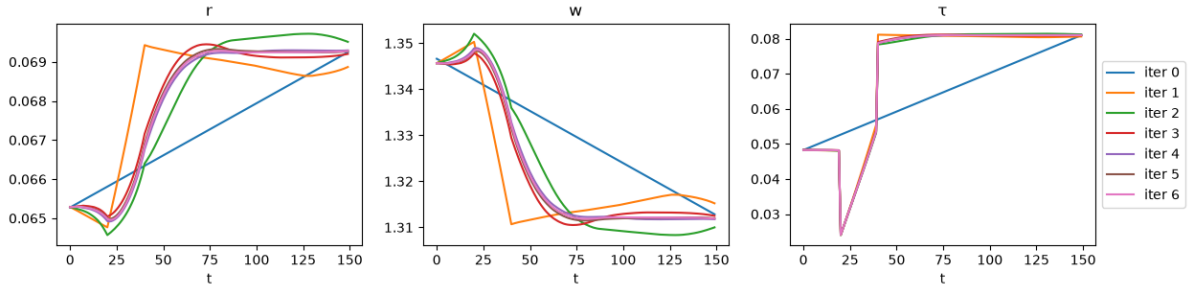
G_seq = jnp.ones(T) * G_ss1

δ_seq = jnp.repeat(δ_ss1, T).reshape((T, δ_ss1.size))
```

```
ss2 = find_ss(hh, firm, [D_seq[-1], G_seq[-1], δ_seq[-1]], Q)
```

```
paths = path_iteration(ss1, ss2, [D_seq, G_seq, δ_seq],
                       hh, firm, Q, verbose=True)
```

```
Iteration   1: error = 5.251274e-02
Iteration   2: error = 1.731133e-02
Iteration   3: error = 3.894093e-03
Iteration   4: error = 7.399297e-04
Iteration   5: error = 2.448352e-04
Iteration   6: error = 4.188173e-05
```




```
V_seq, σ_seq, μ_seq = paths[:3]
K_seq, L_seq = paths[3:5]
r_seq, w_seq = paths[5:7]
τ_seq, D_seq, G_seq, δ_seq = paths[7:11]
```

```
T = σ_seq.shape[0]
J = σ_seq.shape[1]

Cy_seq = np.empty(T)
Co_seq = np.empty(T)

for t in range(T):
    ap = hh.a_grid[σ_seq[t]]
    δ = δ_seq[t].reshape((hh.j_grid.size, 1, 1))

    inc = (1 + r_seq[t]*(1-τ_seq[t])) * a + (1-τ_seq[t]) * w_seq[t] * lj * y - δ
    inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

    c = inc - ap

    Cy_seq[t] = (c[:J//2] * μ_seq[t, :J//2]).sum() / (J // 2)
    Co_seq[t] = (c[J//2:] * μ_seq[t, J//2:]).sum() / (J // 2)
```

Below we plot the transition paths of the economy.

```
fig, axs = plt.subplots(3, 3, figsize=(14, 10))

# Cy (j=0-24)
axs[0, 0].plot(Cy_seq)
axs[0, 0].hlines(Cy_ss1, 0, T, color='r', linestyle='--')
axs[0, 0].set_title('Cy (j < 25)')

# Cy (j=25-49)
```

(continues on next page)

(continued from previous page)

```
axs[0, 1].plot(Co_seq)
axs[0, 1].hlines(Co_ss1, 0, T, color='r', linestyle='--')
axs[0, 1].set_title(r'Co (j $\geq$ 25)')

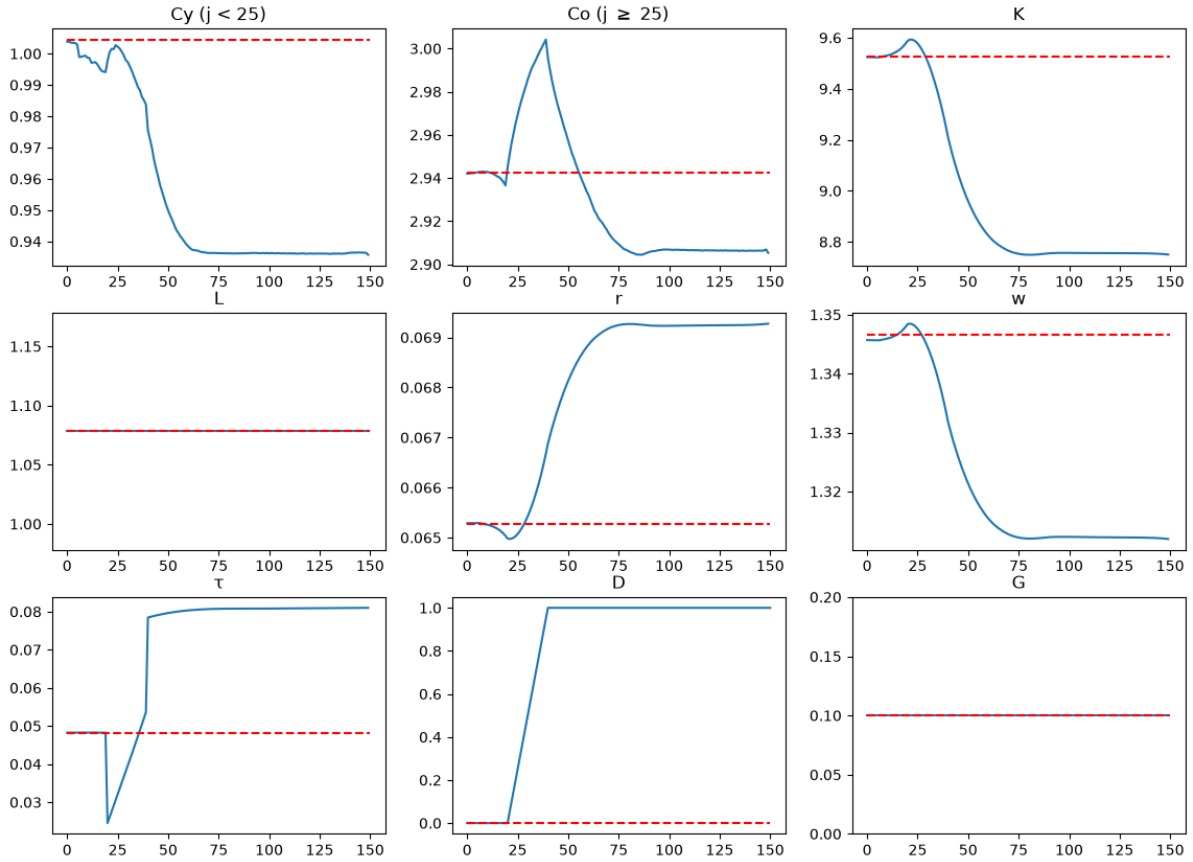
names = ['K', 'L', 'r', 'w', 'τ', 'D', 'G']
for i in range(len(names)):
    i_var = i + 3
    i_axes = i + 2

    row_i = i_axes // 3
    col_i = i_axes % 3

    axs[row_i, col_i].plot(paths[i_var])
    axs[row_i, col_i].hlines(ss1[i_var], 0, T, color='r', linestyle='--')
    axs[row_i, col_i].set_title(names[i])

# ylims
axs[1, 0].set_ylim([ss1[4]-0.1, ss1[4]+0.1])
axs[2, 2].set_ylim([ss1[9]-0.1, ss1[9]+0.1])

plt.show()
```



Notice how prices and quantities respond immediately, well before the policy is implemented at $t = 20$.

Agents who foresee the coming tax cut, and the permanently higher tax rate that follows it, adjust their savings right away.

Let's zoom in on how the capital stock responds.

```
# K
i_var = 3
K_path = paths[i_var][:25]

K_lo = min(float(K_path.min()), float(ss1[i_var]))
K_hi = max(float(K_path.max()), float(ss1[i_var]))
pad = 0.1 * (K_hi - K_lo)

plt.plot(K_path)
plt.hlines(ss1[i_var], 0, 25, color='r', linestyle='--')
plt.axvline(20, color='k', linestyle='--', linewidth=0.5)
plt.text(17, K_lo - 0.5 * pad, r'tax cut')
plt.ylim([K_lo - pad, K_hi + pad])
plt.ylabel("K")
plt.xlabel("t")
plt.show()
```

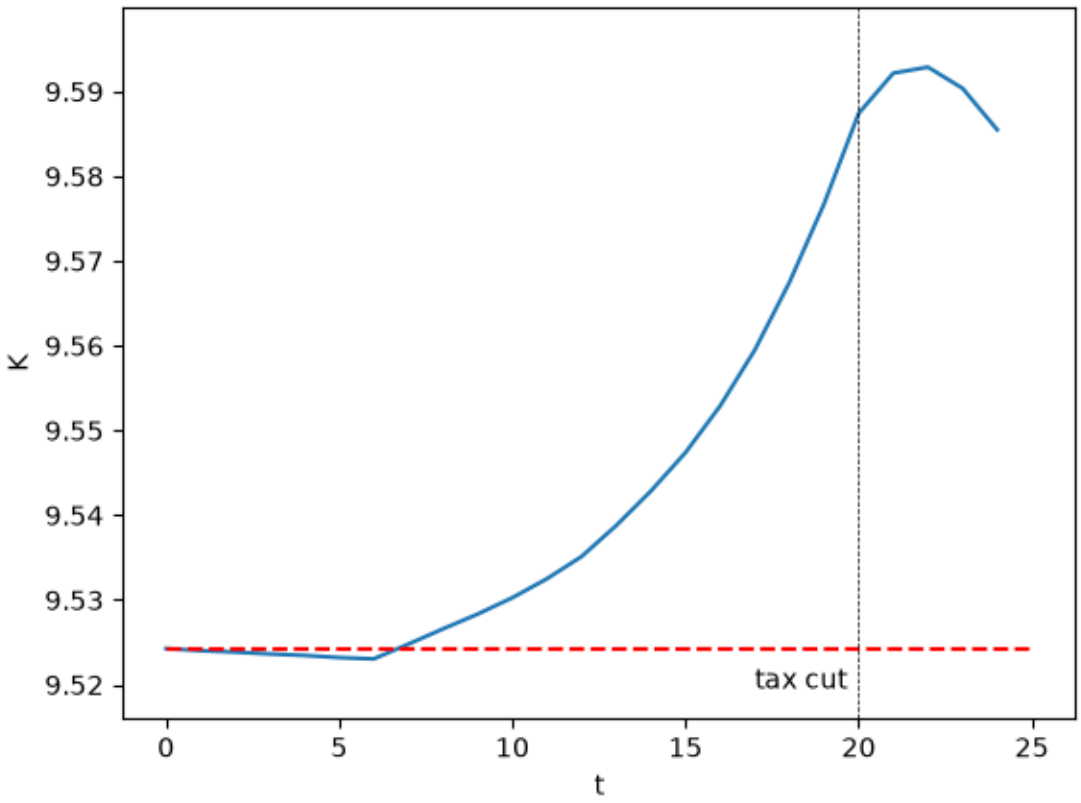


Fig. 126.3: Capital stock around the tax cut

After the tax cut is implemented at $t = 20$, aggregate capital decreases because government debt crowds out private capital.

In the few periods just before $t = 20$, individuals save more.

The reason is that the tax cut raises the *after-tax* return $r_t(1 - \tau_t)$ at $t = 20$, even though the pre-tax return r_t is falling then, so saving into $t = 20$ becomes more attractive.

Because that extra saving raises the capital stock, the pre-tax interest rate dips, reaching its lowest value just before the policy takes effect.

In the first few periods after the announcement the two forces nearly cancel and the capital stock is almost flat, with the saving response building up as the implementation date approaches.

We can also plot evolutions of means and variances of consumption by different cohorts along a transition path.

```
Cmean_seq = np.empty((T, J))
Cvar_seq = np.empty((T, J))

for t in range(T):
    ap = hh.a_grid[σ_seq[t]]
    δ = δ_seq[t].reshape((hh.j_grid.size, 1, 1))

    inc = (1 + r_seq[t]*(1-τ_seq[t])) * a + (1-τ_seq[t]) * w_seq[t] * lj * y - δ
    inc = inc.reshape((hh.j_grid.size, hh.a_grid.size * hh.y_grid.size))

    c = inc - ap

    Cmean_seq[t] = (c * μ_seq[t]).sum(axis=1)
    Cvar_seq[t] = (
        (c - Cmean_seq[t].reshape((J, 1))) ** 2 * μ_seq[t]).sum(axis=1)
```

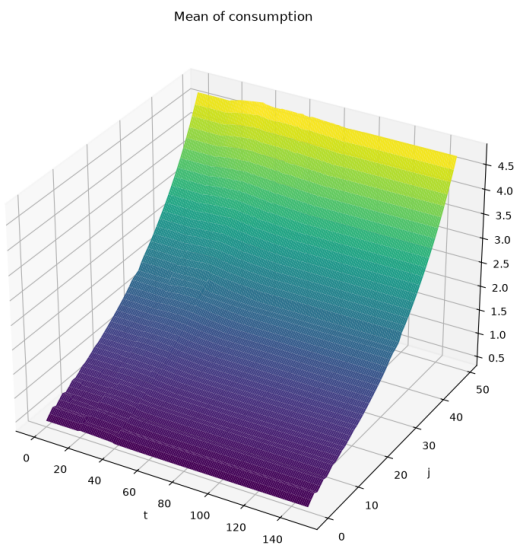
```
J_seq, T_range = np.meshgrid(np.arange(J), np.arange(T))

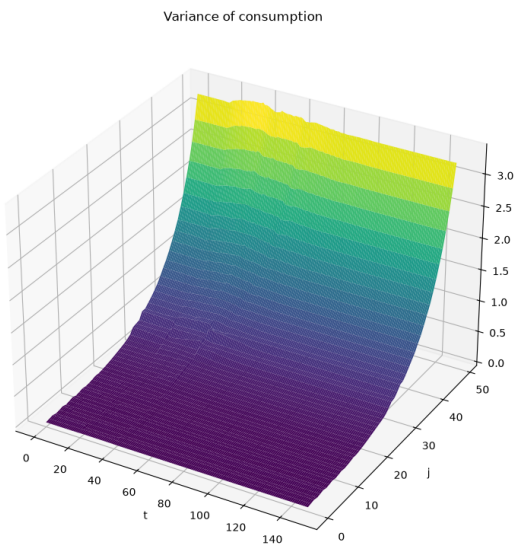
fig = plt.figure(figsize=[20, 20])

# Plot the consumption mean over age and time
ax1 = fig.add_subplot(121, projection='3d')
ax1.plot_surface(T_range, J_seq, Cmean_seq, rstride=1, cstride=1,
                 cmap='viridis', edgecolor='none')
ax1.set_title(r"Mean of consumption")
ax1.set_xlabel(r"t")
ax1.set_ylabel(r"j")

# Plot the consumption variance over age and time
ax2 = fig.add_subplot(122, projection='3d')
ax2.plot_surface(T_range, J_seq, Cvar_seq, rstride=1, cstride=1,
                 cmap='viridis', edgecolor='none')
ax2.set_title(r"Variance of consumption")
ax2.set_xlabel(r"t")
ax2.set_ylabel(r"j")

plt.show()
```





126.15 Exercises

Exercise 126.15.1

Our asset grid runs from 0 to $a_{\max} = 40$.

A grid whose upper bound binds would make the computed equilibrium an artifact of the grid rather than a property of the model.

1. Re-solve the steady state with $a_{\max} \in \{10, 20\}$, holding the number of grid points fixed at 200.
2. For each case report aggregate capital K, the interest rate r, the flat rate tax τ, and the share of the population at the top grid point.
3. What do you conclude about using $a_{\max} = 10$?

Solution

```
def ss_for_a_max(a_max):
    "Solve the steady state for a given upper bound on assets."
    h = create_household(a_max=a_max)
    Q_h = populate_Q(h)
    out = find_ss(h, firm, [0, 0.1, np.zeros(h.j_grid.size)], Q_h)
    μ_h = out[2].reshape((h.j_grid.size, h.a_grid.size, h.y_grid.size))
    top = float(μ_h[:, -1, :].sum() / h.j_grid.size)
    return float(out[3]), float(out[5]), float(out[7]), top

print(f"{'a_max':>6}  {'K':>7}  {'r':>7}  {'τ':>7}  {'mass at a_max':>14}")
for a_max in [10, 20, 40]:
    K_a, r_a, τ_a, top = ss_for_a_max(a_max)
    print(f"{a_max:>6}  {K_a:>7.3f}  {r_a:>7.4f}  {τ_a:>7.4f}  {top:>13.1%}")
```

```
a_max         K        r        τ   mass at a_max
   10     6.622   0.0843   0.0538          30.6%
   20     8.872   0.0686   0.0493           5.8%
   40     9.524   0.0653   0.0482           0.0%
```

With $a_{\max} = 10$, roughly thirty percent of the population is pinned at the upper bound.

Those agents would like to hold more assets than the grid permits, so measured capital is far too low and the interest rate far too high.

Raising the bound to 40 empties the top of the grid and moves K up by about forty percent and r down by nearly two percentage points.

The lesson is that an upper bound must be checked, not assumed: a grid is wide enough only when almost no probability mass reaches its edge.

Exercise 126.15.2

How much of the capital stock is accounted for by precautionary saving against idiosyncratic labor productivity risk?

Hold the mean of γ at one and shrink the spread, comparing $\gamma \in \{0.9, 1.1\}$ and $\gamma \in \{0.75, 1.25\}$ with the baseline $\gamma \in \{0.5, 1.5\}$.

Report K, L, r, and τ in each case, and explain the direction of the effect.

Solution

```
print(f"{'y_grid':>16}  {'K':>7}  {'L':>7}  {'r':>7}  {'τ':>7}")
for y_grid in [[0.9, 1.1], [0.75, 1.25], [0.5, 1.5]]:
    h = create_household(y_grid=y_grid)
    Q_h = populate_Q(h)
    out = find_ss(h, firm, [0, 0.1, np.zeros(h.j_grid.size)], Q_h)
    print(f"{str(y_grid):>16}  {float(out[3]):>7.3f}  {float(out[4]):>7.4f}"
          f"  {float(out[5]):>7.4f}  {float(out[7]):>7.4f}")
```

```
          y_grid        K        L        r        τ
      [0.9, 1.1]    9.088   1.0782   0.0674   0.0489
    [0.75, 1.25]    9.246   1.0782   0.0671   0.0487
      [0.5, 1.5]    9.524   1.0782   0.0653   0.0482
```

Aggregate labor L is the same in all three economies, because the productivity chain is symmetric and its mean is one in each case.

Capital nevertheless rises with the spread, from about 9.1 to about 9.5, and the interest rate falls by roughly twenty basis points.

The extra capital is precautionary saving: an agent who cannot borrow and cannot insure labor income holds a buffer stock of assets against a run of low productivity draws, and a wider spread calls for a bigger buffer.

The effect is non-linear, since most of it appears when we move from $\{0.75, 1.25\}$ to $\{0.5, 1.5\}$.

Note

It is tempting to shut down risk completely by setting $\gamma_l = \gamma_h$.
Doing so is numerically treacherous.
Without idiosyncratic risk, every agent of a given age chooses the same point on the asset grid, so aggregate asset supply becomes a step function of r and the price iteration can cycle rather than converge.
This is one reason for the iteration cap in `find_ss`.

Exercise 126.15.3

This exercise builds an unfunded social security system like the one studied in *Experiment 4 of the two-period model*.

Let the government tax each young agent and pay each old agent, with

$$\delta_j = \begin{cases} d & j < 25 \\ -d & j \geq 25 \end{cases}$$

so that $\sum_j \delta_j = 0$ and the scheme is balanced period by period.

1. Compute steady states for $d \in \{0, 0.1, 0.25\}$ and report K, r, w, and τ.
2. Report mean consumption of the young half and the old half of the population.
3. Compare your findings with the two-period model.

Solution

```
a_vec = hh.a_grid.reshape((1, hh.a_grid.size, 1))
y_vec = hh.y_grid.reshape((1, 1, hh.y_grid.size))
l_vec = l(hh.j_grid).reshape((hh.j_grid.size, 1, 1))
J = hh.j_grid.size

print(f"{'d':>5}  {'K':>7}  {'r':>7}  {'w':>7}  {'τ':>7}"
      f"  {'c young':>8}  {'c old':>7}")
for d in [0.0, 0.1, 0.25]:
    δ_ss = np.zeros(J)
    δ_ss[:J//2], δ_ss[J//2:] = d, -d
    out = find_ss(hh, firm, [0, 0.1, δ_ss], Q)
    σ_d, μ_d = out[1], out[2]
    K_d, r_d, w_d, τ_d = (float(out[3]), float(out[5]),
                          float(out[6]), float(out[7]))

    ap_d = hh.a_grid[σ_d].reshape((J, hh.a_grid.size, hh.y_grid.size))
    c_d = ((1 + r_d * (1 - τ_d)) * a_vec
           + (1 - τ_d) * w_d * l_vec * y_vec
           - δ_ss.reshape((J, 1, 1)) - ap_d)
    c_mean = (c_d * μ_d.reshape(ap_d.shape)).sum(axis=(1, 2))

    print(f"{d:>5.2f}  {K_d:>7.3f}  {r_d:>7.4f}  {w_d:>7.4f}  {τ_d:>7.4f}"
          f"  {c_mean[:J//2].mean():>8.4f}  {c_mean[J//2:].mean():>7.4f}")
```

```
     d         K         r         w         τ   c young    c old
 0.00     9.524    0.0653    1.3466    0.0482    1.0045   2.9427
 0.10     9.212    0.0668    1.3329    0.0487    0.9212   2.9841
 0.25     8.707    0.0696    1.3103    0.0496    0.7918   3.0462
```

The transfer scheme is balanced at every date, so it collects no net revenue and the flat rate tax τ barely moves.

Even so, it contracts the economy: capital falls from 9.5 to 8.7 as d rises to 0.25, the interest rate rises, and the wage falls.

The young consume less because they are taxed, and the old consume more because they are subsidized, but the young also save less, both because their income is lower and because the promised transfer substitutes for their own saving.

This is the same crowding out that *Experiment 4 of the two-period model* displays.

What the long-lived model adds is the observation that the transfer is collected from exactly those agents whose precautionary motive is strongest, namely the young, who hold the fewest assets with which to buffer labor income risk.

CHAPTER

SEVEN

TWO COMPUTATIONS TO FUND SOCIAL SECURITY

Contents

- *Two Computations to Fund Social Security*
  - *Overview*
  - *Model*
  - *Two experiments*
  - *Computation strategy*
  - *Calibration*
  - *Individual optimality*
  - *Transition path computation*
  - *Experiment 1: compensation through debt*
  - *Experiment 2: government capital accumulation*
  - *Distribution surfaces*
  - *Concluding remarks*
  - *Exercises*

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the “play” icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what’s in Anaconda, this lecture will need the following library

```
!pip install jax
```

127.1 Overview

This lecture describes two computational experiments about alternative ways to move gradually from an unfunded (pay-as-you-go) to a fully funded social security system, following Huang *et al.* [1997].

As populations age, pay-as-you-go social security systems have faced financial difficulties.

This situation has led some of today's policy makers and policy advisors to think that today's citizens would be better off if earlier policy makers had set up a fully-funded retirement system.

But starting from where we are today, a transition to a fully funded system creates distributional challenges because older generations who contributed to the unfunded system could lose benefits.

To study possibilities quantitatively, this lecture employs a general equilibrium overlapping generations model that modifies the Auerbach and Kotlikoff [1987] environment by incorporating

- risk-sensitive preferences
- uncertainty about lifetimes
- uninsurable labor income risk
- a theory of consumption distributions within and across cohorts

By employing the discounted risk-sensitive linear-quadratic preferences of Hansen and Sargent [1995], the model yields linear decision rules for individual consumption and savings.

This makes it computationally feasible to track the joint distribution of consumption and wealth across cohorts.

This lecture relates to two other lectures:

- *Transitions in an Overlapping Generations Model* studies how taxes, transfers, and debt affect capital accumulation in a two-period OLG model, introducing the Auerbach and Kotlikoff [1987] framework that this lecture extends.
- *A Long-Lived, Heterogeneous Agents, Overlapping Generations Model* studies how fiscal policy interacts with precautionary savings in a long-lived OLG model, using discrete dynamic programming rather than the linear-quadratic approach adopted here.

We use the following imports and configurations

```
import jax
import jax.numpy as jnp
from jax import jit, vmap
import jax.lax as lax
import numpy as np
import matplotlib.pyplot as plt
from collections import namedtuple

# Enable 64-bit precision
jax.config.update("jax_enable_x64", True)
```

127.2 Model

127.2.1 Environment

The economy consists of overlapping generations of finitely lived individuals who may live up to $T_0 + 1$ years and an infinitely lived government.

Individual consumers and the government can invest at a common risk-free gross rate of return.

That rate is constant when we treat the economy as a small open economy, and it varies over time when we close the economy and let marginal products determine factor prices.

During the first $T_1 + 1$ periods of life, consumers receive labor income that they allocate among consumption, taxes, and asset accumulation.

During the final $T_0 - T_1$ periods of life (retirement), consumers receive social security benefits and dissave by drawing down their assets.

The government taxes income from capital and labor, issues debt, purchases goods, and pays retirement benefits.

For any variable z, we use subscript t to denote age, argument s in parentheses to denote calendar time, and superscript $s - t$ to denote date of birth, so that $z^{s-t}(s) \equiv z_t(s) \equiv z_t^{s-t}(s)$.

Several symbols are close cousins, so we record them here.

Symbol	Meaning
t, T_0, T_1	age, maximum age, retirement age
s, s_1, s_2, s_3	calendar time and the dates that bound the reform
α_t	probability of surviving from age t to $t+1$
$\tilde{\alpha}$, $\tilde{A}$	capital's share and the scale factor in the production function
σ, σ_d	risk-sensitivity parameter and the standard deviation of the income shock
γ_t	preference shock in the one-period return function
ε_t, ϵ_{t+1}	age-efficiency of labor and the random shock
$T_t(s)$	taxes paid by an age-t person at time s
$\mathcal{T}_t$, $\mathcal{D}_t$, $\mathcal{S}_t$	the three operators in the risk-sensitive Riccati recursion

127.2.2 Demographics

At date s, a cohort of measure $N_0(s)$ consumers is born who live during periods $s, s+1, \ldots, s+T_0$.

As a cohort ages, its members face random survival according to age-to-age survival probabilities $\{\alpha_t\}_{t=0}^{T_0}$, where α_t is the probability of surviving from age t to $t+1$.

Let $N_t(s)$ be the number of age-t people alive at time s and let n be the constant gross rate of population growth.

The size of age group t at time s satisfies

$$N_t(s) = \lambda_t \cdot N_0(s-t)$$

where $\lambda_t = \prod_{j=0}^{t-1} \alpha_j$ for $t = 0, \ldots, T_0$, with $\lambda_0 = 1$, and births follow $N_0(s) = n^s \cdot N_0(0)$.

The population fraction of cohort t at each time s is

$$f_t = \frac{\lambda_t \cdot n^{-t}}{\sum_{\tau=0}^{T_0} \lambda_\tau \cdot n^{-\tau}}$$

and total population at time s is

$$N(s) = N_0(0) \cdot n^s \cdot \sum_{t=0}^{T_0} n^{-t} \cdot \lambda_t.$$

127.2.3 Distributions and aggregates

Individuals face life span uncertainty and labor income shocks.

They self-insure by accumulating risk-free assets (government bonds and physical capital), while properly taking into account the social security benefits that they anticipate receiving.

Let ϵ_0^t denote the history of random shocks that an individual has received from birth to age t.

The state vector $x_t(s) = x_t(s; \epsilon_0^t, x_0)$ measures the stock of assets as well as information variables a consumer uses to forecast future preferences or opportunities.

The model delivers consumption as a time- and age-dependent linear function of the state vector

$$c_t(s; \epsilon_0^t, x_0) = \eta_{ct}(s) \cdot x_t(s; \epsilon_0^t, x_0)$$

where the state vector follows a linear law of motion

$$x_{t+1}(s+1; \epsilon_0^{t+1}, x_0) = A_t(s) \cdot x_t(s; \epsilon_0^t, x_0) + C_t(s) \cdot \epsilon_{t+1}$$

with ϵ_{t+1} a martingale difference sequence satisfying $E(\epsilon_{t+1}|J_t) = 0$, $E(\epsilon_{t+1} \cdot \epsilon'_{t+1}|J_t) = I$, and $J_t = (\epsilon_0^t, x_0)$.

The model delivers probability distributions for state vectors.

Let $\mu_t(s) = E[x_t(s)]$ and $\Sigma_t(s) = E[(x_t(s) - \mu_t(s))(x_t(s) - \mu_t(s))']$.

These moments satisfy

$$\mu_{t+1}(s+1) = A_t(s) \cdot \mu_t(s)$$

$$\Sigma_{t+1}(s+1) = A_t(s) \cdot \Sigma_t(s) \cdot A_t(s)' + C_t(s) \cdot C_t(s)'$$

Per capita aggregate consumption is

$$c(s)/N(s) = \sum_{t=0}^{T_0} \mu_{ct}(s) \cdot f_t$$

where $\mu_{ct}(s)$ is the mean consumption of age-t people at time s.

The distribution of consumption within age cohort t has mean $\mu_{ct}(s) = \eta_{ct}(s) \cdot \mu_t(s)$ and variance $\Sigma_{ct}(s) = \eta_{ct}(s) \cdot \Sigma_t(s) \cdot \eta_{ct}(s)'$.

127.2.4 Resource constraint

The economy-wide physical resource constraint is

$$g(s) \cdot N(s) + \sum_{t=0}^{T_0} c_t(s) \cdot N_t^{s-t} + K(s) = R(s-1) \cdot K(s-1) + w(s) \cdot \sum_{t=0}^{T_1} \varepsilon_t \cdot N_t^{s-t} + N_0(s) \cdot k_{-1}(s)$$

where $g(s)$ is per capita government purchases, $K(s-1)$ is physical capital, $R(s-1) = 1 + r(s-1) - \delta$ is the gross return on assets, ε_t is the exogenous efficiency endowment of age-t people, $w(s)$ is the base wage rate, $N_0(s) \cdot k_{-1}(s)$ is capital brought by newborns, and δ is the depreciation rate.

Dividing by the population $N(s)$ and writing $k(s) = K(s)/N(s)$, $\bar{c}(s) = \sum_t \mu_{ct}(s) f_t$, and $\bar{\varepsilon} = \sum_t \varepsilon_t f_t$ puts the constraint in the per capita form that our code iterates on:

$$k(s) = R(s-1)\left[\frac{k(s-1)}{n} + f_0 k_{-1}\right] + w(s)\bar{\varepsilon} - g(s) - \bar{c}(s) \tag{127.1}$$

We also report two ratios below, each built from GDP measured as the sum of capital and labor income,

$$\text{GDP}(s) = r(s-1)\left[\frac{k(s-1)}{n} + f_0 k_{-1}\right] + w(s)\bar{\varepsilon}$$

127.2.5 Factor prices

We consider two alternative assumptions about factor prices:

- *Small open economy:* $r(s-1) = r$ and $w(s) = w$ are exogenous and constant.
- *Closed economy:* Factor prices are determined by marginal products from a Cobb-Douglas production function:

$$r(s-1) = \tilde{A} \cdot \tilde{\alpha} \cdot (K(s-1)/\tilde{N}(s))^{\tilde{\alpha}-1}, \qquad w(s) = (1-\tilde{\alpha}) \cdot \tilde{A} \cdot (K(s-1)/\tilde{N}(s))^{\tilde{\alpha}}$$

where $\tilde{N}(s) = \sum_{t=0}^{T_1} \varepsilon_t \cdot N_t^{s-t}$ is aggregate labor input in efficiency units and $\tilde{\alpha}$ is capital's share of income.

127.2.6 Consumers' problems

Individual consumers face an overlapping-generations version of a classic consumption-saving problem (see *The Permanent Income Model*).

Working-age consumers ($t \leq T_1$) receive labor income $w(s) \cdot \varepsilon_t + d_t$, where d_t is an AR(1) process

$$d_t = \rho_d \cdot d_{t-1} + \xi_t$$

with ξ_t being Gaussian white noise with variance σ_d^2.

The budget constraint at age t and time s is

$$c_t(s) + a_t(s) = R(s-1) \cdot a_{t-1}(s-1) + w(s) \cdot \varepsilon_t + S_t(s) - T_t(s) + d_t \tag{127.2}$$

where $a_{t-1}(s-1)$ is asset holdings at the beginning of age t, $S_t(s)$ is social security benefits (zero while working, $S(s)$ when retired), and taxes are $T_t(s) = \tau_0(s) + \tau_\ell(s)(w(s) \cdot \varepsilon_t + d_t) + \tau_a(s)(R(s-1)-1) \cdot a_{t-1}(s-1)$.

Following Hansen and Sargent [1995], preferences over stochastic consumption processes are defined recursively by

$$U_t = -(\pi \cdot c_t - \gamma_t)^2/2 + \beta_t \cdot \mathcal{R}_t(U_{t+1})$$

where $\mathcal{R}_t(U_{t+1}) = (2/\sigma) \cdot \log E[\exp(\sigma \cdot U_{t+1}/2)|J_t]$, σ is the risk-sensitivity parameter, and $\beta_t = \bar{\beta} \cdot \alpha_t$ is the survival-adjusted discount factor.

This preference specification delivers linear decision rules while allowing a form of risk-sensitivity that induces a type of precautionary savings.

When $\sigma = 0$, the recursion collapses to $U_t = -(\pi c_t - \gamma_t)^2/2 + \beta_t E_t U_{t+1}$, an ordinary discounted quadratic objective whose decision rules satisfy **certainty equivalence**: they do not depend on the volatility σ_d of the income shock.

When $\sigma < 0$, the operator $\mathcal{R}_t$ puts extra weight on bad continuation outcomes, so the consumer behaves as if pessimistic about future shocks.

Certainty equivalence then fails, decision rules depend on the noise statistics CC', and the consumer saves more.

Exercise 127.12.2 measures how much more.

127.2.7 Government

The government purchases goods, pays social security benefits, taxes capital and labor income, confiscates accidental bequests, and issues one-period bonds.

The government budget constraint is

$$g(s) \cdot N(s) + \sum_{t=T_1+1}^{T_0} S_t(s) \cdot N_t^{s-t} + R(s-1) \cdot \sum_{t=1}^{T_0} b_{t-1}(s-1) \cdot N_t^{s-t} =$$

$$\sum_{t=0}^{T_0} N_t^{s-t} \{\tau_a(s)[R(s-1)-1] \cdot a_{t-1}(s-1) + \tau_\ell(s) \cdot w(s) \cdot \varepsilon_t\} + \tau_0 \cdot N(s) +$$

$$\sum_{t=0}^{T_0} b_t(s) \cdot N_t^{s-t} + R(s-1) \cdot \sum_{t=0}^{T_0} (1-\alpha_t) \cdot k_t(s-1) \cdot N_t^{s-t-1}$$

where $b_t(s)$ is government debt held by age-t individuals at time s and the last term represents the bequest tax.

A transition between fiscal regimes is described by dates $0 \leq s_1 < s_2 < s_3$:

- Before $s = 0$: initial stationary equilibrium
- At $s = 0$: government announces a policy change
- $s \in [s_1, s_2)$: fiscal parameters and the social security system change
- $s \geq s_2$: new constant policy parameters
- $s \geq s_3$: final stationary equilibrium

127.2.8 Equilibrium

We define the following objects:

- An **allocation** is a stochastic process for $\{c_t(s), a_t(s)\}_{s=0}^{s_3}$ for $t = 0, \ldots, T_0$, and a sequence $\{K(s)\}_{s=0}^{s_3}$.
- A **government policy** is a sequence $\{b(s), g(s), \tau_\ell(s), S(s), \tau_a(s)\}_{s=0}^{s_3}$.
- A **price system** is a sequence $\{w(s), r(s-1)\}_{s=0}^{s_3}$.

An **equilibrium** is an allocation, a price system, and a government policy such that

1. given the price sequence and government policy, the allocation solves households' optimum problems,
2. the allocation and government policy satisfy the government budget constraint at each date s,
3. the asset market clears, so that household assets equal physical capital plus government debt, $\sum_t a_t(s) N_t^{s-t} = K(s) + \sum_t b_t(s) N_t^{s-t}$, and
4. the goods market clears, which is the resource constraint (127.1).

In the closed economy, condition 4 together with the firm's first-order conditions determines factor prices; in the small open economy, factor prices are given and condition 4 determines how much the economy borrows from or lends to the rest of the world.

In a **stationary equilibrium**, all variables are independent of calendar time s, which simplifies the government budget constraint to

$$g + \sum_{t=T_1+1}^{T_0} S_t \cdot f_t + [R/n - 1] \cdot \bar{b} = \tau_a (R-1) \cdot \sum_{t=0}^{T_0} a_{t-1} \cdot f_t + \tau_\ell \cdot \sum_{t=0}^{T_1} w \cdot \varepsilon_t \cdot f_t + \frac{R}{n} \cdot \sum_{t=0}^{T_0} (1-\alpha_t) \cdot a_t \cdot f_t$$

where $\bar{b}$ is per capita government debt in steady state.

127.2.9 Transition dynamics

The following diagram shows the age-time structure during the transition.

Each horizontal line represents the lifetime of one cohort, and the vertical dashed lines mark the policy change dates s_1 and s_2.

```
fig, ax = plt.subplots(figsize=(8, 5))

birth_range = np.arange(-60, 41, 5)

ax.hlines(birth_range, birth_range, birth_range + 60, 'k', linewidth=1)
ax.vlines([0, 40], -60, 40, 'k', linestyle='--', linewidth=0.5)

ax.set_ylabel("date born")
ax.set_xlabel("time")

ax.text(-20, 25, r"$s_1=0$")
ax.text(43, -50, r"$s_2$")

ax.invert_yaxis()
plt.show()
```

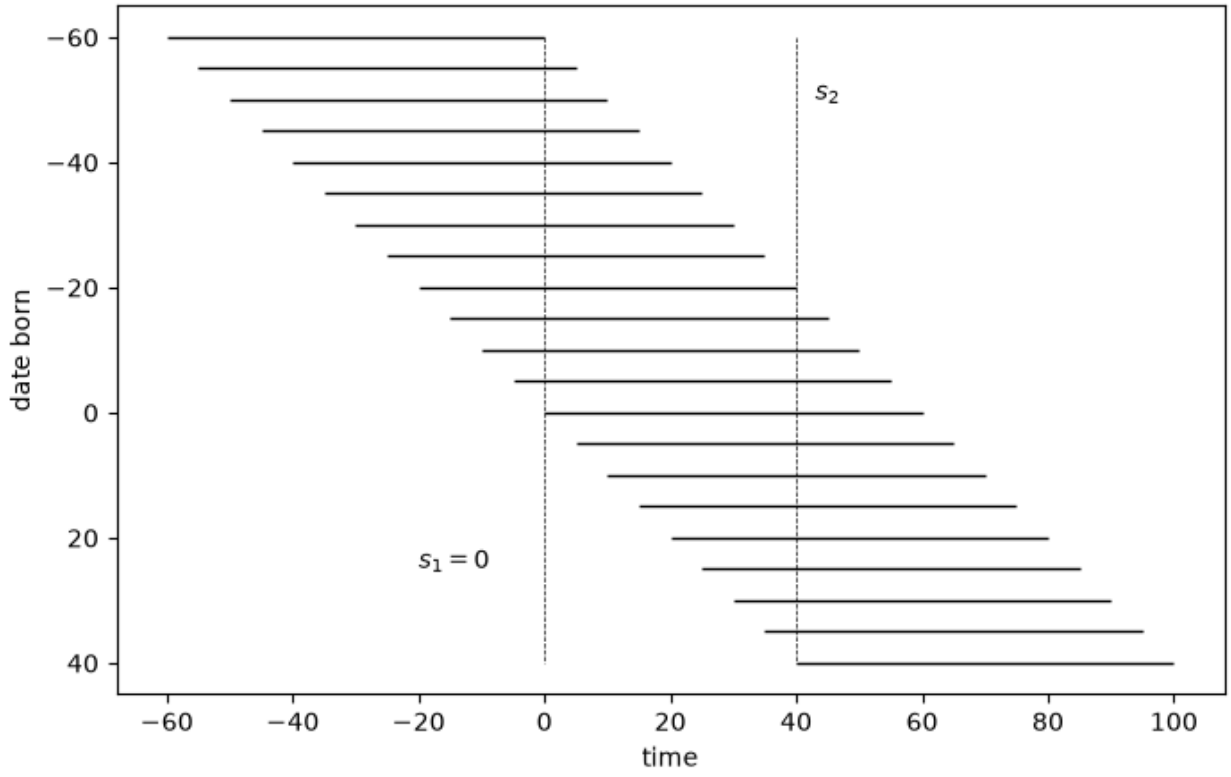


Fig. 127.1: Age-time diagram for overlapping generations

Small open economy

With fixed factor prices, the transition has a clear structure:

- Cohorts born before $s_1 - T_0$ die before the policy change and are unaffected.
- Cohorts alive at s_1 must recalculate their consumption-saving plans for their remaining lifetimes.
- Cohorts born between s_1 and s_2 face time-varying tax and benefit rates, while those born after s_2 face constant parameters.
- The transition ends at $s_3 = s_2 + T_0$ when the last cohort that experienced the policy change has died.

Because factor prices are fixed, we can compute the transition by solving decision rules for cohorts born at dates $s_1 - T_0 - 1, \ldots, s_2$.

For any date s, aggregate consumption is computed by summing across all living cohorts (along a vertical line in the age-time diagram), weighted by their population fractions.

Closed economy

With endogenous factor prices, the transition is more complex:

- Factor prices continue to evolve after policy parameters stabilize at s_2, so we follow Auerbach and Kotlikoff [1987] and truncate at $s_3 = s_2 + 2T_0$.
- The computation requires nested iteration: an inner loop determines labor income tax rates, and an outer loop adjusts interest rates to clear factor markets.
- Changes in saving behavior affect capital accumulation, which alters marginal products and feeds back into household decisions.
- Lower interest rates benefit young workers through higher wages but hurt retirees through lower returns on savings.

127.3 Two experiments

We explore two strategies for transitioning to a fully funded social security system.

In experiment 1, the government terminates social security benefits but compensates entitled generations through a one-time increase in government debt.

In experiment 2, the government retains social security benefits but temporarily raises taxes to accumulate physical capital, the returns from which eventually finance social security payments.

Both proposals aim to finance a transition to fully funded social security without leaving existing generations worse off, but they entail different amounts of intergenerational risk-sharing.

We compute both experiments under fixed and endogenous factor prices and compare outcomes below.

127.4 Computation strategy

127.4.1 Dynamic program

An individual consumer's problem can be formulated as a discounted risk-sensitive linear control problem (see *LQ Control: Foundations*).

Let $x_t = [a_{t-1}, z_t]'$ where z_t is the vector of shocks.

The optimal value function takes the form $U_t = x_t' \cdot P_t \cdot x_t + \xi_t$.

The recursive problem is

$$U_t = \max_{u_t} \left\{ u_t' Q_t u_t + x_t' R_t x_t + \frac{2\beta_t}{\sigma} \log E_t[\exp(\sigma U_{t+1}/2)] \right\}$$

subject to $x_{t+1} = A_t x_t + B_t u_t + C_t w_{t+1}$.

We deploy two operators

$$\begin{aligned}
\mathcal{T}_t(P) &= P + \sigma P C_t \left(I - \sigma C_t' P C_t\right)^{-1} C_t' P \\
\mathcal{D}_t(W) &= R_t + A_t' \left[\beta_t W - \beta_t^2 W B_t \left(Q_t + \beta_t B_t' W B_t\right)^{-1} B_t' W\right] A_t \\
\mathcal{S}_t(k, P) &= \beta_t k - (\beta_t/\sigma) \log\det\left(I - \sigma C_t' P C_t\right)
\end{aligned}$$

that we use to construct a value function recursion $P_t = (\mathcal{D}_t \circ \mathcal{T}_t) P_{t+1}$, $\xi_t = \mathcal{S}_t(\xi_{t+1}, P_{t+1})$, and an optimal control

$$u_t = -F_t x_t, \qquad F_t = \beta_t \left[Q_t + \beta_t B_t' \mathcal{T}_t(P_{t+1}) B_t\right]^{-1} B_t' \mathcal{T}_t(P_{t+1}) A_t.$$

The operator $\mathcal{T}_t$ is where risk sensitivity enters: at $\sigma = 0$ it is the identity, so the decision rule F_t ignores C_t and certainty equivalence holds.

Operators $\mathcal{T}_t$, $\mathcal{D}_t$, $\mathcal{S}_t$ and decision rule F_t are constructed in `solve_riccati_step`.

Given value function parameters (P_{t+1}, ξ_{t+1}) at the next age, it constructs the state-space matrix A_t, applies the cross-product trick, evaluates the Riccati operators, and returns the optimal decision rule F_t, the closed-loop matrix $A_t^o = A_t - BF_t$, and the updated (P_t, ξ_t)

```
def solve_riccati_step(
        ε_t, β_t, Ind_work_t,
        RR, w, τ_l, τ_a, τ_0, benef,
        P_next, ξ_next,
        ρ_d, σ, B, C, R, Q, H):
    """One backward step of the risk-sensitive Riccati recursion."""

    A = jnp.array([
        [RR * (1.0 - τ_a) + τ_a,
         (1.0 - τ_l) * w * ε_t
         - τ_0 + benef * (1.0 - Ind_work_t),
         (1.0 - τ_l) * Ind_work_t],
        [0.0, 1.0, 0.0],
        [0.0, 0.0, ρ_d]
    ])

    Q_scalar = Q[0, 0]
    Q_inv_scalar = 1.0 / Q_scalar
    Q_inv = jnp.array([[Q_inv_scalar]])
```

(continues on next page)

(continued from previous page)

```
    # Cross-product trick: A* = A - B Q^{-1} H
    A = A - B @ Q_inv @ H

    # 𝒞_t operator
    CTP = C.T @ P_next @ C
    PP_scalar = 1.0 - σ * CTP[0, 0]
    PP_inv_scalar = 1.0 / PP_scalar
    PC = P_next @ C
    CP = C.T @ P_next
    TP = P_next + σ * PP_inv_scalar * (PC @ CP)

    # 𝒟_t operator and decision rule F_t
    BTB_scalar = (B.T @ TP @ B)[0, 0]
    Q_BTB_scalar = Q_scalar + β_t * BTB_scalar
    Q_BTB_inv_scalar = 1.0 / Q_BTB_scalar
    BT_TP = B.T @ TP
    BT_TP_A = BT_TP @ A
    F = β_t * Q_BTB_inv_scalar * BT_TP_A

    TP_B = TP @ B
    middle = (β_t * TP
              - β_t**2 * Q_BTB_inv_scalar
              * (TP_B @ BT_TP))
    P = R + A.T @ middle @ A

    # 𝒪_t operator
    log_det_PP = jnp.log(PP_scalar)
    ξ = jnp.where(
        σ != 0.0,
        β_t * (ξ_next - log_det_PP / σ),
        β_t * (ξ_next + CTP[0, 0])
    )

    Ao = A - B @ F
    F = F + Q_inv @ H

    return F.squeeze(), Ao, P, ξ
```

127.4.2 State space preparation

The budget constraint (127.2) and the income process can be written in state-space form.

Let $x_t = [a_{t-1}(s-1),\ 1,\ d_t]'$, $u_t = c_t(s)$, and $w_{t+1} = \epsilon_{t+1}$, so that

$$x_{t+1} = A_t x_t + B u_t + C w_{t+1}$$

where

$$A_t = \begin{bmatrix} R(s-1)(1-\tau_t^a(s)) + \tau_t^a(s) & (1-\tau_\ell(s))w(s)\varepsilon_t - \tau_0(s) + S_t(s) & \mathbf{1}_t^{\text{work}}(1-\tau_\ell(s)) \\ 0 & 1 & 0 \\ 0 & 0 & \rho_d \end{bmatrix},$$

$$B = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix},$$

$$C = \begin{bmatrix} 0 \\ 0 \\ \sigma_d \end{bmatrix}$$

and $\mathbf{1}_t^{\text{work}}$ indicates whether the agent is of working age.

The per-period return $-\frac{1}{2}(\pi c_t - \gamma)^2$ introduces a cross-product term H between the control and the state.

This is eliminated using the cross-product trick (see *Eliminating Cross Products*):

$$A^* = A - BQ^{-1}H, \qquad R^* = R - H'Q^{-1}H.$$

127.4.3 Means and covariances

Define $A_o = A - BF$ as the closed-loop transition matrix, so that $x_{t+1} = A_o x_t + C w_{t+1}$ and unconditional moments satisfy

$$\mu_{t+1} = A_{o,t}\mu_t, \qquad \Sigma_{t+1} = A_{o,t}\Sigma_t A_{o,t}' + CC'.$$

The moment recursion is implemented as `forward_moment_step`, which propagates the mean vector and covariance matrix by one age step and computes consumption statistics as by-products.

```
def forward_moment_step(μx_t, Σx_t, Ao_t, F_t, CCT):
    """One step of the forward moment recursion."""

    μx_next = Ao_t @ μx_t
    μc_t = -F_t @ μx_t
    Σx_next = CCT + Ao_t @ Σx_t @ Ao_t.T
    Vc_t = F_t @ Σx_t @ F_t.T
    return μx_next, μc_t, Σx_next, Vc_t
```

127.4.4 Computing transitions

A cohort born at s lives during $s, s+1, \ldots, s+T_0$ and works during $s, s+1, \ldots, s+T_1$.

Let $0 \leq s_1 < s_2 < s_3$.

- At $s = 0$, the government announces a policy change between $s = s_1$ and $s = s_2$.
- From $s = s_2$, government policies are constant forever.
- From $s = s_3$, convergence to the final stationary equilibrium is achieved (in the small open economy, $s_3 = s_2 + T_0$).

The affected cohorts are those born at $s = s_1 - T_0, s_1 - T_0 + 1, \ldots, s_2$.

In all exercises, we set $T_0 = 65$, $T_1 = 43$, $s_1 = 0$, and $s_2 = 40$.

127.5 Calibration

The model parameters are set as follows.

127.5.1 Preference parameters

Parameter	Description	Value
$\{\alpha_t\}_{t=0}^{T_0}$	Age-to-age survival probabilities	Faber [1982]
π	Consumption preference parameter	1.0
σ	Risk-sensitivity parameter	-0.05
$\bar{\gamma}$	Preference shock parameter	7.0
$\tilde{\beta}$	Discount factor	0.986
T_0	Maximum age	65
T_1	Retirement age	43
n	Gross population growth rate	1.012

127.5.2 Technology parameters

Parameter	Description	Value
k_{-1}	Initial capital endowment	4.0
σ_d	Standard deviation of income shock	0.85
ρ_d	Persistence of income shock	0.8
δ	Depreciation rate	0.06
$\{\varepsilon_t\}_{t=0}^{T_1}$	Age-efficiency profile	Hansen [1993]
w	Base wage rate (exogenous)	5.0147
r	Return on capital (exogenous)	0.1275
$\tilde{A}$	Production function scaling (endogenous)	2.2625
$\tilde{\alpha}$	Capital share (endogenous)	0.40

```
N_GRID_SS = 10
TOL_SS = 1e-10

T0 = 65   # maximum lifespan (ages 21 to 86)
T1 = 43   # working life length (retire at 65)

UNIT_GRID = jnp.linspace(0.0, 1.0, N_GRID_SS)
AGE_INDICES = jnp.arange(T0 + 2)
```

The hidden code cell below defines the age-efficiency profile $\{\varepsilon_t\}$ and the survival probabilities $\{\alpha_t\}$ based on Faber [1982] and Hansen [1993].

```
fig, axs = plt.subplots(1, 2, figsize=(10, 6))

axs[0].plot(ε_arr)
axs[0].set_title("Working efficiency")
axs[0].set_xlabel("age")

axs[1].plot(α_arr)
```

(continues on next page)

(continued from previous page)

```
axs[1].set_title("Survival probability")
axs[1].set_xlabel("age")

plt.tight_layout()
plt.show()
```

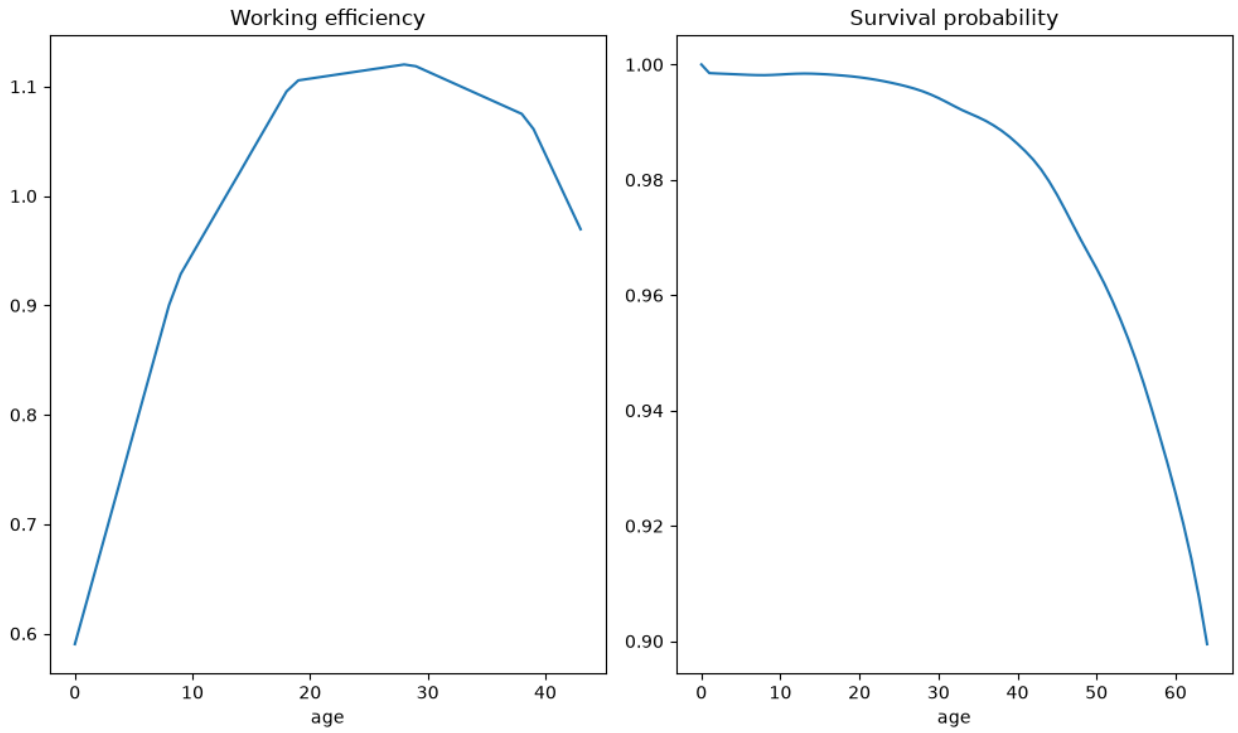


Fig. 127.2: Age-efficiency profile and survival probabilities

We impose a large penalty on terminal asset holdings to enforce the end-of-life condition, and set the initial state to $x_0 = [k_{-1},\ 1,\ 0]'$.

```
P_end = jnp.zeros((3, 3))
P_end = P_end.at[0, 0].set(-2000000.0)
ξ_end = 0.0

x0 = jnp.array([4.0, 1.0, 0.0])
Σ0 = jnp.zeros((3, 3))
```

All household parameters are collected into a named tuple.

```
Household = namedtuple('Household', (
    'α_arr', 'frac', 'n', 'π', 'σ', 'k_init', 'ε_arr', 'Ind_work',
    'σ_d', 'ρ_d', 'γ_bar', 'β_arr',
    'T0', 'T1', 'T2', 'n_x', 'n_w',
    'P_end', 'ξ_end', 'x0', 'Σ0',
    'B', 'C', 'R', 'Q', 'H'
))
```

(continues on next page)

(continued from previous page)

```
def create_household(α_arr=α_arr,     # Age-to-age survival probabilities
                     n=1.012,         # Gross population growth rate
                     π=1,             # Consumption preference parameter
                     σ=-0.05,         # Risk-sensitivity parameter
                     k_init=4,        # initial capital endowment
                     ε_arr=ε_arr,     # age-efficiency profile
                     σ_d=0.85,        # std of income shock
                     ρ_d=0.8,         # persistence of income shock
                     γ_bar=7,         # Preference shock parameter
                     β_tilde=0.986,   # Discount factor
                     T0=65,           # Maximum age
                     T1=43,           # Retirement age
                     n_x=3,           # Number of states
                     n_w=1,           # Number of shocks
                     P_end=P_end,     # Terminal value
                     ξ_end=ξ_end,     # Terminal value
                     x0=x0,           # Initial mean
                     Σ0=Σ0):          # Initial variance
    """Create a Household named tuple with derived arrays."""

    α_arr = np.concatenate([α_arr, np.array([0])])
    T2 = T0 - T1

    frac = np.ones(T0 + 1)
    frac[1:] = np.cumprod(α_arr / n)[:-1]
    frac = frac / frac.sum()

    ε_arr = np.concatenate([ε_arr, np.zeros(T0 + 1 - ε_arr.size)])

    # Indicator for working ages: 1 if working (ε > 0), 0 if retired
    Ind_work = (ε_arr != 0).astype(np.float64)

    β_arr = β_tilde * α_arr
    β_arr[-1] = β_tilde

    B = jnp.array([[-1.0, 0.0, 0.0]]).T
    C = jnp.array([[0.0, 0.0, σ_d]]).T

    Q = jnp.array([[-0.5 * π**2]])
    H = jnp.array([[0.0, 0.5 * π * γ_bar, 0.0]])

    # Apply cross-product trick: R* = R - H'Q^{-1}H
    R_base = np.array([[0.0, 0.0, 0.0],
                       [0.0, -0.5 * γ_bar**2, 0.0],
                       [0.0, 0.0, 0.0]])
    H_np = np.array([[0.0, 0.5 * π * γ_bar, 0.0]])
    Q_inv_np = np.array([[1.0 / (-0.5 * π**2)]])
    R = jnp.array(R_base - H_np.T @ Q_inv_np @ H_np)

    return Household(
        α_arr=jnp.array(α_arr), frac=jnp.array(frac), n=n, π=π, σ=σ,
        k_init=k_init, ε_arr=jnp.array(ε_arr), Ind_work=jnp.array(Ind_work),
        σ_d=σ_d, ρ_d=ρ_d, γ_bar=γ_bar, β_arr=jnp.array(β_arr),
        T0=T0, T1=T1, T2=T2, n_x=n_x, n_w=n_w,
        P_end=P_end, ξ_end=ξ_end, x0=x0, Σ0=Σ0,
        B=B, C=C, R=R, Q=Q, H=H
    )
```

```
hh = create_household()
```

The stationary population distribution follows.

```
fig, ax = plt.subplots()
ax.plot(hh.frac)
ax.set_xlabel("age")
ax.set_ylabel("population fraction")
plt.show()
```

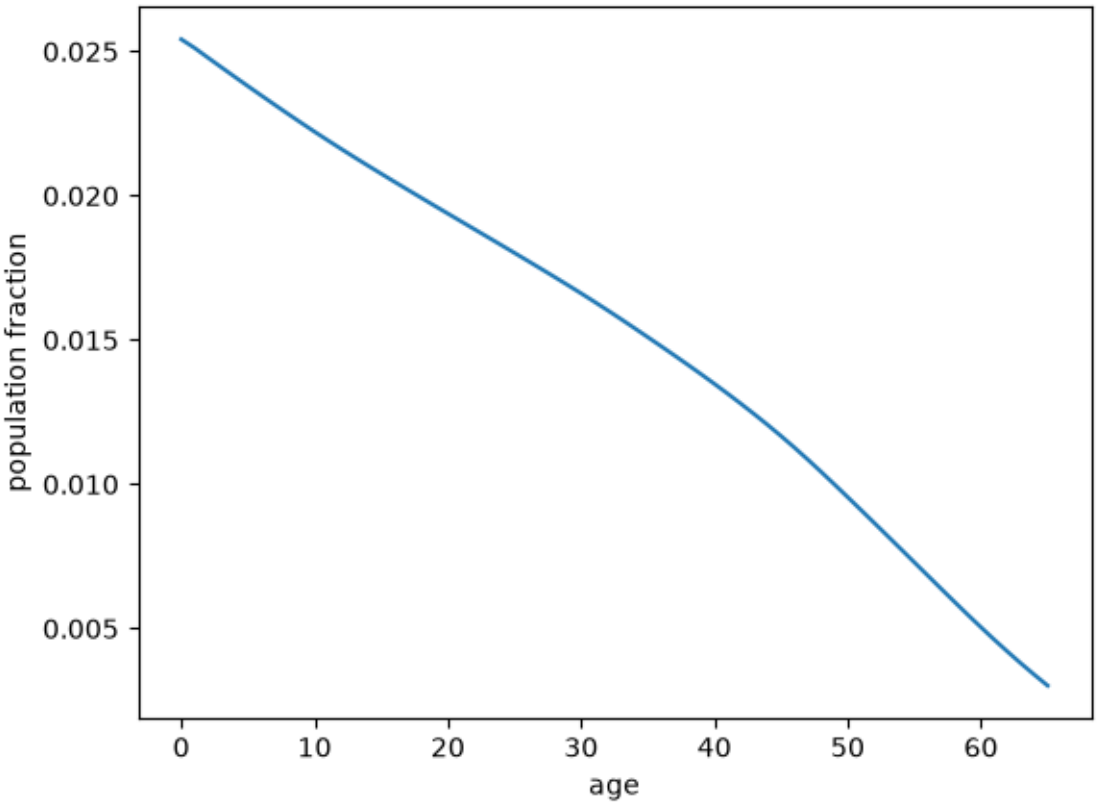


Fig. 127.3: Stationary population distribution over age

Mortality causes the population fraction to decline with age, a demographic pattern central to the intergenerational redistribution that social security reform entails.

Under the small open economy assumption, factor prices are fixed at calibrated values; under the closed economy assumption, they are determined by Cobb-Douglas marginal products

```
Tech = namedtuple('Tech', ('δ', 'w', 'r', 'RR', 'A', 'α_tilde'))

def create_Tech(δ=0.06, w=5.0147, r=0.1275,
                A=2.2625, α_tilde=0.40):
    """Create a Tech named tuple with factor price parameters."""

    RR = 1 + r - δ
```

(continues on next page)

(continued from previous page)

```
    return Tech(δ=δ, w=w, r=r, RR=RR, A=A, α_tilde=α_tilde)
```

```
tech = create_Tech()
```

127.6 Individual optimality

127.6.1 Steady-state computation

With `solve_riccati_step` and `forward_moment_step` in hand, the steady-state computation proceeds in three phases.

Algorithm 127.6.1 (Steady-state computation)

1. **Backward recursion:** scan from age T_0 down to 0. At each age t, call `solve_riccati_step` to obtain decision rule F_t, closed-loop matrix A_t^o, value-function matrix P_t, and certainty-equivalent ξ_t.
2. **Forward simulation:** scan from age 0 up to T_0. At each age t, call `forward_moment_step` to propagate mean $\mu_{x,t}$ and covariance $\Sigma_{x,t}$ of the state vector, and record mean consumption $\mu_{c,t}$ and its variance $V_{c,t}$.
3. **Budget imbalance:** aggregate across cohorts. Sum tax revenues (labor, capital, lump-sum), subtract benefit payments, add accidental bequests, and return the government budget gap.

Phase 1. The backward recursion scans from age T_0 to 0, applying `solve_riccati_step` at each age

```
def _ss_backward_recursion(
        ε_arr, β_arr, Ind_work,
        RR, w, τ_l, τ_a, τ_0, benef,
        P_end, ξ_end,
        ρ_d, σ, B, C, R, Q, H):
    """Backward Riccati scan over all ages."""

    ε_rev = ε_arr[::-1]
    β_rev = β_arr[::-1]
    Ind_work_rev = Ind_work[::-1]

    def backward_step(carry, inputs):
        P_next, ξ_next = carry
        ε_t, β_t, Ind_work_t = inputs
        F, Ao, P, ξ = solve_riccati_step(
            ε_t, β_t, Ind_work_t,
            RR, w, τ_l, τ_a, τ_0, benef,
            P_next, ξ_next,
            ρ_d, σ, B, C, R, Q, H
        )
        return (P, ξ), (F, Ao, P, ξ)

    init_carry = (P_end, ξ_end)
    _, (F_rev, Ao_rev, P_rev, ξ_rev) = lax.scan(
        backward_step, init_carry,
        (ε_rev, β_rev, Ind_work_rev)
    )
```

(continues on next page)

(continued from previous page)

```
    F_arr = F_rev[::-1]
    Ao_arr = Ao_rev[::-1]
    P_inner = P_rev[::-1]
    ξ_inner = ξ_rev[::-1]

    P_arr = jnp.concatenate(
        [P_inner, P_end[None, :, :]], axis=0
    )
    ξ_arr = jnp.concatenate(
        [ξ_inner, jnp.array([ξ_end])]
    )

    return F_arr, Ao_arr, P_arr, ξ_arr
```

Phase 2. The forward recursion propagates means and covariances from age 0 to T_0

```
def _ss_forward_simulation(
        Ao_arr, F_arr, x0, Σ0, C):
    """Forward moment scan using forward_moment_step."""

    CCT = C @ C.T

    def forward_step(carry, inputs):
        μx_t, Σx_t = carry
        Ao_t, F_t = inputs
        result = forward_moment_step(
            μx_t, Σx_t, Ao_t, F_t, CCT
        )
        μx_next, μc_t, Σx_next, Vc_t = result
        return (μx_next, Σx_next), \
            (μx_next, μc_t, Σx_next, Vc_t)

    init_carry = (x0, Σ0)
    _, (μx_scn, μc_arr, Σx_scn, Vc_arr) = lax.scan(
        forward_step, init_carry, (Ao_arr, F_arr)
    )

    μx_arr = jnp.concatenate(
        [x0[None, :], μx_scn], axis=0
    )
    Σx_arr = jnp.concatenate(
        [Σ0[None, :, :], Σx_scn], axis=0
    )

    return μx_arr, μc_arr, Σx_arr, Vc_arr
```

Phase 3. Aggregating tax revenues, benefit payments, and accidental bequests across cohorts gives the government budget gap

```
def _ss_budget_imbalance(
        μx_arr, ε_arr, frac, n, α_arr,
        RR, w, τ_l, τ_a, τ_0, benef,
        G, Gb, Ind_work):
    """Aggregate tax revenues and expenditures."""

    μa_arr = μx_arr[1:, 0]
```

(continues on next page)

(continued from previous page)

```
    μa_last_arr = μx_arr[:-1, 0]

    τ_l_tot = jnp.sum(τ_l * ε_arr * w * frac)
    τ_a_tot = jnp.sum(
        τ_a * (RR - 1.0) * μa_last_arr * frac
    )
    τ_0_tot = jnp.sum(τ_0 * frac)

    retired_mask = 1.0 - Ind_work
    benef_tot = jnp.sum(benef * frac * retired_mask)

    Beq = jnp.sum(
        RR * (1.0 - α_arr) * frac * μa_arr / n
    )

    T_tot = τ_l_tot + τ_a_tot + τ_0_tot + Beq
    diff = (G + benef_tot - T_tot
            + (RR / n - 1.0) * Gb)

    return diff
```

A steady state is found when the budget gap equals zero.

`ss_imbalance` chains the three phases into a single JIT-compiled function: backward recursion, forward simulation, and budget gap

```
@jit
def ss_imbalance(price, policy, α_arr, ε_arr, frac,
                 n, β_arr, ρ_d, σ, B, C, R, Q, H,
                 P_end, ξ_end, x0, Σ0, Ind_work):
    """Backward solve, forward simulate, and return budget gap."""

    RR, w = price
    τ_l, τ_a, τ_0, benef, G, Gb = policy

    F_arr, Ao_arr, P_arr, ξ_arr = \
        _ss_backward_recursion(
            ε_arr, β_arr, Ind_work,
            RR, w, τ_l, τ_a, τ_0, benef,
            P_end, ξ_end,
            ρ_d, σ, B, C, R, Q, H
        )

    μx_arr, μc_arr, Σx_arr, Vc_arr = \
        _ss_forward_simulation(
            Ao_arr, F_arr, x0, Σ0, C
        )

    diff = _ss_budget_imbalance(
        μx_arr, ε_arr, frac, n, α_arr,
        RR, w, τ_l, τ_a, τ_0, benef,
        G, Gb, Ind_work
    )

    return (diff, P_arr, ξ_arr, Ao_arr, F_arr,
            μx_arr, μc_arr, Σx_arr, Vc_arr)
```

A named tuple `SteadyState` collects value-function parameters, decision rules, moments, and aggregate statistics into

a single object that the transition solver can unpack

```
SteadyState = namedtuple("SteadyState", (
    "P_arr",      # Value function matrices by age
    "ξ_arr",      # Certainty equivalent adjustments by age
    "Ao_arr",     # Closed-loop transition matrices by age
    "F_arr",      # Decision rule matrices by age
    "μx_arr",     # Mean state vectors by age
    "μc_arr",     # Mean consumption by age
    "Σx_arr",     # Covariance matrices by age
    "Vc_arr",     # Consumption variances by age
    "debt2gdp",   # Government debt to GDP ratio
    "τ_l",        # Labor income tax rate
    "benef",      # Social security benefit level
    "Gb",         # Per-capita government debt
    "k_bar",      # Per-capita capital stock
    "RR",         # Gross return on assets
    "w",          # Wage rate
    "r",          # Interest rate (before depreciation)
    "k2gdp"       # Capital to GDP ratio
))
```

Given all other fiscal instruments, the equilibrium τ_ℓ is the value that zeroes the budget gap.

We find it by iterative grid refinement: evaluate the gap on a coarse grid, zoom into the best interval, and repeat

```
def _grid_refine(eval_fn, a_init, b_init, unit_grid, tol, max_iter):
    """Iterative grid-refinement root search.

    Must be called inside @jit functions.
    """

    n_grid = unit_grid.shape[0]

    def cond_fn(state):
        a, b, best_val, i = state
        return (jnp.abs(best_val) > tol) & (i < max_iter)

    def body_fn(state):
        a, b, _, i = state
        grid = a + (b - a) * unit_grid
        diffs = vmap(eval_fn)(grid)
        best_idx = jnp.argmin(jnp.abs(diffs))
        best_val = diffs[best_idx]
        idx_lo = jnp.maximum(best_idx - 1, 0)
        idx_hi = jnp.minimum(best_idx + 1, n_grid - 1)
        return (grid[idx_lo], grid[idx_hi], best_val, i + 1)

    grid = a_init + (b_init - a_init) * unit_grid
    diffs = vmap(eval_fn)(grid)
    best_idx = jnp.argmin(jnp.abs(diffs))
    best_val = diffs[best_idx]
    idx_lo = jnp.maximum(best_idx - 1, 0)
    idx_hi = jnp.minimum(best_idx + 1, n_grid - 1)

    init_state = (grid[idx_lo], grid[idx_hi], best_val, 0)
    final_state = lax.while_loop(cond_fn, body_fn, init_state)
    a_final, b_final, _, _ = final_state
    return (a_final + b_final) / 2.0
```

`_ss_diff_for_tau_l` evaluates the budget gap at a given τ_ℓ, and `_find_ss_tau_l` wraps it inside the grid-refine loop.

```
def _ss_diff_for_tau_l(τ_l, price_arr, policy_no_τl, α_arr, ε_arr, frac, n,
                       β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
                       Ind_work):
    """Budget imbalance for a given τ_l."""

    τ_a, τ_0, benef, G, Gb = policy_no_τl
    policy_arr = jnp.array([τ_l, τ_a, τ_0, benef, G, Gb])
    diff, *_ = ss_imbalance(
        price_arr, policy_arr,
        α_arr, ε_arr, frac, n,
        β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
        Ind_work
    )
    return diff
```

```
@jit
def _find_ss_tau_l(price_arr, policy_no_τl, α_arr, ε_arr, frac, n,
                   β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
                   Ind_work, unit_grid):
    """Find τ_l that balances the steady-state budget."""

    def compute_diff(τ_l):
        return _ss_diff_for_tau_l(
            τ_l, price_arr, policy_no_τl, α_arr, ε_arr, frac, n,
            β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
            Ind_work
        )

    return _grid_refine(compute_diff, -0.5, 1.0 - 1e-5, unit_grid, TOL_SS, 10)
```

GDP is the sum of capital and labor income shares

```
def _compute_gdp(μa_arr, frac, ε_arr, Gb, r, w, n, x0_0, frac_0):
    """Compute GDP from aggregates."""

    ε_agg = jnp.sum(frac * ε_arr)
    a_agg = jnp.sum(frac * μa_arr)
    k_agg = a_agg - Gb
    k_share = r * (k_agg / n + frac_0 * x0_0)
    l_share = w * ε_agg
    return k_agg, k_share + l_share
```

`find_ss_exo` ties the pieces together: it solves for τ_ℓ, evaluates the full steady state, and returns a `SteadyState` named tuple

```
def find_ss_exo(price, policy_target, hh, tech):
    """Find steady state with exogenous prices by solving for τ_l."""

    frac, ε_arr, n, x0 = hh.frac, hh.ε_arr, hh.n, hh.x0
    RR, w = price
    r = RR - 1 + tech.δ

    τ_a, τ_0, benef, G, Gb = policy_target
```

(continues on next page)

(continued from previous page)

```
    price_arr = jnp.array([RR, w])
    policy_no_τl = jnp.array([τ_a, τ_0, benef, G, Gb])

    τ_l = _find_ss_tau_l(
        price_arr, policy_no_τl,
        hh.α_arr, hh.ε_arr, hh.frac, hh.n,
        hh.β_arr, hh.ρ_d, hh.σ, hh.B, hh.C, hh.R, hh.Q, hh.H,
        hh.P_end, hh.ξ_end, hh.x0, hh.Σ0,
        hh.Ind_work, UNIT_GRID
    )

    price_arr = jnp.array([RR, w])
    policy_arr = jnp.array([float(τ_l), τ_a, τ_0, benef, G, Gb])
    diff, P_arr, ξ_arr, Ao_arr, F_arr, μx_arr, μc_arr, Σx_arr, Vc_arr = \
        ss_imbalance(
            price_arr, policy_arr,
            hh.α_arr, hh.ε_arr, hh.frac, hh.n,
            hh.β_arr, hh.ρ_d, hh.σ,
            hh.B, hh.C, hh.R, hh.Q, hh.H,
            hh.P_end, hh.ξ_end, hh.x0, hh.Σ0,
            hh.Ind_work
        )

    k_agg, gdp = _compute_gdp(
        μx_arr[1:, 0], frac, ε_arr,
        Gb, r, w, n, x0[0], frac[0]
    )
    debt2gdp = Gb / gdp
    k2gdp = k_agg / gdp

    return SteadyState(
        P_arr=P_arr, ξ_arr=ξ_arr, Ao_arr=Ao_arr, F_arr=F_arr,
        μx_arr=μx_arr, μc_arr=μc_arr, Σx_arr=Σx_arr, Vc_arr=Vc_arr,
        debt2gdp=float(debt2gdp), τ_l=float(τ_l), benef=benef, Gb=Gb,
        k_bar=float(k_agg), RR=RR, w=w, r=float(r), k2gdp=float(k2gdp)
    )
```

The initial fiscal policy sets a social security replacement rate of $\theta = 0.6$

```
aveinc = tech.w * sum(hh.ε_arr) / (hh.T1 + 1)
θ = 0.6
benef_0 = aveinc * θ

G_0 = 1.44          # government purchases
Gb_0 = 2.8 * G_0    # government debt
τ_l_0 = 0.3385      # labor income tax
τ_a_0 = 0.30        # capital income tax
τ_0_0 = 0           # lump-sum tax

RR, w = tech.RR, tech.w
```

127.6.2 Initial and terminal steady states

The initial steady state features a calibrated replacement rate ($\theta = 0.6$), positive social security benefits, and government expenditure and debt set to match targets.

The transition dates are $s_1 = 0$ and $s_2 = 40$, with horizons $S = 140$ (exogenous prices) and $S = 200$ (endogenous prices).

```
S_exo = 140
S_endo = 200
S1, S2 = 0, 40
S3 = S2 + 2 * hh.T0

RR_exo, w_exo = tech.RR, tech.w
```

Two helper functions build the price and policy arrays that the transition solver expects.

Under the small open economy assumption, prices are constant over time.

```
def make_exo_price_seq(S, RR, w):
    """Construct constant price sequence for small open economy."""

    return jnp.column_stack([jnp.full(S + 2, RR), jnp.full(S + 2, w)])
```

The policy sequence sets τ_ℓ to the initial steady-state value before s_1 and to the terminal value after s_2, while holding all other fiscal instruments constant.

```
def make_policy_seq(S, ss0_τl, ss1_τl, S1, S2, τ_a, τ_0, benef, G, Gb):
    """Construct policy sequence with initial/terminal
    τ_l and constant other policies."""

    policy_seq = jnp.empty((S + 2, 6))
    policy_seq = policy_seq.at[:S1 + 1, 0].set(ss0_τl)
    policy_seq = policy_seq.at[S2 + 1:, 0].set(ss1_τl)
    policy_seq = policy_seq.at[:, 1].set(τ_a)
    policy_seq = policy_seq.at[:, 2].set(τ_0)
    policy_seq = policy_seq.at[:, 3].set(benef)
    policy_seq = policy_seq.at[:, 4].set(G)
    policy_seq = policy_seq.at[:, 5].set(Gb)
    return policy_seq
```

```
ss0 = find_ss_exo((RR, w), (τ_a_0, τ_0_0, benef_0, G_0, Gb_0), hh, tech)

print(f"Initial Steady State (s < 0):")
print(f"  Labor tax τ_l = {ss0.τ_l:.4f}")
print(f"  Interest rate r - δ = {ss0.r - tech.δ:.4f}")
print(f"  Capital/GDP = {ss0.k2gdp:.4f}")
print(f"  Debt/GDP = {ss0.debt2gdp:.4f}")
```

```
Initial Steady State (s < 0):
  Labor tax τ_l = 0.3383
  Interest rate r - δ = 0.0675
  Capital/GDP = 3.1615
  Debt/GDP = 0.5899
```

The equilibrium labor tax rate reproduces the calibration target $\tau_\ell = 0.3385$ reported in the table above, which is a useful check on the steady-state solver.

The calibration has a second property that matters for everything that follows.

The production parameters $\tilde{A}$ and $\tilde{\alpha}$ have been chosen so that, at the capital-labor ratio of this initial steady state, the Cobb-Douglas marginal products equal the exogenous prices (r, w) that the small open economy takes as given.

```
ε_bar_check = float(jnp.sum(hh.frac * hh.ε_arr))
K_eff_check = ss0.k_bar / hh.n + float(hh.frac[0] * hh.x0[0])
k_per_eff = K_eff_check / ε_bar_check

print(f"capital share implied by (r, w, K/L): "
      f"{tech.r * K_eff_check / (tech.r * K_eff_check + tech.w * ε_bar_check):.4f}"
      f"   (calibrated α̃ = {tech.α_tilde})")
print(f"A implied by the r equation:  "
      f"{tech.r / (tech.α_tilde * k_per_eff ** (tech.α_tilde - 1)):.4f}")
print(f"A implied by the w equation:  "
      f"{tech.w / ((1 - tech.α_tilde) * k_per_eff ** tech.α_tilde):.4f}"
      f"   (calibrated A = {tech.A})")
```

```
capital share implied by (r, w, K/L): 0.4002   (calibrated α̃ = 0.4)
A implied by the r equation:  2.2639
A implied by the w equation:  2.2620   (calibrated A = 2.2625)
```

The two price regimes therefore coincide at the initial steady state.

That is what makes our fixed-price and endogenous-price experiments comparable, and it is why we may reuse the household objects computed at fixed prices when we construct the initial steady state of the closed economy below.

Exercise 127.12.1 asks you to verify that these parameters are implied by the price targets rather than free.

The following figure traces how the equilibrium labor tax rate varies with government debt in the terminal steady state (no social security)

```
Gb_arr = np.linspace(0.5 * Gb_0, 1.5 * Gb_0, 20)
τl_arr = np.empty_like(Gb_arr)
debt2gdp_arr = np.empty_like(Gb_arr)

for i, Gb in enumerate(Gb_arr):
    ss = find_ss_exo((RR, w), (τ_a_0, τ_0_0, 0, G_0, Gb), hh, tech)
    τl_arr[i] = ss.τ_l
    debt2gdp_arr[i] = ss.debt2gdp

fig, ax = plt.subplots()
ax.plot(τl_arr, debt2gdp_arr)
ax.hlines(ss0.debt2gdp, τl_arr.min(),
          np.maximum(τl_arr.max(), ss0.τ_l), linestyle='--', color='r')
ax.scatter(ss0.τ_l, ss0.debt2gdp)
ax.text(ss0.τ_l * 0.95, ss0.debt2gdp * 0.95, "ss0")
ax.text(0.07, 0.4, r"ss1($G_b$)")
ax.set_xlabel(r'$\tau_\ell$')
ax.set_ylabel('debt/GDP')
plt.show()
```

Higher government debt requires larger interest payments, so the equilibrium labor tax rises.

The marked point shows the initial steady state.

To set the terminal steady state, we need to invert this relationship: given a target debt-to-GDP ratio, find the debt level $\bar{b}$ and the associated τ_ℓ.

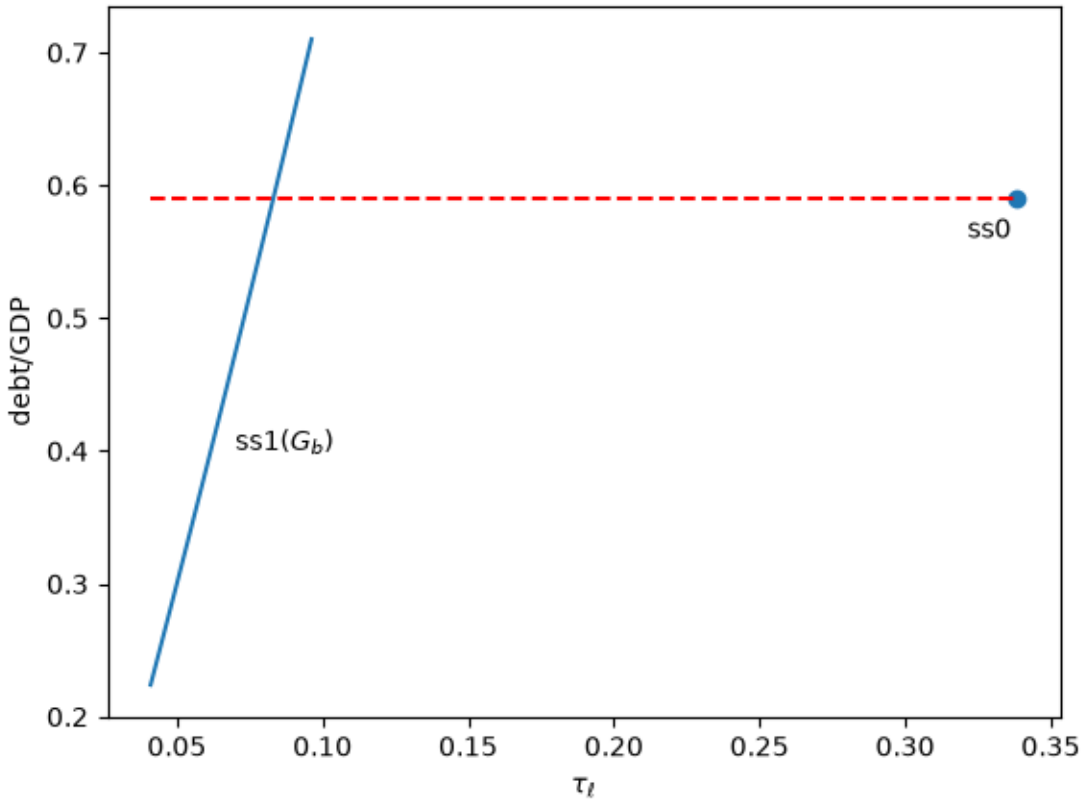


Fig. 127.4: Debt-to-GDP ratio as a function of the labor tax rate

```
def _compute_debt2gdp_for_Gb(
        Gb, price_arr, policy_no_Gb,
        α_arr, ε_arr, frac, n,
        β_arr, ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end, x0, Σ0,
        δ, Ind_work, unit_grid):
    """Compute debt-to-GDP ratio for a given Gb."""

    RR, w = price_arr
    τ_a, τ_0, benef, G = policy_no_Gb
    r = RR - 1 + δ

    policy_no_τl = jnp.array([τ_a, τ_0, benef, G, Gb])

    # Reuse _find_ss_tau_l instead of duplicating grid search
    τ_l = _find_ss_tau_l(
        price_arr, policy_no_τl, α_arr, ε_arr, frac, n,
        β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
        Ind_work, unit_grid
    )

    policy_arr = jnp.array([τ_l, τ_a, τ_0, benef, G, Gb])
    _, _, _, _, _, μx_arr, _, _, _ = ss_imbalance(
        price_arr, policy_arr,
        α_arr, ε_arr, frac, n,
        β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
        Ind_work
    )

    _, gdp = _compute_gdp(
        μx_arr[1:, 0], frac, ε_arr,
        Gb, r, w, n, x0[0], frac[0]
    )
    return Gb / gdp
```

_find_Gb_for_debt2gdp searches over $\bar{b}$ values via grid refinement to match a target debt-to-GDP ratio.

```
@jit
def _find_Gb_for_debt2gdp(
        debt2gdp_target, price_arr, policy_no_Gb,
        α_arr, ε_arr, frac, n,
        β_arr, ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end, x0, Σ0,
        δ, Ind_work, unit_grid, unit_grid_tau):
    """Find Gb consistent with a target debt-to-GDP ratio."""

    RR, w = price_arr
    τ_a, τ_0, benef, G = policy_no_Gb

    def compute_diff_coarse(Gb):
        debt2gdp = _compute_debt2gdp_for_Gb(
            Gb, price_arr, policy_no_Gb, α_arr, ε_arr, frac, n,
            β_arr, ρ_d, σ, B, C, R, Q, H, P_end, ξ_end, x0, Σ0,
            δ, Ind_work, unit_grid_tau
        )
        return debt2gdp - debt2gdp_target

    return _grid_refine(compute_diff_coarse, -40.0, 20.0, unit_grid, TOL_SS, 5)
```

`ss_target_debt2gdp_exo` finds the debt level consistent with the target ratio, then computes the full steady state.

```
def ss_target_debt2gdp_exo(debt2gdp_target, policy_target, price, hh, tech):
    """Find steady state with target debt-to-GDP ratio."""

    τ_a, τ_0, benef, G = policy_target
    RR, w = price

    price_arr = jnp.array([RR, w])
    policy_no_Gb = jnp.array([τ_a, τ_0, benef, G])

    Gb = _find_Gb_for_debt2gdp(
        float(debt2gdp_target), price_arr, policy_no_Gb,
        hh.α_arr, hh.ε_arr, hh.frac, hh.n,
        hh.β_arr, hh.ρ_d, hh.σ, hh.B, hh.C, hh.R, hh.Q, hh.H,
        hh.P_end, hh.ξ_end, hh.x0, hh.Σ0, tech.δ,
        hh.Ind_work, UNIT_GRID, UNIT_GRID
    )

    return find_ss_exo((RR, w), (τ_a, τ_0, benef, G, float(Gb)), hh, tech)
```

The terminal steady state eliminates social security ($\theta = 0$) while matching the initial debt-to-GDP ratio

```
ss1 = ss_target_debt2gdp_exo(
    ss0.debt2gdp, (τ_a_0, τ_0_0, 0, G_0), (RR_exo, w_exo), hh, tech
)

print(f"\nTerminal Steady State (s >= s3):")
print(f"  Labor tax τ_l = {ss1.τ_l:.4f}")
print(f"  Benefits θ = 0")
print(f"  Capital/GDP = {ss1.k2gdp:.4f}")
print(f"  Debt/GDP = {ss1.debt2gdp:.4f}")
```

```
Terminal Steady State (s >= s3):
  Labor tax τ_l = 0.0831
  Benefits θ = 0
  Capital/GDP = 4.1567
  Debt/GDP = 0.5901
```

127.7 Transition path computation

The transition path describes how the economy moves from the initial steady state (with social security) to the terminal steady state (after reform).

This is more complex than the steady-state computation because prices and policies change over time, so each cohort faces a unique lifetime sequence of tax and benefit rates.

`solve_backwards` solves the household problem backward in time during the transition, computing optimal decision rules $F_t(s)$ and closed-loop transition matrices $A_t^o(s)$ at each calendar date s and age t.

```
@jit
def solve_backwards(
        price_seq, policy_seq,
        P_arr_ss1, ξ_arr_ss1,
        ε_arr, β_arr, Ind_work,
```

(continues on next page)

(continued from previous page)

```
        ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end,
        s_indices, ages):
    """Backward Riccati scan over all dates and ages."""

    # Infer dimensions from input arrays
    n_x = P_end.shape[0]
    S = s_indices.shape[0] - 1

    def solve_all_ages(P_next_all, ξ_next_all, RR_s, w_s, τ_l, τ_a, τ_0, benef):
        def solve_one_age(t, P_next, ξ_next):
            ε_t = ε_arr[t]
            β_t = β_arr[t]
            Ind_work_t = Ind_work[t]
            F, Ao, P, ξ = solve_riccati_step(
                ε_t, β_t, Ind_work_t, RR_s, w_s, τ_l, τ_a, τ_0, benef,
                P_next, ξ_next, ρ_d, σ, B, C, R, Q, H
            )
            return F, Ao, P, ξ

        P_next_shifted = P_next_all[1:]
        ξ_next_shifted = ξ_next_all[1:]

        F_all, Ao_all, P_all, ξ_all = vmap(
            solve_one_age
        )(ages, P_next_shifted, ξ_next_shifted)
        return F_all, Ao_all, P_all, ξ_all

    def scan_body(carry, s_inv):
        P_next_seq, ξ_next_seq = carry
        s = S - s_inv

        RR_s = price_seq[s, 0]
        w_s = price_seq[s, 1]
        τ_l = policy_seq[s, 0]
        τ_a = policy_seq[s, 1]
        τ_0 = policy_seq[s, 2]
        benef = policy_seq[s, 3]

        F_s, Ao_s, P_s, ξ_s = solve_all_ages(
            P_next_seq, ξ_next_seq, RR_s, w_s, τ_l, τ_a, τ_0, benef
        )

        # Build P_curr and ξ_curr using the known shapes from input arrays
        P_curr = jnp.zeros_like(P_arr_ss1)
        P_curr = P_curr.at[:T0+1].set(P_s)
        P_curr = P_curr.at[-1].set(P_end)

        ξ_curr = jnp.zeros_like(ξ_arr_ss1)
        ξ_curr = ξ_curr.at[:T0+1].set(ξ_s)
        ξ_curr = ξ_curr.at[-1].set(ξ_end)

        output = (F_s, Ao_s, P_s, ξ_s)
        new_carry = (P_curr, ξ_curr)

        return new_carry, output
```

(continues on next page)

(continued from previous page)

```
    init_carry = (P_arr_ss1, ξ_arr_ss1)

    # s_indices already has the right length
    _, outputs = lax.scan(scan_body, init_carry, s_indices)

    F_seq, Ao_seq, P_seq_inner, ξ_seq_inner = outputs

    F_seq = jnp.flip(F_seq, axis=0)
    Ao_seq = jnp.flip(Ao_seq, axis=0)

    # Build output arrays using shapes from price_seq
    P_seq = jnp.zeros((price_seq.shape[0], P_arr_ss1.shape[0], n_x, n_x))
    ξ_seq = jnp.zeros((price_seq.shape[0], ξ_arr_ss1.shape[0]))

    P_seq_inner = jnp.flip(P_seq_inner, axis=0)
    ξ_seq_inner = jnp.flip(ξ_seq_inner, axis=0)
    P_seq = P_seq.at[:S+1, :T0+1].set(P_seq_inner)
    ξ_seq = ξ_seq.at[:S+1, :T0+1].set(ξ_seq_inner)

    P_seq = P_seq.at[:, -1].set(P_end)
    ξ_seq = ξ_seq.at[:, -1].set(ξ_end)
    P_seq = P_seq.at[-1, :].set(P_arr_ss1)
    ξ_seq = ξ_seq.at[-1, :].set(ξ_arr_ss1)

    return F_seq, Ao_seq, P_seq, ξ_seq
```

`simulate_forwards` takes the computed decision rules and simulates the economy forward from the initial distribution, tracking the evolution of asset means and variances across cohorts through the transition.

```
@jit
def simulate_forwards(
        Ao_seq, F_seq, μx_init, Σx_init,
        C, x0, Σ0, s_indices, ages):
    """Forward moment scan over all dates and ages."""

    # Infer dimensions from input arrays
    n_x = x0.shape[0]
    CCT = C @ C.T
    S = s_indices.shape[0] - 1

    def simulate_all_ages(μx_curr, Σx_curr, Ao_s, F_s):
        def simulate_one_age(t, μx_t, Σx_t, Ao_t, F_t):
            return forward_moment_step(
                μx_t, Σx_t, Ao_t, F_t, CCT
            )

        μx_next_all, μc_all, Σx_next_all, Vc_all = vmap(simulate_one_age)(
            ages, μx_curr[:T0+1], Σx_curr[:T0+1], Ao_s, F_s
        )
        return μx_next_all, μc_all, Σx_next_all, Vc_all

    def scan_body(carry, s):
        μx_curr, Σx_curr = carry

        Ao_s = Ao_seq[s]
        F_s = F_seq[s]
```

(continues on next page)

(continued from previous page)

```
        μx_next_inner, μc_s, Σx_next_inner, Vc_s = simulate_all_ages(
            μx_curr, Σx_curr, Ao_s, F_s
        )

        # Use shapes from μx_init
        μx_next = jnp.zeros_like(μx_init)
        μx_next = μx_next.at[0].set(x0)
        μx_next = μx_next.at[1:T0+2].set(μx_next_inner)

        Σx_next = jnp.zeros_like(Σx_init)
        Σx_next = Σx_next.at[0].set(Σ0)
        Σx_next = Σx_next.at[1:T0+2].set(Σx_next_inner)

        output = (μx_curr, μc_s, Σx_curr, Vc_s)
        new_carry = (μx_next, Σx_next)

        return new_carry, output

    init_carry = (μx_init, Σx_init)
    final_carry, outputs = lax.scan(scan_body, init_carry, s_indices)

    μx_seq_inner, μc_seq, Σx_seq_inner, Vc_seq = outputs

    # Build output arrays using inferred sizes
    μx_seq = jnp.zeros((S + 2, T0 + 2, n_x))
    Σx_seq = jnp.zeros((S + 2, T0 + 2, n_x, n_x))

    μx_seq = μx_seq.at[:S+1].set(μx_seq_inner)
    Σx_seq = Σx_seq.at[:S+1].set(Σx_seq_inner)

    μx_seq = μx_seq.at[S+1].set(final_carry[0])
    Σx_seq = Σx_seq.at[S+1].set(final_carry[1])

    return μx_seq, μc_seq, Σx_seq, Vc_seq
```

Given a candidate transition tax rate τ_ℓ^{trans}, the function `transition_paths` constructs the complete policy sequence, solves backward, simulates forward, computes the capital and debt paths by aggregating across cohorts, and returns the terminal debt carryover that we seek to drive to zero.

```
@jit
def _transition_paths(
        τ_l_trans, price_seq, policy_seq,
        ss1_P_arr, ss1_ξ_arr, ss1_Gb,
        μx_init, Σx_init, k_bar_init,
        s_indices, age_range, S1, S2,
        ε_arr, β_arr, Ind_work,
        ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end, x0, Σ0,
        frac, n):
    """Solve backward and simulate forward for a given τ_l_trans."""

    # Infer dimensions from input arrays
    n_x = x0.shape[0]
    S = s_indices.shape[0] - 2

    # Derive variants via slicing
```

(continues on next page)

(continued from previous page)

```
    s_indices_scan = s_indices[:-1]      # arange(S+1)
    ages = age_range[:-1]                 # arange(T0+1)
    capital_indices = s_indices[1:-1]    # arange(1, S+1)

    # Update policy sequence with transition tax using dynamic indexing
    mask = (s_indices >= S1 + 1) & (s_indices <= S2)
    τ_l_col = jnp.where(mask, τ_l_trans, policy_seq[:, 0])
    policy_seq = policy_seq.at[:, 0].set(τ_l_col)

    # Solve backwards
    F_seq, Ao_seq, P_seq, ξ_seq = solve_backwards(
        price_seq, policy_seq, ss1_P_arr, ss1_ξ_arr,
        ε_arr, β_arr, Ind_work,
        ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end,
        s_indices_scan, ages
    )

    # Simulate forwards
    μx_seq, μc_seq, Σx_seq, Vc_seq = simulate_forwards(
        Ao_seq, F_seq, μx_init, Σx_init, C, x0, Σ0, s_indices_scan, ages
    )

    # Compute capital path
    ε_agg = jnp.sum(ε_arr * frac)
    frac0_x0 = frac[0] * x0[0]

    def capital_step(k_prev, s):
        RR = price_seq[s, 0]
        w = price_seq[s, 1]
        G = policy_seq[s, 4]
        c_agg = jnp.sum(μc_seq[s] * frac)
        k_new = RR * (frac0_x0 + k_prev / n) - G - c_agg + w * ε_agg
        return k_new, k_new

    # capital_indices is pre-created arange(1, S+1)
    _, k_path = lax.scan(capital_step, k_bar_init, capital_indices)
    k_seq = jnp.concatenate([jnp.array([k_bar_init]), k_path])

    # Compute debt path
    a_seq = jnp.sum(μx_seq[1:, 1:, 0] * frac, axis=1)
    Gb_seq = a_seq - k_seq

    carryover = Gb_seq[-1] - ss1_Gb

    return carryover, μx_seq, μc_seq, k_seq, Gb_seq, F_seq, Ao_seq
```

`transition_paths` unpacks the steady-state and household objects and calls the JIT-compiled inner function.

```
def transition_paths(
        τ_l_trans, price_seq, policy_seq,
        ss0, ss1, hh, tech,
        S, S1, S2, μx_init, Σx_init):
    """Compute transition path."""

    policy_seq = jnp.asarray(policy_seq)
    price_seq = jnp.asarray(price_seq)
```

(continues on next page)

(continued from previous page)

```
    # Pre-create iteration arrays (use slicing for variants)
    s_indices = jnp.arange(S + 2)

    carryover, μx_seq, μc_seq, k_seq, Gb_seq, F_seq, Ao_seq = _transition_paths(
        float(τ_l_trans), price_seq, policy_seq,
        ss1.P_arr, ss1.ξ_arr, float(ss1.Gb),
        μx_init, Σx_init, float(ss0.k_bar),
        s_indices, AGE_INDICES,
        S1, S2,
        hh.ε_arr, hh.β_arr, hh.Ind_work,
        hh.ρ_d, hh.σ, hh.B, hh.C,
        hh.R, hh.Q, hh.H,
        hh.P_end, hh.ξ_end, hh.x0, hh.Σ0,
        hh.frac, hh.n
    )

    return (float(carryover), μx_seq, μc_seq,
            k_seq, Gb_seq, F_seq, Ao_seq)
```

127.7.1 Shooting method

To find the correct transition tax rate, we use a shooting method.

If the tax rate is too low, debt explodes; if it is too high, debt falls below the target.

The equilibrium tax rate is found where the terminal debt exactly hits the target.

We start by computing two transition paths with different trial tax rates to illustrate the shooting method

```
price_seq = make_exo_price_seq(S_exo, RR, w)
policy_seq_base = make_policy_seq(S_exo, ss0.τ_l, ss1.τ_l, S1, S2,
                                  τ_a_0, τ_0_0, 0, G_0, Gb_0)

τ_l_low = 0.14
τ_l_high = 0.17

_, μx_seq1, μc_seq1, k_seq1, Gb_seq1, _, _ = transition_paths(
    τ_l_low, price_seq, policy_seq_base, ss0, ss1, hh, tech,
    S_exo, S1, S2, ss0.μx_arr, ss0.Σx_arr)

_, μx_seq2, μc_seq2, k_seq2, Gb_seq2, _, _ = transition_paths(
    τ_l_high, price_seq, policy_seq_base, ss0, ss1, hh, tech,
    S_exo, S1, S2, ss0.μx_arr, ss0.Σx_arr)
```

We can plot the resulting debt paths to see how they differ under the two trial tax rates

```
fig, ax = plt.subplots(figsize=(10, 5))
ax.plot(Gb_seq1, 'b-', linewidth=2,
        label=f'$\\tau_\\ell$ = {τ_l_low:.2f} (too low)')
ax.plot(Gb_seq2, 'r-', linewidth=2,
        label=f'$\\tau_\\ell$ = {τ_l_high:.2f} (too high)')
ax.axhline(ss1.Gb, color='k', linestyle='--',
           label=f'Target $G_b$ = {ss1.Gb:.2f}')
ax.axvspan(S1, S2, alpha=0.1, color='yellow', label='Transition period')
ax.set_xlabel('time')
```

(continues on next page)

(continued from previous page)

```
ax.set_ylabel('government debt')
ax.legend()
plt.show()
```

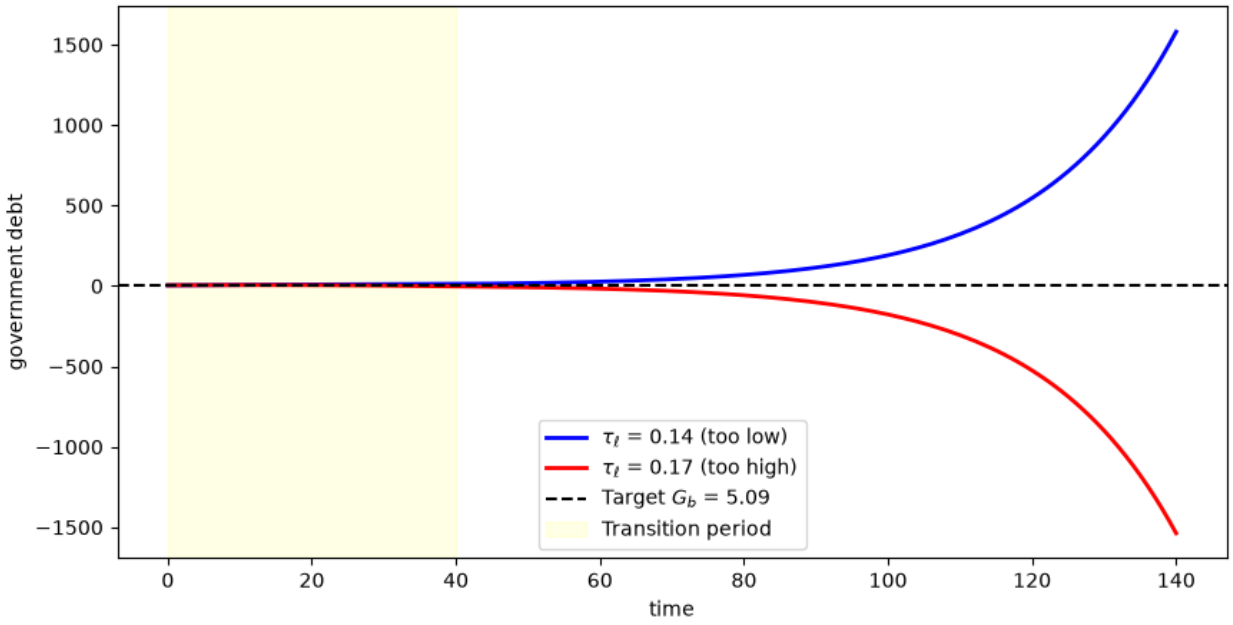


Fig. 127.5: Shooting method for finding the transition tax rate

The blue curve shows government debt increasing over time when the tax rate is too low, while the red curve shows debt falling below the target when the tax rate is too high.

The equilibrium transition tax rate lies between these two extremes, at the value where terminal debt exactly matches the target G_b.

Bisection automates this shooting procedure.

For Experiment 1, each cohort alive at s_1 must also be compensated by the present value of the social security benefits it would have received under the original system.

`_compute_compensation` evaluates this present value for a single cohort identified by its remaining lifetime.

```
def _compute_compensation(
        death_time, τ_l_seq, τ_a_seq,
        benef_diff, RR_seq, w_seq,
        ε_arr, indices, ss0_τ_l):
    """Present value of lost benefits for one cohort."""

    n_periods = death_time + 1
    age_at_0 = T0 - death_time

    time_mask = indices < n_periods

    τ_l_cohort = jnp.where(time_mask, τ_l_seq[:T0 + 1], 0.0)
    τ_a_cohort = jnp.where(time_mask, τ_a_seq[:T0 + 1], 0.0)
    RR_cohort = jnp.where(time_mask, RR_seq[:T0 + 1], 1.0)
```

(continues on next page)

(continued from previous page)

```
    w_cohort = jnp.where(time_mask, w_seq[:T0 + 1], 0.0)

    age_idx = jnp.clip(age_at_0 + indices, 0, T0)
    ε_cohort = jnp.where(
        time_mask, ε_arr[age_idx], 0.0
    )
    benef_cohort = jnp.where(
        time_mask, benef_diff[age_idx], 0.0
    )
    benef_cohort = jnp.where(ε_cohort != 0, 0.0, benef_cohort)

    RR_tilde_seq = RR_cohort - τ_a_cohort * (RR_cohort - 1)
    RR_tilde_seq = jnp.where(time_mask, RR_tilde_seq, 1.0)
    discount_factors = jnp.cumprod(RR_tilde_seq)

    labor_loss = w_cohort * ε_cohort * (τ_l_cohort - ss0_τ_l)

    pv_seq = jnp.where(
        time_mask,
        (benef_cohort + labor_loss) / discount_factors,
        0.0
    )

    valid = (death_time >= 0) & (death_time < T0)
    return jnp.where(valid, jnp.sum(pv_seq), 0.0)
```

`apply_compensation` vectorizes this calculation across all cohorts with `vmap` and adds the result to each cohort's initial asset holdings.

```
@jit
def apply_compensation(
        μx_arr_ss0, Σx_arr_ss0,
        τ_l_seq, τ_a_seq, benef_diff,
        RR_seq, w_seq, ε_arr, ss0_τ_l,
        ages_full, ages, x0, Σ0,
        comp_mult):
    """Vectorize compensation across cohorts and adjust initial assets."""

    def compute_comp_for_age(age):
        death_time = T0 - age
        comp = _compute_compensation(
            death_time,
            τ_l_seq, τ_a_seq, benef_diff,
            RR_seq, w_seq, ε_arr,
            ages, ss0_τ_l
        )
        valid = (age >= 1) & (age <= T0)
        return jnp.where(valid, comp, 0.0)

    compensations = vmap(compute_comp_for_age)(ages_full)

    μx_init = jnp.zeros_like(μx_arr_ss0)
    Σx_init = jnp.zeros_like(Σx_arr_ss0)

    μx_init = μx_init.at[0].set(x0)
    Σx_init = Σx_init.at[0].set(Σ0)
    μx_init = μx_init.at[-1].set(x0)
```

(continues on next page)

(continued from previous page)

```
    Σx_init = Σx_init.at[-1].set(Σ0)

    μx_init = μx_init.at[1:-1].set(μx_arr_ss0[1:-1])
    Σx_init = Σx_init.at[1:-1].set(Σx_arr_ss0[1:-1])

    # comp_mult: 0.0 = no compensation, 1.0 = full
    μx_init = μx_init.at[:, 0].add(comp_mult * compensations)

    return μx_init, Σx_init
```

`_transition_carryover` applies compensation, solves the transition, and returns the terminal debt carryover – the scalar the bisection drives to zero.

```
def _transition_carryover(
        τ_l_trans, price_seq, policy_seq,
        ss1_P_arr, ss1_ξ_arr, ss1_Gb,
        ss0_μx_arr, ss0_Σx_arr, k_bar_init,
        benef_diff, ss0_τ_l, comp_mult,
        s_indices, age_range, S1, S2,
        ε_arr, β_arr, Ind_work,
        ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end, x0, Σ0,
        frac, n):
    """Terminal debt carryover for a given transition τ_l."""

    ages = age_range[:-1]  # arange(T0+1)

    # Update policy sequence with transition tax
    mask = (s_indices >= S1 + 1) & (s_indices <= S2)
    τ_l_col = jnp.where(mask, τ_l_trans, policy_seq[:, 0])
    policy_seq_updated = policy_seq.at[:, 0].set(τ_l_col)

    # Compute initial conditions (compensation zeroed when comp_mult=0.0)
    μx_init, Σx_init = apply_compensation(
        ss0_μx_arr, ss0_Σx_arr,
        policy_seq_updated[:, 0], policy_seq_updated[:, 1], benef_diff,
        price_seq[:, 0], price_seq[:, 1],
        ε_arr, ss0_τ_l,
        age_range, ages, x0, Σ0,
        comp_mult
    )

    carryover, *_ = _transition_paths(
        τ_l_trans, price_seq, policy_seq,
        ss1_P_arr, ss1_ξ_arr, ss1_Gb,
        μx_init, Σx_init, k_bar_init,
        s_indices, age_range, S1, S2,
        ε_arr, β_arr, Ind_work,
        ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end, x0, Σ0,
        frac, n
    )
    return carryover
```

We implement the bisection search in `_find_transition_tau_l`, which repeatedly evaluates `_transition_carryover` at the midpoint of a shrinking interval until the carryover is driven to zero.

```
@jit
def _find_transition_tau_l(
        price_seq, policy_seq, bounds,
        ss1_P_arr, ss1_ξ_arr, ss1_Gb,
        ss0_μx_arr, ss0_Σx_arr, k_bar_init,
        benef_diff, ss0_τ_l, comp_mult,
        s_indices, age_range, S1, S2,
        ε_arr, β_arr, Ind_work,
        ρ_d, σ, B, C, R, Q, H,
        P_end, ξ_end, x0, Σ0,
        frac, n):
    """Find transition τ_l using bisection.

    comp_mult controls compensation (0 or 1).
    """

    a, b = bounds[0], bounds[1]

    def compute_carryover(τ_l_trans):
        return _transition_carryover(
            τ_l_trans, price_seq, policy_seq,
            ss1_P_arr, ss1_ξ_arr, ss1_Gb,
            ss0_μx_arr, ss0_Σx_arr, k_bar_init,
            benef_diff, ss0_τ_l, comp_mult,
            s_indices, age_range,
            S1, S2,
            ε_arr, β_arr, Ind_work,
            ρ_d, σ, B, C, R, Q, H,
            P_end, ξ_end, x0, Σ0,
            frac, n
        )

    def cond_fn(state):
        a, b, fa, fb, i = state
        return (jnp.abs(b - a) > 1e-10) & (i < 100)

    def body_fn(state):
        a, b, fa, fb, i = state
        c = (a + b) / 2.0
        fc = compute_carryover(c)
        a_new = jnp.where(fa * fc > 0, c, a)
        b_new = jnp.where(fa * fc > 0, b, c)
        fa_new = jnp.where(fa * fc > 0, fc, fa)
        fb_new = jnp.where(fa * fc > 0, fb, fc)
        return (a_new, b_new, fa_new, fb_new, i + 1)

    fa, fb = compute_carryover(a), compute_carryover(b)
    init_state = (a, b, fa, fb, 0)
    final_state = lax.while_loop(cond_fn, body_fn, init_state)
    a_final, b_final, _, _, _ = final_state

    return (a_final + b_final) / 2.0
```

The top-level wrapper `find_transition_exo` sets up the compensation parameters and calls the bisection solver, then recomputes the full transition path at the equilibrium tax rate.

```
def find_transition_exo(price_seq, policy_seq_base, ss0, ss1,
```

(continues on next page)

(continued from previous page)

```
                      hh, tech, S, S1, S2,
                      compensation_data=None,
                      τl_bounds=(0.01, 0.6)):
    """Find transition tax rate under exogenous prices."""

    policy_seq = jnp.asarray(policy_seq_base)
    price_seq = jnp.asarray(price_seq)
    bounds = jnp.array([τl_bounds[0], τl_bounds[1]])
    s_indices = jnp.arange(S + 2)

    # Set up compensation parameters (default zeros when not using compensation)
    if compensation_data is not None:
        benef_diff, ss0_τ_l = compensation_data
        comp_mult = 1.0
    else:
        benef_diff = jnp.zeros(hh.T0 + 1)
        ss0_τ_l = ss0.τ_l
        comp_mult = 0.0

    # Find transition tax using unified function
    τ_l_trans = _find_transition_tau_l(
        price_seq, policy_seq, bounds,
        ss1.P_arr, ss1.ξ_arr, float(ss1.Gb),
        ss0.μx_arr, ss0.Σx_arr, float(ss0.k_bar),
        benef_diff, float(ss0_τ_l), comp_mult,
        s_indices, AGE_INDICES,
        S1, S2,
        hh.ε_arr, hh.β_arr, hh.Ind_work,
        hh.ρ_d, hh.σ, hh.B, hh.C,
        hh.R, hh.Q, hh.H,
        hh.P_end, hh.ξ_end, hh.x0, hh.Σ0,
        hh.frac, hh.n
    )
    τ_l_trans = float(τ_l_trans)

    # Compute final results with initial conditions
    mask = (s_indices >= S1 + 1) & (s_indices <= S2)
    τ_l_col = jnp.where(
        mask, τ_l_trans, policy_seq[:, 0]
    )
    policy_seq_final = policy_seq.at[:, 0].set(
        τ_l_col
    )

    μx_init, Σx_init = apply_compensation(
        ss0.μx_arr, ss0.Σx_arr,
        policy_seq_final[:, 0],
        policy_seq_final[:, 1],
        benef_diff,
        price_seq[:, 0], price_seq[:, 1],
        hh.ε_arr, float(ss0_τ_l),
        AGE_INDICES, AGE_INDICES[:-1],
        hh.x0, hh.Σ0, comp_mult
    )

    results = transition_paths(
        τ_l_trans, price_seq, policy_seq,
```

(continues on next page)

(continued from previous page)

```
        ss0, ss1, hh, tech, S, S1, S2,
        μx_init, Σx_init
    )

    return τ_l_trans, results
```

Bisection over the transition tax rate produces the equilibrium path.

```
τ_l_trans, results = find_transition_exo(
    price_seq, policy_seq_base, ss0, ss1,
    hh, tech, S_exo, S1, S2)

carryover, μx_seq, μc_seq, k_seq, Gb_seq, F_seq, Ao_seq = results
```

In the baseline case (no compensation), social security benefits are simply terminated.

```
τ_l_seq = np.zeros(S_exo + 1)
τ_l_seq[:S1 + 1] = ss0.τ_l
τ_l_seq[S1 + 1:S2 + 1] = τ_l_trans
τ_l_seq[S2 + 1:] = ss1.τ_l

fig, axes = plt.subplots(1, 2, figsize=(12, 4))

axes[0].plot(τ_l_seq, 'b-', linewidth=2)
axes[0].axhline(ss0.τ_l, color='r', linestyle=':',
                label=f'Initial $\\tau_\\ell$ = {ss0.τ_l:.4f}')
axes[0].axhline(ss1.τ_l, color='g', linestyle=':',
                label=f'Terminal $\\tau_\\ell$ = {ss1.τ_l:.4f}')
axes[0].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[0].set_xlabel('time')
axes[0].set_ylabel('labor tax rate')
axes[0].set_title('Labor tax rate path')
axes[0].legend()


axes[1].plot(Gb_seq, 'b-', linewidth=2)
axes[1].axhline(ss0.Gb, color='r', linestyle=':',
                label=f'Initial $G_b$ = {ss0.Gb:.2f}')
axes[1].axhline(ss1.Gb, color='g', linestyle=':',
                label=f'Terminal $G_b$ = {ss1.Gb:.2f}')
axes[1].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[1].set_xlabel('time')
axes[1].set_ylabel('government debt')
axes[1].set_title('Government debt path')
axes[1].legend()


plt.tight_layout()
plt.show()
```

The left panel shows the labor tax rate dropping during the transition because the government no longer needs to fund social security benefits.

The right panel shows government debt converging from its initial level to the new steady-state level, with the shaded region marking the transition period $[s_1, s_2)$.

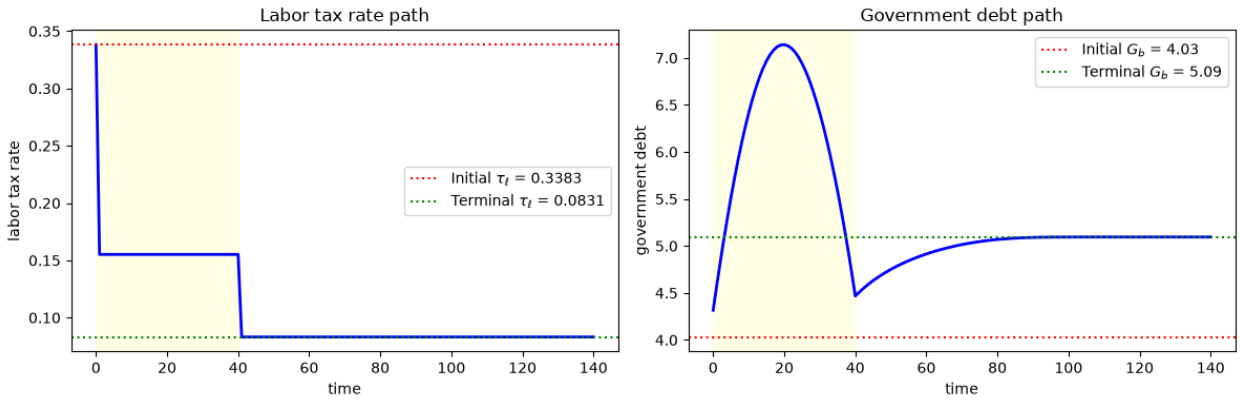


Fig. 127.6: Baseline transition path (no compensation)

127.8 Experiment 1: compensation through debt

In this experiment, the government terminates social security benefits but compensates affected generations.

Each cohort receives a transfer equal to the present value of the benefits they would have received – an actuarially fair buy-out.

Algorithm 127.8.1 (Fixed factor prices – Experiment 1 (buy-out))

Here $s_3 = s_2 + T_0 = 105$.

Step 1. Set up parameters.

Step 2. Solve the initial stationary equilibrium with constant social security benefit S: fix $\tau_a, \tau_0, S, G, \bar{b}$ and solve for $\tau_\ell = \tau_{\ell,0}$ such that the government budget balances:

```
F(τ_ℓ) = government budget imbalance
Find root of F(τ_ℓ) = 0.
```

Step 3. Solve the terminal stationary equilibrium with no social security: search over $\bar{b}$ so that the debt-to-GDP ratio matches a target:

```
H(Gb) = debt-to-GDP given Gb
       (internally solves F(τ_ℓ; Gb) = 0)
Find root of H(Gb) = target.
```

The associated labor tax is $\tau_{\ell,2}$.

Step 4. Solve the transition path: at $s = 0 = s_1$, all cohorts alive lose benefits and a cohort of age t receives a one-time compensation equal to the present value of lost benefits, discounted at the after-tax return $\tilde{R}(s) = R(s)[1-\tau_a(s)] + \tau_a(s)$:

$$\text{comp}_t = \sum_{j=0}^{T_0-t} \left[S_{t+j} + w\varepsilon_{t+j}\left(\tau_\ell(s+j) - \tau_{\ell,0}\right)\right] \prod_{i=0}^{j} \tilde{R}(s+i)^{-1},$$

where S_{t+j} is the benefit that an age-$(t+j)$ person would have received under the old regime, and zero while that person is still working.

The second term inside the brackets compensates a cohort for the change in labor taxes that it faces over the rest of its life, with $\tau_\ell(s+j)$ equal to $\tau_{\ell,1}$ during $[s_1, s_2)$ and $\tau_{\ell,2}$ thereafter, so that the buy-out leaves the present value of

each cohort's lifetime resources unchanged.

The government sets $\tau_{\ell,1}$ during $[s_1, s_2)$ and $\tau_{\ell,2}$ from s_2 onwards, with a one-time expenditure increase of $\sum f_t \, \text{comp}_t$ at s_1.

Find $\tau_{\ell,1}$ such that terminal government debt matches the target:

```
J(τ_ℓ) = terminal debt carryover
Find root of J(τ_ℓ) = Gb_terminal.
```

127.8.1 Fixed prices

We hold factor prices fixed and construct the price and policy sequences for the transition.

The benefit difference vector `benef_diff_exp1` records the per-period benefit loss for each age: retirees lose their old-regime benefits, while workers are unaffected.

```
ss1_exp1_exo = ss1

price_seq_exp1_exo = make_exo_price_seq(S_exo, RR_exo, w_exo)
policy_seq_exp1_exo = make_policy_seq(S_exo, ss0.τ_l, ss1.τ_l, S1, S2,
                                      τ_a_0, τ_0_0, 0, G_0, Gb_0)

benef_diff_exp1 = jnp.zeros(hh.T0 + 1)
benef_diff_exp1 = benef_diff_exp1.at[hh.T1 + 1:].set(ss0.benef)
```

Passing `compensation_data` to `find_transition_exo` turns the buy-out on: that function sets `comp_mult` to one and passes it to `apply_compensation`, which scales each cohort's compensation by it.

We solve for the transition tax rate with and without the buy-out, so that we can compare the two paths.

```
# Solve with buyout
τ_l_exp1_exo_bo, results_exp1_exo_bo = find_transition_exo(
    price_seq_exp1_exo, policy_seq_exp1_exo, ss0, ss1_exp1_exo,
    hh, tech, S_exo, S1, S2,
    compensation_data=(benef_diff_exp1, ss0.τ_l)
)

# Solve without buyout (for comparison)
τ_l_exp1_exo_nb, results_exp1_exo_nb = find_transition_exo(
    price_seq_exp1_exo, policy_seq_exp1_exo, ss0, ss1_exp1_exo,
    hh, tech, S_exo, S1, S2
)
```

We compare the transition paths with and without buy-out compensation.

```
exp1_exo = {
    'ss0': ss0, 'ss1': ss1_exp1_exo,
    'τ_l_buyout': τ_l_exp1_exo_bo, 'τ_l_no_buyout': τ_l_exp1_exo_nb,
    'results_buyout': results_exp1_exo_bo,
    'results_no_buyout': results_exp1_exo_nb,
    'hh': hh, 'tech': tech
}
```

The following figure shows how the buy-out reshapes initial asset holdings across cohorts.

```
# Extract results
_, μx_seq_bo, μc_seq_bo, k_seq_bo, Gb_seq_bo, _, _ = exp1_exo['results_buyout']
results_nb = exp1_exo['results_no_buyout']
_, μx_seq_nb, μc_seq_nb, k_seq_nb, Gb_seq_nb, _, _ = results_nb

# Mean assets by age at time s=0 (with vs without buyout)
μa_bo = μx_seq_bo[0, 1:, 0]  # Assets at s=0 with buyout
μa_nb = μx_seq_nb[0, 1:, 0]  # Assets at s=0 without buyout

# Compensation = difference in initial assets
compensation_by_age = μa_bo - μa_nb

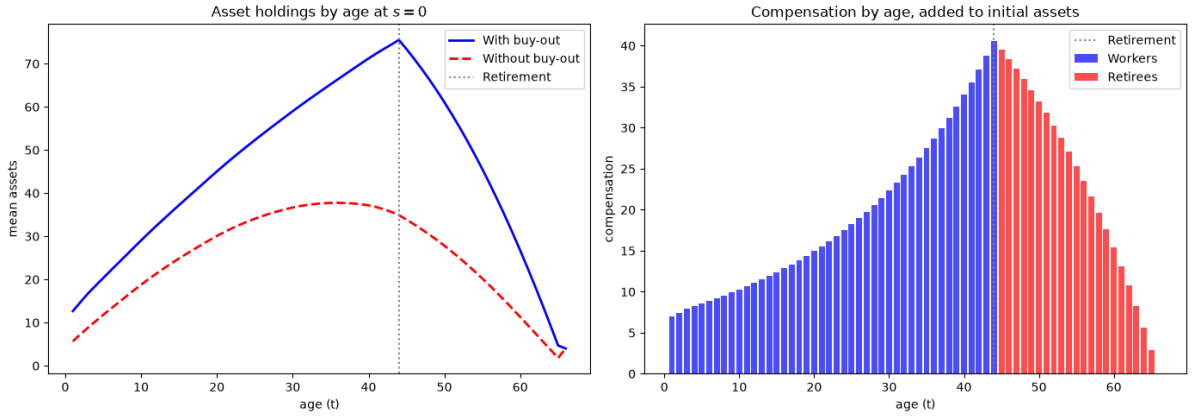
fig, axes = plt.subplots(1, 2, figsize=(14, 5))

# Asset profiles
ages = np.arange(1, hh.T0 + 2)
axes[0].plot(ages, μa_bo, 'b-', linewidth=2, label='With buy-out')
axes[0].plot(ages, μa_nb, 'r--', linewidth=2, label='Without buy-out')
axes[0].axvline(hh.T1 + 1, color='gray', linestyle=':', label='Retirement')
axes[0].set_xlabel('age (t)')
axes[0].set_ylabel('mean assets')
axes[0].set_title('Asset holdings by age at $s=0$')
axes[0].legend()


# Compensation histogram
working_ages = ages[ages <= hh.T1 + 1]
retired_ages = ages[ages > hh.T1 + 1]
comp_working = compensation_by_age[:hh.T1 + 1]
comp_retired = compensation_by_age[hh.T1 + 1:]

axes[1].bar(working_ages, comp_working,
            color='blue', alpha=0.7, label='Workers')
axes[1].bar(retired_ages, comp_retired,
            color='red', alpha=0.7, label='Retirees')
axes[1].axhline(0, color='k', linewidth=0.5)
axes[1].axvline(hh.T1 + 1, color='gray', linestyle=':', label='Retirement')
axes[1].set_xlabel('age (t)')
axes[1].set_ylabel('compensation')
axes[1].set_title('Compensation by age, added to initial assets')
axes[1].legend()


plt.tight_layout()
plt.show()
```



Retirees receive the largest compensation because they were expecting benefits for the remainder of their lives.

Older workers receive significant compensation, while young workers receive little because they have their entire working lives to adjust.

The declining profile among retirees reflects the actuarial calculation: older retirees have fewer remaining years of expected benefits.

We now plot the aggregate transition paths for the labor tax, government debt, capital, and consumption under both schemes.

```
# hh, tech, ss0, ss1 already in scope -- just alias from dict for readability
ss0_exp1 = exp1_exo['ss0']
ss1_exp1 = exp1_exo['ss1']

# Construct τ_l sequences
τ_l_seq_bo = np.zeros(S_exo + 1)
τ_l_seq_bo[:S1 + 1] = ss0_exp1.τ_l
τ_l_seq_bo[S1 + 1:S2 + 1] = exp1_exo['τ_l_buyout']
τ_l_seq_bo[S2 + 1:] = ss1_exp1.τ_l

τ_l_seq_nb = np.zeros(S_exo + 1)
τ_l_seq_nb[:S1 + 1] = ss0_exp1.τ_l
τ_l_seq_nb[S1 + 1:S2 + 1] = exp1_exo['τ_l_no_buyout']
τ_l_seq_nb[S2 + 1:] = ss1_exp1.τ_l

fig, axes = plt.subplots(2, 2, figsize=(14, 10))

# τ_l comparison
axes[0, 0].plot(τ_l_seq_bo, 'b-', linewidth=2, label='With buy-out')
axes[0, 0].plot(τ_l_seq_nb, 'r--', linewidth=2, label='Without buy-out')
axes[0, 0].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[0, 0].set_xlabel('time (s)')
axes[0, 0].set_ylabel(r'$\tau_\ell$')
axes[0, 0].set_title('Labor tax rate path')
axes[0, 0].legend()

# Gb comparison
axes[0, 1].plot(Gb_seq_bo, 'b-', linewidth=2, label='With buy-out')
axes[0, 1].plot(Gb_seq_nb, 'r--', linewidth=2, label='Without buy-out')
axes[0, 1].axhline(ss1_exp1.Gb, color='k', linestyle=':', alpha=0.7)
axes[0, 1].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[0, 1].set_xlabel('time (s)')
axes[0, 1].set_ylabel(r'$G_b$')
```

(continues on next page)

(continued from previous page)

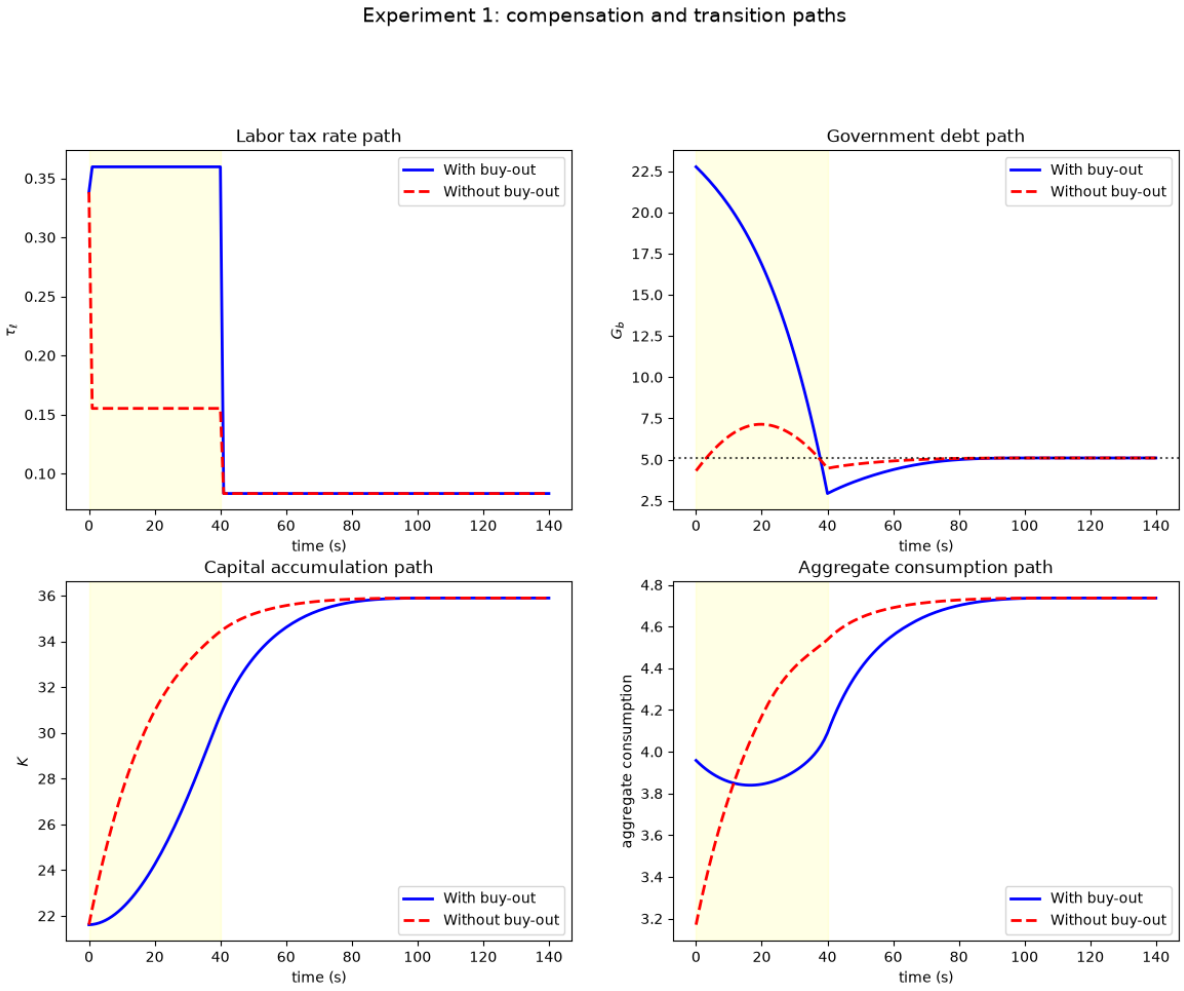
```
axes[0, 1].set_title('Government debt path')
axes[0, 1].legend()

# Capital path
axes[1, 0].plot(k_seq_bo, 'b-', linewidth=2, label='With buy-out')
axes[1, 0].plot(k_seq_nb, 'r--', linewidth=2, label='Without buy-out')
axes[1, 0].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[1, 0].set_xlabel('time (s)')
axes[1, 0].set_ylabel(r'$K$')
axes[1, 0].set_title('Capital accumulation path')
axes[1, 0].legend()

# Aggregate consumption
c_agg_bo = np.array(μc_seq_bo[:S_exo + 1]) @ np.array(hh.frac)
c_agg_nb = np.array(μc_seq_nb[:S_exo + 1]) @ np.array(hh.frac)

axes[1, 1].plot(c_agg_bo, 'b-', linewidth=2, label='With buy-out')
axes[1, 1].plot(c_agg_nb, 'r--', linewidth=2, label='Without buy-out')
axes[1, 1].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[1, 1].set_xlabel('time (s)')
axes[1, 1].set_ylabel('aggregate consumption')
axes[1, 1].set_title('Aggregate consumption path')
axes[1, 1].legend()

plt.suptitle(
    'Experiment 1: compensation and transition paths',
    fontsize=14, y=1.02
)
plt.show()
```



The buy-out scheme leads to a slower initial rise in private capital because the government must make large transfers.

Both schemes converge to the same terminal steady state.

We now examine consumption paths for cohorts at different ages when the reform occurs.

```
fig, axes = plt.subplots(2, 2, figsize=(14, 10))

selected_ages = [0, 20, 40, 60]  # Cohorts at different ages at s=0

for idx, age_at_0 in enumerate(selected_ages):
    ax = axes[idx // 2, idx % 2]

    remaining_life = hh.T0 - age_at_0
    max_time = min(remaining_life + 1, S_exo + 1)

    c_bo = [μc_seq_bo[s, age_at_0 + s]
            for s in range(max_time)
            if age_at_0 + s <= hh.T0]
    c_nb = [μc_seq_nb[s, age_at_0 + s]
            for s in range(max_time)
            if age_at_0 + s <= hh.T0]

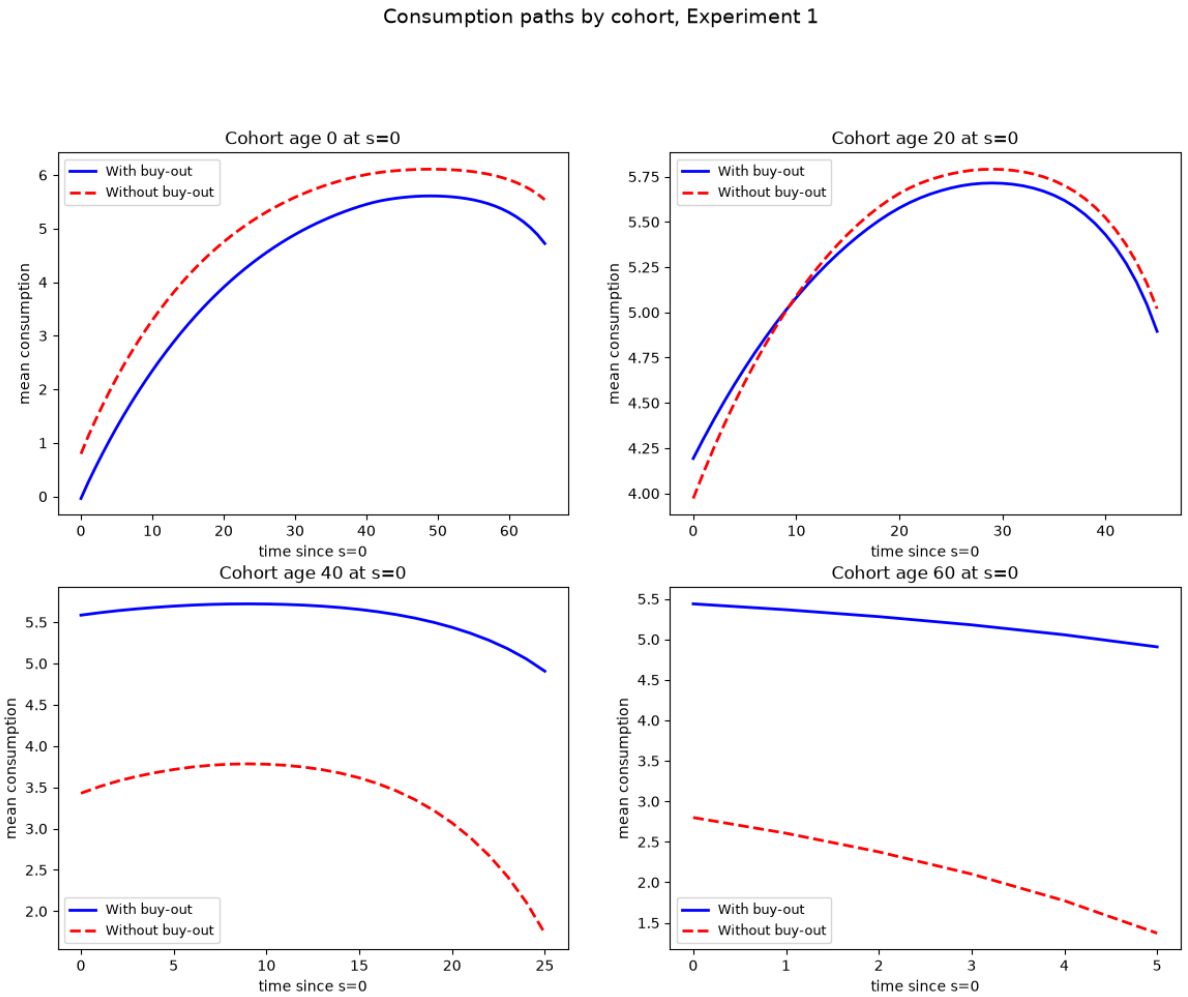
    ax.plot(c_bo, 'b-', linewidth=2, label='With buy-out')
```

(continues on next page)

(continued from previous page)

```
    ax.plot(c_nb, 'r--', linewidth=2, label='Without buy-out')
    ax.set_xlabel('time since s=0')
    ax.set_ylabel('mean consumption')
    ax.set_title(f'Cohort age {age_at_0} at s=0')
    ax.legend(fontsize=9)


plt.suptitle(
    'Consumption paths by cohort, Experiment 1',
    fontsize=14, y=1.02
)
plt.show()
```



Young workers (age 0) show nearly identical consumption paths because they receive little compensation and have decades to adjust their savings.

Near-retirement workers (age 40) exhibit more noticeable differences as the buyout compensation partially offsets their lost benefits.

Retirees (age 60) show the most dramatic difference: without compensation, their consumption drops sharply when benefits end, whereas the buyout scheme maintains higher consumption by replacing the lost income.

127.8.2 Endogenous prices

With endogenous factor prices, changes in saving behavior affect capital accumulation, which alters marginal products and feeds back into household decisions.

Algorithm 127.8.2 (Endogenous factor prices)

Here $s_3 = s_2 + 2T_0 = 170$.

Steps 1–3 are the same as *Algorithm 127.8.1*, with an outer fixed-point loop over factor prices to clear factor markets in steady state.

Step 4. The factor price sequences are now endogenous.

Wrap the fixed-price transition solver in a relaxation loop:

```
T(R_seq):
  1. Compute the wage sequence from R_seq via Cobb-Douglas.
  2. Taking prices as given, find root of J(τ_ℓ) = Gb_terminal.
  3. Compute the implied capital path, then the implied R*.
  Return R* - R_seq.

Iterate T(R_seq) = 0 with relaxation until convergence.
```

We iterate on the price sequence until factor markets clear.

```
@jit
def compute_factor_prices(k_prod, ε_bar, A, α, δ):
    """Compute factor prices from Cobb-Douglas."""

    k_per_eff = k_prod / ε_bar
    r = A * α * (k_per_eff ** (α - 1))
    w = A * (1 - α) * (k_per_eff ** α)
    RR = 1 + r - δ
    return r, w, RR
```

With endogenous prices, finding a steady state requires an outer fixed-point loop: solve the household problem at given prices, compute the implied capital stock, update prices via Cobb-Douglas marginal products, and repeat until convergence.

```
def find_ss_endo(
        debt2gdp_target, policy_target,
        hh, tech, RR_init=None, w_init=None,
        max_iter=50, tol=1e-5, verbose=False):
    """Find steady state with endogenous factor prices."""

    τ_a, τ_0, benef, G = policy_target
    ε_bar = float(jnp.sum(hh.frac * hh.ε_arr))

    RR = RR_init if RR_init else tech.RR
    w = w_init if w_init else tech.w

    relaxation = 0.3

    for iteration in range(max_iter):
        try:
            ss = ss_target_debt2gdp_exo(
```

(continues on next page)

(continued from previous page)

```
                debt2gdp_target,
                (τ_a, τ_0, benef, G),
                (RR, w), hh, tech
            )
        except ValueError:
            RR = RR * 0.99
            continue

        K_eff = ss.k_bar / hh.n + float(hh.frac[0] * hh.x0[0])
        r_new, w_new, RR_new = compute_factor_prices(
            K_eff, ε_bar,
            tech.A, tech.α_tilde, tech.δ
        )
        r_new = float(r_new)
        w_new = float(w_new)
        RR_new = float(RR_new)
        price_diff = abs(RR_new - RR) + abs(w_new - w)

        if verbose and iteration % 5 == 0:
            print(f"    SS iter {iteration}: "
                  f"RR={RR:.6f}, w={w:.4f}, "
                  f"k_bar={ss.k_bar:.4f}")

        if price_diff < tol:
            if verbose:
                print(f"    Converged at iteration {iteration}")
            break

        RR = RR + relaxation * (RR_new - RR)
        w = w + relaxation * (w_new - w)

    return ss_target_debt2gdp_exo(
        debt2gdp_target,
        (τ_a, τ_0, benef, G),
        (RR, w), hh, tech
    )
```

The price iteration also needs an initial guess for the transition price path and a way to update it after each inner solve.

The function `init_price_seq_interp` linearly interpolates between the two steady-state price vectors, while `_update_prices_from_capital` recomputes factor prices from the capital path via Cobb-Douglas marginal products.

```
def init_price_seq_interp(S, S1, S3, ss0_RR, ss0_w, ss1_RR, ss1_w):
    """Linearly interpolate price sequence between steady states."""

    s_indices = jnp.arange(S + 2)
    t_frac = jnp.clip((s_indices - S1) / (S3 - S1), 0.0, 1.0)

    RR_seq = ss0_RR + t_frac * (ss1_RR - ss0_RR)
    w_seq = ss0_w + t_frac * (ss1_w - ss0_w)

    RR_seq = jnp.where(s_indices <= S1, ss0_RR, RR_seq)
    w_seq = jnp.where(s_indices <= S1, ss0_w, w_seq)
    RR_seq = jnp.where(s_indices >= S3, ss1_RR, RR_seq)
    w_seq = jnp.where(s_indices >= S3, ss1_w, w_seq)
```

(continues on next page)

(continued from previous page)

```
    return jnp.column_stack([RR_seq, w_seq])


@jit
def _update_prices_from_capital(
        k_seq, k_bar_ss0, n, frac0_x0,
        ε_bar, A, α, δ,
        s_indices_full, ss1_RR, ss1_w, S3):
    """Compute new price sequence from the capital path."""

    k_prev = jnp.concatenate([jnp.array([k_bar_ss0]), k_seq[:-1]])
    K_eff = k_prev / n + frac0_x0

    k_per_eff = K_eff / ε_bar
    r_new = A * α * (k_per_eff ** (α - 1))
    w_new = A * (1 - α) * (k_per_eff ** α)
    RR_new = 1 + r_new - δ

    price_seq_new = jnp.column_stack([RR_new, w_new])
    price_seq_new = jnp.concatenate([price_seq_new, price_seq_new[-1:]], axis=0)

    terminal_prices = jnp.array([[ss1_RR, ss1_w]])
    mask = s_indices_full >= S3
    price_seq_new = jnp.where(mask[:, None], terminal_prices, price_seq_new)

    return price_seq_new
```

The top-level function `find_transition_endo` wraps everything in a relaxation loop: at each iteration it solves the transition under current prices, computes the implied capital path, updates prices, and checks for convergence.

```
def find_transition_endo(price_seq, policy_seq_base,
                         ss0, ss1, hh, tech, S, S1, S2, S3,
                         compensation_data=None,
                         max_iter=50, tol=1e-3,
                         relaxation=0.5, verbose=False):
    """Find transition with endogenous prices."""

    ε_bar = float(jnp.sum(hh.frac * hh.ε_arr))
    frac0_x0 = float(hh.frac[0] * hh.x0[0])

    price_seq = jnp.asarray(price_seq)
    policy_seq_base = jnp.asarray(policy_seq_base)

    # Pre-create iteration arrays for price update
    s_indices_full = jnp.arange(S + 2)

    if verbose:
        print("  Starting price iteration...")

    for iteration in range(max_iter):
        try:
            τ_l_trans, results = find_transition_exo(
                price_seq, policy_seq_base, ss0, ss1,
                hh, tech, S, S1, S2,
                compensation_data=compensation_data
            )
        except ValueError:
```

(continues on next page)

(continued from previous page)

```
            τ_l_trans = 0.35
            results = transition_paths(
                τ_l_trans, price_seq, policy_seq_base,
                ss0, ss1, hh, tech, S, S1, S2,
                ss0.μx_arr, ss0.Σx_arr
            )

        _, μx_seq, μc_seq, k_seq, Gb_seq, F_seq, Ao_seq = results

        price_seq_new = _update_prices_from_capital(
            k_seq, float(ss0.k_bar), hh.n, float(frac0_x0), float(ε_bar),
            tech.A, tech.α_tilde, tech.δ,
            s_indices_full,
            float(ss1.RR), float(ss1.w), S3
        )

        price_diff = float(jnp.max(jnp.abs(price_seq_new - price_seq)))

        if verbose:
            print(f"  Iter {iteration}: "
                  f"τ_l={τ_l_trans:.4f}, "
                  f"price_diff={price_diff:.6f}")

        if price_diff < tol:
            if verbose:
                print(f"  Converged at iteration {iteration}")
            break

        price_seq = price_seq + relaxation * (price_seq_new - price_seq)

    return τ_l_trans, price_seq, results
```

We now compute the initial and terminal steady states under endogenous prices and solve the transition with price iteration.

```
# Compute endogenous prices for initial SS
ε_bar = float(jnp.sum(hh.frac * hh.ε_arr))
K_eff_0 = ss0.k_bar / hh.n + float(hh.frac[0] * hh.x0[0])
r0_endo, w0_endo, RR0_endo = compute_factor_prices(
    K_eff_0, ε_bar,
    tech.A, tech.α_tilde, tech.δ
)
r0_endo = float(r0_endo)
w0_endo = float(w0_endo)
RR0_endo = float(RR0_endo)

ss0_exp1_endo = SteadyState(
    P_arr=ss0.P_arr, ξ_arr=ss0.ξ_arr, Ao_arr=ss0.Ao_arr, F_arr=ss0.F_arr,
    μx_arr=ss0.μx_arr, μc_arr=ss0.μc_arr, Σx_arr=ss0.Σx_arr, Vc_arr=ss0.Vc_arr,
    debt2gdp=ss0.debt2gdp, τ_l=ss0.τ_l, benef=ss0.benef, Gb=ss0.Gb,
    k_bar=ss0.k_bar, RR=RR0_endo, w=w0_endo, r=r0_endo, k2gdp=ss0.k2gdp
)

ss1_exp1_endo = find_ss_endo(
    ss0.debt2gdp, (τ_a_0, τ_0_0, 0, G_0), hh, tech,
    RR_init=tech.RR, w_init=tech.w, verbose=True
)
```

```
    SS iter 0: RR=1.067500, w=5.0147, k_bar=35.8906
    SS iter 5: RR=1.044957, w=5.7366, k_bar=30.4338
    SS iter 10: RR=1.044407, w=5.7330, k_bar=30.1611
    SS iter 15: RR=1.044405, w=5.7291, k_bar=30.1587
    SS iter 20: RR=1.044406, w=5.7284, k_bar=30.1588
    SS iter 25: RR=1.044406, w=5.7283, k_bar=30.1588
    Converged at iteration 28
```

The initial price guess interpolates linearly between the two steady states.

```
# Initialize price sequence
price_seq_exp1_endo = init_price_seq_interp(
    S_endo, S1, S3,
    float(ss0_exp1_endo.RR), float(ss0_exp1_endo.w),
    float(ss1_exp1_endo.RR), float(ss1_exp1_endo.w)
)

# Policy sequence
policy_seq_exp1_endo = make_policy_seq(
    S_endo,
    ss0_exp1_endo.τ_l, ss1_exp1_endo.τ_l,
    S1, S2,
    τ_a_0, τ_0_0, 0, G_0,
    ss0_exp1_endo.Gb
)
```

The benefit difference vector records the benefit loss at each age, and the price iteration finds the equilibrium transition path with buy-out compensation.

```
# Buyout compensation
benef_diff_exp1_endo = jnp.zeros(hh.T0 + 1)
benef_diff = ss0_exp1_endo.benef - ss1_exp1_endo.benef
benef_diff_exp1_endo = benef_diff_exp1_endo.at[
    hh.T1 + 1:
].set(benef_diff)

# Solve with price iteration
print("\n  Solving transition with endogenous prices...")
endo_result = find_transition_endo(
    price_seq_exp1_endo, policy_seq_exp1_endo,
    ss0_exp1_endo, ss1_exp1_endo,
    hh, tech, S_endo, S1, S2, S3,
    compensation_data=(
        benef_diff_exp1_endo,
        ss0_exp1_endo.τ_l
    ),
    verbose=True
)
τ_l_exp1_endo_bo = endo_result[0]
price_seq_exp1_endo_conv = endo_result[1]
results_exp1_endo = endo_result[2]
(_, μx_seq_exp1_endo, μc_seq_exp1_endo,
 k_seq_exp1_endo, Gb_seq_exp1_endo,
 _, _) = results_exp1_endo
```

```
  Solving transition with endogenous prices...
  Starting price iteration...
```

(continues on next page)

(continued from previous page)

```
  Iter 0: τ_l=0.3391, price_diff=0.490311
  Iter 1: τ_l=0.3549, price_diff=0.184939
  Iter 2: τ_l=0.3637, price_diff=0.064441
  Iter 3: τ_l=0.3676, price_diff=0.021013
  Iter 4: τ_l=0.3692, price_diff=0.006973
  Iter 5: τ_l=0.3698, price_diff=0.002371
  Iter 6: τ_l=0.3701, price_diff=0.000828
  Converged at iteration 6
```

The endogenous-price results are stored for comparison with the fixed-price case.

```
exp1_endo = {
    'ss0': ss0_exp1_endo, 'ss1': ss1_exp1_endo,
    'τ_l_buyout': τ_l_exp1_endo_bo,
    'price_seq': price_seq_exp1_endo_conv,
    'k_seq': k_seq_exp1_endo, 'Gb_seq': Gb_seq_exp1_endo,
    'results': results_exp1_endo,
    'μc_seq': μc_seq_exp1_endo, 'μx_seq': μx_seq_exp1_endo
}
```

The following figure compares the transition paths under fixed and endogenous factor prices, showing how general equilibrium effects alter the tax, debt, interest rate, and wage paths.

```
# Get endogenous price sequences
price_seq_endo = exp1_endo['price_seq']
S_endo = price_seq_endo.shape[0] - 2

# Construct fixed price sequences for comparison
RR_fixed = tech.RR
w_fixed = tech.w

# For fixed prices, construct τ_l sequence
τ_l_seq_fixed = np.zeros(S_exo + 1)
τ_l_seq_fixed[:S1 + 1] = ss0_exp1.τ_l
τ_l_seq_fixed[S1 + 1:S2 + 1] = exp1_exo['τ_l_buyout']
τ_l_seq_fixed[S2 + 1:] = ss1_exp1.τ_l

# For endogenous prices
τ_l_seq_endo = np.zeros(S_endo + 1)
τ_l_seq_endo[:S1 + 1] = exp1_endo['ss0'].τ_l
τ_l_seq_endo[S1 + 1:S2 + 1] = exp1_endo['τ_l_buyout']
τ_l_seq_endo[S2 + 1:] = exp1_endo['ss1'].τ_l

fig, axes = plt.subplots(2, 2, figsize=(14, 10))

# Labor tax comparison
axes[0, 0].plot(τ_l_seq_fixed, 'b-', linewidth=2, label='Fixed prices')
axes[0, 0].plot(τ_l_seq_endo[:len(τ_l_seq_fixed)],
                'r--', linewidth=2,
                label='Endogenous prices')
axes[0, 0].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[0, 0].set_xlabel('time (s)')
axes[0, 0].set_ylabel(r'$\tau_\ell$')
axes[0, 0].set_title('Labor tax rate path')
axes[0, 0].legend()

# Government debt comparison
```

(continues on next page)

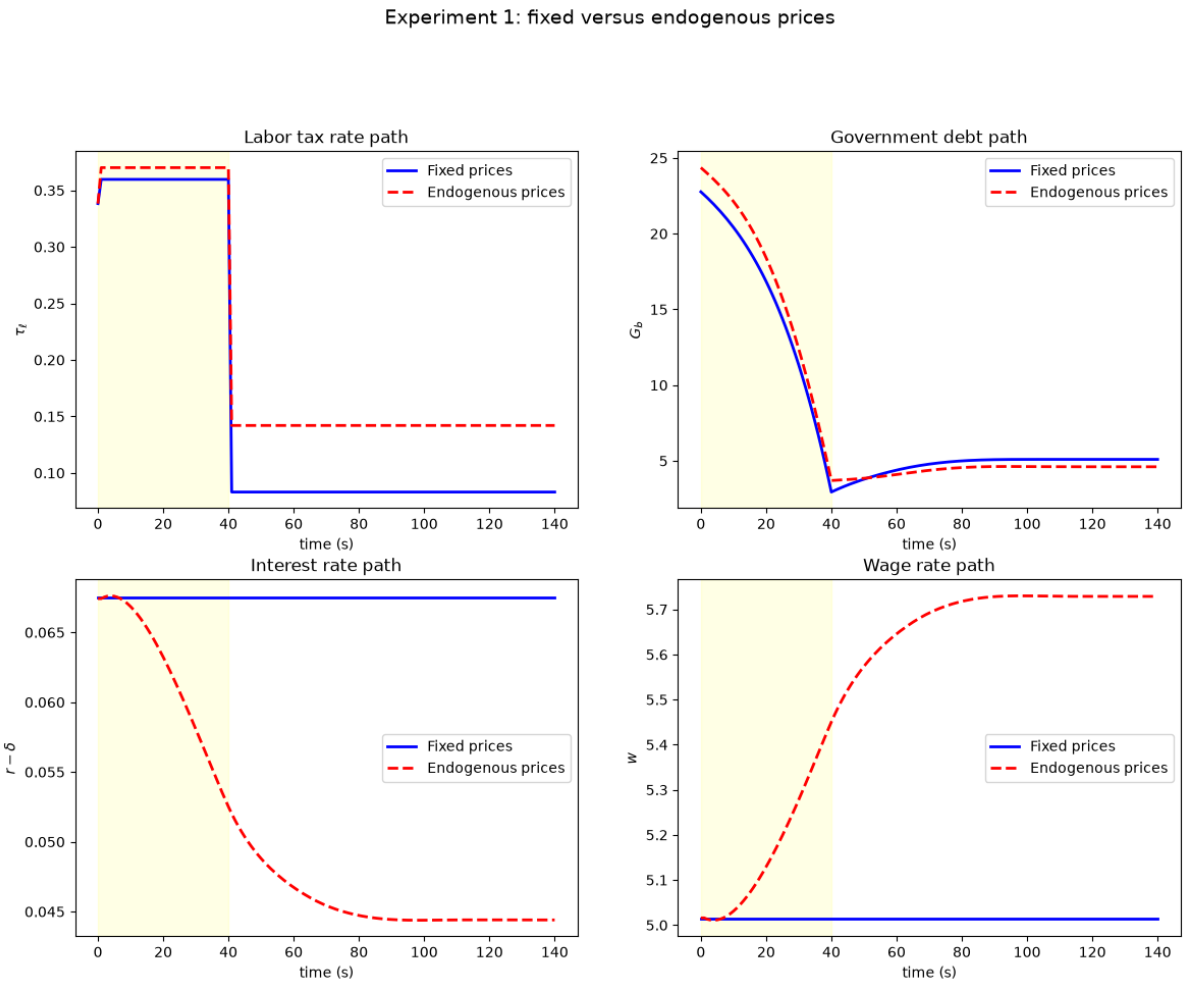
(continued from previous page)

```
Gb_seq_fixed = Gb_seq_bo
Gb_seq_endo_exp1 = exp1_endo['Gb_seq']
axes[0, 1].plot(Gb_seq_fixed, 'b-', linewidth=2, label='Fixed prices')
axes[0, 1].plot(
    Gb_seq_endo_exp1[:len(Gb_seq_fixed)],
    'r--', linewidth=2,
    label='Endogenous prices'
)
axes[0, 1].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[0, 1].set_xlabel('time (s)')
axes[0, 1].set_ylabel(r'$G_b$')
axes[0, 1].set_title('Government debt path')
axes[0, 1].legend()

# Interest rate comparison
r_fixed = np.full(S_exo + 1, tech.r - tech.δ)
r_endo = price_seq_endo[:-1, 0] - 1  # RR - 1 = r - δ
axes[1, 0].plot(r_fixed, 'b-', linewidth=2, label='Fixed prices')
axes[1, 0].plot(
    r_endo[:len(r_fixed)],
    'r--', linewidth=2,
    label='Endogenous prices'
)
axes[1, 0].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[1, 0].set_xlabel('time (s)')
axes[1, 0].set_ylabel(r'$r - \delta$')
axes[1, 0].set_title('Interest rate path')
axes[1, 0].legend()

# Wage rate comparison
w_fixed_seq = np.full(S_exo + 1, tech.w)
w_endo = price_seq_endo[:-1, 1]
axes[1, 1].plot(w_fixed_seq, 'b-', linewidth=2, label='Fixed prices')
axes[1, 1].plot(
    w_endo[:len(w_fixed_seq)],
    'r--', linewidth=2,
    label='Endogenous prices'
)
axes[1, 1].axvspan(S1, S2, alpha=0.1, color='yellow')
axes[1, 1].set_xlabel('time (s)')
axes[1, 1].set_ylabel(r'$w$')
axes[1, 1].set_title('Wage rate path')
axes[1, 1].legend()

plt.suptitle(
    'Experiment 1: fixed versus endogenous prices',
    fontsize=14, y=1.02
)
plt.show()
```



The top-left panel shows that the transition tax rate is similar under both price assumptions, but the bottom panels reveal important differences in factor prices.

As the capital stock rises during the transition, the interest rate falls and the wage rises under endogenous pricing, while these remain constant under the small open economy assumption.

These price effects create additional redistributive consequences beyond those intended by the policy change: lower interest rates benefit young workers through higher wages, but hurt retirees through lower returns on savings.

127.9 Experiment 2: government capital accumulation

In Experiment 2, the government maintains social security benefits but temporarily raises taxes to accumulate physical capital.

The returns from this capital eventually finance the social security payments, so that in the terminal steady state the government is a net creditor rather than a debtor.

Unlike Experiment 1, which eliminates benefits, this approach preserves the social insurance function of social security – namely, insurance against life span risk and partial insurance against labor income volatility.

By having the government save on behalf of households, the economy can achieve higher capital accumulation while maintaining the intergenerational insurance that social security provides.

Algorithm 127.9.1 (Fixed factor prices – Experiment 2 (government funding))

Here $s_3 = s_2 + T_0 = 105$.

Steps 1–3 are the same as *Algorithm 127.8.1*, except that social security benefits are maintained and the target debt-to-GDP ratio is negative (the government becomes a net creditor).

The right target induces the government to accumulate sufficient capital to finance benefits from asset returns.

Step 4 is the same root-finding procedure over $\tau_{\ell,1}$, but without compensation payments.

We first compute the terminal steady state under fixed prices, targeting a negative debt-to-GDP ratio that makes the government a net creditor.

The target is chosen so that the income on the government's assets is enough to pay for social security benefits forever.

In the stationary equilibrium the government's net asset income per capita is $[R/n - 1](-\bar{b})$, and the benefit bill is $\sum_{t>T_1} S_t f_t$, so a first guess at the required asset position solves

$$\left[\frac{R}{n} - 1\right](-\bar{b}) = \sum_{t>T_1} S_t f_t$$

```
benefit_bill = float(jnp.sum(benef_0 * hh.frac * (1.0 - hh.Ind_work)))
Rn_net = tech.RR / hh.n - 1.0

print(f"benefit bill per capita        = {benefit_bill:.4f}")
print(f"net asset income per unit debt = {Rn_net:.4f}")
print(f"implied asset position b_bar   = {-benefit_bill / Rn_net:.4f}")
```

```
benefit bill per capita        = 0.4979
net asset income per unit debt = 0.0548
implied asset position b_bar   = -9.0793
```

The targets below ask for a slightly larger asset position than this calculation, which leaves some room for the government to keep financing its purchases g with distorting taxes.

Exercise 127.12.3 compares the two numbers.

```
debt2gdp_target_exp2_exo = -1.1785
ss1_exp2_exo = ss_target_debt2gdp_exo(
    debt2gdp_target_exp2_exo,
    (τ_a_0, τ_0_0, benef_0, G_0),
    (RR_exo, w_exo), hh, tech
)
```

With the terminal steady state in hand, bisection over the transition tax rate gives the equilibrium path under fixed prices.

```
# Price and policy sequences
price_seq_exp2_exo = make_exo_price_seq(S_exo, RR_exo, w_exo)
policy_seq_exp2_exo = make_policy_seq(
    S_exo, ss0.τ_l, ss1_exp2_exo.τ_l,
    S1, S2,
    τ_a_0, τ_0_0, benef_0, G_0, Gb_0
)

# Solve (no compensation)
τ_l_exp2_exo, results_exp2_exo = find_transition_exo(
```

(continues on next page)

(continued from previous page)

```
    price_seq_exp2_exo, policy_seq_exp2_exo, ss0, ss1_exp2_exo,
    hh, tech, S_exo, S1, S2
)
```

The results are packaged for the cross-experiment comparison below.

```
(carryover_exp2, μx_seq_exp2_exo,
 μc_seq_exp2_exo, k_seq_exp2_exo,
 Gb_seq_exp2_exo, F_seq_exp2_exo,
 Ao_seq_exp2_exo) = results_exp2_exo
exp2_exo = {
    'ss0': ss0, 'ss1': ss1_exp2_exo,
    'τ_l_trans': τ_l_exp2_exo,
    'results': results_exp2_exo,
    'k_seq': k_seq_exp2_exo, 'Gb_seq': Gb_seq_exp2_exo,
    'μc_seq': μc_seq_exp2_exo, 'μx_seq': μx_seq_exp2_exo
}
```

We repeat the computation under endogenous factor prices, using the same initial steady state as Experiment 1.

```
# Compute endogenous prices for initial SS (reuse from Exp 1)
ss0_exp2_endo = ss0_exp1_endo  # Same initial SS

# Terminal steady state with endogenous prices
ss1_exp2_endo = find_ss_endo(
    debt2gdp_target=-1.925,
    policy_target=(τ_a_0, τ_0_0, benef_0, G_0),
    hh=hh, tech=tech,
    RR_init=tech.RR, w_init=tech.w,
    verbose=True
)
```

```
    SS iter 0: RR=1.067500, w=5.0147, k_bar=43.5570
    SS iter 5: RR=1.044472, w=5.7664, k_bar=30.0436
    SS iter 10: RR=1.044623, w=5.7272, k_bar=30.0407
    SS iter 15: RR=1.044648, w=5.7206, k_bar=30.0403
    SS iter 20: RR=1.044651, w=5.7195, k_bar=30.0409
    SS iter 25: RR=1.044651, w=5.7193, k_bar=30.0409
    Converged at iteration 29
```

Price iteration produces the endogenous-price transition path.

```
# Initialize price sequence
price_seq_exp2_endo = init_price_seq_interp(
    S_endo, S1, S3,
    float(ss0_exp2_endo.RR), float(ss0_exp2_endo.w),
    float(ss1_exp2_endo.RR), float(ss1_exp2_endo.w)
)

# Policy sequence
policy_seq_exp2_endo = make_policy_seq(
    S_endo,
    ss0_exp2_endo.τ_l, ss1_exp2_endo.τ_l,
    S1, S2,
    τ_a_0, τ_0_0, benef_0, G_0,
    ss0_exp2_endo.Gb
```

(continues on next page)

(continued from previous page)

```
)

# Solve with price iteration (no compensation)
endo2 = find_transition_endo(
    price_seq_exp2_endo, policy_seq_exp2_endo,
    ss0_exp2_endo, ss1_exp2_endo,
    hh, tech, S_endo, S1, S2, S3,
    verbose=True
)
τ_l_exp2_endo = endo2[0]
price_seq_exp2_endo_conv = endo2[1]
results_exp2_endo = endo2[2]
(_, μx_seq_exp2_endo, μc_seq_exp2_endo,
 k_seq_exp2_endo, Gb_seq_exp2_endo,
 _, _) = results_exp2_endo
```

```
Starting price iteration...
Iter 0: τ_l=0.3660, price_diff=0.394320
Iter 1: τ_l=0.3769, price_diff=0.202210
Iter 2: τ_l=0.3845, price_diff=0.103709
Iter 3: τ_l=0.3891, price_diff=0.050465
Iter 4: τ_l=0.3916, price_diff=0.022852
Iter 5: τ_l=0.3928, price_diff=0.009652
Iter 6: τ_l=0.3933, price_diff=0.003839
Iter 7: τ_l=0.3935, price_diff=0.001457
Iter 8: τ_l=0.3936, price_diff=0.000537
Converged at iteration 8
```

```
exp2_endo = {
    'ss0': ss0_exp2_endo, 'ss1': ss1_exp2_endo,
    'τ_l_trans': τ_l_exp2_endo,
    'price_seq': price_seq_exp2_endo_conv,
    'k_seq': k_seq_exp2_endo, 'Gb_seq': Gb_seq_exp2_endo,
    'results': results_exp2_endo,
    'μc_seq': μc_seq_exp2_endo, 'μx_seq': μx_seq_exp2_endo
}
```

We now compare all four reform scenarios: the buy-out scheme and the government funding scheme, each under fixed and endogenous factor prices.

```
# Get debt sequences for all cases
Gb_buyout_fixed = Gb_seq_bo
Gb_buyout_endo = exp1_endo['Gb_seq']
Gb_accum_fixed = exp2_exo['Gb_seq']
Gb_accum_endo = exp2_endo['Gb_seq']

# Get capital sequences
k_buyout_fixed = k_seq_bo
k_buyout_endo = exp1_endo['k_seq']
k_accum_fixed = exp2_exo['k_seq']
k_accum_endo = exp2_endo['k_seq']

# Common time horizon for plotting
T_plot = min(len(Gb_buyout_fixed), len(Gb_buyout_endo),
             len(Gb_accum_fixed), len(Gb_accum_endo))
```

(continues on next page)

(continued from previous page)

```
fig, axes = plt.subplots(2, 3, figsize=(18, 10))

# Labels for four scenarios
lb = ['Buy-out (fixed)', 'Buy-out (endogenous)',
      'Gov funding (fixed)', 'Gov funding (endogenous)']
ls = ['b-', 'b--', 'r-', 'r--']

# Government Debt paths
ax = axes[0, 0]
for d, s, l in zip(
    [Gb_buyout_fixed, Gb_buyout_endo,
     Gb_accum_fixed, Gb_accum_endo],
    ls, lb
):
    ax.plot(d[:T_plot], s, linewidth=2, label=l)
ax.axhline(0, color='k', linestyle=':', alpha=0.5)
ax.axvspan(0, 40, alpha=0.1, color='yellow')
ax.set_xlabel('time (s)')
ax.set_ylabel(r'$G_b$')
ax.set_title('Government debt paths')
ax.legend(fontsize=9)


# Capital paths
ax = axes[0, 1]
for d, s, l in zip(
    [k_buyout_fixed, k_buyout_endo,
     k_accum_fixed, k_accum_endo],
    ls, lb
):
    ax.plot(d[:T_plot], s, linewidth=2, label=l)
ax.axvspan(0, 40, alpha=0.1, color='yellow')
ax.set_xlabel('time (s)')
ax.set_ylabel(r'$K$')
ax.set_title('Capital accumulation paths')
ax.legend(fontsize=9)


# Aggregate consumption
res2_exo = exp2_exo['results']
res1_endo = exp1_endo['results']
res2_endo = exp2_endo['results']
_, _, μc_exp2_exo, _, _, _, _ = res2_exo
_, _, μc_exp1_endo, _, _, _, _ = res1_endo
_, _, μc_exp2_endo, _, _, _, _ = res2_endo

def _agg_c(μc, T):
    """Aggregate per-capita consumption across cohorts."""

    n = min(T, μc.shape[0])
    return np.array([
        np.sum(μc[s] * hh.frac) for s in range(n)
    ])

c_agg_buyout_fixed = c_agg_bo[:T_plot]
c_agg_buyout_endo = _agg_c(μc_exp1_endo, T_plot)
c_agg_accum_fixed = _agg_c(μc_exp2_exo, T_plot)
```

(continues on next page)

(continued from previous page)

```
c_agg_accum_endo = _agg_c(μc_exp2_endo, T_plot)

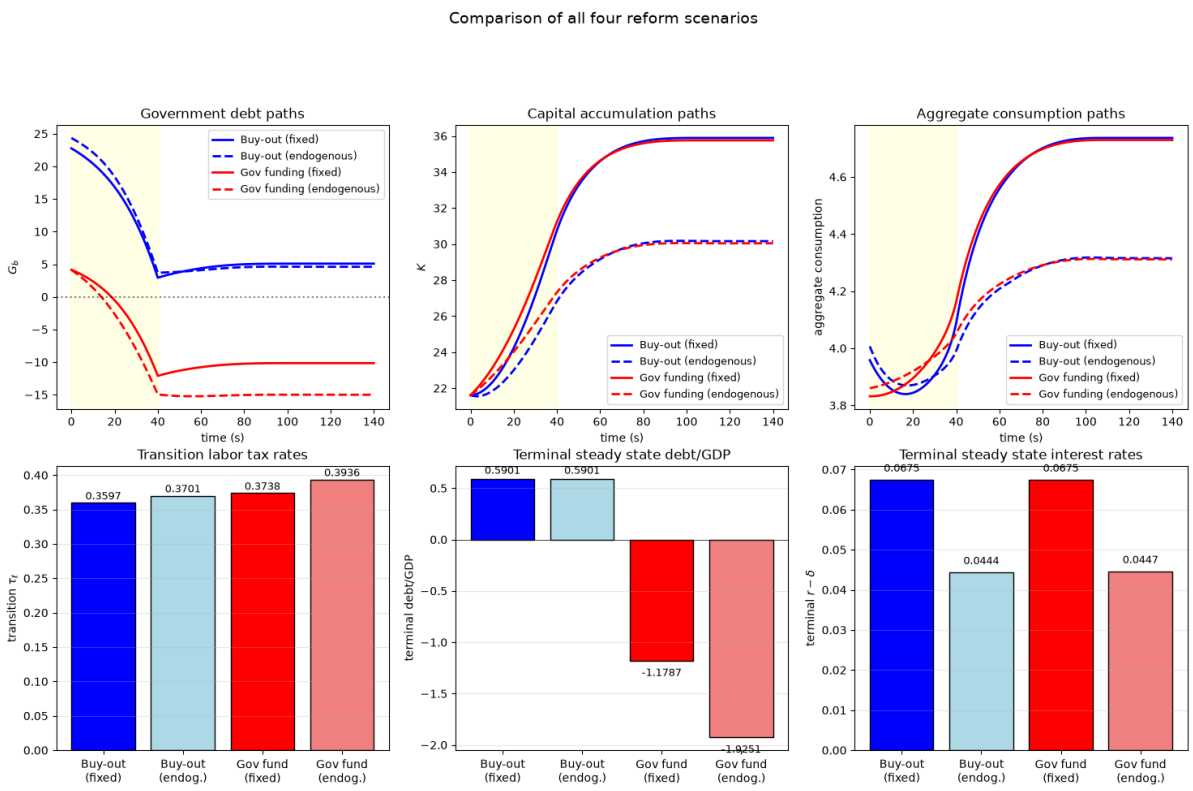
ax = axes[0, 2]
for d, s, l in zip(
    [c_agg_buyout_fixed, c_agg_buyout_endo,
     c_agg_accum_fixed, c_agg_accum_endo],
    ls, lb
):
    ax.plot(d[:T_plot], s, linewidth=2, label=l)
axes[0, 2].axvspan(0, 40, alpha=0.1, color='yellow')
axes[0, 2].set_xlabel('time (s)')
axes[0, 2].set_ylabel('aggregate consumption')
axes[0, 2].set_title('Aggregate consumption paths')
axes[0, 2].legend(fontsize=9)

# Bar chart: Transition tax rates
cases = ['Buy-out\n(fixed)', 'Buy-out\n(endog.)',
         'Gov fund\n(fixed)', 'Gov fund\n(endog.)']
τ_l_values = [exp1_exo['τ_l_buyout'], exp1_endo['τ_l_buyout'],
              exp2_exo['τ_l_trans'], exp2_endo['τ_l_trans']]
colors = ['blue', 'lightblue', 'red', 'lightcoral']
axes[1, 0].bar(cases, τ_l_values, color=colors, edgecolor='black')
axes[1, 0].set_ylabel(r'transition $\tau_\ell$')
axes[1, 0].set_title('Transition labor tax rates')
axes[1, 0].grid(True, alpha=0.3, axis='y')
for i, v in enumerate(τ_l_values):
    axes[1, 0].text(i, v + 0.005, f'{v:.4f}', ha='center', fontsize=9)

# Bar chart: Terminal debt/GDP
debt2gdp_values = [exp1_exo['ss1'].debt2gdp, exp1_endo['ss1'].debt2gdp,
                   exp2_exo['ss1'].debt2gdp, exp2_endo['ss1'].debt2gdp]
axes[1, 1].bar(cases, debt2gdp_values, color=colors, edgecolor='black')
axes[1, 1].axhline(0, color='k', linestyle='-', linewidth=0.5)
axes[1, 1].set_ylabel('terminal debt/GDP')
axes[1, 1].set_title('Terminal steady state debt/GDP')
axes[1, 1].grid(True, alpha=0.3, axis='y')
for i, v in enumerate(debt2gdp_values):
    y = v + 0.05 if v > 0 else v - 0.15
    axes[1, 1].text(
        i, y, f'{v:.4f}',
        ha='center', fontsize=9
    )

# Bar chart: Terminal interest rate
r_values = [exp1_exo['ss1'].r - tech.δ, exp1_endo['ss1'].r - tech.δ,
            exp2_exo['ss1'].r - tech.δ, exp2_endo['ss1'].r - tech.δ]
axes[1, 2].bar(cases, r_values, color=colors, edgecolor='black')
axes[1, 2].set_ylabel(r'terminal $r - \delta$')
axes[1, 2].set_title('Terminal steady state interest rates')
axes[1, 2].grid(True, alpha=0.3, axis='y')
for i, v in enumerate(r_values):
    axes[1, 2].text(i, v + 0.002, f'{v:.4f}', ha='center', fontsize=9)

plt.suptitle('Comparison of all four reform scenarios', fontsize=14, y=1.02)
plt.show()
```



The top row compares transition dynamics: the buy-out scheme (blue) accumulates higher debt during the transition due to compensation payments, while the government funding scheme (red) leads to large negative debt as the government becomes a net creditor.

The bottom-left bar chart shows that government funding requires higher transition tax rates than the buy-out scheme because benefit payments continue alongside capital accumulation.

Under endogenous pricing, the larger capital stock reduces the marginal product of capital and hence interest rates, as shown in the bottom-right panel.

To compare the two reforms we need numbers, not only paths.

The first summary collects the transition tax rate and the terminal steady state of each scenario.

```
summary = [
    ('buy-out, fixed prices',    exp1_exo['τ_l_buyout'],  exp1_exo['ss1']),
    ('buy-out, endogenous',      exp1_endo['τ_l_buyout'], exp1_endo['ss1']),
    ('gov funding, fixed',       exp2_exo['τ_l_trans'],   exp2_exo['ss1']),
    ('gov funding, endogenous',  exp2_endo['τ_l_trans'],  exp2_endo['ss1']),
]

print(f"{'scenario':>24}  {'τ_l trans':>9}  {'τ_l final':>9}"
      f"  {'K/GDP':>7}  {'debt/GDP':>8}  {'r − δ':>7}")
for name, τ_trans, ss in summary:
    print(f"{name:>24}  {τ_trans:>9.4f}  {ss.τ_l:>9.4f}"
          f"  {ss.k2gdp:>7.4f}  {ss.debt2gdp:>8.4f}  {ss.r - tech.δ:>7.4f}")
```

```
                scenario  τ_l trans  τ_l final    K/GDP  debt/GDP     r − δ
  buy-out, fixed prices     0.3597     0.0831   4.1567    0.5901   0.0675
    buy-out, endogenous     0.3701     0.1420   3.8640    0.5901   0.0444
     gov funding, fixed     0.3738     0.1390   4.1493   -1.1787   0.0675
```

(continues on next page)

(continued from previous page)

```
gov funding, endogenous     0.3936     0.2333   3.8549   -1.9251   0.0447
```

The second summary asks which reform a person born into the terminal steady state would prefer.

Because the value function is $U_t = x_t' P_t x_t + \xi_t$, the welfare of a newborn follows directly from objects that we have already computed.

```
def newborn_welfare(ss):
    "Expected lifetime utility of an agent born into a steady state."

    return float(hh.x0 @ ss.P_arr[0] @ hh.x0 + ss.ξ_arr[0])

print(f"{'steady state':>34}  {'U of a newborn':>14}")
for name, ss in [
    ('initial, pay-as-you-go', ss0),
    ('terminal, buy-out (no benefits)', exp1_exo['ss1']),
    ('terminal, government funding', exp2_exo['ss1']),
]:
    print(f"{name:>34}  {newborn_welfare(ss):>14.2f}")
```

```
                    steady state  U of a newborn
          initial, pay-as-you-go         -262.74
  terminal, buy-out (no benefits)         -192.18
     terminal, government funding         -180.88
```

A newborn prefers either reform to the pay-as-you-go system that we started from.

These are fixed-price steady states, so the wage is the same in all three rows; what lifts welfare is the labor tax rate, which falls from 0.34 to 0.08 under the buy-out and to 0.14 under government funding.

Between the two reforms, a newborn prefers government funding.

That scheme taxes labor more heavily than the buy-out does, but it keeps paying social security benefits, which continue to insure against life span risk and to provide partial insurance against labor income risk, and a newborn values those benefits more than the extra taxes cost.

This comparison ranks steady states.

It says nothing about the generations alive during the transition, whose fortunes the cohort consumption paths plotted above describe.

127.10 Distribution surfaces

The 3D surface plots below show how assets and consumption evolve over both the age dimension and calendar time.

```
# Compute variances for 3D plotting
def compute_variances(results, ss0, hh):
    """Compute variance sequences from transition results."""

    _, μx_seq, μc_seq, k_seq, Gb_seq, F_seq, Ao_seq = results

    # Convert to numpy
    μx_seq = np.array(μx_seq)
    μc_seq_full = np.array(μc_seq)
    F_seq = np.array(F_seq)
    Ao_seq = np.array(Ao_seq)
```

(continues on next page)

(continued from previous page)

```
    Σx_arr_ss0 = np.array(ss0.Σx_arr)
    Σ0 = np.array(hh.Σ0)
    C = np.array(hh.C)

    # Get actual dimensions from data
    S_plus_1 = Ao_seq.shape[0]  # S+1
    T0_plus_1 = Ao_seq.shape[1]  # T0+1

    Σx_seq = np.empty((S_plus_1 + 1, T0_plus_1 + 1, hh.n_x, hh.n_x))
    Vc_seq = np.empty((S_plus_1, T0_plus_1))
    Va_seq = np.empty((S_plus_1, T0_plus_1))

    Σx_seq[:, 0] = Σ0
    Σx_seq[0, :] = Σx_arr_ss0[:T0_plus_1 + 1]

    CCT = C @ C.T
    for s in range(S_plus_1):
        Ao_s = Ao_seq[s]                    # (T0+1, n_x, n_x)
        Σx_s = Σx_seq[s, :T0_plus_1]        # (T0+1, n_x, n_x)
        F_s = F_seq[s]                       # (T0+1, n_x)
        Σx_seq[s + 1, 1:] = CCT + Ao_s @ Σx_s @ Ao_s.transpose(0, 2, 1)
        Vc_seq[s] = np.einsum('ti,tij,tj->t', F_s, Σx_s, F_s)
        Va_seq[s] = Σx_s[:, 0, 0]

    # Extract mean assets - match dimensions with Ao_seq
    μa_seq = μx_seq[:S_plus_1, :T0_plus_1, 0]
    μc_seq_out = μc_seq_full[:S_plus_1, :T0_plus_1]

    return μa_seq, Va_seq, μc_seq_out, Vc_seq

# Compute variances for each case
μa_bf, Va_bf, μc_bf, Vc_bf = compute_variances(
    exp1_exo['results_buyout'], exp1_exo['ss0'], hh
)
μa_be, Va_be, μc_be, Vc_be = compute_variances(
    exp1_endo['results'], exp1_endo['ss0'], hh
)
μa_af, Va_af, μc_af, Vc_af = compute_variances(
    exp2_exo['results'], exp2_exo['ss0'], hh
)
μa_ae, Va_ae, μc_ae, Vc_ae = compute_variances(
    exp2_endo['results'], exp2_endo['ss0'], hh
)

case_names = [
    'Buy-out (fixed)', 'Buy-out (endogenous)',
    'Gov funding (fixed)', 'Gov funding (endogenous)'
]


def plot_surface_grid(
        data_cases, case_names, zlabel,
        suptitle, cmap='viridis',
        transform=None):
    """Plot 2x2 grid of 3D surfaces for age-time data."""

    fig = plt.figure(figsize=(16, 12))
```

(continues on next page)

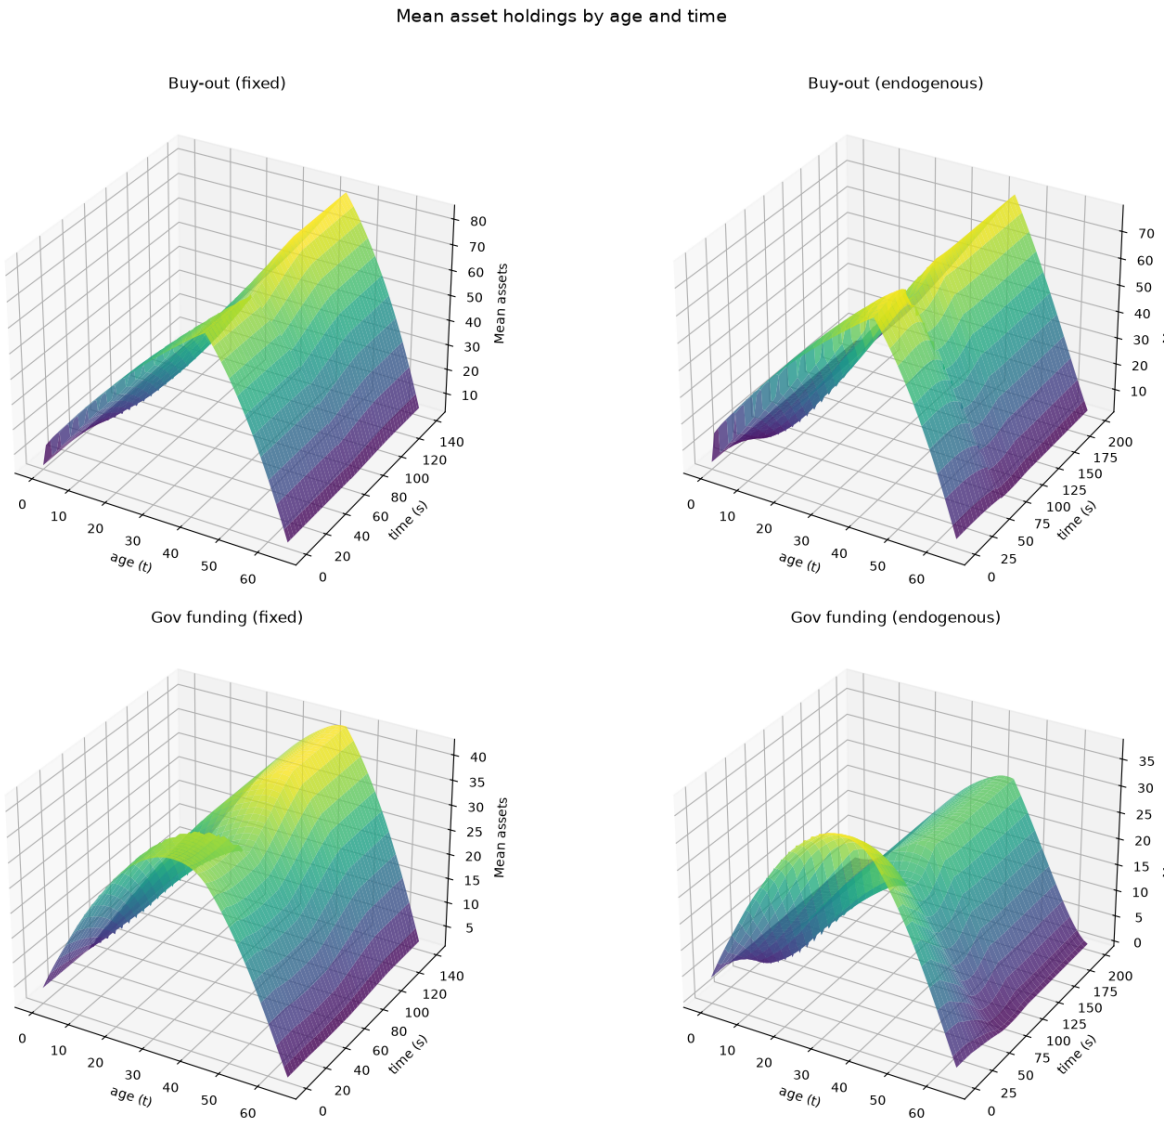
(continued from previous page)

```
    for i, (data, name) in enumerate(zip(data_cases, case_names)):
        Z = transform(data) if transform is not None else data
        n_time, n_age = Z.shape
        X, Y = np.meshgrid(np.arange(n_age), np.arange(n_time))
        ax = fig.add_subplot(2, 2, i + 1, projection='3d')
        ax.plot_surface(X, Y, Z, cmap=cmap, edgecolor='none', alpha=0.8)
        ax.set_xlabel('age (t)')
        ax.set_ylabel('time (s)')
        ax.set_zlabel(zlabel)
        ax.set_title(name)
    plt.suptitle(suptitle, fontsize=14, y=1.02)
    plt.tight_layout()
    plt.show()
```

Each surface shows the joint distribution across age (t) and calendar time (s), revealing life-cycle patterns, transition dynamics, and cross-cohort heterogeneity.

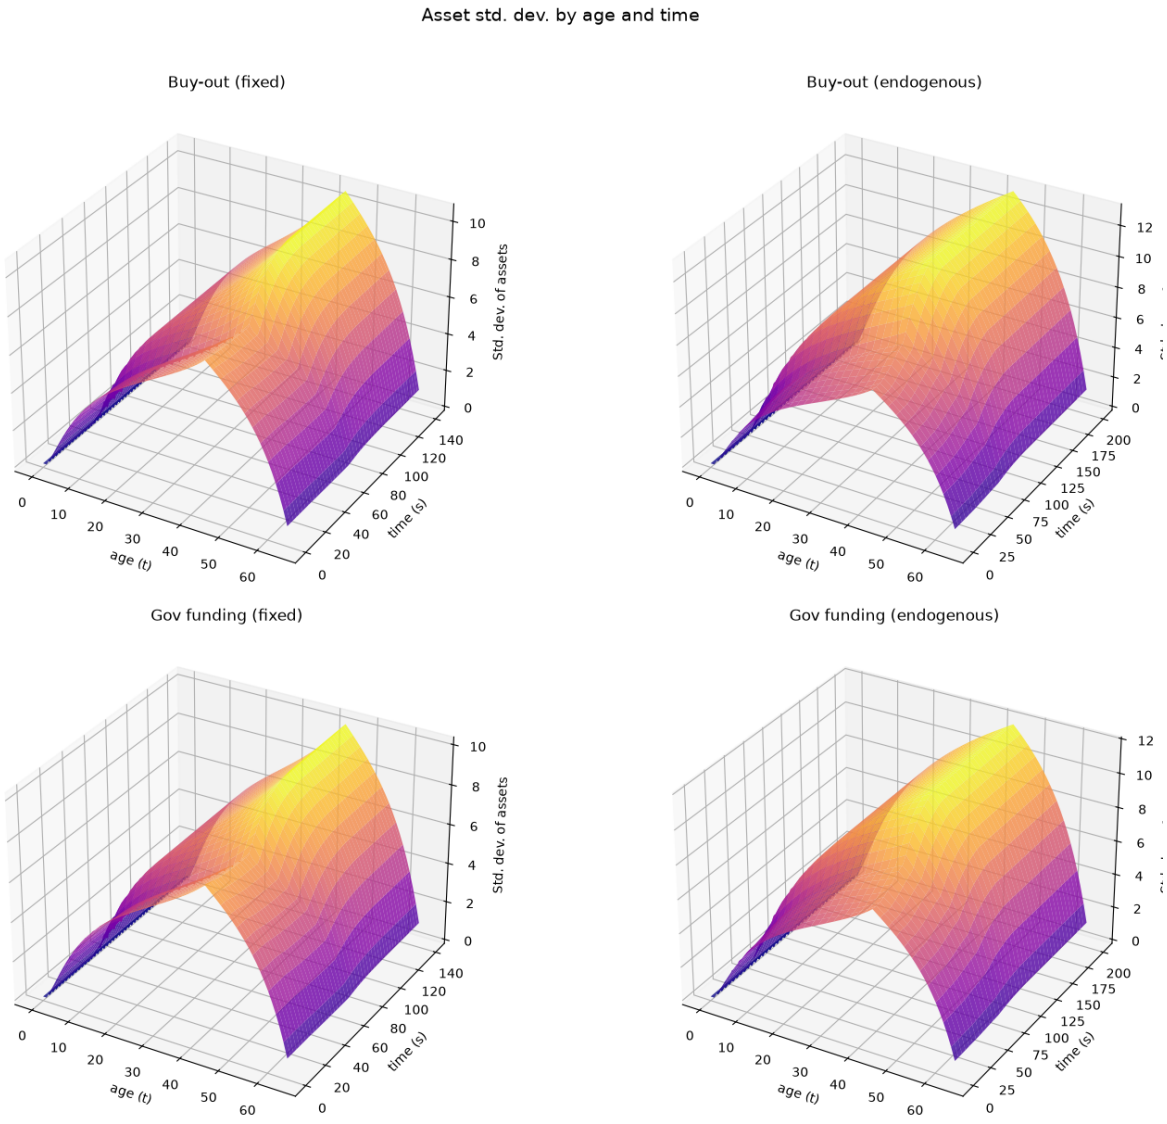
The mean asset surfaces display the hump-shaped life-cycle profile of asset holdings, with peak assets shifting as working generations increase their saving in response to the reform.

```
plot_surface_grid(
    [μa_bf, μa_be, μa_af, μa_ae],
    case_names, 'Mean assets',
    'Mean asset holdings by age and time'
)
```



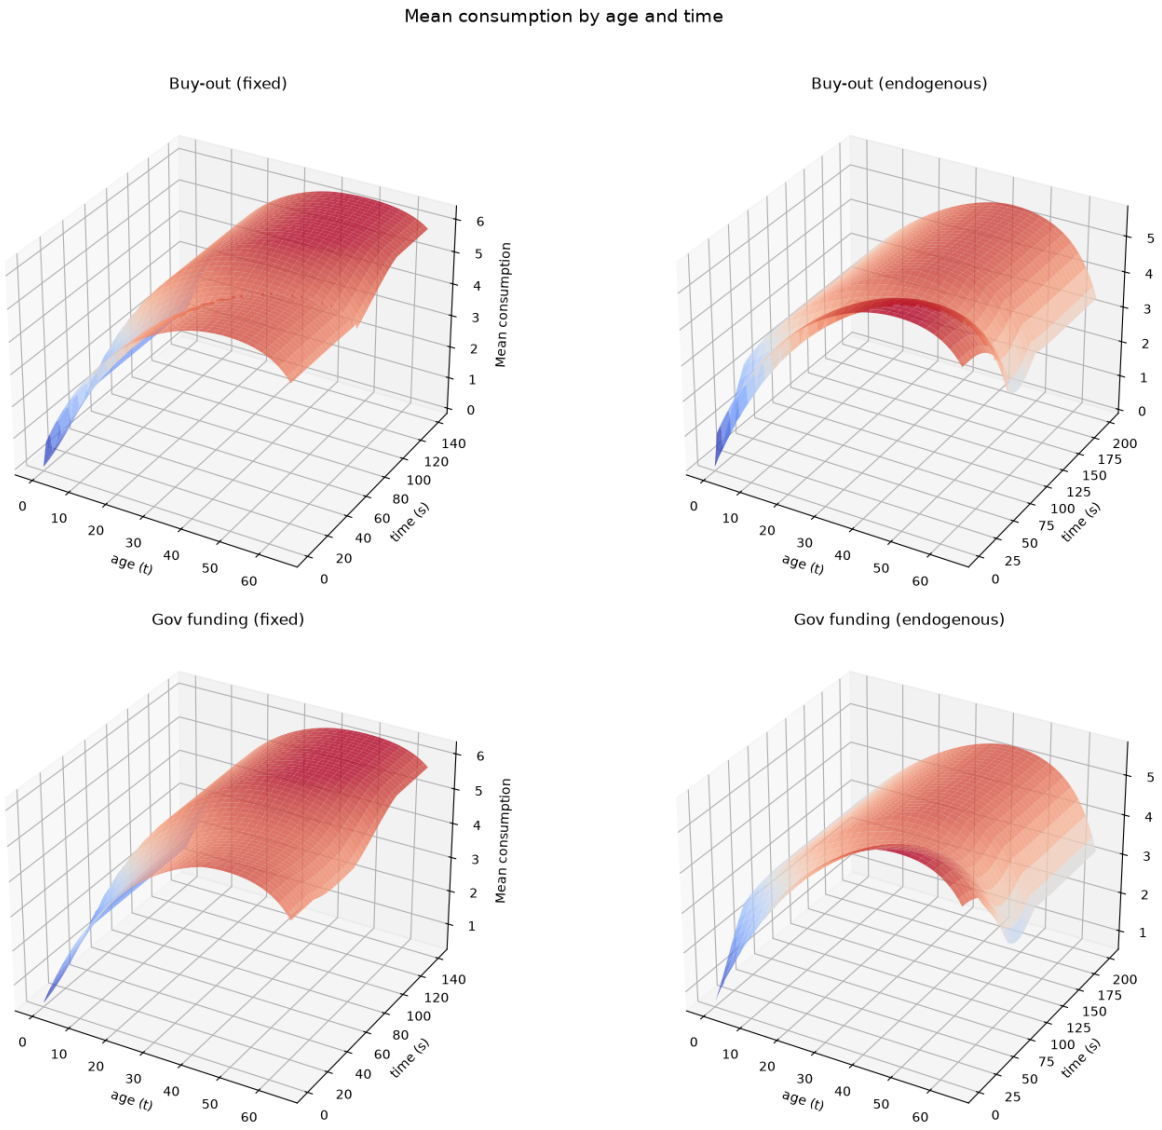
The asset variance surfaces show how cumulative income shocks cause dispersion to increase with age, with the transition potentially altering the rate of dispersion growth.

```
plot_surface_grid(
    [Va_bf, Va_be, Va_af, Va_ae],
    case_names, 'Std. dev. of assets',
    'Asset std. dev. by age and time',
    cmap='plasma', transform=np.sqrt
)
```



The mean consumption surfaces reflect the optimal consumption path, which should be smooth across ages due to the permanent income hypothesis underlying the model.

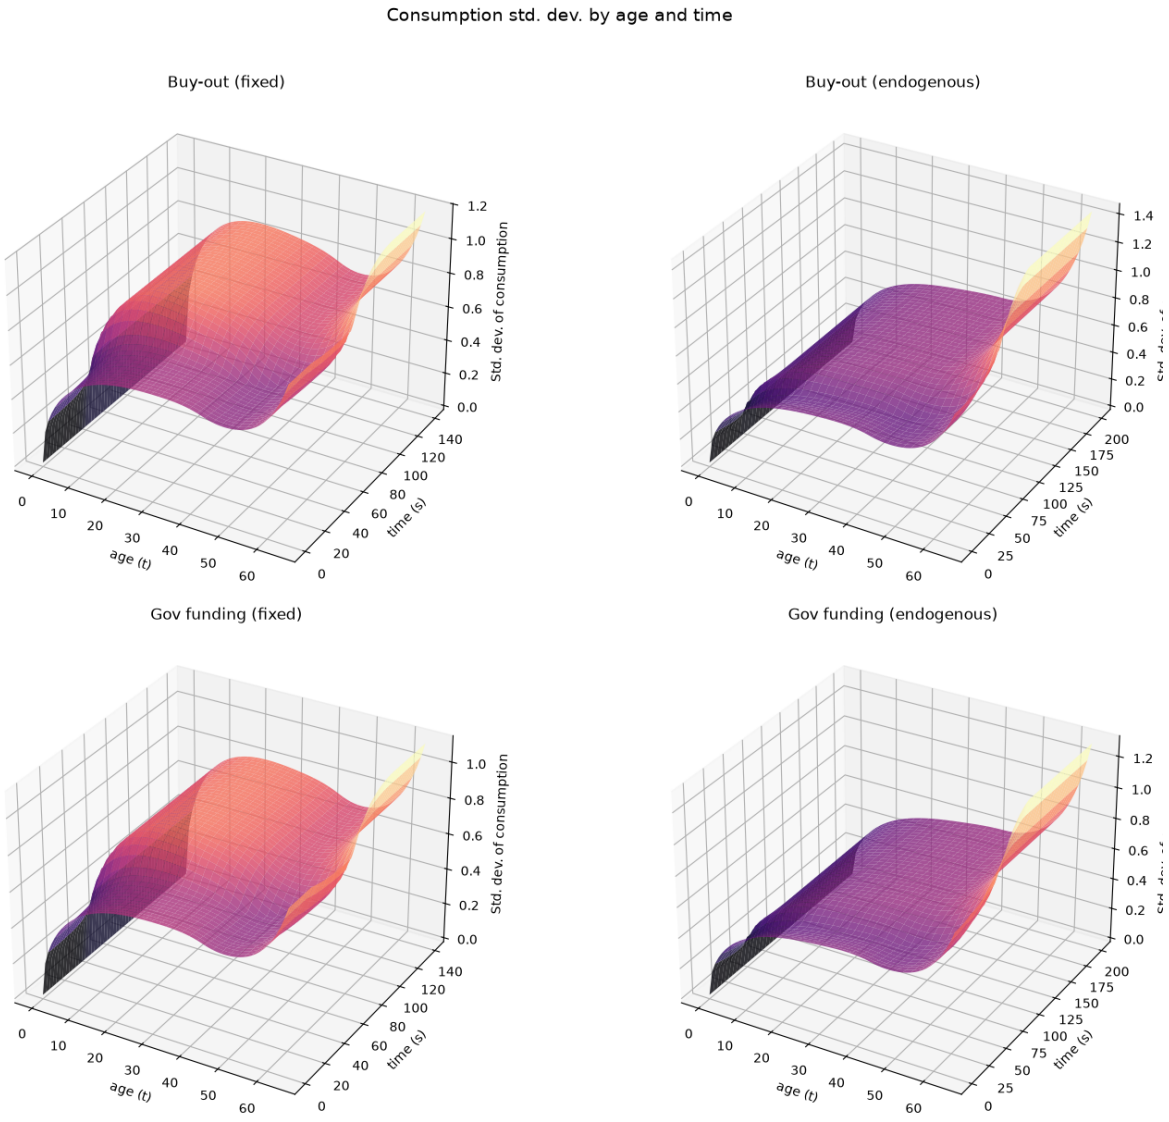
```
plot_surface_grid(
    [μc_bf, μc_be, μc_af, μc_ae],
    case_names, 'Mean consumption',
    'Mean consumption by age and time',
    cmap='coolwarm'
)
```



The consumption variance surfaces show how within-cohort consumption dispersion builds up with age as income shocks accumulate, and how it is reshaped along the transition.

Note that consumption dispersion is not a by-product of certainty equivalence: with $\sigma = -0.05$ the decision rules themselves depend on the volatility of income, as Exercise 127.12.2 confirms.

```
plot_surface_grid(
    [Vc_bf, Vc_be, Vc_af, Vc_ae],
    case_names, 'Std. dev. of consumption',
    'Consumption std. dev. by age and time',
    cmap='magma', transform=np.sqrt
)
```



127.11 Concluding remarks

We set out to compare two ways of moving from an unfunded to a fully funded social security system, following Huang *et al.* [1997].

Both reforms raise the capital stock substantially, lifting capital relative to GDP from about 3.2 to about 4.2 when factor prices are fixed, and to about 3.9 when they are endogenous.

Both also cut the equilibrium labor tax rate, from 0.34 to between 0.08 and 0.23 depending on the reform and the price regime.

They differ in who bears the cost and in what is preserved.

The buy-out of Experiment 1 settles the government's obligations to existing generations by issuing debt, which keeps those generations whole but leaves future ones without social insurance.

The government funding of Experiment 2 keeps paying benefits and accumulates assets instead, which requires higher taxes during the transition but preserves insurance against life span risk and against labor income risk.

A newborn in the terminal steady state prefers the second reform.

Three features of the model do most of the work in these calculations.

Risk-sensitive preferences make saving respond to the volatility of labor income, so that the precautionary motive, rather than the life-cycle motive alone, shapes the aggregate capital stock.

Life span uncertainty gives social security an insurance role that a pure pay-as-you-go transfer scheme would not have.

And the linearity of the decision rules lets us track entire distributions of consumption and wealth, both across cohorts and over calendar time, which is what the surfaces plotted above display.

Readers who want the same questions posed in simpler settings can compare *Transitions in an Overlapping Generations Model*, where transitions can be computed by hand, and *A Long-Lived, Heterogeneous Agents, Overlapping Generations Model*, where precautionary saving arises from a borrowing constraint rather than from risk-sensitive preferences.

127.12 Exercises

Exercise 127.12.1

The calibration table lists $\tilde{A} = 2.2625$ and $\tilde{\alpha} = 0.40$ for the closed economy, and $r = 0.1275$ and $w = 5.0147$ for the small open economy.

This exercise shows that the first pair is implied by the second.

1. Let $K_{\text{eff}} = \bar{k}/n + f_0 k_{-1}$ be the capital available for production per capita in the initial steady state, and $\bar{\varepsilon} = \sum_t \varepsilon_t f_t$ aggregate labor in efficiency units. Explain why $\tilde{\alpha} = rK_{\text{eff}}/(rK_{\text{eff}} + w\bar{\varepsilon})$.
2. Given $\tilde{\alpha}$, compute $\tilde{A}$ from the marginal product of capital, and then again from the marginal product of labor.
3. Suppose instead that we wanted a capital share of $\tilde{\alpha} = 0.33$ while keeping the same (r, w) targets. What goes wrong?

Solution

Part 1. With a Cobb-Douglas technology and constant returns to scale, payments to capital and labor exhaust output, so $Y = rK + wL$.

Capital's share of income is then $rK/Y = rK/(rK + wL)$, and for Cobb-Douglas production that share equals the exponent $\tilde{\alpha}$.

Parts 2 and 3.

```
ε_bar_ex = float(jnp.sum(hh.frac * hh.ε_arr))
K_eff_ex = ss0.k_bar / hh.n + float(hh.frac[0] * hh.x0[0])
k_pe_ex = K_eff_ex / ε_bar_ex

α_implied = tech.r * K_eff_ex / (tech.r * K_eff_ex + tech.w * ε_bar_ex)
print(f"K_eff = {K_eff_ex:.4f},  ε̄ = {ε_bar_ex:.4f},  K/L = {k_pe_ex:.4f}")
print(f"implied capital share = {α_implied:.4f}  (calibrated {tech.α_tilde})")

for α_try in [tech.α_tilde, 0.33]:
    A_from_r = tech.r / (α_try * k_pe_ex ** (α_try - 1))
    A_from_w = tech.w / ((1 - α_try) * k_pe_ex ** α_try)
    print(f"α̃ = {α_try:.2f}: A from r = {A_from_r:.4f}, "
          f"A from w = {A_from_w:.4f}")
```

```
K_eff = 21.4550,  ᾱ = 0.8175,  K/L = 26.2433
implied capital share = 0.4002  (calibrated 0.4)
α̃ = 0.40: A from r = 2.2639, A from w = 2.2620
α̃ = 0.33: A from r = 3.4494, A from w = 2.5462
```

At $\tilde{\alpha} = 0.40$ the two values of $\tilde{A}$ agree with each other and with the calibrated 2.2625.

At $\tilde{\alpha} = 0.33$ they disagree: no single scale factor can reproduce both prices, because the capital-labor ratio of the initial steady state and the two price targets already determine the capital share.

Pinning down (r, w) and the steady-state capital stock therefore leaves no freedom in $(\tilde{A}, \tilde{\alpha})$, which is exactly why the small open economy and the closed economy start from the same place.

Exercise 127.12.2

How much saving does risk sensitivity induce?

Recompute the initial steady state for $\sigma \in \{0, -0.05, -0.1\}$, holding all other parameters at their calibrated values, and report the per capita capital stock, the equilibrium labor tax rate, and capital relative to GDP.

Explain why the case $\sigma = 0$ deserves the name certainty equivalence.

Solution

```
print(f"{'σ':>7}  {'k_bar':>8}  {'τ_l':>7}  {'K/GDP':>7}")
for σ_try in [0.0, -0.05, -0.1]:
    hh_try = create_household(σ=σ_try)
    ss_try = find_ss_exo(
        (tech.RR, tech.w),
        (τ_a_0, τ_0_0, benef_0, G_0, Gb_0),
        hh_try, tech
    )
    print(f"{σ_try:>7.2f}  {ss_try.k_bar:>8.4f}  {ss_try.τ_l:>7.4f}"
          f"  {ss_try.k2gdp:>7.4f}")
```

```
      σ      k_bar       τ_l     K/GDP
   0.00    16.1981    0.3787    2.6324
  -0.05    21.6096    0.3383    3.1615
  -0.10    28.1911    0.2906    3.6782
```

Moving from $\sigma = 0$ to the calibrated $\sigma = -0.05$ raises the capital stock by about a third, and a further doubling of risk sensitivity raises it by another third.

The extra capital widens the tax base, so the labor tax rate that balances the same budget falls by about four percentage points between the first two rows.

At $\sigma = 0$ the operator $\mathcal{T}_t$ is the identity, so P_t and the decision rule F_t never see the shock loading C_t.

Consumption then responds only to the conditional mean of future income, which is the certainty equivalence property of ordinary linear-quadratic problems.

At $\sigma < 0$ the term $\sigma P C_t (I - \sigma C_t' P C_t)^{-1} C_t' P$ makes F_t depend on $C_t C_t'$, and the consumer accumulates a buffer against income risk.

Exercise 127.12.3

This exercise examines the debt target behind Experiment 2 and the long-run welfare ranking of the two reforms.

1. Compute the per capita benefit bill $\sum_{t>T_1} S_t f_t$ and the asset position $\bar{b}$ at which the government's net asset income $[R/n-1](-\bar{b})$ exactly pays for it.
2. Compare that number with the asset position that the lecture's target of -1.1785 for debt relative to GDP actually delivers, and report the net asset income and the labor tax rate in that steady state.
3. Compute the welfare of a newborn in the initial steady state and in the two terminal steady states under fixed prices. Which reform does a newborn prefer, and why?

Solution

```
benefit_bill_ex = float(jnp.sum(benef_0 * hh.frac * (1.0 - hh.Ind_work)))
Rn_ex = tech.RR / hh.n - 1.0

print(f"benefit bill               = {benefit_bill_ex:.4f}")
print(f"R/n - 1                    = {Rn_ex:.4f}")
print(f"b_bar that just funds it   = {-benefit_bill_ex / Rn_ex:.4f}\n")

ss_exp2 = exp2_exo['ss1']
print(f"lecture target -1.1785 gives b_bar = {ss_exp2.Gb:.4f}")
print(f"  net asset income = {-Rn_ex * ss_exp2.Gb:.4f}"
      f"  versus benefits = {benefit_bill_ex:.4f}")
print(f"  labor tax τ_l = {ss_exp2.τ_l:.4f}"
      f"  versus {ss0.τ_l:.4f} initially")
```

```
benefit bill               = 0.4979
R/n - 1                    = 0.0548
b_bar that just funds it   = -9.0793

lecture target -1.1785 gives b_bar = -10.1573
  net asset income = 0.5570  versus benefits = 0.4979
  labor tax τ_l = 0.1390  versus 0.3383 initially
```

The asset position that exactly funds benefits is about -9.1, while the target used in the lecture delivers about -10.2.

The government therefore accumulates a little more than social security alone requires, and the surplus lets it cut the labor tax rate from 0.34 to about 0.14 even though it is still paying benefits and still buying goods.

```
print(f"{'steady state':>34}  {'U of a newborn':>14}")
for name, ss_w in [
    ('initial, pay-as-you-go', ss0),
    ('terminal, buy-out (no benefits)', exp1_exo['ss1']),
    ('terminal, government funding', exp2_exo['ss1']),
]:
    print(f"{name:>34}  {newborn_welfare(ss_w):>14.2f}")
```

```
                      steady state  U of a newborn
           initial, pay-as-you-go         -262.74
  terminal, buy-out (no benefits)         -192.18
     terminal, government funding         -180.88
```

A newborn ranks government funding first, the buy-out second, and the initial pay-as-you-go system last.

Because these are fixed-price steady states, the wage is the same in all three; what lifts welfare above the initial steady

state is the lower labor tax rate that both reforms deliver.

Government funding then does better than the buy-out, despite its higher labor tax, because it continues to pay benefits that insure against outliving one's assets and against low labor income draws.

Remember that this ranking compares steady states, so it leaves out the transition costs borne by generations alive at the time of the reform.

CHAPTER

EIGHT

SOCIAL SECURITY TRUST FUND DEPLETION

Contents

- *Social Security Trust Fund Depletion*
  - *Overview*
  - *The institution*
  - *A model we can solve quickly*
  - *Aggregation and the ledger*
  - *The economy before the news*
  - *Solving for the equilibrium payable path*
  - *Claiming and retirement along the transition*
  - *How this compares with the full model*
  - *Exercises*

128.1 Overview

The Social Security trustees project that the reserves of the Old-Age and Survivors Insurance trust fund will run out during 2032 [Board of Trustees, Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds, 2026].

The program cannot borrow.

Unless Congress acts, current law then reduces every benefit by a common factor $\phi_t \leq 1$ that makes the program's cash flow balance: a worker entitled to a scheduled benefit b^{sch} collects $\phi_t b^{sch}$.

We call ϕ_t the **payable factor**.

It is not a policy parameter that somebody chooses.

It is an equilibrium object, because the quantity it scales depends on decisions that workers take while looking ahead at it:

- a worker who expects a smaller pension works longer and saves more
- a worker inside the claiming window just before the cut has a reason to file early, while benefits are still paid in full

- a worker who expects to live longer has a reason to file late, because the actuarial reward for waiting is fixed by statute

Those decisions determine taxable payroll and the scheduled benefit bill, which are exactly the two quantities that determine ϕ_t.

This lecture studies that fixed point, following Sargent and Shi [2026].

We write down the ledger that current law prescribes, define the map that sends a path for the payable factor into another path for the payable factor, and compute its fixed point in a quantitative overlapping generations model with

- a life cycle from age 20 to age 100 with projected mortality
- permanent heterogeneity in earnings ability and in the taste for work
- the statutory benefit formula, the actuarial adjustment for the claim age, the taxable maximum, and the taxation of benefits
- retirement and benefit claiming as **separate** irreversible decisions

The exercises then put the machinery to work: on the oscillation that a decreasing map produces, on raising the full retirement age, on what it would cost to guarantee scheduled benefits instead, and on how much of the long-run cut is the work of rising life expectancy.

Several symbols recur, so we record them here.

Symbol	Meaning
ϕ_t	payable factor: the common fraction of scheduled benefits actually paid
F_t, r^F	trust fund balance and the rate credited on it
N_t	population; every other aggregate is per capita
Pay_t, B_t^{sch}	per capita taxable payroll and the per capita scheduled benefit bill
τ^{ss}, κ_b, $\bar{y}$	payroll tax rate, tax rate on benefits received, taxable maximum
AIME, PIA, ψ	the earnings record, the benefit it buys, and the rate at which the record accrues
AF(χ), FRA, ς	actuarial factor at claim age χ, the full retirement age, and a benefit scale
R, χ	a worker's retirement age and claim age
j, t	age index (0 at age 20) and calendar year
e, h	permanent earnings ability and taste for work
ℓ_j, Ω	survival from age 20 to age $20+j$, and $\sum_j \beta^j \ell_j$
$\tilde{R}$, a_0, $\mathcal{W}$	after-tax gross return, assets on hand, and lifetime resources

This lecture is a companion to *Two Computations to Fund Social Security*, which studies two ways of moving from an unfunded to a fully funded system in the model of Huang *et al.* [1997].

Here nobody moves to a funded system: we ask what current law does on its own, and what it would cost to stop it.

Two other lectures supply background: *Transitions in an Overlapping Generations Model* introduces the Auerbach and Kotlikoff [1987] overlapping generations framework, and *A Long-Lived, Heterogeneous Agents, Overlapping Generations Model* adds uninsurable idiosyncratic risk to it.

We use the following imports.

```
import numpy as np
import matplotlib.pyplot as plt
from collections import namedtuple
from scipy.optimize import fsolve, brentq
```

128.2 The institution

Social Security's ledger is a stock, F_t, and two flows.

The fund earns an administered rate r^F on the special issue Treasury securities it holds, it receives the payroll tax, and it pays benefits.

Writing N_t for the population, Pay_t for per capita taxable payroll, and B_t^{sch} for the per capita scheduled benefit bill, the ledger evolves as

$$F_{t+1} = (1 + r^F)F_t + N_t \left(Rev_t - \phi_t B_t^{sch} \right), \qquad F_t \geq 0, \tag{128.1}$$

where program revenue is the payroll tax plus the tax that returns to the program on the benefits it pays,

$$Rev_t = \tau^{ss} Pay_t + \kappa_b \phi_t B_t^{sch}. \tag{128.2}$$

The restriction $F_t \geq 0$ in (128.1) is the law: the program may not borrow.

It is what forces ϕ_t below one.

The program pays the largest common fraction of scheduled benefits that its resources allow,

$$\phi_t = \min \left\{ 1, \ \frac{(1 + r^F)F_t + N_t \tau^{ss} Pay_t}{N_t (1 - \kappa_b) B_t^{sch}} \right\}. \tag{128.3}$$

The numerator is what the program can draw on: the balance it carries in plus the payroll tax it collects.

The denominator is what it owes, net of the benefit tax that comes back to it.

We date depletion at the first year the fund is empty,

$$t^{dep} = \min\{t : F_t = 0\},$$

and we call $t^{dep} - 1$ the **first shortfall year**, the year in which the program pays out its last reserves and the first year in which $\phi_t < 1$.

> **Note**
>
> The trustees date "reserve depletion" to the year the reserves run out, which is the first shortfall year here.
>
> Their projected 2032 is therefore our $t^{dep} - 1$, and the first year in which the whole benefit is cut is 2033.

From t^{dep} on the fund is empty and (128.3) collapses to a cash flow condition,

$$\phi_t = \frac{\tau^{ss} Pay_t}{(1 - \kappa_b) B_t^{sch}}, \qquad t \geq t^{dep}. \tag{128.4}$$

128.2.1 The payable factor is a fixed point

Equation (128.4) looks like a formula for ϕ_t, but both Pay_t and B_t^{sch} are built from decisions that workers make while looking at the whole future path $\{\phi_s\}$.

Let Φ be the map that takes a candidate path $\{\phi_s\}_{s \geq 0}$, hands it to workers, collects the taxable payroll and scheduled benefits that their optimal decisions generate, runs the ledger (128.1), and returns the payable path (128.3) that the ledger implies.

An equilibrium payable path is a fixed point,

$$\{\phi_t\} = \Phi(\{\phi_t\}). \tag{128.5}$$

Important

Φ is a **decreasing** map.

Raise the candidate path and benefits become more valuable, so workers claim earlier and retire earlier; taxable payroll falls and the scheduled bill rises; (128.3) then returns a *smaller* number.

Successive guesses therefore alternate around the fixed point rather than approaching it from one side.

Whether the alternation dies out depends on how strongly workers respond, and Exercise 128.9.1 finds dates at which it does not: the undamped iteration settles into a two-cycle and stays there.

Every computation below therefore takes a damped step.

One consequence of (128.4) is worth recording before we compute anything.

In a long run in which the fund is empty and the economy has settled down, the payable factor is pinned by the program's **cost rate**, the ratio of scheduled benefits to taxable payroll:

$$\phi_\infty = \frac{\tau^{ss}}{(1-\kappa_b)} \cdot \frac{Pay_\infty}{B^{sch}_\infty}. \tag{128.6}$$

Everything demographic and behavioral that matters for the long-run benefit cut reaches it through that one ratio.

128.3 A model we can solve quickly

Sargent and Shi [2026] compute general equilibrium transition paths of an economy with idiosyncratic earnings risk, a neoclassical firm, a consolidated government budget constraint, and a cross-section of 80,000 workers per age.

That calculation takes about 22 hours on an A100 GPU.

The model below keeps the institution exactly and strips the economy down until the whole fixed point runs in about a minute on a laptop.

What this lecture leaves out

- **Prices are fixed.** The interest rate and the wage do not respond. In the paper, capital deepens as workers save against a smaller pension, the capital–output ratio rises from 3.00 to 3.19, and the interest rate falls from 6.0 to 5.3 percent.
- **There is no government budget constraint outside the program.** In the paper, depletion removes the Social Security deficit from the consolidated budget, so the labor tax *falls* from 0.220 to 0.190: current law delivers a benefit cut to retirees and a tax cut to workers in the same year.
- **There is no idiosyncratic earnings risk and no borrowing constraint**, so the life cycle has a closed form. Workers differ only in permanent ability and in their taste for work.
- **Annuity markets are complete**, so there are no accidental bequests and no bequest motive.
- **Benefits are scaled rather than extended.** The program's spousal and survivor benefits go to additional beneficiaries; here a single scale factor on each worker's own benefit matches the aggregate cost instead, which raises individual replacement rates by about a seventh.

- **Demography is stylized**, a smooth birth path and a Gompertz mortality table calibrated to the trustees' dependency ratios, rather than the published population projection.

What survives is the part this lecture is about: an endogenous payable factor, and claiming and retirement decisions that respond to it.

128.3.1 Parameters

Units are the average wage index of 2025, so one model unit is \$72,025 and mean earnings per worker are one.

The statutory amounts – the taxable maximum, the two bend points of the benefit formula, the payroll tax rate, the actuarial factors, the trust fund balance – take their legislated or recorded 2025 values.

The tax rates on labor, capital and consumption follow Kitao [2014], and the discount factor, the interest rate, the dispersion of earnings ability and the rate credited on the fund are those of Sargent and Shi [2026].

Two parameters are chosen inside this model, and we say which below.

```
AWI = 72025.07          # average wage index, 2025: dollars per model unit
AGE0, J = 20, 81        # ages 20 ... 100
Y0, T = 2026, 120       # first transition year, horizon
YEARS = np.arange(Y0, Y0+T+1)

Params = namedtuple('Params', [
    'β', 'η', 'r', 'τ_l', 'τ_k', 'τ_c', 'τ_ss', 'κ_b', 'r_F',
    'ycap', 'bend', 'prate', 'FRA', 'ψ', 'σ_a', 'σ_η', 'g_m', 'bscale', 'F_ratio'])

par = Params(
    β=0.9739,            # discount factor
    η=0.81,              # disutility of work (calibrated)
    r=0.060,             # interest rate
    τ_l=0.22, τ_k=0.30, τ_c=0.05,          # labor, capital, consumption taxes
    τ_ss=0.106,          # OASI payroll tax rate
    κ_b=0.03891,         # tax on benefits received
    r_F=0.02,            # rate credited on the trust fund
    ycap=176100/AWI,     # taxable maximum
    bend=np.array([1226.0, 7391.0])*12/AWI,    # PIA bend points
    prate=np.array([0.90, 0.32, 0.15]),        # PIA replacement rates
    FRA=67,              # full retirement age
    ψ=0.055,             # AIME accrual (a 35 year moving average)
    σ_a=0.315,           # dispersion of earnings ability
    σ_η=0.16,            # dispersion of the taste for work
    g_m=0.007,           # annual mortality improvement
    bscale=1.145,        # benefit scale ς (calibrated)
    F_ratio=1.6255)      # 2025 reserves relative to benefit payments

def after_tax_R(p):
    "The gross after-tax return on saving."
    return 1 + p.r*(1 - p.τ_k)
```

The benefit formula has two pieces.

The **primary insurance amount** is a concave, piecewise linear function of a worker's average indexed earnings, with marginal replacement rates 0.90, 0.32 and 0.15 on the three segments.

The **actuarial factor** $\mathrm{AF}(\chi)$ multiplies it, and depends on the age χ at which the worker files: the statutory rule reduces benefits by $5/9$ of one percent per month for the first 36 months of early claiming and $5/12$ of one percent per month beyond that, and adds a delayed retirement credit of 8 percent a year after the full retirement age.

```
def claim_factors(fra=67, lo=62, hi=70):
    "Statutory actuarial factors AF(χ) for claim ages lo ... hi."
    out = []
    for a in range(lo, hi+1):
        if a <= fra:
            m = (fra - a)*12
            out.append(round(1 - (min(m, 36)*(5/9) + max(m-36, 0)*(5/12))/100, 4))
        else:
            out.append(round(1 + 0.08*(a - fra), 4))
    return np.array(out)

def pia(aime, p):
    "The primary insurance amount: a concave piecewise-linear function of AIME."
    s0 = np.minimum(aime, p.bend[0])
    s1 = np.clip(aime - p.bend[0], 0.0, p.bend[1] - p.bend[0])
    s2 = np.maximum(aime - p.bend[1], 0.0)
    return p.prate[0]*s0 + p.prate[1]*s1 + p.prate[2]*s2

X_GRID = np.arange(62, 71)
print("claim age     ", X_GRID)
print("AF(χ) at FRA 67", claim_factors(67))
print("AF(χ) at FRA 69", claim_factors(69))
```

```
claim age      [62 63 64 65 66 67 68 69 70]
AF(χ) at FRA 67 [0.7    0.75   0.8    0.8667 0.9333 1.     1.08   1.16   1.24  ]
AF(χ) at FRA 69 [0.6    0.65   0.7    0.75   0.8    0.8667 0.9333 1.     1.08  ]
```

Fig. 128.1 shows the progressivity of the program on the left and its price of waiting on the right.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))

grid = np.linspace(0, 3, 400)
axes[0].plot(grid, pia(grid, par), lw=2)
for b in par.bend:
    axes[0].axvline(b, color='k', ls=':', lw=1)
axes[0].set_xlabel('AIME, multiples of the average wage')
axes[0].set_ylabel('primary insurance amount')
axes[0].set_title('the benefit formula, with its bend points')

axes[1].plot(X_GRID, claim_factors(67), 'o-', lw=2, label='full retirement age 67')
axes[1].plot(X_GRID, claim_factors(69), 's-', lw=2, label='full retirement age 69')
axes[1].axhline(1.0, color='k', lw=0.8)
axes[1].set_xlabel('claim age $\\chi$')
axes[1].set_ylabel('actuarial factor $\\mathrm{AF}(\\chi)$')
axes[1].set_title('the reward for waiting')
axes[1].legend()
plt.tight_layout()
plt.show()
```

A worker whose earnings record is below the first bend point replaces 90 cents of every dollar; one above the second replaces 15.

Waiting from 62 to 70 raises the benefit by 77 percent, and the schedule is statutory: it does not move when mortality improves, and it does not move when the payable factor falls.

Exercise 128.9.2 raises the full retirement age, which is the orange line.

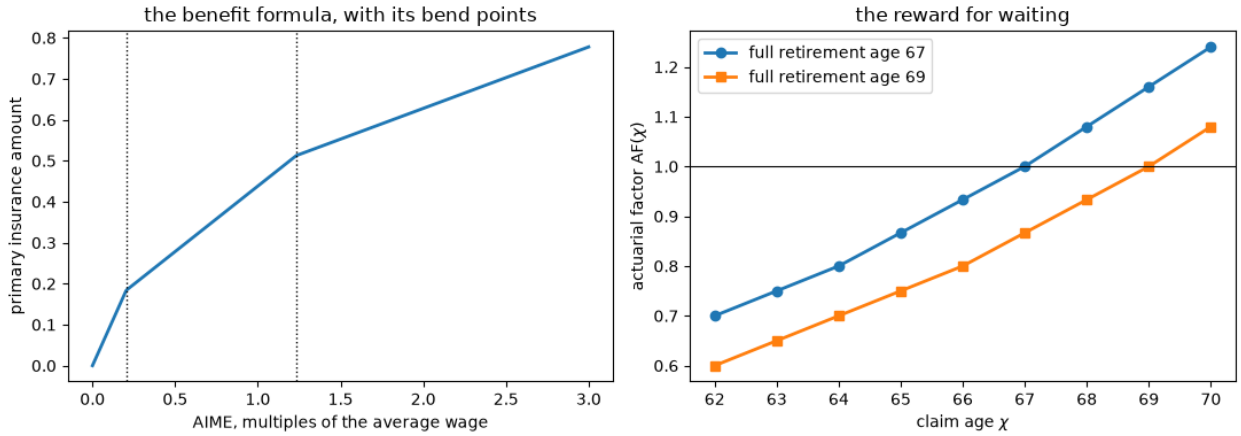


Fig. 128.1: The benefit formula is progressive and the reward for waiting is statutory

128.3.2 Demography

Mortality follows a Gompertz–Makeham force of mortality that improves at 0.7 percent a year until 2085 and is frozen afterwards, so that the economy settles into a stationary equilibrium inside the horizon.

Cohort sizes grow at one rate before 1990 and another after, and we pick the two rates so that the model reproduces the trustees' aged dependency ratio in 2025 and in 2085: the population aged 65 and over per person aged 20 to 64 rises from 0.32 to 0.50.

```
def mortality(age, year, p):
    "Force of mortality by age and calendar year; improvement stops in 2085."
    yr = np.minimum(year, 2085)
    return (5e-4 + 2.5e-5*np.exp(0.095*age))*np.exp(-p.g_m*(yr - 2025))

def cohort_survival(entry_years, p):
    "ℓ[c, j]: probability of reaching age 20+j, cohort by cohort."
    ent = np.asarray(entry_years, float).reshape(-1, 1)
    j = np.arange(J).reshape(1, -1)
    s = np.exp(-mortality(AGE0 + j, ent + j, p))
    s[:, -1] = 0.0
    return np.concatenate([np.ones((len(ent), 1)),
                           np.cumprod(s[:, :-1], axis=1)], axis=1)

YS = 1990           # the year the growth rate of entering cohorts changes

def entry_sizes(g, years):
    "Size of the cohort that turns 20 in each year."
    y = np.asarray(years, float)
    return np.where(y <= YS, np.exp(g[0]*(y - YS)), np.exp(g[1]*(y - YS)))

def population(g, years, p):
    "N[j, t]: people of age 20+j in each calendar year."
    years = np.atleast_1d(years)
    out = np.zeros((J, len(years)))
    for j in range(J):
        ent = years - j
        out[j] = entry_sizes(g, ent)*cohort_survival(ent, p)[:, j]
    return out
```

(continues on next page)

(continued from previous page)

```
def dependency(g, t, p):
    "People aged 65 and over per person aged 20 to 64."
    n = population(g, [t], p)[:, 0]
    return n[45:].sum()/n[:45].sum()

g_pop = fsolve(lambda g: [dependency(g, 2025, par) - 0.32,
                          dependency(g, 2085, par) - 0.50], [0.02, 0.0])

print(f"entry cohorts grow at {g_pop[0]:.4f} before {YS} and {g_pop[1]:.4f} after")
for t in (2025, 2035, 2055, 2085):
    print(f"   dependency ratio in {t}: {dependency(g_pop, t, par):.3f}")
print(f"   life expectancy at 20 of the cohort of 2026: "
      f"{20 + cohort_survival([2026], par)[0].sum():.1f}")
```

```
entry cohorts grow at 0.0103 before 1990 and -0.0020 after
   dependency ratio in 2025: 0.320
   dependency ratio in 2035: 0.370
   dependency ratio in 2055: 0.449
   dependency ratio in 2085: 0.500
   life expectancy at 20 of the cohort of 2026: 84.3
```

Fig. 128.2 is the engine of everything that follows, so it is worth looking at.

```
span = np.arange(2000, 2121)
fig, axes = plt.subplots(1, 2, figsize=(11, 4))

axes[0].plot(span, [dependency(g_pop, t, par) for t in span], lw=2)
axes[0].plot([2025, 2085], [0.32, 0.50], 'o', ms=7, color='C1',
             label='trustees')
axes[0].set_xlabel('year')
axes[0].set_ylabel('people 65 and over per person 20 to 64')
axes[0].legend()

for t, style in ((2026, '-'), (2085, '--')):
    n_t = population(g_pop, [t], par)[:, 0]
    axes[1].plot(AGE0 + np.arange(J), n_t/n_t.sum(), style, lw=2, label=str(t))
axes[1].set_xlabel('age')
axes[1].set_ylabel('share of the population')
axes[1].legend()
plt.tight_layout()
plt.show()
```

Two forces push the dependency ratio up, and the model carries both.

Cohorts entering the labor force stopped growing, so the base of the pyramid stops widening.

Mortality keeps improving, so its top keeps filling in.

Exercise 128.9.4 separates them.

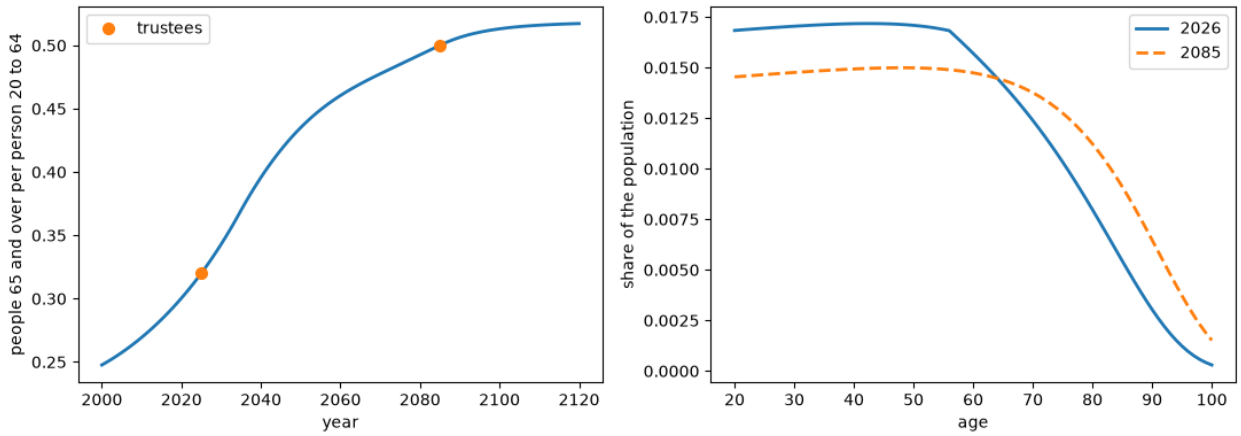


Fig. 128.2: Cohorts stopped growing and lives keep lengthening

128.3.3 Earnings and types

Earnings are the wage times a hump shaped age-efficiency profile that peaks at age 50, times a permanent ability draw.

Workers also differ in their taste for work.

Both are five point discretizations of lognormal distributions, so there are 25 types.

```
def efficiency():
    "Age-efficiency profile, normalized to mean one over ages 20 to 64."
    j = np.arange(J)
    e = np.exp(0.0583*j - 0.000972*j**2)
    return e/e[:45].mean()

ε = efficiency()

def types(p, n=5):
    "Permanent earnings ability and taste for work, on a tensor grid."
    x, w = np.polynomial.hermite_e.hermegauss(n)
    w = w/w.sum()
    e = np.exp(p.σ_a*x); e = e/(w*e).sum()
    h = np.exp(p.σ_η*x)
    E, H = np.meshgrid(e, h, indexing='ij')
    return E.ravel(), H.ravel(), np.outer(w, w).ravel()
```

Earnings rise with experience, peak at age 50, and decline thereafter.

```
fig, ax = plt.subplots(figsize=(7, 3.6))
ax.plot(AGE0 + np.arange(J), ε, lw=2)
ax.set_xlabel('age')
ax.set_ylabel('efficiency units of labor')
ax.set_xlim(20, 80)
plt.tight_layout()
plt.show()
```

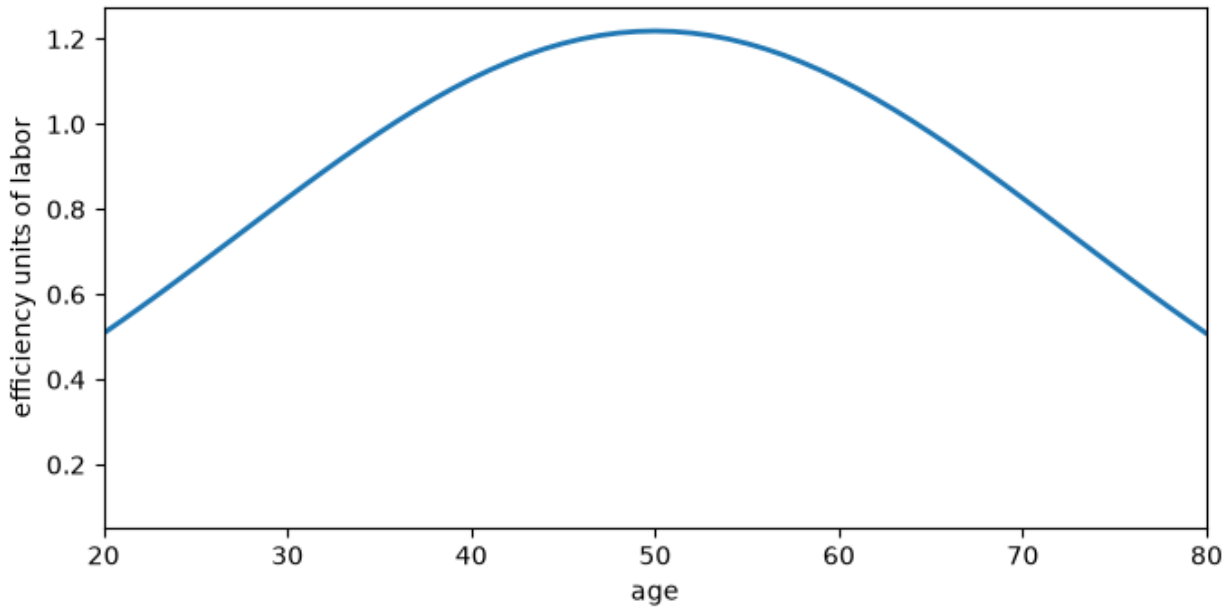


Fig. 128.3: Earnings peak at age 50 and decline thereafter

128.3.4 The worker's problem

A worker chooses when to stop working and when to file for benefits.

Both decisions are irreversible, and they are separate: Coile *et al.* [2002] and the literature that follows them establish that leaving the labor force and claiming a benefit are distinct events.

Write R for the retirement age, χ for the claim age, and ℓ_j for the probability of surviving from age 20 to age $20 + j$.

A worker of type (e, h) ranks plans by

$$\sum_j \beta^j \ell_j [\log c_j - \eta\, h\, n_j], \qquad n_j = \mathbf{1}\{j < R - 20\}, \tag{128.7}$$

where e is permanent earnings ability, h is the taste for work, and ηh is therefore the worker's own disutility of a year of work.

We normalize the wage to one.

While working, a worker earns $y_j = \varepsilon_j e$, pays the labor tax and the payroll tax on earnings up to the taxable maximum $\bar{y}$, and accumulates the earnings record

$$\text{AIME}_{j+1} = \text{AIME}_j + \psi\, \mathbf{1}\{j < R - 20\}\big(\min(y_j, \bar{y}) - \text{AIME}_j\big),$$

an exponentially weighted moving average whose effective window of $2/\psi - 1 \approx 35$ years approximates the statutory highest-35-years rule.

Once filed, the worker collects

$$b_j = \varsigma\, \phi_t\, \text{AF}(\chi)\, \text{PIA}(\text{AIME}_j),$$

taxed at κ_b, where ς is a benefit scale that the calibration below sets.

It stands in for the spousal and survivor benefits that a model of single workers leaves out.

With complete annuity markets and log utility the consumption problem has a closed form.

Annuities pay the gross return $\tilde{R}/s_j$ to a survivor, where $\tilde{R} = 1 + r(1-\tau_k)$ and s_j is the one-period survival probability, because the balances of those who die are shared among those who live.

Compounding those returns turns survival into a discount factor: a unit of income received at age $20+j$ is worth $\ell_j \tilde{R}^{-j}$ at age 20, and the lifetime budget constraint is

$$(1+\tau_c)\sum_j \ell_j \tilde{R}^{-j} c_j = a_0 + \sum_j \ell_j \tilde{R}^{-j} m_j(R,\chi) \equiv \mathcal{W}(R,\chi), \tag{128.8}$$

where m_j is net income at age j: earnings net of the two taxes while working, and the payable benefit net of its own tax once filed.

Because the annuity return and the survival probability cancel, the Euler equation is the one from a model without mortality,

$$c_j = c_0(\beta\tilde{R})^j.$$

Substituting it into (128.8) gives $(1+\tau_c)c_0\Omega = \mathcal{W}$ with $\Omega = \sum_j \beta^j \ell_j$, and substituting it into (128.7) gives

$$\sum_j \beta^j \ell_j \log c_j = \Omega \log c_0 + \log(\beta\tilde{R}) \sum_j j\,\beta^j \ell_j.$$

The second term does not depend on the plan, and neither does the $\log[(1+\tau_c)\Omega]$ that comes out of $c_0 = \mathcal{W}/[(1+\tau_c)\Omega]$.

Dropping both leaves a two-dimensional discrete choice:

$$\max_{R,\,\chi} \;\; \Omega \log \mathcal{W}(R,\chi) \;-\; \eta\, h \sum_{j<R-20} \beta^j \ell_j. \tag{128.9}$$

The first term is the present value of everything the plan delivers, including the whole future path of the payable factor.

The second is the cost of the extra years of work.

A worker who is already alive when the news arrives solves the same problem from where it stands.

Write j_0 for its age index in 2026.

Then both sums in (128.9) run from j_0 rather than from 0, with ℓ_j and β^j measured from j_0; a_0 in (128.8) is the wealth it has accumulated by then; and the decisions it has already taken are not up for reconsideration.

That is what the arguments `j0`, `a0` and `plan0` carry in the code below.

Nothing there is special to 2026: Exercise 128.9.5 uses the same code to let workers re-plan in the year a surprise arrives.

Note

The paper also models the **retirement earnings test**, which withholds the benefits of a claimant who keeps working before the full retirement age and then credits them back by advancing the effective claim age.

For a worker who leaves work at or before the full retirement age the credit matches the withholding one for one, so filing earlier than retirement is payoff neutral in this deterministic model.

We therefore impose $\chi \geq \min(R, 70)$ and drop the test.

Claiming remains a separate decision: a worker may file at retirement or wait, up to age 70.

The restriction does bind on one margin.

A worker who means to work past the full retirement age could file at 62, have the benefits withheld until the test stops applying, and then collect at the credited factor while still working.

That option is not quite equivalent to filing later, and this model rules it out.

We enumerate the admissible plans and index every cohort alive along the transition.

```
R_GRID = np.arange(58, 73)
# a worker files no earlier than it stops working, and no later than 70
PLANS = np.array([(R, X) for R in R_GRID for X in X_GRID if X >= min(R, 70)])
R_P, X_P = PLANS[:, 0], PLANS[:, 1]
NP = len(PLANS)

ENTRY = np.arange(Y0-J+1, Y0+T+1)       # entry year of every cohort ever alive
NC = len(ENTRY)
ENT_IX = {y: i for i, y in enumerate(ENTRY)}
IDX = np.array([[ENT_IX[t-j] for t in YEARS] for j in range(J)])

print(f"{NP} admissible plans, {NC} cohorts")
```

```
83 admissible plans, 201 cohorts
```

The retirement grid runs from 58 to 72, and it binds for the two most extreme tastes for work, which carry 1.1 percent of the population each.

Widening it to 50 to 80 moves those two types to 56 and 73 and moves nothing else: the depletion year, both mean ages, and the payable factor to four decimals are the same.

We keep the narrow grid, which is about twice as fast.

Everything that does not depend on the payable path is built once.

```
class World:
    "Demography, types, survival and discounting: fixed along the iteration."

    def __init__(self, p, g):
        self.p = p
        e, h, w = types(p)
        self.NE = len(e)
        self.C = NC*self.NE
        self.entry = np.repeat(ENTRY, self.NE)
        self.e = np.tile(e, NC)
        self.η = np.tile(h, NC)
        self.w = np.tile(w, NC)
        self.j0 = np.clip(Y0 - self.entry, 0, J-1)       # age index in 2026
        self.ell = np.repeat(cohort_survival(ENTRY, p), self.NE, axis=0)
        self.disc = self.ell/after_tax_R(p)**np.arange(J)
        self.AF = claim_factors(p.FRA)
        bet_ell = (p.β**np.arange(J))*self.ell
        self.tail = np.cumsum(bet_ell[:, ::-1], axis=1)[:, ::-1]
        self.pre = np.concatenate([np.zeros((self.C, 1)),
                                   np.cumsum(bet_ell, axis=1)], axis=1)
        self.N = population(g, YEARS, p)
        self.Npop = self.N.sum(axis=0)
        self.share = self.N/self.Npop

def phi_of(ψ, year):
    "The payable factor: one before 2026, constant beyond the horizon."
    k = np.clip(year - Y0, 0, len(ψ)-1)
    return np.where(year < Y0, 1.0, np.asarray(ψ)[k])
```

The life cycle itself is one loop over ages.

```
def lifecycle(W, ψ, Rj, Xa, j0, disc, keep=False):
    """Discounted net income from age index j0 on, for plans (Rj, Xa).
```

(continues on next page)

(continued from previous page)

```
    With keep=True also returns taxable payroll and scheduled benefits at
    every age, which is what aggregation needs.
    """
    p = W.p
    nd = np.ndim(Rj)
    e_row = W.e.reshape(-1, *[1]*(nd-1))
    ent = W.entry.reshape(-1, *[1]*(nd-1))
    shape = np.broadcast(Rj, Xa, e_row).shape
    aime = np.zeros(shape)
    S = np.zeros(shape)
    AFx = np.interp(Xa, X_GRID, W.AF)*np.ones(shape)
    pay, ben = [], []
    for j in range(J):
        work = (j < Rj)*np.ones(shape)
        y = ε[j]*e_row*work
        capped = np.minimum(y, p.ycap)*work
        b = p.bscale*AFx*pia(aime, p)*((AGE0 + j >= Xa)*np.ones(shape))
        inc = ((1-p.τ_l)*y - p.τ_ss*capped
               + (1-p.κ_b)*phi_of(ψ, ent + j)*b)
        S = S + np.where(j >= j0, disc[..., j]*inc, 0.0)
        if keep:
            pay.append(capped)
            ben.append(b)
        aime = aime + p.ψ*work*(capped - aime)
    return (S, np.array(pay), np.array(ben)) if keep else S

def solve_plans(W, ψ, j0, plan0=None, a0=0.0, at_year=Y0):
    """The plan that maximizes the worker's objective, for every cohort and type.

    Workers re-plan in at_year, from the age index j0 they have reached, holding
    the wealth a0 they carry and the decisions in plan0 they cannot undo.
    """
    p, C = W.p, W.C
    j0 = np.asarray(j0)
    S = lifecycle(W, ψ, (R_P-AGE0).reshape(1, NP), X_P.reshape(1, NP),
                  j0.reshape(C, 1), W.disc.reshape(C, 1, J))
    scale = (after_tax_R(p)**j0)/W.ell[np.arange(C), j0]
    wealth = a0 + S*scale.reshape(C, 1)
    norm = (p.β**j0)*W.ell[np.arange(C), j0]
    Om = W.tail[np.arange(C), j0]/norm
    Dw = (W.pre[np.arange(C)[:, None], np.clip(R_P-AGE0, 0, J)[None, :]]
          - W.pre[np.arange(C), j0][:, None])/norm[:, None]
    V = (Om[:, None]*np.log(np.maximum(wealth, 1e-9))
         - (p.η*W.η)[:, None]*np.maximum(Dw, 0.0))
    if plan0 is not None:
        # decisions already executed by at_year cannot be undone
        R0, X0 = plan0
        age_now = np.asarray(j0) + AGE0
        alive = W.entry < at_year
        ok = np.where((alive & (R0 <= age_now))[:, None],
                      R_P[None, :] == R0[:, None],
                      R_P[None, :] >= np.where(alive, age_now, 0)[:, None])
        ok &= np.where((alive & (X0 <= age_now))[:, None],
                       X_P[None, :] == X0[:, None],
                       X_P[None, :] >= np.where(alive, np.maximum(62, age_now), 62)[:,
↪ None])
```

(continues on next page)

(continued from previous page)

```
        V = np.where(ok, V, -np.inf)
    k = V.argmax(axis=1)
    return R_P[k], X_P[k], wealth[np.arange(C), k], Om, Dw[np.arange(C), k]

def flows(W, ψ, R, X, j0):
    "Taxable payroll and scheduled benefits at the chosen plans, shape (J, C)."
    _, pay, ben = lifecycle(W, ψ, (R-AGE0), X, np.asarray(j0), W.disc, keep=True)
    return pay, ben
```

128.4 Aggregation and the ledger

Aggregates are population weighted averages across ages, and the operator Φ is the fund law (128.1) together with (128.3).

```
def aggregate(W, pay, ben):
    "Per capita taxable payroll and scheduled benefits in each calendar year."
    wt = W.w.reshape(NC, W.NE)[0]           # the type weights, tiled across cohorts
    Pj = pay.reshape(J, NC, W.NE) @ wt
    Bj = ben.reshape(J, NC, W.NE) @ wt
    Pay = (np.take_along_axis(Pj, IDX, axis=1)*W.share).sum(axis=0)
    Ben = (np.take_along_axis(Bj, IDX, axis=1)*W.share).sum(axis=0)
    return Pay, Ben

def fund_and_phi(W, Pay, Ben, F0):
    "The ledger and the payable factor: one pass of the operator Φ."
    p = W.p
    F = np.zeros(len(YEARS)+1)
    F[0] = F0
    ψ = np.ones(len(YEARS))
    for t in range(len(YEARS)):
        obligations = W.Npop[t]*(1-p.κ_b)*Ben[t]
        resources = (1+p.r_F)*F[t] + W.Npop[t]*p.τ_ss*Pay[t]
        ψ[t] = min(1.0, resources/obligations) if obligations > 0 else 1.0
        rev = p.τ_ss*Pay[t] + p.κ_b*ψ[t]*Ben[t]
        F[t+1] = max((1+p.r_F)*F[t] + W.Npop[t]*(rev - ψ[t]*Ben[t]), 0.0)
    return F, ψ
```

128.5 The economy before the news

Workers alive in 2026 made their plans expecting benefits to be paid in full forever.

We solve that economy first.

It supplies two things that the transition needs: the retirement and claim decisions that cohorts have **already executed** by 2026, and the assets they carry into 2026.

It also supplies the scale for the third, the opening trust fund balance, which we set below to the recorded end-of-2025 reserves of 1.63 times annual benefit payments.

```
def income_path(W, ψ, R, X, rows=slice(None)):
    "Net income at every age at the chosen plans, for all rows or a subset."
    p = W.p
    e, entry, R, X = W.e[rows], W.entry[rows], R[rows], X[rows]
```

(continues on next page)

(continued from previous page)

```
    aime = np.zeros(len(e))
    out = np.zeros((J, len(e)))
    AFx = np.interp(X, X_GRID, W.AF)
    for j in range(J):
        work = (j < R - AGE0).astype(float)
        y = ε[j]*e*work
        capped = np.minimum(y, p.ycap)*work
        b = p.bscale*AFx*pia(aime, p)*((AGE0 + j >= X).astype(float))
        out[j] = ((1-p.τ_l)*y - p.τ_ss*capped
                  + (1-p.κ_b)*phi_of(ψ, entry + j)*b)
        aime = aime + p.ψ*work*(capped - aime)
    return out

def pre_transition(W):
    "Plans, assets at 2026, and aggregates when ψ ≡ 1 forever."
    p = W.p
    ψ1 = np.ones(T+1)
    z = np.zeros(W.C, int)
    R, X, wealth, Om, _ = solve_plans(W, ψ1, z)
    pay, ben = flows(W, ψ1, R, X, z)
    Pay, Ben = aggregate(W, pay, ben)
    c0 = wealth/((1+p.τ_c)*Om)
    inc = income_path(W, ψ1, R, X)
    a = np.zeros(W.C)
    A = np.zeros((J, W.C))
    for j in range(J-1):
        A[j] = a
        s = W.ell[:, j+1]/np.maximum(W.ell[:, j], 1e-12)
        a = (after_tax_R(p)/np.maximum(s, 1e-12))*(
            a + inc[j] - (1+p.τ_c)*c0*(p.β*after_tax_R(p))**j)
    A[J-1] = a
    return R, X, A[W.j0, np.arange(W.C)], Pay, Ben, wealth, Om

W = World(par, g_pop)
R_pre, X_pre, a26, Pay_pre, Ben_pre, wealth_pre, Om_pre = pre_transition(W)

mask = W.entry == 2026
wts = W.w[mask]/W.w[mask].sum()
print(f"cohort entering in 2026:  mean retirement age {(R_pre[mask]*wts).sum():.2f}, "
      f"mean claim age {(X_pre[mask]*wts).sum():.2f}")
print(f"cost rate in 2026: {Ben_pre[0]/Pay_pre[0]:.4f}"
      f"   (trustees for 2025: 0.1362)")
```

```
cohort entering in 2026:  mean retirement age 64.76, mean claim age 66.49
cost rate in 2026: 0.1350   (trustees for 2025: 0.1362)
```

Two parameters were chosen to make those two numbers come out.

The disutility of work η sets the retirement age: the cohort entering in 2026 retires at 64.8 on average, against the median retirement age of 64.5 that Sargent and Shi [2026] target.

The benefit scale ς sets the cost rate, and brings it within about a tenth of a percentage point of the trustees' 2025 figure.

The claim age is not targeted, and the model misses it by more: this cohort files at 66.5 on average, a year and a half later than the 65.0 the paper matches.

Claiming later than the data is a bias to keep in mind when reading the transition below, where the claim age is one of the two margins doing the work.

Two more statistics that nothing was fitted to are worth printing, because both bear on how hard the model pushes the two margins.

```
def replacement_rate(W, R, X):
    "The first benefit received, relative to earnings in the last year of work."
    p = W.p
    aime, last, first = (np.zeros(W.C) for _ in range(3))
    AFx = np.interp(X, X_GRID, W.AF)
    for j in range(J):
        work = (j < R - AGE0).astype(float)
        y = ε[j]*W.e*work
        capped = np.minimum(y, p.ycap)*work
        last = np.where(R - AGE0 - 1 == j, y, last)
        first = np.where(X - AGE0 == j, p.bscale*AFx*pia(aime, p), first)
        aime = aime + p.ψ*work*(capped - aime)
    return first/np.maximum(last, 1e-12)

rr = (replacement_rate(W, R_pre, X_pre)[mask]*wts).sum()
print(f"replacement rate, cohort entering in 2026:      {rr:.3f}")
print(f"   at the statutory formula, without ς:         {rr/par.bscale:.3f}")

shares = {int(a): float(W.w[mask][X_pre[mask] == a].sum()/W.w[mask].sum())
          for a in X_GRID}
print("   claim ages chosen:",
      {a: round(v, 3) for a, v in shares.items() if v > 0.005})
```

```
replacement rate, cohort entering in 2026:      0.492
   at the statutory formula, without ς:         0.430
   claim ages chosen: {66: 0.767, 68: 0.222, 70: 0.009}
```

The first number exposes something that the benefit scale ς does quietly.

Spousal and survivor benefits are paid to *additional people*; here the whole 14.5 percent lands on each worker's own benefit.

So a worker in this model collects about a seventh more than the statutory formula would pay it, and the incentives to retire and to file are correspondingly stronger.

Matching the program's aggregate cost with a model of single workers has to distort something, and this is where it goes.

The second number exposes the other limitation.

Claiming piles up on a single age, and nobody files at 62.

In the data, and in the paper, there are large spikes at 62 and at the full retirement age; the forces that produce the spike at 62 – liquidity constraints, health, the bad earnings draw that ends a career – are exactly the ones this version leaves out.

128.6 Solving for the equilibrium payable path

Now give workers the news.

Starting from $\phi_t \equiv 1$, we apply Φ, take a damped step toward the result, and repeat.

The step length starts at one half and anneals toward a floor, as in the solver of Sargent and Shi [2026].

```
def transition(W, F0, plan0, a0, iters=40, tol=1e-4,
               v0=0.5, anneal=0.92, floor=0.08, verbose=True):
    "Damped iteration on the payable path."
```

(continues on next page)

(continued from previous page)

```
    ψ = np.ones(T+1)
    for k in range(iters):
        R, X, _, _, _ = solve_plans(W, ψ, W.j0, plan0=plan0,
                                    a0=a0.reshape(W.C, 1))
        pay, ben = flows(W, ψ, R, X, W.j0)
        Pay, Ben = aggregate(W, pay, ben)
        F, ψ_new = fund_and_phi(W, Pay, Ben, F0)
        gap = np.abs(ψ_new - ψ).max()
        ν = max(ν0*anneal**k, floor)
        ψ = (1-ν)*ψ + ν*ψ_new
        if verbose and k % 5 == 0:
            print(f"   iteration {k:>3}   step {ν:.3f}   gap {gap:.2e}"
                  f"   ψ in 2033 {ψ[2033-Y0]:.4f}")
        if gap < tol:
            break
    return dict(ψ=ψ, F=F, R=R, X=X, Pay=Pay, Ben=Ben)

F0 = par.F_ratio*W.Npop[0]*Ben_pre[0]
base = transition(W, F0, (R_pre, X_pre), a26)
```

```
   iteration   0   step 0.500   gap 5.23e-01   ψ in 2033 0.9725
   iteration   5   step 0.330   gap 4.31e-02   ψ in 2033 0.7717
   iteration  10   step 0.217   gap 1.51e-02   ψ in 2033 0.7578
   iteration  15   step 0.143   gap 1.72e-02   ψ in 2033 0.7560
   iteration  20   step 0.094   gap 1.18e-02   ψ in 2033 0.7555
   iteration  25   step 0.080   gap 1.44e-02   ψ in 2033 0.7553
   iteration  30   step 0.080   gap 1.41e-02   ψ in 2033 0.7551
   iteration  35   step 0.080   gap 1.48e-02   ψ in 2033 0.7550
```

The converged path gives the two dates and the size of the cut.

```
ψ, F = base['ψ'], base['F']
depletion = YEARS[np.argmax(F[:-1] <= 1e-12)]
shortfall = YEARS[np.argmax(ψ < 1)]
print(f"first shortfall year: {shortfall}")
print(f"depletion year:       {depletion}")
print(f"ψ in {shortfall}: {ψ[shortfall-Y0]:.3f},  in {depletion}: {ψ[depletion-Y0]:.
↪3f},"
      f"  in the long run: {ψ[-1]:.3f}")
```

```
first shortfall year: 2032
depletion year:       2033
ψ in 2032: 0.952,  in 2033: 0.755,  in the long run: 0.482
```

The dates are the trustees' dates, and the first full-year cut is close to theirs: they project that continuing income will cover about 78 percent of scheduled benefits in the first year the fund is empty.

Neither date was targeted.

The calibration fixed the cost rate of a single year, 2025, and took the opening balance from the record; the year in which that balance runs out is then whatever the cash flows of the intervening years make it, and those depend on the claiming and retirement decisions that the falling payable factor itself provokes.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))
axes[0].plot(YEARS, F[:-1]/(W.Npop*base['Ben']), lw=2)
axes[0].set_title('trust fund reserves / annual scheduled benefits')
```

(continues on next page)

(continued from previous page)

```
axes[1].plot(YEARS, ψ, lw=2)
axes[1].set_title('payable factor $\\phi_t$')
axes[1].set_ylim(0, 1.05)
for ax in axes:
    ax.axvline(depletion, color='k', ls='--', lw=1, alpha=0.7)
    ax.set_xlabel('year')
    ax.set_xlim(Y0, 2120)
plt.tight_layout()
plt.show()
```

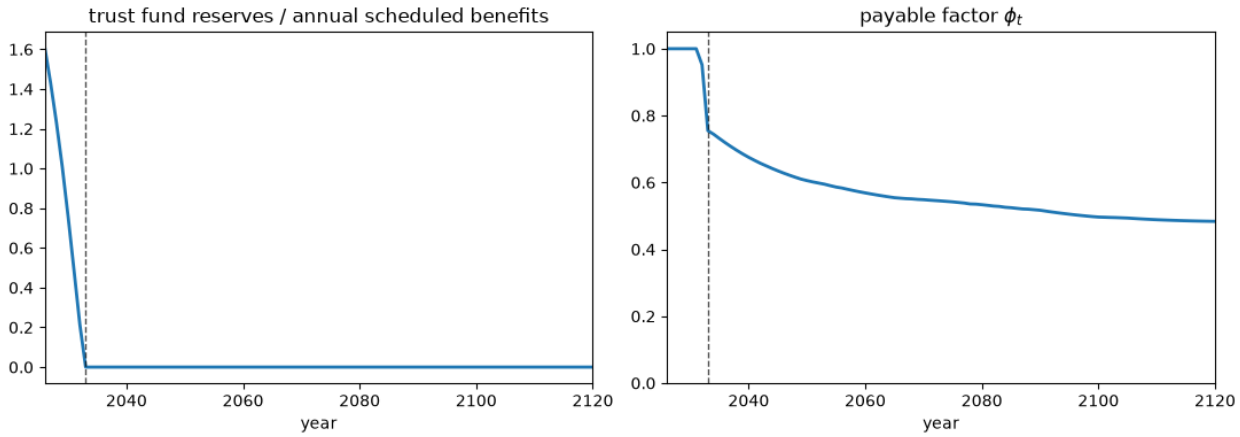


Fig. 128.4: The reserves run out in 2032 and the payable factor keeps falling after the cut

The payable factor does not settle at the level it reaches on the day the fund empties.

It falls further for the rest of the century, because the dependency ratio keeps rising.

By (128.6) the long-run cut is governed entirely by the program's cost rate, and we can check that identity on the computed path.

```
cost = base['Ben']/base['Pay']
print(f"terminal cost rate: {cost[-1]:.4f}")
print(f"τ_ss/[(1-κ_b)·cost rate] = {par.τ_ss/((1-par.κ_b)*cost[-1]):.4f}"
      f"   against ψ_∞ = {ψ[-1]:.4f}")
```

```
terminal cost rate: 0.2291
τ_ss/[(1-κ_b)·cost rate] = 0.4814   against ψ_∞ = 0.4816
```

128.7 Claiming and retirement along the transition

Both decisions respond to the path of ϕ_t, and both reverse direction at the cut.

```
def flow_ages(W, R, X):
    "Population weighted mean age of new retirees and of new claimants."
    wt = W.w.reshape(NC, W.NE)[0]
    age = (AGE0 + np.arange(J)).reshape(J, 1, 1)
```

(continues on next page)

(continued from previous page)

```
    hitR = (R.reshape(1, NC, W.NE) == age)*wt
    hitX = (X.reshape(1, NC, W.NE) == age)*wt
    nR = np.take_along_axis(hitR.sum(axis=2), IDX, axis=1)*W.share
    nX = np.take_along_axis(hitX.sum(axis=2), IDX, axis=1)*W.share
    ages = (AGE0 + np.arange(J)).reshape(J, 1)
    return ((nR*ages).sum(axis=0)/np.maximum(nR.sum(axis=0), 1e-12),
            (nX*ages).sum(axis=0)/np.maximum(nX.sum(axis=0), 1e-12))

mean_R, mean_X = flow_ages(W, base['R'], base['X'])
```

The two ages move together, and both turn at the cut.

```
fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(YEARS, mean_R, lw=2, label='mean age of new retirees')
ax.plot(YEARS, mean_X, lw=2, label='mean age of new claimants')
ax.axvline(depletion, color='k', ls='--', lw=1, alpha=0.7)
ax.set_xlabel('year')
ax.set_ylabel('age')
ax.set_xlim(Y0, 2075)
ax.legend()
plt.tight_layout()
plt.show()
```

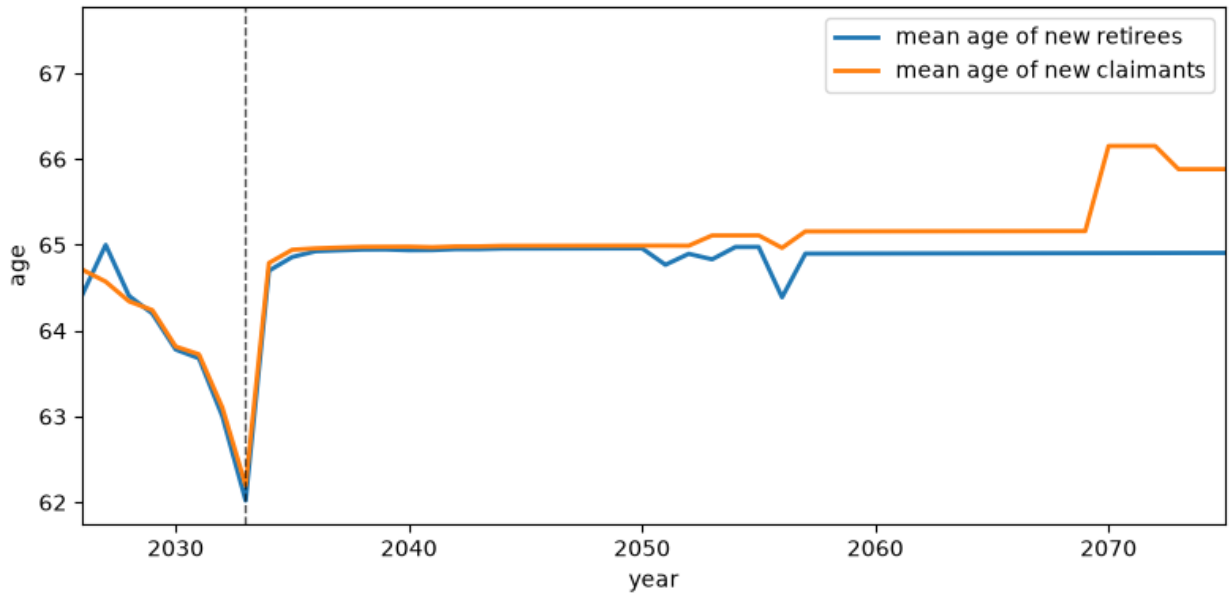


Fig. 128.5: Claiming and retirement rush forward before the cut and reverse after it

As the fund runs down there is a **rush**.

Both ages fall year after year from 2026, and the mean age of new claimants reaches 62.2 in the first year the fund is empty, against 64.7 at the start.

Workers inside the claiming window file as early as they can, because filing fixes the actuarial factor while the payments made in those years are still paid in full or nearly so.

A worker who waits instead collects under the reduced factors of later years.

The year after, the pattern reverses: there is nothing left to rush for, and both ages jump back above 64.5 and then drift up, to 65.3 and 66.6 in the long run.

The pension is smaller, so careers lengthen; and mortality keeps improving while the actuarial schedule stays at its statutory values, so a later claim buys the higher benefit for more years.

Past about 2075 the mean claim age chatters between two adjacent integers, which is why Fig. 128.5 stops there; the note on computation below explains why.

```
fig, axes = plt.subplots(1, 2, figsize=(11, 4))
axes[0].plot(YEARS, base['Pay'], lw=2, label='taxable payroll')
axes[0].plot(YEARS, base['Ben'], lw=2, label='scheduled benefits')
axes[0].legend()
axes[0].set_title('per capita, model units')
axes[1].plot(YEARS, cost, lw=2)
axes[1].set_title('cost rate: scheduled benefits / taxable payroll')
for ax in axes:
    ax.axvline(depletion, color='k', ls='--', lw=1, alpha=0.7)
    ax.set_xlabel('year')
    ax.set_xlim(Y0, 2120)
plt.tight_layout()
plt.show()
```

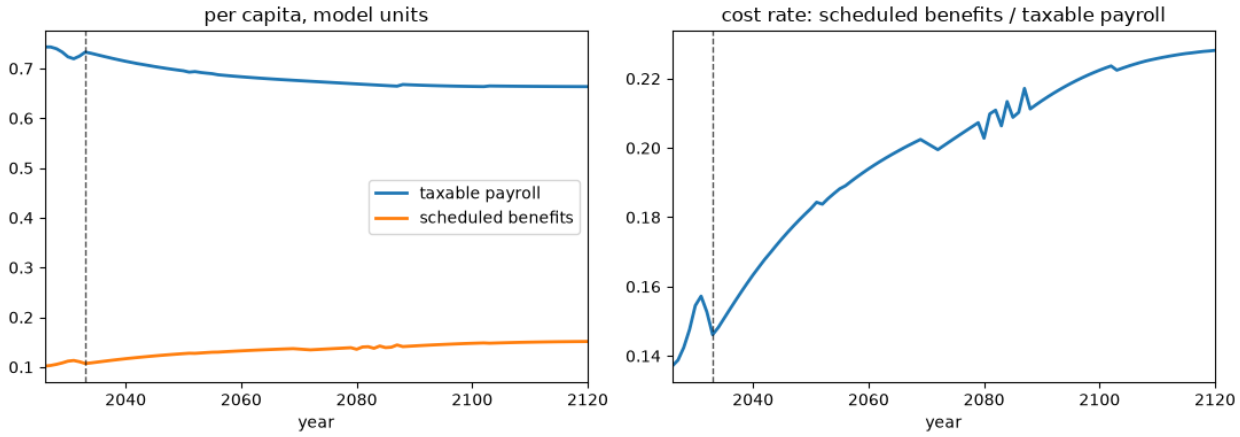


Fig. 128.6: Payroll gains arrive at once, the benefits of later claiming decades later

The benefit formula runs on two clocks, and the right panel of Fig. 128.6 shows both.

Longer careers add contributors to taxable payroll **immediately**.

Later filing lowers the scheduled bill immediately, by reducing the number of claimants, and raises it only **decades later**, when the delayed filers collect benefits carrying higher actuarial factors.

The cost rate therefore bumps up in the years before the cut, as the rush adds claimants, and then falls back for a year or two afterwards, as filing is deferred again, before resuming its climb.

Once the fund is empty, (128.4) makes the payable factor exactly proportional to the inverse of the cost rate, so that dip is a benefit cut that does not happen.

At the cost rate of 2031 the formula would pay $\phi = 0.702$; what workers' own responses deliver in 2033 is 0.755.

Sargent and Shi [2026] find the asymmetry strong enough that the payable factor recovers for a few years after the cut before resuming its decline.

Here the recovery shows up in the cost rate but not in ϕ_t itself: the dip is not large enough to outrun a dependency ratio that is still rising.

128.7.1 Who rushes

The flow means hide who is doing the rushing, so it is worth looking at the decisions one cohort at a time.

Fig. 128.7 plots, against a worker's age in 2026, the plan it had made before the news and the plan it chooses after.

```
ages26 = np.arange(21, 91)

def by_age26(vals):
    "Weighted mean of a decision across the types of each cohort alive in 2026."
    out = np.zeros(len(ages26))
    for i, a in enumerate(ages26):
        m = W.entry == Y0 - (a - AGE0)
        out[i] = (vals[m]*W.w[m]).sum()/W.w[m].sum()
    return out

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
for ax, before, after, title in (
        (axes[0], R_pre, base['R'], 'retirement age'),
        (axes[1], X_pre, base['X'], 'claim age')):
    ax.plot(ages26, by_age26(before), lw=2, label='planned before the news')
    ax.plot(ages26, by_age26(after), lw=2, ls='--', label='chosen after')
    ax.set_xlabel('age in 2026')
    ax.set_ylabel('age')
    ax.set_title(title)
axes[0].legend()
plt.tight_layout()
plt.show()
```

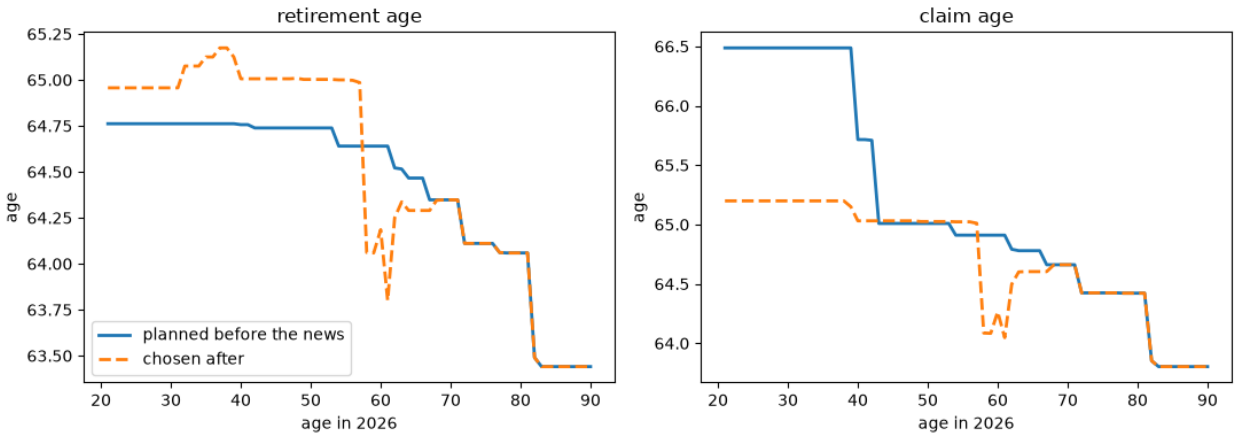


Fig. 128.7: Irreversibility binds the old, while the young work longer but file sooner

Three groups appear in Fig. 128.7.

Workers who have already retired and filed by 2026 sit on top of their old plans, because both decisions are irreversible: the two lines coincide at the right of each panel.

Workers in or near the claiming window move first and move earliest: they file sooner than they had meant to, and some of them leave work sooner too, to collect while benefits are still paid in full.

Workers in their forties and younger go the other way on one margin and with the crowd on the other.

They work longer, because their pension is smaller.

But they file **sooner** than they had planned, by more than a year: with the payable factor falling from one year to the next, a benefit deferred is a benefit collected at a lower factor, and that headwind works against the actuarial reward for waiting.

128.7.2 A worker's life cycle

What the news does to one worker is easier to see than what it does to an aggregate.

Fig. 128.8 follows the median type of the cohort entering in 2026 through its whole life, before and after.

```
R_b, X_b, wealth_b, Om_b, _ = solve_plans(W, ψ, W.j0, plan0=(R_pre, X_pre),
                                          a0=a26.reshape(W.C, 1))

def life_profile(W_, ψ_, R, X, wealth, Om, row):
    "Net income, consumption and assets by age for a single worker."
    p = W_.p
    m = income_path(W_, ψ_, R, X, rows=[row])[:, 0]
    c = wealth[row]/((1+p.τ_c)*Om[row])*(p.β*after_tax_R(p))**np.arange(J)
    a = np.zeros(J)
    for j in range(J-1):
        s = W_.ell[row, j+1]/W_.ell[row, j]
        a[j+1] = (after_tax_R(p)/s)*(a[j] + m[j] - (1+p.τ_c)*c[j])
    return m, c, a

row = np.where(W.entry == 2026)[0][12]          # the median type
before = life_profile(W, np.ones(T+1), R_pre, X_pre, wealth_pre, Om_pre, row)
after = life_profile(W, ψ, R_b, X_b, wealth_b, Om_b, row)
```

Three panels: what the worker receives, what it spends, and what it holds.

```
age = AGE0 + np.arange(J)
fig, axes = plt.subplots(1, 3, figsize=(13, 3.8))
for k, ttl in enumerate(('net income', 'consumption', 'assets')):
    axes[k].plot(age, before[k], lw=2, label='before the news')
    axes[k].plot(age, after[k], lw=2, ls='--', label='after')
    axes[k].set_title(ttl)
    axes[k].set_xlabel('age')
    axes[k].set_xlim(20, 90)
axes[0].legend()
plt.tight_layout()
plt.show()

print(f"retirement age {R_pre[row]} -> {R_b[row]},  claim age {X_pre[row]} -> {X_
 ↪b[row]}")
print(f"peak assets    {before[2].max():.2f} -> {after[2].max():.2f}")
```

```
retirement age 65 -> 65,  claim age 66 -> 65
peak assets    7.53 -> 9.86
```

This worker's retirement age does not move – the choice is an integer, and an aggregate that moves by a year moves because different workers switch at different dates – but it files a year earlier, it consumes less at every age, and it carries about a third more assets into retirement.

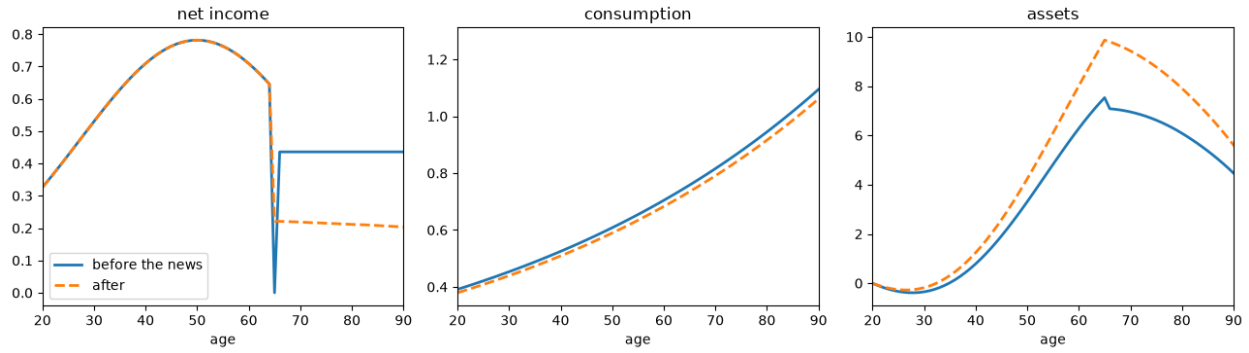


Fig. 128.8: A smaller pension means less consumption, more saving, and no bridge year

The income panel of Fig. 128.8 shows one more thing that the plan before the news was willing to pay for.

Retiring at 65 and filing at 66 leaves a year with no income at all, bridged out of savings, in exchange for a permanently higher actuarial factor.

That bridge is the first thing the news takes away.

The upward tilt of consumption is $\beta\tilde{R} > 1$: this worker is patient enough, at an after-tax return of 4.2 percent, to want a rising profile.

The third panel is the saving response that, in general equilibrium, would deepen the capital stock and raise the wage, and with it taxable payroll: the channel this lecture shuts off and the paper keeps.

128.8 How this compares with the full model

The table collects what the stripped-down model delivers next to the general equilibrium numbers of Sargent and Shi [2026].

```
rows = [("first shortfall year",            f"{shortfall}",            "2032"),
        ("fund empty from",                 f"{depletion}",            "2033"),
        ("ψ in the first full-cut year",    f"{ψ[depletion-Y0]:.3f}",  "0.773"),
        ("ψ in the long run",               f"{ψ[-1]:.3f}",            "0.546"),
        ("mean claim age, 2026",            f"{mean_X[0]:.1f}",        "65.0"),
        ("mean claim age, long run",        f"{mean_X[-1]:.1f}",       "69.2"),
        ("mean retirement age, 2026",       f"{mean_R[0]:.1f}",        "62.7"),
        ("mean retirement age, long run",   f"{mean_R[-1]:.1f}",       "66.2"),
        ("cost rate, 2026",                 f"{cost[0]:.3f}",          "0.136")]

print(f"{'':34}{'this lecture':>14}{'the paper':>12}")
for name, mine, theirs in rows:
    print(f"{name:34}{mine:>14}{theirs:>12}")
```

```
                                  this lecture   the paper
first shortfall year                      2032        2032
fund empty from                           2033        2033
ψ in the first full-cut year             0.755       0.773
ψ in the long run                        0.482       0.546
mean claim age, 2026                      64.7        65.0
mean claim age, long run                  66.6        69.2
mean retirement age, 2026                 64.4        62.7
```

(continues on next page)

(continued from previous page)

```
mean retirement age, long run                65.3        66.2
cost rate, 2026                             0.137       0.136
```

The dates agree and the direction of every response agrees.

Two things do not, and both are worth seeing plainly.

The long-run cut is deeper here, 0.482 against 0.546.

Prices are fixed in this lecture, so the capital deepening that raises the wage, and with it taxable payroll, is missing; and the demography here is a smooth stylization rather than the trustees' projection.

The behavioral responses are also **smaller** here than in the paper, not merely centered differently.

Over the transition the mean retirement age rises by 0.9 years against the paper's 3.5, and the mean claim age by 1.9 against 4.2.

Three features of the stripped-down model mute them: workers face no idiosyncratic earnings risk, so nobody is pushed across a threshold by a bad draw; each worker's choice is an integer, so a type responds only when the gain crosses a whole year; and prices do not move, so the wage never rewards the longer careers.

The retirement *levels* need one more caution: the paper's 62.7 is a mean pulled down by a left tail of early retirements that idiosyncratic shocks produce, and its median, which is what the calibration here was matched to, is 64.5.

What the general equilibrium model adds

A tax cut arrives with the benefit cut. Depletion removes the Social Security deficit from the consolidated federal budget, so the labor tax falls from 0.220 to 0.190 in the long run. Current law cuts retirees' benefits and workers' taxes in the same year.

Capital deepens. Every working cohort saves against a smaller pension, the capital–output ratio rises from 3.00 to 3.19, the interest rate falls from 6.0 to 5.3 percent, and the wage rises.

Financing a bailout is expensive, and the ranking is not the ranking of capital. Paying every scheduled benefit in full costs a newborn in the terminal steady state 4.8 percent of lifetime consumption under consumption taxation, 8.7 percent under capital taxation and 11.6 percent under wage taxation. The capital tax contracts capital the most and is still cheaper than the wage tax, which distorts the retirement margin through which the economy had relieved the pressure on the program, and shrinks its own base.

One program reform changes the arithmetic. Raising the full retirement age from 67 to 69 delays depletion by ten years and raises the long-run payable factor from 0.546 to 0.621. Raising the taxable maximum and cutting the top replacement rate move the incidence and leave the depletion year where it is.

The announcement date redistributes among the living. A bailout is worth about $+4.5$ percent of remaining lifetime consumption to a worker aged 60 in 2026 and about -2 percent to a worker aged 25.

Computing the full model

The paper's solver has the same shape as the one above and differs in scale.

The worker's problem is a dynamic program with assets, an idiosyncratic productivity shock, permanent ability, the earnings record and the claim status as states, solved by backward induction over 81 ages and 150 calendar years; the cross-section is a panel of $80{,}000$ workers per age cell, simulated forward with common random numbers so that every scenario faces the same workers and the same shocks.

Stationary equilibria are found with Anderson acceleration, and the transition is a damped iteration on prices, the bequest transfer, the payable path and a transitory tax rate, as here.

The code is written in JAX and runs on an A100 GPU: about 22 hours from a cold start for the full set of experiments, or about an hour to replay the stored converged paths.

The paper also reports a residual that damping does not remove.

Because the period is a year, the iteration can cycle between two adjacent first shortfall years, each implying the other.

Our own iteration stalls in the same way, for a related reason: retirement and claim ages are integers, so a mass of workers switches discontinuously, and Φ is not continuous.

In the output above, the gap stops falling at about 10^{-2} while the payable factor printed beside it settles to the fourth decimal.

128.9 Exercises

Exercise 128.9.1

The box above claims that Φ is a decreasing map, and that is the reason every iteration here is damped.

Check it, in two steps.

1. Take the converged path, raise it by five percentage points in every year from 2033 on, and apply Φ once. Does the image rise or fall?
2. Run the iteration **undamped**, that is, with $\phi_{k+1} = \Phi(\phi_k)$, for twenty steps starting from $\phi \equiv 1$, and plot the payable factor in 2040 and in 2080 against the iteration count. Compare both with a damped iteration. Does the undamped iteration converge?

Solution

For part 1, one application of the operator is one call each to `solve_plans`, `flows`, `aggregate` and `fund_and_phi`.

```
def Phi(W, ψ, F0, plan0, a0):
    "One application of the operator Φ."
    R, X, _, _, _ = solve_plans(W, ψ, W.j0, plan0=plan0, a0=a0.reshape(W.C, 1))
    pay, ben = flows(W, ψ, R, X, W.j0)
    Pay, Ben = aggregate(W, pay, ben)
    return fund_and_phi(W, Pay, Ben, F0)[1]

ψ_up = ψ.copy()
ψ_up[depletion-Y0:] += 0.05

im_0 = Phi(W, ψ, F0, (R_pre, X_pre), a26)
im_up = Phi(W, ψ_up, F0, (R_pre, X_pre), a26)

k = 2050 - Y0
print(f"ψ in 2050:            {ψ[k]:.4f}  ->  raised to {ψ_up[k]:.4f}")
print(f"Φ(ψ) in 2050:         {im_0[k]:.4f}  ->  {im_up[k]:.4f}")
print(f"mean change in Φ(ψ) over 2033-2100: "
      f"{(im_up - im_0)[depletion-Y0:2100-Y0].mean():+.4f}")
```

```
ψ in 2050:            0.6056  ->  raised to 0.6556
Φ(ψ) in 2050:         0.6044  ->  0.6004
mean change in Φ(ψ) over 2033-2100: -0.0027
```

The image moves the other way: a more generous promise brings earlier claiming and earlier retirement, which lowers taxable payroll and raises the scheduled bill, so the program can pay a smaller fraction.

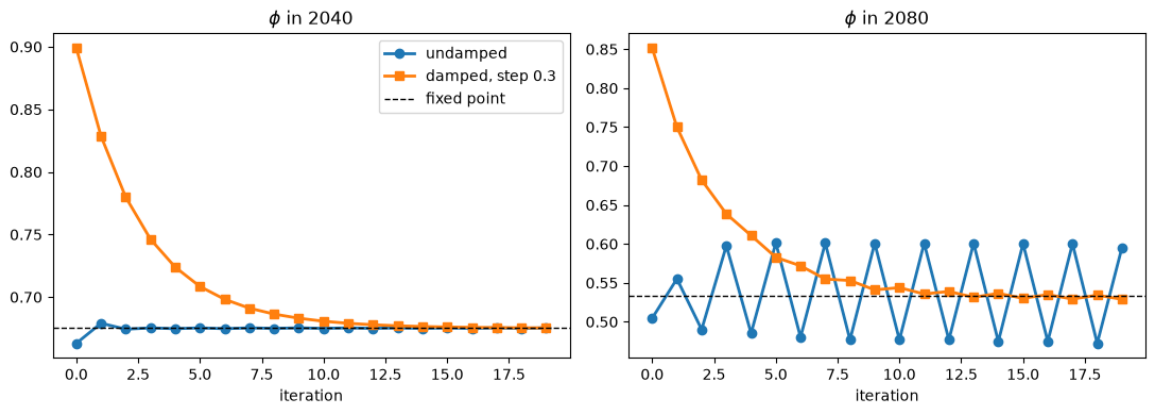
For part 2, an iteration on a decreasing map alternates around its fixed point, and whether the alternation dies out depends on how strong the response is.

```
def trace(ν, years=(2040, 2080), n=20):
    "Payable factor at selected dates along an iteration with step ν."
    φ_k = np.ones(T+1)
    out, gaps = [], []
    for _ in range(n):
        φ_new = Phi(W, φ_k, F0, (R_pre, X_pre), a26)
        gaps.append(np.abs(φ_new - φ_k).max())
        φ_k = (1-ν)*φ_k + ν*φ_new
        out.append([φ_k[y-Y0] for y in years])
    return np.array(out), np.array(gaps)

un, gap_un = trace(1.0)
dp, gap_dp = trace(0.3)

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
for k, y in enumerate((2040, 2080)):
    axes[k].plot(un[:, k], 'o-', lw=2, label='undamped')
    axes[k].plot(dp[:, k], 's-', lw=2, label='damped, step 0.3')
    axes[k].axhline(φ[y-Y0], color='k', ls='--', lw=1, label='fixed point')
    axes[k].set_title(f'$\\phi$ in {y}')
    axes[k].set_xlabel('iteration')
axes[0].legend()
plt.tight_layout()
plt.show()

print(f"undamped, largest residual after 20 steps: {gap_un[-1]:.3f}")
print(f"damped,   largest residual after 20 steps: {gap_dp[-1]:.3f}")
```




```
undamped, largest residual after 20 steps: 0.168
damped,   largest residual after 20 steps: 0.025
```

In 2040 the undamped iteration alternates and settles: the response of claiming and retirement to the payable factor is weak enough at that date that Φ is a contraction there.

In 2080 it is not.

The undamped sequence locks into a two-cycle between about 0.48 and 0.60 and stays there: each payable factor calls forth decisions that justify the other one, and the iteration never stops swinging between them.

The damped sequence settles between the two.

This is why every solver in this lecture takes a fractional step, and why Sargent and Shi [2026] anneal theirs toward a floor.

Exercise 128.9.2

Of the reforms that Sargent and Shi [2026] study, only one moves the depletion date: raising the full retirement age from 67 to 69.

The reform is announced in 2026 and applies to everyone at once, current beneficiaries included.

It works entirely through the actuarial factor: at a full retirement age of 69 the factor is about 13 percent lower at every claim age, which is the gap between the two lines in the right panel of Fig. 128.1, and the claim window still closes at 70, so a worker can recover only part of the cut by filing later.

Compute the reformed transition and report the depletion year and the long-run payable factor.

By how many years is depletion delayed, and how much of the long-run benefit cut does the reform undo?

Solution

Only the actuarial schedule changes, so we build a new `World` with `FRA=69` and run the same transition from the same initial conditions.

```
W69 = World(par._replace(FRA=69), g_pop)
base69 = transition(W69, F0, (R_pre, X_pre), a26, verbose=False)

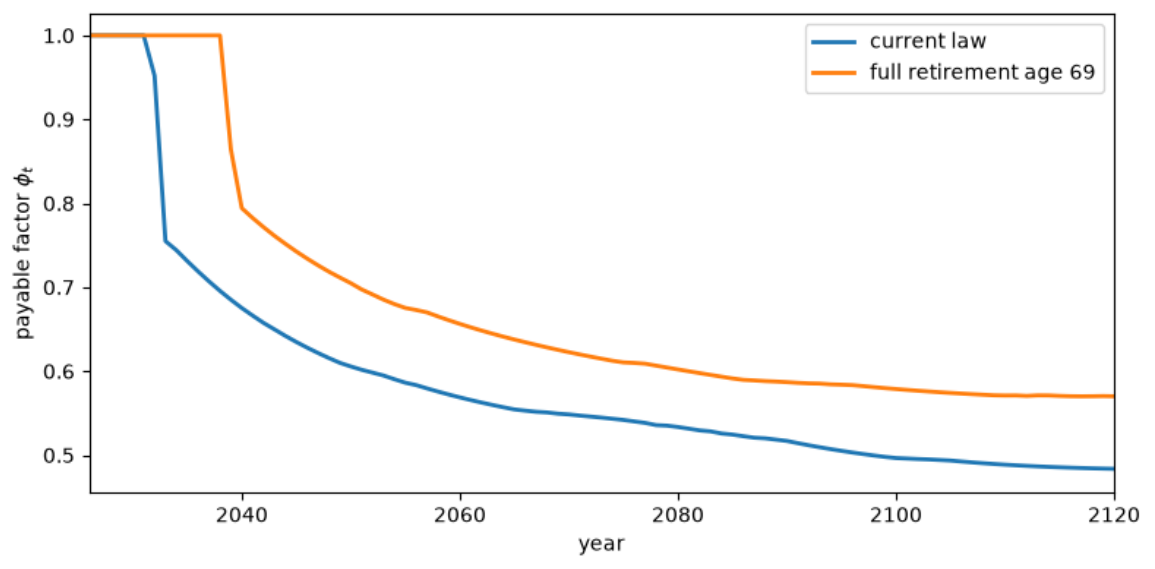
ψ69, F69 = base69['ψ'], base69['F']
dep69 = YEARS[np.argmax(F69[:-1] <= 1e-12)]

print(f"{'':28}{'FRA 67':>10}{'FRA 69':>10}")
print(f"{'depletion year':28}{depletion:>10}{dep69:>10}")
print(f"{'ψ in the long run':28}{ψ[-1]:>10.3f}{ψ69[-1]:>10.3f}")
print(f"{'cost rate in the long run':28}"
      f"{cost[-1]:>10.3f}{(base69['Ben']/base69['Pay'])[-1]:>10.3f}")
```

```
                            FRA 67    FRA 69
depletion year                2033      2040
ψ in the long run            0.482     0.568
cost rate in the long run    0.229     0.195
```

The reform lifts the whole payable path, and the gap between the two never closes.

```
fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(YEARS, ψ, lw=2, label='current law')
ax.plot(YEARS, ψ69, lw=2, label='full retirement age 69')
ax.set_xlabel('year')
ax.set_ylabel('payable factor $\\phi_t$')
ax.set_xlim(Y0, 2120)
ax.legend()
plt.tight_layout()
plt.show()
```



Depletion is delayed by seven years, from 2033 to 2040, against the ten years the paper reports.

The long-run cut falls from 51.8 to 43.2 percent of scheduled benefits, so the reform undoes about a sixth of it.

It does that work because it cuts the scheduled bill at once, for current beneficiaries as well as future ones, so the fund lasts longer and the cash flow condition (128.4) clears at a higher payable factor forever after.

It does not restore scheduled benefits; it lowers them.

A worker's *paid* benefit is the product of the two factors, and the reform raises one by cutting the other.

Exercise 128.9.3

Suppose that instead of letting the cut happen, Congress guarantees every scheduled benefit, $\phi_t = 1$ forever, and pays for it by raising the payroll tax τ^{ss} to a new constant from 2026 on.

1. Find the rate that balances the program's cash flow in the long run, taking into account that workers respond to it.
2. Compute the consumption equivalent variation of the guarantee for cohorts entering between 2026 and 2096. Under the closed form (128.9), a cohort's lifetime value is $V = \Omega \log \mathcal{W} - \eta h D$, so the compensation that equates the two economies is $\lambda = \exp\left(\sum_i w_i \Delta V_i / \Omega\right) - 1$.
3. Do the same for the workers who are **already alive** in 2026, as a function of their age. Who gains from the guarantee and who pays for it?

Solution

For part 1, with $\phi_t \equiv 1$ there is no fixed point in ϕ left to solve: we need only the payroll tax at which the program's long-run cash flow balances, and workers' responses to that tax are inside *Pay* and B^{sch}.

```
def program_gap(τ, plan0, a0):
    "Long-run cash flow of the program when every benefit is paid in full."
    W_τ = World(par._replace(τ_ss=τ), g_pop)
    R, X, _, _, _ = solve_plans(W_τ, np.ones(T+1), W_τ.j0,
                                plan0=plan0, a0=a0.reshape(W_τ.C, 1))
```

```
    pay, ben = flows(W_τ, np.ones(T+1), R, X, W_τ.j0)
    Pay, Ben = aggregate(W_τ, pay, ben)
    return τ*Pay[-1] - (1-par.κ_b)*Ben[-1]

τ_star = brentq(lambda τ: program_gap(τ, (R_pre, X_pre), a26), 0.10, 0.40, xtol=1e-
↪4)
print(f"payroll tax that funds the guarantee: {τ_star:.4f}"
      f"   (current law: {par.τ_ss})")
```

```
payroll tax that funds the guarantee: 0.2229   (current law: 0.106)
```

For part 2, we read the lifetime values off `solve_plans` in the two economies.

```
def cohort_values(W_, ψ_, plan0, a0):
    "Lifetime value of every row, and the weights and Ω needed to aggregate."
    R, X, wealth, Om, Dw = solve_plans(W_, ψ_, W_.j0, plan0=plan0,
                                       a0=a0.reshape(W_.C, 1))
    V = Om*np.log(wealth) - (W_.p.η*W_.η)*Dw
    return V, Om

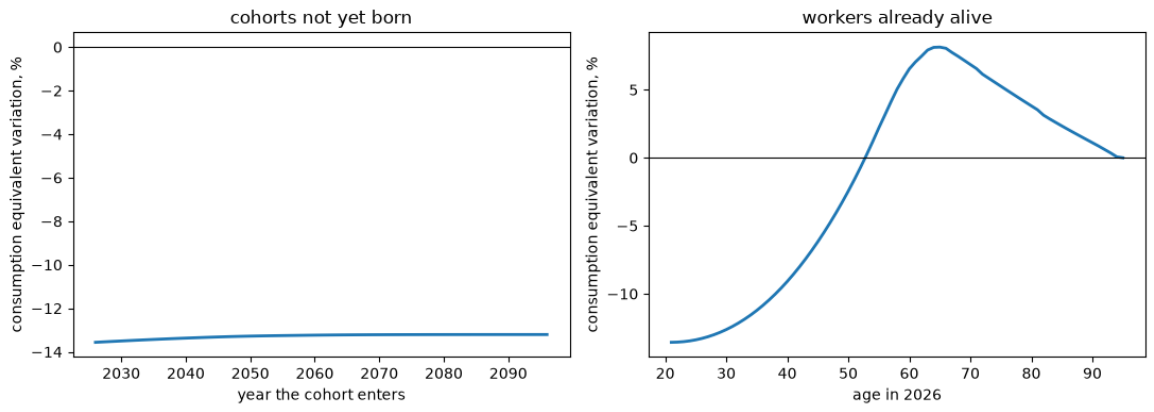
V_dep, Om_dep = cohort_values(W, ψ, (R_pre, X_pre), a26)
W_bail = World(par._replace(τ_ss=τ_star), g_pop)
V_bail, _ = cohort_values(W_bail, np.ones(T+1), (R_pre, X_pre), a26)

def cev(entry_year):
    "Consumption equivalent variation of the guarantee, cohort by cohort."
    m = W.entry == entry_year
    wts = W.w[m]/W.w[m].sum()
    return np.exp(((V_bail[m] - V_dep[m])*wts).sum()/Om_dep[m][0]) - 1

entries = np.arange(2026, 2097)
λ_new = np.array([cev(y) for y in entries])
ages = np.arange(21, 96)
λ_old = np.array([cev(Y0 - (a - AGE0)) for a in ages])

fig, axes = plt.subplots(1, 2, figsize=(11, 4))
axes[0].plot(entries, 100*λ_new, lw=2)
axes[0].set_xlabel('year the cohort enters')
axes[0].set_title('cohorts not yet born')
axes[1].plot(ages, 100*λ_old, lw=2)
axes[1].set_xlabel('age in 2026')
axes[1].set_title('workers already alive')
for ax in axes:
    ax.axhline(0, color='k', lw=0.8)
    ax.set_ylabel('consumption equivalent variation, %')
plt.tight_layout()
plt.show()

print(f"entering in 2026: {100*λ_new[0]:+.2f} percent")
print(f"entering in 2096: {100*λ_new[-1]:+.2f} percent")
for a in (25, 45, 60, 65, 80):
    print(f"   aged {a} in 2026: {100*λ_old[a-21]:+.2f} percent")
```




```
entering in 2026: −13.56 percent
entering in 2096: −13.20 percent
   aged 25 in 2026: −13.39 percent
   aged 45 in 2026: −6.18 percent
   aged 60 in 2026: +6.54 percent
   aged 65 in 2026: +8.14 percent
   aged 80 in 2026: +3.82 percent
```

Every entering cohort is worse off, by about thirteen percent of lifetime consumption, and the cost barely depends on the entry date.

An entering cohort pays the higher payroll tax across its whole working life, to restore the benefits of the retirees of those years, and collects its own restored benefit only decades later, by which time the number of retirees per worker is higher still.

In Sargent and Shi [2026] the losses grow with the entry date, because their financing tax steps up from a transitory to a higher terminal rate partway through the transition, so later cohorts spend more of their lives under the higher rate.

Here the financing rate is one constant from 2026 on, so every entering cohort faces the same schedule and the profile is nearly flat.

That is the finding of Sargent and Shi [2026] for all three of their financing instruments, and the mechanism here is the same one.

The right panel is the other half of the story, and it is the paper's result on timing and incidence.

The cohorts already alive in 2026 divide sharply.

A worker aged 25 loses about as much as a cohort that has yet to enter.

A worker in the late fifties or sixties **gains**: the restored benefit arrives within a few years, and the higher payroll tax falls on the few working years that remain.

Retirees gain without paying at all, except through the general equilibrium channels this model suppresses.

The crossover is a little above age 50.

Sargent and Shi [2026] find the same pattern with a smaller spread, because their wage tax is a broader instrument than a doubled payroll tax, and because their general equilibrium prices move.

What the lecture model cannot say is which instrument is cheapest, because that comparison runs through those prices and through the taxes the guarantee displaces.

Exercise 128.9.4

Two forces raise the dependency ratio in this model: cohorts entering the labor force stopped growing, and mortality keeps improving.

Separate them.

Hold the birth path fixed and switch the mortality improvement off, `par._replace(g_m=0.0)`, so that the life table of 2025 applies forever.

1. What happens to the dependency ratio in 2085, and to the long-run payable factor?
2. What happens to the depletion year? Explain why it moves so much less than the long-run cut does.
3. What happens to the claim age, and what does that tell you about why claiming drifts later in the baseline?

Solution

We rerun the transition in a `World` that keeps the 2025 life table forever.

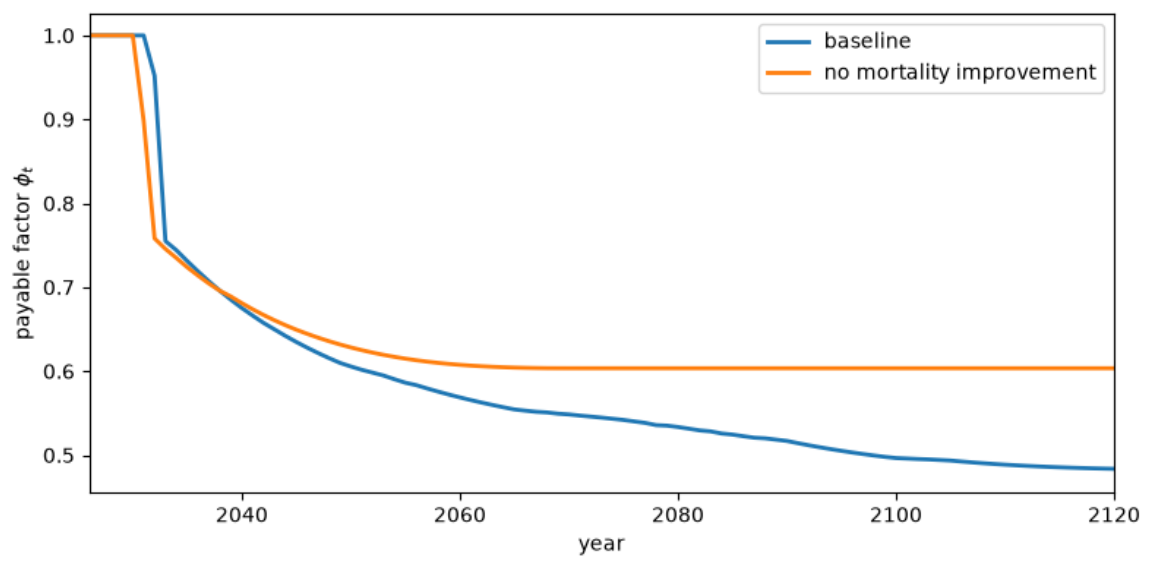
```
W0 = World(par._replace(g_m=0.0), g_pop)        # same cohorts, the 2025 life table␣
↪forever
flat = transition(W0, F0, (R_pre, X_pre), a26, verbose=False)
dep0 = YEARS[np.argmax(flat['F'][:-1] <= 1e-12)]
mR0, mX0 = flow_ages(W0, flat['R'], flat['X'])

print(f"{'':28}{'baseline':>11}{'no gains':>11}")
print(f"{'life expectancy at 20, 2086':28}"
      f"{20+cohort_survival([2086], par)[0].sum():>11.1f}"
      f"{20+cohort_survival([2086], par._replace(g_m=0.0))[0].sum():>11.1f}")
print(f"{'dependency ratio, 2085':28}{dependency(g_pop, 2085, par):>11.3f}"
      f"{dependency(g_pop, 2085, par._replace(g_m=0.0)):>11.3f}")
print(f"{'depletion year':28}{depletion:>11}{dep0:>11}")
print(f"{'ψ in the long run':28}{ψ[-1]:>11.3f}{flat['ψ'][-1]:>11.3f}")
print(f"{'cost rate in the long run':28}{cost[-1]:>11.3f}"
      f"{(flat['Ben']/flat['Pay'])[-1]:>11.3f}")
print(f"{'mean claim age, long run':28}{mean_X[-1]:>11.2f}{mX0[-1]:>11.2f}")
```

```
                            baseline   no gains
life expectancy at 20, 2086     85.1       80.8
dependency ratio, 2085         0.500      0.428
depletion year                  2033       2032
ψ in the long run              0.482      0.603
cost rate in the long run      0.229      0.183
mean claim age, long run       66.65      64.98
```

The two paths separate only slowly, which is the point: longevity does its work in the long run.

```
fig, ax = plt.subplots(figsize=(8, 4))
ax.plot(YEARS, ψ, lw=2, label='baseline')
ax.plot(YEARS, flat['ψ'], lw=2, label='no mortality improvement')
ax.set_xlabel('year')
ax.set_ylabel('payable factor $\\phi_t$')
ax.set_xlim(Y0, 2120)
ax.legend()
plt.tight_layout()
plt.show()
```



Longer lives are a large part of the long-run problem.

Holding the birth path fixed, they account for about a quarter of the long-run benefit cut: the payable factor settles at 0.60 rather than 0.48.

The depletion year hardly moves, and the reason is visible in (128.3).

The date at which the reserves run out is set by the cash flow of the next few years, which the cohorts already alive determine; improvements in mortality accumulate slowly, and most of what they will do to the cost rate still lies decades ahead when the fund empties.

The long run is where they do their work.

The claim age settles at 65.0 rather than 66.6.

That is the mechanism the lecture asserted for later claiming and had not tested.

The actuarial schedule $\mathrm{AF}(\chi)$ is statutory and frozen, so when mortality improves, waiting buys the higher benefit for more years, and workers wait.

Switch the improvement off and the drift disappears.

Exercise 128.9.5

Exercise 128.9.3 guaranteed scheduled benefits from 2026, with workers told in 2026 and taxed from 2026.

Sargent and Shi [2026] also study the opposite timing, which they call the **unanticipated** scenario.

The government says nothing and does nothing.

Workers believe current law, plan against the cut they expect, and save against it.

Then, in the year the reserves are spent, the government announces the guarantee and raises the financing tax at once.

The cut that workers had expected never happens.

Reproduce that experiment here, and compare it with the immediate bailout of Exercise 128.9.3.

1. Roll the baseline economy forward to the first shortfall year and collect the state it carries there: the assets each worker has accumulated and the retirement and claim decisions it has already executed.

2. Let workers re-plan in that year under $\phi_t = 1$ and the payroll tax τ^*, using the `at_year` argument of `solve_plans`.
3. Value the three economies for the workers alive in 2026 and report the consumption equivalent variation of each bailout against the depletion baseline. Who prefers the surprise?

Solution

Consumption is $c_0(\beta\tilde{R})^{j-j_0}$ from whenever a worker last re-planned, so assets follow a one-line recursion along the baseline path.

```
t_star = shortfall                          # the year the reserves are spent
j_star = np.clip(t_star - W.entry, 0, J-1)

c0_b = wealth_b/((1+par.τ_c)*Om_b)          # the baseline plan, re-optimized in 2026
m_b = income_path(W, ψ, R_b, X_b)

a, A = np.zeros(W.C), np.zeros((J, W.C))
for j in range(J-1):
    a = np.where(j == W.j0, a26, a)         # the 2026 cross-section enters here
    A[j] = a
    surv = W.ell[:, j+1]/np.maximum(W.ell[:, j], 1e-12)
    c = c0_b*(par.β*after_tax_R(par))**np.maximum(j - W.j0, 0)
    a = np.where(j >= W.j0,
                 (after_tax_R(par)/np.maximum(surv, 1e-12))
                 * (a + m_b[j] - (1+par.τ_c)*c), 0.0)
A[J-1] = a
a_star = A[j_star, np.arange(W.C)]

old = W.entry == Y0 - 40                    # the cohort aged 60 in 2026
print(f"assets of a worker aged 60 in 2026: {a26[old].mean():.2f} in 2026,"
      f" {a_star[old].mean():.2f} in {t_star}")
```

```
assets of a worker aged 60 in 2026: 8.18 in 2026, 9.25 in 2032
```

Now the surprise.

The same `solve_plans` re-plans at `t_star`, freezing the decisions already executed by then.

```
R_n, X_n, wealth_n, Om_n, _ = solve_plans(
    W_bail, np.ones(T+1), j_star, plan0=(R_b, X_b),
    a0=a_star.reshape(W.C, 1), at_year=t_star)
c0_n = wealth_n/((1+par.τ_c)*Om_n)

R_i, X_i, wealth_i, Om_i, _ = solve_plans(          # the immediate bailout again
    W_bail, np.ones(T+1), W.j0, plan0=(R_pre, X_pre), a0=a26.reshape(W.C, 1))
c0_i = wealth_i/((1+par.τ_c)*Om_i)
```

We value the three economies directly, by summing utility along each worker's actual consumption and work path rather than by the closed form of Exercise 128.9.3.

```
ages_j = np.arange(J)
βR = par.β*after_tax_R(par)

def consumption(c0, j0, j_switch=None, c0_new=None):
    "Consumption at every age, with an optional replan at j_switch."
    c = c0[None, :]*βR**np.maximum(ages_j[:, None] - j0[None, :], 0)
    if j_switch is not None:
        c_new = c0_new[None, :]*βR**np.maximum(ages_j[:, None] - j_switch[None, :],
 ↪ 0)
        c = np.where(ages_j[:, None] >= j_switch[None, :], c_new, c)
    return c

def working(R, R_new=None, j_switch=None):
    w = (ages_j[:, None] < (R - AGE0)[None, :]).astype(float)
    if R_new is not None:
        w = np.where(ages_j[:, None] >= j_switch[None, :],
                     (ages_j[:, None] < (R_new - AGE0)[None, :]).astype(float), w)
    return w

def lifetime_value(W, c, work, j0):
    "Σ_{j≥j0} β^{j-j0} ℓ_{j|j0} [log c_j - η h n_j], for every row."
    p = W.p
    wt = (p.β**ages_j)[:, None]*W.ell.T
    u = np.log(np.maximum(c, 1e-12)) - (p.η*W.η)[None, :]*work
    keep = ages_j[:, None] >= j0[None, :]
    return (np.where(keep, wt*u, 0.0).sum(axis=0)
            / ((p.β**j0)*W.ell[np.arange(W.C), j0]))

V_dep = lifetime_value(W, consumption(c0_b, W.j0), working(R_b), W.j0)
V_imm = lifetime_value(W, consumption(c0_i, W.j0), working(R_i), W.j0)
V_sur = lifetime_value(W, consumption(c0_b, W.j0, j_star, c0_n),
                       working(R_b, R_n, j_star), W.j0)

def cev_age(V1, V0, age):
    m = W.entry == Y0 - (age - AGE0)
    wts = W.w[m]/W.w[m].sum()
    return np.exp(((V1 - V0)[m]*wts).sum()/Om_b[m][0]) - 1
```

Before reading the answer, it is worth checking the machinery against Exercise 128.9.3, which computed the immediate bailout a different way.

```
print(f"{'age in 2026':>12}{'this exercise':>15}{'exercise 3':>12}")
for a in (25, 45, 60, 80):
    print(f"{a:>12}{100*cev_age(V_imm, V_dep, a):>15.2f}"
          f"{100*λ_old[a-21]:>12.2f}")
```

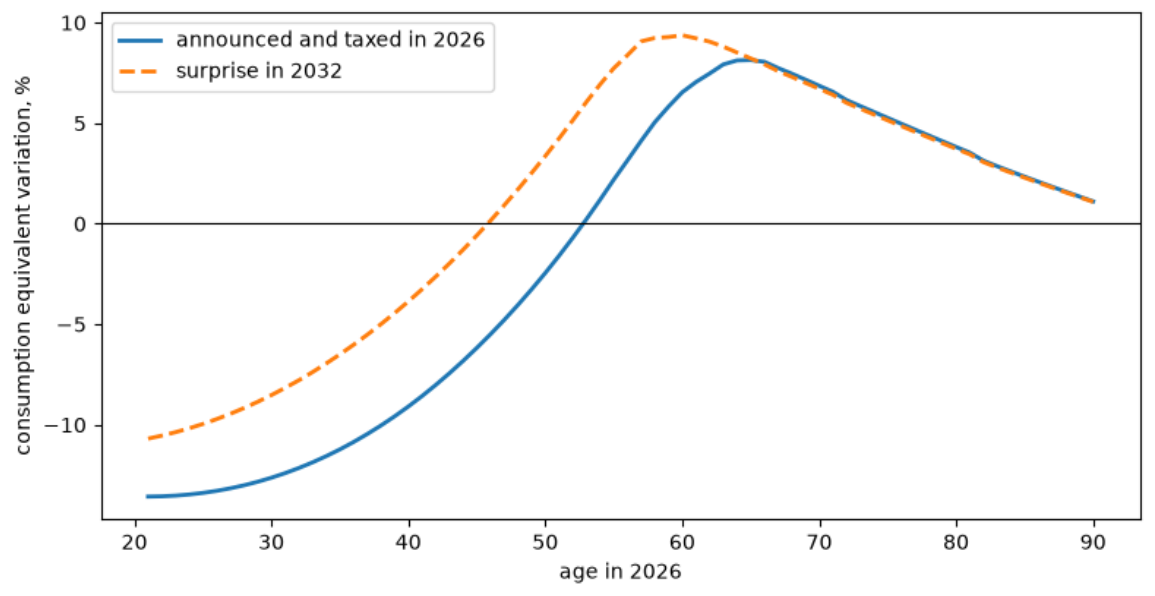
```
  age in 2026  this exercise  exercise 3
           25         -13.39      -13.39
           45          -6.18       -6.18
           60           6.54        6.54
           80           3.82        3.82
```

The two agree, so we can compare the timings.

```
ages_alive = np.arange(21, 91)
λ_imm = np.array([cev_age(V_imm, V_dep, a) for a in ages_alive])
λ_sur = np.array([cev_age(V_sur, V_dep, a) for a in ages_alive])

fig, ax = plt.subplots(figsize=(8, 4.2))
ax.plot(ages_alive, 100*λ_imm, lw=2, label='announced and taxed in 2026')
ax.plot(ages_alive, 100*λ_sur, lw=2, ls='--', label=f'surprise in {t_star}')
ax.axhline(0, color='k', lw=0.8)
ax.set_xlabel('age in 2026')
ax.set_ylabel('consumption equivalent variation, %')
ax.legend()
plt.tight_layout()
plt.show()

print(f"{'age in 2026':>12}{'immediate':>11}{'surprise':>10}{'gain':>8}")
for a in (25, 35, 45, 55, 65, 75, 85):
    i, u = 100*λ_imm[a-21], 100*λ_sur[a-21]
    print(f"{a:>12}{i:>11.2f}{u:>10.2f}{u-i:>8.2f}")
```




```
age in 2026  immediate  surprise    gain
         25     -13.39     -9.96    3.43
         35     -11.22     -6.50    4.72
         45      -6.18     -0.58    5.59
         55       2.19      7.73    5.54
         65       8.14      8.22    0.08
         75       5.26      5.14   -0.12
         85       2.34      2.29   -0.05
```

Everyone with working years left prefers the surprise, and those in middle age prefer it by a lot.

The reason is arithmetic: both economies end up paying every scheduled benefit in full, and the surprise collects six fewer years of the higher tax.

The age at which a worker stops losing from the guarantee falls from 53 to 46.

Nobody is ever actually cut, because the rescue lands in the year the reserves are spent, the year before the first full cut would have fallen.

The gain is largest in middle age, 5.6 percent at 45, where many taxed years lie ahead and retirement is near enough that the mis-planning costs little.

It is smaller for the young, who pay the higher tax for most of their lives either way.

For workers in their seventies and beyond, who have already retired and filed, the surprise is very slightly **worse**.

They pay no payroll tax under either timing, so deferring it gains them nothing, and they spent six years consuming as though their benefits were about to be cut.

That consumption they did not take is a loss that the announcement would have spared them.

Sargent and Shi [2026] find the same ranking under wage-tax financing, with one addition this model cannot produce: in general equilibrium the assets that workers accumulate against a cut that never comes *are* the economy's capital stock when the rescue arrives, so waiting also leaves a higher wage behind it.

Here prices do not move, and the whole gain is the deferred tax.

Exercise 128.9.6

Everything in this lecture rests on one calibration.

Two of its numbers deserve a sensitivity check: the demographic target, and the interest rate at which workers discount.

1. Recalibrate the birth path to an aged dependency ratio in 2085 of 0.45 and of 0.55, instead of the trustees' 0.50, and recompute the long run.
2. Recompute it at interest rates of 5 and 7 percent, holding the preference parameters at their calibrated values.
3. Which of the two moves the long-run benefit cut, and which moves behavior? Explain what you find using (128.6).

Solution

Each case is a fresh economy: new demography or new prices change the plans made before the news as well as the transition, so we re-solve both.

Twenty-five damped steps are enough for three decimals here.

```
def long_run(p, g, label):
    "Solve one economy end to end and report its long run."
    W_ = World(p, g)
    R0_, X0_, a0_, _, Ben0_, _, _ = pre_transition(W_)
    res = transition(W_, p.F_ratio*W_.Npop[0]*Ben0_[0], (R0_, X0_), a0_,
                     iters=25, verbose=False)
    mR_, mX_ = flow_ages(W_, res['R'], res['X'])
    print(f"{label:26}{YEARS[np.argmax(res['F'][:-1] <= 1e-12)]:>7}"
          f"{res['ψ'][-1]:>8.3f}{(res['Ben']/res['Pay'])[-1]:>8.3f}"
          f"{res['Pay'][-1]:>9.3f}{res['Ben'][-1]:>10.4f}{mR_[-1]:>8.1f}")

print(f"{'':26}{'deplete':>7}{'ψ_∞':>8}{'cost':>8}{'payroll':>9}{'benefits':>10}{
 ↪'retire':>8}")
long_run(par, g_pop, "baseline")
for target in (0.45, 0.55):
    g_t = fsolve(lambda g: [dependency(g, 2025, par) - 0.32,
                            dependency(g, 2085, par) - target], [0.02, 0.0])
    long_run(par, g_t, f"dependency 2085 = {target}")
for r_alt in (0.05, 0.07):
    long_run(par._replace(r=r_alt), g_pop, f"interest rate = {r_alt}")
```

	deplete	ψ_∞	cost	payroll	benefits	retire
baseline	2033	0.482	0.229	0.663	0.1519	65.3
dependency 2085 = 0.45	2034	0.531	0.208	0.677	0.1408	64.9
dependency 2085 = 0.55	2033	0.440	0.251	0.647	0.1622	65.5
interest rate = 0.05	2033	0.488	0.227	0.705	0.1597	68.5
interest rate = 0.07	2033	0.500	0.221	0.613	0.1355	62.0

Demography moves the long-run cut, and the interest rate essentially does not.

Five points of the dependency ratio are worth about five points of the payable factor, in the opposite direction.

The interest rate moves the payable factor by one or two points while moving the average retirement age by **six and a half years**, from 68.5 at five percent to 62.0 at seven.

Equation (128.6) says why.

The long-run cut depends on behavior only through the cost rate, the ratio of the scheduled bill to taxable payroll, and the two columns beside it show what each experiment does to the numerator and the denominator.

A higher dependency ratio pushes them in opposite directions: fewer workers to tax, more retirees to pay.

A higher interest rate pushes them in the same direction, and by almost the same proportion: shorter careers shrink taxable payroll, and the earlier claiming that comes with them buys a lower actuarial factor, so the scheduled bill shrinks too.

The ratio hardly notices.

This is the two-clock structure of the benefit formula seen from the long run rather than from the transition: the same career decision that sets what a worker contributes also sets what the worker is owed, and in a stationary economy the two very nearly cancel.

CHAPTER

NINE

LINEAR QUADRATIC MEAN FIELD GAMES

Contents

- *Linear Quadratic Mean Field Games*
  - *Overview*
  - *The environment*
  - *Solving the individual problem*
  - *A state-costate system*
  - *The equilibrium Riccati equation*
  - *Computing equilibria*
  - *The scalar case*
  - *An industry equilibrium with capital accumulation*
  - *Aggregate shocks and identification*
  - *Persistence*
  - *Multiproduct price setting with Kimball demand*
  - *The planner*
  - *Summary*
  - *Exercises*
  - *Further reading*

129.1 Overview

A **mean field game** describes a continuum of small agents, each of whom solves a dynamic optimization problem whose payoff depends on what everybody else is doing, summarized by the cross-sectional distribution of states.

The framework was introduced by Lasry and Lions [2007] and, independently, by Huang *et al.* [2006].

It pairs two partial differential equations:

- a **Hamilton-Jacobi-Bellman** equation that runs backward in time and describes an individual's optimal choice, given paths for the aggregates

- a **Kolmogorov forward** equation that runs forward in time and describes how the distribution of individual states evolves, given those choices

An equilibrium requires that the aggregates that agents take as given are the ones that their own decisions generate.

Readers of *Rational Expectations Equilibrium* will recognize that requirement: it is the "Big Y, little y" idea, now applied to an entire distribution rather than to a single number.

The same pair of equations is the workhorse of continuous-time heterogeneous-agent macroeconomics; see Achdou *et al.* [2022].

This lecture studies a tractable special case in which the payoff is quadratic and the state evolves linearly, following Alvarez and Argente [2026].

Agents care about the cross-sectional average *state* X, through a matrix Θ_X, and about the cross-sectional average *action* $\mathcal{A}$, through a matrix $\Theta_{\mathcal{A}}$.

The main result is a striking simplification:

Note

The equilibrium of the mean field game solves the algebraic Riccati equation of a *single-agent* linear quadratic regulator problem, in which the curvature matrices Q and Γ of the return in an agent's own state and own action are replaced by the **effective curvatures**

$$Q + \Theta_X \qquad \text{and} \qquad \Gamma + \Theta_{\mathcal{A}}.$$

Everything we know about the linear regulator can therefore be brought to bear on the equilibrium: existence conditions, uniqueness, comparative statics, and numerical methods.

Four lessons organize the lecture:

- **Effective curvature** (*Effective curvature*). Aggregate dynamics depend on strategic interaction only through the effective curvatures, which measure how fast an agent's marginal returns fall when everybody moves together.
- **Individual versus aggregate** (*Aggregate and individual dynamics*). An agent's deviation from the cross-sectional mean depends only on Q and Γ, which measure how fast her marginal returns fall when she moves alone, so individual (micro) and aggregate (macro) adjustment speeds differ: three years against five in our capital example.
- **The missing intercept** (*Aggregate shocks and identification*). Regressions with time effects see only deviations from the mean, while strategic interaction lives entirely in the time effects.
- **Marginal versus total** (*The planner*). A planner responds to how the *total* return changes when everybody moves together, not to how an agent's marginal return changes, and that can make the planner's allocation either slower or faster than the equilibrium.

We illustrate them with two economic examples from Alvarez and Argente [2026]: an industry equilibrium with capital accumulation, and a multiproduct pricing problem with Kimball demand.

A sequel, *Linear Quadratic Mean Field Games: Extensions*, adds cross terms between states and actions, forcing processes, and interactions through higher moments of the cross-sectional distribution.

Riccati equations appear in several other QuantEcon lectures:

- *LQ Control: Foundations* introduces the linear regulator and its Riccati equation
- *Lagrangian for LQ Control* studies the state-costate system and its stable invariant subspace, which is exactly the structure we meet below
- *Markov Perfect Equilibrium* studies dynamic games with *finitely many* players, where each player has a Riccati equation and the equations are coupled

- *A First Look at the Kalman Filter* presents the Riccati equation that is dual to the control problem

Let's start with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from matplotlib.patches import Patch
from matplotlib.lines import Line2D
from scipy.linalg import solve_continuous_are, expm
from scipy.optimize import brentq
```

129.2 The environment

There is a continuum of agents.

An individual has state $x \in \mathbb{R}^n$ and takes an action $\alpha \in \mathbb{R}^k$.

The cross-sectional average state and action are

$$X = \int x\, m(x) dx, \qquad \mathcal{A} = \int \alpha^*(x)\, m(x) dx,$$

where m is the density of the state across agents and α^* is the optimal policy.

The period return is

$$F(x, X) + R(\alpha, \mathcal{A}) = -\tfrac{1}{2} x^\top Q x - X^\top \Theta_X x - \tfrac{1}{2} \alpha^\top \Gamma \alpha - \mathcal{A}^\top \Theta_{\mathcal{A}} \alpha. \tag{129.1}$$

This is not the most general quadratic return.

A second-order approximation to a general return also produces cross terms between states and actions, which *Linear Quadratic Mean Field Games: Extensions* takes up.

It also produces terms that involve only the aggregates, which do not affect an individual's choices but do matter for a planner, as we will see in *The planner*.

We omit constants, which do not affect choices, and linear terms, which affect steady-state levels but not the dynamics of deviations from them.

The individual state follows

$$dx = (B\alpha - Ax) dt + \Sigma^{1/2} dW, \tag{129.2}$$

where W is an n-dimensional Brownian motion and shocks are purely idiosyncratic for now.

Agents discount at rate $\rho > 0$.

We assume that Q and Γ are positive definite, that Σ is positive semi-definite, that B has full row rank, and that Θ_X and $\Theta_{\mathcal{A}}$ are symmetric.

Symmetry of the interaction matrices is a genuine restriction.

The own-curvature matrix Q can always be taken to be symmetric, because $x^\top Q x$ depends only on the symmetric part of Q, but $X^\top \Theta_X x$ pairs two different vectors, and the antisymmetric part of Θ_X does affect choices.

It is, however, a restriction on payoffs rather than on behavior: the pricing example below has asymmetric best responses but a symmetric Θ_X.

129.2.1 Complements and substitutes

From (129.1), the cross derivative between an agent's own state and the average state is $-\Theta_X$, and between her own action and the average action it is $-\Theta_{\mathcal{A}}$.

So

- states are **strategic complements** when $-\Theta_X$ is positive definite, i.e. when Θ_X is negative definite
- states are **strategic substitutes** when Θ_X is positive definite

and similarly for actions.

We compare symmetric matrices with the Loewner order, in which $\Theta_1 \preceq \Theta_2$ means that $\Theta_2 - \Theta_1$ is positive semi-definite.

In that order, *smaller* Θ means *more* complementarity.

Note

The monotonicity condition that Lasry and Lions [2007] impose to obtain uniqueness corresponds here to Θ_X being positive semi-definite, that is, to strategic *substitutability* in states.

We will not need it: under the symmetry assumptions above, and provided $\Gamma + \Theta_{\mathcal{A}}$ is positive definite, there is at most one equilibrium whether interactions are complements or substitutes.

Linear Quadratic Mean Field Games: Extensions shows how uniqueness can fail once interactions between states and actions are allowed to be asymmetric.

129.2.2 Equilibrium

Taking the paths $\{X(t), \mathcal{A}(t)\}$ as given, an agent's value function satisfies the HJB equation

$$\rho u(x,t) = -\tfrac{1}{2} x^\top Q x - X(t)^\top \Theta_X x + H(u_x(x,t), x, \mathcal{A}(t)) + \tfrac{1}{2} \operatorname{tr}(\Sigma u_{xx}(x,t)) + u_t(x,t), \tag{129.3}$$

where the Hamiltonian is

$$H(p, x, \mathcal{A}) = \max_\alpha \left\{ -\tfrac{1}{2} \alpha^\top \Gamma \alpha - \mathcal{A}^\top \Theta_{\mathcal{A}} \alpha + p^\top (B\alpha - Ax) \right\},$$

with maximizer $\alpha^*(p, \mathcal{A}) = \Gamma^{-1}(B^\top p - \Theta_{\mathcal{A}} \mathcal{A})$.

The density evolves according to the Kolmogorov forward equation

$$m_t(x,t) = -\operatorname{div}\left(H_p(u_x(x,t), x, \mathcal{A}(t)) m(x,t)\right) + \tfrac{1}{2} \operatorname{tr}(\Sigma m_{xx}(x,t)), \tag{129.4}$$

and an **equilibrium** is a value function, a density, and paths for X and $\mathcal{A}$ that satisfy (129.3), (129.4), and the consistency requirements that X and $\mathcal{A}$ really are the cross-sectional averages implied by m and by the optimal policy.

129.3 Solving the individual problem

Given the aggregate paths, an individual faces a time-varying linear quadratic regulator problem, so her value function is quadratic:

$$u(x,t) = \beta_0(t) + \beta_1(t)^\top x + \tfrac{1}{2} x^\top \beta_2(t) x.$$

Substituting into (129.3) and matching terms of each order gives three differential equations.

The one for β_2 is

$$\dot{\beta}_2 = Q - \beta_2 B\Gamma^{-1}B^\top\beta_2 + \beta_2 A + A^\top\beta_2 + \rho\beta_2. \tag{129.5}$$

Notice what is *absent* from (129.5): neither interaction matrix appears.

The curvature of an individual's value function is therefore the same as it would be if she were alone in the world.

Because the individual problem is concave and stationary, $\beta_2(t)$ equals the constant $\bar{\beta}_2$, the negative definite solution of

$$\bar{\beta}_2 B\Gamma^{-1}B^\top\bar{\beta}_2 = Q + \rho\bar{\beta}_2 + \bar{\beta}_2 A + A^\top\bar{\beta}_2. \tag{129.6}$$

This is the familiar algebraic Riccati equation of *LQ Control: Foundations*, written in continuous time.

The optimal action is

$$\alpha^*(x,t) = \Gamma^{-1}\left[B^\top(\beta_1(t) + \bar{\beta}_2 x) - \Theta_{\mathcal{A}}\mathcal{A}(t)\right].$$

Averaging across agents and solving the resulting fixed point in $\mathcal{A}$ gives

$$\mathcal{A}(t) = (\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top\left(\beta_1(t) + \bar{\beta}_2 X(t)\right). \tag{129.7}$$

Equation (129.7) is where the action interaction first bites: each agent responds to the average action, and solving for the average that is consistent with everyone doing so replaces Γ by $\Gamma + \Theta_{\mathcal{A}}$.

129.4 A state-costate system

Two objects now remain: the linear coefficient $\beta_1(t)$ of the value function, and the aggregate state $X(t)$.

Matching the terms in the HJB equation that are linear in x gives

$$\dot{\beta}_1 = (\rho I + A^\top)\beta_1 + \Theta_X X - \bar{\beta}_2 B\Gamma^{-1}\left(B^\top\beta_1 - \Theta_{\mathcal{A}}\mathcal{A}\right),$$

and averaging the state equation across agents gives $\dot{X} = B\mathcal{A} - AX$, because the idiosyncratic shocks average out.

Substituting (129.7) into both equations gives a pair of linear differential equations,

$$\begin{bmatrix}\dot{\beta}_1 \\ \dot{X}\end{bmatrix} = \mathcal{H}\begin{bmatrix}\beta_1 \\ X\end{bmatrix}, \qquad \mathcal{H} = \begin{bmatrix}\rho I + A^\top - \bar{\beta}_2 D & \Theta_X + \bar{\beta}_2(B\Gamma^{-1}B^\top - D)\bar{\beta}_2 \\ D & -A + D\bar{\beta}_2\end{bmatrix}, \tag{129.8}$$

where we abbreviate

$$D \equiv B(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top.$$

This is a **state-costate** system of exactly the kind studied in *Lagrangian for LQ Control*.

The aggregate state X has an initial condition, namely the mean of the initial distribution.

The costate β_1 does not: it must be chosen so that the solution does not violate the agent's transversality condition, which here requires the eigenvalues governing the path to have real parts below $\rho/2$.

Because Θ_X, $B\Gamma^{-1}B^\top$ and D are symmetric, $\mathcal{H} - \frac{\rho}{2}I$ is a Hamiltonian matrix, so its eigenvalues are symmetric about the origin.

It follows that:

Proposition 129.4.1

If λ is an eigenvalue of $\mathcal{H}$, then so are $\rho - \lambda$, $\bar{\lambda}$, and $\rho - \bar{\lambda}$.

So at most n eigenvalues can have real parts below $\rho/2$, and exactly n do when none lies on the line $\operatorname{Re}\lambda = \rho/2$.

That pins down a unique stable invariant subspace, and when D is positive definite that subspace can be written as $\beta_1 = SX$ for a single matrix S, so there is at most one equilibrium.

This is the standard connection between algebraic Riccati equations and invariant subspaces of Hamiltonian matrices, treated at length by Lancaster and Rodman [1995].

129.5 The equilibrium Riccati equation

We look for a saddle path along which the costate is a linear function of the state, $\beta_1(t) = SX(t)$.

Substituting into (129.8) gives a quadratic matrix equation for S whose coefficients involve $\bar{\beta}_2$, and a change of variable simplifies it.

Define

$$P \equiv S + \bar{\beta}_2.$$

The new matrix has a natural meaning.

An agent at the mean has marginal value of the state $u_x = \beta_1 + \bar{\beta}_2 X = (S + \bar{\beta}_2)X = PX$, so P maps the aggregate state into the costate of the average agent.

Proposition 129.5.1

An equilibrium is characterized by a symmetric matrix P solving

$$P\,B(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top\,P = Q + \Theta_X + \rho P + PA + A^\top P,$$

with aggregate dynamics

$$\dot{X} = \left(B(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top P - A\right)X.$$

The equilibrium requires all eigenvalues of the closed-loop matrix to have real parts below $\rho/2$.

Compare this with the single-agent equation (129.6).

They have *the same form*.

The only difference is that Q has become $Q + \Theta_X$ and Γ has become $\Gamma + \Theta_{\mathcal{A}}$.

Proposition 129.5.2

Suppose $Q + \Theta_X$ and $\Gamma + \Theta_{\mathcal{A}}$ are positive definite.

Then the equilibrium of a linear quadratic mean field game with interaction matrices Θ_X and $\Theta_{\mathcal{A}}$, and its aggregate law of motion, coincide with the solution of a single-agent linear quadratic control problem whose state curvature is $Q + \Theta_X$ and whose action curvature is $\Gamma + \Theta_{\mathcal{A}}$.

This is the organizing result of the lecture.

129.5.1 Effective curvature

We call $Q + \Theta_X$ and $\Gamma + \Theta_{\mathcal{A}}$ the **effective curvatures**.

They have a simple reading.

An agent's marginal return to her state is $-Qx - \Theta_X X$.

If she moves alone, holding X fixed, it falls with slope Q.

If everybody moves together, so that $x = X$, it equals $-(Q + \Theta_X)X$ and falls with slope $Q + \Theta_X$.

The first slope governs an agent's deviations from the crowd, and the second governs the crowd.

The same reading applies to actions, with Γ and $\Gamma + \Theta_{\mathcal{A}}$.

Fig. 129.2 below draws the two slopes for our capital accumulation example.

So strategic interaction does not change the *form* of the problem that determines aggregate dynamics; it changes the *curvatures* that enter it.

An immediate implication is that two models with very different strategic interactions can generate identical aggregate dynamics, provided their effective curvatures agree, and *Aggregate shocks and identification* shows what that means for estimation.

129.5.2 Existence and uniqueness

Since P solves a standard Riccati equation, standard conditions apply.

Define

$$E \equiv Q + \Theta_X + (A^\top + \tfrac{\rho}{2} I)\left[B(\Gamma + \Theta_{\mathcal{A}})^{-1} B^\top\right]^{-1}(A + \tfrac{\rho}{2} I)\,. \tag{129.9}$$

Proposition 129.5.3

Suppose $B(\Gamma + \Theta_{\mathcal{A}})^{-1} B^\top$ is positive definite.

1. There is at most one equilibrium.
2. A necessary condition for a symmetric solution of the equilibrium Riccati equation is that E be positive semi-definite.
3. If $Q + \Theta_X$ is positive definite, an equilibrium exists.

The sufficient condition has a clean reading: *effective* curvature must remain positive in both states and actions.

For states that is $Q + \Theta_X \succ 0$; for actions it is the maintained assumption on $B(\Gamma + \Theta_{\mathcal{A}})^{-1} B^\top$, which holds whenever $\Gamma + \Theta_{\mathcal{A}} \succ 0$.

Complementarity is therefore permissible, but only up to a point.

The two interactions enter E through separate terms, and a larger $\Theta_{\mathcal{A}}$ raises the second, so substitutability in actions can offset complementarity in states.

129.5.3 Aggregate and individual dynamics

The equilibrium splits each agent's state into the cross-sectional mean X and her deviation from it, $z = x - X$.

Using $\beta_1 = (P - \bar{\beta}_2)X$ in the optimal action shows that $\alpha^* - \mathcal{A} = \Gamma^{-1}B^\top \bar{\beta}_2 z$.

Subtracting the aggregate law of motion from the individual one then gives

$$\begin{aligned} \dot{X} &= (DP - A)\, X, \\ dz &= \left(B\Gamma^{-1}B^\top \bar{\beta}_2 - A\right) z\, dt + \Sigma^{1/2} dW, \\ \alpha^* &= (\Gamma + \Theta_{\mathcal{A}})^{-1} B^\top P X + \Gamma^{-1} B^\top \bar{\beta}_2 z. \end{aligned} \tag{129.10}$$

The aggregate follows the equilibrium matrix P.

An agent's deviation from the mean follows the single-agent matrix $\bar{\beta}_2$, and neither interaction matrix appears, because strategic interaction operates through X and $\mathcal{A}$, which drop out of deviations from the mean.

So the deviation of any agent from the crowd moves exactly as the whole economy would move if there were no interactions.

From now on we call these **individual** dynamics, as opposed to **aggregate** dynamics, and we label them "no interactions" in tables and figures.

This split is also the source of the missing intercept that we study in *Aggregate shocks and identification*: data on deviations from the mean reveal only the individual dynamics.

129.6 Computing equilibria

`scipy.linalg.solve_continuous_are(A, B, Q, R)` returns the stabilizing solution $\tilde{P}$ of

$$A^\top \tilde{P} + \tilde{P}A - \tilde{P}BR^{-1}B^\top \tilde{P} + Q = 0.$$

Our equation differs in two ways: our P is negative definite when the effective curvatures are positive definite, and we discount.

Writing $P = -\tilde{P}$ and collecting terms shows that our equation is the standard one with A replaced by $-(A + \frac{\rho}{2} I)$.

The discount rate enters exactly as it does in the "$\rho/2$ shift" familiar from continuous-time control.

When no admissible solution exists, `solve_continuous_are` usually raises an error, but not always: it can return a matrix that does not solve the equation at all.

So our function checks both the Riccati residual and the eigenvalue condition, and raises an error if either fails.

```
def closed_loop(A, B, Γ_eff, P):
    "The matrix governing aggregate dynamics, Ẋ = (B Γ_eff^{-1} B' P - A) X."
    return B @ np.linalg.solve(Γ_eff, B.T @ P) - A

def mfg_riccati(A, B, Q, Γ, ρ, tol=1e-8):
    """
    Solve  P B Γ^{-1} B' P = Q + ρ P + P A + A' P  for the solution P whose
    closed-loop matrix B Γ^{-1} B' P - A has eigenvalues with real parts
    below ρ/2, and raise ValueError if there is none.

    Passing Q + 0_X and Γ + 0_A gives the equilibrium of the mean field game;
    passing Q and Γ gives the single agent's value function curvature.
    """
```

(continues on next page)

(continued from previous page)

```
    n = A.shape[0]
    try:
        P = -solve_continuous_are(-(A + ρ/2 * np.eye(n)), B, Q, Γ)
    except (np.linalg.LinAlgError, ValueError):
        raise ValueError("no admissible solution of the Riccati equation")
    quad = P @ B @ np.linalg.solve(Γ, B.T) @ P
    residual = quad - (Q + ρ*P + P @ A + A.T @ P)
    scale = 1 + np.abs(Q).max() + np.abs(quad).max()
    if (np.abs(residual).max() > tol*scale
            or np.linalg.eigvals(closed_loop(A, B, Γ, P)).real.max() >= ρ/2):
        raise ValueError("no admissible solution of the Riccati equation")
    return P
```

Let's set up a two-dimensional example with complementarity in states and substitutability in actions, the configuration that Alvarez and Argente [2026] obtain from an industry equilibrium with capital accumulation.

```
ρ = 0.05
A = np.array([[0.4, 0.1],
              [0.0, 0.3]])
B = np.eye(2)
Q = np.array([[1.0, 0.2],
              [0.2, 0.8]])
Γ = np.array([[1.0, 0.1],
              [0.1, 1.2]])
Θ_X = np.array([[-0.3, 0.05],      # negative definite: complements in states
                [0.05, -0.2]])
Θ_A = np.array([[0.2, 0.0],        # positive definite: substitutes in actions
                [0.0, 0.1]])

β2 = mfg_riccati(A, B, Q, Γ, ρ)                  # single agent
P = mfg_riccati(A, B, Q + Θ_X, Γ + Θ_A, ρ)       # equilibrium

print("individual curvature β₂ =\n", β2.round(4))
print("\nequilibrium matrix P =\n", P.round(4))
```

```
individual curvature β₂ =
 [[-0.6614 -0.0869]
 [-0.0869 -0.654 ]]

equilibrium matrix P =
 [[-0.5353 -0.134 ]
 [-0.134  -0.5337]]
```

Let's compare aggregate dynamics with individual dynamics, that is, with the dynamics of an agent's deviation from the mean.

```
D = B @ np.linalg.solve(Γ + Θ_A, B.T)

U_ind = closed_loop(A, B, Γ, β2)          # individual dynamics
U_agg = closed_loop(A, B, Γ + Θ_A, P)     # aggregate dynamics

print("individual eigenvalues:", np.linalg.eigvals(U_ind).round(4))
print("aggregate eigenvalues: ", np.linalg.eigvals(U_agg).round(4))
```

```
individual eigenvalues: [-1.0698 -0.8321]
aggregate eigenvalues:  [-0.9025 -0.6423]
```

The aggregate eigenvalues are closer to zero, so the economy as a whole adjusts more slowly than an agent's deviation from it.

This example has complementarity in states and substitutability in actions, and *Persistence* shows that both raise the sum of the eigenvalues; here each eigenvalue moves toward zero as well.

129.6.1 Checking the Hamiltonian structure

Proposition 129.4.1 says the eigenvalues of $\mathcal{H}$ come in pairs $\{\lambda, \rho - \lambda\}$, and that the n eigenvalues with real parts below $\rho/2$ are the ones that govern equilibrium dynamics.

Let's check both claims.

```
BΓB = B @ np.linalg.solve(Γ, B.T)
n = A.shape[0]

Ham = np.block([[ρ*np.eye(n) + A.T - β2 @ D, Θ_X + β2 @ (BΓB - D) @ β2],
                [D,                          -A + D @ β2]])

ev = np.linalg.eigvals(Ham)
print("eigenvalues of ℋ:", np.sort(ev.real).round(4))
print("paired as λ and ρ - λ:",
      np.allclose(np.sort(ev.real), np.sort(ρ - ev.real)))

stable = np.sort(ev.real[ev.real < ρ/2])
print("\nstable half of ℋ:      ", stable.round(4))
print("closed-loop eigenvalues:", np.sort(np.linalg.eigvals(U_agg).real).round(4))
```

```
eigenvalues of ℋ: [-0.9025 -0.6423  0.6923  0.9525]
paired as λ and ρ - λ: True

stable half of ℋ:       [-0.9025 -0.6423]
closed-loop eigenvalues: [-0.9025 -0.6423]
```

The saddle path is the stable invariant subspace of $\mathcal{H}$, exactly as in *Lagrangian for LQ Control*.

129.7 The scalar case

With $n = k = 1$ everything is explicit.

Write $q, a, b, \gamma, \theta_X, \theta_{\mathcal{A}}$ for the scalars, and assume $\gamma + \theta_{\mathcal{A}} > 0$.

The Riccati equation becomes a quadratic, and the admissible root gives the aggregate eigenvalue

$$\lambda = \frac{\rho}{2} - \sqrt{\left(\frac{\rho}{2} + a\right)^2 + \frac{b^2(q + \theta_X)}{\gamma + \theta_{\mathcal{A}}}}. \tag{129.11}$$

When the initial mean $X(0)$ is not zero, an equilibrium exists if and only if the term under the square root is positive,

$$q + \theta_X + \left(\frac{\rho}{2} + a\right)^2 \frac{\gamma + \theta_{\mathcal{A}}}{b^2} > 0, \tag{129.12}$$

and the equilibrium is stable, meaning $\lambda < 0$, if and only if

$$q + \theta_X + a(a + \rho)\frac{\gamma + \theta_{\mathcal{A}}}{b^2} > 0. \tag{129.13}$$

When $q + \theta_X > 0$, as in both of our examples, formula (129.11) displays the two comparative statics at a glance.

More complementarity in *states* (a smaller θ_X) raises λ and makes aggregate dynamics *more* persistent: when others stay away from the steady state, each agent has less reason to return to it.

More complementarity in *actions* (a smaller $\theta_{\mathcal{A}}$) lowers λ and makes dynamics *less* persistent: when others adjust, each agent wants to adjust too.

If $q + \theta_X < 0$, the second comparative static reverses, which is why the general results below assume positive effective curvatures.

129.7.1 Four regions

Conditions (129.12) and (129.13) define two thresholds for the state interaction,

$$\theta^* \equiv -q - a(a+\rho)\frac{\gamma + \theta_{\mathcal{A}}}{b^2}, \qquad \theta^{**} \equiv -q - \left(\frac{\rho}{2} + a\right)^2 \frac{\gamma + \theta_{\mathcal{A}}}{b^2} = \theta^* - \left(\frac{\rho}{2}\right)^2 \frac{\gamma + \theta_{\mathcal{A}}}{b^2},$$

so that $\theta^{**} < \theta^*$.

As θ_X falls, that is, as complementarity in states increases, the economy passes through four regions:

1. if $\theta_X > \theta^*$, the equilibrium is stable and $X(t)$ converges to zero
2. if $\theta_X = \theta^*$, then $\lambda = 0$ and $X(t)$ stays where it starts
3. if $\theta^{**} < \theta_X < \theta^*$, an equilibrium exists but $0 < \lambda < \rho/2$, so $X(t)$ moves away from zero, slowly enough to keep discounted payoffs finite
4. if $\theta_X \leq \theta^{**}$, there is no equilibrium

So stability is lost before existence.

Formulas (129.11) to (129.13) depend on the two interactions only through the ratio $(q + \theta_X)/(\gamma + \theta_{\mathcal{A}})$.

In the $(\theta_X, \theta_{\mathcal{A}})$ plane, every curve along which that ratio is constant is a ray from the point $(-q, -\gamma)$.

So contours of the aggregate eigenvalue, the stability boundary and the existence boundary are all rays from that point.

Let's draw them, taking a large discount rate, $\rho = 0.5$, to widen the band between the two thresholds.

We first write the closed form, and check it against the solver.

```
def scalar_λ(ρ, a, b, q, γ, θ_X, θ_A):
    "Closed-form aggregate eigenvalue in the scalar case."
    return ρ/2 - np.sqrt((ρ/2 + a)**2 + b**2*(q + θ_X)/(γ + θ_A))

def half_life(λ):
    "Time for a state with eigenvalue λ < 0 to close half of a gap."
    return np.log(2)/(-λ)

ρ_s, a, b, q, γ = 0.5, 0.3, 1.0, 1.0, 1.0

errors = []
for θ_X, θ_A in ((0.0, 0.0), (-0.5, 0.0), (0.0, 0.5), (-0.5, 0.5)):
    P_s = mfg_riccati(np.array([[a]]), np.array([[b]]),
                      np.array([[q + θ_X]]), np.array([[γ + θ_A]]), ρ_s)
    λ_num = closed_loop(np.array([[a]]), np.array([[b]]),
                        np.array([[γ + θ_A]]), P_s)[0, 0]
    errors.append(abs(λ_num - scalar_λ(ρ_s, a, b, q, γ, θ_X, θ_A)))
print(f"largest gap between solver and closed form: {max(errors):.1e}")
```

```
largest gap between solver and closed form: 2.2e-16
```

The next figure shades the four regions and draws contours of the half-life of the aggregate state.

```
θ_X_grid = np.linspace(-2.0, 1.0, 400)
θ_A_grid = np.linspace(-0.99, 2.0, 400)
TX, TA = np.meshgrid(θ_X_grid, θ_A_grid)

inside = (ρ_s/2 + a)**2 + b**2*(q + TX)/(γ + TA)
λ_plane = np.where(inside >= 0, ρ_s/2 - np.sqrt(np.maximum(inside, 0)), np.nan)
region = np.where(inside < 0, 2, np.where(λ_plane >= 0, 1, 0))
hl_plane = np.where(λ_plane < 0, np.log(2)/np.abs(λ_plane), np.nan)

def θ_X_on_contour(h, θ_A):
    "The value of θ_X at which the half-life equals h, given θ_A."
    λ = -np.log(2)/h
    ratio = ((ρ_s/2 - λ)**2 - (ρ_s/2 + a)**2)/b**2
    return ratio*(γ + θ_A) - q

levels = [0.75, 1, 1.5, 2, 3, 5]
colors = ['white', '#fdd0a2', '#d9d9d9']
fig, ax = plt.subplots(figsize=(10, 5.5))
ax.contourf(TX, TA, region, levels=[-0.5, 0.5, 1.5, 2.5], colors=colors)
cs = ax.contour(TX, TA, hl_plane, levels=levels, colors='C0', linewidths=1.2)
ax.clabel(cs, fmt='%g', fontsize=9,
          manual=[(θ_X_on_contour(h, 0.6), 0.6) for h in levels])
θ_A_line = np.linspace(-γ, 2.0, 50)
ax.plot(-q - a*(a + ρ_s)*(γ + θ_A_line)/b**2, θ_A_line, 'k', lw=1.5)
ax.axhspan(-1.2, -γ, facecolor='none', hatch='//', edgecolor='grey', lw=0)
ax.plot(-q, -γ, 'ko', ms=5)
ax.annotate(r'$(-q, -\gamma)$', (-q, -γ), xytext=(-q + 0.06, -γ + 0.06))
handles = [Patch(facecolor='white', edgecolor='k', lw=0.5, label='stable'),
           Line2D([0], [0], color='k', lw=1.5, label=r'unit root, $\lambda = 0$'),
           Patch(facecolor=colors[1], label='divergent but admissible'),
           Patch(facecolor=colors[2], label='no equilibrium'),
           Patch(facecolor='none', hatch='//', edgecolor='grey',
                 label=r'excluded, $\gamma + \theta_{\mathcal{A}} \leq 0$'),
           Line2D([0], [0], color='C0', lw=1.2, label='half-life contours, years')]
ax.legend(handles=handles, loc='center left', bbox_to_anchor=(1.02, 0.5))
ax.set_xlim(-2.0, 1.0)
ax.set_ylim(-1.2, 2.0)
ax.set_xlabel(r'$\theta_X$')
ax.set_ylabel(r'$\theta_{\mathcal{A}}$')
plt.tight_layout()
plt.show()
```

Because the half-life contours are rays from $(-q, -\gamma)$, a move toward complementarity in states can be exactly offset by a move toward substitutability in actions.

The black line is the stability boundary $\theta_X = \theta^*$, and the thin orange band between it and the existence boundary is the third region, in which equilibrium paths diverge.

Our solver respects the boundaries: with $\theta_{\mathcal{A}} = 0$, it returns an admissible solution with $\lambda > 0$ inside the band and raises an error beyond it.

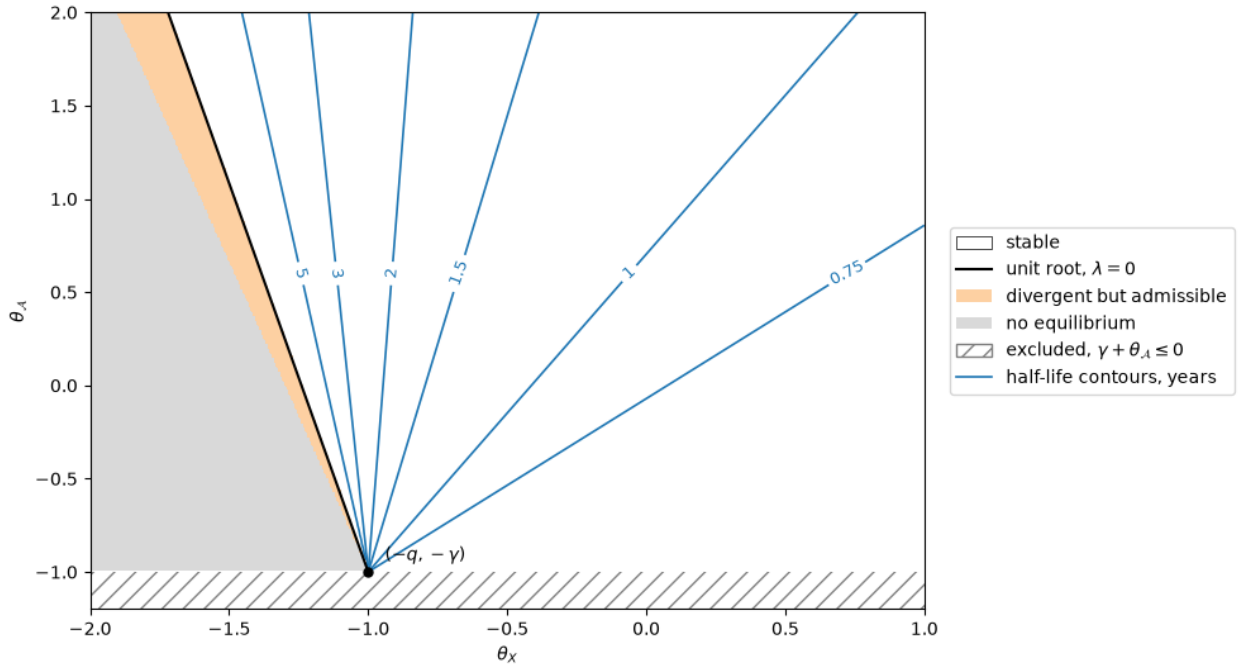


Fig. 129.1: Scalar model regions and half-lives

```
θ_star = -q - a*(a + ρ_s)*γ/b**2
θ_2star = θ_star - (ρ_s/2)**2*γ/b**2

for θ_X in ((θ_star + θ_2star)/2, θ_2star - 0.02):
    try:
        P_s = mfg_riccati(np.array([[a]]), np.array([[b]]),
                          np.array([[q + θ_X]]), np.array([[γ]]), ρ_s)
        λ_num = closed_loop(np.array([[a]]), np.array([[b]]), np.array([[γ]]), P_s)[0,
↪ 0]
        print(f"θ_X = {θ_X:.4f}: λ = {λ_num:.4f}")
    except ValueError as e:
        print(f"θ_X = {θ_X:.4f}: {e}")
```

```
θ_X = -1.2712: λ = 0.0732
θ_X = -1.3225: no admissible solution of the Riccati equation
```

Exercise 129.14.1 asks you to derive (129.11), (129.12) and (129.13).

129.8 An industry equilibrium with capital accumulation

The scalar model is not a toy.

Alvarez and Argente [2026] show that it describes an industry equilibrium with capital accumulation in which *both* kinds of interaction appear, with opposite signs.

There is a continuum of monopolistically competitive firms.

A firm with capital k produces $y = k^\nu$ with $0 < \nu < 1$, and a constant returns sector aggregates the differentiated goods with elasticity of substitution $\eta > 1$.

With the final good as numeraire, a firm that produces y when industry output is Y earns revenue proportional to $y^{1-1/\eta} Y^{1/\eta}$.

So if all other firms hold capital K, operating profit is proportional to

$$\Pi(k, K) = k^{\nu(1-1/\eta)} K^{\nu/\eta}. \tag{129.14}$$

Capital evolves according to

$$dk = (i - \delta k)dt + k\sigma dW.$$

The firm buys investment goods at a price $\mathcal{P}(I)$ that depends on aggregate investment I, and it pays a convex adjustment cost $\psi(i)$, so it chooses investment to maximize

$$\mathbb{E} \int_0^\infty e^{-\rho t} \left[\Pi(k, K) - \mathcal{P}(I)\, i - \psi(i) \right] dt.$$

Let $\bar{i} = \delta \bar{k}$ be steady-state investment, normalize $\mathcal{P}(\bar{i}) = 1$ and $\psi'(\bar{i}) = 0$, and write percentage deviations from the deterministic steady state as

$$x = \frac{k - \bar{k}}{\bar{k}}, \qquad X = \frac{K - \bar{k}}{\bar{k}}, \qquad \alpha = \frac{i - \bar{i}}{\bar{i}}, \qquad \mathcal{A} = \frac{I - \bar{i}}{\bar{i}}.$$

A second-order expansion of the return around the steady state, with the objective normalized by $\bar{\Pi} \equiv \Pi(\bar{k}, \bar{k})$, delivers exactly (129.1) with

$$q = -\frac{\bar{k}^2 \Pi_{kk}}{\bar{\Pi}}, \qquad \theta_X = -\frac{\bar{k}^2 \Pi_{kK}}{\bar{\Pi}}, \qquad \gamma = \frac{\delta^2 \bar{k}^2 \psi''(\bar{i})}{\bar{\Pi}}, \qquad \theta_{\mathcal{A}} = \frac{\delta^2 \bar{k}^2 \mathcal{P}'(\bar{i})}{\bar{\Pi}},$$

while the state equation becomes $dx = \delta(\alpha - x)dt + \sigma dW$ to first order, so that

$$a = b = \delta.$$

Differentiating (129.14) gives closed forms:

$$\begin{aligned} q &= \nu \frac{\eta - 1}{\eta^2} \left[\eta(1 - \nu) + \nu \right] > 0, \\ \theta_X &= -\frac{\eta - 1}{\eta} \frac{\nu^2}{\eta} < 0, \\ q + \theta_X &= \frac{\eta - 1}{\eta} \nu (1 - \nu) > 0. \end{aligned} \tag{129.15}$$

Three features of (129.15) deserve emphasis.

First, $\theta_X < 0$: capital stocks are strategic *complements*, because a larger industry capital stock raises the demand shifter $Y^{1/\eta}$ and hence the marginal profitability of a firm's own capital.

Second, if the supply of investment goods slopes upward then $\mathcal{P}'(\bar{i}) > 0$ and so $\theta_{\mathcal{A}} > 0$: investment rates are strategic *substitutes*, because everyone investing at once bids up the price of capital goods.

Third, $q + \theta_X > 0$ for every $\eta > 1$ and every $\nu \in (0, 1)$, so by *Proposition 129.5.3* an equilibrium exists and is unique no matter how strong market power is and no matter how close returns to scale come to constant.

Let's put these formulas into code, and check them against numerical derivatives of the profit function itself.

```
def cap_q(η, ν):
    "Own-state curvature in the capital accumulation example."
    return ν*(η - 1)/η**2*(η*(1 - ν) + ν)

def cap_θ_X(η, ν):
    "State interaction in the capital accumulation example."
    return -(η - 1)/η*ν**2/η

η_c, ν_c = 4.0, 0.7

Π = lambda k, K: k**(ν_c*(1 - 1/η_c))*K**(ν_c/η_c)

h = 1e-5
Π_kk = (Π(1+h, 1) - 2*Π(1, 1) + Π(1-h, 1))/h**2
Π_kK = (Π(1+h, 1+h) - Π(1+h, 1-h) - Π(1-h, 1+h) + Π(1-h, 1-h))/(4*h**2)

print(f"{'':>10}{'finite difference':>20}{'closed form':>15}")
print(f"{'q':>10}{-Π_kk/Π(1, 1):>20.6f}{cap_q(η_c, ν_c):>15.6f}")
print(f"{'θ_X':>10}{-Π_kK/Π(1, 1):>20.6f}{cap_θ_X(η_c, ν_c):>15.6f}")
print(f"{'q + θ_X':>10}{-(Π_kk + Π_kK)/Π(1, 1):>20.6f}"
      f"{(η_c - 1)/η_c*ν_c*(1 - ν_c):>15.6f}")
```

```
             finite difference    closed form
         q            0.249374       0.249375
       θ_X           -0.091875      -0.091875
   q + θ_X            0.157499       0.157500
```

The effective curvature of *Effective curvature* has a concrete meaning here.

The next figure plots the marginal return to a firm's capital, measured relative to $\bar{\Pi}/\bar{k}$, in three experiments.

In the first, one firm changes its capital while the rest of the industry stays at $\bar{k}$, and the marginal return falls with slope $-q$.

In the second, every firm changes its capital together, and the marginal return falls with the smaller slope $-(q+\theta_X)$, because a larger industry raises the demand for each firm's good.

The third line is the marginal revenue of the whole industry when every firm holds the same capital, $\frac{d}{dk}\Pi(k,k) = \nu k^{\nu-1}$.

The second line is exactly $(\eta-1)/\eta$ times the third, because a firm's marginal revenue is only that fraction of the value of its marginal product, so the markup lowers both the level and the slope of the private return.

The third line will matter when we come to the planner.

```
Π_k = lambda k, K: ν_c*(1 - 1/η_c)*k**(ν_c*(1 - 1/η_c) - 1)*K**(ν_c/η_c)
k_grid = np.linspace(0.8, 1.2, 200)       # k̄ = 1, so that Π = 1

fig, ax = plt.subplots(figsize=(8, 4.8))
ax.plot(k_grid, Π_k(k_grid, 1.0), lw=2,
        label=r'$\Pi_k(k, \bar k)$: one firm moves, slope $-q$')
ax.plot(k_grid, Π_k(k_grid, k_grid), lw=2,
        label=r'$\Pi_k(k, k)$: all firms move, slope $-(q+\theta_X)$')
ax.plot(k_grid, ν_c*k_grid**(ν_c - 1), lw=2, ls='--',
        label=r'$\frac{d}{dk}\Pi(k, k)$: industry, slope $-\nu(1-\nu)$')
ax.set_xlabel(r'capital relative to steady state, $k/\bar k$')
ax.set_ylabel(r'marginal return, relative to $\bar\Pi/\bar k$')
ax.set_ylim(0.45, 0.88)
ax.legend(loc='upper right')
```

(continues on next page)

(continued from previous page)

```
plt.tight_layout()
plt.show()
```

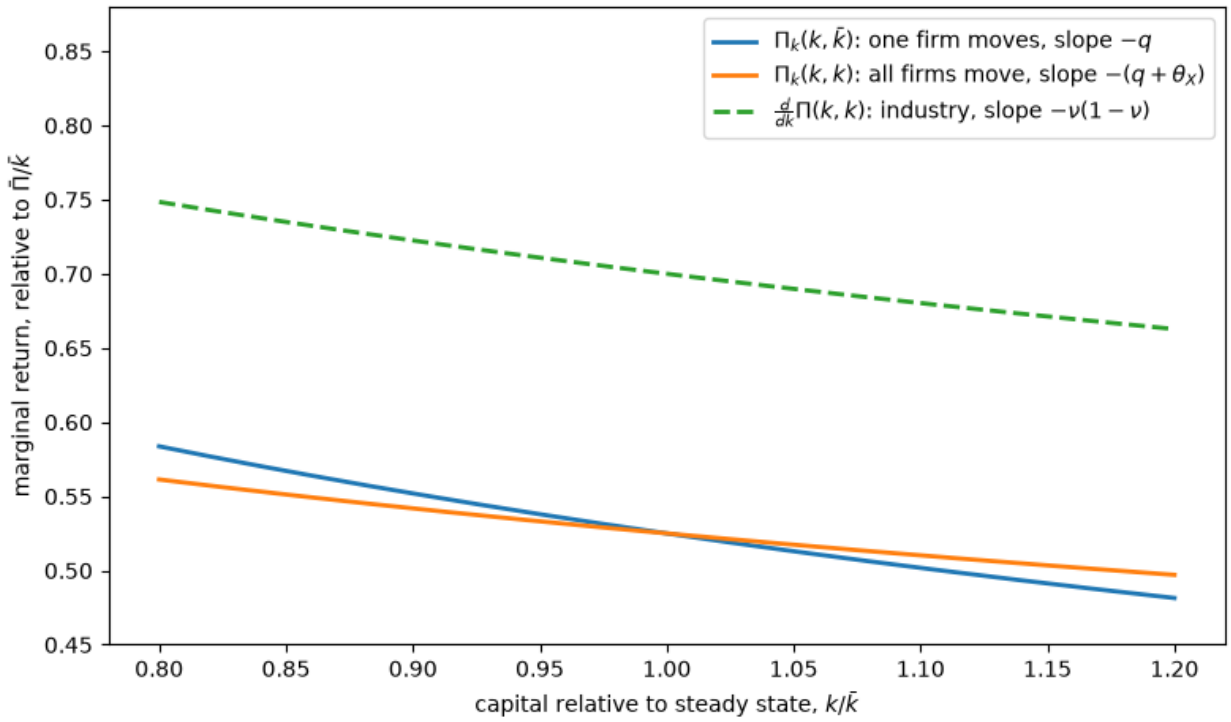


Fig. 129.2: Marginal returns to capital

129.8.1 Calibration

Take annual units, $\delta = 0.10$, $\rho = 0.05$, an elasticity of substitution $\eta = 4$, and returns to scale $\nu = 0.7$.

For investment costs we use a quadratic adjustment cost and a constant-elasticity supply of investment goods,

$$\psi(i) = \frac{\phi}{2}\,\bar{i}\left(\frac{i - \bar{i}}{\bar{i}}\right)^2, \qquad \mathcal{P}(I) = \left(\frac{I}{\bar{i}}\right)^{1/\varepsilon_s},$$

which satisfy the normalizations $\psi'(\bar{i}) = 0$ and $\mathcal{P}(\bar{i}) = 1$.

The steady-state Euler equation $\Pi_k(\bar{k}, \bar{k}) = \rho + \delta$ gives $\bar{k}^{1-\nu} = \nu(\eta - 1)/[\eta(\rho + \delta)]$, and substituting into the formulas for γ and $\theta_{\mathcal{A}}$ gives

$$\gamma = \phi\,\delta\bar{k}^{1-\nu}, \qquad \theta_{\mathcal{A}} = \frac{\delta\bar{k}^{1-\nu}}{\varepsilon_s}.$$

We choose $\phi = 1/7$, so that a firm that ignores the industry closes half of a capital gap in three years, and $\varepsilon_s = 35/9 \approx 3.9$, a moderately elastic supply of investment goods chosen so that $\theta_{\mathcal{A}} = 0.09$.

With these values $\gamma = 0.05$.

Exercise 129.14.2 asks you to derive these formulas.

```
ρ_c, δ_c = 0.05, 0.10
ψ_c, ε_s_c = 1/7, 35/9

def cap_scale(η, ν, ρ=ρ_c, δ=δ_c):
    "δ k̄^{1-ν}, which converts ψ and 1/ε_s into γ and θ_A."
    return δ*ν*(η - 1)/(η*(ρ + δ))

γ_c = ψ_c*cap_scale(η_c, ν_c)
θ_A_c = cap_scale(η_c, ν_c)/ε_s_c
q_c, θ_X_c = cap_q(η_c, ν_c), cap_θ_X(η_c, ν_c)

λ_c = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, θ_X_c, θ_A_c)

θ_star_c = -q_c - δ_c*(δ_c + ρ_c)*(γ_c + θ_A_c)/δ_c**2   # stability threshold

print(f"γ = {γ_c:.4f},  θ_A = {θ_A_c:.4f}")
print(f"q = {q_c:.4f},  θ_X = {θ_X_c:.4f},  q + θ_X = {q_c + θ_X_c:.4f}")
print(f"stability threshold θ* = {θ_star_c:.4f}")
print(f"half-life of aggregate capital = {half_life(λ_c):.2f} years")
```

```
γ = 0.0500,  θ_A = 0.0900
q = 0.2494,  θ_X = -0.0919,  q + θ_X = 0.1575
stability threshold θ* = -0.4594
half-life of aggregate capital = 4.99 years
```

Since the technology delivers $q + \theta_X > 0$ and an upward-sloping supply of investment goods delivers $\theta_{\mathcal{A}} > 0$, this example is always in the first of the four regions.

Complementarity in states would have to be five times stronger than the technology implies before the industry's capital stock stopped converging.

129.8.2 How much do the interactions matter?

Because (129.11) depends on θ_X and $\theta_{\mathcal{A}}$ separately, we can switch each interaction off and read the answer.

```
cases = {'no interactions':              (0.0,   0.0),
         'state complementarity only':   (θ_X_c, 0.0),
         'action substitutability only': (0.0,   θ_A_c),
         'equilibrium':                  (θ_X_c, θ_A_c)}

print(f"{'':>30}{'λ':>10}{'half-life':>12}")
for label, (tx, ta) in cases.items():
    λ_case = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, tx, ta)
    print(f"{label:>30}{λ_case:>10.4f}{half_life(λ_case):>12.2f}")
```

```
                                    λ   half-life
                no interactions   -0.2309        3.00
     state complementarity only   -0.1921        3.61
   action substitutability only   -0.1579        4.39
                    equilibrium   -0.1389        4.99
```

Both interactions slow aggregate capital adjustment, and they do so for different reasons.

Complementarity in states means that a firm has less reason to rebuild its capital while the rest of the industry is still below its steady state.

Substitutability in actions means that a burst of industry investment is expensive, so firms spread their investment over time.

Together they stretch the half-life of the industry's capital stock from three years to five.

We will see in *The planner in the capital accumulation example* that a planner would adjust slightly faster than the industry does, once all the terms in the planner's objective are counted.

129.8.3 Individual and aggregate adjustment speeds

By (129.10), a firm's deviation from the industry mean adjusts at the speed of a firm that ignores the industry, so the first row of the table is also the individual half-life: three years, against five for the industry, a ratio of 1.66.

Because state complementarity and action substitutability both slow aggregate adjustment here, individual capital reverts to the mean faster than aggregate capital does.

A researcher who estimated the speed of capital adjustment from firm-level data, and used it to predict the industry's response to an industry-wide shock, would predict a half-life of three years rather than five.

That is the missing intercept of *Aggregate shocks and identification* at work.

The gap is a pure interaction effect: the same technology and the same adjustment costs generate both numbers.

129.8.4 Market power and returns to scale

How does the industry's adjustment speed depend on the two technological parameters?

Because γ and $\theta_{\mathcal{A}}$ are measured relative to profits, they depend on η and ν through $\bar{k}$, so we hold the structural parameters ϕ and ε_s fixed rather than γ and $\theta_{\mathcal{A}}$.

That embeds an assumption of our calibration: the cost of a given percentage change in investment is proportional to the scale of investment $\bar{i}$.

Then

$$\frac{q + \theta_X}{\gamma + \theta_{\mathcal{A}}} = \frac{(1-\nu)(\rho+\delta)}{\delta(\phi + 1/\varepsilon_s)},$$

in which η does not appear: the effective curvature and the investment costs, both measured relative to profits, are proportional to $(\eta - 1)/\eta$.

So by (129.11), market power has no effect on the industry's adjustment speed.

The invariance depends on the scaling of adjustment costs: with a cost $\frac{c}{2}(i - \bar{i})^2$ and c fixed, adjustment costs relative to profits would grow with the steady state, which rises with η, and the industry's half-life would rise from 4.5 years at $\eta = 1.5$ to 5.3 years at $\eta = 12$.

It does affect a single firm's, because

$$\frac{q}{\gamma} = \frac{\rho + \delta}{\phi\,\delta}\left[(1-\nu) + \frac{\nu}{\eta}\right]$$

falls as η rises: with more elastic demand, a firm's revenue is closer to linear in its own capital, so its marginal return falls more slowly as it expands.

```
η_grid = np.linspace(1.2, 12, 200)
ν_grid = np.linspace(0.05, 0.995, 200)

def macro_micro(η, ν):
```

(continues on next page)

(continued from previous page)

```
    "Aggregate and individual half-lives, holding ψ and ε_s fixed."
    γ, θ_A = ψ_c*cap_scale(η, ν), cap_scale(η, ν)/ε_s_c
    q, θ_X = cap_q(η, ν), cap_θ_X(η, ν)
    λ_agg = scalar_λ(ρ_c, δ_c, δ_c, q, γ, θ_X, θ_A)
    λ_ind = scalar_λ(ρ_c, δ_c, δ_c, q, γ, 0.0, 0.0)
    return half_life(λ_agg), half_life(λ_ind)

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))

hl_agg, hl_ind = np.array([macro_micro(η, ν_c) for η in η_grid]).T
axes[0].plot(η_grid, hl_agg, lw=2, label='aggregate (industry)')
axes[0].plot(η_grid, hl_ind, lw=2, ls='--', label='individual (firm)')
axes[0].set_xlabel('elasticity of substitution $\\eta$')
axes[0].set_title('varying $\\eta$, with $\\nu = 0.7$')

hl_agg, hl_ind = np.array([macro_micro(η_c, ν) for ν in ν_grid]).T
axes[1].plot(ν_grid, hl_agg, lw=2, label='aggregate (industry)')
axes[1].plot(ν_grid, hl_ind, lw=2, ls='--', label='individual (firm)')
axes[1].set_title('varying $\\nu$, with $\\eta = 4$')
axes[1].axhline(np.log(2)/δ_c, color='k', lw=1, alpha=0.6)
axes[1].set_xlabel('returns to scale $\\nu$')
axes[1].annotate('$\\ln 2/\\delta$', (0.08, np.log(2)/δ_c - 0.35))

for ax in axes:
    ax.set_ylabel('half-life in years')
    ax.legend(loc='lower right')
plt.tight_layout()
plt.show()
```

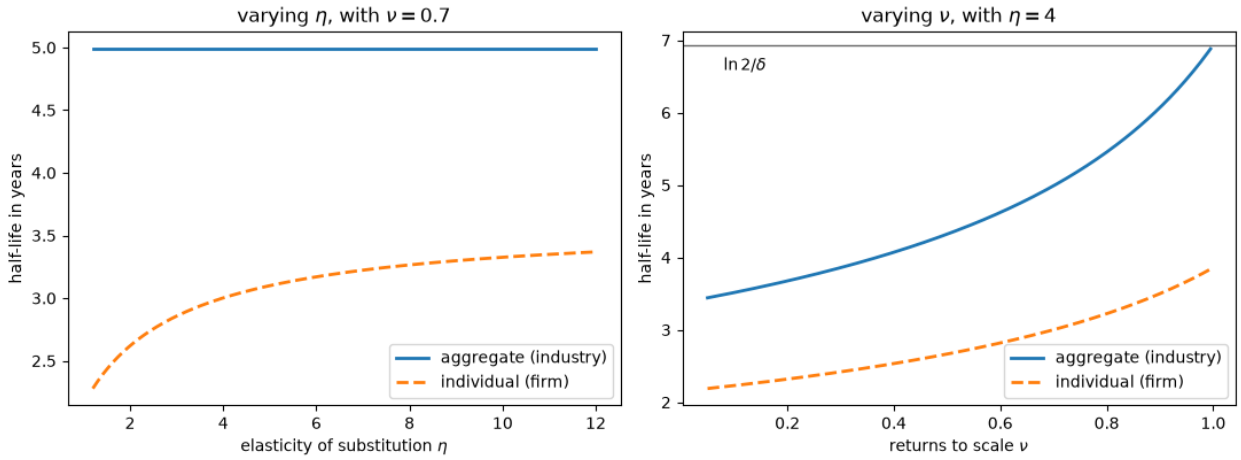


Fig. 129.3: Half-lives of industry and firm capital

The left panel shows both results: the industry's half-life does not move with η, while a single firm's rises from about 2.3 to 3.4 years as competition intensifies.

So the gap between individual and aggregate adjustment is widest when market power is greatest.

The right panel shows that both slow down as returns to scale rise.

As ν approaches 1 the effective curvature $q + \theta_X$ vanishes, and λ approaches $-\delta$: the industry's capital stock then returns

to its steady state almost entirely through depreciation, and aggregate investment barely responds.

Fig. 129.2 shows why: as ν approaches 1, the marginal return $\Pi_k(k, k)$ becomes flat along the diagonal, so an industry-wide capital gap creates almost no incentive to invest differently.

Exercise 129.14.3 asks you to derive these results.

129.8.5 Transition paths

Finally, let's trace out the industry's response to a capital stock that starts ten percent above its steady state.

Aggregate investment follows from the aggregate law of motion: $\dot{X} = \delta(\mathcal{A} - X) = \lambda X$ gives $\mathcal{A}(t) = \frac{\lambda+\delta}{\delta} X(t)$.

```
t_c = np.linspace(0, 25, 300)
X0_c = 0.10

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))
for label in ('no interactions', 'equilibrium'):
    tx, ta = cases[label]
    λ_case = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, tx, ta)
    path = X0_c*np.exp(λ_case*t_c)
    axes[0].plot(t_c, 100*path, lw=2, label=label)
    axes[1].plot(t_c, 100*(λ_case + δ_c)/δ_c*path, lw=2, label=label)

axes[0].set_ylabel('capital, % above steady state')
axes[1].set_ylabel('investment, % above steady state')
for ax in axes:
    ax.set_xlabel('years')
    ax.axhline(0, color='k', lw=0.8)
axes[0].legend()
plt.tight_layout()
plt.show()
```

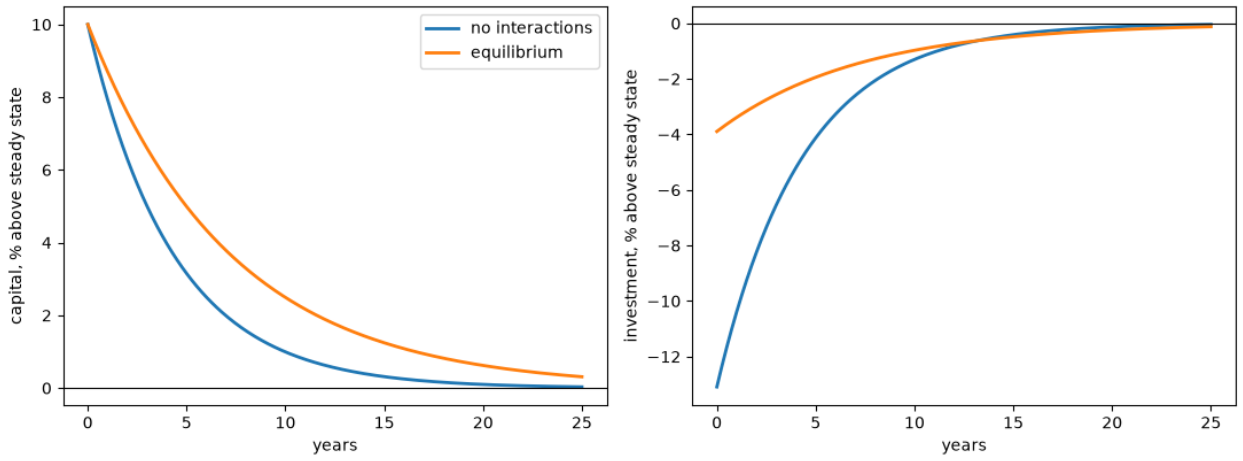


Fig. 129.4: Transition paths of capital and investment

The right panel is the industry's investment response, and it is where the two interactions show up most clearly.

A firm that ignored the industry would cut investment on impact by thirteen percent of its steady-state level; in equilibrium the cut is under four percent, and it is undone much more slowly.

The costate $\beta_1(t) + \bar{\beta}_2 x$ that supports these paths is the marginal value of installed capital, measured as a deviation from its steady-state value; divided by the price of investment goods, it is Tobin's marginal q, not to be confused with the curvature q.

Optimal Growth Under Uncertainty and Tobin's q studies the marginal value of capital when the irreversibility constraint $i \geq 0$ sometimes binds, and shows that the value function remains differentiable even then; the quadratic approximation here sets that nonlinearity aside.

129.9 Aggregate shocks and identification

Now add a shock that hits everyone at once.

Let $\mathcal{J}$ be a compensated jump process, the difference between two independent Poisson processes with equal intensities, so that its increments have mean zero, and suppose

$$dx = (B\alpha - Ax)dt + \Sigma^{1/2} dW + \Upsilon \, d\mathcal{J}.$$

A common Brownian motion would work the same way.

With common noise the value function must keep track of the aggregate state as well as the individual state, so it becomes $v(x, X)$.

Nevertheless, matching coefficients in the recursive HJB equation delivers a result worth emphasizing.

Proposition 129.9.1

With common noise, the matrix P governing aggregate dynamics solves the *same* equilibrium Riccati equation as before.

This is a certainty-equivalence result of the kind familiar from linear quadratic control.

The decomposition (129.10) continues to hold, with $\Upsilon \, d\mathcal{J}$ added to the law of motion of X, and the common shock cancels from the law of motion of an agent's deviation z.

129.9.1 The missing intercept

Researchers often estimate how agents respond to shocks from panel data with time fixed effects, which control for aggregate conditions.

Such regressions identify how an agent responds to her position relative to others, but not the common response of all agents to an aggregate shift, which the time effects absorb; this is the "missing intercept" [Wolf, 2023].

The decomposition (129.10) makes the point exact.

Writing an agent's action in terms of her own state and the mean gives

$$\alpha_{it} = U_{\alpha,z} \, x_{it} + \left(U_{\alpha,X} - U_{\alpha,z}\right) X_t, \qquad U_{\alpha,X} \equiv (\Gamma + \Theta_{\mathcal{A}})^{-1} B^\top P, \qquad U_{\alpha,z} \equiv \Gamma^{-1} B^\top \bar{\beta}_2.$$

A regression of actions on states with time effects estimates the slope $U_{\alpha,z}$, which contains no interaction matrix, while the time effects absorb $(U_{\alpha,X} - U_{\alpha,z})X_t$, which contains all of them.

Proposition 129.9.2

Data on $x(t) - X(t)$ and $\alpha(t) - \mathcal{A}(t)$ contain no information about Θ_X or $\Theta_{\mathcal{A}}$.

> **Note**
>
> Removing time effects from micro data removes exactly the variation that identifies strategic interaction.

To see this in the capital example, consider two industries that are identical except for the strength of complementarity in states: one has $\theta_X = 0$, and the other has twice the calibrated amount.

Let's compute the two policy coefficients for each.

```
def capital_policy(θ_X, θ_A=θ_A_c):
    "Coefficients of investment on X and on z in the capital example."
    one = np.array([[1.0]])
    P_s = mfg_riccati(δ_c*one, δ_c*one, (q_c + θ_X)*one, (γ_c + θ_A)*one, ρ_c)[0, 0]
    β2_s = mfg_riccati(δ_c*one, δ_c*one, q_c*one, γ_c*one, ρ_c)[0, 0]
    return δ_c*P_s/(γ_c + θ_A), δ_c*β2_s/γ_c

industries = (('θ_X = 0', 0.0), ('θ_X = 2 × calibrated', 2*θ_X_c))
print(f"{'':>22}{'slope on z':>12}{'slope on X':>12}")
for label, θ_X_use in industries:
    U_αX, U_αz = capital_policy(θ_X_use)
    print(f"{label:>22}{U_αz:>12.4f}{U_αX:>12.4f}")
```

```
                       slope on z  slope on X
              θ_X = 0     -1.3093     -0.5786
 θ_X = 2 × calibrated     -1.3093     -0.1752
```

The slope on a firm's deviation from the mean is the same in the two industries, and only the slope on the industry mean differs.

Now let's simulate panels of 400 firms from the two industries, using the same shocks for both.

Both start at their steady state with some dispersion across firms, and both are hit at $t = 2$ by a common shock that lowers every firm's capital by ten percent.

```
rng = np.random.default_rng(0)
N, T, dt, σ = 400, 20.0, 0.02, 0.05
steps = int(T/dt)
dW = rng.normal(size=(steps, N))*np.sqrt(dt)      # the same draws for both
x0 = 0.05*rng.normal(size=N)
t_sim = np.arange(steps)*dt

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))
deviations = []
for k, (label, θ_X_sim) in enumerate(industries):
    U_αX, U_αz = capital_policy(θ_X_sim)
    x = x0.copy()
    X_path, z_path = np.empty(steps), np.empty(steps)
    for s in range(steps):
        if s == int(2.0/dt):
            x = x - 0.10                          # common shock
        X = x.mean()
        α = U_αX*X + U_αz*(x - X)
        X_path[s], z_path[s] = X, x[0] - X
        x = x + δ_c*(α - x)*dt + σ*dW[s]
    deviations.append(z_path)
    tex_label = label.replace('θ_X', r'$\theta_X$')
    axes[0].plot(t_sim, 100*z_path, lw=2.5 - 1.5*k, ls='-' if k == 0 else '--',
```

(continues on next page)

(continued from previous page)

```
                  label=tex_label)
    axes[1].plot(t_sim, 100*X_path, lw=2, label=tex_label)

axes[0].set_title("one firm's deviation from the industry mean")
axes[1].set_title('industry mean')
for ax in axes:
    ax.set_xlabel('years')
    ax.set_ylabel('percent')
    ax.legend()
plt.tight_layout()
plt.show()

print("demeaned paths identical:", np.allclose(*deviations))
```

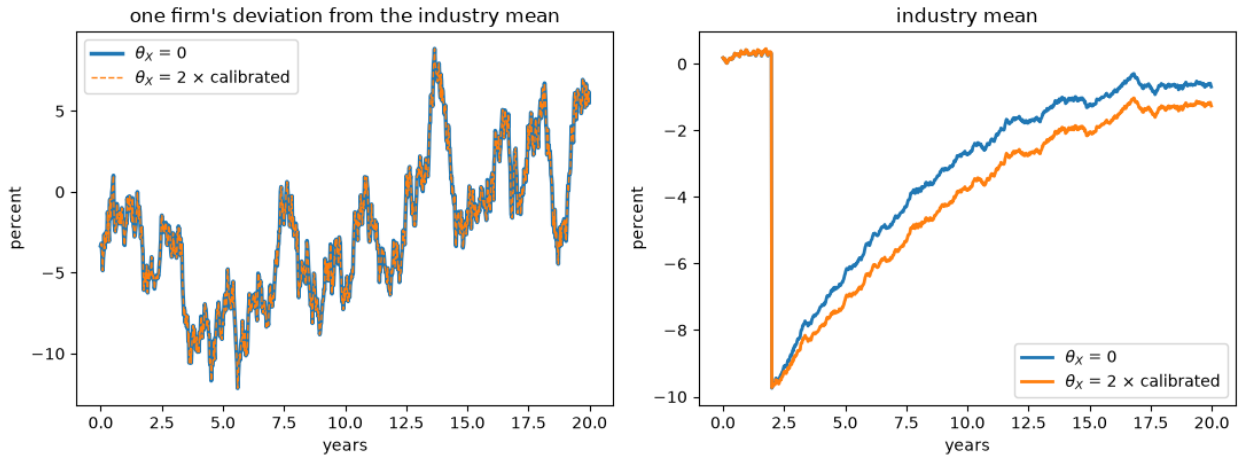


Fig. 129.5: Simulated panels from two industries

```
demeaned paths identical: True
```

The left panel plots one firm's deviation from its industry mean in each economy, and the two lines coincide exactly, so a regression on demeaned data would return identical estimates for the two industries.

The industry means in the right panel differ: after the common shock, the industry with stronger complementarity returns to its steady state more slowly.

The small wiggles in the industry means come from averaging the idiosyncratic shocks of only 400 firms.

The demeaned paths coincide by construction, because the slope on z does not involve θ_X, so the simulation illustrates *Proposition 129.9.2* rather than testing it: it shows what an econometrician would see.

129.9.2 What aggregate data reveal

Aggregate data, by contrast, do carry information about strategic interaction, and can be combined with micro data to recover it.

Collect the four reduced-form matrices that link the aggregate state, the deviations from it and the actions:

$$U_{\dot{X},X} = DP - A, \qquad U_{\dot{z},z} = B\Gamma^{-1}B^\top\bar{\beta}_2 - A, \qquad U_{\alpha,X} = (\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top P, \qquad U_{\alpha,z} = \Gamma^{-1}B^\top\bar{\beta}_2.$$

Proposition 129.9.3

Suppose $\Theta_{\mathcal{A}}$, Γ and ρ are known, that D and Σ are nonsingular, and that $\mathbb{E}[X(t)X(t)^\top]$ is nonsingular for some $t > 0$.

1. If Q, A and B are also known, a time series for X identifies Θ_X.
2. If B is also known, time series for X, z and $\mathcal{A}$ identify Q, Θ_X and A.
3. If $n = k$ and $U_{\alpha,z} - U_{\alpha,X}$ is nonsingular, time series for X, z, α and $\mathcal{A}$ identify Q, Θ_X, A and B.

In every case $\Theta_{\mathcal{A}}$ must be known.

Exercise 129.14.4 walks through the recovery, and shows why Θ_X and $\Theta_{\mathcal{A}}$ cannot be told apart without a normalization.

129.10 Persistence

How do the interactions change aggregate adjustment when $n > 1$?

Throughout this section we follow Alvarez and Argente [2026] and assume that $Q + \Theta_X$ and $\Gamma + \Theta_{\mathcal{A}}$ are positive definite in every economy we compare.

As the scalar case showed, the comparative statics for actions can reverse sign without that assumption.

Proposition 129.10.1

Compare two economies that differ in one interaction matrix only.

1. More complementarity in states makes the stabilizing solution P less negative definite and raises the trace of the closed-loop matrix.
2. More complementarity in actions also makes P less negative definite, but lowers the trace of the closed-loop matrix.

The two kinds of interaction move P in the same direction and the trace in opposite directions because they reach the closed-loop matrix $DP - A$ in different ways.

States act only through P.

Actions also change $D = B(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top$: more complementarity in actions raises D, so agents respond more strongly to a given marginal value, and the net effect lowers the trace.

The trace is the sum of the eigenvalues, and also the rate at which a set of initial aggregate states contracts in volume, so in a stable equilibrium the first force slows aggregate adjustment on average and the second speeds it up.

Let's check the first part in our two-dimensional example.

```
print(f"{'scale on Θ_X':>14}{'trace':>10}   eigenvalues")
for scale in (0.0, 0.5, 1.0, 1.5):
    P_s = mfg_riccati(A, B, Q + scale*Θ_X, Γ + Θ_A, ρ)
    U_s = closed_loop(A, B, Γ + Θ_A, P_s)
    print(f"{scale:>14}{np.trace(U_s):>10.4f}   {np.linalg.eigvals(U_s).real.round(4)}
 ↪")
```

```
  scale on Θ_X     trace   eigenvalues
           0.0   -1.7933   [-0.9956 -0.7977]
           0.5   -1.6747   [-0.9485 -0.7262]
           1.0   -1.5448   [-0.9025 -0.6423]
           1.5   -1.3988   [-0.857  -0.5418]
```

Raising the scale makes Θ_X more negative, that is, complementarity stronger, and the trace rises as claimed.

A rise in the trace says that the sum of the eigenvalues moves, not that each of them does.

Alvarez and Argente [2026] give a two-dimensional example in which a strict Loewner ordering of the Riccati solutions coexists with eigenvalues that move in opposite directions.

It sets $\rho = 0$, holds D fixed, and compares two effective state curvatures Q_1 and Q_2 with $Q_2 - Q_1$ positive definite, so that the first economy has more complementarity in states.

```
A_ce = np.array([[1.0, 2.0],
                 [2.0, 5.0]])
D_ce = np.array([[1.0, -3.0],
                 [-3.0, 10.0]])
Q_1 = np.array([[2.0, -1.0],
                [-1.0, 1.0]])
Q_2 = np.array([[3.0, 5.0],
                [5.0, 51.0]])

# with B = I, D = (Γ + Θ_A)^{-1}
Γ_ce = np.linalg.inv(D_ce)
P_ce = {}
for label, Q_i in (('Q_1', Q_1), ('Q_2', Q_2)):
    P_ce[label] = mfg_riccati(A_ce, np.eye(2), Q_i, Γ_ce, 0.0)
    U_i = closed_loop(A_ce, np.eye(2), Γ_ce, P_ce[label])
    print(f"{label}: trace {np.trace(U_i):.2f}, "
          f"closed-loop eigenvalues {np.sort(np.linalg.eigvals(U_i).real).round(2)}")

print("\neigenvalues of P_1 - P_2:", np.linalg.eigvalsh(P_ce['Q_1'] - P_ce['Q_2']).
 ↪round(3))
print("A and D commute:", np.allclose(A_ce @ D_ce, D_ce @ A_ce))
```

```
Q_1: trace -9.57, closed-loop eigenvalues [-6.55 -3.02]
Q_2: trace -24.25, closed-loop eigenvalues [-22.68  -1.56]

eigenvalues of P_1 - P_2: [0.281 1.833]
A and D commute: False
```

The Riccati solutions are strictly ordered, $P_1 \succ P_2$, and the first economy has the larger trace, as the theory requires, but only one of its eigenvalues is larger: one mode becomes more persistent and the other less.

What goes wrong is that A and D do not commute.

Since both are symmetric here, that is the same as saying that $D^{-1}A$ is not symmetric, so the closed-loop matrix $DP - A$ need not be similar to a symmetric matrix.

Two restrictions repair this.

Proposition 129.10.2

Hold D fixed and suppose that $D^{-1}A$ is symmetric.

Then more complementarity in states raises every eigenvalue of the closed-loop matrix.

The reason is that $DP - A = D(P - D^{-1}A)$ is similar to the symmetric matrix $D^{1/2}(P - D^{-1}A)D^{1/2}$, whose eigenvalues rise one by one when P rises in the Loewner order.

Proposition 129.10.3

Suppose $A = aI$.

Then the eigenvalues of the closed-loop matrix are real and given by

$$\lambda_j = \frac{\rho}{2} - \sqrt{\left(\frac{\rho}{2} + a\right)^2 + \mu_j},$$

where μ_j are the eigenvalues of $D^{1/2}(Q + \Theta_X)D^{1/2}$.

So more complementarity in states raises every λ_j, more complementarity in actions lowers every λ_j, and mode j is stable if and only if $\mu_j + a(a + \rho) > 0$.

The second result is still genuinely multidimensional.

Aggregate dynamics do decouple into scalar modes, but the coordinates that achieve this mix the eigenvectors of $Q + \Theta_X$ and D, and they differ from the coordinates that decouple individual dynamics, so the model as a whole does not separate into independent one-dimensional decisions.

Exercise 129.14.5 asks you to measure how often eigenvalue-by-eigenvalue monotonicity fails in random economies, and to confirm *Proposition 129.10.2*.

129.11 Multiproduct price setting with Kimball demand

Our second example is multidimensional, and it delivers a surprise.

Alvarez and Argente [2026] study a continuum of "stores", each selling n products with constant marginal costs c_j.

Products within a store are aggregated by a CES price index with elasticity $\bar{\eta}_d$, and stores are aggregated by a symmetric Kimball aggregator [Kimball, 1995].

Let $\eta_D(y)$ be the elasticity of a store's demand with respect to its relative price y, and let $\bar{\eta}_D > 1$ and $\bar{\eta}'_D$ denote its level and its derivative at the symmetric point $y = 1$.

The derivative $\bar{\eta}'_D$ is the **superelasticity** of demand, and it is the reason Kimball demand is so widely used in models of price setting: a positive superelasticity means that a store which raises its prices above the average faces a more elastic demand, which discourages it from moving away from the crowd.

That is strategic complementarity in prices, and it is the mechanism that Klenow and Willis [2016] and much of the subsequent literature rely on for real rigidity.

Writing x and X for log deviations (equivalently, to first order, percentage deviations) of a store's prices and of the average store's prices from the flexible-price level $\bar{p}_i = \bar{c}_i\bar{\eta}_D/(\bar{\eta}_D - 1)$, and $\bar{s}$ for the vector of steady-state expenditure shares,

the curvature matrices are

$$\begin{aligned} Q &= (\bar{\eta}_D - 1)\left[\left(\bar{\eta}_D - \bar{\eta}_d + \frac{\bar{\eta}'_D}{\bar{\eta}_D - 1}\right)\bar{s}\bar{s}^\top + \bar{\eta}_d \operatorname{diag}(\bar{s})\right], \\ \Theta_X &= -\bar{\eta}'_D\,\bar{s}\bar{s}^\top, \\ Q + \Theta_X &= (\bar{\eta}_D - 1)\left[(\bar{\eta}_D - \bar{\eta}_d)\bar{s}\bar{s}^\top + \bar{\eta}_d \operatorname{diag}(\bar{s})\right]. \end{aligned} \tag{129.16}$$

Stare at the third line.

The superelasticity appears in Q and in Θ_X, but it *cancels* from $Q + \Theta_X$.

By *Proposition 129.5.1*, the sum $Q + \Theta_X$ is the only channel through which either matrix reaches aggregate dynamics.

Actions are price changes, and stores pay quadratic Rotemberg costs of changing prices.

The matrix Γ lets that cost depend on which bundle of prices is changed, with negative off-diagonal elements representing economies of scope in repricing of the kind emphasized by Midrigan [2011] and Alvarez and Lippi [2014].

There is no interaction through aggregate actions, so $\Theta_{\mathcal{A}} = 0$.

With constant marginal costs $A = 0$, and to obtain a nontrivial drift we follow Alvarez and Argente [2026] in letting marginal costs, and hence flexible prices, grow at a common rate π.

Let $\bar{p}_i(t)$ be the trend price, which grows at rate π, and measure a price by its percentage deviation from trend, $x_i = p_i/\bar{p}_i - 1$, and the action by the change in the price minus the change in the trend, in units of the trend price, $\alpha_i = (\dot{p}_i - \dot{\bar{p}}_i)/\bar{p}_i$.

Then $\dot{x}_i = \alpha_i - \pi x_i$, so $A = \pi I$: a store that leaves its price change equal to the trend change sees its gap from trend eroded at rate π.

```
def kimball(η_d, η_D, η_D_prime, s):
    "Curvature and state-interaction matrices under Kimball demand."
    s = np.asarray(s, dtype=float)
    ss = np.outer(s, s)
    Q = (η_D - 1)*((η_D - η_d + η_D_prime/(η_D - 1))*ss + η_d*np.diag(s))
    Θ_X = -η_D_prime*ss
    return Q, Θ_X

s_bar = np.array([0.5, 0.3, 0.2])        # three products, unequal shares
η_d, η_D = 6.0, 4.0                      # within-store and across-store elasticities
ρ_K, π_K = 0.04, 0.02                    # discount rate and cost inflation

n_K = len(s_bar)
A_K = π_K*np.eye(n_K)
B_K = np.eye(n_K)
Γ_K = 20.0*(np.eye(n_K) - 0.25*(np.ones((n_K, n_K)) - np.eye(n_K)))

print("Rotemberg cost matrix Γ =\n", Γ_K)
print("\neigenvalues of Γ:", np.linalg.eigvalsh(Γ_K).round(4))
```

```
Rotemberg cost matrix Γ =
 [[20. -5. -5.]
 [-5. 20. -5.]
 [-5. -5. 20.]]

eigenvalues of Γ: [10. 25. 25.]
```

129.11.1 The superelasticity and aggregate dynamics

Now sweep the superelasticity across a wide range, including a negative value, and watch what changes and what does not.

```
print(f"{'η_D'':>6}  {'eigenvalues of Q':>28}  {'eigenvalues of Q + Θ_X':>28}")
for η_Dp in (-3.0, 0.0, 3.0, 10.0):
    Q_K, Θ_K = kimball(η_d, η_D, η_Dp, s_bar)
    print(f"{η_Dp:>6}  {str(np.linalg.eigvalsh(Q_K).round(3)):>28}"
          f"  {str(np.linalg.eigvalsh(Q_K + Θ_K).round(3)):>28}")
```

```
  η_D'              eigenvalues of Q        eigenvalues of Q + Θ_X
  -3.0           [2.598 4.5   7.482]           [3.127 4.766 7.827]
   0.0           [3.127 4.766 7.827]           [3.127 4.766 7.827]
   3.0           [3.435 5.096 8.329]           [3.127 4.766 7.827]
  10.0        [ 3.714  5.694 10.112]           [3.127 4.766 7.827]
```

The own-curvature matrix Q moves a great deal, and Θ_X moves with it, but the effective curvature does not move at all.

Aggregate price dynamics are therefore invariant to the superelasticity.

```
print(f"{'η_D'':>6}  {'aggregate eigenvalues':>34}  {'individual eigenvalues':>34}")
for η_Dp in (-3.0, 0.0, 3.0, 10.0):
    Q_K, Θ_K = kimball(η_d, η_D, η_Dp, s_bar)
    P_K = mfg_riccati(A_K, B_K, Q_K + Θ_K, Γ_K, ρ_K)
    β2_K = mfg_riccati(A_K, B_K, Q_K, Γ_K, ρ_K)
    ev_agg = np.sort(np.linalg.eigvals(closed_loop(A_K, B_K, Γ_K, P_K)).real)
    ev_ind = np.sort(np.linalg.eigvals(closed_loop(A_K, B_K, Γ_K, β2_K)).real)
    print(f"{η_Dp:>6}  {str(ev_agg.round(4)):>34}  {str(ev_ind.round(4)):>34}")
```

```
  η_D'              aggregate eigenvalues             individual eigenvalues
  -3.0           [-0.6527 -0.4882 -0.3836]          [-0.5813 -0.4752 -0.3815]
   0.0           [-0.6527 -0.4882 -0.3836]          [-0.6527 -0.4882 -0.3836]
   3.0           [-0.6527 -0.4882 -0.3836]          [-0.7234 -0.4928 -0.3844]
  10.0           [-0.6527 -0.4882 -0.3836]          [-0.8713 -0.4966 -0.3853]
```

The left block of numbers is identical in every row, and the right block is not.

The superelasticity changes how an individual store behaves without changing how the industry behaves, just as market power did in the capital accumulation example.

Almost all of the movement is in a single mode, because Θ_X is rank one and acts only through the store's price index $\bar{s}^\top x$.

The remaining modes, which involve relative prices within the store, barely move.

Because $A_K = \pi I$, the aggregate eigenvalues are also given in closed form by *Proposition 129.10.3*, with μ_j the eigenvalues of $(\Gamma + \Theta_{\mathcal{A}})^{-1}(Q + \Theta_X)$, which equal those of $D^{1/2}(Q + \Theta_X)D^{1/2}$ when $B = I$.

Let's confirm that the formula matches the solver.

```
Q_K, Θ_K = kimball(η_d, η_D, 3.0, s_bar)
P_K = mfg_riccati(A_K, B_K, Q_K + Θ_K, Γ_K, ρ_K)
ev_solver = np.sort(np.linalg.eigvals(closed_loop(A_K, B_K, Γ_K, P_K)).real)
μ = np.linalg.eigvals(np.linalg.solve(Γ_K, Q_K + Θ_K)).real
ev_formula = np.sort(ρ_K/2 - np.sqrt((ρ_K/2 + π_K)**2 + μ))
print(f"largest gap between formula and solver: {np.abs(ev_solver - ev_formula).
↪max():.1e}")
```

```
largest gap between formula and solver: 5.6e-16
```

129.11.2 Static complementarity

It is worth seeing how much static complementarity we are varying.

Maximizing the static profit function gives the best response $x^*(X) = -Q^{-1}\Theta_X X$, and Alvarez and Argente [2026] show that

$$\frac{\partial x_i^*(X)}{\partial X_j} = \bar{s}_j \frac{\kappa}{1+\kappa}, \qquad \kappa \equiv \frac{\bar{\eta}_D'}{(\bar{\eta}_D - 1)\bar{\eta}_D}. \tag{129.17}$$

We call the row sum of this matrix, $\kappa/(1+\kappa)$, **static pass-through**: it is the response of a store's prices to a uniform one percent rise in its competitors' prices.

The next cell confirms the formula and reports pass-through for each superelasticity.

```
print(f"{'η_D'':>6}{'pass-through':>14}   best response matrix, first row")
for η_Dp in (-3.0, 0.0, 3.0, 10.0):
    Q_K, Θ_K = kimball(η_d, η_D, η_Dp, s_bar)
    BR = -np.linalg.solve(Q_K, Θ_K)
    κ = η_Dp/((η_D - 1)*η_D)
    BR_closed = κ/(1 + κ)*np.outer(np.ones(n_K), s_bar)
    assert np.allclose(BR, BR_closed)
    print(f"{η_Dp:>6}{BR.sum(axis=1)[0]:>14.4f}   {BR[0].round(4)}")

print("\nΘ_X symmetric:           ", np.allclose(Θ_K, Θ_K.T))
print("best response symmetric:", np.allclose(BR, BR.T))
```

```
  η_D'  pass-through   best response matrix, first row
  -3.0       -0.3333   [-0.1667 -0.1    -0.0667]
   0.0        0.0000   [0. 0. 0.]
   3.0        0.2000   [0.1  0.06 0.04]
  10.0        0.4545   [0.2273 0.1364 0.0909]

Θ_X symmetric:           True
best response symmetric: False
```

Static pass-through runs from -33% to $+45\%$ across these rows, and it changes sign with the superelasticity, yet every one of these economies has exactly the same aggregate dynamics.

Note

Stronger static complementarity need not mean more aggregate persistence.

What matters for aggregate dynamics is the effective curvature $Q + \Theta_X$, not the interaction Θ_X or the static best responses on their own.

Notice also that the best response matrix is *not* symmetric, because shares differ across products, while Θ_X is symmetric, as promised in the description of the environment.

129.11.3 Aggregate and individual price paths

The decomposition (129.10) splits a store's prices into the industry average X and its own deviation $z = x - X$.

The first follows the equilibrium matrix P, the second the single-agent matrix $\bar{\beta}_2$.

So the invariance we found should show up as identical paths for the industry and different paths for a store that is out of line.

```
t_K = np.linspace(0, 8, 200)
shock = 0.10*np.ones(n_K)       # ten percent above target, all products

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))
for k, η_Dp in enumerate((-3.0, 0.0, 3.0, 10.0)):
    Q_K, Θ_K = kimball(η_d, η_D, η_Dp, s_bar)
    P_K = mfg_riccati(A_K, B_K, Q_K + Θ_K, Γ_K, ρ_K)
    β2_K = mfg_riccati(A_K, B_K, Q_K, Γ_K, ρ_K)
    U_agg_K = closed_loop(A_K, B_K, Γ_K, P_K)
    U_ind_K = closed_loop(A_K, B_K, Γ_K, β2_K)
    agg = np.array([s_bar @ expm(U_agg_K*t) @ shock for t in t_K])
    dev = np.array([s_bar @ expm(U_ind_K*t) @ shock for t in t_K])
    label = f"$\\bar\\eta_D' = {η_Dp}$"
    # decreasing line widths, so that four coincident curves remain visible
    axes[0].plot(t_K, 100*agg, lw=6 - 1.5*k, label=label)
    axes[1].plot(t_K, 100*dev, lw=2, label=label)

axes[0].set_title('industry price index, $\\bar s^\\top X(t)$')
axes[1].set_title("one store's deviation, $\\bar s^\\top z(t)$")
for ax in axes:
    ax.set_xlabel('years')
    ax.set_ylabel('percent above target')
    ax.legend()
plt.tight_layout()
plt.show()
```

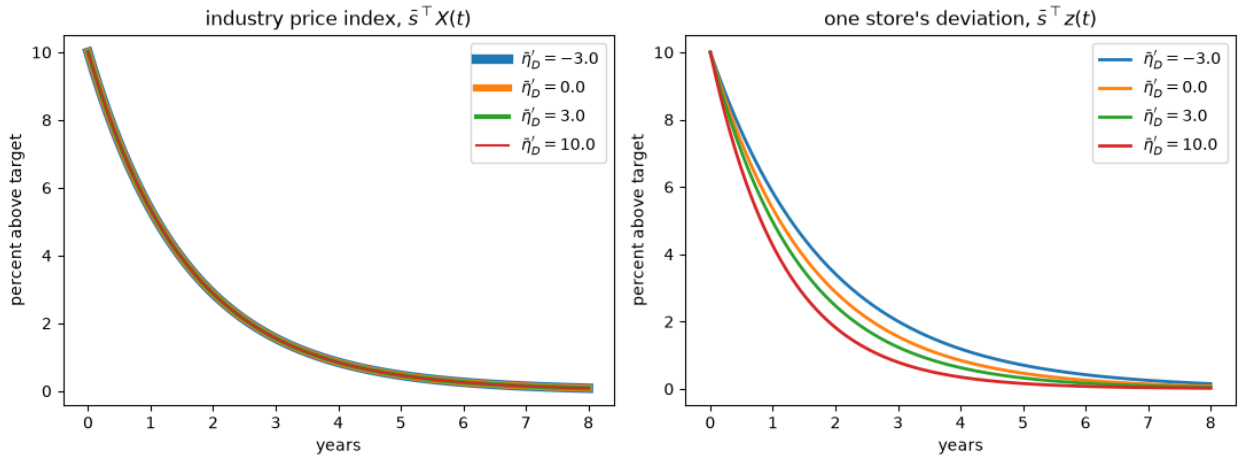


Fig. 129.6: Industry and store price paths

The four curves in the left panel lie exactly on top of one another.

The four curves in the right panel do not: the higher the superelasticity, the faster a store closes a gap between its own

prices and the industry's.

Individual and aggregate price flexibility are governed by different objects, and only the individual one responds to the superelasticity.

129.12 The planner

Is the equilibrium speed of adjustment efficient?

To answer, we compare it with the allocation of a utilitarian planner, who chooses an action $\alpha(x,t)$ for every agent and date to maximize the discounted value of the cross-sectional average of the period return, subject to the Kolmogorov forward equation.

A second-order approximation to a general return also contains terms in the aggregates alone,

$$-\tfrac{1}{2}X^\top \Lambda_X X - \tfrac{1}{2}\mathcal{A}^\top \Lambda_{\mathcal{A}} \mathcal{A},$$

with Λ_X and $\Lambda_{\mathcal{A}}$ symmetric.

An individual takes X and $\mathcal{A}$ as given, so these terms do not affect her choices or the decentralized equilibrium, and we left them out of (129.1).

The planner cannot ignore them, and its period objective is

$$-\tfrac{1}{2}\int x^\top Q x\, m(x)dx - X^\top\left(\Theta_X + \tfrac{1}{2}\Lambda_X\right)X - \tfrac{1}{2}\int \alpha(x)^\top \Gamma \alpha(x)\, m(x)dx - \mathcal{A}^\top\left(\Theta_{\mathcal{A}} + \tfrac{1}{2}\Lambda_{\mathcal{A}}\right)\mathcal{A},$$

where $X = \int x\, m(x)dx$ and $\mathcal{A} = \int \alpha(x) m(x)dx$.

The planner internalizes both sides of each interaction.

Moving one agent's state changes her own payoff and, through X, everybody else's, so the derivative of $-X^\top \Theta_X X$ with respect to the density contributes $-X^\top(\Theta_X + \Theta_X^\top)x$, twice what an individual takes into account.

The aggregate-only term adds $-X^\top \Lambda_X x$, and the same reasoning applies to actions.

Proposition 129.12.1

Whenever the planner's problem has a finite optimum, the planner's allocation coincides with the decentralized equilibrium of an economy whose interaction matrices are $\Theta_X + \Theta_X^\top + \Lambda_X$ and $\Theta_{\mathcal{A}} + \Theta_{\mathcal{A}}^\top + \Lambda_{\mathcal{A}}$.

With symmetric interactions these are $2\Theta_X + \Lambda_X$ and $2\Theta_{\mathcal{A}} + \Lambda_{\mathcal{A}}$, and the planner's Riccati equation is

$$P^{\text{pl}} B(\Gamma + 2\Theta_{\mathcal{A}} + \Lambda_{\mathcal{A}})^{-1} B^\top P^{\text{pl}} = Q + 2\Theta_X + \Lambda_X + \rho P^{\text{pl}} + P^{\text{pl}} A + A^\top P^{\text{pl}}.$$

A finite optimum exists if $Q + 2\Theta_X + \Lambda_X$ and $\Gamma + 2\Theta_{\mathcal{A}} + \Lambda_{\mathcal{A}}$ are positive definite.

Because each interaction matrix appears together with its transpose, the planner never needs them to be symmetric.

Symmetry is a restriction we need for the decentralized equilibrium, not for the planner.

The planner's curvatures have a reading that complements the one in *Effective curvature*.

When everybody moves together, so that $x = X$, the state part of the return including the aggregate-only term is

$$-\tfrac{1}{2}X^\top Q X - X^\top \Theta_X X - \tfrac{1}{2}X^\top \Lambda_X X = -\tfrac{1}{2}X^\top(Q + 2\Theta_X + \Lambda_X)X.$$

So $Q + \Theta_X$ is the slope of an agent's *marginal* return when everybody moves together, while $Q + 2\Theta_X + \Lambda_X$ is the curvature of the *total* return when everybody moves together.

The equilibrium responds to the first and the planner to the second, and the same is true of actions.

The comparative statics of *Persistence* now apply directly, under the assumptions stated there.

Internalizing complementarity in states raises the trace of the closed-loop matrix, which slows aggregate adjustment on average, while internalizing complementarity in actions does the opposite.

The aggregate-only terms can reinforce or offset these forces.

A negative semi-definite Λ_X or a positive semi-definite $\Lambda_{\mathcal{A}}$ slows adjustment further, and the opposite signs speed it up.

So the planner may adjust more slowly or more quickly than the decentralized equilibrium, and statements about individual eigenvalues need the restrictions of *Proposition 129.10.2* or *Proposition 129.10.3*.

Our two-dimensional example has no aggregate-only terms, so we set $\Lambda_X = \Lambda_{\mathcal{A}} = 0$ when we compute its planner's allocation.

```
P_planner = mfg_riccati(A, B, Q + 2*Θ_X, Γ + 2*Θ_A, ρ)
U_pl = closed_loop(A, B, Γ + 2*Θ_A, P_planner)

print("individual eigenvalues:", np.sort(np.linalg.eigvals(U_ind).real).round(4))
print("aggregate eigenvalues: ", np.sort(np.linalg.eigvals(U_agg).real).round(4))
print("planner eigenvalues:   ", np.sort(np.linalg.eigvals(U_pl).real).round(4))
```

```
individual eigenvalues: [-1.0698 -0.8321]
aggregate eigenvalues:  [-0.9025 -0.6423]
planner eigenvalues:    [-0.7786 -0.4039]
```

With complementarity in states and substitutability in actions, both of which the planner internalizes, every planner eigenvalue is closer to zero than the corresponding equilibrium eigenvalue, so the planner's allocation is the most persistent of the three.

Exercise 129.14.6 compares the planner with the equilibrium in other configurations of the interactions.

129.12.1 The planner in the capital accumulation example

The capital accumulation example does have aggregate-only terms, and they change the answer.

Alvarez and Argente [2026] do not work out the planner for this example, so we do it here.

A second-order expansion of the firm's return produces terms in K and I alone, which did not matter for the firm's problem.

They are

$$\Lambda_X = -\frac{\bar{k}^2 \Pi_{KK}}{\bar{\Pi}} = \frac{\nu}{\eta}\left(1 - \frac{\nu}{\eta}\right), \qquad \Lambda_{\mathcal{A}} = \frac{\bar{\imath}^3\, \mathcal{P}''(\bar{\imath})}{\bar{\Pi}} = \theta_{\mathcal{A}}\, \frac{\bar{\imath}\, \mathcal{P}''(\bar{\imath})}{\mathcal{P}'(\bar{\imath})},$$

where the second comes from the term $\frac{1}{2}\bar{\imath}\,\mathcal{P}''(\bar{\imath})(I - \bar{\imath})^2$ in the expansion of investment spending $\imath\,\mathcal{P}(I)$.

Adding the state interaction twice gives a clean result,

$$q + 2\theta_X + \Lambda_X = \nu(1 - \nu),$$

which is the slope of the dashed line in Fig. 129.2, the marginal revenue of the whole industry when every firm holds the same capital.

The equilibrium's effective curvature $\frac{\eta-1}{\eta}\nu(1-\nu)$ is smaller by the factor $(\eta - 1)/\eta$.

A firm's marginal revenue is only $(\eta - 1)/\eta$ of the marginal value of its output, so an industry-wide capital gap changes the private return to investing by less than it changes the value of the industry's output.

For actions, the constant-elasticity supply curve of the calibration gives $\bar{i}\,\mathcal{P}''(\bar{i})/\mathcal{P}'(\bar{i}) = 1/\varepsilon_s - 1$, so the planner's action curvature is $\gamma + (1 + 1/\varepsilon_s)\theta_{\mathcal{A}}$.

That exceeds the equilibrium's $\gamma + \theta_{\mathcal{A}}$, because the planner recognizes that a burst of industry investment raises the price that every firm pays.

So the two sides of the planner's problem pull in opposite directions: on states the planner is faster than the equilibrium, on actions slower.

Let's compute the net effect, and also what we would have found had we doubled the interactions and ignored the aggregate-only terms.

```
def cap_Λ_X(η, ν):
    "Aggregate-only state curvature in the capital accumulation example."
    return ν/η*(1 - ν/η)

Λ_X_c = cap_Λ_X(η_c, ν_c)
Λ_A_c = (1/ε_s_c - 1)*θ_A_c

print(f"Λ_X = {Λ_X_c:.4f},  Λ_A = {Λ_A_c:.4f}")
print(f"q + 2θ_X + Λ_X = {q_c + 2*θ_X_c + Λ_X_c:.4f},  ν(1 - ν) = {ν_c*(1 - ν_c):.4f}
↪")

planner_cases = {'no interactions': (0.0,              0.0),
                 'equilibrium':     (θ_X_c,            θ_A_c),
                 'doubling only':   (2*θ_X_c,          2*θ_A_c),
                 'planner':         (2*θ_X_c + Λ_X_c,  2*θ_A_c + Λ_A_c)}

print(f"\n{'':>16}{'λ':>10}{'half-life':>12}{'investment on impact':>23}")
for label, (tx, ta) in planner_cases.items():
    λ_case = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, tx, ta)
    impact = 100*X0_c*(λ_case + δ_c)/δ_c
    print(f"{label:>16}{λ_case:>10.4f}{half_life(λ_case):>12.2f}{impact:>22.2f}%")
```

```
Λ_X = 0.1444,  Λ_A = -0.0669
q + 2θ_X + Λ_X = 0.2100,  ν(1 - ν) = 0.2100

                       λ   half-life   investment on impact
 no interactions   -0.2309        3.00                -13.09%
     equilibrium   -0.1389        4.99                 -3.89%
   doubling only   -0.1109        6.25                 -1.09%
         planner   -0.1438        4.82                 -4.38%
```

The last column is the change in industry investment, relative to its steady-state level, when the capital stock starts ten percent above its steady state.

At this calibration the state side wins narrowly.

The planner's half-life is 4.8 years against the equilibrium's 5.0, and the planner cuts investment by more on impact.

Doubling the interactions without the aggregate-only terms would have given the opposite answer, a half-life above six years.

Doubling θ_X strengthens the complementarity in states, while Λ_X more than undoes it.

The next figure shows the corresponding paths of investment.

```
fig, ax = plt.subplots(figsize=(8, 4.5))
for label in ('equilibrium', 'doubling only', 'planner'):
    tx, ta = planner_cases[label]
```

(continues on next page)

(continued from previous page)

```
    λ_case = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, tx, ta)
    ls = '--' if label == 'doubling only' else '-'
    ax.plot(t_c, 100*X0_c*(λ_case + δ_c)/δ_c*np.exp(λ_case*t_c), lw=2, ls=ls,␣
↪label=label)
ax.axhline(0, color='k', lw=0.8)
ax.set_xlabel('years')
ax.set_ylabel('investment, % above steady state')
ax.legend()
plt.tight_layout()
plt.show()
```

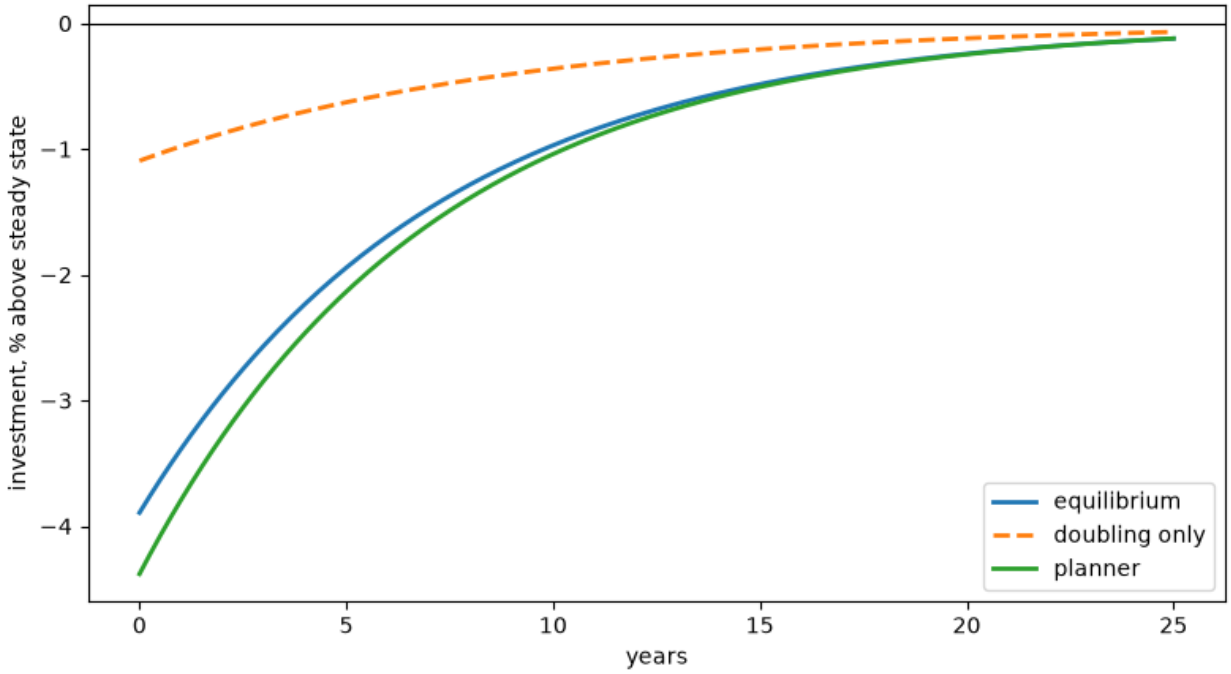


Fig. 129.7: Equilibrium and planner investment paths

Two caveats are worth keeping in mind.

First, the planner maximizes the sum of the firms' payoffs, so it counts the higher price of investment goods as a cost to firms; a planner who also counted the surplus of the suppliers of investment goods would treat part of it as a transfer.

Second, the comparison is local.

Within the quadratic model, the planner's different steady state comes from linear terms, which do not affect the dynamics of deviations.

But as an approximation to the nonlinear economy, the planner's own steady state differs from the equilibrium's (at this calibration it has more capital), and an expansion around it would give somewhat different curvatures.

Exercise 129.14.7 splits the gap between the planner and the equilibrium into its state and action parts, and finds the supply elasticity at which the ranking flips.

129.12.2 The planner in the Kimball example

Alvarez and Argente [2026] point out that in the Kimball example doubling the interaction alone would give a planner the state curvature $Q + 2\Theta_X = (Q + \Theta_X) - \bar{\eta}'_D \bar{s}\bar{s}^\top$, which depends on the superelasticity, and that the aggregate-only term Λ_X restores the cancellation.

The reading above explains why: $Q + 2\Theta_X + \Lambda_X$ is the curvature of a store's profit when its prices equal the industry's, where its relative price is one, so the superelasticity drops out, although the elasticity $\bar{\eta}_D$ still matters through the markup.

With household spending held fixed, Exercise 129.14.8 shows that the planner's state curvature is

$$Q + 2\Theta_X + \Lambda_X = (\bar{\eta}_D - 1)\left[\bar{s}\bar{s}^\top + (2\bar{\eta}_d - 1)V\right], \qquad V \equiv \mathrm{diag}(\bar{s}) - \bar{s}\bar{s}^\top,$$

against $Q + \Theta_X = (\bar{\eta}_D - 1)\left[\bar{\eta}_D \bar{s}\bar{s}^\top + \bar{\eta}_d V\right]$ for the equilibrium.

The planner's curvature is lower in the direction of a proportional increase in all of a store's prices, which changes the store's price level, and higher in directions that change only relative prices.

Let's see what that does to the modes of aggregate dynamics, measuring how much each mode moves the price level by the cosine of the angle between its eigenvector and a vector of ones.

```
ss_bar = np.outer(s_bar, s_bar)
V = np.diag(s_bar) - ss_bar
M_act = (η_D - 1)*(ss_bar + (2*η_d - 1)*V)
Q_K, Θ_K = kimball(η_d, η_D, 3.0, s_bar)

for label, M_use in (('equilibrium', Q_K + Θ_K), ('planner', M_act)):
    U_use = closed_loop(A_K, B_K, Γ_K, mfg_riccati(A_K, B_K, M_use, Γ_K, ρ_K))
    ev, W = np.linalg.eig(U_use)
    for j in np.argsort(ev.real):
        w = W[:, j].real
        cosine = abs(w.sum())/(np.sqrt(n_K)*np.linalg.norm(w))
        print(f"{label:>12}: eigenvalue {ev[j].real:.4f}, "
              f"alignment with the price level {cosine:.2f}")
```

```
equilibrium: eigenvalue -0.6527, alignment with the price level 0.94
equilibrium: eigenvalue -0.4882, alignment with the price level 0.58
equilibrium: eigenvalue -0.3836, alignment with the price level 0.37
    planner: eigenvalue -0.7001, alignment with the price level 0.07
    planner: eigenvalue -0.5372, alignment with the price level 0.10
    planner: eigenvalue -0.2959, alignment with the price level 1.00
```

The price-level mode is the fastest mode in equilibrium and the slowest for the planner: within the quadratic approximation, a cartel of stores would adjust its price level more slowly, and its relative prices more quickly, than the industry does.

Two qualifications apply.

This planner maximizes the stores' joint profits and ignores the households who buy from them, so it is a collusive benchmark rather than an efficient one.

And with household spending fixed, the stores' joint profit keeps rising as all prices rise together, so the cartel has no interior optimum and the comparison is purely local.

Exercise 129.14.8 also shows that with the profit normalization used to compute Q and Θ_X, which drops a factor that does not depend on the store's own prices, the planner's Riccati equation has no symmetric solution, so that planner's problem has no finite optimum.

129.13 Summary

The four lessons of the lecture, with the numbers from our examples:

- **Effective curvature.** Aggregate dynamics solve a single-agent Riccati equation in which Q and Γ become $Q+\Theta_X$ and $\Gamma+\Theta_{\mathcal{A}}$; in the pricing example the superelasticity cancels from $Q+\Theta_X$, so aggregate price dynamics do not depend on it.
- **Individual versus aggregate.** In the capital example a firm's deviation from the industry mean has a half-life of three years, while the industry's capital stock has a half-life of five.
- **The missing intercept.** In simulated panels from two industries that differ only in complementarity, the demeaned paths of firms are identical, and only the industry means differ.
- **Marginal versus total.** The planner's half-life in the capital example is 4.8 years, against 5.0 in equilibrium, while doubling the interactions and ignoring the aggregate-only terms would have given 6.25.

129.14 Exercises

Exercise 129.14.1

This exercise derives the formulas of the scalar case.

1. Write the equilibrium Riccati equation in the scalar case, solve the quadratic, and show that the root satisfying the admissibility condition gives (129.11).
2. Derive the existence condition (129.12) and the stability condition (129.13).
3. With the parameters of the regions figure and $\theta_{\mathcal{A}}=0$, plot λ against θ_X and mark the two thresholds.

Solution

For part 1, with $g\equiv\gamma+\theta_{\mathcal{A}}$ the Riccati equation is $\frac{b^2}{g}p^2=q+\theta_X+(\rho+2a)p$, and the closed-loop eigenvalue is $\lambda=\frac{b^2}{g}p-a$.

Substituting $p=g(\lambda+a)/b^2$ turns the Riccati equation into

$$(\lambda+a)^2-(\rho+2a)(\lambda+a)=\frac{b^2(q+\theta_X)}{g},$$

a quadratic in λ whose roots are $\frac{\rho}{2}\pm\sqrt{(\frac{\rho}{2}+a)^2+b^2(q+\theta_X)/g}$.

The two roots are symmetric about $\rho/2$, as *Proposition 129.4.1* requires, and only the smaller one satisfies the admissibility condition $\lambda<\rho/2$.

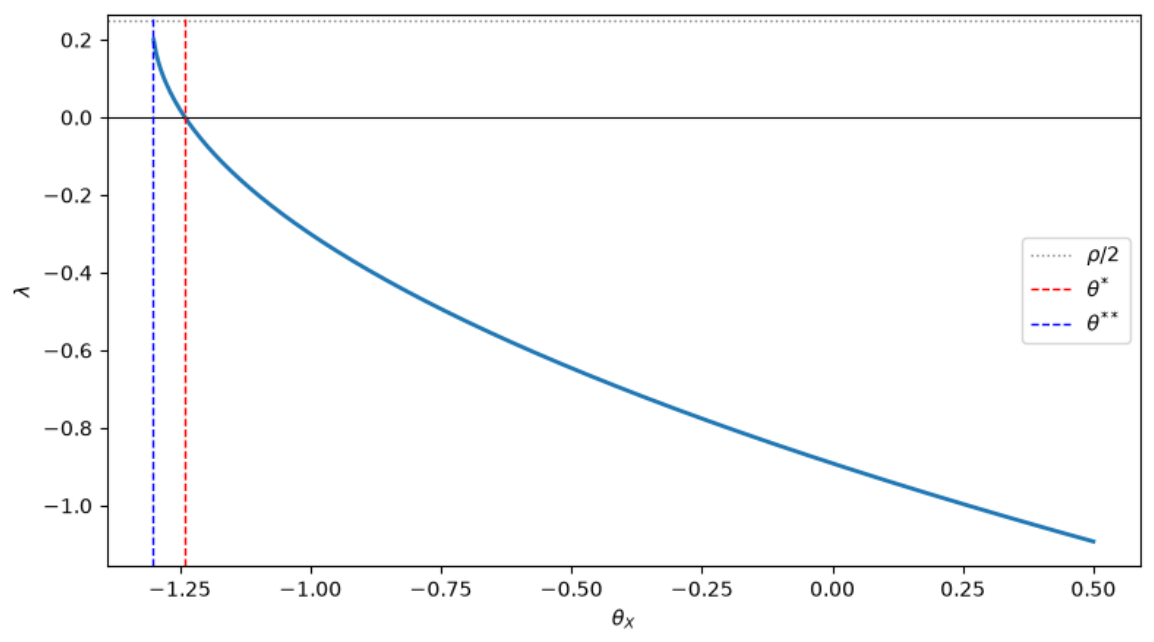
For part 2, a real root exists when the term under the square root is nonnegative, and the admissible root lies strictly below $\rho/2$ when it is positive, which gives (129.12).

The admissible root is negative when $\sqrt{(\rho/2+a)^2+b^2(q+\theta_X)/g}>\rho/2$, which after squaring gives (129.13).

For part 3, we evaluate the closed form along a line of θ_X values, using the thresholds `θ_star` and `θ_2star` computed in the lecture.

```
θ_line = np.linspace(-1.30, 0.5, 400)
inside_line = (ρ_s/2 + a)**2 + b**2*(q + θ_line)/γ
λ_line = np.where(inside_line >= 0,
                  ρ_s/2 - np.sqrt(np.maximum(inside_line, 0)), np.nan)

fig, ax = plt.subplots(figsize=(8, 4.5))
ax.plot(θ_line, λ_line, lw=2)
ax.axhline(0, color='k', lw=0.8)
ax.axhline(ρ_s/2, color='grey', ls=':', lw=1, label=r'$\rho/2$')
ax.axvline(θ_star, color='r', ls='--', lw=1, label=r'$\theta^{*}$')
ax.axvline(θ_2star, color='b', ls='--', lw=1, label=r'$\theta^{**}$')
ax.set_xlabel(r'$\theta_X$')
ax.set_ylabel(r'$\lambda$')
ax.legend()
plt.tight_layout()
plt.show()
```



Moving left along the horizontal axis means stronger complementarity in states.

Persistence rises smoothly until θ^*, where the aggregate state stops converging; between θ^* and θ^{**} equilibrium paths diverge, though slowly enough to keep discounted payoffs finite; below θ^{**} no equilibrium exists at all.

Exercise 129.14.2

This exercise derives the calibration of the capital accumulation example.

1. Use the steady-state Euler equation $\Pi_k(\bar{k}, \bar{k}) = \rho + \delta$ to show that $\bar{k}^{1-\nu} = \nu(\eta - 1)/[\eta(\rho + \delta)]$, and then that $\gamma = \phi\, \delta \bar{k}^{1-\nu}$ and $\theta_{\mathcal{A}} = \delta \bar{k}^{1-\nu}/\varepsilon_s$.
2. Show that all pairs (ϕ, ε_s) with the same value of $\phi + 1/\varepsilon_s$ generate identical *aggregate* dynamics, and verify numerically that they nevertheless imply different individual investment rules.
3. After reading *The planner in the capital accumulation example*, show that these pairs also imply different planner

allocations.

Solution

For part 1, $\Pi_k(k, K) = \nu \frac{\eta-1}{\eta} k^{\nu(1-1/\eta)-1} K^{\nu/\eta}$, so on the diagonal $\Pi_k(\bar{k}, \bar{k}) = \nu \frac{\eta-1}{\eta} \bar{k}^{\nu-1}$.

Setting this equal to $\rho + \delta$, which is the steady-state user cost when $\mathcal{P}(\bar{i}) = 1$ and $\psi'(\bar{i}) = 0$, gives the stated expression for $\bar{k}^{1-\nu}$.

Since $\bar{\Pi} = \bar{k}^\nu$, $\bar{i} = \delta \bar{k}$, $\psi''(\bar{i}) = \phi/\bar{i}$ and $\mathcal{P}'(\bar{i}) = 1/(\varepsilon_s \bar{i})$,

$$\gamma = \frac{\delta^2 \bar{k}^2}{\bar{k}^\nu} \frac{\phi}{\delta \bar{k}} = \phi \, \delta \bar{k}^{1-\nu}, \qquad \theta_{\mathcal{A}} = \frac{\delta^2 \bar{k}^2}{\bar{k}^\nu} \frac{1}{\varepsilon_s \delta \bar{k}} = \frac{\delta \bar{k}^{1-\nu}}{\varepsilon_s}.$$

For part 2, $\gamma + \theta_{\mathcal{A}} = \delta \bar{k}^{1-\nu}(\phi + 1/\varepsilon_s)$, and by (129.11) only this sum matters for λ.

But the individual's Riccati equation (129.6) involves γ alone, so it separates the pairs.

For part 3, the planner's Riccati equation involves $\gamma + 2\theta_{\mathcal{A}} + \Lambda_{\mathcal{A}} = \gamma + (1 + 1/\varepsilon_s)\theta_{\mathcal{A}} = \delta \bar{k}^{1-\nu}(\phi + 1/\varepsilon_s + 1/\varepsilon_s^2)$, so it separates them too.

The table confirms all three claims, reporting half-lives in years.

```
scale = cap_scale(η_c, ν_c)
print(f"{'ψ':>6}{'1/ε_s':>8}{'aggregate':>12}{'individual':>12}{'planner':>10}")
for ψ, inv_ε in ((0.40, 0.00), (0.30, 0.10), (1/7, 9/35), (0.05, 0.35)):
    γ_case, θ_A_case = ψ*scale, inv_ε*scale
    λ_agg = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_case, θ_X_c, θ_A_case)
    λ_ind = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_case, 0.0, 0.0)
    λ_pl = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_case,
                    2*θ_X_c + Λ_X_c, (1 + inv_ε)*θ_A_case)
    print(f"{ψ:>6.3f}{inv_ε:>8.3f}{half_life(λ_agg):>12.4f}"
          f"{half_life(λ_ind):>12.4f}{half_life(λ_pl):>10.4f}")
```

```
     ψ   1/ε_s   aggregate  individual   planner
 0.400   0.000      4.9890      4.3909    4.6210
 0.300   0.100      4.9890      3.9967    4.6535
 0.143   0.257      4.9890      3.0016    4.8198
 0.050   0.350      4.9890      1.8601    4.9636
```

Aggregate dynamics are identical down to the last digit, while the individual half-life ranges from about 4.4 years to under 2.

Aggregate data alone cannot tell a congested market for capital goods from an expensive installation technology.

Data on firms' deviations from the industry mean can, because they reveal γ alone, which is the mirror image of *Proposition 129.9.2*.

Exercise 129.14.3

This exercise derives the results in the subsection on market power and returns to scale.

1. Holding ϕ and ε_s fixed, show that $(q + \theta_X)/(\gamma + \theta_{\mathcal{A}}) = (1 - \nu)(\rho + \delta)/[\delta(\phi + 1/\varepsilon_s)]$ and $q/\gamma = (\rho + \delta)[(1 - \nu) + \nu/\eta]/(\phi\delta)$, and explain why market power affects individual but not aggregate adjustment.
2. Show that $\lambda \to -\delta$ as $\nu \to 1$, for every $\eta > 1$, and that the response of aggregate investment to a capital gap,

$(\lambda+\delta)/\delta$, goes to zero in the same limit. Explain the limit.

Solution

For part 1, from (129.15) and the calibration formulas,

$$\frac{q+\theta_X}{\gamma+\theta_{\mathcal{A}}} = \frac{\frac{\eta-1}{\eta}\nu(1-\nu)}{\frac{\delta\nu(\eta-1)}{\eta(\rho+\delta)}(\phi+1/\varepsilon_s)} = \frac{(1-\nu)(\rho+\delta)}{\delta(\phi+1/\varepsilon_s)}.$$

Both numerator and denominator are proportional to $(\eta-1)/\eta$, the ratio of a firm's marginal revenue to the value of its marginal product.

For a single firm, $q = e(1-e)$ with $e = \nu(\eta-1)/\eta$ the elasticity of a firm's revenue with respect to its own capital, so

$$\frac{q}{\gamma} = \frac{e(1-e)}{\phi e\delta/(\rho+\delta)} = \frac{(\rho+\delta)(1-e)}{\phi\delta} = \frac{\rho+\delta}{\phi\,\delta}\left[(1-\nu)+\frac{\nu}{\eta}\right].$$

Market power changes how fast a single firm's marginal revenue falls as it expands alone, but along the diagonal the industry's revenue is K^ν whatever η is, and the $(\eta-1)/\eta$ factors cancel.

For part 2, $q+\theta_X = \frac{\eta-1}{\eta}\nu(1-\nu) \to 0$, and with $a=b=\delta$,

$$\lambda \to \frac{\rho}{2} - \sqrt{\left(\frac{\rho}{2}+\delta\right)^2} = -\delta.$$

The aggregate law of motion gives $\mathcal{A} = \frac{\lambda+\delta}{\delta}X$, which vanishes as $\lambda \to -\delta$, as the table confirms.

```
print(f"{'ν':>8}{'q + θ_X':>10}{'λ':>10}{'half-life':>11}{'A/X':>9}")
for ν in (0.7, 0.9, 0.99, 0.999, 0.99999):
    γ_ν, θ_A_ν = ψ_c*cap_scale(η_c, ν), cap_scale(η_c, ν)/ε_s_c
    λ_ν = scalar_λ(ρ_c, δ_c, δ_c, cap_q(η_c, ν), γ_ν, cap_θ_X(η_c, ν), θ_A_ν)
    print(f"{ν:>8}{cap_q(η_c,ν) + cap_θ_X(η_c,ν):>10.5f}{λ_ν:>10.5f}"
          f"{half_life(λ_ν):>11.3f}{(λ_ν + δ_c)/δ_c:>9.4f}")
```

```
       ν   q + θ_X         λ  half-life      A/X
     0.7   0.15750  -0.13894      4.989  -0.3894
     0.9   0.06750  -0.11419      6.070  -0.1419
    0.99   0.00743  -0.10149      6.830  -0.0149
   0.999   0.00075  -0.10015      6.921  -0.0015
 0.99999   0.00001  -0.10000      6.931  -0.0000
```

As ν approaches 1 the profit function (129.14) approaches one that is homogeneous of degree one in (k, K) jointly, along which $\Pi_k(k,k)$ does not depend on k.

An industry-wide capital gap then creates almost no incentive to invest differently, and the gap closes almost entirely through depreciation.

The effective curvature approaches zero, the boundary of the sufficient condition $Q+\Theta_X \succ 0$ of *Proposition 129.5.3*, while the existence condition (129.12) continues to hold with room to spare.

The limit $\nu = 1$ itself is not an economy we can study this way: with constant returns the steady-state Euler equation has no solution, and $\bar{k}$ grows without bound as $\nu \to 1$.

Exercise 129.14.4

This exercise works through the identification results of *Aggregate shocks and identification*.

1. Suppose an econometrician knows ρ, Q, A, B, Γ and $\Theta_{\mathcal{A}}$, and estimates the aggregate drift matrix $U_{\dot{X},X}$. Show how to recover Θ_X, and verify your procedure numerically in the two-dimensional example.
2. In the scalar model, show that θ_X and $\theta_{\mathcal{A}}$ are not separately identified: construct a family of pairs that generate identical aggregate dynamics *and* identical policy coefficients.
3. Now suppose the econometrician knows only ρ, Γ and $\Theta_{\mathcal{A}}$, but observes X, z, α and $\mathcal{A}$, so that all four reduced-form matrices are available. Replace B by $\begin{bmatrix} 1 & 0.3 \\ 0 & 0.8 \end{bmatrix}$ in the two-dimensional example, and recover B, A, Q and Θ_X.

Solution

For part 1, invert the definition of the closed-loop matrix to get P, then read Θ_X off the equilibrium Riccati equation:

$$P = D^{-1}\left(U_{\dot{X},X} + A\right), \qquad \Theta_X = PDP - Q - \rho P - PA - A^\top P.$$

Let's apply this to the aggregate drift of the two-dimensional example.

```
U_obs = closed_loop(A, B, Γ + Θ_A, P)       # what the econometrician estimates

P_hat = np.linalg.solve(D, U_obs + A)
Θ_X_hat = P_hat @ D @ P_hat - Q - ρ*P_hat - P_hat @ A - A.T @ P_hat

print("true Θ_X =\n", Θ_X.round(6))
print("\nrecovered Θ_X =\n", Θ_X_hat.round(6))
print(f"\nmaximum error: {np.abs(Θ_X - Θ_X_hat).max():.2e}")
```

```
true Θ_X =
 [[-0.3   0.05]
 [ 0.05 -0.2 ]]

recovered Θ_X =
 [[-0.3   0.05]
 [ 0.05 -0.2 ]]

maximum error: 1.67e-16
```

For part 2, formula (129.11) shows that λ depends on the two interactions only through the ratio $(q+\theta_X)/(\gamma+\theta_{\mathcal{A}})$.

Holding that ratio fixed traces out a one-parameter family of observationally equivalent economies, one of the rays in Fig. 129.1.

The policy coefficient cannot break the tie, because $\lambda = b\,U_{\alpha,X} - a$ ties it to λ.

```
ratio = (q + (-0.4))/(γ + 0.3)        # baseline θ_X = -0.4, θ_A = 0.3

print(f"{'θ_A':>7}{'θ_X':>12}{'λ':>12}{'policy coeff':>15}")
for θ_A_alt in (0.3, 0.0, 0.6, 1.0):
    θ_X_alt = ratio*(γ + θ_A_alt) - q
    λ_alt = scalar_λ(ρ_s, a, b, q, γ, θ_X_alt, θ_A_alt)
    print(f"{θ_A_alt:>7.2f}{θ_X_alt:>12.6f}{λ_alt:>12.6f}{(λ_alt + a)/b:>15.6f}")
```

```
      θ_A         θ_X           λ   policy coeff
    0.30   -0.400000   -0.624093      -0.324093
    0.00   -0.538462   -0.624093      -0.324093
    0.60   -0.261538   -0.624093      -0.324093
    1.00   -0.076923   -0.624093      -0.324093
```

Every row describes a different economy, with a different amount of strategic complementarity in states and in actions, yet all of them generate exactly the same aggregate dynamics and the same policy rule.

Distinguishing them requires normalizing one interaction, or bringing outside information to bear.

For part 3, the two drift matrices share A, so

$$U_{\dot{z},z} - U_{\dot{X},X} = B\left(U_{\alpha,z} - U_{\alpha,X}\right)$$

identifies B.

Then $A = BU_{\alpha,X} - U_{\dot{X},X}$, $P = (B^\top)^{-1}(\Gamma + \Theta_{\mathcal{A}})U_{\alpha,X}$ and $\bar{\beta}_2 = (B^\top)^{-1}\Gamma U_{\alpha,z}$.

Finally, the single-agent Riccati equation (129.6) gives Q and the equilibrium one gives Θ_X.

```
B_true = np.array([[1.0, 0.3],
                   [0.0, 0.8]])
Γ_eq = Γ + Θ_A

# the reduced-form matrices the econometrician estimates
β2_true = mfg_riccati(A, B_true, Q, Γ, ρ)
P_true = mfg_riccati(A, B_true, Q + Θ_X, Γ_eq, ρ)
U_XX = closed_loop(A, B_true, Γ_eq, P_true)
U_zz = closed_loop(A, B_true, Γ, β2_true)
U_αX = np.linalg.solve(Γ_eq, B_true.T @ P_true)
U_αz = np.linalg.solve(Γ, B_true.T @ β2_true)

# recovery
B_hat = (U_zz - U_XX) @ np.linalg.inv(U_αz - U_αX)
A_hat = B_hat @ U_αX - U_XX
P_hat = np.linalg.solve(B_hat.T, Γ_eq @ U_αX)
β2_hat = np.linalg.solve(B_hat.T, Γ @ U_αz)
Q_hat = (β2_hat @ B_hat @ np.linalg.solve(Γ, B_hat.T) @ β2_hat
         - ρ*β2_hat - β2_hat @ A_hat - A_hat.T @ β2_hat)
D_hat = B_hat @ np.linalg.solve(Γ_eq, B_hat.T)
Θ_X_hat = P_hat @ D_hat @ P_hat - Q_hat - ρ*P_hat - P_hat @ A_hat - A_hat.T @ P_hat

for name, est, true in (('B', B_hat, B_true), ('A', A_hat, A),
                        ('Q', Q_hat, Q), ('Θ_X', Θ_X_hat, Θ_X)):
    print(f"maximum error in {name:>3}: {np.abs(est - true).max():.1e}")
```

```
maximum error in   B: 3.3e-16

maximum error in   A: 1.7e-16
maximum error in   Q: 6.7e-16
maximum error in Θ_X: 3.9e-16
```

Every object is recovered to machine precision, but only because $\Theta_{\mathcal{A}}$ was known: part 2 shows that it cannot be recovered along with the rest.

Exercise 129.14.5

Proposition 129.10.1 says that stronger complementarity in states raises the *trace* of the closed-loop matrix, and the example of Alvarez and Argente [2026] shows that it need not raise every eigenvalue.

How common is that?

1. Draw many random two-dimensional economies: take $B = I$ and $\Theta_{\mathcal{A}} = 0$, let Q and Γ be random positive definite matrices, let Θ_X be random negative definite, and set $\rho = 0.05$. For each draw, compare the equilibrium with interaction Θ_X to the one with $1.5\,\Theta_X$, keeping only draws in which $Q + 1.5\,\Theta_X$ is positive definite. Report how often the trace rises and how often *every* eigenvalue rises.
2. Repeat the experiment with $A = \Gamma^{-1}S$ for a random symmetric S, so that $D^{-1}A = S$ is symmetric, and confirm *Proposition 129.10.2*.

Solution

We write one function for both experiments, taking the rule that draws A as an argument.

```
def compare_draws(draw_A, n_draws=400, seed=3, ρ=0.05):
    """
    Compare equilibria with interactions Θ_X and 1.5 Θ_X in random
    two-dimensional economies with B = I and Θ_A = 0.  draw_A(rng, Γ)
    returns the drift matrix.  Returns the number of draws kept, the number
    in which the trace rose and the number in which every eigenvalue rose.
    """
    rng = np.random.default_rng(seed)
    I2 = np.eye(2)
    kept = trace_rose = eig_rose = 0
    for _ in range(n_draws):
        M = rng.normal(size=(2, 2)); Q_r = M @ M.T + 0.5*I2
        M = rng.normal(size=(2, 2)); Γ_r = M @ M.T + 0.5*I2
        M = rng.normal(size=(2, 2)); Θ_r = -0.2*(M @ M.T)
        A_r = draw_A(rng, Γ_r)
        if np.linalg.eigvalsh(Q_r + 1.5*Θ_r).min() <= 0:
            continue
        J0 = closed_loop(A_r, I2, Γ_r, mfg_riccati(A_r, I2, Q_r + Θ_r, Γ_r, ρ))
        J1 = closed_loop(A_r, I2, Γ_r, mfg_riccati(A_r, I2, Q_r + 1.5*Θ_r, Γ_r, ρ))
        e0 = np.sort(np.linalg.eigvals(J0).real)
        e1 = np.sort(np.linalg.eigvals(J1).real)
        kept += 1
        trace_rose += np.trace(J1) > np.trace(J0) - 1e-10
        eig_rose += np.all(e1 >= e0 - 1e-10)
    return kept, trace_rose, eig_rose

def general_drift(rng, Γ):
    "A random drift matrix with no particular structure."
    return rng.normal(size=(2, 2))*0.3 + 0.4*np.eye(2)

def symmetric_drift(rng, Γ):
    "A random drift matrix with Γ A symmetric, so that D^{-1} A is symmetric."
    M = rng.normal(size=(2, 2))
    return np.linalg.solve(Γ, 0.3*(M + M.T) + 0.4*np.eye(2))

for label, draw_A in (('general A', general_drift), ('D⁻¹A symmetric', symmetric_
↪drift)):
    kept, trace_rose, eig_rose = compare_draws(draw_A)
```

```
    print(f"{label:>16}: draws kept {kept}, trace rose {trace_rose}, "
          f"every eigenvalue rose {eig_rose}")
```

```
          general A: draws kept 266, trace rose 266, every eigenvalue rose 228
    D^-1 A symmetric: draws kept 266, trace rose 266, every eigenvalue rose 266
```

The trace result holds in every draw, as the theory says it must.

With a general drift matrix, eigenvalue-by-eigenvalue monotonicity fails in a minority of draws.

The obstacle is that $D^{-1}A$ need not be symmetric (here A itself is not symmetric), so the closed-loop matrix $DP - A$ need not be similar to a symmetric matrix, and a change that slows adjustment overall can still speed it up along some direction.

When $D^{-1}A$ is symmetric, every eigenvalue rises in every draw, as *Proposition 129.10.2* says.

Note that $Q + \Theta_X$ and D need not commute in either experiment: by *Proposition 129.10.3*, that kind of non-commutativity does not by itself break monotonicity.

Exercise 129.14.6

Compare the planner with the decentralized equilibrium in the two-dimensional example from the lecture, with no aggregate-only terms.

Compute the eigenvalues of the closed-loop matrix for the equilibrium and for the planner in four cases: only complementarity in states (Θ_X as in the lecture, $\Theta_{\mathcal{A}} = 0$); only substitutability in actions ($\Theta_X = 0$, $\Theta_{\mathcal{A}} = 0.2I$); only complementarity in actions ($\Theta_X = 0$, $\Theta_{\mathcal{A}} = -0.2I$); and complementarity in both.

In which cases is the planner's allocation more persistent than the equilibrium? Explain.

Solution

We compute both sets of eigenvalues in each case.

```
def eigen_pair(Θ_X_use, Θ_A_use):
    "Closed-loop eigenvalues for the equilibrium and for the planner."
    P_eq = mfg_riccati(A, B, Q + Θ_X_use, Γ + Θ_A_use, ρ)
    P_pl = mfg_riccati(A, B, Q + 2*Θ_X_use, Γ + 2*Θ_A_use, ρ)
    e_eq = np.sort(np.linalg.eigvals(closed_loop(A, B, Γ + Θ_A_use, P_eq)).real)
    e_pl = np.sort(np.linalg.eigvals(closed_loop(A, B, Γ + 2*Θ_A_use, P_pl)).real)
    return e_eq, e_pl

zero = np.zeros((2, 2))
planner_tests = {'states, complements':  (Θ_X, zero),
                 'actions, substitutes': (zero,  0.2*np.eye(2)),
                 'actions, complements': (zero, -0.2*np.eye(2)),
                 'both complements':     (Θ_X, -0.2*np.eye(2))}

print(f"{'case':24}{'equilibrium':>22}{'planner':>22}   planner slower?")
for label, (tx, ta) in planner_tests.items():
    e_eq, e_pl = eigen_pair(tx, ta)
    slower = bool(np.all(e_pl >= e_eq - 1e-12))
    print(f"{label:24}{str(e_eq.round(4)):>22}{str(e_pl.round(4)):>22}   {slower}")
```

```
case                                equilibrium                planner   planner␣
 ↪slower?
states, complements          [-0.9584 -0.673 ]      [-0.8522 -0.4287]   True
actions, substitutes         [-0.9933 -0.7741]      [-0.9344 -0.7284]   True
actions, complements         [-1.1748 -0.9088]      [-1.3325 -1.0151]   False
both complements             [-1.0426 -0.7346]      [-1.0137 -0.4958]   True
```

With only complementarity in states, the planner's allocation is more persistent.

The planner recognizes that when one agent stays away from the steady state, others are content to stay away too, so there is less reason to hurry back.

With only substitutability in actions the planner is also slower, because it recognizes that a burst of adjustment by everyone at once is costly.

With only complementarity in actions the comparison reverses and the planner adjusts faster: it internalizes that when one agent adjusts, others want to adjust as well.

When complementarity is present in both states and actions, the two forces work against each other.

In this calibration the state interaction dominates and the planner is still slower, but that ranking is a quantitative accident, not a theorem.

Exercise 129.14.7

In *The planner in the capital accumulation example* the planner adjusts slightly faster than the equilibrium, because the state side and the action side of its problem pull in opposite directions.

1. Compute the planner's half-life when only the state side is corrected, with state curvature $q + 2\theta_X + \Lambda_X$ and action curvature $\gamma + \theta_{\mathcal{A}}$, and when only the action side is corrected, with $q + \theta_X$ and $\gamma + (1 + 1/\varepsilon_s)\theta_{\mathcal{A}}$.
2. Hold ϕ fixed and let the supply elasticity ε_s vary. Find the elasticity at which the planner and the equilibrium have the same half-life.
3. Repeat part 2 for $\eta \in \{2, 4, 8\}$ and explain how the threshold depends on market power.

Solution

For part 1, we switch on one side of the planner's problem at a time.

```
sides = {'equilibrium':       (θ_X_c,           θ_A_c),
         'state side only':   (2*θ_X_c + Λ_X_c, θ_A_c),
         'action side only':  (θ_X_c,           2*θ_A_c + Λ_A_c),
         'planner':           (2*θ_X_c + Λ_X_c, 2*θ_A_c + Λ_A_c)}

for label, (tx, ta) in sides.items():
    λ_case = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, tx, ta)
    print(f"{label:>18}: half-life {half_life(λ_case):.2f} years")
```

```
      equilibrium: half-life 4.99 years
  state side only: half-life 4.62 years
 action side only: half-life 5.17 years
          planner: half-life 4.82 years
```

On its own, the state side would cut the half-life to 4.6 years and the action side would raise it to 5.2 years.

For parts 2 and 3, we write the gap between the two half-lives as a function of $1/\varepsilon_s$ and find its root.

```
def planner_gap(inv_ε, η):
    "Planner's half-life minus the equilibrium's, as a function of 1/ε_s."
    scale_η = cap_scale(η, ν_c)
    γ_η, θ_A_η = ψ_c*scale_η, inv_ε*scale_η
    q_η, θ_X_η = cap_q(η, ν_c), cap_θ_X(η, ν_c)
    λ_eq = scalar_λ(ρ_c, δ_c, δ_c, q_η, γ_η, θ_X_η, θ_A_η)
    λ_pl = scalar_λ(ρ_c, δ_c, δ_c, q_η, γ_η,
                    2*θ_X_η + cap_Λ_X(η, ν_c), (1 + inv_ε)*θ_A_η)
    return half_life(λ_pl) - half_life(λ_eq)

for η in (2.0, 4.0, 8.0):
    inv_star = brentq(planner_gap, 0.01, 20.0, args=(η,))
    print(f"η = {η}: half-lives coincide at ε_s = {1/inv_star:.2f}")
```

```
η = 2.0: half-lives coincide at ε_s = 0.89
η = 4.0: half-lives coincide at ε_s = 2.27
η = 8.0: half-lives coincide at ε_s = 4.33
```

With $\eta = 4$, the planner is faster whenever the supply of investment goods is more elastic than about 2.3, which includes the calibrated 3.9, and slower when it is less elastic.

The threshold rises with η.

The planner's advantage on the state side is $\nu(1-\nu) - \frac{\eta-1}{\eta}\nu(1-\nu) = \nu(1-\nu)/\eta$, which shrinks as market power falls.

So in a more competitive industry, a smaller penalty on the action side, which means a more elastic supply of investment goods, is enough to make the planner slower.

Exercise 129.14.8

This exercise derives the planner's state curvature in the Kimball example, and checks it against the profit function itself.

Normalize flexible prices to one, so that the CES weights equal the expenditure shares $\bar{s}$ and marginal costs equal $(\bar{\eta}_D - 1)/\bar{\eta}_D$, and use a Kimball demand whose elasticity at relative price y is $\bar{\eta}_D + \bar{\eta}'_D(y-1)$.

Let $\Phi(p)$ be the CES price index of a store that charges prices p, let $\tilde{p}$ be the other stores' prices, and let $\mathcal{Y}$ be household spending, which we hold fixed.

A store's actual profits are proportional to $\mathcal{Y}/\Phi(\tilde{p})$; Alvarez and Argente [2026] drop this factor, which does not depend on the store's own prices, when they compute Q and Θ_X.

Define

$$V \equiv \operatorname{diag}(\bar{s}) - \bar{s}\bar{s}^\top, \qquad \ell(X) \equiv \frac{1}{1-\bar{\eta}_d} \log \sum_j \bar{s}_j e^{(1-\bar{\eta}_d)X_j},$$

so that ℓ is the log of the store's price index.

1. Show that along the diagonal, where a store's prices equal the industry's, actual profit is proportional to $1 - \frac{\bar{\eta}_D - 1}{\bar{\eta}_D} e^{(\bar{\eta}_d - 1)\ell(X)} \sum_j \bar{s}_j e^{-\bar{\eta}_d X_j}$, and that profit with the factor dropped is proportional to $e^{\ell(X)}$ times that expression.
2. Using $\nabla\ell(0) = \bar{s}$ and $\nabla^2\ell(0) = (1-\bar{\eta}_d)V$, show that $Q + 2\Theta_X + \Lambda_X$ equals $(\bar{\eta}_D - 1)[\bar{s}\bar{s}^\top + (2\bar{\eta}_d - 1)V]$ with actual profits and $[\bar{\eta}_d(2\bar{\eta}_D - 1) - \bar{\eta}_D]V - \bar{\eta}_D\bar{s}\bar{s}^\top$ with the factor dropped.
3. Check both formulas, and the formulas for Q and Θ_X, against finite-difference Hessians of the profit function, for several values of the superelasticity.

4. Using $V\mathbb{1} = 0$, where $\mathbb{1}$ is a vector of ones, show that the first matrix is positive definite whenever $\bar{\eta}_d > 1/2$ and the second is not positive definite. Use the matrix E of (129.9) to show that with the factor dropped the planner's Riccati equation has no symmetric solution.
5. Compare the planner's curvature with the equilibrium's, $Q + \Theta_X = (\bar{\eta}_D - 1)[\bar{\eta}_D \bar{s}\bar{s}^\top + \bar{\eta}_d V]$, in the direction $\mathbb{1}$ and in directions v with $\bar{s}^\top v = 0$.

Solution

For part 1, along the diagonal the relative price $\Phi(p)/\Phi(\tilde{p})$ equals one, so the Kimball demand function is evaluated at a constant.

Revenue from the store's products is $\sum_j p_j \bar{s}_j (p_j/\Phi(p))^{-\bar{\eta}_d} = \Phi(p)^{\bar{\eta}_d} \sum_j \bar{s}_j p_j^{1-\bar{\eta}_d} = \Phi(p) = e^{\ell(X)}$, and costs are $\frac{\bar{\eta}_D - 1}{\bar{\eta}_D} e^{\bar{\eta}_d \ell(X)} \sum_j \bar{s}_j e^{-\bar{\eta}_d X_j}$.

Dividing by $\Phi(\tilde{p}) = e^{\ell(X)}$ gives actual profit.

For part 2, let $g(X) = \sum_j \bar{s}_j e^{-\bar{\eta}_d X_j}$, so that $\log g$ has gradient $-\bar{\eta}_d \bar{s}$ and Hessian $\bar{\eta}_d^2 V$ at zero.

For any function f with $f(0) = 0$, the Hessian of e^f at zero is $\nabla f \nabla f^\top + \nabla^2 f$.

For $e^{(\bar{\eta}_d - 1)\ell} g$, the logarithm has gradient $(\bar{\eta}_d - 1)\bar{s} - \bar{\eta}_d \bar{s} = -\bar{s}$ and Hessian $(\bar{\eta}_d - 1)(1 - \bar{\eta}_d)V + \bar{\eta}_d^2 V = (2\bar{\eta}_d - 1)V$, so its Hessian is $\bar{s}\bar{s}^\top + (2\bar{\eta}_d - 1)V$.

For e^ℓ, the Hessian is $\bar{s}\bar{s}^\top + (1 - \bar{\eta}_d)V$, and for $e^{\bar{\eta}_d \ell} g$ the logarithm has gradient zero and Hessian $\bar{\eta}_d(1 - \bar{\eta}_d)V + \bar{\eta}_d^2 V = \bar{\eta}_d V$.

Both profit functions equal $1/\bar{\eta}_D$ at zero, and dividing the negative Hessians by that value gives the two matrices.

For part 3, we write the profit function and a finite-difference Hessian.

```
def hessian(f, v, h=1e-4):
    "Central finite-difference Hessian of the scalar function f at v."
    m = len(v)
    E = h*np.eye(m)
    hess = np.empty((m, m))
    for i in range(m):
        for j in range(m):
            hess[i, j] = (f(v + E[i] + E[j]) - f(v + E[i] - E[j])
                          - f(v - E[i] + E[j]) + f(v - E[i] - E[j]))/(4*h**2)
    return hess

def store_profit(x, X, η_d, η_D, η_Dp, s, actual=True):
    """
    Profit of a store charging log prices x when all other stores charge X,
    with flexible prices normalized to one.  With actual=False the factor
    Φ(p̃), which does not depend on x, is dropped.
    """
    p, p_others = np.exp(x), np.exp(X)
    Φ = lambda p: (s @ p**(1 - η_d))**(1/(1 - η_d))
    y = Φ(p)/Φ(p_others)
    demand = y**(η_Dp - η_D)*np.exp(-η_Dp*(y - 1))
    mc = (η_D - 1)/η_D
    profit = np.sum((p - mc)*s*(p/Φ(p))**(-η_d))*demand
    return profit/Φ(p_others) if actual else profit

def profit_curvatures(η_d, η_D, η_Dp, s, actual=True):
    "Q, Θ_X and Λ_X from the Hessian of store profits at the steady state."
    n = len(s)
    f = lambda v: store_profit(v[:n], v[n:], η_d, η_D, η_Dp, s, actual)
    v0 = np.zeros(2*n)
    hess = -hessian(f, v0)/f(v0)
    return hess[:n, :n], hess[:n, n:], hess[n:, n:]
```

Now we compare the finite differences with the closed forms.

```
M_norm = (η_d*(2*η_D - 1) - η_D)*V - η_D*ss_bar

for actual, M_closed in ((True, M_act), (False, M_norm)):
    label = "actual" if actual else "factor dropped"
    for η_Dp in (-3.0, 3.0, 10.0):
        Q_f, Θ_f, Λ_f = profit_curvatures(η_d, η_D, η_Dp, s_bar, actual)
        Q_K, Θ_K = kimball(η_d, η_D, η_Dp, s_bar)
        err_QΘ = max(np.abs(Q_f - Q_K).max(), np.abs(Θ_f - Θ_K).max())
        err_M = np.abs(Q_f + 2*Θ_f + Λ_f - M_closed).max()
        print(f"{label:>15}, η_D' = {η_Dp:>5}: error in Q and Θ_X {err_QΘ:.0e}, "
              f"error in Q + 2Θ_X + Λ_X {err_M:.0e}")
```

```
         actual, η_D' =  -3.0: error in Q and Θ_X 1e-06, error in Q + 2Θ_X + Λ_
↪X 7e-07
         actual, η_D' =   3.0: error in Q and Θ_X 9e-07, error in Q + 2Θ_X + Λ_
↪X 6e-07
         actual, η_D' =  10.0: error in Q and Θ_X 9e-07, error in Q + 2Θ_X + Λ_
↪X 7e-07
factor dropped, η_D' =  -3.0: error in Q and Θ_X 1e-06, error in Q + 2Θ_X + Λ_
↪X 6e-07
factor dropped, η_D' =   3.0: error in Q and Θ_X 9e-07, error in Q + 2Θ_X + Λ_
↪X 5e-07
factor dropped, η_D' =  10.0: error in Q and Θ_X 9e-07, error in Q + 2Θ_X + Λ_
↪X 6e-07
```

The formulas hold for every superelasticity, and neither planner matrix depends on it.

Dropping the factor leaves Q and Θ_X unchanged, because the store's first-order condition holds at the steady state, but it changes Λ_X.

For part 4, $v^\top V v = \sum_j \bar{s}_j v_j^2 - (\bar{s}^\top v)^2 \geq 0$, with equality only when v is proportional to $\mathbb{1}$, and then $v^\top \bar{s}\bar{s}^\top v > 0$.

So the actual matrix is positive definite whenever $2\bar{\eta}_d - 1 > 0$.

With household spending fixed, an industry-wide price increase lowers quantities, so joint profit is concave in the level of prices.

For the other matrix, $V\mathbb{1} = 0$ and $\bar{s}^\top \mathbb{1} = 1$ give $\mathbb{1}^\top M \mathbb{1} = -\bar{\eta}_D < 0$.

Dropping the factor holds quantities fixed, so raising all prices in proportion raises profit at an increasing rate.

The planner's problem has $\Gamma + 2\Theta_{\mathcal{A}} + \Lambda_{\mathcal{A}} = \Gamma$, so its version of E is $M + (\pi + \rho/2)^2\Gamma$.

```
for label, M_closed in (('actual', M_act), ('factor dropped', M_norm)):
    E_pl = M_closed + (π_K + ρ_K/2)**2*Γ_K
    print(f"{label:>15}: eigenvalues of E {np.linalg.eigvalsh(E_pl).round(3)}")
    try:
        mfg_riccati(A_K, B_K, M_closed, Γ_K, ρ_K)
        print(" "*17 + "the planner's Riccati equation has an admissible solution")
    except ValueError as e:
        print(" "*17 + str(e))
```

```
         actual: eigenvalues of E [ 1.001  7.747 12.948]
                 the planner's Riccati equation has an admissible solution
factor dropped: eigenvalues of E [-1.336  8.795 14.677]
                 no admissible solution of the Riccati equation
```

With the factor dropped, E has a negative eigenvalue, so by *Proposition 129.5.3* no symmetric solution exists, the planner's problem has no finite optimum, and our solver duly raises an error.

For part 5, the equilibrium's curvature in the direction $\mathbb{1}$ is $(\bar{\eta}_D - 1)\bar{\eta}_D$, against $\bar{\eta}_D - 1$ for the planner, smaller by the factor $\bar{\eta}_D$.

A single store that raises its prices loses customers to other stores at elasticity $\bar{\eta}_D$, while stores that raise their prices together keep their revenue and only save on costs.

In directions with $\bar{s}^\top v = 0$, which leave the store's price index unchanged, the equilibrium's curvature is $(\bar{\eta}_D - 1)\bar{\eta}_d\, v^\top \operatorname{diag}(\bar{s}) v$ and the planner's is $(\bar{\eta}_D - 1)(2\bar{\eta}_d - 1)\, v^\top \operatorname{diag}(\bar{s}) v$, which is larger because $\bar{\eta}_d > 1$.

```
ones = np.ones(n_K)
v = np.array([1.0, -1.0, 0.5])
v = v - (s_bar @ v)*ones                    # now s̄'v = 0
Q_K, Θ_K = kimball(η_d, η_D, 3.0, s_bar)

for name, M_use in (('equilibrium', Q_K + Θ_K), ('planner', M_act)):
    ratio_v = (v @ M_use @ v)/(v @ np.diag(s_bar) @ v)
    print(f"{name:>12}: 1'M1 = {ones @ M_use @ ones:7.3f},   "
          f"v'Mv / v'diag(s̄)v = {ratio_v:7.3f}")
```

```
equilibrium: 1'M1 =  12.000,   v'Mv / v'diag(s̄)v =  18.000
    planner: 1'M1 =   3.000,   v'Mv / v'diag(s̄)v =  33.000
```

This is why the planner adjusts its price level more slowly than the industry does and its relative prices more quickly.

129.15 Further reading

Alvarez and Argente [2026] develop the results in this lecture, along with the two economic examples; their appendices contain the proofs and the derivations behind both examples, and they discuss the planner in the pricing example.

The planner in the capital example, the analysis of market power with structural parameters held fixed, the finite-difference check of the pricing planner, the observation that with their profit normalization that planner has no finite optimum, and the simulated panels are our own illustrations of their results.

Linear Quadratic Mean Field Games: Extensions continues with cross terms between states and actions, forcing processes, and interactions through higher moments.

Carmona and Delarue [2018] give a comprehensive probabilistic treatment of mean field games, and Achdou *et al.* [2022] describe the numerical methods used to solve the coupled partial differential equations when the model is not linear quadratic.

CHAPTER

LINEAR QUADRATIC MEAN FIELD GAMES: EXTENSIONS

Contents

- *Linear Quadratic Mean Field Games: Extensions*
  - *Overview*
  - *What we need from the previous lecture*
  - *State-action cross terms*
  - *Forcing processes and pass-through*
  - *Coupling through higher moments*
  - *Exercises*
  - *Further reading*

130.1 Overview

This lecture continues *Linear Quadratic Mean Field Games*, and it assumes familiarity with that lecture's theory, through the split into individual and aggregate dynamics, and with its capital accumulation example.

There we studied a continuum of agents whose quadratic payoffs depend on the cross-sectional averages of states and actions, and found that equilibrium aggregate dynamics solve the Riccati equation of a single-agent problem in which strategic interaction changes only the curvatures.

Here we take up three extensions developed by Alvarez and Argente [2026]:

- **Cross terms between states and actions.** The equivalence result survives when the mixed interactions between states and actions satisfy a symmetry restriction, with the mixed interaction subtracted from the own cross term, and uniqueness can fail when they do not.
- **Forcing processes.** States that agents cannot control, such as costs, can be added without changing the Riccati equation of the controlled block, and a pricing example shows that firms pass idiosyncratic and common cost shocks through at different speeds and to different extents.
- **Higher moments.** When payoffs depend on higher moments of the cross-sectional distribution, idiosyncratic risk affects aggregate persistence.

Let's start with some imports:

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.linalg import solve_continuous_are, expm
```

130.2 What we need from the previous lecture

Throughout, lower-case letters denote an agent's own state x and action α, and upper-case or calligraphic letters denote their cross-sectional averages X and $\mathcal{A}$.

In *Linear Quadratic Mean Field Games* an agent has period return and state equation

$$-\tfrac{1}{2}x^\top Q x - X^\top \Theta_X x - \tfrac{1}{2}\alpha^\top \Gamma \alpha - \mathcal{A}^\top \Theta_{\mathcal{A}} \alpha, \qquad dx = (B\alpha - Ax)dt + \Sigma^{1/2} dW,$$

and discounts at rate ρ.

Given the aggregate paths, her value function is $u(x,t) = \beta_0(t) + \beta_1(t)^\top x + \frac{1}{2}x^\top \bar{\beta}_2 x$, where $\bar{\beta}_2$ solves the single-agent Riccati equation and involves neither interaction matrix.

In equilibrium the aggregate state follows $\dot{X} = \left(B(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top P - A\right) X$, where the symmetric matrix P solves

$$PB(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top P = Q + \Theta_X + \rho P + PA + A^\top P,$$

and the eigenvalues of the closed-loop matrix must have real parts below $\rho/2$.

That equation comes from a state-costate system whose eigenvalues come in pairs λ and $\rho - \lambda$, so at most half of them lie below $\rho/2$; with $B(\Gamma + \Theta_{\mathcal{A}})^{-1}B^\top$ positive definite, that is what makes the equilibrium unique.

An agent's deviation from the cross-sectional mean, $z = x - X$, follows $dz = (B\Gamma^{-1}B^\top \bar{\beta}_2 - A)z\,dt + \Sigma^{1/2}dW$, in which no interaction matrix appears.

So data on deviations from the mean contain no information about strategic interaction, which is the missing intercept.

Shocks common to all agents, modeled with a compensated jump process $\mathcal{J}$, enter the law of motion of X but leave P unchanged.

We reuse two tools: the closed form for the aggregate eigenvalue in the scalar case,

$$\lambda = \frac{\rho}{2} - \sqrt{\left(\frac{\rho}{2} + a\right)^2 + \frac{b^2(q + \theta_X)}{\gamma + \theta_{\mathcal{A}}}},$$

and the calibration of the capital accumulation example.

In that example a firm with capital k earns $\Pi(k,K) = k^{\nu(1-1/\eta)}K^{\nu/\eta}$ when the rest of the industry holds K, pays an adjustment cost $\psi(i) = \frac{\phi}{2}\bar{i}\left(\frac{i-\bar{i}}{\bar{i}}\right)^2$, and buys investment goods at price $\mathcal{P}(I) = (I/\bar{i})^{1/\varepsilon_s}$.

In percentage deviations from the steady state it maps into the scalar model with $a = b = \delta$, $q + \theta_X = \frac{\eta-1}{\eta}\nu(1-\nu)$, $\gamma = \phi\,\delta\bar{k}^{1-\nu}$ and $\theta_{\mathcal{A}} = \delta\bar{k}^{1-\nu}/\varepsilon_s$, and the calibration gives half-lives of three years for a firm's deviation from the industry mean and five years for the industry.

```
def scalar_λ(ρ, a, b, q, γ, θ_X, θ_A):
    "Closed-form aggregate eigenvalue in the scalar case."
    return ρ/2 - np.sqrt((ρ/2 + a)**2 + b**2*(q + θ_X)/(γ + θ_A))

def half_life(λ):
    "Time for a state with eigenvalue λ < 0 to close half of a gap."
    return np.log(2)/(-λ)

def cap_q(η, ν):
    "Own-state curvature in the capital accumulation example."
    return ν*(η - 1)/η**2*(η*(1 - ν) + ν)

def cap_θ_X(η, ν):
```

(continues on next page)

(continued from previous page)

```
    "State interaction in the capital accumulation example."
    return -(η - 1)/η*ν**2/η

ρ_c, δ_c, η_c, ν_c = 0.05, 0.10, 4.0, 0.7
ψ_c, ε_s_c = 1/7, 35/9

def cap_scale(η, ν, ρ=ρ_c, δ=δ_c):
    "δ k̄^{1-ν}, which converts ψ and 1/ε_s into γ and θ_A."
    return δ*ν*(η - 1)/(η*(ρ + δ))

γ_c = ψ_c*cap_scale(η_c, ν_c)
θ_A_c = cap_scale(η_c, ν_c)/ε_s_c
q_c, θ_X_c = cap_q(η_c, ν_c), cap_θ_X(η_c, ν_c)

λ_c = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, θ_X_c, θ_A_c)          # industry
λ_micro = scalar_λ(ρ_c, δ_c, δ_c, q_c, γ_c, 0.0, 0.0)          # single firm
print(f"γ = {γ_c:.3f}, θ_A = {θ_A_c:.3f}")
print(f"half-lives: firm {half_life(λ_micro):.2f} years, industry {half_life(λ_c):.2f}
↪ years")
```

```
γ = 0.050, θ_A = 0.090
half-lives: firm 3.00 years, industry 4.99 years
```

130.3 State-action cross terms

The return in *Linear Quadratic Mean Field Games* was not the most general quadratic return.

A second-order approximation to a general period return produces

$$\begin{aligned}\tilde{F}(x, X, \alpha, \mathcal{A}) = &- \tfrac{1}{2} x^\top Q x - X^\top \Theta_X x - \tfrac{1}{2}\alpha^\top \Gamma \alpha - \mathcal{A}^\top \Theta_{\mathcal{A}} \alpha \\ &+ x^\top C \alpha - X^\top \Theta_{X\alpha} \alpha - x^\top \Theta_{x\mathcal{A}} \mathcal{A} \\ &- \tfrac{1}{2} X^\top \Lambda_X X - \tfrac{1}{2}\mathcal{A}^\top \Lambda_{\mathcal{A}} \mathcal{A} - X^\top \Lambda_{X\mathcal{A}} \mathcal{A}.\end{aligned} \tag{130.1}$$

The previous lecture used the first line throughout, and the terms in the aggregates alone in the third line when it studied the planner.

Here we restore the second line.

The own cross term C links an agent's state and action.

It appears when adjustment costs depend on the state as well as on the action, as with investment costs that depend on the investment *rate* [Cooper and Haltiwanger, 2006, Hayashi, 1982].

The mixed interactions appear when the average state affects the marginal return to an agent's action ($\Theta_{X\alpha}$), or the average action affects the marginal return to her state ($\Theta_{x\mathcal{A}}$).

130.3.1 The individual problem

Taking the aggregate paths as given, the agent's first-order condition is now

$$\alpha^*(x,t) = \Gamma^{-1}\left[B^\top u_x(x,t) + C^\top x - \Theta_{\mathcal{A}}\mathcal{A}(t) - \Theta_{X\alpha}^\top X(t)\right],$$

and the curvature of her value function solves

$$(\bar{\beta}_2 B + C)\Gamma^{-1}(B^\top \bar{\beta}_2 + C^\top) = Q + \rho\bar{\beta}_2 + \bar{\beta}_2 A + A^\top \bar{\beta}_2. \tag{130.2}$$

The own cross term enters the individual's curvature, but, as before, no interaction matrix does.

Completing the square in the action,

$$-\tfrac{1}{2}x^\top Q x - \tfrac{1}{2}\alpha^\top\Gamma\alpha + x^\top C\alpha = -\tfrac{1}{2}x^\top(Q - C\Gamma^{-1}C^\top)x - \tfrac{1}{2}(\alpha - \Gamma^{-1}C^\top x)^\top\Gamma(\alpha - \Gamma^{-1}C^\top x),$$

shows that the counterpart of $Q \succ 0$ is $Q - C\Gamma^{-1}C^\top \succ 0$, which we assume.

130.3.2 The equilibrium

Aggregating the first-order condition gives

$$(\Gamma + \Theta_{\mathcal{A}})\mathcal{A} = B^\top(\beta_1 + \bar{\beta}_2 X) + (C - \Theta_{X\alpha})^\top X.$$

It is convenient to track $w \equiv \beta_1 + \bar{\beta}_2 X$, the marginal value of the state for an agent at the mean.

Differentiating, aggregating, and using (130.2) gives the state-costate system

$$\begin{aligned} \dot{w} &= (\rho I + A^\top)w + (Q + \Theta_X)X - (C - \Theta_{x\mathcal{A}})\mathcal{A}, \\ \dot{X} &= B\mathcal{A} - AX. \end{aligned} \tag{130.3}$$

The costate equation involves $\Theta_{x\mathcal{A}}$, through which the average action shifts the marginal value of an agent's state.

The aggregate action involves $\Theta_{X\alpha}$, through which the average state shifts the marginal return to her action.

The two need not be equal, but when they are, the system keeps the Hamiltonian structure of the benchmark.

Proposition 130.3.1

Suppose $\Theta_{X\alpha} = \Theta_{x\mathcal{A}}$, and write $\tilde{C} \equiv C - \Theta_{X\alpha}$.

An equilibrium is characterized by a symmetric matrix P solving

$$(PB + \tilde{C})(\Gamma + \Theta_{\mathcal{A}})^{-1}(B^\top P + \tilde{C}^\top) = Q + \Theta_X + \rho P + PA + A^\top P,$$

with aggregate dynamics

$$\dot{X} = \left[B(\Gamma + \Theta_{\mathcal{A}})^{-1}(B^\top P + \tilde{C}^\top) - A\right] X.$$

This is the Riccati equation of a single agent whose return is

$$-\tfrac{1}{2}X^\top(Q + \Theta_X)X - \tfrac{1}{2}\mathcal{A}^\top(\Gamma + \Theta_{\mathcal{A}})\mathcal{A} + X^\top\tilde{C}\mathcal{A}.$$

If $\Gamma + \Theta_{\mathcal{A}} \succ 0$ and $Q + \Theta_X - \tilde{C}(\Gamma + \Theta_{\mathcal{A}})^{-1}\tilde{C}^\top \succ 0$, an equilibrium exists and is unique.

So the equivalence result of *Linear Quadratic Mean Field Games* survives, with one more adjustment: the mixed interaction shifts the effective cross term from C to $C - \Theta_{X\alpha}$.

Completing the square in the effective problem shows that a cross term does two things.

It lowers the effective state curvature to $Q + \Theta_X - \tilde{C}(\Gamma + \Theta_{\mathcal{A}})^{-1}\tilde{C}^\top$, and it changes the effective drift matrix from A to $A - B(\Gamma + \Theta_{\mathcal{A}})^{-1}\tilde{C}^\top$.

The comparative statics of the previous lecture apply to these transformed objects rather than to the interaction matrices themselves.

The dynamics of an agent's deviation from the mean become

$$dz = \left[B\Gamma^{-1}(B^\top \bar{\beta}_2 + C^\top) - A\right] z\, dt + \Sigma^{1/2} dW,$$

which involve the own cross term C but still no interaction matrix, so data on deviations from the mean still contain no information about strategic interaction.

The planner is handled in the same way.

Its effective cross term is $C - \Theta_{X\alpha} - \Theta_{x\mathcal{A}} - \Lambda_{X\mathcal{A}}$, so it needs no restriction on the mixed interactions, which enter only through their sum.

To compute equilibria with cross terms, we extend the solver of *Linear Quadratic Mean Field Games*.

The `s` argument of `solve_continuous_are` replaces $\tilde{P}B$ by $\tilde{P}B + s$ in the quadratic term of the standard Riccati equation, so setting $s = -\tilde{C}$ handles an effective cross term $\tilde{C}$.

As in the previous lecture, the function checks the Riccati residual and the eigenvalue condition, and raises an error if either fails.

```
def closed_loop(A, B, Γ_eff, P, C=None):
    "The matrix governing aggregate dynamics, Ẋ = (B Γ_eff^{-1} (B' P + C') - A) X."
    K = B.T @ P if C is None else B.T @ P + C.T
    return B @ np.linalg.solve(Γ_eff, K) - A

def mfg_riccati(A, B, Q, Γ, ρ, C=None, tol=1e-8):
    """
    Solve  (P B + C) Γ^{-1} (B' P + C') = Q + ρ P + P A + A' P, with C = 0 by
    default, for the solution P whose closed-loop matrix has eigenvalues with
    real parts below ρ/2, and raise ValueError if there is none.
    """
    n, k = B.shape
    C_use = np.zeros((n, k)) if C is None else C
    try:
        P = -solve_continuous_are(-(A + ρ/2 * np.eye(n)), B, Q, Γ, s=-C_use)
    except (np.linalg.LinAlgError, ValueError):
        raise ValueError("no admissible solution of the Riccati equation")
    quad = (P @ B + C_use) @ np.linalg.solve(Γ, B.T @ P + C_use.T)
    residual = quad - (Q + ρ*P + P @ A + A.T @ P)
    scale = 1 + np.abs(Q).max() + np.abs(quad).max()
    if (np.abs(residual).max() > tol*scale
            or np.linalg.eigvals(closed_loop(A, B, Γ, P, C)).real.max() >= ρ/2):
        raise ValueError("no admissible solution of the Riccati equation")
    return P
```

130.3.3 Investment rates and Hayashi adjustment costs

In the capital accumulation example the adjustment cost depends on investment alone.

A common alternative makes it depend on the investment rate,

$$\psi(i,k) = \frac{\phi}{2}\,\delta k\left(\frac{i}{\delta k} - 1\right)^2,$$

which is homogeneous of degree one, as in Hayashi [1982], and which reduces to the cost $\psi(i)$ above when $k = \bar{k}$.

Since $i - \delta k = \delta\bar{k}(\alpha - x)$ in percentage deviations, the normalized cost is $\frac{\gamma}{2}(\alpha - x)^2$ to second order, with the same $\gamma = \phi\,\delta\bar{k}^{1-\nu}$ as before.

Expanding the square shows that the cost adds γ to the own-state curvature and creates an own cross term,

$$q \to q + \gamma, \qquad C = \gamma,$$

while $Q - C\Gamma^{-1}C^\top = q$ is unchanged.

There are no mixed interactions, so $\tilde{C} = \gamma$.

Completing the square in the effective problem of *Proposition 130.3.1* turns the scalar problem into one without a cross term, with the drift $a = \delta$ replaced by

$$\tilde{a} = \delta - \frac{\delta\gamma}{\gamma + \theta_{\mathcal{A}}} = \frac{\delta\,\theta_{\mathcal{A}}}{\gamma + \theta_{\mathcal{A}}}$$

and the effective state curvature $q + \theta_X + \gamma - \gamma^2/(\gamma + \theta_{\mathcal{A}}) = q + \theta_X + \gamma\theta_{\mathcal{A}}/(\gamma + \theta_{\mathcal{A}})$.

So the aggregate eigenvalue is

$$\lambda = \frac{\rho}{2} - \sqrt{\left(\frac{\rho}{2} + \tilde{a}\right)^2 + \frac{\delta^2}{\gamma + \theta_{\mathcal{A}}}\left(q + \theta_X + \frac{\gamma\,\theta_{\mathcal{A}}}{\gamma + \theta_{\mathcal{A}}}\right)},$$

and setting $\theta_X = \theta_{\mathcal{A}} = 0$, so that $\tilde{a} = 0$, gives the individual's.

With Hayashi costs, investing at the replacement rate is free whatever the capital stock, so depreciation no longer pushes a firm back toward its steady state.

The drift δ that sat inside the square root in the formula without cross terms disappears from the individual's problem, and it returns in the industry's problem only through the price of investment goods.

Let's compare the half-lives with and without Hayashi costs, holding γ at its calibrated value.

```
def λ_cross(ρ, a, b, q_eff, γ_eff, c_eff):
    "Scalar aggregate eigenvalue with an effective cross term, by completing the␣
 ↪square."
    a_t = a - b*c_eff/γ_eff
    q_t = q_eff - c_eff**2/γ_eff
    return ρ/2 - np.sqrt((ρ/2 + a_t)**2 + b**2*q_t/γ_eff)

λ_ind_H = λ_cross(ρ_c, δ_c, δ_c, q_c + γ_c, γ_c, γ_c)
λ_agg_H = λ_cross(ρ_c, δ_c, δ_c, q_c + γ_c + θ_X_c, γ_c + θ_A_c, γ_c)

# the general solver agrees
one = np.array([[1.0]])
P_H = mfg_riccati(δ_c*one, δ_c*one, (q_c + γ_c + θ_X_c)*one,
                  (γ_c + θ_A_c)*one, ρ_c, C=γ_c*one)
λ_agg_solver = closed_loop(δ_c*one, δ_c*one, (γ_c + θ_A_c)*one, P_H, C=γ_c*one)[0, 0]
```

(continues on next page)

(continued from previous page)

```
print(f"aggregate λ: closed form {λ_agg_H:.6f}, solver {λ_agg_solver:.6f}\n")

print(f"{'':>20}{'individual':>12}{'industry':>10}{'ratio':>8}")
for label, (λ_i, λ_a) in {'investment costs': (λ_micro, λ_c),
                          'Hayashi costs':    (λ_ind_H, λ_agg_H)}.items():
    print(f"{label:>20}{half_life(λ_i):>12.2f}{half_life(λ_a):>10.2f}"
          f"{half_life(λ_a)/half_life(λ_i):>8.2f}")
```

```
aggregate λ: closed form -0.121690, solver -0.121690

                     individual  industry   ratio
    investment costs        3.00      4.99    1.66
       Hayashi costs        3.47      5.70    1.64
```

Hayashi costs slow both the individual firm and the industry, and leave the ratio of the two half-lives about where it was.

Exercise 130.6.1 recalibrates ϕ to restore a three-year individual half-life, and revisits the limit of constant returns to scale.

130.3.4 When the restriction fails

Alvarez and Argente [2026] point out that the restriction $\Theta_{X\alpha} = \Theta_{x\mathcal{A}}$ has no general economic justification.

Consider a country with capital k and investment i that takes world capital K and world investment I as given.

Suppose its consumption is $c = f(k, K) - i$ and its period utility is $u(c)$, with $u'' < 0$ and $f_K > 0$.

Then $\partial^2 u / \partial i \, \partial K = -u''(c) f_K > 0$, because higher world capital raises consumption and so lowers the marginal utility cost of investing.

But $\partial^2 u / \partial k \, \partial I = 0$, because world investment does not affect the return to capital.

So the two mixed interactions differ, and in the notation of (130.1), $\Theta_{X\alpha} < 0 = \Theta_{x\mathcal{A}}$.

When they differ, (130.3) is still a linear state-costate system, but the two outer factors of the Riccati equation are no longer transposes of one another, and the matrix of the system is no longer Hamiltonian.

Its eigenvalues then need not come in pairs λ and $\rho - \lambda$, and the argument that delivered at most one equilibrium in *Linear Quadratic Mean Field Games* breaks down.

Even with a single state, the two eigenvalues can both lie below $\rho/2$, giving a continuum of equilibria of the kind studied by Benhabib and Farmer [1994], or both lie above it, so that no equilibrium exists.

To see the pairing fail, we write a function for the matrix of (130.3) that allows the two mixed interactions to differ.

```
def state_costate(A, B, Q_eff, Γ_eff, ρ, C1, C2):
    """
    Matrix of the state-costate system for (w, X).  C1 = C - Θ_xA enters the
    costate equation and C2 = C - Θ_Xa enters the aggregate action.
    """
    n = A.shape[0]
    G_B = np.linalg.solve(Γ_eff, B.T)       # (Γ + Θ_A)^{-1} B'
    G_C = np.linalg.solve(Γ_eff, C2.T)      # (Γ + Θ_A)^{-1} C2'
    return np.block([[ρ*np.eye(n) + A.T - C1 @ G_B, Q_eff - C1 @ G_C],
                     [B @ G_B,                      B @ G_C - A]])
```

We apply it to the Hayashi economy, first as it is and then with a small mixed interaction of the open-economy sign, $\Theta_{X\alpha} = -0.05$ and $\Theta_{x\mathcal{A}} = 0$.

```
for label, θ_Xα, θ_xA in (('restriction holds', 0.0, 0.0),
                          ('restriction fails', -0.05, 0.0)):
    M_sc = state_costate(δ_c*one, δ_c*one, (q_c + y_c + θ_X_c)*one, (y_c + θ_A_c)*one,
                         ρ_c, (y_c - θ_xA)*one, (y_c - θ_Xα)*one)
    ev = np.sort(np.linalg.eigvals(M_sc).real)
    print(f"{label}: eigenvalues {ev.round(4)}, sum {ev.sum():.4f}, ρ = {ρ_c}")
```

```
restriction holds: eigenvalues [-0.1217  0.1717], sum 0.0500, ρ = 0.05
restriction fails: eigenvalues [-0.0889  0.1747], sum 0.0857, ρ = 0.05
```

When the restriction holds, the eigenvalues sum to ρ, so they are symmetric about $\rho/2$.

When it fails, they sum to $\rho + b(\theta_{x\mathcal{A}} - \theta_{X\alpha})/(\gamma + \theta_{\mathcal{A}})$, and the open-economy sign pushes them up, toward the case in which no equilibrium exists.

Exercise 130.6.2 finds parameters at which both eigenvalues lie on the same side of $\rho/2$.

The planner is immune to all of this, because it internalizes both mixed interactions and so sees only their sum.

130.4 Forcing processes and pass-through

The same machinery handles states that agents cannot control, such as a marginal cost.

Partition the state as $x = (x^{(1)}, x^{(2)})$, where $x^{(1)}$ is controlled and $x^{(2)}$ is a forcing process with its own law of motion,

$$A = \begin{bmatrix} A_{11} & A_{12} \\ 0 & A_{22} \end{bmatrix}, \qquad B = \begin{bmatrix} B_1 \\ 0 \end{bmatrix}.$$

Now B no longer has full row rank and $B\Gamma^{-1}B^\top$ is singular, so the Riccati equations are not of the standard full-rank kind, but they are block recursive.

Following Alvarez and Argente [2026], we keep interactions through states only, $\Theta_{\mathcal{A}} = 0$.

Proposition 130.4.1

Suppose $\Theta_{\mathcal{A}} = 0$, that $B_1\Gamma^{-1}B_1^\top$, Q_{11} and $Q_{11} + \Theta_{X,11}$ are positive definite, and that the eigenvalues of $-A_{22}$ have real parts below $\rho/2$.

Then the upper-left block of P solves the equilibrium Riccati equation of the controlled block alone,

$$P_{11}B_1\Gamma^{-1}B_1^\top P_{11} = Q_{11} + \Theta_{X,11} + \rho P_{11} + P_{11}A_{11} + A_{11}^\top P_{11},$$

the remaining blocks solve linear equations, and the eigenvalues of the aggregate closed-loop matrix are those of $B_1\Gamma^{-1}B_1^\top P_{11} - A_{11}$ together with those of $-A_{22}$.

The blocks of Q and Θ_X that involve only the forcing process affect the value of the problem but not the policy, so they can be chosen freely to keep both matrices positive definite.

Here is a pricing example in the spirit of the real rigidities of Ball and Romer [1990].

A store's log price $x^{(1)}$ has a target that puts weight $1 - \chi$ on its own log marginal cost $x^{(2)}$ and weight $\chi \in [0, 1)$ on the average price $X^{(1)}$, and the store pays a Rotemberg cost to change its price:

$$-\tfrac{q}{2}\left(x^{(1)} - (1-\chi)x^{(2)} - \chi X^{(1)}\right)^2 - \tfrac{\gamma}{2}\alpha^2, \qquad dx^{(1)} = \alpha\, dt.$$

The weight χ is the static pass-through of competitors' prices, so it measures strategic complementarity, and $\Theta_{X,11} = -q\chi$.

Terms that pair the forcing process with the average price affect neither the policy nor the dynamics, so we drop them.

Costs follow $dx^{(2)} = -a_{22}x^{(2)}dt + \sigma dW + \upsilon\, d\mathcal{J}$, where W is idiosyncratic and $\mathcal{J}$ is a common jump process.

An idiosyncratic cost shock moves a store away from the mean, so the single-agent matrix $\bar{\beta}_2$ governs its response, while a common cost shock moves the industry, which P governs.

Let's trace both responses to a permanent one percent increase in costs, $a_{22} = 0$.

```
def pricing_with_costs(χ, a_22, ρ=0.04, q=1.0, γ=1.0):
    "Aggregate and deviation drift matrices in the pricing example with cost shocks."
    A_f = np.array([[0.0, 0.0],
                    [0.0, a_22]])
    B_f = np.array([[1.0],
                    [0.0]])
    Γ_f = np.array([[γ]])
    Q_f = q*np.array([[1.0, -(1 - χ)],
                      [-(1 - χ), 2.0]])       # (2, 2) element chosen freely
    Θ_f = np.array([[-q*χ, 0.0],
                    [0.0,  0.0]])
    β2_f = mfg_riccati(A_f, B_f, Q_f, Γ_f, ρ)
    P_f = mfg_riccati(A_f, B_f, Q_f + Θ_f, Γ_f, ρ)
    return closed_loop(A_f, B_f, Γ_f, P_f), closed_loop(A_f, B_f, Γ_f, β2_f)
```

The first matrix moves the industry's average price after a common cost shock, and the second moves a store's price relative to the average after an idiosyncratic one.

```
t_f = np.linspace(0, 6, 200)
cost_shock = np.array([0.0, 1.0])

fig, axes = plt.subplots(1, 2, figsize=(11, 4.2))
for χ in (0.0, 0.5, 0.8):
    U_agg, U_dev = pricing_with_costs(χ, 0.0)
    agg = [(expm(U_agg*t) @ cost_shock)[0] for t in t_f]
    dev = [(expm(U_dev*t) @ cost_shock)[0] for t in t_f]
    axes[0].plot(t_f, agg, lw=2, label=f"$\\chi = {χ}$")
    axes[1].plot(t_f, dev, lw=2, label=f"$\\chi = {χ}$")

axes[0].set_title('average price, common cost shock')
axes[1].set_title("store's relative price, idiosyncratic cost shock")
for ax in axes:
    ax.set_xlabel('years')
    ax.set_ylabel('percent')
    ax.legend()
plt.tight_layout()
plt.show()
```

A store passes an idiosyncratic cost increase through quickly but only partially, because its competitors' prices do not move: in the long run it passes through $1 - \chi$ of it.

The industry passes an aggregate cost increase through completely, but slowly, because each store waits for the others.

Let's compare the half-lives.

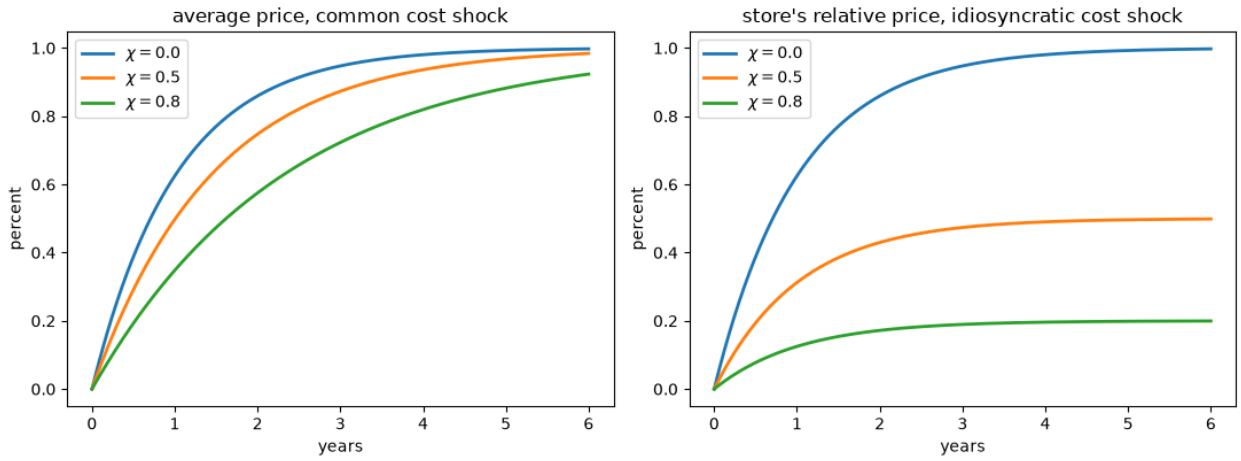


Fig. 130.1: Common and idiosyncratic cost pass-through

```
for χ in (0.0, 0.5, 0.8):
    U_agg, U_dev = pricing_with_costs(χ, 0.0)
    λ_agg_f, λ_dev_f = U_agg[0, 0], U_dev[0, 0]
    print(f"χ = {χ}: store half-life {half_life(λ_dev_f):.2f} years, "
          f"industry half-life {half_life(λ_agg_f):.2f} years")
```

```
χ = 0.0: store half-life 0.71 years, industry half-life 0.71 years
χ = 0.5: store half-life 0.71 years, industry half-life 1.01 years
χ = 0.8: store half-life 0.71 years, industry half-life 1.62 years
```

With $\chi = 0.8$ a store closes half of the gap to its new target in about 0.7 years, while the industry takes about 1.6 years.

The store's half-life is the same for every χ, because a store's deviation from the average is governed by its own curvature q alone.

So how fast stores respond to their own cost shocks reveals nothing about χ, which is the missing intercept of *Linear Quadratic Mean Field Games* in another guise.

How far they respond is a different matter: in the long run a store passes through $1 - \chi$ of its own cost shock, because χ also enters the store's own return, through the weight on its cost.

When a model ties the own curvature Q to the interaction Θ_χ in this way, micro data can be informative after all; the result of the previous lecture is that, for given Q, they are not.

Exercise 130.6.3 verifies *Proposition 130.4.1* in this example and studies transitory cost shocks.

130.5 Coupling through higher moments

So far agents have cared about the distribution of states only through its mean.

Alvarez and Argente [2026] also let payoffs depend on higher moments.

Take a scalar state and a scalar action, set $\theta_{\mathcal{A}} = 0$, and write $M_r(t) = \int x^r m(x,t)dx$ for the r-th moment of the cross-sectional distribution.

The period return is

$$-\tfrac{1}{2}qx^2 - \left(\theta_0 + \sum_{r=1}^{\bar r} \theta_r M_r\right) x - \tfrac{1}{2}\gamma\alpha^2,$$

the state follows $dx = (b\alpha - ax)dt + \sigma dW$, and the benchmark of *Linear Quadratic Mean Field Games* is the case $\bar r = 1$ and $\theta_0 = 0$.

A natural notion of strategic complementarity is that a stochastically higher distribution raises the marginal return to x, which here means that $-\sum_r \theta_r x^r$ is nondecreasing in x.

Our example below, with θ_1 and θ_3 negative and the other coefficients zero, satisfies it.

The value function is still quadratic, its curvature $\bar\beta_2$ is still that of the single-agent problem, and the individual eigenvalue

$$\lambda_{\text{ind}} \equiv \frac{b^2}{\gamma}\bar\beta_2 - a = \frac{\rho}{2} - \sqrt{\left(a + \frac{\rho}{2}\right)^2 + \frac{b^2 q}{\gamma}}$$

is unchanged.

What changes is aggregation.

The Kolmogorov forward equation implies

$$\dot M_r = r\lambda_{\text{ind}} M_r + r\frac{b^2}{\gamma}\beta_1 M_{r-1} + \frac{r(r-1)\sigma^2}{2} M_{r-2}, \qquad \dot\beta_1 = (\rho - \lambda_{\text{ind}})\beta_1 + \theta_0 + \sum_{r=1}^{\bar r} \theta_r M_r,$$

so the moments obey a finite system, but a nonlinear one, because of the products $\beta_1 M_{r-1}$.

When $\sigma > 0$, $\lambda_{\text{ind}} < 0$, and $\theta_0 = -\sum_{r \text{ even}} \theta_r \bar M_r$, there is a steady state with $\beta_1 = 0$ in which the invariant distribution is normal with mean zero and variance $v = \sigma^2/(-2\lambda_{\text{ind}})$.

Its odd moments vanish and its even moments are $\bar M_r = (r-1)!!\, v^{r/2}$; as $\sigma \to 0$ it collapses to a point mass at zero, which we include below as a limiting case.

Linearizing the system around this steady state gives $\dot y = \tilde{\mathcal{H}} y$ for $y = (\hat\beta_1, \hat M_1, \dots, \hat M_{\bar r})$, where hats denote deviations from the steady state.

Proposition 130.5.1

Let $\bar r$ be even and define

$$\bar\theta \equiv \sum_{r \text{ odd}} r\, \bar M_{r-1}\, \theta_r.$$

The eigenvalues of $\tilde{\mathcal{H}}$ are $\tilde\lambda_r = r\lambda_{\text{ind}}$ for $r = 2, \dots, \bar r$, together with

$$\tilde\lambda_1 = \frac{\rho}{2} - \sqrt{\left(a + \frac{\rho}{2}\right)^2 + \frac{b^2(q + \bar\theta)}{\gamma}}, \qquad \tilde\lambda_{\bar r + 1} = \frac{\rho}{2} + \sqrt{\left(a + \frac{\rho}{2}\right)^2 + \frac{b^2(q + \bar\theta)}{\gamma}}.$$

If $q + \bar\theta + \gamma a(a + \rho)/b^2 > 0$, the steady state is locally determinate, with one unstable and $\bar r$ stable eigenvalues.

Three features stand out.

First, the interaction coefficients enter only through $\bar\theta$, and only those on odd moments enter it, so $\bar\theta$ plays exactly the role of θ_X in the scalar formula of *Linear Quadratic Mean Field Games*.

Second, $\bar\theta$ weights $\theta_3, \theta_5, \dots$ by even moments of the invariant distribution, which grow with σ^2.

In the benchmark, the variance of idiosyncratic shocks never enters the Riccati equations, so it has no effect on aggregate dynamics; with higher moments, idiosyncratic risk affects aggregate persistence.

Third, coefficients on even moments do not affect the local eigenvalues, although they can affect transitional dynamics, as Exercise 130.6.4 shows.

Let's build $\tilde{\mathcal{H}}$ and check *Proposition 130.5.1*, using $\bar{r} = 4$ and interactions through the first and third moments.

```
def moment_system(ρ, a, b, q, γ, σ, θ):
    """
    Linearized matrix for (β̂_1, M̂_1, ..., M̂_r̄) around the zero-mean steady
    state, with θ = (θ_1, ..., θ_r̄).  Also returns the steady-state moments
    M̄_0, ..., M̄_r̄ and the individual eigenvalue.
    """
    r_bar = len(θ)
    λ_ind = ρ/2 - np.sqrt((a + ρ/2)**2 + b**2*q/γ)
    v = σ**2/(-2*λ_ind)                          # invariant variance
    M_bar = np.zeros(r_bar + 1)
    M_bar[0] = 1.0
    for r in range(2, r_bar + 1, 2):
        M_bar[r] = (r - 1)*v*M_bar[r - 2]
    H_tilde = np.zeros((r_bar + 1, r_bar + 1))
    H_tilde[0, 0] = ρ - λ_ind
    H_tilde[0, 1:] = θ
    for r in range(1, r_bar + 1):
        H_tilde[r, 0] = r*M_bar[r - 1]*b**2/γ
        H_tilde[r, r] = r*λ_ind
        if r >= 3:                               # M̂_0 is identically zero
            H_tilde[r, r - 2] = r*(r - 1)*σ**2/2
    return H_tilde, M_bar, λ_ind

ρ_m, a_m, b_m, q_m, γ_m = 0.05, 0.3, 1.0, 1.0, 1.0
θ_m = np.array([-0.3, 0.0, -0.2, 0.0])       # complementarity through M_1 and M_3

for σ in (0.0, 0.5, 1.0):
    H_m, M_bar, λ_ind = moment_system(ρ_m, a_m, b_m, q_m, γ_m, σ, θ_m)
    θ_bar = θ_m[0] + 3*M_bar[2]*θ_m[2]
    root = np.sqrt((a_m + ρ_m/2)**2 + b_m**2*(q_m + θ_bar)/γ_m)
    predicted = np.r_[ρ_m/2 - root, ρ_m/2 + root, [r*λ_ind for r in range(2, 5)]]
    print(f"σ = {σ}: eigenvalues {np.sort(np.linalg.eigvals(H_m).real).round(4)}")
    print(f"        predicted   {np.sort(predicted).round(4)}")
```

```
σ = 0.0: eigenvalues [-4.1059 -3.0795 -2.053  -0.8726  0.9226]
        predicted   [-4.1059 -3.0795 -2.053  -0.8726  0.9226]
σ = 0.5: eigenvalues [-4.1059 -3.0795 -2.053  -0.8309  0.8809]
        predicted   [-4.1059 -3.0795 -2.053  -0.8309  0.8809]
σ = 1.0: eigenvalues [-4.1059 -3.0795 -2.053  -0.6915  0.7415]
        predicted   [-4.1059 -3.0795 -2.053  -0.6915  0.7415]
```

To trace the equilibrium path we keep the $\bar{r}$ eigenvalues below $\rho/2$, express $\hat{\beta}_1$ as a linear function of the moment deviations on the stable subspace, and substitute.

That gives a law of motion for the moment deviations alone, so an initial condition specifies only those deviations, and $\hat{\beta}_1(0)$ follows from the saddle path.

```
def moment_dynamics(H_tilde, ρ):
    "Law of motion of the moment deviations on the stable subspace of H_tilde."
    ev, V = np.linalg.eig(H_tilde)
```

(continues on next page)

(continued from previous page)

```
    V_s = V[:, ev.real < ρ/2]
    s = np.real(V_s[0] @ np.linalg.inv(V_s[1:]))      # β̂₁ = s · M̂
    return H_tilde[1:, 1:] + np.outer(H_tilde[1:, 0], s)
```

As an initial condition we shift the invariant distribution to the right by Δ, which to first order raises M_1 by Δ and M_3 by $3v\Delta$ and leaves the even moments unchanged.

The next figure shows the path of the mean for several levels of idiosyncratic volatility.

```
t_m = np.linspace(0, 6, 200)
Δ = 0.1

fig, ax = plt.subplots(figsize=(8, 4.5))
for σ in (0.0, 0.6, 1.0, 1.3):
    H_m, M_bar, λ_ind = moment_system(ρ_m, a_m, b_m, q_m, γ_m, σ, θ_m)
    U_m = moment_dynamics(H_m, ρ_m)
    M_hat0 = np.array([Δ, 0.0, 3*M_bar[2]*Δ, 0.0])
    path = [(expm(U_m*t) @ M_hat0)[0] for t in t_m]
    θ_bar = θ_m[0] + 3*M_bar[2]*θ_m[2]
    λ_1 = ρ_m/2 - np.sqrt((a_m + ρ_m/2)**2 + b_m**2*(q_m + θ_bar)/γ_m)
    ax.plot(t_m, path, lw=2,
            label=f"$\\sigma = {σ}$, slowest half-life {half_life(λ_1):.2f}")
ax.set_xlabel('$t$')
ax.set_ylabel('cross-sectional mean $M_1(t)$')
ax.legend()
plt.tight_layout()
plt.show()
```

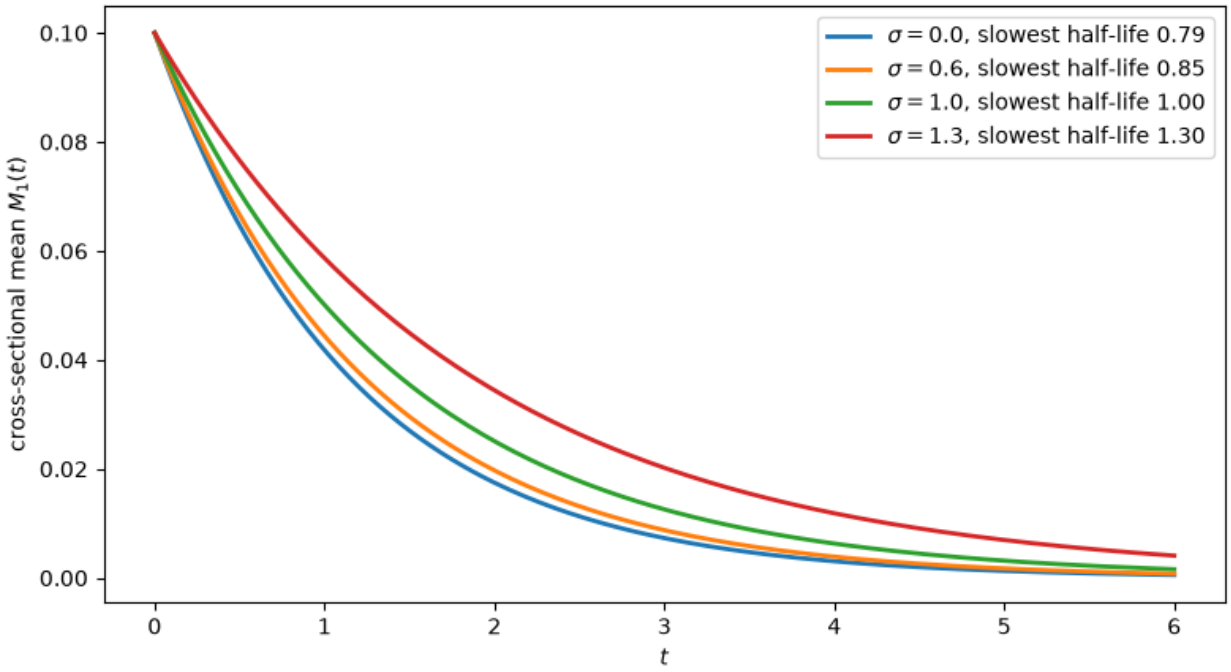


Fig. 130.2: Idiosyncratic risk and mean persistence

Greater idiosyncratic volatility makes the mean more persistent here.

A larger variance puts more weight on the third moment, through which complementarity operates, so the effective interaction $\bar{\theta}$ is stronger.

130.6 Exercises

Exercise 130.6.1

This exercise continues the Hayashi example.

1. Choose ϕ so that a firm that ignores the industry again closes half of a capital gap in three years under Hayashi costs, and recompute the industry's half-life and the ratio of the two half-lives.
2. Holding ϕ from part 1 and ε_s fixed, compute the aggregate eigenvalue in the limit of constant returns to scale, $\nu \to 1$, where $q + \theta_X \to 0$, once with no price effect and once with the calibrated supply elasticity. Explain why the answers differ from the limit $\lambda \to -\delta$ that holds without Hayashi costs.

Solution

For part 1, the individual's eigenvalue is $\rho/2 - \sqrt{\rho^2/4 + \delta^2 q/\gamma}$, so a target λ requires $\gamma = \delta^2 q/[(\rho/2-\lambda)^2 - \rho^2/4]$.

```
λ_target = -np.log(2)/3
γ_H = δ_c**2*q_c/((ρ_c/2 - λ_target)**2 - ρ_c**2/4)

λ_ind_H = λ_cross(ρ_c, δ_c, δ_c, q_c + γ_H, γ_H, γ_H)
λ_agg_H = λ_cross(ρ_c, δ_c, δ_c, q_c + γ_H + θ_X_c, γ_H + θ_A_c, γ_H)

print(f"γ = {γ_H:.4f}, so ψ = {γ_H/cap_scale(η_c, ν_c):.4f}")
print(f"individual half-life {half_life(λ_ind_H):.2f}, "
      f"industry half-life {half_life(λ_agg_H):.2f}, "
      f"ratio {half_life(λ_agg_H)/half_life(λ_ind_H):.2f}")
```

```
γ = 0.0384, so ψ = 0.1097
individual half-life 3.00, industry half-life 5.42, ratio 1.81
```

Matching the three-year individual half-life requires a cheaper installation technology than before, $\phi \approx 0.11$ instead of $1/7$.

The industry's half-life is then 5.4 years, so the gap between micro and macro adjustment widens: the ratio rises from 1.66 to 1.81.

For part 2, when $q + \theta_X = 0$ the effective state curvature in *Proposition 130.3.1* is γ, and the industry's effective return is

$$-\tfrac{1}{2}\gamma X^2 - \tfrac{1}{2}(\gamma + \theta_{\mathcal{A}})\mathcal{A}^2 + \gamma X \mathcal{A} = -\tfrac{\gamma}{2}(\mathcal{A} - X)^2 - \tfrac{\theta_{\mathcal{A}}}{2}\mathcal{A}^2.$$

In the limit $\nu \to 1$, the scale $\delta \bar{k}^{1-\nu}$ that converts ϕ and $1/\varepsilon_s$ into γ and $\theta_{\mathcal{A}}$ tends to $\delta(\eta - 1)/[\eta(\rho + \delta)]$.

Let's compute the limit with and without a price effect.

```
ψ_H = γ_H/cap_scale(η_c, ν_c)
scale_lim = cap_scale(η_c, 1.0)
γ_lim = ψ_H*scale_lim
for inv_ε in (0.0, 1/ε_s_c):
    θ_A_lim = inv_ε*scale_lim
    λ_lim = λ_cross(ρ_c, δ_c, δ_c, γ_lim, γ_lim + θ_A_lim, γ_lim)
    print(f"1/ε_s = {inv_ε:.3f}: θ_A = {θ_A_lim:.4f}, λ = {λ_lim:.5f}")
```

```
1/ε_s = 0.000: θ_A = 0.0000, λ = 0.00000
1/ε_s = 0.257: θ_A = 0.1286, λ = -0.08054
```

With $\theta_{\mathcal{A}} = 0$ the only cost is the cost of net investment, so any capital stock can be maintained for free and nothing pulls the industry back: the aggregate state has a unit root.

Without Hayashi costs, maintaining a capital stock above its steady state requires gross investment above $\bar{\imath}$, which $\psi(i)$ penalizes, so the industry lets depreciation do the work and $\lambda \to -\delta$.

With $\theta_{\mathcal{A}} > 0$ the price of investment goods penalizes gross investment away from $\bar{\imath}$, which restores some mean reversion.

Exercise 130.6.2

This exercise shows why the restriction $\Theta_{X\alpha} = \Theta_{x\mathcal{A}}$ matters, even with a single state.

Take $\rho = 0.05$, $a = 0.3$, $b = q = \gamma = 1$, $\theta_X = \theta_{\mathcal{A}} = 0$ and no own cross term, and use `state_costate` with $C_1 = -\theta_{x\mathcal{A}}$ and $C_2 = -\theta_{X\alpha}$.

1. Show that the trace of the state-costate matrix is $\rho + b(\theta_{x\mathcal{A}} - \theta_{X\alpha})/(\gamma + \theta_{\mathcal{A}})$, so that its eigenvalues can be symmetric about $\rho/2$ only when $\theta_{X\alpha} = \theta_{x\mathcal{A}}$.
2. Compute the eigenvalues for $(\theta_{X\alpha}, \theta_{x\mathcal{A}})$ equal to $(0, 0)$, $(-2, -2)$, $(-1.5, -2.5)$ and $(-2.5, -1.5)$, and count how many have real parts below $\rho/2$. What does each count imply about the number of equilibria?

Solution

For part 1, with $g = \gamma + \theta_{\mathcal{A}}$ the diagonal elements of the matrix are $\rho + a - bC_1/g$ and $bC_2/g - a$.

Their sum is $\rho + b(C_2 - C_1)/g = \rho + b(\theta_{x\mathcal{A}} - \theta_{X\alpha})/g$.

Eigenvalues symmetric about $\rho/2$ must sum to ρ, which requires $\theta_{X\alpha} = \theta_{x\mathcal{A}}$.

For part 2, we compute the eigenvalues in each case.

```
ρ_e, a_e = 0.05, 0.3

print(f"{'(θ_Xα, θ_xA)':>14}{'eigenvalues':>34}{'trace':>8}{'below ρ/2':>11}")
for θ_Xα_e, θ_xA_e in ((0.0, 0.0), (-2.0, -2.0), (-1.5, -2.5), (-2.5, -1.5)):
    M_e = state_costate(a_e*one, one, one, one, ρ_e,
                        -θ_xA_e*one, -θ_Xα_e*one)
    ev = np.linalg.eigvals(M_e)
    print(f"{str((θ_Xα_e, θ_xA_e)):>14}{str(np.round(ev, 4)):>34}"
          f"{ev.real.sum():>8.3f}{np.sum(ev.real < ρ_e/2):>11}")
```

```
  (θ_Xα, θ_xA)                      eigenvalues   trace  below ρ/2
    (0.0, 0.0)                [ 1.0765 -1.0265]   0.050          1
  (-2.0, -2.0)     [0.025+0.4409j 0.025-0.4409j]   0.050          0
  (-1.5, -2.5)                [-0.7108 -0.2392]  -0.950          2
  (-2.5, -1.5)                  [0.2892 0.7608]   1.050          0
```

With no mixed interactions there is one eigenvalue on each side of $\rho/2$ and a unique equilibrium.

At $(-2, -2)$ the restriction holds, but the sufficient condition $q - \tilde{C}^2/g > 0$ of *Proposition 130.3.1* fails, since $1 - 4 < 0$.

The eigenvalues are then complex with real part exactly $\rho/2$, so for $X(0) \neq 0$ there is no admissible path and no equilibrium.

At $(-1.5, -2.5)$ both eigenvalues are below $\rho/2$, so every initial costate $w(0)$ generates an admissible path.

Each of these paths is an equilibrium.

Given aggregate paths that grow more slowly than $e^{\rho t/2}$, the individual's costate equation has a unique admissible solution, because its own root $\rho - \lambda_{\text{ind}}$ exceeds $\rho/2$, and the w on any path of the system is that solution.

So there is a continuum of equilibria for each initial distribution.

At $(-2.5, -1.5)$ both eigenvalues are above $\rho/2$, and again no equilibrium exists for $X(0) \neq 0$.

Exercise 130.6.3

This exercise works with the pricing example of *Forcing processes and pass-through*.

1. Take $\chi = 0.8$ and $a_{22} = 0.5$. Check that the eigenvalues of the aggregate closed-loop matrix are those of $B_1 \Gamma^{-1} B_1^\top P_{11} - A_{11}$ together with $-a_{22}$, as *Proposition 130.4.1* says. Check also that changing the $(2,2)$ elements of Q and Θ_X leaves the first row of P, and hence the pricing policy, unchanged.
2. With transitory cost shocks, $a_{22} = 0.5$, plot the industry's response to a common one percent cost increase for $\chi \in \{0, 0.5, 0.8\}$. How do the size and timing of the peak response depend on χ?

Solution

For part 1, we rebuild the example with the $(2,2)$ elements as arguments.

```
χ_e, a_22_e, ρ_f, q_f, γ_f = 0.8, 0.5, 0.04, 1.0, 1.0
A_f = np.array([[0.0, 0.0],
                [0.0, a_22_e]])
B_f = np.array([[1.0],
                [0.0]])
Γ_f = np.array([[γ_f]])

def P_forcing(q_22, θ_22):
    "Equilibrium P when the (2, 2) elements of Q and Θ_X are q_22 and θ_22."
    Q_f = q_f*np.array([[1.0, -(1 - χ_e)],
                        [-(1 - χ_e), 0.0]])
    Q_f[1, 1] = q_22
    Θ_f = np.array([[-q_f*χ_e, 0.0],
                    [0.0, θ_22]])
    return mfg_riccati(A_f, B_f, Q_f + Θ_f, Γ_f, ρ_f)

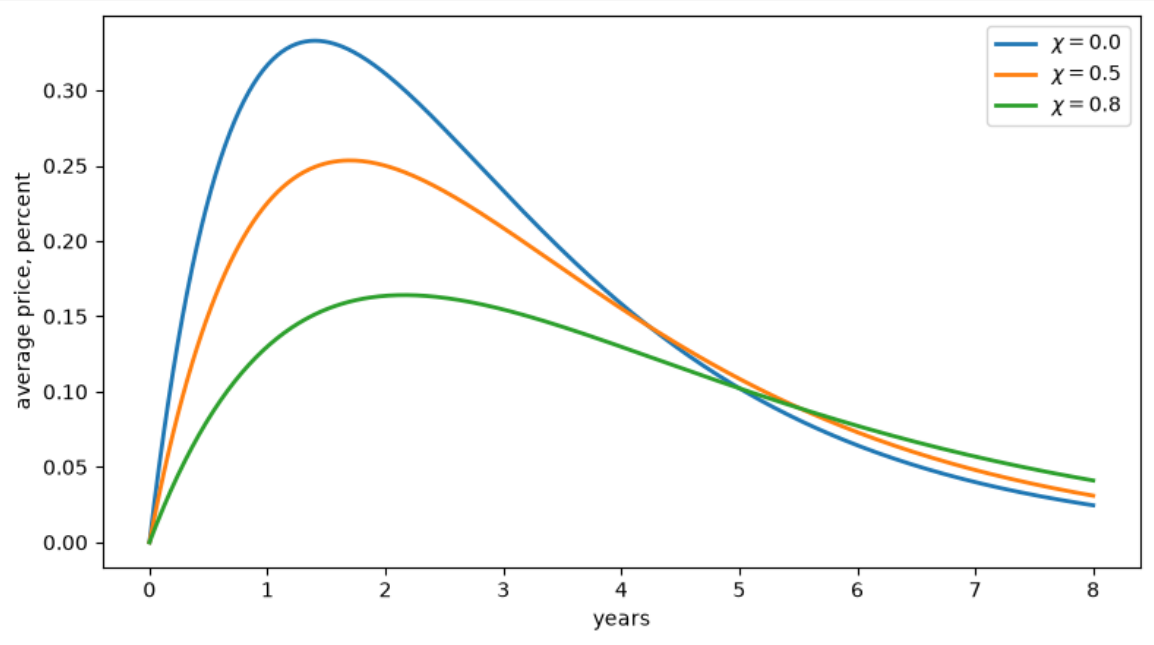
P_a, P_b = P_forcing(2.0, 0.0), P_forcing(7.0, 3.0)
λ_11 = P_a[0, 0]/γ_f          # B_1 Γ^{-1} B_1' P_11 - A_11
ev_cl = np.sort(np.linalg.eigvals(closed_loop(A_f, B_f, Γ_f, P_a)).real)
print("closed-loop eigenvalues:   ", ev_cl.round(5))
print("controlled block and -a_22:", np.sort([λ_11, -a_22_e]).round(5))
print("first row of P unchanged:", np.allclose(P_a[0], P_b[0]))
```

```
closed-loop eigenvalues:    [-0.5     -0.42766]
controlled block and -a_22: [-0.5     -0.42766]
first row of P unchanged: True
```

For part 2, we plot the response of the average price.

```
t_e = np.linspace(0, 8, 400)
fig, ax = plt.subplots(figsize=(8, 4.5))
for χ in (0.0, 0.5, 0.8):
    U_agg, _ = pricing_with_costs(χ, 0.5)
    agg = np.array([(expm(U_agg*t) @ np.array([0.0, 1.0]))[0] for t in t_e])
    k = agg.argmax()
    ax.plot(t_e, agg, lw=2, label=f"$\\chi = {χ}$")
    print(f"χ = {χ}: peak {agg[k]:.3f} percent after {t_e[k]:.2f} years")
ax.set_xlabel('years')
ax.set_ylabel('average price, percent')
ax.legend()
plt.tight_layout()
plt.show()
```

```
χ = 0.0: peak 0.333 percent after 1.40 years
χ = 0.5: peak 0.253 percent after 1.70 years
χ = 0.8: peak 0.164 percent after 2.17 years
```



Stronger complementarity makes the peak response smaller and later.

Each store raises its price more slowly because the others do, so by the time the industry's price has risen much, the cost shock has already started to fade.

Exercise 130.6.4

This exercise shows how an interaction through an even moment affects transitional dynamics without changing the local eigenvalues.

Take $\bar{r} = 2, \theta_1 = -0.3, \sigma = 1$ and the other parameters of the moment example in *Coupling through higher moments*.

1. Verify that θ_2 changes none of the eigenvalues of $\tilde{\mathcal{H}}$.

2. Suppose the cross-sectional variance rises by 0.1 while the mean stays at zero. Trace the path of the mean for $\theta_2 \in \{-0.5, 0, 0.5\}$ and explain what you see.

3. Alvarez and Argente [2026] show that along the saddle path $\dot{\hat{M}}_1 = \tilde{\lambda}_1 \hat{M}_1 + \zeta \hat{M}_2$ with $\zeta = b^2\theta_2/[\gamma(\tilde{\lambda}_1 + 2\lambda_{\text{ind}} - \rho)]$. Check this formula.

Solution

For part 1, we compute the eigenvalues for each θ_2.

```
θ_1 = -0.3
for θ_2 in (-0.5, 0.0, 0.5):
    H_2, _, _ = moment_system(ρ_m, a_m, b_m, q_m, γ_m, 1.0, np.array([θ_1, θ_2]))
    print(f"θ_2 = {θ_2:>4}: eigenvalues {np.sort(np.linalg.eigvals(H_2).real).
↪round(5)}")
```

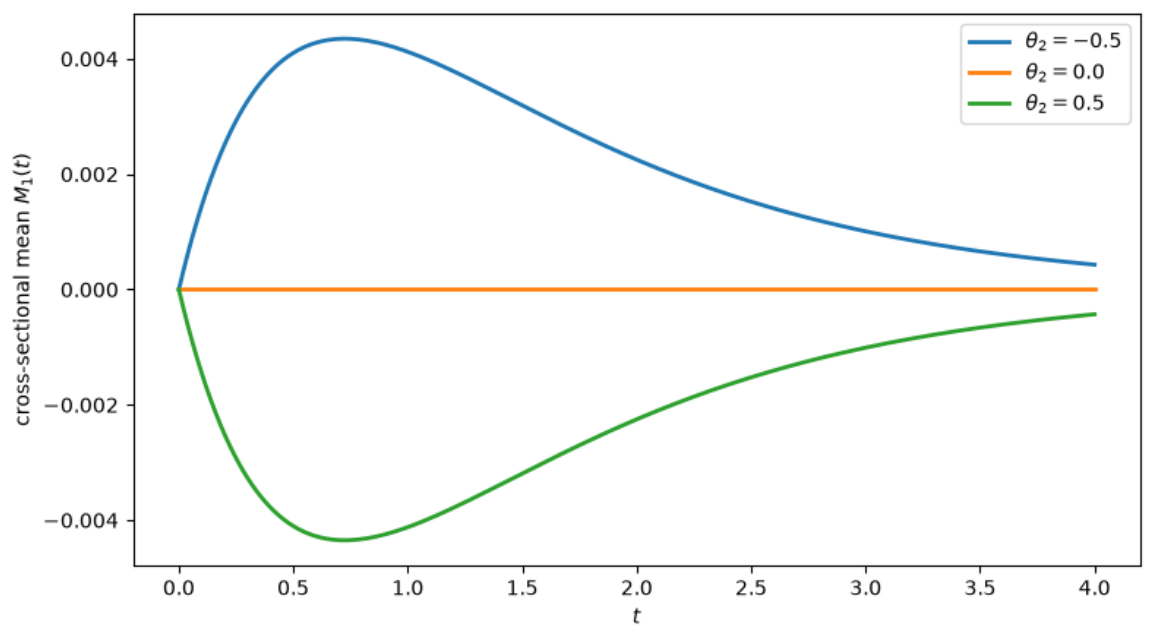
```
θ_2 = -0.5: eigenvalues [-2.05297 -0.87257  0.92257]
θ_2 =  0.0: eigenvalues [-2.05297 -0.87257  0.92257]
θ_2 =  0.5: eigenvalues [-2.05297 -0.87257  0.92257]
```

They do not move.

At the zero-mean steady state the second moment evolves on its own, $\dot{\hat{M}}_2 = 2\lambda_{\text{ind}} \hat{M}_2$, because the term $2(b^2/\gamma)\beta_1 M_1$ in its law of motion vanishes to first order.

For part 2, we start from $\hat{M}_1 = 0$ and $\hat{M}_2 = 0.1$.

```
t_e = np.linspace(0, 4, 200)
fig, ax = plt.subplots(figsize=(8, 4.5))
for θ_2 in (-0.5, 0.0, 0.5):
    H_2, _, _ = moment_system(ρ_m, a_m, b_m, q_m, γ_m, 1.0, np.array([θ_1, θ_2]))
    U_2 = moment_dynamics(H_2, ρ_m)
    path = [(expm(U_2*t) @ np.array([0.0, 0.1]))[0] for t in t_e]
    ax.plot(t_e, path, lw=2, label=f"$\\theta_2 = {θ_2}$")
ax.set_xlabel('$t$')
ax.set_ylabel('cross-sectional mean $M_1(t)$')
ax.legend()
plt.tight_layout()
plt.show()
```



With $\theta_2 > 0$, a higher second moment lowers everybody's marginal return to x through the term $-\theta_2 M_2 x$, so agents lower x while the variance is high, and the mean dips before returning as the variance returns to its invariant level.

With $\theta_2 < 0$ the mean rises instead, and with $\theta_2 = 0$ it does not move at all.

So a shock to dispersion moves the average, even though the eigenvalues that govern local persistence are unchanged.

For part 3, we compare the coefficients of the saddle-path law of motion with the formula.

```
for θ_2 in (-0.5, 0.5):
    H_2, _, λ_ind_2 = moment_system(ρ_m, a_m, b_m, q_m, γ_m, 1.0, np.array([θ_1, θ_
↪2]))
    U_2 = moment_dynamics(H_2, ρ_m)
    λ_1 = ρ_m/2 - np.sqrt((a_m + ρ_m/2)**2 + b_m**2*(q_m + θ_1)/γ_m)
    ζ = b_m**2*θ_2/(γ_m*(λ_1 + 2*λ_ind_2 - ρ_m))
    print(f"θ_2 = {θ_2}: coefficients ({U_2[0, 0]:.5f}, {U_2[0, 1]:.5f}), "
          f"formula ({λ_1:.5f}, {ζ:.5f})")
```

```
θ_2 = -0.5: coefficients (-0.87257, 0.16804), formula (-0.87257, 0.16804)
θ_2 = 0.5: coefficients (-0.87257, -0.16804), formula (-0.87257, -0.16804)
```

130.7 Further reading

Alvarez and Argente [2026] develop the cross terms and the forcing processes in the appendices of their paper, and the coupling through higher moments in its main text.

The Hayashi and pass-through examples, and the one-state analysis of the mixed interactions in Exercise 130.6.2, are our own illustrations of their results.

Part XVI

Asset Pricing and Finance

CHAPTER

ONE

ASSET PRICING: FINITE STATE MODELS

Contents

- *Asset Pricing: Finite State Models*
  - *Overview*
  - *Pricing Models*
  - *Prices in the Risk-Neutral Case*
  - *Risk Aversion and Asset Prices*
  - *Exercises*

> "A little knowledge of geometric series goes a long way" – Robert E. Lucas, Jr.
>
> "Asset pricing is all about covariances" – Lars Peter Hansen

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install quantecon
```

131.1 Overview

An asset is a claim on one or more future payoffs.

The spot price of an asset depends primarily on

- the anticipated income stream
- attitudes about risk
- rates of time preference

In this lecture, we consider some standard pricing models and dividend stream specifications.

We study how prices and dividend-price ratios respond in these different scenarios.

We also look at creating and pricing *derivative* assets that repackage income streams.

Key tools for the lecture are

- Markov processses
- formulas for predicting future values of functions of a Markov state

- a formula for predicting the discounted sum of future values of a Markov state

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
import quantecon as qe
from numpy.linalg import eigvals, solve
```

131.2 Pricing Models

Let $\{d_t\}_{t \geq 0}$ be a stream of dividends

- A time-t **cum-dividend** asset is a claim to the stream $d_t, d_{t+1}, ...$
- A time-t **ex-dividend** asset is a claim to the stream $d_{t+1}, d_{t+2}, ...$

Let's look at some equations that we expect to hold for prices of assets under ex-dividend contracts (we will consider cum-dividend pricing in the exercises).

131.2.1 Risk-Neutral Pricing

Our first scenario is risk-neutral pricing.

Let $\beta = 1/(1 + \rho)$ be an intertemporal discount **factor**, where ρ is the **rate** at which agents discount the future.

The basic risk-neutral asset pricing equation for pricing one unit of an ex-dividend asset is

$$p_t = \beta \mathbb{E}_t[d_{t+1} + p_{t+1}] \tag{131.1}$$

This is a simple "cost equals expected benefit" relationship.

Here $\mathbb{E}_t[y]$ denotes the best forecast of y, conditioned on information available at time t.

More precisely, $\mathbb{E}_t[y]$ is the mathematical expectation of y conditional on information available at time t.

131.2.2 Pricing with Random Discount Factor

What happens if for some reason traders discount payouts differently depending on the state of the world?

Michael Harrison and David Kreps [Harrison and Kreps, 1979] and Lars Peter Hansen and Scott Richard [Hansen and Richard, 1987] showed that in quite general settings the price of an ex-dividend asset obeys

$$p_t = \mathbb{E}_t \left[m_{t+1} (d_{t+1} + p_{t+1}) \right] \tag{131.2}$$

for some **stochastic discount factor** m_{t+1}.

Here the fixed discount factor β in (131.1) has been replaced by the random variable m_{t+1}.

How anticipated future payoffs are evaluated now depends on statistical properties of m_{t+1}.

The stochastic discount factor can be specified to capture the idea that assets that tend to have good payoffs in bad states of the world are valued more highly than other assets whose payoffs don't behave that way.

This is because such assets pay well when funds are more urgently wanted.

We give examples of how the stochastic discount factor has been modeled below.

131.2.3 Asset Pricing and Covariances

Recall that, from the definition of a conditional covariance $\text{cov}_t(x_{t+1}, y_{t+1})$, we have

$$\mathbb{E}_t(x_{t+1}y_{t+1}) = \text{cov}_t(x_{t+1}, y_{t+1}) + \mathbb{E}_t x_{t+1}\mathbb{E}_t y_{t+1} \tag{131.3}$$

If we apply this definition to the asset pricing equation (131.2) we obtain

$$p_t = \mathbb{E}_t m_{t+1}\mathbb{E}_t(d_{t+1} + p_{t+1}) + \text{cov}_t(m_{t+1}, d_{t+1} + p_{t+1}) \tag{131.4}$$

It is useful to regard equation (131.4) as a generalization of equation (131.1)

- In equation (131.1), the stochastic discount factor $m_{t+1} = \beta$, a constant.
- In equation (131.1), the covariance term $\text{cov}_t(m_{t+1}, d_{t+1} + p_{t+1})$ is zero because $m_{t+1} = \beta$.
- In equation (131.1), $\mathbb{E}_t m_{t+1}$ can be interpreted as the reciprocal of the one-period risk-free gross interest rate.
- When m_{t+1} covaries more negatively with the payout $p_{t+1} + d_{t+1}$, the price of the asset is lower.

Equation (131.4) asserts that the covariance of the stochastic discount factor with the one period payout $d_{t+1} + p_{t+1}$ is an important determinant of the price p_t.

We give examples of some models of stochastic discount factors that have been proposed later in this lecture and also in a later lecture.

131.2.4 The Price-Dividend Ratio

Aside from prices, another quantity of interest is the **price-dividend ratio** $v_t := p_t/d_t$.

Let's write down an expression that this ratio should satisfy.

We can divide both sides of (131.2) by d_t to get

$$v_t = \mathbb{E}_t\left[m_{t+1}\frac{d_{t+1}}{d_t}(1 + v_{t+1})\right] \tag{131.5}$$

Below we'll discuss the implication of this equation.

131.3 Prices in the Risk-Neutral Case

What can we say about price dynamics on the basis of the models described above?

The answer to this question depends on

1. the process we specify for dividends
2. the stochastic discount factor and how it correlates with dividends

For now we'll study the risk-neutral case in which the stochastic discount factor is constant.

We'll focus on how an asset price depends on a dividend process.

131.3.1 Example 1: Constant Dividends

The simplest case is risk-neutral price of a constant, non-random dividend stream $d_t = d > 0$.

Removing the expectation from (131.1) and iterating forward gives

$$\begin{aligned} p_t &= \beta(d + p_{t+1}) \\ &= \beta(d + \beta(d + p_{t+2})) \\ &\quad \vdots \\ &= \beta(d + \beta d + \beta^2 d + \cdots + \beta^{k-2} d + \beta^{k-1} p_{t+k}) \end{aligned}$$

If $\lim_{k \to +\infty} \beta^{k-1} p_{t+k} = 0$, this sequence converges to

$$\bar{p} := \frac{\beta d}{1 - \beta} \tag{131.6}$$

This is the equilibrium price in the constant dividend case.

Indeed, simple algebra shows that setting $p_t = \bar{p}$ for all t satisfies the difference equation $p_t = \beta(d + p_{t+1})$.

131.3.2 Example 2: Dividends with Deterministic Growth Paths

Consider a growing, non-random dividend process $d_{t+1} = g d_t$ where $0 < g\beta < 1$.

While prices are not usually constant when dividends grow over time, a price dividend-ratio can be.

If we guess this, substituting $v_t = v$ into (131.5) as well as our other assumptions, we get $v = \beta g(1 + v)$.

Since $\beta g < 1$, we have a unique positive solution:

$$v = \frac{\beta g}{1 - \beta g}$$

The price is then

$$p_t = \frac{\beta g}{1 - \beta g} d_t$$

If, in this example, we take $g = 1 + \kappa$ and let $\rho := 1/\beta - 1$, then the price becomes

$$p_t = \frac{1 + \kappa}{\rho - \kappa} d_t$$

This is called the **Gordon formula**.

131.3.3 Example 3: Markov Growth, Risk-Neutral Pricing

Next, we consider a dividend process

$$d_{t+1} = g_{t+1} d_t \tag{131.7}$$

The stochastic growth factor $\{g_t\}$ is given by

$$g_t = g(X_t), \quad t = 1, 2, ...$$

where

1. $\{X_t\}$ is a finite Markov chain with state space S and transition probabilities

$$P(x, y) := \mathbb{P}\{X_{t+1} = y \,|\, X_t = x\} \qquad (x, y \in S)$$

2. g is a given function on S taking nonnegative values

You can think of

- S as n possible "states of the world" and X_t as the current state.
- g as a function that maps a given state X_t into a growth of dividends factor $g_t = g(X_t)$.
- $\ln g_t = \ln(d_{t+1}/d_t)$ is the growth rate of dividends.

(For a refresher on notation and theory for finite Markov chains see *this lecture*)

The next figure shows a simulation, where

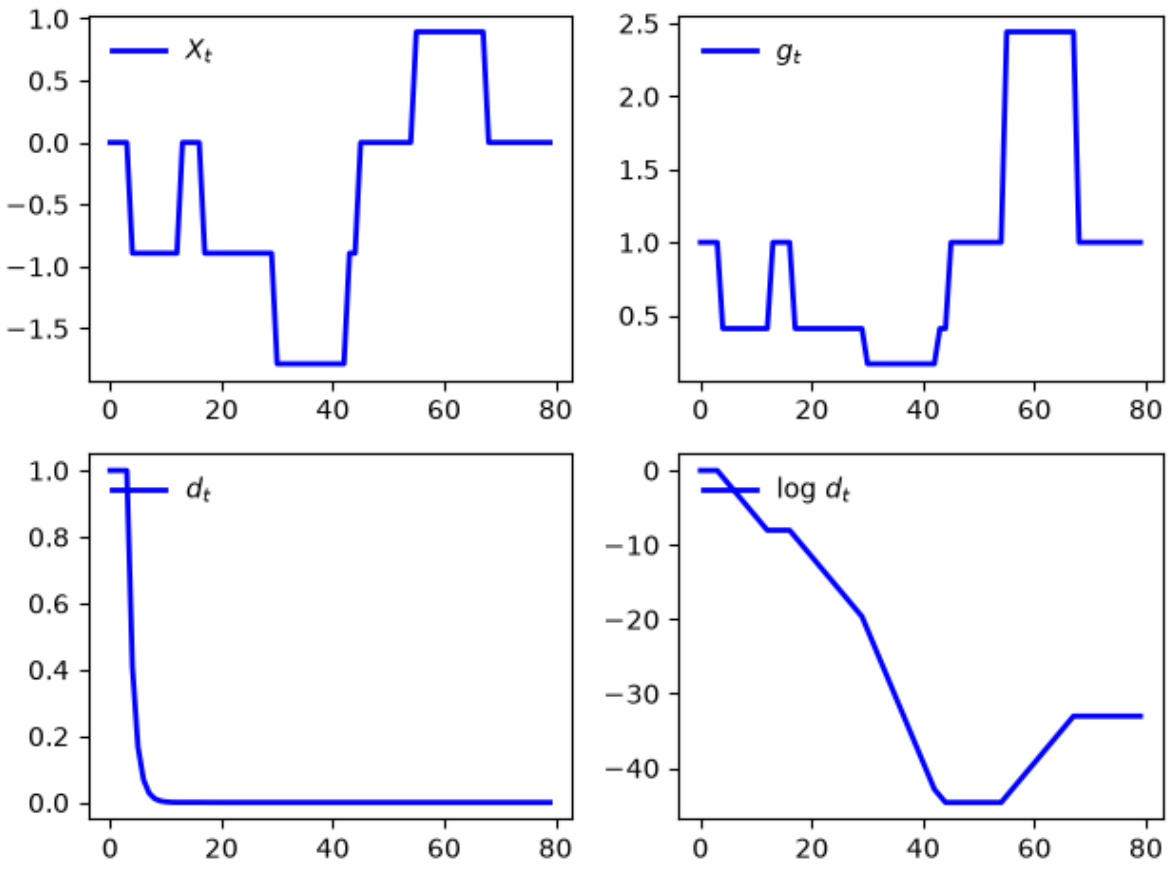
- $\{X_t\}$ evolves as a discretized AR1 process produced using *Tauchen's method*.
- $g_t = \exp(X_t)$, so that $\ln g_t = X_t$ is the growth rate.

```
n = 7
mc = qe.tauchen(n, 0.96, 0.25)
sim_length = 80

x_series = mc.simulate(sim_length, init=np.median(mc.state_values))
g_series = np.exp(x_series)
d_series = np.cumprod(g_series) # Assumes d_0 = 1

series = [x_series, g_series, d_series, np.log(d_series)]
labels = ['$X_t$', '$g_t$', '$d_t$', r'$\log \, d_t$']

fig, axes = plt.subplots(2, 2)
for ax, s, label in zip(axes.flatten(), series, labels):
    ax.plot(s, 'b-', lw=2, label=label)
    ax.legend(loc='upper left', frameon=False)
plt.tight_layout()
plt.show()
```



Pricing Formula

To obtain asset prices in this setting, let's adapt our analysis from the case of deterministic growth.

In that case, we found that v is constant.

This encourages us to guess that, in the current case, v_t is a fixed function of the state X_t.

We seek a function v such that the price-dividend ratio satisfies $v_t = v(X_t)$.

We can substitute this guess into (131.5) to get

$$v(X_t) = \beta \mathbb{E}_t[g(X_{t+1})(1 + v(X_{t+1}))]$$

If we condition on $X_t = x$, this becomes

$$v(x) = \beta \sum_{y \in S} g(y)(1 + v(y))P(x, y)$$

or

$$v(x) = \beta \sum_{y \in S} K(x, y)(1 + v(y)) \quad \text{where} \quad K(x, y) := g(y)P(x, y) \tag{131.8}$$

Suppose that there are n possible states $x_1, \ldots, x_n$.

We can then think of (131.8) as n stacked equations, one for each state, and write it in matrix form as

$$v = \beta K(\mathbb{1} + v) \tag{131.9}$$

Here

- v is understood to be the column vector $(v(x_1), \ldots, v(x_n))'$.
- K is the matrix $(K(x_i, x_j))_{1 \leq i,j \leq n}$.
- $\mathbb{1}$ is a column vector of ones.

When does equation (131.9) have a unique solution?

From the *Neumann series lemma* and Gelfand's formula, equation (131.9) has a unique solution when βK has spectral radius strictly less than one.

Thus, we require that the eigenvalues of K be strictly less than β^{-1} in modulus.

The solution is then

$$v = (I - \beta K)^{-1} \beta K \mathbb{1} \tag{131.10}$$

131.3.4 Code

Let's calculate and plot the price-dividend ratio at some parameters.

As before, we'll generate $\{X_t\}$ as a *discretized AR1 process* and set $g_t = \exp(X_t)$.

Here's the code, including a test of the spectral radius condition

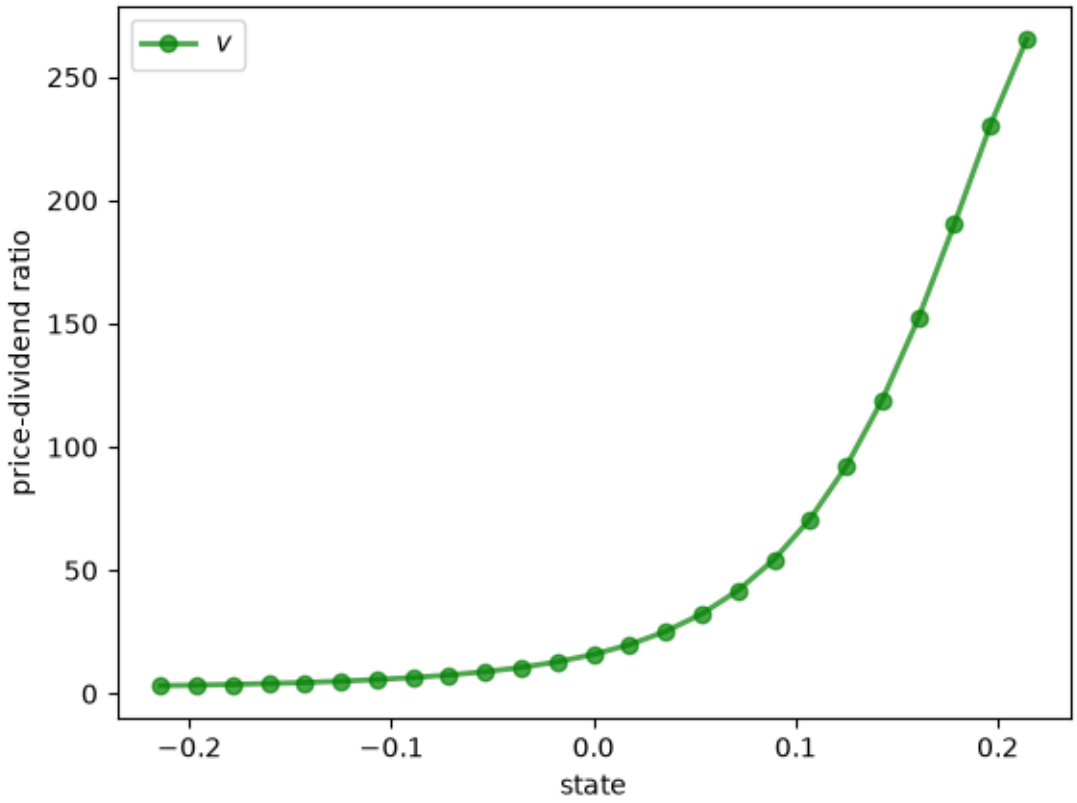
```
n = 25  # Size of state space
β = 0.9
mc = qe.tauchen(n, 0.96, 0.02)

K = mc.P * np.exp(mc.state_values)

warning_message = "Spectral radius condition fails"
assert np.max(np.abs(eigvals(K))) < 1 / β,  warning_message

I = np.identity(n)
v = solve(I - β * K, β * K @ np.ones(n))

fig, ax = plt.subplots()
ax.plot(mc.state_values, v, 'g-o', lw=2, alpha=0.7, label='$v$')
ax.set_ylabel("price-dividend ratio")
ax.set_xlabel("state")
ax.legend(loc='upper left')
plt.show()
```



Why does the price-dividend ratio increase with the state?

The reason is that this Markov process is positively correlated, so high current states suggest high future states.

Moreover, dividend growth is increasing in the state.

The anticipation of high future dividend growth leads to a high price-dividend ratio.

131.4 Risk Aversion and Asset Prices

Now let's turn to the case where agents are risk averse.

We'll price several distinct assets, including

- An endowment stream
- A consol (a type of bond issued by the UK government in the 19th century)
- Call options on a consol

131.4.1 Pricing a Lucas Tree

Let's start with a version of the celebrated asset pricing model of Robert E. Lucas, Jr. [Lucas, 1978].

Lucas considered an abstract pure exchange economy with these features:

- a single non-storable consumption good
- a Markov process that governs the total amount of the consumption good available each period
- a single *tree* that each period yields *fruit* that equals the total amount of consumption available to the economy
- a competitive market in *shares* in the tree that entitles their owners to corresponding shares of the *dividend* stream, i.e., the *fruit* stream, yielded by the tree
- a representative consumer who in a competitive equilibrium
  - consumes the economy's entire endowment each period
  - owns 100 percent of the shares in the tree

As in [Lucas, 1978], we suppose that the stochastic discount factor takes the form

$$m_{t+1} = \beta \frac{u'(c_{t+1})}{u'(c_t)} \tag{131.11}$$

where u is a concave utility function and c_t is time t consumption of a representative consumer.

(A derivation of this expression is given in a later lecture)

Assume the existence of an endowment that follows growth process (131.7).

The asset being priced is a claim on the endowment process, i.e., the **Lucas tree** described above.

Following [Lucas, 1978], we suppose that in equilibrium the representative consumer's consumption equals the aggregate endowment, so that $d_t = c_t$ for all t.

For utility, we'll assume the **constant relative risk aversion** (CRRA) specification

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma} \text{ with } \gamma > 0 \tag{131.12}$$

When $\gamma = 1$ we let $u(c) = \ln c$.

Inserting the CRRA specification into (131.11) and using $c_t = d_t$ gives

$$m_{t+1} = \beta \left(\frac{c_{t+1}}{c_t} \right)^{-\gamma} = \beta g_{t+1}^{-\gamma} \tag{131.13}$$

Substituting this into (131.5) gives the price-dividend ratio formula

$$v(X_t) = \beta \mathbb{E}_t \left[g(X_{t+1})^{1-\gamma} (1 + v(X_{t+1})) \right] \tag{131.14}$$

Conditioning on $X_t = x$, we can write this as

$$v(x) = \beta \sum_{y \in S} g(y)^{1-\gamma} (1 + v(y)) P(x, y)$$

If we let

$$J(x, y) := g(y)^{1-\gamma} P(x, y)$$

then we can rewrite equation (131.14) in vector form as

$$v = \beta J(\mathbb{1} + v)$$

Assuming that the spectral radius of J is strictly less than β^{-1}, this equation has the unique solution

$$v = (I - \beta J)^{-1} \beta J \mathbb{1} \tag{131.15}$$

We will define a function tree_price to compute v given parameters stored in the class AssetPriceModel

```
class AssetPriceModel:
    """
    A class that stores the primitives of the asset pricing model.

    Parameters
    ----------
    β : scalar, float
        Discount factor
    mc : MarkovChain
        Contains the transition matrix and set of state values for the state
        process
    γ : scalar(float)
        Coefficient of risk aversion
    g : callable
        The function mapping states to growth rates

    """
    def __init__(self, β=0.96, mc=None, γ=2.0, g=np.exp):
        self.β, self.γ = β, γ
        self.g = g

        # A default process for the Markov chain
        if mc is None:
            self.ρ = 0.9
            self.σ = 0.02
            self.mc = qe.tauchen(n, self.ρ, self.σ)
        else:
            self.mc = mc

        self.n = self.mc.P.shape[0]

    def test_stability(self, Q):
        """
        Stability test for a given matrix Q.
        """
        sr = np.max(np.abs(eigvals(Q)))
        if not sr < 1 / self.β:
            msg = f"Spectral radius condition failed with radius = {sr}"
            raise ValueError(msg)


def tree_price(ap):
    """
    Computes the price-dividend ratio of the Lucas tree.

    Parameters
    ----------
    ap: AssetPriceModel
        An instance of AssetPriceModel containing primitives

    Returns
    -------
```

(continues on next page)

(continued from previous page)

```
    v : array_like(float)
        Lucas tree price-dividend ratio

    """
    # Simplify names, set up matrices
    β, γ, P, y = ap.β, ap.γ, ap.mc.P, ap.mc.state_values
    J = P * ap.g(y)**(1 - γ)

    # Make sure that a unique solution exists
    ap.test_stability(J)

    # Compute v
    I = np.identity(ap.n)
    Ones = np.ones(ap.n)
    v = solve(I - β * J, β * J @ Ones)

    return v
```

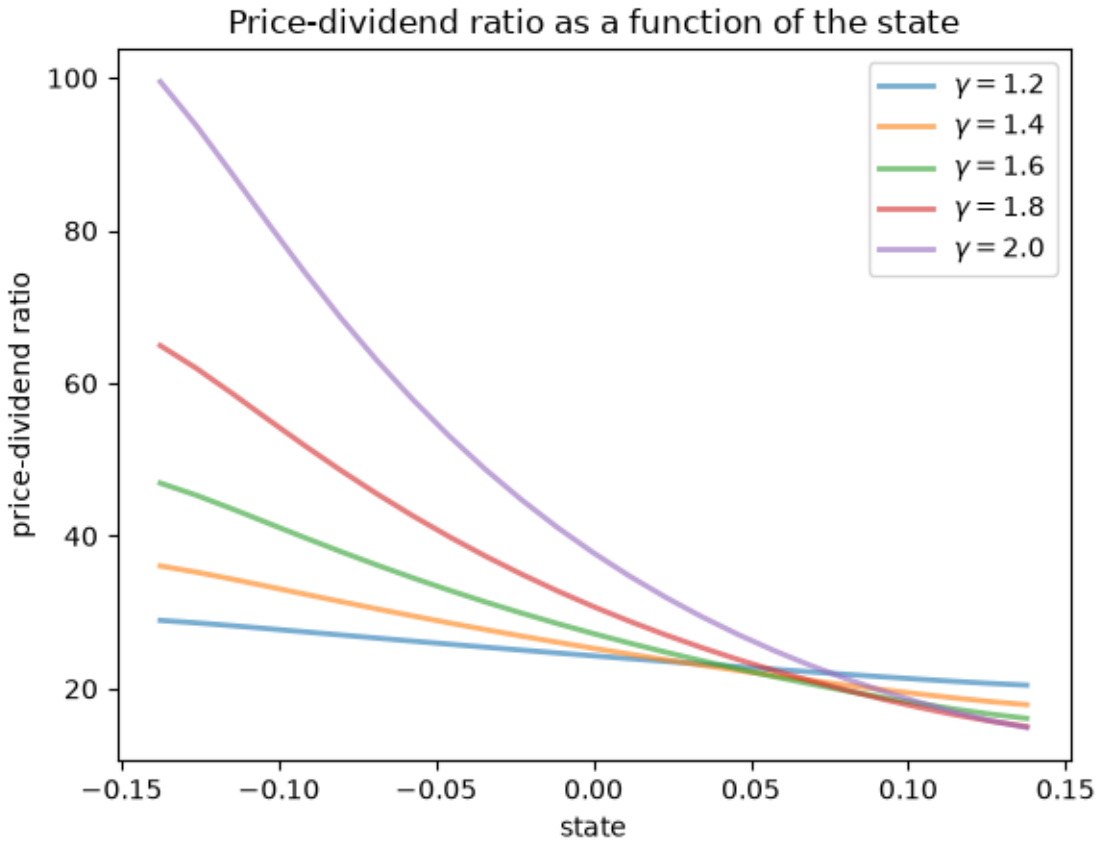
Here's a plot of v as a function of the state for several values of γ, with a positively correlated Markov process and $g(x) = \exp(x)$

```
ys = [1.2, 1.4, 1.6, 1.8, 2.0]
ap = AssetPriceModel()
states = ap.mc.state_values

fig, ax = plt.subplots()

for γ in ys:
    ap.γ = γ
    v = tree_price(ap)
    ax.plot(states, v, lw=2, alpha=0.6, label=rf"$\gamma = {γ}$")

ax.set_title('Price-dividend ratio as a function of the state')
ax.set_ylabel("price-dividend ratio")
ax.set_xlabel("state")
ax.legend(loc='upper right')
plt.show()
```



Notice that v is decreasing in each case.

This is because, with a positively correlated state process, higher states indicate higher future consumption growth.

With the stochastic discount factor (131.13), higher growth decreases the discount factor, lowering the weight placed on future dividends.

Special Cases

In the special case $\gamma = 1$, we have $J = P$.

Recalling that $P^i \mathbb{1} = \mathbb{1}$ for all i and applying *Neumann's geometric series lemma*, we are led to

$$v = \beta(I - \beta P)^{-1} \mathbb{1} = \beta \sum_{i=0}^{\infty} \beta^i P^i \mathbb{1} = \beta \frac{1}{1-\beta} \mathbb{1}$$

Thus, with log preferences, the price-dividend ratio for a Lucas tree is constant.

Alternatively, if $\gamma = 0$, then $J = K$ and we recover the risk-neutral solution (131.10).

This is as expected, since $\gamma = 0$ implies $u(c) = c$ (and hence agents are risk-neutral).

131.4.2 A Risk-Free Consol

Consider the same pure exchange representative agent economy.

A risk-free consol promises to pay a constant amount $\zeta > 0$ each period.

Recycling notation, let p_t now be the price of an ex-coupon claim to the consol.

An ex-coupon claim to the consol entitles an owner at the end of period t to

- ζ in period $t+1$, plus
- the right to sell the claim for p_{t+1} next period

The price satisfies (131.2) with $d_t = \zeta$, or

$$p_t = \mathbb{E}_t \left[m_{t+1} (\zeta + p_{t+1}) \right]$$

With the stochastic discount factor (131.13), this becomes

$$p_t = \mathbb{E}_t \left[\beta g_{t+1}^{-\gamma} (\zeta + p_{t+1}) \right] \tag{131.16}$$

Guessing a solution of the form $p_t = p(X_t)$ and conditioning on $X_t = x$, we get

$$p(x) = \beta \sum_{y \in S} g(y)^{-\gamma} (\zeta + p(y)) P(x, y)$$

Letting $M(x, y) = P(x, y) g(y)^{-\gamma}$ and rewriting in vector notation yields the solution

$$p = (I - \beta M)^{-1} \beta M \zeta \mathbb{1} \tag{131.17}$$

The above is implemented in the function consol_price.

```
def consol_price(ap, ζ):
    """
    Computes price of a consol bond with payoff ζ

    Parameters
    ----------
    ap: AssetPriceModel
        An instance of AssetPriceModel containing primitives

    ζ : scalar(float)
        Coupon of the console

    Returns
    -------
    p : array_like(float)
        Console bond prices

    """
    # Simplify names, set up matrices
    β, γ, P, y = ap.β, ap.γ, ap.mc.P, ap.mc.state_values
    M = P * ap.g(y)**(- γ)

    # Make sure that a unique solution exists
    ap.test_stability(M)

    # Compute price
    I = np.identity(ap.n)
```

(continues on next page)

(continued from previous page)

```
    Ones = np.ones(ap.n)
    p = solve(I - β * M, β * ζ * M @ Ones)

    return p
```

131.4.3 Pricing an Option to Purchase the Consol

Let's now price options of various maturities.

We'll study an option that gives the owner the right to purchase a consol at a price p_S.

An Infinite Horizon Call Option

We want to price an **infinite horizon** option to purchase a consol at a price p_S.

The option entitles the owner at the beginning of a period either

1. to purchase the bond at price p_S now, or
2. not to exercise the option to purchase the asset now but to retain the right to exercise it later

Thus, the owner either *exercises* the option now or chooses *not to exercise* and wait until next period.

This is termed an infinite-horizon **call option** with **strike price** p_S.

The owner of the option is entitled to purchase the consol at price p_S at the beginning of any period, after the coupon has been paid to the previous owner of the bond.

The fundamentals of the economy are identical with the one above, including the stochastic discount factor and the process for consumption.

Let $w(X_t, p_S)$ be the value of the option when the time t growth state is known to be X_t but *before* the owner has decided whether to exercise the option at time t (i.e., today).

Recalling that $p(X_t)$ is the value of the consol when the initial growth state is X_t, the value of the option satisfies

$$w(X_t, p_S) = \max \left\{ \beta \, \mathbb{E}_t \frac{u'(c_{t+1})}{u'(c_t)} w(X_{t+1}, p_S), \; p(X_t) - p_S \right\}$$

The first term on the right is the value of waiting, while the second is the value of exercising now.

We can also write this as

$$w(x, p_S) = \max \left\{ \beta \sum_{y \in S} P(x, y) g(y)^{-\gamma} w(y, p_S), \; p(x) - p_S \right\} \tag{131.18}$$

With $M(x, y) = P(x, y) g(y)^{-\gamma}$ and w as the vector of values $(w(x_i), p_S)_{i=1}^n$, we can express (131.18) as the nonlinear vector equation

$$w = \max\{\beta M w, \; p - p_S \mathbb{1}\} \tag{131.19}$$

To solve (131.19), form an operator T that maps vector w into vector Tw via

$$Tw = \max\{\beta M w, \; p - p_S \mathbb{1}\}$$

Start at some initial w and iterate with T to convergence .

We can find the solution with the following function call_option

```
def call_option(ap, ζ, p_s, ϵ=1e-7):
    """
    Computes price of a call option on a consol bond.

    Parameters
    ----------
    ap: AssetPriceModel
        An instance of AssetPriceModel containing primitives

    ζ : scalar(float)
        Coupon of the console

    p_s : scalar(float)
        Strike price

    ϵ : scalar(float), optional(default=1e-8)
        Tolerance for infinite horizon problem

    Returns
    -------
    w : array_like(float)
        Infinite horizon call option prices

    """
    # Simplify names, set up matrices
    β, γ, P, y = ap.β, ap.γ, ap.mc.P, ap.mc.state_values
    M = P * ap.g(y)**(- γ)

    # Make sure that a unique consol price exists
    ap.test_stability(M)

    # Compute option price
    p = consol_price(ap, ζ)
    w = np.zeros(ap.n)
    error = ϵ + 1
    while error > ϵ:
        # Maximize across columns
        w_new = np.maximum(β * M @ w, p - p_s)
        # Find maximal difference of each component and update
        error = np.amax(np.abs(w - w_new))
        w = w_new

    return w
```

Here's a plot of w compared to the consol price when $P_S = 40$

```
ap = AssetPriceModel(β=0.9)
ζ = 1.0
strike_price = 40

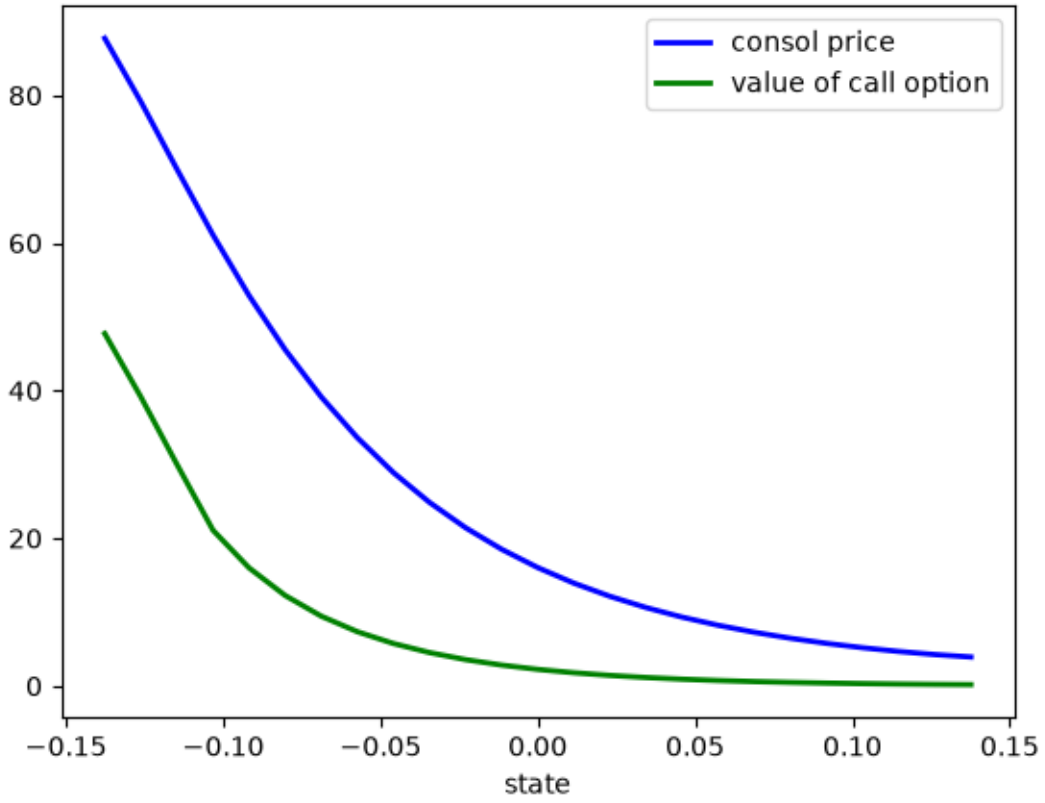
x = ap.mc.state_values
p = consol_price(ap, ζ)
w = call_option(ap, ζ, strike_price)

fig, ax = plt.subplots()
ax.plot(x, p, 'b-', lw=2, label='consol price')
ax.plot(x, w, 'g-', lw=2, label='value of call option')
```

(continues on next page)

(continued from previous page)

```
ax.set_xlabel("state")
ax.legend(loc='upper right')
plt.show()
```



In high values of the Markov growth state, the value of the option is close to zero.

This is despite the facts that the Markov chain is irreducible and that low states — where the consol prices are high — will be visited recurrently.

The reason for low valuations in high Markov growth states is that $\beta = 0.9$, so future payoffs are discounted substantially.

131.4.4 Risk-Free Rates

Let's look at risk-free interest rates over different periods.

The One-period Risk-free Interest Rate

As before, the stochastic discount factor is $m_{t+1} = \beta g_{t+1}^{-\gamma}$.

It follows that the reciprocal R_t^{-1} of the gross risk-free interest rate R_t in state x is

$$\mathbb{E}_t m_{t+1} = \beta \sum_{y \in S} P(x,y) g(y)^{-\gamma}$$

We can write this as

$$m_1 = \beta M \mathbb{1}$$

where the i-th element of m_1 is the reciprocal of the one-period gross risk-free interest rate in state x_i.

Other Terms

Let m_j be an $n \times 1$ vector whose i th component is the reciprocal of the j -period gross risk-free interest rate in state x_i.

Then $m_1 = \beta M$, and $m_{j+1} = M m_j$ for $j \geq 1$.

131.5 Exercises

Exercise 131.5.1

In the lecture, we considered **ex-dividend assets**.

A **cum-dividend** asset is a claim to the stream $d_t, d_{t+1}, \ldots$.

Following (131.1), find the risk-neutral asset pricing equation for one unit of a cum-dividend asset.

With a constant, non-random dividend stream $d_t = d > 0$, what is the equilibrium price of a cum-dividend asset?

With a growing, non-random dividend process $d_t = g d_t$ where $0 < g\beta < 1$, what is the equilibrium price of a cum-dividend asset?

Solution

For a cum-dividend asset, the basic risk-neutral asset pricing equation is

$$p_t = d_t + \beta \mathbb{E}_t[p_{t+1}]$$

With constant dividends, the equilibrium price is

$$p_t = \frac{1}{1-\beta} d_t$$

With a growing, non-random dividend process, the equilibrium price is

$$p_t = \frac{1}{1-\beta g} d_t$$

Exercise 131.5.2

Consider the following primitives

```
n = 5  # Size of State Space
P = np.full((n, n), 0.0125)
P[range(n), range(n)] += 1 - P.sum(1)
# State values of the Markov chain
s = np.array([0.95, 0.975, 1.0, 1.025, 1.05])
γ = 2.0
β = 0.94
```

Let g be defined by $g(x) = x$ (that is, g is the identity map).

Compute the price of the Lucas tree.

Do the same for

- the price of the risk-free consol when $\zeta = 1$
- the call option on the consol when $\zeta = 1$ and $p_S = 150.0$

Solution

First, let's enter the parameters:

```
n = 5
P = np.full((n, n), 0.0125)
P[range(n), range(n)] += 1 - P.sum(1)
s = np.array([0.95, 0.975, 1.0, 1.025, 1.05])  # State values
mc = qe.MarkovChain(P, state_values=s)

γ = 2.0
β = 0.94
ζ = 1.0
p_s = 150.0
```

Next, we'll create an instance of `AssetPriceModel` to feed into the functions

```
apm = AssetPriceModel(β=β, mc=mc, γ=γ, g=lambda x: x)
```

Now we just need to call the relevant functions on the data:

```
tree_price(apm)
```

```
array([29.47401578, 21.93570661, 17.57142236, 14.72515002, 12.72221763])
```

```
consol_price(apm, ζ)
```

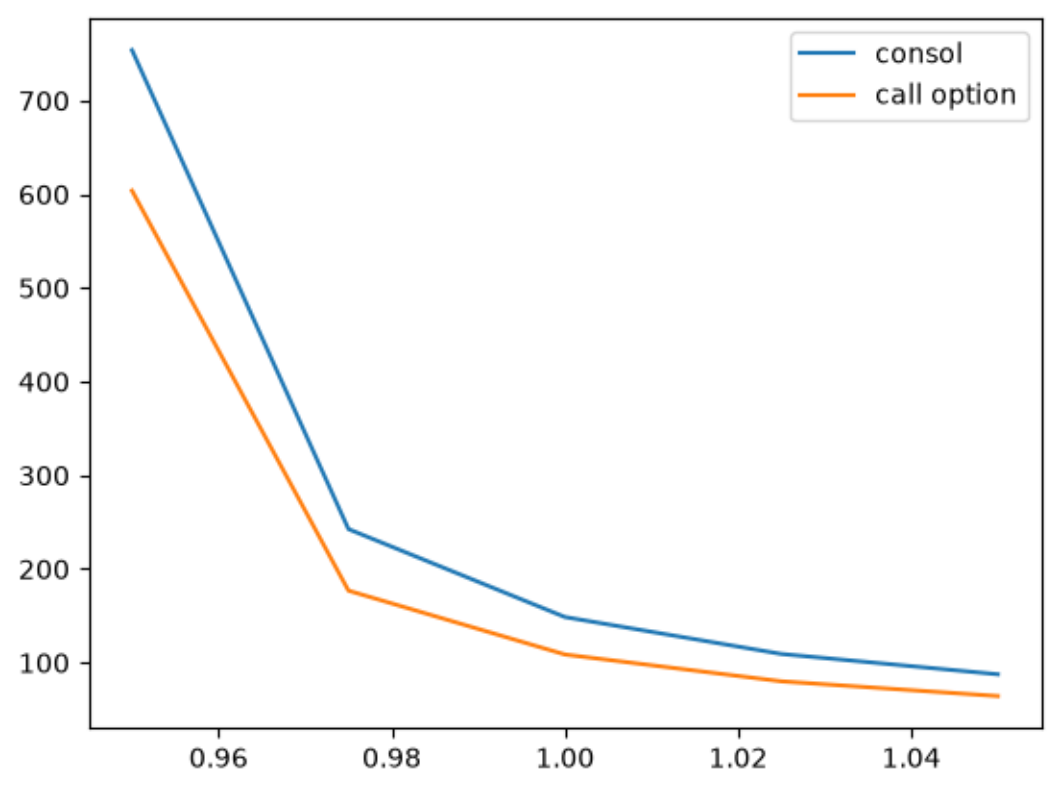
```
array([753.87100476, 242.55144082, 148.67554548, 109.25108965,
       87.56860139])
```

```
call_option(apm, ζ, p_s)
```

```
array([603.87100476, 176.8393343 , 108.67734499,  80.05179254,
       64.30843748])
```

Let's show the last two functions as a plot

```
fig, ax = plt.subplots()
ax.plot(s, consol_price(apm, ζ), label='consol')
ax.plot(s, call_option(apm, ζ, p_s), label='call option')
ax.legend()
plt.show()
```


consol
call option
700
600
500
400
300
200
100
0.96
0.98
1.00
1.02
1.04

Exercise 131.5.3

Let's consider finite horizon call options, which are more common than infinite horizon ones.

Finite horizon options obey functional equations closely related to (131.18).

A k period option expires after k periods.

If we view today as date zero, a k period option gives the owner the right to exercise the option to purchase the risk-free consol at the strike price p_S at dates $0, 1, \ldots, k-1$.

The option expires at time k.

Thus, for $k = 1, 2, \ldots$, let $w(x, k)$ be the value of a k-period option.

It obeys

$$w(x,k) = \max\left\{\beta \sum_{y \in S} P(x,y) g(y)^{-\gamma} w(y, k-1),\ p(x) - p_S\right\}$$

where $w(x, 0) = 0$ for all x.

We can express this as a sequence of nonlinear vector equations

$$w_k = \max\{\beta M w_{k-1},\ p - p_S \mathbb{1}\} \quad k = 1, 2, \ldots \quad \text{with } w_0 = 0$$

Write a function that computes w_k for any given k.

Compute the value of the option with `k = 5` and `k = 25` using parameter values as in Exercise 131.5.1.

Is one higher than the other? Can you give intuition?

Solution

Here's a suitable function:

```
def finite_horizon_call_option(ap, ζ, p_s, k):
    """
    Computes k period option value.
    """
    # Simplify names, set up matrices
    β, γ, P, y = ap.β, ap.γ, ap.mc.P, ap.mc.state_values
    M = P * ap.g(y)**(- γ)

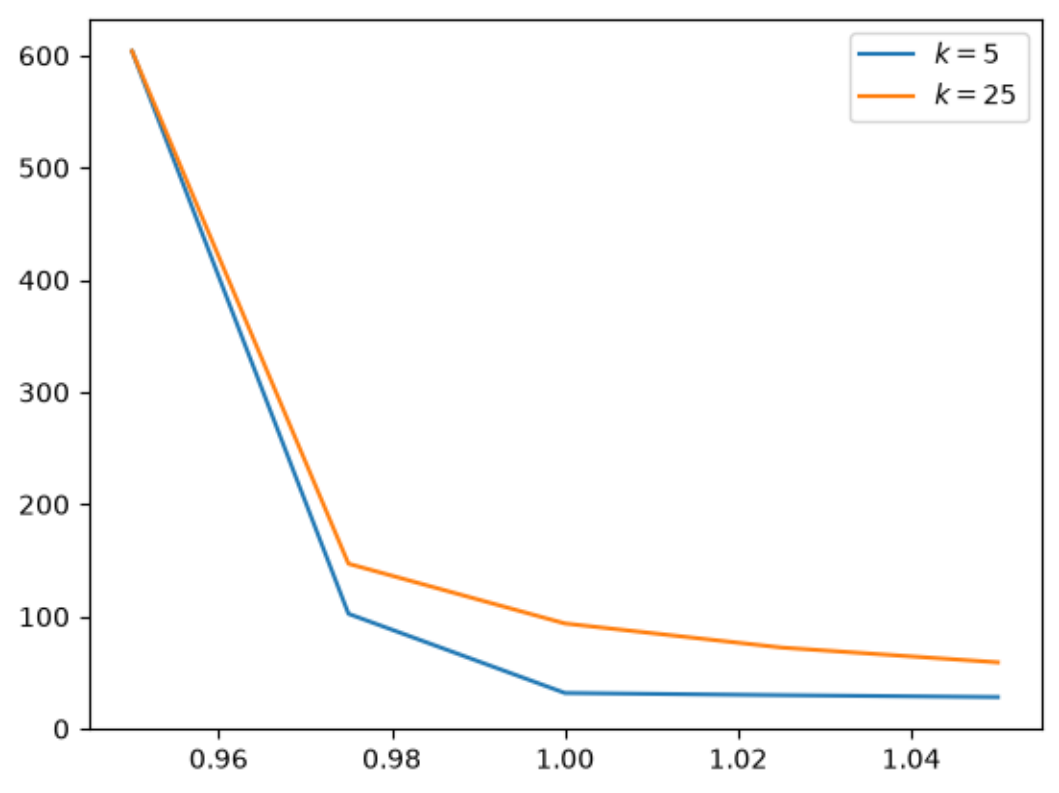
    # Make sure that a unique solution exists
    ap.test_stability(M)


    # Compute option price
    p = consol_price(ap, ζ)
    w = np.zeros(ap.n)
    for i in range(k):
        # Maximize across columns
        w = np.maximum(β * M @ w, p - p_s)

    return w
```

Now let's compute the option values at `k=5` and `k=25`

```
fig, ax = plt.subplots()
for k in [5, 25]:
    w = finite_horizon_call_option(apm, ζ, p_s, k)
    ax.plot(s, w, label=rf'$k = {k}$')
ax.legend()
plt.show()
```



Not surprisingly, options with larger k are worth more.

This is because an owner has a longer horizon over which the option can be exercised.

CHAPTER

TWO

STOCHASTIC CONSUMPTION, RISK AVERSION, AND THE TEMPORAL BEHAVIOR OF ASSET RETURNS

Contents

- *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*
  - *Overview*
  - *Euler equation*
  - *The Euler equation under lognormality*
  - *The restricted system and its likelihood*
  - *Likelihood implementation*
  - *Preference parameters and likelihood ratio tests*
  - *Predictability and the R-squared restriction*
  - *Return-difference tests*
  - *Empirical MLE estimation*
  - *Residual diagnostics*
  - *Connection to the equity premium puzzle*
  - *Another approach*

Evans and Honkapohja: What were the profession's most important responses to the Lucas Critique?

Sargent: There were two. The first and most optimistic response was complete rational expectations econometrics. A rational expectations equilibrium is a likelihood function. Maximize it.

– An Interview with Thomas J. Sargent [Evans and Honkapohja, 2005]

132.1 Overview

This lecture describes how Hansen and Singleton [1983] formulated a complete statistical model of asset returns and consumption growth, then estimated its parameters by the method of maximum likelihood.

They detect a defects in their model, one of which Mehra and Prescott [1985] later called the **equity premium puzzle**.

Hansen and Singleton [1983] study a consumption-based asset pricing model in which a representative consumer with CRRA preferences chooses how to allocate wealth across traded assets.

First-order conditions for asset holdings are stochastic Euler equations that connect consumption growth, asset returns, and preference parameters.

Hansen and Singleton [1983] assume that consumption growth and asset returns are *jointly lognormal*.

Then Euler equations imply a set of restrictions on a the joint distribution of asset prices and returns.

These restrict a linear time-series model for logarithms of consumption growth and returns in which predictable movements in log returns are proportional to predictable movements in log consumption growth.

Hansen and Singleton estimated their model by *Maximum Likelihood Estimation*.

The empirical findings of Hansen and Singleton [1983] constitute what Mehra and Prescott [1985] would later call the **equity premium puzzle**.

To keep lecture this lecture narrowly focused, we estimate one return at a time (either a market proxy or a T-bill) rather than the full multi-asset systems studied by Hansen and Singleton [1983].

- we use only monthly nondurable consumption (`ND`).

```
from itertools import combinations

import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from IPython.display import Latex
from scipy import stats
from scipy.linalg import LinAlgError, cholesky, solve_triangular
from scipy.optimize import minimize
from statsmodels.stats.stattools import durbin_watson
```

We also define a helper to display DataFrames as LaTeX arrays in the hidden cell below

132.2 Euler equation

Consider a single-good economy of identical consumers whose utility functions are in the CRRA form

$$U(c_t) = c_t^\gamma / \gamma, \quad \gamma < 1, \tag{132.1}$$

where c_t is aggregate real per capita consumption and $U(\cdot)$ is the period utility function.

The representative consumer chooses a stochastic consumption plan to maximize the expected value of a time-additive utility function,

$$E_0 \sum_{t=0}^{\infty} \beta^t U(c_t), \quad 0 < \beta < 1. \tag{132.2}$$

Consumers substitute present for future consumption by trading the ownership rights of N financial and capital assets.

Let $\mathbf{w}_t$ denote the holdings of the N assets at date t, $\mathbf{q}_t$ the vector of asset prices, $\mathbf{d}_t$ the vector of dividends, and y_t real labor income.

A feasible consumption and investment plan $\{c_t, \mathbf{w}_t\}$ must satisfy the sequence of budget constraints

$$c_t + \mathbf{q}_t \cdot \mathbf{w}_{t+1} \leq (\mathbf{q}_t + \mathbf{d}_t) \cdot \mathbf{w}_t + y_t, \tag{132.3}$$

where $(\mathbf{q}_t + \mathbf{d}_t) \cdot \mathbf{w}_t$ is the cum-dividend value of the portfolio carried into period t.

To derive the first-order conditions, attach a Lagrange multiplier λ_t to the budget constraint (132.3) at each date t and form the Lagrangian

$$\mathcal{L} = E_0 \sum_{t=0}^{\infty} \beta^t \left\{ U(c_t) + \lambda_t \left[(\mathbf{q}_t + \mathbf{d}_t) \cdot \mathbf{w}_t + y_t - c_t - \mathbf{q}_t \cdot \mathbf{w}_{t+1} \right] \right\}. \tag{132.4}$$

Differentiating $\mathcal{L}$ with respect to c_t gives

$$\frac{\partial \mathcal{L}}{\partial c_t} = 0 \implies \lambda_t = U'(c_t).$$

Differentiating with respect to $w_{i,t+1}$, the holdings of asset i carried from date t into date $t+1$, collects two terms: $w_{i,t+1}$ appears in the date-t budget constraint as well as in the date-$(t+1)$ constraint:

$$\frac{\partial \mathcal{L}}{\partial w_{i,t+1}} = 0 \implies E_0 \left[\beta^t \left(-\lambda_t \, q_{it} + \beta \, \lambda_{t+1} (q_{i,t+1} + d_{i,t+1}) \right) \right] = 0.$$

By the law of iterated expectations, this becomes

$$E_0 \left[\beta^t \, E_t \left(-\lambda_t \, q_{it} + \beta \, \lambda_{t+1} (q_{i,t+1} + d_{i,t+1}) \right) \right] = 0.$$

Hence,

$$\lambda_t \, q_{it} = \beta \, E_t \left[\lambda_{t+1} (q_{i,t+1} + d_{i,t+1}) \right].$$

Substituting $\lambda_t = U'(c_t)$ and $\lambda_{t+1} = U'(c_{t+1})$ yields

$$q_{it} \, U'(c_t) = \beta \, E_t \left[U'(c_{t+1}) (q_{i,t+1} + d_{i,t+1}) \right].$$

Dividing both sides by q_{it} and defining the gross return $r_{it+1} := (q_{i,t+1} + d_{i,t+1})/q_{it}$ yields the stochastic Euler equation:

$$U'(c_t) = \beta E_t [U'(c_{t+1}) \, r_{it+1}], \quad i = 1, \ldots, N, \tag{132.5}$$

where r_{it+1} is the gross real return on asset i.

Substituting the CRRA marginal utility $U'(c_t) = c_t^{\gamma - 1} = c_t^{\alpha}$ with $\alpha := \gamma - 1$ into (132.5) and rearranging gives

$$E_t \left[\beta \left(\frac{c_{t+1}}{c_t} \right)^{\alpha} r_{it+1} \right] = 1, \quad i = 1, \ldots, N. \tag{132.6}$$

The coefficient of relative risk aversion is $-\alpha$.

132.3 The Euler equation under lognormality

Using the Euler equation (132.6) derived above, we now impose the distributional assumptions of Hansen and Singleton [1983].

Let $x_t = c_t / c_{t-1}$ denote the consumption ratio, and define $u_{it} = x_t^{\alpha} r_{it}$ where r_{it} is the gross real return on asset i.

The Euler equation (132.6) states $E_{t-1}[u_{it}] = 1/\beta$.

Define log variables $X_t = \log x_t$, $R_{it} = \log r_{it}$, and $U_{it} = \log u_{it}$, so that

$$U_{i,t} = \alpha X_t + R_{i,t}. \tag{132.7}$$

Assume the vector process $\{Y_t\} = \{(X_t, R_{1t}, \dots, R_{nt})^\top\}$ is jointly stationary and Gaussian.

Under this assumption, U_{it} conditional on the information set ψ_{t-1} is normal with constant variance σ_i^2 and a conditional mean $\mu_{i,t-1}$ that is a linear function of past Y's.

Because $u_{it} = \exp(U_{it})$ is conditionally lognormal,

$$E_{t-1}[u_{it}] = \exp\left(\mu_{i,t-1} + \tfrac{1}{2}\sigma_i^2\right). \tag{132.8}$$

Setting $E_{t-1}[u_{it}] = 1/\beta$ and taking logs gives $\mu_{i,t-1} + \sigma_i^2/2 = -\log\beta$.

Now define the innovation

$$V_{i,t} := U_{i,t} - \mu_{i,t-1} = \alpha X_t + R_{i,t} + \log\beta + \frac{\sigma_i^2}{2}, \quad i = 1, \dots, N. \tag{132.9}$$

Then $E_{t-1}[V_{i,t}] = 0$, where $\sigma_i^2 = \mathrm{Var}_{t-1}(\alpha X_t + R_{it})$ is constant under the stationarity and Gaussian assumptions.

Setting $E_{t-1}[V_{i,t}] = 0$ gives the conditional mean restriction

$$E_{t-1}[R_{i,t}] = -\alpha\, E_{t-1}[X_t] - \log\beta - \frac{\sigma_i^2}{2}. \tag{132.10}$$

Equation (132.10) is the central result of Hansen and Singleton [1983].

The predictable component of each asset's log return is proportional to the predictable component of log consumption growth, with proportionality factor $-\alpha$.

The intercept absorbs the discount factor β and the lognormal variance term $\sigma_i^2/2$.

Let's consider three special cases to better understand the implications of (132.10):

- Risk neutrality ($\alpha = 0$): Each asset's log return equals a constant plus a serially uncorrelated error, so returns are serially uncorrelated.
- Log utility ($\alpha = -1$): The difference $R_{it} - X_t$ has no time-varying predictable component, so returns and consumption growth share the same predictable component.
- Risk aversion ($\alpha < 0$): The time-varying part of $E_{t-1}[R_{it}]$ equals $-\alpha$ times $E_{t-1}[X_t]$, so a larger $|\alpha|$ amplifies the sensitivity of expected returns to expected consumption growth.

For a given consumption–return covariance structure, higher relative risk aversion ($-\alpha$) widens the gap between risky-asset returns and the risk-free return.

The equity premium puzzle arises because the observed spread is large but estimated $|\alpha|$ is small as we will see soon in our estimation.

132.4 The restricted system and its likelihood

To build a likelihood, we need to parameterize the conditional expectation $E_{t-1}[X_t]$.

In the single-return case, write $\mathbf{Y}_t = (X_t, R_t)^\top$ and assume that the predictable component of X_t is a finite-order linear function of past observations:

$$E(X_t \mid \psi_{t-1}) = \mathbf{a}(L)^\top \mathbf{Y}_{t-1} + \mu_x, \tag{132.11}$$

where $\mathbf{a}(L)$ is a vector of lag polynomial coefficients in past (X, R) and μ_x is a constant.

The consumption-growth equation is unrestricted, so X_t depends freely on its own lags and on lagged returns.

The return equation, however, is restricted by the Euler equation.

Combining (132.10) with (132.11) forces the predictable part of R_t to be $-\alpha$ times the predictable part of X_t, plus a constant.

This gives the restricted system

$$\mathbf{A}_0 \mathbf{Y}_t = \mathbf{A}_1(L)\mathbf{Y}_{t-1} + \mu + \mathbf{V}_t, \tag{132.12}$$

with

$$\mathbf{A}_0 = \begin{bmatrix} 1 & 0 \\ \alpha & 1 \end{bmatrix}, \quad \mathbf{A}_1(L) = \begin{bmatrix} \mathbf{a}(L)^\top \\ 0 \end{bmatrix}, \quad \mu = \begin{bmatrix} \mu_x \\ -\log\beta - \sigma_U^2/2 \end{bmatrix}, \tag{132.13}$$

where $\sigma_U^2 := \mathrm{Var}_{t-1}(\alpha X_t + R_t) = \alpha^2\sigma_{XX} + \sigma_{RR} + 2\alpha\sigma_{XR}$ under conditional homoskedasticity.

The sign in the second element of μ follows directly from (132.10).

The innovations $\mathbf{V}_t$ are assumed to be Gaussian with density $f_V(\mathbf{v})$.

Since $\mathbf{V}_t = \mathbf{A}_0\mathbf{Y}_t - \mathbf{A}_1(L)\mathbf{Y}_{t-1} - \mu$, the observables $\mathbf{Y}_t$ are a linear transformation of $\mathbf{V}_t$ with $\mathbf{Y}_t = \mathbf{A}_0^{-1}(\mathbf{V}_t + \mathbf{A}_1(L)\mathbf{Y}_{t-1} + \mu)$.

The change-of-variables formula for densities gives

$$f_Y(\mathbf{y}_t \mid \mathbf{Y}_{t-1}) = f_V(\mathbf{A}_0\mathbf{y}_t - \mathbf{A}_1(L)\mathbf{Y}_{t-1} - \mu) \left|\det\left(\frac{\partial \mathbf{V}_t}{\partial \mathbf{Y}_t}\right)\right| = f_V(\mathbf{v}_t) \,|\det(\mathbf{A}_0)|.$$

Because $\mathbf{A}_0$ is unit lower triangular, $\det(\mathbf{A}_0) = 1$, so the Jacobian term disappears and the log-likelihood of $\mathbf{Y}_t$ equals the log-likelihood of $\mathbf{V}_t$ evaluated at the residuals.

Since the Jacobian is unity, the conditional density of $\mathbf{Y}_t$ is the Gaussian density of $\mathbf{V}_t$ evaluated at $\mathbf{v}_t = \mathbf{A}_0\mathbf{Y}_t - \mathbf{A}_1(L)\mathbf{Y}_{t-1} - \mu$.

For a single observation, $\mathbf{V}_t \sim N(\mathbf{0}, \Sigma)$ has density

$$f_V(\mathbf{v}_t) = (2\pi)^{-1}|\Sigma|^{-1/2} \exp\left(-\tfrac{1}{2}\,\mathbf{v}_t^\top \Sigma^{-1}\mathbf{v}_t\right).$$

Taking logarithms gives

$$\ell_t(\theta) = -\log(2\pi) - \tfrac{1}{2}\log|\Sigma| - \tfrac{1}{2}\,\mathbf{v}_t^\top\Sigma^{-1}\,\mathbf{v}_t.$$

Summing over T observations and dropping the constant $-T\log(2\pi)$ yields the log-likelihood function

$$L(\theta) = -\frac{T}{2}\log|\Sigma| - \frac{1}{2}\sum_{t=1}^{T}(\mathbf{A}_0\mathbf{Y}_t - \mathbf{A}_1(L)\mathbf{Y}_{t-1} - \mu)^\top\Sigma^{-1}(\mathbf{A}_0\mathbf{Y}_t - \mathbf{A}_1(L)\mathbf{Y}_{t-1} - \mu), \tag{132.14}$$

where Σ is the covariance matrix of the innovation $\mathbf{V}_t$ and θ collects all free parameters including α, β, the covariance parameters, the first-row intercept μ_x, and the first-row lag coefficients.

The restrictions imposed by (132.6) enter (132.14) through the structure of $\mathbf{A}_0$, $\mathbf{A}_1(L)$, and μ in (132.13).

The second row of (132.12) has no free lag coefficients and its intercept is determined by α, β, and Σ.

An alternative unrestricted bivariate VAR(p) would estimate both rows of (132.12) freely.

Each row has 1 intercept plus $2p$ lag coefficients (p lags $\times$ 2 variables), giving $2(1 + 2p)$ mean parameters.

Adding 3 free covariance parameters ($\sigma_{XX}, \sigma_{RR}, \sigma_{XR}$) gives a total of $5 + 4p$.

The restricted system (132.12) has only $6 + 2p$ free parameters: the first row contributes $1 + 2p$ (its intercept μ_x and $2p$ lag coefficients), plus α, β, and the 3 covariance parameters.

The second row adds nothing because its lag structure and intercept are pinned down by α, β, and Σ via (132.10).

The difference $(5 + 4p) - (6 + 2p) = 2p - 1$ gives the degrees of freedom for the likelihood ratio tests we will perform soon.

132.5 Likelihood implementation

Now let's implement the likelihood (132.14).

The building blocks are:

- a function to construct lagged data matrices $(\mathbf{Y}_t, \mathbf{Y}_{t-1}, \dots, \mathbf{Y}_{t-p})$,
- a function to map the parameter vector into the matrices $\mathbf{A}_0$, $\mathbf{A}_1$, μ, Σ,
- a function to compute the restricted-system residuals $\mathbf{V}_t = \mathbf{A}_0 \mathbf{Y}_t - \mathbf{A}_1(L)\mathbf{Y}_{t-1} - \mu$, and
- the Gaussian log-likelihood itself.

Since we are working with log-transformed data, we define a helper for the transformation below

```
def to_mle_array(data):
    valid = (data[:, 0] > 0.0) & (data[:, 1] > 0.0)
    return np.column_stack(
        [np.log(data[valid, 1]), np.log(data[valid, 0])])
```

First we build the lagged data matrices, which are the inputs to the likelihood function

```
def build_lagged_data(data, n_lags):
    """
    Build Y_t and lag stacks [Y_{t-1}, ..., Y_{t-p}] for bivariate data.
    """
    if data.ndim != 2 or data.shape[1] != 2:
        raise ValueError("data must be T x 2.")
    if data.shape[0] <= n_lags:
        raise ValueError("Sample size must exceed n_lags.")

    t_obs = data.shape[0]
    n_obs = t_obs - n_lags
    y_t = data[n_lags:, :]
    y_lag = np.empty((n_obs, 2 * n_lags))

    for lag in range(1, n_lags + 1):
        y_lag[:, 2 * (lag - 1) : 2 * lag] = data[n_lags - lag : t_obs - lag, :]

    return y_t, y_lag
```

Next, we validate and unpack the parameter vector while enforcing feasibility conditions

```
def unpack_parameters(params, n_lags):
    """
    Validate and unpack parameter vector.
    """
    if len(params) != 6 + 2 * n_lags:
        return None
```

(continues on next page)

(continued from previous page)

```
    α, β, σ_x, σ_r, cov_xr, μ_x = params[:6]
    a_lags = params[6:]

    tol = 1e-8
    if not np.isfinite(α):
        return None
    if not (tol < β):
        return None
    if not (σ_x > tol and σ_r > tol):
        return None

    Σ = np.array(
        [
            [σ_x ** 2, cov_xr],
            [cov_xr, σ_r ** 2],
        ]
    )

    try:
        cholesky(Σ, lower=True)
    except (LinAlgError, ValueError):
        return None

    return {
        "α": np.array(α),
        "β": np.array(β),
        "σ_x": np.array(σ_x),
        "σ_r": np.array(σ_r),
        "cov_xr": np.array(cov_xr),
        "μ_x": np.array(μ_x),
        "a_lags": a_lags,
    }
```

The next step maps parameters and lagged data into restricted-system residuals

```
def restricted_residuals(
    params,
    y_t,
    y_lag,
    n_lags,
):
    """
    Compute V_t implied by A0 Y_t - A1 Y_{t-1} - mu.
    """
    parsed = unpack_parameters(params, n_lags)
    if parsed is None:
        return None

    α = float(parsed["α"])
    β = float(parsed["β"])
    σ_x = float(parsed["σ_x"])
    σ_r = float(parsed["σ_r"])
    cov_xr = float(parsed["cov_xr"])
    μ_x = float(parsed["μ_x"])
    a_lags = np.asarray(parsed["a_lags"])
```

(continues on next page)

(continued from previous page)

```
    A0 = np.array([[1.0, 0.0], [α, 1.0]])
    A1 = np.zeros((2, 2 * n_lags))
    A1[0, :] = a_lags
    σ_u2 = α ** 2 * σ_x ** 2 + σ_r ** 2 + 2.0 * α * cov_xr
    μ = np.array([μ_x, -np.log(β) - 0.5 * σ_u2])

    resid = y_t @ A0.T - y_lag @ A1.T - μ[None, :]
    if np.any(np.abs(resid) > 1e10):
        return None
    return resid
```

The recursive structure also lets us simulate data from the model.

We can generate data from the model by drawing innovations $\mathbf{V}_t \sim N(\mathbf{0}, \Sigma)$ and compute $\mathbf{Y}_t = \mathbf{A}_0^{-1}(\mathbf{A}_1(L)\mathbf{Y}_{t-1} + \mu + \mathbf{V}_t)$ forward in time.

This is useful for checking the likelihood implementation through Monte Carlo

```
def simulate_restricted_var(
    params,
    n_obs,
    n_lags,
    burn_in=200,
    seed=0,
):
    """
    Simulate [log consumption growth, log return] from the restricted model.
    """
    rng = np.random.default_rng(seed)

    if len(params) != 6 + 2 * n_lags:
        raise ValueError("Parameter vector length must be 6 + 2 * n_lags.")

    α, β, σ_x, σ_r, cov_xr, μ_x = params[:6]
    a_lags = params[6:]

    Σ_e = np.array(
        [
            [σ_x ** 2, cov_xr],
            [cov_xr, σ_r ** 2],
        ]
    )

    A0 = np.array([[1.0, 0.0], [α, 1.0]])
    Σ_v = A0 @ Σ_e @ A0.T

    eigvals = np.linalg.eigvals(Σ_v)
    if np.min(eigvals) <= 0.0:
        Σ_v += np.eye(2) * 1e-6

    A1 = np.zeros((2, 2 * n_lags))
    A1[0, :] = a_lags
    σ_u2 = α ** 2 * σ_x ** 2 + σ_r ** 2 + 2.0 * α * cov_xr
    μ = np.array([μ_x, -np.log(β) - 0.5 * σ_u2])

    total_n = n_obs + burn_in
    y = np.zeros((total_n, 2))
```

(continues on next page)

(continued from previous page)

```
    for t in range(n_lags, total_n):
        lag_stack = []
        for lag in range(1, n_lags + 1):
            lag_stack.append(y[t - lag, :])
        lag_vec = np.concatenate(lag_stack)
        shock = rng.multivariate_normal(np.zeros(2), Σ_v)
        y[t, :] = np.linalg.solve(A0, A1 @ lag_vec + μ + shock)

    return y[burn_in:, :]
```

Next, we encode the Gaussian log-likelihood implied by the residual covariance matrix

```
def log_likelihood_mle(
    params,
    y_t,
    y_lag,
    n_lags,
    include_const=True,
):
    """
    Evaluate Gaussian log-likelihood for the restricted system.
    """
    parsed = unpack_parameters(params, n_lags)
    if parsed is None:
        return -np.inf

    resid = restricted_residuals(params, y_t, y_lag, n_lags)
    if resid is None:
        return -np.inf

    α = float(parsed["α"])
    σ_x = float(parsed["σ_x"])
    σ_r = float(parsed["σ_r"])
    cov_xr = float(parsed["cov_xr"])

    Σ_e = np.array(
        [
            [σ_x ** 2, cov_xr],
            [cov_xr, σ_r ** 2],
        ]
    )

    A0 = np.array([[1.0, 0.0], [α, 1.0]])
    Σ_v = A0 @ Σ_e @ A0.T

    try:
        chol = cholesky(Σ_v, lower=True)
        log_det = 2.0 * np.sum(np.log(np.diag(chol) + 1e-16))
        std_resid = solve_triangular(chol, resid.T, lower=True).T
        quad_form = np.sum(std_resid ** 2)
    except (LinAlgError, ValueError):
        return -np.inf

    sample_size = y_t.shape[0]
    ll = -0.5 * sample_size * log_det - 0.5 * quad_form
    if include_const:
```

(continues on next page)

(continued from previous page)

```
        ll -= sample_size * np.log(2.0 * np.pi)

    if np.isnan(ll) or np.isinf(ll):
        return -np.inf
    return float(ll)
```

To pass this into a numerical optimizer, we wrap the log-likelihood as a minimization objective

```
def negative_log_likelihood(
    params,
    y_t,
    y_lag,
    n_lags,
):
    """
    Return negative log-likelihood for minimization.
    """
    ll = log_likelihood_mle(params, y_t, y_lag, n_lags, include_const=False)
    if np.isfinite(ll):
        return -ll
    return 1e20
```

We set parameter bounds and generate data-driven starting values for multi-start optimization.

We keep β positive but do not force $\beta < 1$ in estimation

```
def parameter_bounds(n_lags):
    """
    Bounds for optimization.
    """
    bounds = [
        (-200.0, 200.0),    # α (= -risk aversion)
        (1e-8, 2.0),      # β (discount factor)
        (1e-8, None),     # σ_x (std dev of consumption innovation)
        (1e-8, None),     # σ_r (std dev of return innovation)
        (None, None),     # cov_xr (covariance)
        (None, None),     # μ_x (consumption growth intercept)
    ]
    bounds += [(-0.99, 0.99)] * (2 * n_lags)  # VAR lag coefficients
    return bounds
```

We use multiple starting vectors to help local solvers escape poor initializations

```
def starting_values(y_t, y_lag, n_lags, n_starts=10):
    """
    Generate multiple starting values.
    """
    rng = np.random.default_rng(123)
    starts = []
    n_params = 6 + 2 * n_lags

    base = np.zeros(n_params)
    base[0] = -0.5
    base[1] = 0.999
    base[2] = max(float(np.std(y_t[:, 0])), 1e-3)
    base[3] = max(float(np.std(y_t[:, 1])), 1e-3)
    base[4] = float(np.cov(y_t.T)[0, 1])
```

(continues on next page)

(continued from previous page)

```
    base[5] = float(np.mean(y_t[:, 0]))
    base[6:] = 0.1
    starts.append(base.copy())

    # OLS seed from unrestricted VAR
    n_obs = y_t.shape[0]
    x = np.column_stack([np.ones(n_obs), y_lag])
    coef = np.linalg.lstsq(x, y_t, rcond=None)[0]
    resid = y_t - x @ coef
    Σ_e = resid.T @ resid / max(1, n_obs)

    a_lags_ols = coef[1:, 0]
    r_lags_ols = coef[1:, 1]
    denom = float(a_lags_ols @ a_lags_ols)
    if denom > 1e-10:
        α_ols = -float((a_lags_ols @ r_lags_ols) / denom)
    else:
        α_ols = -0.5

    μ_x_ols = float(coef[0, 0])
    μ_r_ols = float(coef[0, 1])
    σ_x_ols = float(np.sqrt(max(Σ_e[0, 0], 1e-8)))
    σ_r_ols = float(np.sqrt(max(Σ_e[1, 1], 1e-8)))
    cov_xr_ols = float(Σ_e[0, 1])
    σ_u2_ols = (
        α_ols ** 2 * σ_x_ols ** 2
        + σ_r_ols ** 2
        + 2.0 * α_ols * cov_xr_ols
    )
    β_ols = float(np.exp(-(μ_r_ols + α_ols * μ_x_ols + 0.5 * σ_u2_ols)))
    β_ols = float(np.clip(β_ols, 1e-6, 2.0))

    ols_seed = np.zeros(n_params)
    ols_seed[0] = α_ols
    ols_seed[1] = β_ols
    ols_seed[2] = σ_x_ols
    ols_seed[3] = σ_r_ols
    ols_seed[4] = cov_xr_ols
    ols_seed[5] = μ_x_ols
    ols_seed[6:] = a_lags_ols
    starts.append(ols_seed.copy())

    seeds = [base, ols_seed]
    while len(starts) < n_starts:
        seed = seeds[len(starts) % len(seeds)]
        trial = seed.copy()
        trial[:2] += rng.normal(0.0, 0.2, 2)
        trial[2:6] *= 1.0 + rng.normal(0.0, 0.15, 4)
        trial[6:] += rng.normal(0.0, 0.08, 2 * n_lags)
        trial[1] = max(trial[1], 1e-6)
        trial[2] = max(trial[2], 1e-6)
        trial[3] = max(trial[3], 1e-6)
        starts.append(trial)

    return starts
```

Standard errors come from an outer-product-of-gradients (OPG) approximation to the information matrix, computed by

finite differences of per-observation log-likelihood contributions.

This tends to be more numerically stable than finite-difference Hessians in this application.

```
def log_likelihood_contributions(
    params,
    y_t,
    y_lag,
    n_lags,
    include_const=False,
):
    """
    Vector of per-observation Gaussian log-likelihood contributions.

    Returns an array of length T, or None if the parameter vector is infeasible.
    """
    parsed = unpack_parameters(params, n_lags)
    if parsed is None:
        return None

    resid = restricted_residuals(params, y_t, y_lag, n_lags)
    if resid is None:
        return None

    α = float(parsed["α"])
    σ_x = float(parsed["σ_x"])
    σ_r = float(parsed["σ_r"])
    cov_xr = float(parsed["cov_xr"])

    Σ_e = np.array(
        [
            [σ_x ** 2, cov_xr],
            [cov_xr, σ_r ** 2],
        ]
    )
    A0 = np.array([[1.0, 0.0], [α, 1.0]])
    Σ_v = A0 @ Σ_e @ A0.T

    try:
        chol = cholesky(Σ_v, lower=True)
        log_det = 2.0 * np.sum(np.log(np.diag(chol) + 1e-16))
        std_resid = solve_triangular(chol, resid.T, lower=True).T
    except (LinAlgError, ValueError):
        return None

    quad = np.sum(std_resid ** 2, axis=1)
    ll_t = -0.5 * log_det - 0.5 * quad
    if include_const:
        ll_t -= np.log(2.0 * np.pi)
    if not np.all(np.isfinite(ll_t)):
        return None
    return ll_t


def opg_standard_errors(
    params,
    y_t,
    y_lag,
```

(continues on next page)

(continued from previous page)

```
    n_lags,
    step=1e-6,
    max_step_shrink=12,
    eig_floor=1e-12,
):
    """
    Standard errors via OPG approximation to the information matrix.
    """
    n = len(params)
    ll0 = log_likelihood_contributions(params, y_t, y_lag, n_lags, include_
↪const=False)
    if ll0 is None:
        return np.full(n, np.nan)

    n_obs = int(ll0.shape[0])
    scores = np.empty((n_obs, n))

    for i in range(n):
        base_step = step * (abs(params[i]) + 1.0)
        hi = base_step
        ll_plus = None
        ll_minus = None

        for _ in range(max_step_shrink + 1):
            p_plus = params.copy()
            p_minus = params.copy()
            p_plus[i] += hi
            p_minus[i] -= hi
            ll_plus = log_likelihood_contributions(
                p_plus, y_t, y_lag, n_lags, include_const=False
            )
            ll_minus = log_likelihood_contributions(
                p_minus, y_t, y_lag, n_lags, include_const=False
            )
            if ll_plus is not None and ll_minus is not None:
                break
            hi *= 0.5

        if ll_plus is None or ll_minus is None:
            return np.full(n, np.nan)

        scores[:, i] = (ll_plus - ll_minus) / (2.0 * hi)

    if not np.all(np.isfinite(scores)):
        return np.full(n, np.nan)

    # Center scores to mitigate numerical drift away
    scores = scores - scores.mean(axis=0, keepdims=True)

    opg = scores.T @ scores
    if not np.all(np.isfinite(opg)):
        return np.full(n, np.nan)
    opg = 0.5 * (opg + opg.T)

    try:
        eigvals, eigvecs = np.linalg.eigh(opg)
    except (LinAlgError, ValueError):
```

(continues on next page)

(continued from previous page)

```
        return np.full(n, np.nan)

    floor = float(eig_floor) * max(1.0, float(np.max(eigvals)))
    eigvals = np.clip(eigvals, floor, None)
    cov = eigvecs @ np.diag(1.0 / eigvals) @ eigvecs.T
    se = np.sqrt(np.maximum(np.diag(cov), 0.0))
    se[~np.isfinite(se)] = np.nan
    return se
```

The multi-start MLE estimator below combines these pieces and returns parameters, fit criteria, and residuals

```
def estimate_mle(data, n_lags, verbose=False):
    """
    Estimate the restricted model by multi-start local optimization.
    """
    y_t, y_lag = build_lagged_data(data, n_lags)
    bounds = parameter_bounds(n_lags)
    starts = starting_values(y_t, y_lag, n_lags)

    best_result = None
    best_ll = -np.inf

    for i, x0 in enumerate(starts):
        try:
            result = minimize(
                negative_log_likelihood,
                x0=x0,
                args=(y_t, y_lag, n_lags),
                method="L-BFGS-B",
                bounds=bounds,
            )
        except Exception:
            continue

        if np.isfinite(result.fun):
            ll_val = log_likelihood_mle(result.x, y_t, y_lag, n_lags)
            if np.isfinite(ll_val) and ll_val > best_ll:
                best_ll = ll_val
                best_result = result
                if verbose:
                    print(
                        f"start={i}, success={result.success}, "
                        f"status={result.status}, loglike={ll_val:.2f}"
                    )

    n_params = 6 + 2 * n_lags

    if best_result is None:
        return {
            "params": np.full(n_params, np.nan),
            "se": np.full(n_params, np.nan),
            "loglike": -np.inf,
            "converged": False,
            "optimizer_success": False,
            "residuals": None,
            "n_obs": int(y_t.shape[0]),
        }
```

(continues on next page)

(continued from previous page)

```
    params = best_result.x
    se = opg_standard_errors(params, y_t, y_lag, n_lags)
    resid = restricted_residuals(params, y_t, y_lag, n_lags)
    ll_val = log_likelihood_mle(params, y_t, y_lag, n_lags)

    return {
        "params": params,
        "se": se,
        "loglike": ll_val,
        "converged": bool(np.isfinite(ll_val)),
        "optimizer_success": bool(best_result.success),
        "residuals": resid,
        "n_obs": int(y_t.shape[0]),
    }
```

Residual diagnostics below summarize normality and serial-correlation checks

```
def residual_diagnostics(resid):
    """
    Compute basic residual diagnostics.
    """
    out = {}

    for i, label in enumerate(["consumption", "return"]):
        jb_stat, jb_pval = stats.jarque_bera(resid[:, i])
        out[f"{label}_jb_stat"] = float(jb_stat)
        out[f"{label}_jb_pval"] = float(jb_pval)
        out[f"{label}_dw"] = float(durbin_watson(resid[:, i]))

    return out
```

Finally, a lag-loop wrapper runs MLE across different lag lengths and collects results in a DataFrame

```
def run_mle_by_lag(
    data,
    lags=(2, 4, 6),
    verbose=False,
):
    """
    Estimate restricted MLE models by lag length.
    """
    rows = []
    fits = {}

    for lag in lags:
        fit = estimate_mle(data, n_lags=lag, verbose=verbose)
        fits[lag] = fit

        rows.append(
            {
                "NLAG": lag,
                "α_hat": fit["params"][0],
                "se_α": fit["se"][0],
                "β_hat": fit["params"][1],
                "se_β": fit["se"][1],
                "loglike": fit["loglike"],
```

(continues on next page)

(continued from previous page)

```
                "n_obs": fit["n_obs"],
            }
        )

    table = pd.DataFrame(rows).set_index("NLAG")
    return table, fits
```

132.5.1 Simulation

Before applying the likelihood to real data, we verify that it recovers known parameters from simulated data generated by the restricted system.

For `n_lags = 1`, the parameter vector is

$$\theta = (\alpha,\ \beta,\ \sigma_x,\ \sigma_r,\ \sigma_{xr},\ \mu_x,\ a_{x,1},\ a_{r,1}),$$

where the last two entries are the coefficients on X_{t-1} and R_{t-1} in the first-row regression for X_t.

More generally, for `n_lags = p` we pack `a_lags` in the order $[a_{x,1}, a_{r,1}, \dots, a_{x,p}, a_{r,p}]$, so the full parameter vector has length $6 + 2p$.

The covariance parameters $(\sigma_x, \sigma_r, \sigma_{xr})$ describe the reduced-form shocks to (X_t, R_t): if $\varepsilon_t = (\varepsilon_{x,t}, \varepsilon_{r,t})^\top \sim N(0, \Sigma_\varepsilon)$, then $\Sigma_\varepsilon = \begin{bmatrix} \sigma_x^2 & \sigma_{xr} \\ \sigma_{xr} & \sigma_r^2 \end{bmatrix}$.

The restricted-system innovation is $\mathbf{V}_t = \mathbf{A}_0 \varepsilon_t$, so its covariance is $\Sigma_V = \mathbf{A}_0 \Sigma_\varepsilon \mathbf{A}_0^\top$, which is what enters the likelihood.

In the simulation recursion we draw $\mathbf{V}_t \sim N(0, \Sigma_V)$ and compute $\mathbf{Y}_t = \mathbf{A}_0^{-1}(\mathbf{A}_1(L)\mathbf{Y}_{t-1} + \mu + \mathbf{V}_t)$ forward.

The following table compares the true parameters to the MLE estimates from a large simulated sample

```
α_true = -1.00
β_true = 0.993
σ_x_true = 0.015
σ_r_true = 0.020
cov_xr_true = 0.0001
μ_x_true = 0.002
a_x1_true = 0.40
a_r1_true = 0.10

true_params = np.array(
    [
        α_true,
        β_true,
        σ_x_true,
        σ_r_true,
        cov_xr_true,
        μ_x_true,
        a_x1_true,
        a_r1_true,
    ]
)

sim_mle_data = simulate_restricted_var(
    params=true_params,
    n_obs=50000,
```

(continues on next page)

(continued from previous page)

```
    n_lags=1,
    burn_in=5000,
    seed=0,
)

fit_sim = estimate_mle(sim_mle_data, n_lags=1, verbose=False)
```

	true	estimate	se	t (H_0: true)
α	-1.0000	-0.9757	0.011291	2.15
β	0.9930	0.9929	0.000107	-0.78

The point estimates are close to the true parameters, and the t-statistics for the null hypothesis that the true value is correct are small in magnitude, consistent with sampling variation.

132.6 Preference parameters and likelihood ratio tests

The restricted system embeds preference parameters through cross-equation restrictions.

The parameter α links predictable variation in returns to predictable variation in consumption growth.

Under the model, the return equation's dependence on lagged variables is entirely determined by $-\alpha$ times the consumption equation's lag coefficients.

The parameter β shifts the return-equation intercept through $-\log\beta - \sigma_U^2/2$.

Hansen and Singleton [1983] test these restrictions by comparing the restricted system to an unrestricted bivariate VAR estimated on the same sample.

If the restrictions imposed by the Euler-equation are correct, the restricted model should fit nearly as well as the unrestricted model.

The standard test is the likelihood ratio statistic

$$LR = 2(\ell_u - \ell_r) \Rightarrow \chi^2_d, \tag{132.15}$$

where ℓ_u and ℓ_r are the maximized log-likelihoods of the unrestricted and restricted models, and d is the difference in the number of free parameters.

Both likelihoods must be evaluated on the same effective sample for the LR distribution to be valid.

Hansen and Singleton [1983] report that for the value-weighted aggregate stock return, the χ^2 test statistics provide little evidence against the model.

In the tables below we report both `chi2.cdf(LR, df)` (corresponding to what HS report in parentheses) and the usual right-tail `p(LR) = 1 - chi2.cdf`

```
def likelihood_ratio_test(
    fit_restricted,
    fit_unrestricted,
    df_diff,
):
    """
    Compare nested specifications with an LR test.
    """
    if not (fit_restricted["converged"] and fit_unrestricted["converged"]):
```

(continues on next page)

(continued from previous page)

```
        return {"lr_stat": np.nan, "p_value": np.nan, "chi2_cdf": np.nan}
    if fit_restricted.get("n_obs") != fit_unrestricted.get("n_obs"):
        return {"lr_stat": np.nan, "p_value": np.nan, "chi2_cdf": np.nan}

    lr_stat = 2.0 * (fit_unrestricted["loglike"] - fit_restricted["loglike"])
    chi2_cdf = stats.chi2.cdf(lr_stat, df=df_diff)
    p_value = 1.0 - chi2_cdf
    return {
        "lr_stat": float(lr_stat),
        "p_value": float(p_value),
        "chi2_cdf": float(chi2_cdf),
    }
```

The unrestricted benchmark is a Gaussian VAR for $\mathbf{Y}_t = (X_t, R_t)^\top$ with free coefficients on all lags in both equations.

```
def estimate_unrestricted_var(data, n_lags):
    """
    Estimate an unrestricted Gaussian VAR for Y_t = [X_t, R_t].
    """
    y_t, y_lag = build_lagged_data(data, n_lags)
    n_obs = y_t.shape[0]
    x = np.column_stack([np.ones(n_obs), y_lag])
    coef = np.linalg.lstsq(x, y_t, rcond=None)[0]
    resid = y_t - x @ coef
    Σ = resid.T @ resid / n_obs

    try:
        chol = cholesky(Σ, lower=True)
        log_det = 2.0 * np.sum(np.log(np.diag(chol) + 1e-16))
        std_resid = solve_triangular(chol, resid.T, lower=True).T
        quad_form = np.sum(std_resid ** 2)
    except (LinAlgError, ValueError):
        return {
            "coef": coef,
            "Σ": np.full((2, 2), np.nan),
            "residuals": resid,
            "loglike": -np.inf,
            "converged": False,
            "n_obs": int(n_obs),
        }

    d = y_t.shape[1]
    loglike = float(-0.5 * n_obs * d * np.log(2.0 * np.pi)
            - 0.5 * n_obs * log_det - 0.5 * quad_form)

    return {
        "coef": coef,
        "Σ": Σ,
        "residuals": resid,
        "loglike": loglike,
        "converged": True,
        "n_obs": int(n_obs),
    }
```

```
def run_unrestricted_var_by_lag(data, lags=(2, 4, 6)):
    """
```

(continues on next page)

(continued from previous page)

```
    Estimate unrestricted VAR models by lag length.
    """
    rows = []
    fits = {}

    for lag in lags:
        fit = estimate_unrestricted_var(data, n_lags=lag)
        fits[lag] = fit
        rows.append(
            {
                "NLAG": lag,
                "loglike": fit["loglike"],
                "n_obs": fit["n_obs"],
            }
        )

    table = pd.DataFrame(rows).set_index("NLAG")
    return table, fits
```

The following function computes the LR statistic at each lag length, replicating the testing strategy of Table 1 in Hansen and Singleton [1983].

```
def restricted_vs_unrestricted_lr(mle_fits, unrestricted_fits, lags=(2, 4, 6)):
    """
    Compute LR tests of restricted model versus unrestricted VAR.
    """
    rows = []

    for lag in lags:
        fit_r = mle_fits[lag]
        fit_u = unrestricted_fits[lag]
        df_diff = (2 * (1 + 2 * lag) + 3) - (6 + 2 * lag)
        lr = likelihood_ratio_test(fit_restricted=fit_r, fit_unrestricted=fit_u, df_
↪diff=df_diff)
        rows.append(
            {
                "NLAG": lag,
                "lr_stat": lr["lr_stat"],
                "p_value": lr["p_value"],
                "chi2_cdf": lr["chi2_cdf"],
                "df": df_diff,
                "T": fit_r.get("n_obs", np.nan),
            }
        )

    return pd.DataFrame(rows).set_index("NLAG")
```

132.7 Predictability and the R-squared restriction

Section II of Hansen and Singleton [1983] emphasizes an implication of the restricted system for return predictability.

From (132.10) and (132.11), the predictable component of the log return is

$$E(R_t \mid \psi_{t-1}) = -\alpha\, E(X_t \mid \psi_{t-1}) - \log\beta - \frac{\sigma_U^2}{2}. \tag{132.16}$$

Since the predictable return is a linear function of the predictable consumption growth, the variance of the predictable return component is exactly α^2 times the variance of the predictable consumption-growth component:

$$\operatorname{Var}[E(R_t \mid \psi_{t-1})] = \alpha^2 \operatorname{Var}[E(X_t \mid \psi_{t-1})]. \tag{132.17}$$

Hansen and Singleton [1983] derive the implied R^2 of the return projection onto ψ_{t-1}:

$$R_R^2 = \frac{\alpha^2 \operatorname{Var}[E(X_t \mid \psi_{t-1})]}{\operatorname{Var}(R_t \mid \psi_{t-1}) + \alpha^2 \operatorname{Var}[E(X_t \mid \psi_{t-1})]}. \tag{132.18}$$

If $\alpha = 0$ (risk neutrality), then $R_R^2 = 0$ and asset returns are unpredictable.

If $\alpha = -1$ (log utility), then $R_t - X_t$ is unpredictable, so returns and consumption growth share the same predictable component.

More generally, the R_R^2 for returns will be small whenever the variance of the unpredictable return component $\operatorname{Var}(R_t \mid \psi_{t-1})$ is large relative to the predictable variance, which is the case for stock returns.

The function below reports:

- restriction-side predictable-variance terms implied by the Euler equation, and
- R_X^2 and R_R^2 from the unrestricted VAR.

```
def predictability_metrics(data, restricted_fit, unrestricted_fit, n_lags):
    """
    Compute predictable-component metrics and unrestricted VAR R^2 values.
    """
    y_t, y_lag = build_lagged_data(data, n_lags)
    parsed = unpack_parameters(restricted_fit["params"], n_lags)
    α = float(parsed["α"])
    β = float(parsed["β"])
    σ_x = float(parsed["σ_x"])
    σ_r = float(parsed["σ_r"])
    cov_xr = float(parsed["cov_xr"])
    μ_x = float(parsed["μ_x"])
    a_lags = np.asarray(parsed["a_lags"])

    pred_x = y_lag @ a_lags + μ_x
    σ_u2 = α ** 2 * σ_x ** 2 + σ_r ** 2 + 2.0 * α * cov_xr
    pred_r = -α * pred_x - np.log(β) - 0.5 * σ_u2

    x = y_t[:, 0]
    r = y_t[:, 1]
    resid_x = x - pred_x
    resid_r = r - pred_r

    r2_x = 1.0 - np.var(resid_x) / np.var(x)
    r2_r = 1.0 - np.var(resid_r) / np.var(r)
```

(continues on next page)

(continued from previous page)

```
    if unrestricted_fit["converged"] and unrestricted_fit.get("coef") is not None:
        n_obs = y_t.shape[0]
        x_u = np.column_stack([np.ones(n_obs), y_lag])
        pred_u = x_u @ unrestricted_fit["coef"]
        resid_u = y_t - pred_u
        var_x = np.var(y_t[:, 0])
        var_r = np.var(y_t[:, 1])
        r2_x_unres = np.nan if var_x <= 0.0 else 1.0 - np.var(resid_u[:, 0]) / var_x
        r2_r_unres = np.nan if var_r <= 0.0 else 1.0 - np.var(resid_u[:, 1]) / var_r
    else:
        r2_x_unres = np.nan
        r2_r_unres = np.nan

    return {
        "alpha_hat": α,
        "var_pred_x": float(np.var(pred_x)),
        "var_pred_r": float(np.var(pred_r)),
        "alpha2_var_pred_x": float(α ** 2 * np.var(pred_x)),
        "r2_x_restricted": float(r2_x),
        "r2_r_restricted": float(r2_r),
        "r2_x_unrestricted": float(r2_x_unres),
        "r2_r_unrestricted": float(r2_r_unres),
        "T": int(y_t.shape[0]),
    }
```

132.8 Return-difference tests

Hansen and Singleton [1983] also propose tests based on differences in log returns across assets.

From (132.10), the conditional mean of asset i's log return is $E_{t-1}[R_{it}] = -\alpha\, E_{t-1}[X_t] - \log\beta - \sigma_i^2/2$.

The term $-\alpha\, E_{t-1}[X_t] - \log\beta$ is common to all assets, so it cancels in the difference:

$$E_{t-1}[R_{it} - R_{jt}] = \frac{\sigma_j^2 - \sigma_i^2}{2},$$

which is a constant that does not depend on time-$(t-1)$ information.

Return differences should therefore be unpredictable if the model is correct, regardless of the values of α and β.

These tests avoid the need to measure consumption, at the cost of losing the ability to identify α and β.

Hansen and Singleton [1983] report that the return-difference restrictions are strongly rejected for models with multiple stock returns, providing substantial evidence against the CRRA-lognormal specification even when consumption measurement problems are eliminated.

The code below is an illustration of this logic on simulated data.

Reproducing the paper's empirical return-difference tables requires estimating multi-asset systems that are outside this lecture's scope

```
def simulate_multi_asset_nominal_returns(
    n_obs,
    n_assets=3,
    α_true=-1.0,
    β_true=0.993,
```

(continues on next page)

(continued from previous page)

```
    seed=0,
):
    """
    Simulate log nominal returns satisfying E_t[beta * exp(alpha X) * r_i] = 1.
    """
    if n_assets < 2:
        raise ValueError("n_assets must be at least 2.")

    rng = np.random.default_rng(seed)
    x = np.empty(n_obs)
    x[0] = 0.001
    for t in range(1, n_obs):
        x[t] = 0.001 + 0.4 * (x[t - 1] - 0.001) + 0.006 * rng.standard_normal()

    sigmas = np.linspace(0.03, 0.06, n_assets)
    eps = rng.standard_normal((n_obs, n_assets)) * sigmas
    log_returns = -np.log(β_true) - α_true * x[:, None] + eps \
            - 0.5 * sigmas[None, :] ** 2
    return x, log_returns


def return_difference_test(log_returns, n_lags=2):
    """
    Test predictability of pairwise log-return differences.
    """
    if log_returns.ndim != 2 or log_returns.shape[1] < 2:
        raise ValueError("log_returns must be T x m with m >= 2.")
    if log_returns.shape[0] <= n_lags + 1:
        raise ValueError("Sample size must exceed n_lags + 1.")

    t_obs, n_assets = log_returns.shape
    pairs = list(combinations(range(n_assets), 2))
    n_obs = t_obs - n_lags - 1
    z = np.empty((n_obs, 1 + n_assets * n_lags))
    z[:, 0] = 1.0

    for j in range(n_lags):
        z[:, 1 + j * n_assets : 1 + (j + 1) * n_assets] = log_returns[
            n_lags - j : t_obs - 1 - j, :
        ]

    rows = []
    for i, j in pairs:
        y = log_returns[n_lags + 1 :, i] - log_returns[n_lags + 1 :, j]
        coef = np.linalg.lstsq(z, y, rcond=None)[0]
        resid = y - z @ coef
        sigma2 = float((resid @ resid) / max(1, n_obs - z.shape[1]))
        cov = sigma2 * np.linalg.pinv(z.T @ z)
        slopes = coef[1:]
        cov_slopes = cov[1:, 1:]
        stat = float(slopes @ np.linalg.pinv(cov_slopes) @ slopes)
        p_value = float(1.0 - stats.chi2.cdf(stat, df=slopes.shape[0]))
        rows.append(
            {
                "pair": f"{i+1}-{j+1}",
                "wald_chi2": stat,
                "p_value": p_value,
```

(continues on next page)

(continued from previous page)

```
                "mean_spread": float(np.mean(y)),
            }
        )

    return pd.DataFrame(rows).set_index("pair")
```

We run `return_difference_test` on simulated data with $m = 3$ assets, giving $\binom{3}{2} = 3$ pairs.

For each pair, the function regresses the return spread on a constant and `n_lags` lags of all asset returns, then tests whether the slope coefficients are jointly zero using a Wald χ^2 statistic

```
_, sim_log_returns = simulate_multi_asset_nominal_returns(
    n_obs=1500,
    n_assets=3,
    α_true=-1.0,
    β_true=0.993,
    seed=0,
)

spread_test = return_difference_test(sim_log_returns, n_lags=2)
spread_pretty = spread_test.rename(columns={
    "wald_chi2": r"\chi^2", "p_value": "p", "mean_spread": r"\overline{\Delta R}",
})
display_table(spread_pretty, fmt={
    r"\chi^2": "{:.3f}",
    "p": "{:.3f}",
    r"\overline{\Delta R}": "{:.5f}",
})
```

pair	χ^2	p	$\overline{\Delta R}$
$1-2$	9.445	0.150	0.00132
$1-3$	8.338	0.214	0.00154
$2-3$	4.272	0.640	0.00023

Large p-values confirm that return differences are unpredictable in this simulation, exactly as the model predicts when $\alpha = -1$.

132.9 Empirical MLE estimation

We now apply the maximum likelihood estimator of Hansen and Singleton [1983] to real data.

132.9.1 Data

Both this lecture and the companion lecture *Estimating Euler Equations by Generalized Method of Moments* use the same data construction.

Hansen and Singleton [1982] and Hansen and Singleton [1983] use monthly data on real per capita consumption (nondurables) and stock returns from CRSP for the period 1959:2 through 1978:12.

To align with the paper, we set the default sample to 1959:2–1978:12.

You can pass different `start` and `end` dates to study later periods.

This lecture pulls stock-market and one-month bill returns from the Ken French data library (`F-F_Research_Data_Factors`) and constructs gross nominal returns as `1 + (Mkt-RF + RF)/100` for the market and `1 + RF/100` for T-bills.

`Mkt-RF` is the value-weighted return on all CRSP firms minus the risk-free rate, and `RF` is the one-month Treasury bill return (see here for details).

While Hansen-Singleton use CRSP value-weighted NYSE returns, we use the Ken French market factor as the closest open-access alternative.

The consumption series is constructed from consumption of nondurables (`ND`) with the nondurables deflator.

The hidden cell below loads a monthly dataset of returns and consumption series from QuantEcon/data-lectures. The data are built from the FRED and Ken French data libraries by the maintenance script at `builders/hansen_singleton_1983_data.py` and read here directly from GitHub.

132.9.2 MLE estimation and predictability summaries

With the data in hand, we can run the MLE estimation and compute the predictability summaries.

Lognormality makes the MLE tractable when the assumption is correct.

As Hansen and Singleton [1983] stress, however, α is estimated with relatively large standard errors, and precise inference on risk aversion cannot be expected from aggregate monthly data alone.

```
lags = (2, 4, 6)

emp_frame, emp_data = get_estimation_data()
```

The following table reports the restricted MLE estimates of $\hat{\alpha}$ and $\hat{\beta}$ by lag length

```
emp_log_data = to_mle_array(emp_data)
mle_table, mle_fits = run_mle_by_lag(
    emp_log_data, lags=lags, verbose=False
)
mle_pretty = mle_table.rename(columns={
    "α_hat": r"\hat{\alpha}", "se_α": r"\mathrm{se}(\hat{\alpha})",
    "β_hat": r"\hat{\beta}", "se_β": r"\mathrm{se}(\hat{\beta})",
    "loglike": "logL", "n_obs": "T",
})
display_table(mle_pretty, fmt={
    r"\hat{\alpha}": "{:.4f}", r"\mathrm{se}(\hat{\alpha})": "{:.4f}",
    r"\hat{\beta}": "{:.4f}", r"\mathrm{se}(\hat{\beta})": "{:.4f}",
    "logL": "{:.1f}", "T": "{:.0f}",
})
```

NLAG	$\hat{\alpha}$	se($\hat{\alpha}$)	$\hat{\beta}$	se($\hat{\beta}$)	logL	T
2	−0.3228	0.7916	0.9976	0.0030	1209.3	237
4	−0.6402	0.6902	0.9980	0.0028	1206.4	235
6	−0.6155	0.6448	0.9979	0.0028	1199.2	233

The table reports $\hat{\alpha}$ and $\hat{\beta}$ by lag length for the sample used in the code cell above.

For comparison, Hansen and Singleton [1983] report $\hat{\alpha}$ values of -0.32 to -1.25 (standard errors 0.65 to 0.83) for the value-weighted return with nondurables consumption.

Our numbers fall into those ranges.

In risk-aversion units, this corresponds to $-\hat{\alpha}$ between 0.32 and 1.25.

We now compute the predictability summaries

```
unres_table, unres_fits = run_unrestricted_var_by_lag(
    emp_log_data,
    lags=lags,
)
```

NLAG	$\hat{\alpha}$	$\mathrm{Var}(\hat{E}[X_t \mid \psi_{t-1}])$	$\mathrm{Var}(\hat{E}[R_t \mid \psi_{t-1}])$	$\hat{\alpha}^2\,\mathrm{Var}(\hat{E}[X_t \mid \psi_{t-1}])$	R^2_X	R^2_R	T
2	−0.3228	0.000014	0.000001	0.000001	0.1630	0.0246	237
4	−0.6402	0.000016	0.000007	0.000007	0.1966	0.0459	235
6	−0.6155	0.000018	0.000007	0.000007	0.2270	0.0602	233

The column $\hat{\alpha}^2\,\mathrm{Var}(\hat{E}[X_t \mid \psi_{t-1}])$ should equal $\mathrm{Var}(\hat{E}[R_t \mid \psi_{t-1}])$ if the restriction (132.17) holds.

The columns R^2_X and R^2_R match the paper convention of reporting the values from the unrestricted VAR projections.

In Hansen and Singleton [1983], R^2_R is small (0.02 to 0.06) even when R^2_X is nontrivial: most stock-return variation is unpredictable.

Our estimates show the same pattern.

We now test the Euler-equation restrictions by comparing the restricted system to an unrestricted VAR.

```
lr_hs83 = restricted_vs_unrestricted_lr(mle_fits, unres_fits, lags=lags)
lr_pretty = lr_hs83.rename(columns={
    "lr_stat": "LR",
    "chi2_cdf": "chi2.cdf(LR,df)",
    "p_value": "p(LR)",
    "df": "df",
    "T": "T",
})
display_table(lr_pretty, fmt={
    "LR": "{:.3f}",
    "chi2.cdf(LR,df)": "{:.3f}",
    "p(LR)": "{:.3f}",
    "df": "{:.0f}",
    "T": "{:.0f}",
})

sig_level = 0.05
rejected_lags = [int(lag) for lag in lr_hs83.index if lr_hs83.loc[lag, "p_value"] <␣
↪sig_level]
not_rejected_lags = [int(lag) for lag in lr_hs83.index if lr_hs83.loc[lag, "p_value"]␣
↪>= sig_level]
print(f"Not rejected at 5%: {not_rejected_lags if not_rejected_lags else 'none'}")
```

NLAG	LR	p(LR)	chi2.cdf(LR,df)	df	T
2	5.977	0.113	0.887	3	237
4	10.023	0.187	0.813	7	235
6	12.519	0.326	0.674	11	233

```
Not rejected at 5%: [2, 4, 6]
```

The LR test does not reject the model for the value-weighted return, consistent with Hansen and Singleton [1983].

132.9.3 Treasury bill estimation

We now repeat the estimation using the 1-month Treasury bill return in place of the stock return.

Hansen and Singleton [1983] find that the model is strongly rejected for Treasury bills (Table 4 of their paper).

Because the nominally risk-free T-bill return is much more predictable than stock returns, the proportionality restriction has more bite, and the LR test has greater power to detect violations

```
tbill_frame, tbill_data = get_tbill_estimation_data()
```

The following table reports the restricted MLE estimates for Treasury bill returns by lag length.

```
tbill_log_data = to_mle_array(tbill_data)
tbill_mle_table, tbill_mle_fits = run_mle_by_lag(
    tbill_log_data, lags=lags, verbose=False
)
```

NLAG	$\hat{\alpha}$	se($\hat{\alpha}$)	$\hat{\beta}$	se($\hat{\beta}$)	logL	T
2	−0.3060	0.0600	0.9996	0.0003	1817.5	237
4	−0.2882	0.0612	0.9996	0.0003	1804.9	235
6	−0.1991	0.0571	0.9996	0.0003	1791.4	233

The LR test below remains the main specification check in the paper

```
tbill_unres_table, tbill_unres_fits = run_unrestricted_var_by_lag(
    tbill_log_data,
    lags=lags,
)

tbill_unres_pretty = tbill_unres_table.rename(columns={"loglike": "logL", "n_obs": "T
↪"})
display_table(tbill_unres_pretty, fmt={
    "logL": "{:.1f}", "T": "{:.0f}",
})
```

NLAG	logL	T
2	1838.8	237
4	1828.2	235
6	1824.8	233

NLAG	$\hat{\alpha}$	Var($\hat{E}[X_t \mid \psi_{t-1}]$)	Var($\hat{E}[R_t \mid \psi_{t-1}]$)	$\hat{\alpha}^2$ Var($\hat{E}[X_t \mid \psi_{t-1}]$)	R^2_X	R^2_R	T
2	−0.3060	0.000019	0.000002	0.000002	0.2309	0.1618	237
4	−0.2882	0.000018	0.000001	0.000001	0.2223	0.1870	235
6	−0.1991	0.000019	0.000001	0.000001	0.2333	0.2560	233

The following table reports the likelihood ratio tests for the Treasury bill model.

```
tbill_lr = restricted_vs_unrestricted_lr(tbill_mle_fits, tbill_unres_fits, lags=lags)
```

NLAG	LR	p(LR)	chi2.cdf(LR,df)	df	T
2	42.503	0.000	1.000	3	237
4	46.617	0.000	1.000	7	235
6	66.656	0.000	1.000	11	233

```
T-bill model rejected at 5% for lags: [2, 4, 6]
```

Hansen and Singleton [1983] find the same qualitative pattern in their 1959–1978 sample: the Treasury bill model is rejected much more strongly than the value-weighted stock-return model.

Why does the LR test not reject the model for stocks?

As we hinted earlier, one reason may be limited test power when return predictability is small (as reflected in the low R_R^2 for stocks).

When aggregate stock returns are nearly unpredictable, there is almost no predictable variation to constrain.

132.10 Residual diagnostics

As a final check, let's inspect the residual paths, histograms, and diagnostic statistics from the restricted model to assess whether the normality and serial independence assumptions hold.

```
diag_lag = 2
diag_fit = mle_fits[diag_lag]
resid = diag_fit["residuals"]

if diag_fit["converged"] and resid is not None:
    fig, axes = plt.subplots(2, 2, figsize=(12, 8))
    axes[0, 0].plot(resid[:, 0], lw=2)
    axes[0, 0].axhline(0.0, color="black", lw=2)
    axes[0, 0].set_ylabel("consumption residual")
    axes[0, 0].set_xlabel("observation")

    axes[0, 1].plot(resid[:, 1], lw=2)
    axes[0, 1].axhline(0.0, color="black", lw=2)
    axes[0, 1].set_ylabel("return residual")
    axes[0, 1].set_xlabel("observation")

    axes[1, 0].hist(resid[:, 0], bins=30, edgecolor="white")
    axes[1, 0].set_xlabel("consumption residual")
    axes[1, 0].set_ylabel("count")

    axes[1, 1].hist(resid[:, 1], bins=30, edgecolor="white")
    axes[1, 1].set_xlabel("return residual")
    axes[1, 1].set_ylabel("count")
    plt.tight_layout()
    plt.show()
```

The following table reports Jarque-Bera normality tests and Durbin-Watson serial correlation statistics for the restricted-model residuals.

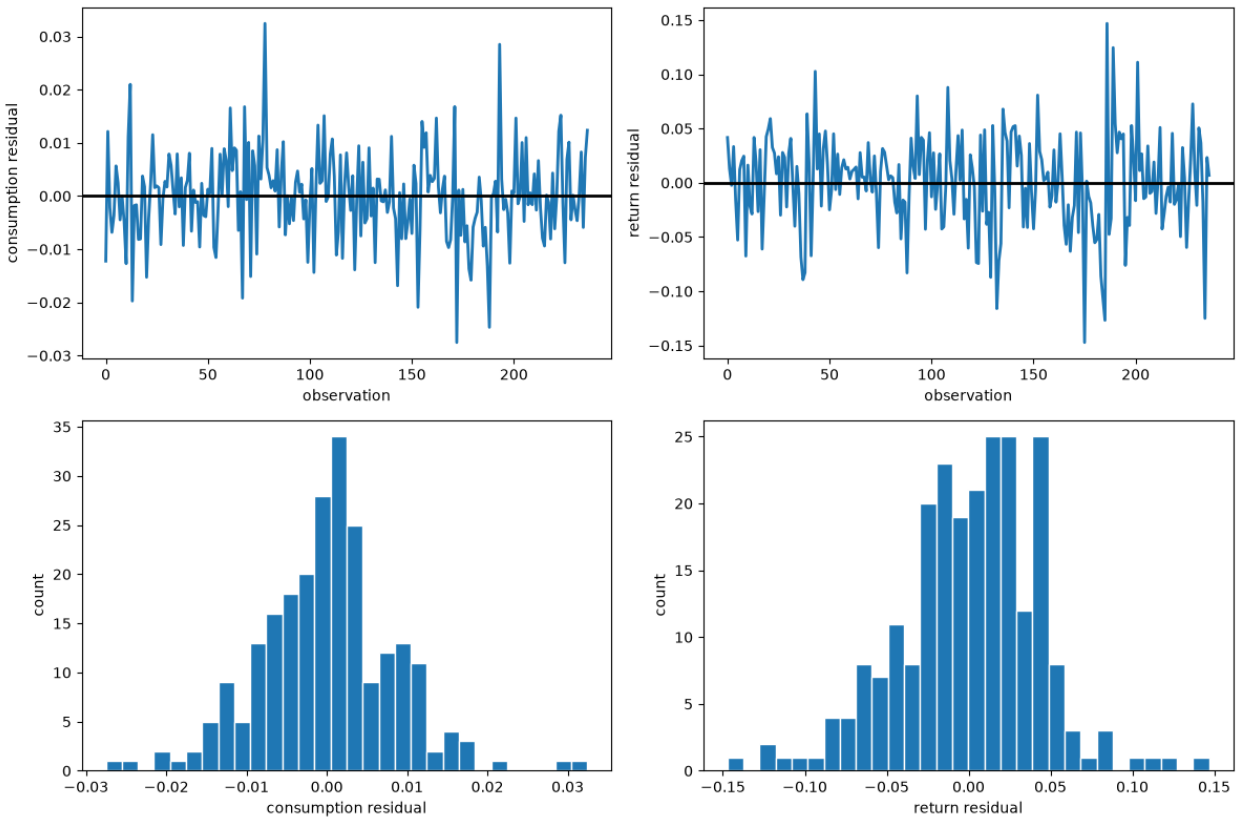


Fig. 132.1: Restricted-model residual diagnostics

```
if diag_fit["converged"] and resid is not None:
    diag = residual_diagnostics(resid)
    diag_df = pd.DataFrame({
        "JB stat": [diag["consumption_jb_stat"], diag["return_jb_stat"]],
        "JB p-val": [diag["consumption_jb_pval"], diag["return_jb_pval"]],
        "DW": [diag["consumption_dw"], diag["return_dw"]],
    }, index=pd.Index(["consumption", "return"], name="series"))
    display_table(diag_df, fmt={
        "JB stat": "{:.2f}", "JB p-val": "{:.4f}", "DW": "{:.3f}",
    })
```

series	JB stat	JB p-val	DW
consumption	18.16	0.0001	1.969
return	13.27	0.0013	1.766

The residual time-series plots reveal periods of volatility clustering, while the histograms show departures from bell curve with fatter tails.

The Jarque-Bera statistics are large for both series, rejecting the normality assumption that underlies the likelihood.

The Durbin-Watson statistics are close to 2 for both series, so serial correlation is not a concern.

This motivates the companion lecture *Estimating Euler Equations by Generalized Method of Moments* where we discuss how GMM avoids the normality assumption.

132.11 Connection to the equity premium puzzle

Our estimates reproduce the pattern that Mehra and Prescott [1985] later called the **equity premium puzzle**.

- *Low estimated risk aversion:* The estimated $\hat{\alpha}$ values (and thus risk aversion $-\hat{\alpha}$) from the table above are similar to those in Hansen and Singleton [1983], who report $\hat{\alpha}$ between -0.32 and -1.25.
- *Tiny return predictability:* The unrestricted-VAR R^2_R values are comparable to the 0.02 to 0.06 range in Hansen and Singleton [1983] – the predictable component of stock returns is small relative to the unpredictable component.
- *Strong rejection for Treasury bills:* The Euler-equation restrictions are decisively rejected for the nominally risk-free Treasury bill return, just as in Table 4 of Hansen and Singleton [1983].

The restrictions have more bite when the return series is more predictable (the unrestricted-VAR R^2_R for bills is much larger than for stocks), a precursor to what Weil [1989] would later call the **risk-free rate puzzle**.

Mehra and Prescott [1985] presented closely related findings.

In a calibrated version of the consumption-based model with CRRA utility, they showed that for relative risk aversion γ_{MP} in the range of 0 to 10, the model could not simultaneously match:

1. the average annual equity premium of about 6%,
2. the average annual risk-free rate of about 1%.

Matching the equity premium requires γ_{MP} well above 10, while matching the risk-free rate requires low γ_{MP}.

The same difficulty appears in our estimates: the implied risk aversion $-\hat{\alpha}$ is too low to generate a large equity premium.

132.12 Another approach

A companion lecture *Estimating Euler Equations by Generalized Method of Moments* describes a closely related paper in which Hansen and Singleton specified less about the joint distribution of returns and consumption growth.

They formulated an incomplete probability model that stops short of specifying a likelihood function.

To proceed, they used a generalized methods of moments (GMM) estimator to estimate key parameters that appear in the Euler equation.

CHAPTER

THREE

ESTIMATING EULER EQUATIONS BY GENERALIZED METHOD OF MOMENTS

Contents

- *Estimating Euler Equations by Generalized Method of Moments*
  - *Overview*
  - *The economic model*
  - *From conditional to unconditional moments*
  - *GMM criterion and asymptotic theory*
  - *Covariance estimation and the choice of instruments*
  - *Empirical GMM estimation*
  - *Summary*

133.1 Overview

This lecture implements the generalized instrumental variables estimator of Hansen and Singleton [1982] for nonlinear rational expectations models.

The preceding lecture *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns* derives the consumption Euler equation from the representative consumer's problem with CRRA preferences and estimates it by maximum likelihood under normality assumptions.

That approach requires specifying the joint distribution of consumption and returns, and its validity depends on lognormality being correct.

However, as we saw in *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*, the lognormal model is rejected by the data.

Moreover, outside of linear-quadratic environments, closed-form solutions for equilibrium typically require strong assumptions about the stochastic properties of forcing variables, the nature of preferences, or the production technology.

Hansen and Singleton [1982] propose an estimation strategy that circumvents this requirement.

The key idea is that the Euler equations from economic agents' optimization problems imply a set of population orthogonality conditions that depend on observable variables and unknown preference parameters.

By making sample counterparts of these orthogonality conditions close to zero, the parameters can be estimated without explicitly solving for the stochastic equilibrium and without specifying the distribution of the observable variables.

Though maximum likelihood estimators (such as the MLE in *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*) will be asymptotically more efficient when the distributional assumptions are correctly specified.

Relative to the full paper, we only estimate one return at a time (value-weighted stock returns), using only monthly nondurable consumption (`ND`), and omitting their maximum-likelihood comparison (Table II) and multi-return systems (Table III).

```
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
from IPython.display import Latex
from numba import njit
from scipy import stats
from scipy.optimize import minimize
from statsmodels.sandbox.regression import gmm
from statsmodels.tsa.stattools import acf
```

We also define a helper to display DataFrames as LaTeX arrays in the hidden cell below

133.2 The economic model

We consider a single-good economy with a representative consumer whose preferences are of the CRRA type, following Hansen and Singleton [1982] and Hansen and Singleton [1983].

Note

The following discussion is very close to the setup in *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*, but it is more general in that it allows for multiple assets with different maturities and does not assume lognormality.

The representative consumer chooses stochastic consumption and investment plans to maximize

$$\max E_0 \sum_{t=0}^{\infty} \beta^t u(C_t) \tag{133.1}$$

where C_t is consumption in period t, $\beta \in (0,1)$ is a subjective discount factor, and $u(\cdot)$ is a strictly concave period utility function.

The consumer trades N assets with potentially different maturities.

Let Q_{jt} denote the quantity of asset j held at the end of date t, P_{jt} its price at date t, and R_{jt} the date-t payoff from holding one unit of an M_j-period asset purchased at date $t - M_j$.

Feasible consumption and investment plans must satisfy the sequence of budget constraints

$$C_t + \sum_{j=1}^{N} P_{jt} Q_{jt} \leq \sum_{j=1}^{N} R_{jt} Q_{jt-M_j} + W_t, \tag{133.2}$$

where W_t is real labor income at date t.

We specialize to CRRA preferences

$$u(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}, \quad \gamma > 0, \tag{133.3}$$

where γ is the coefficient of relative risk aversion.

The maximization of (133.1) subject to (133.2) gives the first-order necessary conditions (see Lucas [1978], Brock [1982], Prescott and Mehra [1980], and *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*):

$$P_{jt} u'(C_t) = \beta^{M_j} E_t \left[R_{jt+M_j} u'(C_{t+M_j}) \right], \quad j = 1, \ldots, N. \tag{133.4}$$

When asset j is a one-period stock ($M_j = 1$) with payoff $R_{jt+1} = P_{jt+1} + D_{jt+1}$ where D_{jt} is the dividend, the **gross real return** is $R^i_{t+1} = (P_{i,t+1} + D_{i,t+1})/P_{i,t}$.

Substituting the CRRA marginal utility into (133.4) and dividing both sides by $P_{jt} u'(C_t)$ yields the Euler equation

$$E_t \left[\beta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} R^i_{t+1} \right] = 1, \tag{133.5}$$

where R^i_{t+1} is the gross real return on asset i.

We define the **stochastic discount factor** $M_{t+1}(\theta) = \beta (C_{t+1}/C_t)^{-\gamma}$ with parameter vector $\theta = (\gamma, \beta)$.

In this notation the Euler equation becomes $E_t[M_{t+1}(\theta) R^i_{t+1} - 1] = 0$.

As we have seen and will see, equation (133.5) is the central object of both *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns* and this lecture.

It holds for every traded asset for which the agent's optimality conditions apply (interior solution, no binding portfolio constraints or transaction costs).

It depends on observable quantities (consumption growth and returns) and unknown preference parameters (γ and β).

The challenge that Hansen and Singleton [1982] address is how to estimate θ from (133.5) without specifying the rest of the economic environment.

133.3 From conditional to unconditional moments

A natural approach to estimating θ from (133.5) would be to specify the entire economic environment, solve for equilibrium, and apply maximum likelihood.

Just like what we did together in *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*,

Hansen and Singleton [1982] argue that this is impractical for most nonlinear models because it requires strong assumptions about the stochastic properties of "forcing variables" and the production technology.

Their alternative is to work directly with the Euler equation's implications for observable moments.

The Euler equation (133.5) states that $E_t[M_{t+1}(\theta_0) R^i_{t+1} - 1] = 0$.

Let z_t denote any q-dimensional vector of variables that are in the agent's time-t information set and observed by the econometrician.

Because z_t is known at date t,

$$E\left[(M_{t+1}(\theta_0) R^i_{t+1} - 1) \otimes z_t \right] = E\left[z_t \otimes \underbrace{E_t[M_{t+1}(\theta_0) R^i_{t+1} - 1]}_{=0 \text{ by the Euler equation}} \right] = 0.$$

This yields the moment restriction

$$E\left[(M_{t+1}(\theta_0)R^i_{t+1} - 1) \otimes z_t\right] = 0, \tag{133.6}$$

where $\otimes$ denotes the Kronecker product.

The vector z_t plays the role of **instruments**.

The conditional Euler equation $E_t[M_{t+1}R^i_{t+1} - 1] = 0$ says that the pricing error is unpredictable given *everything* in the agent's time-t information set.

That is a very strong restriction – it says the pricing error is orthogonal to every time-t measurable random variable.

We cannot use the entire information set in practice, but we can pick any finite collection of time-t observable variables z_t and the orthogonality must still hold.

Each variable we include in z_t gives us one sample moment condition $\frac{1}{T}\sum_t (M_{t+1}R^i_{t+1} - 1)\, z_{kt} \approx 0$ that we can compute from data.

More instruments means more orthogonality conditions to match, which can improve efficiency and provides more overidentifying restrictions to test the model against.

With q instruments and m Euler equations we obtain mq moment restrictions for estimating the parameter vector θ.

For one return and p lags of instruments, Hansen and Singleton [1982] use

$$z_t = \left[1, R_t, g_t, R_{t-1}, g_{t-1}, \ldots, R_{t-p+1}, g_{t-p+1}\right]^\top, \tag{133.7}$$

where $g_t = C_t / C_{t-1}$ is gross consumption growth.

Note

Hansen and Singleton [1982] describe z_t as "lagged values of x_{t+1}" where $x_{t+1} = (R_{t+1}, g_{t+1})$.

The constant 1 is not stated explicitly but is implied by the degrees of freedom reported in their Table I.

More generally, Hansen and Singleton [1982] write the first-order condition as $E_t[h(x_{t+n}, b_0)] = 0$, where x_{t+n} is a vector of observables dated $t+n$ and b_0 is the true parameter vector.

The disturbance $u_{t+1} = h(x_{t+1}, b_0)$ is serially uncorrelated in the one-period stock case.

Stacking moment conditions via the Kronecker product gives $f(x_{t+n}, z_t, b) = h(x_{t+n}, b) \otimes z_t$, a vector of dimension mq.

The resulting unconditional restriction $E[f(x_{t+n}, z_t, b_0)] = 0$ nests both the single-return one-period Euler equation and the multi-maturity asset pricing restrictions.

The orthogonality condition and lagged instrument vector follow equations (133.6) and (133.7)

```
def euler_error_horizon(params, exog, horizon=1):
    """
    Compute Euler-equation pricing errors for (γ, β) at a given horizon.
    """
    if horizon < 1:
        raise ValueError("horizon must be at least one.")
    γ, β = params
    gross_return = exog[:, 0]
    gross_cons_growth = exog[:, 1]
    return (β ** horizon) * gross_cons_growth ** (-γ) * gross_return - 1.0
```

(continues on next page)

(continued from previous page)

```
def euler_error_grad_horizon(params, exog, horizon=1):
    """
    Gradient of the Euler error wrt (γ, β) at a given horizon.
    """
    if horizon < 1:
        raise ValueError("horizon must be at least one.")
    γ, β = params
    gross_return = exog[:, 0]
    gross_cons_growth = exog[:, 1]

    g_pow = gross_cons_growth ** (-γ)
    common = (β ** horizon) * g_pow * gross_return

    dγ = -common * np.log(gross_cons_growth)
    dβ = horizon * (β ** (horizon - 1)) * g_pow * gross_return
    return np.column_stack([dγ, dβ])


def euler_error(params, exog):
    """
    One-period Euler-equation pricing errors for (γ, β).
    """
    return euler_error_horizon(params, exog, horizon=1)
```

A helper aligns outcomes and lagged instruments for nonlinear IV-GMM

```
def build_gmm_arrays(data, n_lags):
    """
    Build endog, exog, and instruments for nonlinear IV-GMM.
    """
    if n_lags < 1:
        raise ValueError("n_lags must be at least one.")
    if data.shape[0] <= n_lags:
        raise ValueError("Sample size must exceed n_lags.")

    t_obs = data.shape[0]
    exog = data[n_lags:, :]
    endog = np.zeros(exog.shape[0])
    n_obs = t_obs - n_lags
    n_instr = 2 * n_lags + 1

    instruments = np.empty((n_obs, n_instr))
    instruments[:, 0] = 1.0

    for j in range(n_lags):
        left = 2 * j + 1
        right = left + 2
        instruments[:, left:right] = data[n_lags - 1 - j : t_obs - 1 - j, :]

    return endog, exog, instruments
```

When the asset has maturity $n > 1$, the Euler equation involves n-period compounded returns and consumption growth.

For CRRA preferences, the n-period Euler restriction is

$$E_t\left[\beta^n \left(\frac{C_{t+n}}{C_t}\right)^{-\gamma} R^i_{t,t+n}\right] = 1. \tag{133.8}$$

For estimation, the n-period exog can either use directly observed n-period returns/payoffs, or be constructed by compounding one-period returns and consumption growth over n consecutive periods.

Again, a key requirement from Hansen and Singleton [1982] is that instruments z_t must lie in the agent's time-t information set $\mathcal{I}_t$.

For the multi-period case, instruments must be measurable with respect to $\mathcal{I}_t$.

In particular, one should avoid any lagged multi-period aggregates that still include periods after t, since these would embed realizations not in $\mathcal{I}_t$.

The n-period exog is constructed by compounding one-period data, and instruments are timed to lie in $\mathcal{I}_t$.

```
def build_gmm_arrays_horizon(one_period_data, n_lags, horizon):
    """
    Build endog, exog, and instruments for multi-period GMM.

    Exog contains n-period compounded returns and consumption growth.
    """
    if horizon < 1:
        raise ValueError("horizon must be at least one.")
    if n_lags < 1:
        raise ValueError("n_lags must be at least one.")
    T = one_period_data.shape[0]
    if T <= n_lags + horizon:
        raise ValueError("Sample size too small for given n_lags and horizon.")

    # Each observation starts at index t (the first period in the window)
    # The window spans one_period_data[t : t + horizon]
    # Instruments use one_period_data[t - 1], ..., one_period_data[t - n_lags]
    starts = np.arange(n_lags, T - horizon + 1)
    n_obs = len(starts)

    exog = np.empty((n_obs, 2))
    n_instr = 2 * n_lags + 1
    instruments = np.empty((n_obs, n_instr))
    instruments[:, 0] = 1.0

    for i, t in enumerate(starts):
        window = one_period_data[t : t + horizon, :]
        exog[i, 0] = np.prod(window[:, 0])    # n-period return
        exog[i, 1] = np.prod(window[:, 1])    # n-period consumption growth
        for j in range(n_lags):
            instruments[
                i, 2 * j + 1 : 2 * j + 3] = one_period_data[t - 1 - j, :]

    endog = np.zeros(n_obs)
    return endog, exog, instruments
```

When $n > 1$, the Euler equation involves variables dated $t + n$, and the disturbance $u_{t+n} = h(x_{t+n}, b_0)$ will generally be serially correlated.

As Hansen and Singleton [1982] note, if the m assets are all one-period stocks, then u is serially uncorrelated because observations on x_{t-s}, $s \geq 0$, are in the agent's time-t information set and $E_t[h(x_{t+1}, b_0)] = 0$.

But if $n_j > 1$ for some asset j, the condition $E_t[h(x_{t+n}, b_0)] = 0$ does not preclude serial correlation in u, since x_{t+n-1} is not necessarily in I_t when $n > 1$.

The number of population autocovariances in the long-run covariance S_0 is determined by n, the order of the moving average disturbance term u_t.

We implement this directly as a finite-order covariance estimator

```
def finite_ma_covariance(moment_series, ma_order):
    """
    Estimate
    S = Gamma_0 + sum_{j=1}^{ma_order}(Gamma_j + Gamma_j.T) for moment vectors.
    """
    if ma_order < 0:
        raise ValueError("ma_order must be nonnegative.")
    if moment_series.ndim != 2:
        raise ValueError("moment_series must be 2D.")

    t_obs, n_mom = moment_series.shape
    if t_obs <= ma_order:
        raise ValueError("Need more observations than ma_order.")

    # Use the *uncentered* cross products
    # T^{-1} sum_t f_t f_{t-j}' and then add the symmetric lag terms.
    s_hat = moment_series.T @ moment_series / t_obs

    for j in range(1, ma_order + 1):
        gamma_j = moment_series[j:, :].T @ moment_series[:-j, :] / t_obs
        s_hat += gamma_j + gamma_j.T

    ridge = 1e-8 * np.eye(n_mom)
    return s_hat + ridge
```

The estimation procedure in Hansen and Singleton [1982] is a two-step generalized instrumental variables procedure.

In the first step, we minimize the GMM criterion with a suboptimal weighting matrix (the identity) to obtain a consistent preliminary estimate b_T.

In the second step, we use b_T to estimate the covariance matrix of the sample moment conditions and invert it to form the optimal weighting matrix, then re-minimize the criterion.

Let's implement this algorithm

```
def two_step_gmm(data, n_lags, ma_order=0, horizon=1, start_params=None):
    """
    Two-step GMM with finite-order covariance.

    The Euler error uses β**horizon.
    """
    if start_params is None:
        start_params = np.array([1.0, 0.99])
    else:
        start_params = np.asarray(start_params, dtype=float)

    if horizon == 1:
        _, exog, instruments = build_gmm_arrays(data, n_lags)
    else:
        _, exog, instruments = build_gmm_arrays_horizon(data, n_lags, horizon)
    n_obs = exog.shape[0]

    def sample_moments(params):
        err = euler_error_horizon(params, exog, horizon=horizon)
        return err[:, None] * instruments

    def objective(params, weight_matrix):
```

(continues on next page)

(continued from previous page)

```
    g_bar = sample_moments(params).mean(axis=0)
    return float(g_bar @ weight_matrix @ g_bar)

def objective_grad(params, weight_matrix):
    g_bar = sample_moments(params).mean(axis=0)
    grad_err = euler_error_grad_horizon(params, exog, horizon=horizon)
    d_bar = (instruments.T @ grad_err) / n_obs
    return 2.0 * d_bar.T @ weight_matrix @ g_bar

q = instruments.shape[1]
w_identity = np.eye(q)

bounds = [(-2.0, 10.0), (0.85, 1.5)]

def coarse_starts(weight_matrix, n_best=5):
    γ_grid = np.linspace(bounds[0][0], bounds[0][1], 33)
    β_grid = np.linspace(bounds[1][0], bounds[1][1], 33)
    scored = []
    for γ0 in γ_grid:
        for β0 in β_grid:
            params0 = np.array([γ0, β0])
            val = objective(params0, weight_matrix)
            if np.isfinite(val):
                scored.append((val, params0))
    scored.sort(key=lambda item: item[0])
    return [params for _, params in scored[:n_best]] or [start_params]

def best_local_minimize(weight_matrix, starts):
    best = None
    for x0 in starts:
        res = minimize(
            objective,
            x0=x0,
            args=(weight_matrix,),
            jac=objective_grad,
            method="L-BFGS-B",
            bounds=bounds,
            options={"maxiter": 25_000},
        )
        if not np.isfinite(res.fun):
            continue
        if best is None or (res.fun < best.fun):
            best = res
    return best if best is not None else minimize(
        objective,
        x0=start_params,
        args=(weight_matrix,),
        jac=objective_grad,
        method="L-BFGS-B",
        bounds=bounds,
        options={"maxiter": 25_000},
    )

step1_starts = [start_params] + coarse_starts(w_identity, n_best=5)
step1 = best_local_minimize(w_identity, step1_starts)
params1 = step1.x
```

(continues on next page)

(continued from previous page)

```
    m1 = sample_moments(params1)
    s_hat = finite_ma_covariance(m1, ma_order=ma_order)
    w_opt = np.linalg.pinv(s_hat)

    step2_starts = [params1] + coarse_starts(w_opt, n_best=5)
    step2 = best_local_minimize(w_opt, step2_starts)
    params2 = step2.x

    m2 = sample_moments(params2)
    s_hat2 = finite_ma_covariance(m2, ma_order=ma_order)
    w_opt2 = np.linalg.pinv(s_hat2)
    g2 = m2.mean(axis=0)
    j_stat = float(n_obs * (g2 @ w_opt2 @ g2))
    df = instruments.shape[1] - len(params2)
    j_prob = float(stats.chi2.cdf(j_stat, df=df)) if df > 0 else np.nan
    p_value = float(1.0 - j_prob) if df > 0 else np.nan

    # Asymptotic covariance under optimal weighting: (D' S^{-1} D)^{-1} / T.
    grad_err = euler_error_grad_horizon(params2, exog, horizon=horizon)
    d_hat = (instruments.T @ grad_err) / n_obs
    cov_hat = np.linalg.pinv(d_hat.T @ w_opt2 @ d_hat) / n_obs
    se_hat = np.sqrt(np.diag(cov_hat))

    return {
        "params_step1": params1,
        "params_step2": params2,
        "se_step2": se_hat,
        "weight_opt": w_opt2,
        "j_stat": j_stat,
        "j_df": int(df),
        "j_prob": j_prob,
        "j_pval": p_value,
        "n_obs": int(n_obs),
        "success": bool(step2.success),
    }
```

133.4 GMM criterion and asymptotic theory

We now formalize the estimation procedure.

Let $m_t(\theta) = (M_{t+1}(\theta) R^i_{t+1} - 1) \otimes z_t$ denote the vector of moment conditions at date t, and define the sample mean

$$g_T(\theta) = \frac{1}{T} \sum_{t=1}^{T} m_t(\theta). \tag{133.9}$$

If the model is correctly specified, $g_T(\theta_0)$ should be close to zero for large T.

We estimate θ by choosing the parameter vector that makes g_T as close to zero as possible, measured by a quadratic form:

$$\hat{\theta} = \arg\min_{\theta} g_T(\theta)^\top W_T g_T(\theta) \tag{133.10}$$

where W_T is a symmetric positive-definite weighting matrix.

Under regularity conditions given in Hansen [1982], the GMM estimator is consistent, asymptotically normal, and has the sandwich covariance matrix

$$\sqrt{T}(\hat{\theta} - \theta_0) \Rightarrow N\left(0, (D^\top W D)^{-1} D^\top W S W D (D^\top W D)^{-1}\right), \tag{133.11}$$

where $D = E[\partial m_t(\theta_0)/\partial\theta^\top]$ is the Jacobian of the moment conditions, S is the long-run covariance matrix of $m_t(\theta_0)$, and W is the probability limit of W_T.

Hansen [1982] shows that the optimal weighting matrix is $W^* = S^{-1}$, which yields the smallest asymptotic covariance matrix among all choices of W.

Under $W = S^{-1}$ the sandwich simplifies to $(D^\top S^{-1} D)^{-1}$.

When the number of moment conditions r (e.g., $r = mq$ for m Euler equations and q instruments) exceeds the number of parameters k, the model is overidentified and we can test whether the data are consistent with the maintained restrictions.

Hansen and Singleton [1982] test the overidentifying restrictions using a result from Hansen [1982]:

$$J_T = T\, g_T(\hat{\theta})^\top \hat{S}^{-1} g_T(\hat{\theta}) \Rightarrow \chi^2_{r-k}, \tag{133.12}$$

where $\hat{S}$ is a consistent estimator of S.

A large J_T relative to χ^2_{r-k} critical values leads to rejection of the model's overidentifying restrictions.

For the multi-period case ($n > 1$), the disturbance is at most MA($n-1$) as discussed above, so Hansen and Singleton [1982] obtain the optimal weighting matrix by setting $W_0 = S_0^{-1}$ where

$$S_0 = \sum_{j=-n+1}^{n-1} E\left[f(x_{t+n}, z_t, b_0)\, f(x_{t+n-j}, z_{t-j}, b_0)^\top\right]. \tag{133.13}$$

133.5 Covariance estimation and the choice of instruments

For the one-period Euler equation ($n = 1$), the disturbance $u_{t+1} = M_{t+1}(\theta_0) R_{t+1} - 1$ is a martingale difference sequence.

In this case the moment vector $m_t(\theta_0) = u_{t+1} \otimes z_t$ is serially uncorrelated, and the long-run covariance S equals the contemporaneous variance $E[m_t m_t^\top]$.

The covariance estimator therefore requires no kernel or bandwidth choice.

We simply use the sample analogue $\hat{S} = T^{-1} \sum_t m_t m_t^\top$.

In our implementation below, we use a HAC (heteroskedasticity and autocorrelation consistent) estimator against possible mild serial dependence from time aggregation or measurement timing.

This is a modern precaution, not part of the original Hansen and Singleton [1982] procedure, which exploits the known MA order directly as in (133.13).

The number of instrument lags p determines how many orthogonality conditions we use and hence the power of the J test.

Hansen and Singleton [1982] report results for NLAG $= 1, 2, 4, 6$ and note (footnote 12) that using more orthogonality conditions may lead to estimators with less desirable small-sample properties.

Below we implement the estimation procedure and allow the user to choose the number of lags and whether to use a HAC estimator for the covariance matrix

```
def estimate_gmm(
    data,
    n_lags,
    start_params=None,
    use_hac=True,
    hac_maxlag=None,
    maxiter=2,
```

(continues on next page)

(continued from previous page)

```
):
    """
    Estimate Euler-equation parameters with nonlinear IV-GMM.
    """
    if start_params is None:
        start_params = np.array([1.0, 0.99])

    endog, exog, instruments = build_gmm_arrays(data, n_lags)
    model = gmm.NonlinearIVGMM(endog, exog, instruments, euler_error)

    if use_hac:
        if hac_maxlag is None:
            hac_maxlag = max(
                1, int(
                    np.floor(4.0 * (endog.shape[0] / 100.0) ** (2.0 / 9.0))))
        result = model.fit(
            start_params=start_params,
            maxiter=maxiter,
            optim_method="bfgs",
            optim_args={"disp": False},
            weights_method="hac",
            wargs={"maxlag": hac_maxlag},
        )
    else:
        result = model.fit(
            start_params=start_params,
            maxiter=maxiter,
            optim_method="bfgs",
            optim_args={"disp": False},
        )

    return result
```

Next we include a helper to run GMM estimation across lag lengths and summarize results in a table

```
def run_gmm_by_lag(
    data,
    lags=(1, 2, 4, 6),
    use_hac=True,
    hac_maxlag=None,
):
    """
    Estimate GMM models by lag length and return a summary table.
    """
    rows = []
    results = {}

    for lag in lags:
        res = estimate_gmm(data, n_lags=lag, use_hac=use_hac, hac_maxlag=hac_maxlag)
        results[lag] = res
        j_stat, j_pval, j_df = res.jtest()
        j_prob = float(stats.chi2.cdf(j_stat, df=j_df)) if j_df > 0 else np.nan
        rows.append(
            {
                "NLAG": lag,
                "y_hat": res.params[0],
                "se_y": res.bse[0],
```

(continues on next page)

(continued from previous page)

```
                "β_hat": res.params[1],
                "se_β": res.bse[1],
                "j_stat": j_stat,
                "j_prob": j_prob,
                "j_pval": j_pval,
                "j_df": int(j_df),
                "n_obs": int(res.nobs),
            }
        )

    table = pd.DataFrame(rows).set_index("NLAG")
    return table, results


def run_two_step_by_lag(
    data,
    lags=(1, 2, 4, 6),
    horizon=1,
):
    """
    Two-step GMM with exact S0 (MA order 0) across lag lengths.
    """
    rows = []
    start_params = None
    for lag in lags:
        res = two_step_gmm(
            data,
            n_lags=lag,
            ma_order=0,
            horizon=horizon,
            start_params=start_params,
        )
        start_params = res["params_step2"]
        rows.append(
            {
                "NLAG": lag,
                "γ_hat": res["params_step2"][0],
                "se_γ": res["se_step2"][0],
                "β_hat": res["params_step2"][1],
                "se_β": res["se_step2"][1],
                "j_stat": res["j_stat"],
                "j_prob": res["j_prob"],
                "j_pval": res["j_pval"],
                "j_df": res["j_df"],
                "n_obs": res["n_obs"],
            }
        )
    return pd.DataFrame(rows).set_index("NLAG")
```

133.5.1 Simulation

Before applying the estimator to real data, we verify that GMM recovers known parameters from simulated data.

We build a simulator that generates synthetic return-growth pairs satisfying the Euler equation by construction.

We generate log consumption growth from a stationary AR(1), compute the stochastic discount factor at known true parameters, and construct gross returns as $R_{t+1} = \xi_{t+1}/M_{t+1}(\theta_0)$ where ξ_{t+1} is an iid lognormal shock with mean one.

```
@njit
def _ar1_simulate(mu_c, phi_c, sigma_c, shocks_c, total_n):
    """
    Simulate AR(1) log consumption growth.
    """
    delta_c = np.empty(total_n)
    delta_c[0] = mu_c
    for t in range(1, total_n):
        delta_c[t] = mu_c * (1.0 - phi_c) + phi_c * delta_c[t - 1] + sigma_c * shocks_
↪c[t]
    return delta_c


def simulate_euler_sample(
    n_obs,
    γ_true=0.8,
    β_true=0.993,
    seed=0,
):
    """
    Simulate [gross real return, gross consumption growth].
    """
    rng = np.random.default_rng(seed)
    mu_c = 0.0015
    sigma_c = 0.006
    phi_c = 0.4
    sigma_eta = 0.02
    burn_in = 200

    total_n = n_obs + burn_in
    shocks_c = rng.standard_normal(total_n)
    delta_c = _ar1_simulate(mu_c, phi_c, sigma_c, shocks_c, total_n)

    cons_growth = np.exp(delta_c[burn_in:])
    sdf = β_true * cons_growth ** (-γ_true)

    # Positive mean-one return shock: E[ξ]=1 so E[M R]=1 by construction.
    eps = rng.standard_normal(n_obs)
    xi = np.exp(sigma_eta * eps - 0.5 * sigma_eta**2)
    gross_return = xi / sdf

    return np.column_stack([gross_return, cons_growth])
```

We set $\gamma = 2$ and $\beta = 0.995$ as the true parameters and generate 700 monthly observations from the Euler-consistent DGP.

```
γ_true = 2.0
β_true = 0.995
sim_data = simulate_euler_sample(
```

(continues on next page)

(continued from previous page)

```
    n_obs=5000,
    γ_true=γ_true,
    β_true=β_true,
    seed=0,
)

print(f"Simulation sample size: {sim_data.shape[0]}")
print(f"True γ: {γ_true:.3f}")
print(f"True β: {β_true:.3f}")
```

```
Simulation sample size: 5000
True γ: 2.000
True β: 0.995
```

We now estimate GMM across lag lengths, following the format of Table I in Hansen and Singleton [1982].

```
sim_table = run_two_step_by_lag(sim_data, lags=(1, 2, 4, 6), horizon=1)
```

NLAG	$\hat{\gamma}$	se($\hat{\gamma}$)	$\hat{\beta}$	se($\hat{\beta}$)	J	df	Prob(J)
1	2.0806	0.1551	0.9948	0.0004	0.135	1	0.286
2	2.0927	0.1546	0.9949	0.0004	2.493	3	0.523
4	2.0930	0.1471	0.9949	0.0004	3.936	7	0.213
6	2.0991	0.1443	0.9949	0.0004	8.125	11	0.298

GMM recovers the true γ and β across lag specifications pretty closely.

For hypothesis testing, the right-tail p value is $1 - \text{Prob}(J)$.

The large standard errors visible in both papers' tables suggest that the preference parameters γ and β can be weakly identified.

To visualize this, we plot the GMM criterion over a (γ, β) grid using the simulated data

```
def gmm_objective_surface(
    data,
    n_lags=2,
    γ_grid=None,
    β_grid=None,
):
    """
    Compute identity-weighted GMM objective on a parameter grid.
    """
    _, exog, instruments = build_gmm_arrays(data, n_lags)

    if γ_grid is None:
        γ_grid = np.linspace(-1.0, 8.0, 70)
    if β_grid is None:
        β_grid = np.linspace(0.96, 1.02, 70)

    objective = np.empty((len(β_grid), len(γ_grid)))

    for i, β_val in enumerate(β_grid):
        for j, γ_val in enumerate(γ_grid):
            err = euler_error(np.array([γ_val, β_val]), exog)
            moments = (err[:, None] * instruments).mean(axis=0)
```

(continues on next page)

(continued from previous page)

```
            objective[i, j] = moments @ moments

    return γ_grid, β_grid, objective
```

```
γ_grid, β_grid, objective = gmm_objective_surface(sim_data, n_lags=2)
log_obj = np.log10(objective + 1e-12)

fig, ax = plt.subplots()
contours = ax.contourf(γ_grid, β_grid, log_obj, levels=30, cmap="viridis")
ax.set_xlabel(r"$\gamma$")
ax.set_ylabel(r"$\beta$")
ax.plot(γ_true, β_true, "k*", ms=12, lw=2, label="true values")
ax.legend()
plt.colorbar(contours, ax=ax)
plt.tight_layout()
plt.show()
```

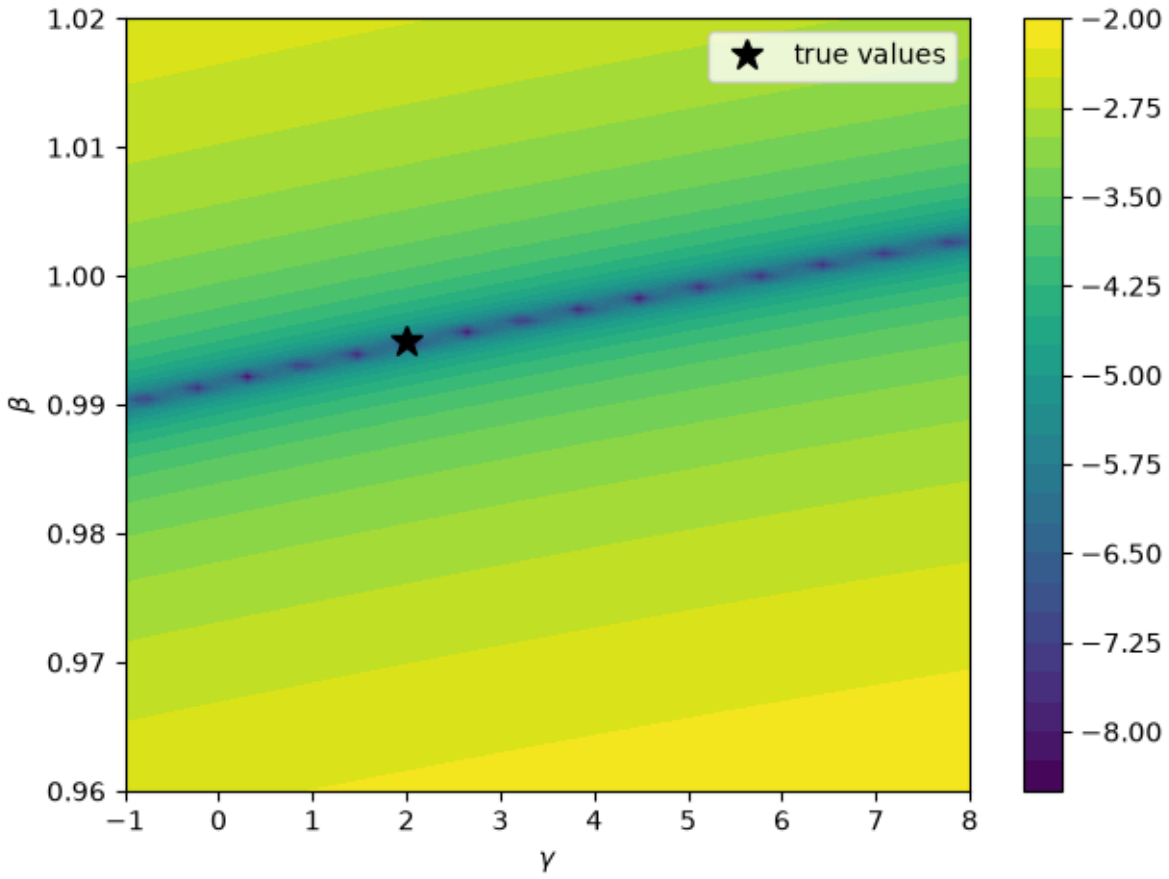


Fig. 133.1: GMM objective contour surface (simulated data)

The criterion surface may have elongated valleys where many parameter combinations fit the moments nearly equally well.

To illustrate the multi-period case from Section 2 of Hansen and Singleton [1982], we estimate the three-period Euler

restriction using overlapping-horizon returns and consumption growth, with instruments formed from one-period data dated t or earlier and the finite-order covariance appropriate for MA(2) disturbances.

```
horizon_n = 3
two_step = two_step_gmm(
    sim_data,
    n_lags=2,
    ma_order=horizon_n - 1,
    horizon=horizon_n,
)

print(f"Horizon n: {horizon_n}")
print(f"Step-2 converged: {two_step['success']}")
print(f"Step-2 gamma: {two_step['params_step2'][0]:.4f}")
print(f"Step-2 beta (one-period): {two_step['params_step2'][1]:.4f}")
print(
    f"J({two_step['j_df']}): {two_step['j_stat']:.3f}, "
    f"Prob={two_step['j_prob']:.3f}, p={two_step['j_pval']:.3f}"
)

_, exog_n, _ = build_gmm_arrays_horizon(sim_data, n_lags=2, horizon=horizon_n)
acf_n = acf(
    euler_error_horizon(two_step["params_step2"], exog_n, horizon=horizon_n),
    nlags=6,
    fft=True,
)
print("Euler-error ACF lags 1-3:", ", ".join([f"{v:.3f}" for v in acf_n[1:4]]))
```

```
Horizon n: 3
Step-2 converged: True
Step-2 gamma: 2.1095
Step-2 beta (one-period): 0.9949
J(3): 2.518, Prob=0.528, p=0.472
Euler-error ACF lags 1-3: 0.678, 0.356, 0.023
```

The low-lag ACF is consistent with the MA(2) dependence implied by the 3-period horizon.

We now run a Monte Carlo exercise with 500 replications to visualize the finite-sample distribution of $\hat{\gamma}$, $\hat{\beta}$, and the J statistic and verify that the asymptotic theory from Section 3 of Hansen and Singleton [1982] provides a reasonable approximation.

```
n_rep = 500
estimates = []
j_stats = []

for rep in range(n_rep):
    rep_data = simulate_euler_sample(
        n_obs=900,
        γ_true=γ_true,
        β_true=β_true,
        seed=rep,
    )
    rep_res = estimate_gmm(rep_data, n_lags=2, use_hac=True, maxiter=2)
    estimates.append(rep_res.params)
    j_stats.append(rep_res.jval)

estimates = np.asarray(estimates)
j_stats = np.asarray(j_stats)
```

(continues on next page)

(continued from previous page)

```
fig, axes = plt.subplots(1, 3, figsize=(15, 4))
axes[0].hist(estimates[:, 0], bins=20, edgecolor="white")
axes[0].axvline(γ_true, color="red", ls="--", lw=2)
axes[0].set_xlabel(r"$\hat{\gamma}$")
axes[1].hist(estimates[:, 1], bins=20, edgecolor="white")
axes[1].axvline(β_true, color="red", ls="--", lw=2)
axes[1].set_xlabel(r"$\hat{\beta}$")

df_j = 2 * 2 + 1 - 2
axes[2].hist(j_stats, bins=20, density=True, edgecolor="white")
grid = np.linspace(0.0, max(j_stats.max(), 1.0), 200)
axes[2].plot(grid, stats.chi2.pdf(grid, df_j), "r-", lw=2)
axes[2].set_xlabel("j-statistic")
plt.tight_layout()
plt.show()
```

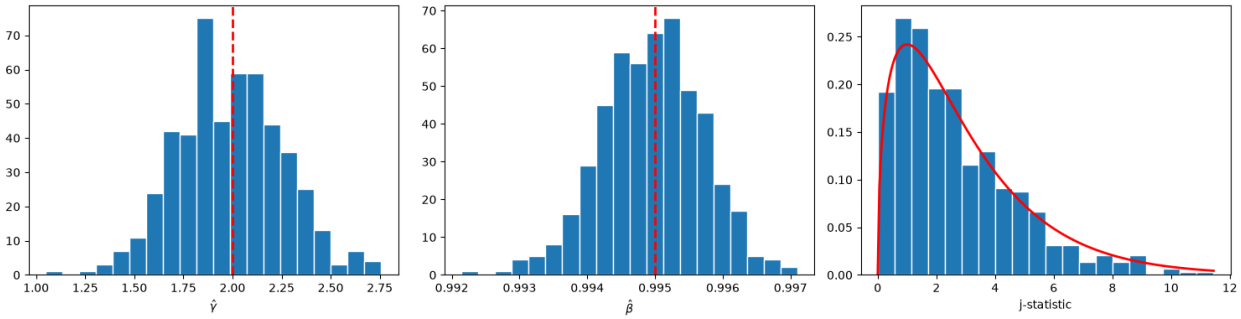


Fig. 133.2: Monte Carlo GMM sampling distributions

Both $\hat{\gamma}$ and $\hat{\beta}$ are centered near their true values, and the J histogram tracks the χ^2 density, supporting the asymptotic approximation at this sample size.

133.6 Empirical GMM estimation

We now apply GMM to observed data, following the empirical strategy of Section 5 in Hansen and Singleton [1982].

Hansen and Singleton [1982] use monthly per capita consumption of nondurables (ND) and nondurables plus services (NDS) paired with the equally-weighted (EWR) and value-weighted (VWR) aggregate stock returns from CRSP, for 1959:2 through 1978:12.

We focus on their ND+VWR specification using FRED nondurables consumption and the Ken French value-weighted market return as a proxy for CRSP, on the same 1959:2–1978:12 sample period.

Because the Ken French return is not identical to the original CRSP NYSE value-weighted return, we only want to match the paper qualitatively.

Both this lecture and the companion lecture *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns* use the same data construction.

The hidden cell below loads a monthly dataset of gross real returns and gross consumption growth from QuantEcon/data-lectures. The data are built from the FRED and Ken French data libraries by the maintenance script at `builders/hansen_singleton_1982_data.py` and read here directly from GitHub.

We first examine the raw data moments.

```
LAGS = (1, 2, 4, 6)

emp_frame, emp_data = get_estimation_data()

print(f"Mean net real return: {(emp_data[:, 0].mean() - 1.0) * 100:.3f}%")
print(f"Std net real return: {emp_data[:, 0].std() * 100:.3f}%")
print(f"Mean net consumption growth: {(emp_data[:, 1].mean() - 1.0) * 100:.3f}%")
print(f"Std net consumption growth: {emp_data[:, 1].std() * 100:.3f}%")
print(f"Std log consumption growth: {np.log(emp_data[:, 1]).std() * 100:.3f}%")
```

```
Mean net real return: 0.308%
Std net real return: 4.335%
Mean net consumption growth: 0.093%
Std net consumption growth: 0.915%
Std log consumption growth: 0.915%
```

One feature of these data is the large gap between the volatility of returns and the volatility of consumption growth.

This is again an empirical fact underlying the equity premium puzzle of Mehra and Prescott [1985]: matching the observed equity premium with CRRA preferences requires implausibly high risk aversion.

We now estimate the Euler equation using the two-step generalized instrumental variables (GIV) / GMM procedure in Hansen and Singleton [1982].

For the one-period stock-return Euler equation ($n = 1$), the disturbance is a martingale difference sequence, so the optimal weighting matrix uses the contemporaneous covariance $S_0 = E[m_t m_t^\top]$.

To match Table I, we report the paper's exponent parameter α in $E_t[\beta(C_{t+1}/C_t)^\alpha R_{t+1} - 1] = 0$.

The two-step GMM estimates of $\hat{\alpha}$ and $\hat{\beta}$ by lag length are

```
gmm_raw = run_two_step_by_lag(emp_data, lags=LAGS, horizon=1)
gmm_raw.index.name = "NLAG"

table_i = pd.DataFrame(index=gmm_raw.index)
table_i.index.name = "NLAG"
table_i[r"\hat{\alpha}"] = -gmm_raw["γ_hat"]
table_i[r"SE(\hat{\alpha})"] = gmm_raw["se_γ"]
table_i[r"\beta"] = gmm_raw["β_hat"]
table_i[r"\mathrm{SE}(\beta)"] = gmm_raw["se_β"]
table_i[r"\chi^2"] = gmm_raw["j_stat"]
table_i["DF"] = gmm_raw["j_df"]
table_i["Prob"] = gmm_raw["j_prob"]

display_table(
    table_i,
    fmt={
        r"\hat{\alpha}": "{:.4f}",
        r"SE(\hat{\alpha})": "{:.4f}",
        r"\beta": "{:.4f}",
        r"\mathrm{SE}(\beta)": "{:.4f}",
        r"\chi^2": "{:.4f}",
        "DF": "{:.0f}",
        "Prob": "{:.4f}",
    },
)
```

NLAG	$\hat{\alpha}$	$SE(\hat{\alpha})$	β	SE(β)	χ^2	DF	Prob
1	−1.1974	0.8428	0.9966	0.0026	1.7644	1	0.8159
2	−0.5494	0.7844	0.9969	0.0026	6.6667	3	0.9167
4	−0.5441	0.7040	0.9968	0.0026	9.0237	7	0.7490
6	−0.8718	0.6457	0.9973	0.0026	11.0602	11	0.5618

For comparison, Table I of Hansen and Singleton [1982] (as corrected in the 1984 *Econometrica* errata) reports the following ND+VWR values for 1959:2–1978:12:

NLAG	α	$SE(\alpha)$	β	SE(β)	χ^2	DF	Prob
1	−1.2028	0.7789	0.9976	0.0027	1.4570	1	0.7726
2	−0.5761	0.7067	0.9975	0.0027	5.8190	3	0.8792
4	−0.6565	0.6896	0.9978	0.0027	7.9230	7	0.6606
6	−0.9638	0.6425	0.9985	0.0027	10.5220	11	0.5159

They are very close to our estimates.

133.7 Summary

The GMM estimator requires only the orthogonality conditions implied by the Euler equation and a set of conditioning variables.

It does not require assumptions about the joint distribution of consumption and returns, the production technology, or any other part of the economic environment beyond the representative agent's first-order conditions.

So GMM is provides a way to estimate some objects of interest while not estimating all of the parameters that Hansen and Singleton [1983] estimated in their loglinear model of consumption growth and returns.

- If the complete model of Hansen and Singleton [1983] is correctly specified, then their maximum likelihood estimators promise to be more efficient that are the GMM estimators described in this lecture.
- The theme of estimating *something* while not estimating *everything* runs through much of Lars Peter Hansen's work. See Hansen [2014]

CHAPTER

FOUR

COMPETITIVE EQUILIBRIA WITH ARROW SECURITIES

Contents

- *Competitive Equilibria with Arrow Securities*
  - *Overview*
  - *The setting*
  - *Markov asset prices*
  - *Recursive formulation*
  - *A rational expectations equilibrium*
  - *State variable degeneracy*
  - *Computing a competitive equilibrium*
  - *Finite horizon*
  - *Python code*
  - *Examples*
  - *Concluding remarks*
  - *Related lectures*
  - *Exercises*

134.1 Overview

This lecture presents Python code for experimenting with competitive equilibria of an infinite-horizon pure exchange economy with

- heterogeneous agents,
- endowments of a single consumption good that are person-specific functions of a common Markov state,
- complete markets in one-period Arrow state-contingent securities,
- discounted expected utility preferences of a kind often used in macroeconomics and finance,
- identical preferences across agents, with a common discount factor and a constant relative risk aversion (CRRA) one-period utility function, and

- common beliefs among agents.

Differences in their endowments make individuals want to reallocate consumption goods across time and Markov states.

Identical CRRA preferences imply that equilibrium consumption shares are constant, so we can compute equilibrium prices from the aggregate endowment *before* we compute the equilibrium distribution of wealth.

We impose restrictions that allow us to **Bellmanize** competitive equilibrium prices and quantities.

We use Bellman equations to describe

- asset prices,
- continuation wealth levels for each person, and
- state-by-state natural debt limits for each person.

In the course of presenting the model we shall encounter these important ideas:

- a **resolvent operator** widely used in this class of models,
- the absence of **borrowing limits** in finite-horizon economies,
- state-by-state **borrowing limits** required in infinite-horizon economies,
- a counterpart of the **law of iterated expectations** known as a **law of iterated values**, and
- a **state variable degeneracy** that prevails within a competitive equilibrium and that opens the way to various appearances of resolvent operators.

The lecture implements a Python version of the model presented in section 9.3.3 of [Ljungqvist and Sargent, 2018].

The material is older than it might appear.

Both the sequential trading arrangement studied here and the time 0 arrangement with a complete set of history-contingent claims, as well as the proof that the two support the same allocation, appear in a paper that Kenneth Arrow read in Paris in May 1952 and published in French in 1953 [Arrow, 1964].

Arrow's sequential arrangement deserves a second name as well.

Because a household that trades securities today must act on a forecast of the prices at which goods will trade tomorrow, and because the equivalence of the two arrangements holds only when markets confirm that forecast, Arrow's sequential equilibrium is a **rational expectations** equilibrium in the sense that Muth [1961] gave the term almost a decade later.

We develop this reading, which we learned from Kihlstrom [2019], in the section on *A rational expectations equilibrium* below.

Readers will find it helpful to know the finite-state Markov asset pricing formulas of *Asset Pricing: Finite State Models* and the Markov chain concepts of *Finite Markov Chains*.

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import root

np.set_printoptions(suppress=True)
```

134.2 The setting

134.2.1 Preferences and endowments

In each period $t \geq 0$, a stochastic event $s_t \in \mathbf{S}$ is realized.

Let the history of events up until time t be denoted $s^t = [s_0, s_1, \ldots, s_{t-1}, s_t]$.

The unconditional probability of observing a particular sequence of events s^t is given by a probability measure $\pi_t(s^t)$.

For $t > \tau$, we write the probability of observing s^t conditional on the realization of s^τ as $\pi_t(s^t \mid s^\tau)$.

We assume that trading occurs after observing s_0, which we capture by setting $\pi_0(s_0) = 1$ for the initially given value of s_0.

In this lecture we follow much of macroeconomics and econometrics and assume that $\pi_t(s^t)$ is induced by a Markov process.

There are K consumers named $k = 1, \ldots, K$.

Consumer k owns a stochastic endowment of one good $y_t^k(s^t)$ that depends on the history s^t.

The history s^t is publicly observable.

Consumer k purchases a history-dependent consumption plan $c^k = \{c_t^k(s^t)\}_{t=0}^\infty$.

All consumers order consumption plans by

$$U(c^k) = \sum_{t=0}^{\infty} \sum_{s^t} \beta^t u[c_t^k(s^t)] \pi_t(s^t),$$

where $0 < \beta < 1$.

The right side equals $E_0 \sum_{t=0}^\infty \beta^t u(c_t^k)$, where E_0 is the mathematical expectation operator conditioned on s_0.

Here $u(c)$ is an increasing, twice continuously differentiable, strictly concave function of consumption $c \geq 0$ of one good.

The utility function satisfies the Inada condition

$$\lim_{c \downarrow 0} u'(c) = +\infty.$$

This condition implies that each agent chooses strictly positive consumption for every date-history pair (t, s^t) whenever the present value of its endowment is positive.

Those interior solutions allow us to confine our analysis to Euler equations that hold with equality, and they guarantee that **natural debt limits** don't bind in economies like ours with sequential trading of Arrow securities.

We adopt the assumption, routinely employed in much of macroeconomics, that consumers share probabilities $\pi_t(s^t)$ for all t and s^t.

A **feasible allocation** satisfies

$$\sum_{k=1}^{K} c_t^k(s^t) \leq \sum_{k=1}^{K} y_t^k(s^t)$$

for all t and for all s^t.

Until we reach the section on computing an equilibrium, u need only satisfy the properties listed above.

From then on we specialize to CRRA utility.

134.3 Markov asset prices

Before setting up the equilibrium, we summarize formulas for computing asset prices in a Markov setting.

These formulas are developed at greater length in *Asset Pricing: Finite State Models*.

The setup assumes the following infrastructure:

- Markov states $s \in \mathbf{S} = \{\bar{s}_1, \ldots, \bar{s}_n\}$ governed by an n-state Markov chain with transition probability
$$P_{ij} = \Pr\left\{s_{t+1} = \bar{s}_j \mid s_t = \bar{s}_i\right\};$$
- a collection $h = 1, \ldots, H$ of assets, where asset h pays $d^h(s)$ in state s, so that d^h is an $n \times 1$ vector; and
- an $n \times n$ pricing kernel Q for one-period Arrow securities, where Q_{ij} is the price at time t in state $s_t = \bar{s}_i$ of one unit of consumption delivered at time $t+1$ if $s_{t+1} = \bar{s}_j$.

The price in state $\bar{s}_i$ of a one-period risk-free bond that pays one unit of consumption in every state is $\sum_j Q_{ij}$.

The gross rate of return on that bond is therefore

$$R_i = \Big(\sum_j Q_{ij}\Big)^{-1}.$$

134.3.1 An exogenous pricing kernel

For now we take the pricing kernel Q as exogenous, that is, determined outside the model.

Two examples are

- $Q = \beta P$, where $\beta \in (0, 1)$, and
- $Q_{ij} = m_{ij} P_{ij}$, where $m_{ij} > 0$ is the value of a **stochastic discount factor** when the Markov state moves from $\bar{s}_i$ to $\bar{s}_j$.

The second example multiplies P element by element, not as a matrix product.

We now describe the prices of two types of assets.

The first is a **cum-dividend** stock that entitles its owner to the time t dividend and to the option to sell the asset at time $t+1$.

Its price satisfies $p^h(\bar{s}_i) = d^h(\bar{s}_i) + \sum_j Q_{ij} p^h(\bar{s}_j)$, so the vector p^h satisfies $p^h = d^h + Q p^h$.

Provided every eigenvalue of Q has modulus less than one, this implies

$$p^h = (I - Q)^{-1} d^h.$$

The second is an **ex-dividend** stock purchased at the end of time t, which entitles its owner to the time $t+1$ dividend and to the option to sell the stock at time $t+1$.

Its price is

$$p^h = (I - Q)^{-1} Q d^h.$$

> **Note**
>
> The matrix geometric sum $(I - Q)^{-1} = I + Q + Q^2 + \cdots$ is an example of a **resolvent operator**.
>
> It converges when the spectral radius of Q is less than one.

Below we describe an equilibrium model with trading of one-period Arrow securities in which the pricing kernel is endogenous.

In constructing that model, we'll repeatedly encounter formulas that remind us of these asset pricing formulas.

134.3.2 Multi-step transition probabilities and pricing kernels

The (i, j) component of the j-step-ahead transition matrix P^j is

$$\Pr(s_{t+j} = \bar{s}_{j'} \mid s_t = \bar{s}_i) = (P^j)_{ij'}.$$

To keep the notation light below, we write $P_j(s_{t+j} \mid s_t)$ for these j-step transition probabilities, so that P_j is represented by the matrix P^j.

In the same way, the price at time t in state s_t of one unit of consumption delivered at time $t + j$ in state s_{t+j} is $Q_j(s_{t+j} \mid s_t)$, represented by the matrix Q^j.

We'll use these objects to state a useful property of asset pricing theory.

134.3.3 Laws of iterated expectations and iterated values

A **law of iterated values** has a mathematical structure that parallels a **law of iterated expectations**.

We can describe its structure readily in the Markov setting of this lecture.

Recall the following recursion satisfied by j-step-ahead transition probabilities for our finite-state Markov chain:

$$P_j(s_{t+j} \mid s_t) = \sum_{s_{t+1}} P_{j-1}(s_{t+j} \mid s_{t+1}) P(s_{t+1} \mid s_t).$$

We can use this recursion to verify the law of iterated expectations applied to the conditional expectation of a random variable $d(s_{t+j})$ conditioned on s_t:

$$\begin{aligned}
E\big[E[d(s_{t+j}) \mid s_{t+1}] \,\big|\, s_t\big] &= \sum_{s_{t+1}} \left[\sum_{s_{t+j}} d(s_{t+j}) P_{j-1}(s_{t+j} \mid s_{t+1}) \right] P(s_{t+1} \mid s_t) \\
&= \sum_{s_{t+j}} d(s_{t+j}) \left[\sum_{s_{t+1}} P_{j-1}(s_{t+j} \mid s_{t+1}) P(s_{t+1} \mid s_t) \right] \\
&= \sum_{s_{t+j}} d(s_{t+j}) P_j(s_{t+j} \mid s_t) \\
&= E[d(s_{t+j}) \mid s_t].
\end{aligned}$$

The pricing kernel for j-step-ahead Arrow securities satisfies the recursion

$$Q_j(s_{t+j} \mid s_t) = \sum_{s_{t+1}} Q_{j-1}(s_{t+j} \mid s_{t+1}) Q(s_{t+1} \mid s_t).$$

The time t **value** in Markov state s_t of a time $t + j$ payout $d(s_{t+j})$ is

$$W[d(s_{t+j}) \mid s_t] = \sum_{s_{t+j}} d(s_{t+j}) Q_j(s_{t+j} \mid s_t).$$

The **law of iterated values** states that

$$W\big[W[d(s_{t+j}) \mid s_{t+1}] \,\big|\, s_t\big] = W[d(s_{t+j}) \mid s_t].$$

We verify it with a string of equalities that are counterparts to those we used to verify the law of iterated expectations:

$$\begin{aligned} W[W[d(s_{t+j}) \mid s_{t+1}] \mid s_t] &= \sum_{s_{t+1}} \left[\sum_{s_{t+j}} d(s_{t+j}) Q_{j-1}(s_{t+j} \mid s_{t+1}) \right] Q(s_{t+1} \mid s_t) \\ &= \sum_{s_{t+j}} d(s_{t+j}) \left[\sum_{s_{t+1}} Q_{j-1}(s_{t+j} \mid s_{t+1}) Q(s_{t+1} \mid s_t) \right] \\ &= \sum_{s_{t+j}} d(s_{t+j}) Q_j(s_{t+j} \mid s_t) \\ &= W[d(s_{t+j}) \mid s_t]. \end{aligned}$$

134.4 Recursive formulation

Following section 9.3.3 of [Ljungqvist and Sargent, 2018], we now set up a competitive equilibrium of a pure exchange economy with complete markets in one-period Arrow securities.

When endowments $y^k(s)$ are all functions of a common Markov state s, the pricing kernel takes the form $Q(s' \mid s)$, the price of one unit of consumption in state s' at date $t+1$ when the Markov state at date t is s.

This lets us formulate a consumer's optimization problem recursively.

Consumer k's state at time t is its financial wealth a_t^k and the Markov state s_t.

Let $v^k(a, s)$ be the optimal value of consumer k's problem starting from state (a, s).

Thus $v^k(a, s)$ is the maximum expected discounted utility that consumer k with current financial wealth a can attain in Markov state s.

The optimal value function satisfies the Bellman equation

$$v^k(a, s) = \max_{c, \hat{a}(s')} \left\{ u(c) + \beta \sum_{s'} v^k[\hat{a}(s'), s'] \pi(s' \mid s) \right\},$$

where the maximization is subject to the budget constraint

$$c + \sum_{s'} \hat{a}(s') Q(s' \mid s) \leq y^k(s) + a$$

and the constraints

$$\begin{aligned} c &\geq 0, \\ -\hat{a}(s') &\leq \bar{A}^k(s'), \quad \forall s' \in \mathbf{S}. \end{aligned}$$

The second set of constraints is a collection of state-by-state debt limits.

The value function and decision rule that solve the Bellman equation depend on the pricing kernel $Q(\cdot \mid \cdot)$ because it appears in the budget constraint.

The first-order conditions for the problem on the right of the Bellman equation, together with a Benveniste–Scheinkman formula, imply

$$Q(s_{t+1} \mid s_t) = \frac{\beta u'(c_{t+1}^k) \pi(s_{t+1} \mid s_t)}{u'(c_t^k)},$$

where it is understood that $c_t^k = c^k(s_t)$ and $c_{t+1}^k = c^k(s_{t+1})$.

134.4.1 Recursive competitive equilibrium

A **recursive competitive equilibrium** is an initial distribution of wealth $\vec{a}_0$, a set of borrowing limits $\{\bar{A}^k(s)\}_{k=1}^K$, a pricing kernel $Q(s' \mid s)$, sets of value functions $\{v^k(a,s)\}_{k=1}^K$, and decision rules $\{c^k(s), \hat{a}^k(s)\}_{k=1}^K$ such that

1. the state-by-state borrowing limits satisfy the recursion
$$\bar{A}^k(s) = y^k(s) + \sum_{s'} Q(s' \mid s)\bar{A}^k(s');$$
2. for all k, given a_0^k, $\bar{A}^k(s)$, and the pricing kernel, the value functions and decision rules solve the consumers' problems;
3. for all realizations of $\{s_t\}_{t=0}^\infty$, the consumption and asset portfolios $\{\{c_t^k, \{\hat{a}_{t+1}^k(s')\}_{s'}\}_k\}_t$ satisfy $\sum_k c_t^k = \sum_k y^k(s_t)$ and $\sum_k \hat{a}_{t+1}^k(s') = 0$ for all t and s'; and
4. the initial financial wealth vector $\vec{a}_0$ satisfies $\sum_{k=1}^K a_0^k = 0$.

The third condition asserts that goods markets clear and that there are zero net aggregate claims in all Markov states.

The fourth condition asserts that the economy is closed and starts from a situation in which there are zero net aggregate claims.

134.5 A rational expectations equilibrium

The equilibrium that we have just defined asks each consumer to choose a portfolio $\hat{a}(s')$ of claims today while knowing the pricing kernel $Q(s' \mid s)$ at which those claims trade.

But the value of a portfolio to a consumer depends on what goods will cost tomorrow.

So an equilibrium of this kind embeds a forecast, and it is an equilibrium only when markets confirm that forecast.

That is what Radner [1972] later formalized as an equilibrium of plans, prices, and price expectations.

134.5.1 What a single good conceals

In our economy the only price that a sequential trader must forecast is the one-period pricing kernel $Q(s' \mid s)$.

Within a date and Markov state there are no relative prices left to forecast, because one good is the numeraire by default.

In Arrow's economy with $C > 1$ goods, a household choosing a portfolio must forecast the entire vector of spot prices $\bar{p}_s = (\bar{p}_{s1}, \ldots, \bar{p}_{sC})$ that will prevail in each state s, because only those prices tell it what a dollar delivered in state s will be worth.

The equivalence of the two trading arrangements then holds *if and only if those forecasts are correct*, state by state.

That is the precise sense in which Arrow's sequential equilibrium is a rational expectations equilibrium, and it is a requirement that our one-good formulation satisfies vacuously.

Exercise 134.13.4 asks you to verify the equivalence in a two-period economy with several goods, where the requirement has content.

134.5.2 Securities denominated in a unit of account

That Arrow's equivalence turns *only* on the correctness of those forecasts is itself a consequence of his having denominated securities in a unit of account.

Because a unit of security s pays one dollar if state s occurs, the set of state-contingent wealth distributions that a portfolio can deliver is settled before any forecast is made.

Hart [1975] asked what happens when securities instead pay *goods*.

If security f delivers the bundle $a_f(s)$ in state s, its payoff in dollars is $\bar{p}_s \cdot a_f(s)$, so the forecast $\bar{p}$ determines the span of the asset market, the span determines consumers' budget sets, and those budget sets determine the prices being forecast.

Completeness becomes an equilibrium object rather than an assumption.

Two things can then go wrong that cannot go wrong in our one-good economy.

There can be several equilibria, all of them with correct forecasts, that differ in which markets are *effectively* open and that are strictly Pareto ranked.

And there can be no equilibrium at all, even with strictly concave utility functions, strictly positive endowments, and no transaction costs.

Exercise 134.13.5 works through Hart's examples of both.

Note

Hart [1975] also delimits when things go right.

An equilibrium of the kind defined by Radner [1972] is Pareto optimal when the market structure is complete, and it is Pareto optimal relative to the set of such equilibria when there is one good and the market structure is complete up to the next-to-last date, which generalizes a result of Diamond [1967].

Outside those two cases, utility functions and endowments can generally be chosen so that an equilibrium fails even that weaker test.

Both of Hart's examples are knife edge, so small perturbations of preferences, endowments, or dividends destroy them.

134.5.3 Beliefs about nature versus expectations of prices

Arrow drew a distinction here that is worth preserving.

His consumers act on *subjective* probabilities that are allowed to differ across people, and nothing in his theorems requires agreement about how likely the states are.

What his sequential arrangement does require is that consumers agree, and be right, about something else: the prices that will rule in each state.

Beliefs about exogenous nature may be heterogeneous; expectations about endogenous prices may not.

This lecture imposes the stronger assumption of common beliefs, which lets us write a single probability $\pi(s' \mid s)$ in each consumer's Euler equation.

Heterogeneous Beliefs and Bubbles and *Heterogeneous Beliefs and Financial Markets* relax that assumption while keeping the expectations of prices correct.

134.5.4 Economizing on markets

Arrow observed that the sequential arrangement *permits economizing on markets*.

In our infinite-horizon economy, time 0 trading requires a market in a claim for every date and every history, while sequential trading requires only n one-period markets at each date.

In the two-period, S-state, C-good economy of Exercise 134.13.4, time 0 trading requires SC contingent claims markets, while the sequential arrangement requires S security markets followed by C spot markets in whichever state occurs.

No counting of markets, however, would reveal the extra assumption that the sequential arrangement needs, which is that the forecasts of spot prices are correct.

Note

Kihlstrom [2019] remarks that when Arrow presented this material in 1952, the dynamic programming machinery that makes the equivalence transparent did not yet exist, because Bellman was in the midst of creating it.

The step that decomposes a household's once-and-for-all problem into a portfolio choice followed by a sequence of within-period problems is exactly the argument that we used to Bellmanize the equilibrium above.

134.6 State variable degeneracy

[Ljungqvist and Sargent, 2018] and *Cass-Koopmans Competitive Equilibrium* describe a different timing protocol, in which there is a complete menu of history-contingent claims on consumption at all dates and all trades occur once and for all at time 0.

For the allocation and pricing kernel of a recursive competitive equilibrium to coincide with those of that time 0 arrangement, we must impose $a_0^k = 0$ for $k = 1, \ldots, K$.

That initial condition ensures that at time 0 the present value of each consumer's consumption equals the present value of its endowment stream, which is the single budget constraint of the time 0 arrangement.

Starting the system with $a_0^k = 0$ for all k has a striking implication that we call **state variable degeneracy**.

Although two state variables a and s appear in the value function $v^k(a, s)$, within a recursive competitive equilibrium that starts from $a_0^k = 0$ for all k at initial Markov state s_0, two outcomes prevail:

- financial wealth a_t^k is an exact function of the Markov state s_t, which we compute below, and
- $a_t^k = 0$ for all k whenever the Markov state s_t returns to s_0.

The first finding asserts that within a competitive equilibrium the exogenous Markov state is all we require to track an individual, because financial wealth is redundant.

The second finding asserts that each household returns to the zero financial wealth with which it began life whenever the Markov state returns to its initial value.

That happens infinitely often if s_0 is a recurrent state of the Markov chain, but it need not happen at all if s_0 is transient; see *Finite Markov Chains*.

This outcome depends critically on there being complete markets in Arrow securities.

For example, it does not prevail in the incomplete markets setting of *The Aiyagari Model*, where a household's wealth depends on its whole history of shocks.

134.7 Computing a competitive equilibrium

Now we are ready to do some fun calculations.

We find it useful to think in terms of analytical **inputs** into and **outputs** from our general equilibrium theorizing.

134.7.1 Inputs and outputs

The inputs are

- Markov states $s \in \mathbf{S} = \{\bar{s}_1, \ldots, \bar{s}_n\}$ governed by an n-state Markov chain with transition matrix P;
- K vectors of individual endowments y^k, each of dimension $n \times 1$ with components $y^k(\bar{s}_i)$;
- the $n \times 1$ vector of aggregate endowments $y(s) \equiv \sum_{k=1}^K y^k(s)$; and
- preferences given by the common utility functional $E_0 \sum_{t=0}^\infty \beta^t u(c_t^k)$ with discount factor $\beta \in (0,1)$ and CRRA one-period utility function

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \qquad u'(c) = c^{-\gamma}.$$

Feasibility requires

$$c(s) = \sum_{k=1}^K c^k(s) \leq y(s).$$

The outputs are

- an $n \times n$ pricing kernel Q for one-period Arrow securities;
- the aggregate allocation, which in a pure exchange economy is $c(s) = y(s)$;
- a $K \times 1$ distribution of wealth α with $\alpha_k \geq 0$ and $\sum_{k=1}^K \alpha_k = 1$; and
- K vectors of individual consumptions c^k, each of dimension $n \times 1$.

134.7.2 The pricing kernel

For any agent $k \in \{1, \ldots, K\}$, at the equilibrium allocation, the one-period Arrow securities pricing kernel satisfies

$$Q_{ij} = \beta \left(\frac{c^k(\bar{s}_j)}{c^k(\bar{s}_i)} \right)^{-\gamma} P_{ij}.$$

This follows from agent k's first-order necessary conditions.

Because all agents face the same pricing kernel, the Euler equations of any two agents k and m imply

$$\left(\frac{c^k(\bar{s}_j)}{c^k(\bar{s}_i)} \right)^{-\gamma} = \left(\frac{c^m(\bar{s}_j)}{c^m(\bar{s}_i)} \right)^{-\gamma} \quad \text{whenever } P_{ij} > 0.$$

So the ratio $c^k(s)/c^m(s)$ is the same in any two states connected by a positive transition probability, and hence in every state reachable from s_0.

Consumption shares are therefore constant, and feasibility gives

$$c^k(s) = \alpha_k c(s) = \alpha_k y(s)$$

for a **distribution of wealth** α that satisfies $\alpha_k \geq 0$ and $\sum_{k=1}^{K} \alpha_k = 1$.

> **Note**
>
> Identical CRRA preferences also satisfy the conditions for **Gorman aggregation**, since Engel curves are linear with a common slope, so a representative consumer exists.
>
> The constancy of consumption shares, however, follows directly from the Euler equations above.

This means that we can compute the pricing kernel from

$$Q_{ij} = \beta \left(\frac{y_j}{y_i} \right)^{-\gamma} P_{ij}. \tag{134.1}$$

This is the pricing kernel of the Lucas tree economy studied in *Asset Pricing: Finite State Models*, in which a representative consumer eats the aggregate endowment.

The pricing kernel Q does not depend on the vector α.

Key finding: We can compute competitive equilibrium **prices** prior to computing a **distribution of wealth**.

The wealth distribution α is not arbitrary.

It is pinned down by the initial condition $a_0^k = 0$, as we show below.

Formula (134.1) has a useful matrix form.

Let $D = \text{diag}(u'(y_1), \ldots, u'(y_n))$.

Then $Q = \beta D^{-1} P D$, so Q is similar to βP.

Its eigenvalues are β times those of P, and because P is a stochastic matrix, the spectral radius of Q equals $\beta < 1$.

This guarantees that the resolvent $(I - Q)^{-1}$ used below exists.

The factorization is an instance of the **transition independence** structure exploited in *The Recovery Theorem*, a connection we pursue in Exercise 134.13.2.

134.7.3 Natural debt limits

Having computed an equilibrium pricing kernel Q, we can compute several **values** that are required to pose or represent the solution of an individual household's optimization problem.

For each individual k, let $\bar{A}^k$ be the $n \times 1$ vector with components $\bar{A}^k(\bar{s}_i)$.

The recursion in the definition of equilibrium implies

$$\bar{A}^k = [I - Q]^{-1} y^k. \tag{134.2}$$

In a competitive equilibrium of an **infinite-horizon** economy with sequential trading of one-period Arrow securities, $\bar{A}^k(s)$ is a state-by-state limit on the quantity of one-period Arrow securities paying off in state s at time $t+1$ that individual k can issue at time t.

These are often called **natural debt limits**.

They equal the maximum amount that individual k can repay in state s even if it consumes nothing forevermore.

Note

If utility satisfies an Inada condition at zero consumption, or if consumption is simply required to be nonnegative, then a **finite-horizon** economy with sequential trading of one-period Arrow securities needs no natural debt limits.

See the section on a finite-horizon economy below.

134.7.4 Continuation wealth and optimal portfolios

Continuation wealth plays an important role in Bellmanizing a competitive equilibrium with sequential trading of a complete set of one-period Arrow securities.

For each individual k, let ψ^k be the $n \times 1$ vector with components $\psi^k(\bar{s}_i)$, the financial wealth that consumer k holds when the Markov state is $\bar{s}_i$.

Continuation wealth satisfies

$$\psi^k = [I - Q]^{-1} [\alpha_k y - y^k] . \tag{134.3}$$

To see why, note that with consumption $c^k = \alpha_k y$ and a portfolio $\hat{a}^k(s') = \psi^k(s')$, the budget constraint holds with equality in every state exactly when $\psi^k = \alpha_k y - y^k + Q\psi^k$.

Summing over k shows that $\sum_{k=1}^K \psi^k = 0_{n\times 1}$, so Arrow security markets clear.

A nifty feature of the model is that an optimal portfolio of a type k agent equals the continuation wealth that we just computed.

Thus, agent k's purchases of Arrow securities that pay off next period depend only on next period's Markov state and equal

$$\hat{a}^k(s) = \psi^k(s), \quad s \in \{\bar{s}_1, \ldots, \bar{s}_n\}. \tag{134.4}$$

134.7.5 The equilibrium wealth distribution

With the initial state being a particular state $s_0 \in \{\bar{s}_1, \ldots, \bar{s}_n\}$, we must have

$$\psi^k(s_0) = 0, \quad k = 1, \ldots, K,$$

so that every agent starts debt-free and holding no financial assets.

This means that the equilibrium distribution of wealth satisfies

$$\alpha_k = \frac{V_z y^k}{V_z y}, \tag{134.5}$$

where $V \equiv [I - Q]^{-1}$ and V_z is the row of V corresponding to the initial state s_0.

Since $\sum_{k=1}^K V_z y^k = V_z y$, we have $\sum_{k=1}^K \alpha_k = 1$.

Comparing (134.5) with (134.2) gives a revealing interpretation,

$$\alpha_k = \frac{\bar{A}^k(s_0)}{\sum_{m=1}^K \bar{A}^m(s_0)}.$$

Each consumer's share of aggregate consumption equals its share of the value, in the initial state, of the aggregate endowment.

Because α depends on s_0 through V_z, the same economy started in different Markov states delivers different wealth distributions.

134.7.6 Value functions

We can also compute optimal value functions in a competitive equilibrium with trades in a complete set of one-period state-contingent Arrow securities.

Call the optimal value function of consumer k the $n \times 1$ vector J^k.

For the infinite-horizon economy now under study,

$$J^k = (I - \beta P)^{-1} u(\alpha_k y),$$

where $u(\alpha_k y)$ is the $n \times 1$ vector with components $u(\alpha_k y_i)$.

134.7.7 Summary of the algorithm

Here is the logical flow of an algorithm to compute a competitive equilibrium:

1. compute Q from the aggregate endowment using formula (134.1);
2. compute the distribution of wealth α from formula (134.5);
3. using α, assign each consumer k the share α_k of the aggregate endowment in each state;
4. compute continuation wealths from the α-dependent formula (134.3);
5. set agent k's portfolio equal to its continuation wealth state by state, as in (134.4); and
6. compute value functions J^k.

134.8 Finite horizon

We now describe a finite-horizon version of the economy that operates for $T+1$ periods $t \in \mathbf{T} = \{0, 1, \ldots, T\}$.

We'll want time-dependent counterparts of the objects described above, with one important exception: we won't need **borrowing limits**.

- Borrowing limits aren't required in a finite-horizon economy in which the one-period utility function $u(c)$ satisfies an Inada condition that sends the marginal utility of consumption to infinity as consumption approaches zero.
- Nonnegativity of consumption at all $t \in \mathbf{T}$ automatically limits borrowing, because no one can end period T in debt.

For each individual k and date t, let ψ_t^k be the $n \times 1$ vector of continuation wealths.

At the terminal date there is no future to finance, so $\psi_T^k = \alpha_k y - y^k$.

Working backward with the budget constraint $\psi_t^k = \alpha_k y - y^k + Q\psi_{t+1}^k$ gives

$$\psi_t^k = \left[I + Q + Q^2 + \cdots + Q^{T-t}\right]\left[\alpha_k y - y^k\right], \quad t = 0, 1, \ldots, T. \tag{134.6}$$

As before, $\sum_{k=1}^K \psi_t^k = 0_{n\times 1}$ for all $t \in \mathbf{T}$.

With the initial state being a particular state s_0, we must have

$$\psi_0^k(s_0) = 0, \quad k = 1, \ldots, K,$$

which means the equilibrium distribution of wealth satisfies

$$\alpha_k = \frac{V_z y^k}{V_z y}, \tag{134.7}$$

where now

$$V = [I + Q + Q^2 + \cdots + Q^T] \tag{134.8}$$

and V_z is the row of V corresponding to the initial state s_0.

Note

In the finite-horizon economy, continuation wealth depends on calendar time as well as on the Markov state.

The initial condition sets $\psi_0^k(s_0) = 0$, but when the Markov state returns to s_0 at a later date t, fewer periods remain, the geometric sum in (134.6) is truncated sooner, and in general $\psi_t^k(s_0) \neq 0$.

The strong form of state variable degeneracy, in which wealth is a function of the Markov state alone, is special to the infinite horizon.

Exercise 134.13.3 explores this.

To compute a competitive equilibrium with Arrow securities in the finite-horizon Markov economy,

1. compute Q from the aggregate endowment using formula (134.1);
2. compute the distribution of wealth α from formulas (134.7) and (134.8);
3. using α, assign each consumer k the share α_k of the aggregate endowment in each state;
4. compute continuation wealths from formula (134.6); and
5. set agent k's portfolio equal to its continuation wealth state by state.

The value function of consumer k at time t is

$$J_t^k = [I + \beta P + \cdots + (\beta P)^{T-t}]\, u(\alpha_k y).$$

134.9 Python code

We now create a Python class to compute the objects that comprise a competitive equilibrium with sequential trading of one-period Arrow securities.

The class handles both infinite-horizon economies and finite-horizon economies indexed by horizon T.

Every geometric sum in the lecture has the form $I + M + \cdots$, with $M = Q$ for prices and wealth and $M = \beta P$ for values, so a single helper method computes them all.

In the finite-horizon case the helper works backward with the recursion $S_t = I + MS_{t+1}$, starting from $S_T = I$.

With $M = Q$, this recursion is the law of iterated values at work: the time t value of payouts from t through T is the payout at t plus the time t value of the time $t + 1$ value of the remaining payouts.

With $M = \beta P$, it is the law of iterated expectations applied to discounted utility.

The class also has a method that prices an asset with dividend vector d by applying these sums, which delivers the cum-dividend price $p = d + Qd + Q^2 d + \cdots$ and the ex-dividend price $p - d$.

In the finite-horizon case, arrays that depend on time are ordered from $t = 0$ to $t = T$, so that `ψ[t]` is ψ_t and `J[t]` is J_t.

In the infinite-horizon case they have a single leading element.

```
class RecurCompetitive:
    """
    A competitive equilibrium with complete markets in one-period
    Arrow securities.

    Parameters
    ----------
    s : array of length n
        Markov states
    P : n x n array
        Markov transition matrix
    ys : n x K array
        endowments, with column k holding agent k's endowment
    γ : float
        coefficient of relative risk aversion
    β : float
        discount factor
    T : int or None
        time horizon, None for an infinite horizon
    """

    def __init__(self, s, P, ys, γ=0.5, β=0.98, T=None):

        self.s, self.P, self.ys = s, P, ys
        self.γ, self.β, self.T = γ, β, T
        self.n, self.K = ys.shape
        self.y = ys.sum(axis=1)                  # aggregate endowment

        self.Q = self.pricing_kernel()
        self.PRF = self.Q.sum(axis=1)            # price of a risk-free bond
        self.R = 1 / self.PRF                    # gross risk-free rate

        # V[t] = I + Q + ... + Q^(T-t), or [(I - Q)^(-1)] if T is None
        self.V = self.geometric_sums(self.Q)

        # time-0 values of endowments, the natural debt limits
        self.A = self.asset_price(ys)

    def u(self, c):
        "CRRA utility"
        return c ** (1 - self.γ) / (1 - self.γ)

    def u_prime(self, c):
        "Marginal utility"
        return c ** (-self.γ)

    def pricing_kernel(self):
        "Pricing kernel Q from equation (eq:Qformula)"
        mu = self.u_prime(self.y)
        return self.β * self.P * mu[None, :] / mu[:, None]

    def geometric_sums(self, M):
        """
        Return [(I - M)^(-1)] if T is None; otherwise return the sequence
        S[0], ..., S[T] with S[t] = I + M + ... + M^(T-t).
        """
        n, T = self.n, self.T
```

(continues on next page)

(continued from previous page)

```
        if T is None:
            return np.linalg.inv(np.eye(n) - M)[None, :, :]
        S = np.empty((T+1, n, n))
        S[T] = np.eye(n)
        for t in range(T-1, -1, -1):
            S[t] = np.eye(n) + M @ S[t+1]      # law of iterated values
        return S

    def asset_price(self, d, ex_dividend=False):
        """
        Time-0 price of an asset with dividend vector d (n or n x K):
        cum-dividend p = d + Q d + Q^2 d + ..., or ex-dividend p - d.
        """
        p = self.V[0] @ d
        return p - d if ex_dividend else p

    def wealth_distribution(self, s0_idx):
        "Wealth distribution α when the initial state has index s0_idx"
        self.s0_idx = s0_idx
        V_z = self.V[0, s0_idx, :]
        self.α = V_z @ self.ys / (V_z @ self.y)
        return self.α

    def continuation_wealths(self):
        "Continuation wealths ψ, with ψ[t, i, k] = ψ_t^k(s_i)"
        excess = np.outer(self.y, self.α) - self.ys     # α_k y - y^k
        self.ψ = self.V @ excess
        return self.ψ

    def value_functions(self):
        "Value functions J, with J[t, i, k] = J_t^k(s_i)"
        flow = self.u(np.outer(self.y, self.α))         # u(α_k y)
        self.J = self.geometric_sums(self.β * self.P) @ flow
        return self.J
```

134.10 Examples

We'll use our code to construct equilibrium objects in several example economies.

Our first several examples are infinite-horizon economies.

Our final example is a finite-horizon economy.

Unless we say otherwise, examples use the default parameter values $\gamma = 0.5$ and $\beta = 0.98$.

134.10.1 Example 1: a constant aggregate endowment

Two agents have perfectly negatively correlated endowments, so the aggregate endowment is constant.

```
s = np.array([0, 1])
P = np.array([[.5, .5],
              [.5, .5]])

ys = np.empty((2, 2))
ys[:, 0] = 1 - s        # agent 1
ys[:, 1] = s            # agent 2

ex1 = RecurCompetitive(s, P, ys)
```

```
print("aggregate endowment y =", ex1.y)
print("pricing kernel Q = \n", ex1.Q)
print("risk-free rate R =", ex1.R)
print("natural debt limits A = \n", ex1.A)
```

```
aggregate endowment y = [1. 1.]
pricing kernel Q =
 [[0.49 0.49]
 [0.49 0.49]]
risk-free rate R = [1.02040816 1.02040816]
natural debt limits A =
 [[25.5 24.5]
 [24.5 25.5]]
```

Because the aggregate endowment is constant, marginal utility is constant, so $Q = \beta P$ and the risk-free rate is β^{-1} in both states.

```
# initial state is state 1
print(f'α = {ex1.wealth_distribution(s0_idx=0)}')
print(f'ψ = \n{ex1.continuation_wealths()}')
print(f'J = \n{ex1.value_functions()}')
print(f'share of natural debt limits in s0: {ex1.A[0] / ex1.A[0].sum()}')
```

```
α = [0.51 0.49]
ψ =
[[-0.  0.]
  [ 1. -1.]]]
J =
[[[71.41428429 70.        ]
  [71.41428429 70.        ]]]
share of natural debt limits in s0: [0.51 0.49]
```

When the economy starts in state 1, agent 1 receives slightly more than half of aggregate consumption.

Its endowment arrives in the initial period, and consumption received sooner is discounted less.

As the last line confirms, each agent's consumption share equals its share of the natural debt limits in the initial state.

In state 2, where agent 1 has no endowment, agent 1 holds financial wealth of one unit, which finances its consumption, and agent 2 owes exactly that amount.

```
# initial state is state 2
print(f'α = {ex1.wealth_distribution(s0_idx=1)}')
```

(continues on next page)

(continued from previous page)

```
print(f'ψ = \n{ex1.continuation_wealths()}')
print(f'J = \n{ex1.value_functions()}')
```

```
α = [0.49 0.51]
ψ =
[[[-1.  1.]
  [ 0. -0.]]]
J =
[[[70.         71.41428429]
  [70.         71.41428429]]]
```

Starting in state 2 simply swaps the roles of the two agents.

134.10.2 Example 2: a fluctuating aggregate endowment

Now agent 1 has a constant endowment while agent 2's endowment fluctuates, so the aggregate endowment fluctuates.

```
s = np.array([1, 2])
P = np.array([[.5, .5],
              [.5, .5]])

ys = np.empty((2, 2))
ys[:, 0] = 1.5          # agent 1
ys[:, 1] = s            # agent 2

ex2 = RecurCompetitive(s, P, ys)

print("aggregate endowment y =", ex2.y)
print("pricing kernel Q = \n", ex2.Q)
print("risk-free rate R =", ex2.R)
print("natural debt limits A = \n", ex2.A)
```

```
aggregate endowment y = [2.5 3.5]
pricing kernel Q =
 [[0.49       0.41412558]
 [0.57977582 0.49      ]]
risk-free rate R = [1.10604104 0.93477529]
natural debt limits A =
 [[69.30941886 66.91255848]
 [81.73318641 79.98879094]]
```

The pricing kernels in examples 1 and 2 differ because the aggregate endowment is constant in example 1 but differs across states in example 2.

We can check two off-diagonal entries of Q directly against formula (134.1).

```
print(ex2.β * ex2.u_prime(3.5) / ex2.u_prime(2.5) * ex2.P[0, 1], ex2.Q[0, 1])
print(ex2.β * ex2.u_prime(2.5) / ex2.u_prime(3.5) * ex2.P[1, 0], ex2.Q[1, 0])
```

```
0.4141255848169731 0.4141255848169731
0.5797758187437624 0.5797758187437624
```

A claim to consumption in the high-endowment state is cheap, because marginal utility is low there.

The risk-free rate is high in the low-endowment state, where consumption is expected to rise, and low in the high-endowment state, where consumption is expected to fall.

Now let's price some risky assets with the formulas $p^h = (I - Q)^{-1} d^h$ and $p^h = (I - Q)^{-1} Q d^h$ from the section on Markov asset prices.

We price a Lucas tree, which pays the aggregate endowment as its dividend, and claims to each agent's endowment stream.

```
p_tree = ex2.asset_price(ex2.y)
p_tree_ex = ex2.asset_price(ex2.y, ex_dividend=True)

print("cum-dividend tree price     p =", p_tree)
print("ex-dividend tree price      p =", p_tree_ex)
print("price-dividend ratio (ex)     =", p_tree_ex / ex2.y)
print("Bellman residual |p - d - Qp| =",
      np.abs(p_tree - ex2.y - ex2.Q @ p_tree).max())
print("claims to endowments = \n", ex2.asset_price(ex2.ys))
```

```
cum-dividend tree price     p = [136.22197734 161.72197734]
ex-dividend tree price      p = [133.72197734 158.22197734]
price-dividend ratio (ex)     = [53.48879094 45.20627924]
Bellman residual |p - d - Qp| = 0.0
claims to endowments =
 [[69.30941886 66.91255848]
 [81.73318641 79.98879094]]
```

The cum-dividend price satisfies the one-step Bellman equation $p = d + Qp$ to machine precision.

The ex-dividend price-dividend ratio is higher in the low-endowment state, where the current dividend is low relative to expected future dividends.

The last array reproduces the natural debt limits computed above: agent k's natural debt limit in state s is the cum-dividend price of a claim to agent k's own endowment stream.

That is why an agent can always repay a debt no larger than $\bar{A}^k(s)$: it could sell the claim to its endowment and consume nothing forever.

```
# initial state is state 1
print(f'α = {ex2.wealth_distribution(s0_idx=0)}')
print(f'ψ = \n{ex2.continuation_wealths()}')
print(f'J = \n{ex2.value_functions()}')
```

```
α = [0.50879763 0.49120237]
ψ =
[[[ 0.          0.        ]
  [ 0.55057195 -0.55057195]]]
J =
[[[122.907875   120.76397493]
  [123.32114686 121.17003803]]]
```

```
# initial state is state 2
print(f'α = {ex2.wealth_distribution(s0_idx=1)}')
print(f'ψ = \n{ex2.continuation_wealths()}')
print(f'J = \n{ex2.value_functions()}')
```

```
α = [0.50539319 0.49460681]
ψ =
[[[-0.46375886  0.46375886]
  [-0.         -0.        ]]]
J =
```

(continues on next page)

(continued from previous page)

```
[[[122.49598809 121.18174895]
  [122.907875   121.58921679]]]
```

134.10.3 Example 3: an absorbing state

In this example state 2 is absorbing, so state 1 is transient.

```
s = np.array([1, 2])

λ = 0.9
P = np.array([[1-λ, λ],
              [0, 1]])

ys = np.empty((2, 2))
ys[:, 0] = [1, 0]      # agent 1
ys[:, 1] = [0, 1]      # agent 2

ex3 = RecurCompetitive(s, P, ys)

print("pricing kernel Q = \n", ex3.Q)
print("natural debt limits A = \n", ex3.A)
```

```
pricing kernel Q =
 [[0.098 0.882]
 [0.    0.98 ]]
natural debt limits A =
 [[ 1.10864745 48.89135255]
 [ 0.         50.        ]]
```

The natural debt limit for agent 1 in state 2 is 0.

Once the economy enters the absorbing state, agent 1 never receives another unit of endowment, so it cannot credibly promise to repay anything.

```
# initial state is state 1
print(f'α = {ex3.wealth_distribution(s0_idx=0)}')
print(f'ψ = \n{ex3.continuation_wealths()}')
print(f'J = \n{ex3.value_functions()}')
```

```
α = [0.02217295 0.97782705]
ψ =
[[[ 0.         -0.        ]
  [ 1.10864745 -1.10864745]]]
J =
[[[14.89058394 98.88513796]
  [14.89058394 98.88513796]]]
```

Starting in state 1, agent 1 receives only a small share of aggregate consumption, because its endowment arrives only while the economy remains in the transient state.

Because state 1 is transient, the economy eventually leaves it for good, and agents' wealths do not recurrently return to zero.

```
# initial state is state 2
print(f'α = {ex3.wealth_distribution(s0_idx=1)}')
print(f'ψ = \n{ex3.continuation_wealths()}')
print(f'J = \n{ex3.value_functions()}')
```

```
α = [0. 1.]
ψ =
[[[-1.10864745  1.10864745]
  [ 0.          0.        ]]]
J =
[[[  0. 100.]
  [  0. 100.]]]
```

Starting in the absorbing state, agent 1 owns nothing of value, so $\alpha_1 = 0$ and agent 1 consumes nothing forever.

This corner is consistent with the discussion of the Inada condition above, which guarantees interior consumption only for agents whose endowments have positive value.

For the specification of the Markov chain in example 3, let's see how the equilibrium wealth distribution varies with the transition probability λ.

```
λ_seq = np.linspace(0, 0.99, 100)

# prepare containers
αs0_seq = np.empty((len(λ_seq), 2))
αs1_seq = np.empty((len(λ_seq), 2))

for i, λ in enumerate(λ_seq):
    P_λ = np.array([[1-λ, λ],
                    [0, 1]])
    ex3_λ = RecurCompetitive(s, P_λ, ys)

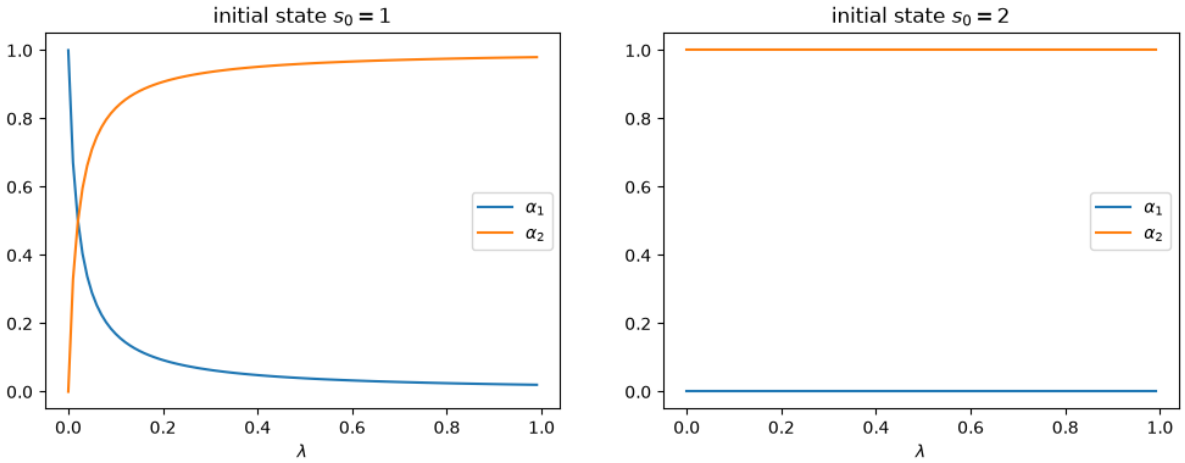
    # initial state s0 = 1
    αs0_seq[i, :] = ex3_λ.wealth_distribution(s0_idx=0)

    # initial state s0 = 2
    αs1_seq[i, :] = ex3_λ.wealth_distribution(s0_idx=1)
```

```
fig, axs = plt.subplots(1, 2, figsize=(12, 4))

for i, αs_seq in enumerate([αs0_seq, αs1_seq]):
    for j in range(2):
        axs[i].plot(λ_seq, αs_seq[:, j], label=f'$\\alpha_{j+1}$')
    axs[i].set_xlabel(r'$\lambda$')
    axs[i].set_title(f'initial state $s_0 = {s[i]}$')
    axs[i].legend()

plt.show()
```



When the economy starts in state 1, a higher probability λ of leaving that state for good shortens the expected duration of agent 1's endowment and lowers its wealth share.

When the economy starts in the absorbing state 2, λ is irrelevant and agent 2 owns everything.

134.10.4 Example 4: prosperity, a moderate state, and recession

Our last infinite-horizon example has three Markov states, which we interpret as prosperity, a moderate state, and recession.

```
s = np.array([1, 2, 3])

λ = .9
μ = .9
δ = .05

# prosperous, moderate, and recession states
P = np.array([[1-λ, λ, 0],
              [(1-μ)/2, μ, (1-μ)/2],
              [(1-δ)/2, (1-δ)/2, δ]])

ys = np.empty((3, 2))
ys[:, 0] = [.25, .75, .2]       # agent 1
ys[:, 1] = [1.25, .25, .2]      # agent 2

ex4 = RecurCompetitive(s, P, ys)

print("rows of P sum to", P.sum(axis=1))
print("aggregate endowment y =", ex4.y)
print("pricing kernel Q = \n", ex4.Q)
print("risk-free rate R =", ex4.R)
print("natural debt limits A = \n", ex4.A)
```

```
rows of P sum to [1. 1. 1.]
aggregate endowment y = [1.5 1.  0.4]
pricing kernel Q =
 [[0.098      1.08022498 0.        ]
 [0.04000833 0.882      0.0774758 ]
 [0.24038317 0.29440805 0.049     ]]
```

(continues on next page)

(continued from previous page)

```
risk-free rate R = [0.84873434 1.00051613 1.71294115]
natural debt limits A =
 [[41.66943407 19.83317897]
 [34.56301265 15.40376106]
 [21.44296495  9.99216994]]
```

The moderate state is highly persistent, and the economy moves out of recession quickly.

The gross risk-free rate is below one in prosperity, where aggregate consumption is expected to fall, and well above one in recession, where it is expected to recover.

```
for i in range(3):
    print(f"initial state is state {i+1}")
    print(f'α = {ex4.wealth_distribution(s0_idx=i)}')
    print(f'ψ = \n{ex4.continuation_wealths()}')
    print(f'J = \n{ex4.value_functions()}\n')
```

```
initial state is state 1
α = [0.67752299 0.32247701]
ψ =
[[[-0.         -0.        ]
  [-0.7093749   0.7093749 ]
  [-0.14493847  0.14493847]]]
J =
[[[82.66847975 57.03314431]
  [82.25708619 56.7493231 ]
  [81.82336679 56.45009924]]]

initial state is state 2
α = [0.69171992 0.30828008]
ψ =
[[[ 0.87314842 -0.87314842]
  [ 0.         -0.        ]
  [ 0.30134401 -0.30134401]]]
J =
[[[83.53011415 55.76358242]
  [83.11443273 55.48607909]
  [82.67619277 55.19351598]]]

initial state is state 3
α = [0.6821337 0.3178663]
ψ =
[[[ 0.28357106 -0.28357106]
  [-0.47899231  0.47899231]
  [-0.         -0.        ]]]
J =
[[[82.9492928  56.62395182]
  [82.53650179 56.34216692]
  [82.1013091  56.04508989]]]
```

Agent 1, whose endowment is concentrated in the persistent moderate state, receives about two thirds of aggregate consumption whichever state the economy starts in.

134.10.5 A finite-horizon example

We now revisit the economy defined in example 1, but set the time horizon to $T = 10$.

```
s = np.array([0, 1])
P = np.array([[.5, .5],
              [.5, .5]])

ys = np.empty((2, 2))
ys[:, 0] = 1 - s        # agent 1
ys[:, 1] = s            # agent 2

ex1_finite = RecurCompetitive(s, P, ys, T=10)
```

```
# I + Q + Q^2 + ... + Q^T
ex1_finite.V[0]
```

```
array([[5.48171623, 4.48171623],
       [4.48171623, 5.48171623]])
```

In the finite-horizon case, ψ and $\mathcal{J}$ are returned as sequences ordered from $t = 0$ to $t = T$.

```
# initial state is state 1
print(f'α = {ex1_finite.wealth_distribution(s0_idx=0)}')
print(f'ψ = \n{ex1_finite.continuation_wealths()}\n')
print(f'J = \n{ex1_finite.value_functions()}')
```

```
α = [0.55018351 0.44981649]
ψ =
[[[ 0.         -0.        ]
  [ 1.         -1.        ]]

 [[-0.04100358  0.04100358]
  [ 0.95899642 -0.95899642]]

 [[-0.08284397  0.08284397]
  [ 0.91715603 -0.91715603]]

 [[-0.12553824  0.12553824]
  [ 0.87446176 -0.87446176]]

 [[-0.16910383  0.16910383]
  [ 0.83089617 -0.83089617]]

 [[-0.21355851  0.21355851]
  [ 0.78644149 -0.78644149]]

 [[-0.25892042  0.25892042]
  [ 0.74107958 -0.74107958]]

 [[-0.30520809  0.30520809]
  [ 0.69479191 -0.69479191]]

 [[-0.35244041  0.35244041]
  [ 0.64755959 -0.64755959]]

 [[-0.40063665  0.40063665]
```

(continues on next page)

(continued from previous page)

```
  [ 0.59936335 -0.59936335]]

 [[-0.44981649  0.44981649]
  [ 0.55018351 -0.55018351]]]

J =
[[[14.78062373 13.3646215 ]
  [14.78062373 13.3646215 ]]

 [[13.56850674 12.26862684]
  [13.56850674 12.26862684]]

 [[12.33165268 11.15026494]
  [12.33165268 11.15026494]]

 [[11.06955669 10.00907933]
  [11.06955669 10.00907933]]

 [[ 9.78170364  8.84460421]
  [ 9.78170364  8.84460421]]

 [[ 8.46756788  7.6563643 ]
  [ 8.46756788  7.6563643 ]]

 [[ 7.12661302  6.44387459]
  [ 7.12661302  6.44387459]]

 [[ 5.75829174  5.20664019]
  [ 5.75829174  5.20664019]]

 [[ 4.36204553  3.94415611]
  [ 4.36204553  3.94415611]]

 [[ 2.9373045   2.65590706]
  [ 2.9373045   2.65590706]]

 [[ 1.48348712  1.3413672 ]
  [ 1.48348712  1.3413672 ]]]
```

```
# initial state is state 2
print(f'α = {ex1_finite.wealth_distribution(s0_idx=1)}')
print(f'ψ = \n{ex1_finite.continuation_wealths()}\n')
print(f'J = \n{ex1_finite.value_functions()}')
```

```
α = [0.44981649 0.55018351]
ψ =
[[[-1.          1.        ]
  [ 0.         -0.        ]]

 [[-0.95899642  0.95899642]
  [ 0.04100358 -0.04100358]]

 [[-0.91715603  0.91715603]
  [ 0.08284397 -0.08284397]]

 [[-0.87446176  0.87446176]
```

(continues on next page)

(continued from previous page)

```
  [ 0.12553824 -0.12553824]]

 [[-0.83089617  0.83089617]
  [ 0.16910383 -0.16910383]]

 [[-0.78644149  0.78644149]
  [ 0.21355851 -0.21355851]]

 [[-0.74107958  0.74107958]
  [ 0.25892042 -0.25892042]]

 [[-0.69479191  0.69479191]
  [ 0.30520809 -0.30520809]]

 [[-0.64755959  0.64755959]
  [ 0.35244041 -0.35244041]]

 [[-0.59936335  0.59936335]
  [ 0.40063665 -0.40063665]]

 [[-0.55018351  0.55018351]
  [ 0.44981649 -0.44981649]]]

J =
[[[13.3646215  14.78062373]
  [13.3646215  14.78062373]]

 [[12.26862684 13.56850674]
  [12.26862684 13.56850674]]

 [[11.15026494 12.33165268]
  [11.15026494 12.33165268]]

 [[10.00907933 11.06955669]
  [10.00907933 11.06955669]]

 [[ 8.84460421  9.78170364]
  [ 8.84460421  9.78170364]]

 [[ 7.6563643   8.46756788]
  [ 7.6563643   8.46756788]]

 [[ 6.44387459  7.12661302]
  [ 6.44387459  7.12661302]]

 [[ 5.20664019  5.75829174]
  [ 5.20664019  5.75829174]]

 [[ 3.94415611  4.36204553]
  [ 3.94415611  4.36204553]]

 [[ 2.65590706  2.9373045 ]
  [ 2.65590706  2.9373045 ]]

 [[ 1.3413672   1.48348712]
  [ 1.3413672   1.48348712]]]
```

The wealth distribution is further from equal than in example 1, because with a short horizon the endowment received in the initial period is a larger fraction of the total value of an agent's endowment.

Let's check why this economy needs no borrowing limits.

At date t, the most that agent k could ever repay from its own endowment is the value of its remaining endowment stream, $[I + Q + \cdots + Q^{T-t}]\, y^k$.

Unlike in the infinite-horizon economy, these bounds need not be imposed.

At date T no securities are traded, so an agent cannot roll debt over, and nonnegative consumption at T limits what it can owe to $y^k(s_T)$.

Working backward, nonnegative consumption at each earlier date implies the bound for that date.

The distance between an agent's bound and its debt, $\psi_t^k + [I + Q + \cdots + Q^{T-t}]\, y^k$, equals $[I + Q + \cdots + Q^{T-t}]\, \alpha_k y$, the value of the agent's remaining consumption.

That value is positive because consumption is positive, so the implied bounds never bind.

The following code computes the bounds for every date and state and confirms this.

```
ex1_finite.wealth_distribution(s0_idx=0)
ψ_finite = ex1_finite.continuation_wealths()
bounds = ex1_finite.V @ ex1_finite.ys      # bounds[t] = (I + ... + Q^(T-t)) y^k

print("implied bounds at t = T (rows: states, columns: agents):\n", bounds[-1])
print("smallest slack ψ_t + bound over all t, states, agents:",
      (ψ_finite + bounds).min().round(4))
```

```
implied bounds at t = T (rows: states, columns: agents):
 [[1. 0.]
 [0. 1.]]
smallest slack ψ_t + bound over all t, states, agents: 0.4498
```

At $t = T$, the bounds are just current endowments, and the slack is smallest there because only one period of consumption remains to be valued.

In the infinite-horizon economy there is no last date at which debts must be settled, so without explicit limits an agent could roll over ever larger debts forever, and the natural debt limits (134.2) must be imposed.

We can check that as $T \to \infty$ the finite-horizon results converge to those of the infinite-horizon economy.

Both economies below start in state 2, and we compare time-0 objects.

```
ex1_large = RecurCompetitive(s, P, ys, T=10000)
ex1.wealth_distribution(s0_idx=1)
ex1_large.wealth_distribution(s0_idx=1)

print("V:", np.abs(ex1.V[0] - ex1_large.V[0]).max())
print("ψ:", np.abs(ex1.continuation_wealths()[0]
                   - ex1_large.continuation_wealths()[0]).max())
print("J:", np.abs(ex1.value_functions()[0]
                   - ex1_large.value_functions()[0]).max())
```

```
V: 8.526512829121202e-14
ψ: 1.1857181902996411e-15
J: 2.2737367544323206e-13
```

The maximum absolute differences are negligible.

134.11 Concluding remarks

We began by promising a computable account of competitive equilibrium in an infinite-horizon exchange economy with heterogeneous endowments, complete markets in one-period Arrow securities, identical CRRA preferences, and common beliefs.

Here is how the lecture delivered on that promise.

134.11.1 Prices before the wealth distribution

Because all agents face the same pricing kernel, their Euler equations force consumption shares to be constant, so that $c^k(s) = \alpha_k y(s)$.

The pricing kernel (134.1) therefore depends only on the aggregate endowment, and it coincides with the kernel of a representative-agent Lucas tree economy.

The wealth distribution α comes second, pinned down by the requirement that every agent start with zero financial wealth, which makes α_k equal to agent k's share of the value of aggregate endowments in the initial state.

134.11.2 Bellmanizing the equilibrium

Each object that the overview promised to describe with a Bellman equation satisfies a one-step recursion of the form $x = b + Qx$:

- asset prices satisfy $p^h = d^h + Qp^h$, as we verified for a Lucas tree in example 2;
- natural debt limits satisfy $\bar{A}^k = y^k + Q\bar{A}^k$, which makes them the prices of claims to agents' endowment streams; and
- continuation wealth satisfies $\psi^k = (\alpha_k y - y^k) + Q\psi^k$.

Value functions satisfy the analogous recursion $J^k = u(\alpha_k y) + \beta P J^k$, with βP in place of Q.

These recursions are what let a single Python class, `RecurCompetitive`, compute an entire equilibrium from a few matrix operations.

134.11.3 The five ideas

- **Resolvent operators.** Solving each recursion gives $(I - Q)^{-1}$ or $(I - \beta P)^{-1}$, and the similarity $Q = \beta D^{-1} P D$ guarantees that the spectral radius of Q is $\beta < 1$, so these resolvents exist.
- **State-by-state borrowing limits in infinite horizons.** Natural debt limits $\bar{A}^k = (I - Q)^{-1} y^k$ are the largest debts that agent k could repay from its own endowment, and example 3 showed that they can be zero in states where an agent's future endowment is worthless.
- **No borrowing limits in finite horizons.** When the economy ends at T, no one can roll debt over past the last date, and nonnegative consumption implies bounds $[I + Q + \cdots + Q^{T-t}]\, y^k$ that we computed and showed never bind, so no separate debt limits need to be imposed.
- **The law of iterated values.** Multi-period Arrow prices compound as Q^j in the same way that multi-period transition probabilities compound as P^j, and the backward recursion $S_t = I + QS_{t+1}$ used in `RecurCompetitive` values payouts one period at a time.
- **State variable degeneracy.** Starting from zero financial wealth, each agent's wealth is a function of the Markov state alone and returns to zero whenever the state returns to s_0, as Exercise 134.13.1 verified along a simulated path;

Exercise 134.13.3 showed that this strong form of degeneracy fails in finite horizons, where wealth also depends on calendar time.

134.11.4 What the assumptions bought

Complete markets, identical CRRA preferences, and common beliefs together are what make prices independent of the wealth distribution and make wealth a function of the current state alone.

Relaxing any of them breaks at least one of these properties, which is the subject of several of the lectures listed below.

134.11.5 How old this is

Before leaving the model, it is worth recording how much of it was in place at the start.

Both trading arrangements, the proof that they support the same allocation, and the observation that the sequential one presumes correct forecasts of future spot prices all appear in the paper that Arrow read in 1952, nine years before Muth [1961] named the hypothesis that such forecasts embody.

Arrow did not use the term, and the dynamic programming arguments that make the equivalence transparent were not yet available to him.

His argument also has a knife edge that his unit of account conceals and that Hart [1975] exposed: when securities pay goods rather than dollars, consumers' forecasts determine which risks the asset market can carry, and correct forecasts no longer pin down a unique, or even an existing, equilibrium.

The single good per date and history that we assumed throughout is exactly the case in which the forecasting requirement has no bite.

134.12 Related lectures

This lecture assumes that all agents share beliefs about nature, and that their expectations of prices are correct.

Arrow's distinction between those two uses of probability organizes the first two entries below.

- *Heterogeneous Beliefs and Bubbles* studies an economy in which agents disagree about probabilities and short sales are constrained.
- *Heterogeneous Beliefs and Financial Markets* studies complete markets when agents hold different beliefs, in which case wealth shares drift with likelihood ratios instead of staying constant at α.
- *Consumption Smoothing with Incomplete and Complete Markets* shows, in a linear-quadratic setting, how complete markets in Arrow securities deliver a time-invariant cross-section distribution of consumption.
- *The Recovery Theorem* and *Long-term Risk: An Operator Approach* study what the Perron–Frobenius eigenvalue and eigenvector of a pricing kernel like Q reveal.
- *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns* confronts Euler equations like those used here with data.

134.13 Exercises

Exercise 134.13.1

This exercise verifies that the objects computed by `RecurCompetitive` constitute a recursive competitive equilibrium.

Use example 4 and start the economy in state 1.

1. Check that every agent's Euler equation holds at the allocation $c^k(s) = \alpha_k y(s)$ and the pricing kernel Q.
2. Check that agent k's budget constraint $c^k(s) + \sum_{s'} Q(s' \mid s)\,\psi^k(s') = y^k(s) + \psi^k(s)$ holds in every state when the agent holds the portfolio $\hat{a}^k(s') = \psi^k(s')$.
3. Check that Arrow security markets clear, $\sum_k \psi^k(s) = 0$, and that no natural debt limit binds.
4. Simulate 200 periods of the Markov chain and confirm that each agent's financial wealth is exactly zero at every visit to the initial state.

Solution

Here is one solution.

```
ex = RecurCompetitive(ex4.s, ex4.P, ex4.ys)

α = ex.wealth_distribution(s0_idx=0)
ψ = ex.continuation_wealths()[0]
c = np.outer(ex.y, α)

euler = max(np.abs(ex.Q - ex.β * ex.P * ex.u_prime(c[:, k])[None, :]
                   / ex.u_prime(c[:, k])[:, None]).max()
            for k in range(ex.K))
budget = np.abs(c + ex.Q @ ψ - ex.ys - ψ).max()

print(f"largest Euler equation residual    {euler:.1e}")
print(f"largest budget constraint residual {budget:.1e}")
print(f"largest net supply of any security {np.abs(ψ.sum(axis=1)).max():.1e}")
print(f"all natural debt limits slack:     {np.all(ψ + ex.A > 0)}")
```

```
largest Euler equation residual    2.2e-16
largest budget constraint residual 3.1e-16
largest net supply of any security 6.5e-15
all natural debt limits slack:     True
```

```
rng = np.random.default_rng(0)
T_sim = 200
states = np.empty(T_sim, dtype=int)
states[0] = 0
for t in range(T_sim - 1):
    states[t+1] = rng.choice(ex.n, p=ex.P[states[t]])

a = ψ[states]            # financial wealth of each agent along the path
print(f"visits to the initial state: {np.sum(states == 0)}")
print(f"largest |wealth| at those visits: {np.abs(a[states == 0]).max():.1e}")
```

```
visits to the initial state: 16
largest |wealth| at those visits: 5.9e-15
```

All four conditions hold to machine precision.

The last check is **state variable degeneracy** in action: financial wealth is a function of the Markov state alone, so it returns to its initial value of zero whenever the state does.

Exercise 134.13.2

This exercise connects the lecture to *The Recovery Theorem*.

An outside observer sees the equilibrium pricing kernel Q and the aggregate endowment y, but not β, γ, or the transition matrix P.

1. Show that the vector with components $y(\bar{s}_j)^\gamma$ is a right eigenvector of Q with eigenvalue β.
2. Using example 4, compute the Perron–Frobenius eigenvalue and eigenvector of Q, and use them to recover β, γ, and P.

Solution

Here is one solution.

Let $v_j = y_j^\gamma$.

Using (134.1),

$$\sum_j Q_{ij} v_j = \sum_j \beta \left(\frac{y_j}{y_i}\right)^{-\gamma} P_{ij}\, y_j^\gamma = \beta\, y_i^\gamma \sum_j P_{ij} = \beta\, v_i,$$

because the rows of P sum to one.

Since v is strictly positive, it is the Perron–Frobenius eigenvector of the nonnegative matrix Q, and β is its largest eigenvalue.

Given β and v, the transition matrix is $P_{ij} = Q_{ij} v_j / (\beta v_i)$, and γ is the slope of $\log v$ on $\log y$.

```
eigvals, eigvecs = np.linalg.eig(ex4.Q)
i = np.argmax(eigvals.real)
β_hat = eigvals[i].real
v = np.abs(eigvecs[:, i].real)

γ_hat = np.polyfit(np.log(ex4.y), np.log(v), 1)[0]
P_hat = ex4.Q * v[None, :] / (β_hat * v[:, None])

print(f"recovered β = {β_hat:.6f}   (true {ex4.β})")
print(f"recovered γ = {γ_hat:.6f}   (true {ex4.γ})")
print(f"max |P_hat - P| = {np.abs(P_hat - ex4.P).max():.1e}")
```

```
recovered β = 0.980000   (true 0.98)
recovered γ = 0.500000   (true 0.5)
max |P_hat - P| = 3.3e-16
```

Recovery is exact because the equilibrium kernel has precisely the **transition independence** structure studied in *The Recovery Theorem*: $Q = \beta D^{-1} P D$, with D built from the marginal utility of the aggregate endowment.

Exercise 134.13.3

In the infinite-horizon economy, continuation wealth depends only on the Markov state.

This exercise shows that in the finite-horizon economy it depends on calendar time as well.

1. For the finite-horizon version of example 1 with $T = 10$ and initial state 1, report $\psi_t^1(\bar{s}_1)$ for $t = 0, 1, \ldots, 10$.
2. Explain why agent 1's wealth does not return to zero when the Markov state returns to $\bar{s}_1$.
3. Compute the wealth distribution α for horizons $T = 1, \ldots, 300$ and show that it converges to the infinite-horizon distribution at a geometric rate close to β.

Solution

Here is one solution.

```
s = np.array([0, 1])
P = np.array([[.5, .5],
              [.5, .5]])
ys = np.array([[1., 0.],
               [0., 1.]])

ex_T = RecurCompetitive(s, P, ys, T=10)
ex_T.wealth_distribution(s0_idx=0)
ψ_T = ex_T.continuation_wealths()

print("ψ_t^1(s_1), t = 0,...,10:", ψ_T[:, 0, 0].round(4))
```

```
ψ_t^1(s_1), t = 0,...,10: [ 0.     -0.041  -0.0828 -0.1255 -0.1691 -0.2136 -0.
↪2589 -0.3052 -0.3524
 -0.4006 -0.4498]
```

From (134.6), $\psi_t^k = [I + Q + \cdots + Q^{T-t}][\alpha_k y - y^k]$.

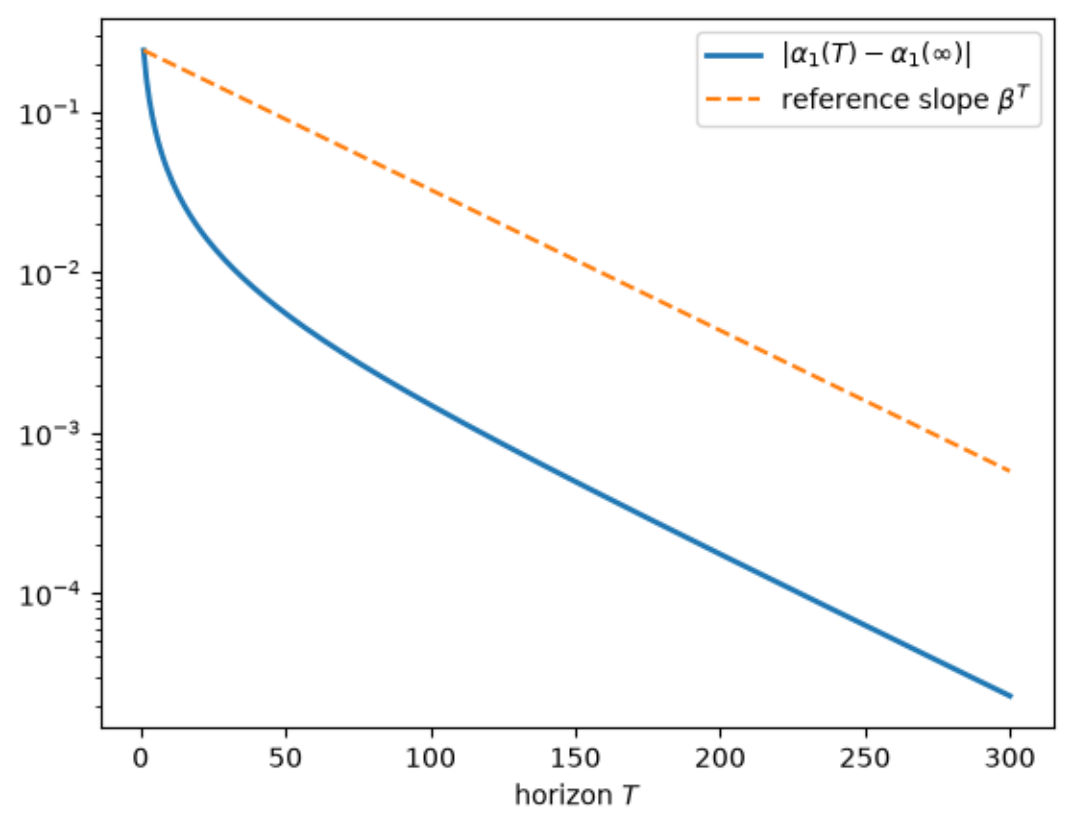
The initial condition pins $\psi_0^k(\bar{s}_1) = 0$, but at a later date t fewer periods remain, so the geometric sum is truncated sooner and $\psi_t^k(\bar{s}_1) \neq 0$.

In the infinite horizon the sum is never truncated, which is why wealth there depends on the state alone.

```
α_inf = RecurCompetitive(s, P, ys).wealth_distribution(s0_idx=0)[0]

T_grid = np.arange(1, 301)
gaps = np.array([abs(RecurCompetitive(s, P, ys, T=T).wealth_distribution(s0_
↪idx=0)[0]
                     - α_inf)
                 for T in T_grid])

fig, ax = plt.subplots()
ax.semilogy(T_grid, gaps, lw=2, label=r'$|\alpha_1(T) - \alpha_1(\infty)|$')
ax.semilogy(T_grid, gaps[0] * ex_T.β ** (T_grid - 1), '--', lw=1.5,
            label=r'reference slope $\beta^{T}$')
ax.set_xlabel('horizon $T$')
ax.legend()
plt.show()
```



The gap shrinks geometrically, at a rate governed by the spectral radius of Q, which equals β.

Exercise 134.13.4

This exercise asks you to verify, in a setting with more than one good, the equivalence that *A rational expectations equilibrium* describes for our one-good economy.

Following Arrow [1964], consider a two-period pure exchange economy with n consumers, S states, and C goods.

Consumer i is endowed with the state-contingent vector $\omega_i = (\omega_{i1}, \ldots, \omega_{iS})$ and ranks state-contingent consumption plans x_i by

$$U_i(x_i) = \sum_{s=1}^{S} u_i(x_{is})\,\pi_s, \qquad u_i(x) = a_{i1}\sqrt{x_1} + a_{i2}\sqrt{x_2}$$

In a **contingent claims equilibrium**, a price vector p^* for claims to goods in states and an allocation $\{x_i^*\}$ are such that each consumer maximizes U_i subject to $p^* \cdot x_i = p^* \cdot \omega_i$ and all SC markets clear.

In the **sequential** arrangement, consumer i enters a securities market with an endowment M_i of dollars, a unit of security s pays one dollar if state s occurs and sells for q_s, and after the state is realized goods trade on spot markets at prices $\bar{p}_s$.

1. Show that consumer i's spot problem has indirect utility $V_i(y, \bar{p}_s) = \sqrt{y}\, G_i(\bar{p}_s)$ with $G_i(\bar{p}_s) = \left(\sum_c a_{ic}^2 / \bar{p}_{sc}\right)^{1/2}$, and derive the within-state demands. At which point does the definition of a **rational expectations equilibrium** require consumers to forecast $\bar{p}_s$ correctly?

2. Compute the contingent claims equilibrium of the two-state, two-good, two-consumer economy specified in the solution below.
3. Define

$$M_i = p^* \cdot \omega_i, \qquad q_s^* = \frac{\sum_i p_s^* \cdot x_{is}^*}{\sum_\sigma \sum_i p_\sigma^* \cdot x_{i\sigma}^*}, \qquad \bar{p}_s = \frac{p_s^*}{q_s^*}, \qquad y_{is}^* = \bar{p}_s \cdot x_{is}^*$$

and verify that $\{y_i^*, x_i^*\}, q^*, \bar{p}$ is a rational expectations equilibrium: check that $\sum_s q_s^* y_{is}^* = M_i$, that $\sum_i y_{is}^* = \sum_i M_i$ for every s, that the spot demands reproduce x_i^*, and that the portfolio satisfies its own first-order conditions.

4. Verify that $\sum_s q_s^* = 1$, and explain why a consumer who could hold cash instead of securities would arbitrage against any price system that violated this equality.
5. Going the other way, define $p_s^* = q_s^* \bar{p}_s$ and verify that you recover the contingent claims equilibrium.
6. Finally, show that an incorrect forecast is not an equilibrium: perturb the forecast of relative prices in state 1, recompute portfolios, and compare the spot prices that then clear the goods markets with the forecast.

Solution

Part 1. With $u_i(x) = \sum_c a_{ic}\sqrt{x_c}$ and a spot budget $\bar{p}_s \cdot x = y$, the first-order conditions give $x_c \propto a_{ic}^2/\bar{p}_{sc}^2$, so

$$x_c = \frac{a_{ic}^2/\bar{p}_{sc}^2}{\sum_{c'} a_{ic'}^2/\bar{p}_{sc'}}\, y, \qquad V_i(y, \bar{p}_s) = \sqrt{y}\left(\sum_c a_{ic}^2/\bar{p}_{sc}\right)^{1/2}$$

The portfolio problem is max $\sum_s \pi_s V_i(y_{is}, \bar{p}_s)$ subject to $\sum_s q_s y_{is} = M_i$.

The forecast $\bar{p}_s$ enters here, through $G_i(\bar{p}_s)$, before any state is realized.

The definition of equilibrium requires the forecast used in this portfolio problem to be the price vector that later clears the spot market in state s.

Since V_i is increasing in y and concave, and G_i depends on $\bar{p}_s$, a consumer who misforecasts relative prices chooses a portfolio that is optimal for prices that never materialize.

Parts 2 to 5.

```
π = np.array([0.4, 0.6])                  # probabilities of the two states
A = np.array([[2.0, 1.0],                 # taste parameters a_{ic}
              [1.0, 3.0]])
ω = np.array([[[2.0, 0.5], [0.5, 1.5]],   # consumer 1, by state
              [[1.0, 1.5], [1.0, 1.0]]]) # consumer 2, by state
Ω = ω.sum(axis=0)                         # aggregate endowment by state
S, C = Ω.shape

def spot_demand(p_s, y, a):
    "Within-state demands given spot prices p_s and spending y."
    w = a**2 / p_s**2
    return w * y / (w * p_s).sum()

def G(p_s, a):
    "Indirect utility coefficient: V = sqrt(y) G(p_s)."
    return np.sqrt((a**2 / p_s).sum())

def spending_shares(p, a):
    "Shares of time 0 wealth spent in each state."
    g = np.array([π[s] * G(p[s], a) for s in range(S)])
    return g**2 / (g**2).sum()
```

```
def cc_excess_demand(p_flat):
    "Excess demand in the SC contingent claims markets."
    p = p_flat.reshape((S, C))
    exc = np.zeros((S, C))
    for i in range(len(A)):
        M_i = (p * ω[i]).sum()
        share = spending_shares(p, A[i])
        for s in range(S):
            exc[s] += spot_demand(p[s], share[s] * M_i, A[i])
    return (exc - Ω).ravel()

sol = root(cc_excess_demand, np.ones(S * C), tol=1e-13)
p_star = sol.x.reshape((S, C))
p_star = p_star / p_star[0, 0]           # numeraire: good 1 in state 1

M = np.array([(p_star * ω[i]).sum() for i in range(len(A))])
x_star = np.array([[spot_demand(p_star[s], spending_shares(p_star, A[i])[s] * M[i],
                                A[i]) for s in range(S)] for i in range(len(A))])

print("p* =\n", p_star.round(5))
print("largest excess demand:", np.abs(cc_excess_demand(p_star.ravel())).max())
print("M =", M.round(5))
print("x1* =\n", x_star[0].round(5), "\nx2* =\n", x_star[1].round(5))
```

```
p* =
 [[1.      1.42351]
 [2.12132 1.90983]]
largest excess demand: 1.3322676295501878e-15
M = [6.63716 7.16641]
x1* =
 [[2.62254 0.32355]
 [1.31127 0.40444]]
x2* =
 [[0.37746 1.67645]
 [0.18873 2.09556]]
```

Now build the sequential arrangement from these objects.

```
q = np.array([(p_star[s] * Ω[s]).sum() for s in range(S)]) / M.sum()
p_bar = p_star / q[:, None]
y_star = np.array([[p_bar[s] @ x_star[i, s] for s in range(S)]
                   for i in range(len(A))])

print("q* =", q.round(5), " and Σ_s q*_s =", q.sum().round(10))
print("p_bar =\n", p_bar.round(5))
print("y* =\n", y_star.round(5))

print("\nΣ_s q*_s y*_is = M_i:          ", np.allclose(y_star @ q, M))
print("Σ_i y*_is = Σ_i M_i for each s:", np.allclose(y_star.sum(axis=0), M.sum()))
print("spot demands reproduce x*:     ",
      all(np.allclose(spot_demand(p_bar[s], y_star[i, s], A[i]), x_star[i, s])
          for i in range(len(A)) for s in range(S)))

# portfolio first-order condition: sqrt(y_is) q_s / (π_s G_i) is constant across s
for i in range(len(A)):
    ratio = np.sqrt(y_star[i]) * q / np.array([π[s] * G(p_bar[s], A[i])
                                               for s in range(S)])
    print(f"consumer {i+1} portfolio FOC, across states:", ratio.round(8))

print("\nq*_s p_bar_s recovers p*:", np.allclose(q[:, None] * p_bar, p_star))
```

```
q* = [0.42359 0.57641]  and Σ_s q*_s = 1.0
p_bar =
 [[2.36079 3.3606 ]
 [3.68021 3.31331]]
y* =
 [[7.27861 6.16579]
 [6.52497 7.63779]]

Σ_s q*_s y*_is = M_i:           True
Σ_i y*_is = Σ_i M_i for each s: True
spot demands reproduce x*:      True
consumer 1 portfolio FOC, across states: [2.02428387 2.02428387]
consumer 2 portfolio FOC, across states: [1.53593595 1.53593595]

q*_s p_bar_s recovers p*: True
```

The securities prices sum to one because a unit of each security, bought in equal numbers, pays one dollar for sure.

A portfolio of one unit of every security is therefore a riskless claim to one dollar, so it must cost one dollar.

If $\sum_s q_s < 1$, a consumer could buy that bundle and hold a sure dollar for less than a dollar; if $\sum_s q_s > 1$, selling the bundle and holding cash would do the same in reverse.

Part 6.

```
def realized_spot_prices(y):
    "Spot prices that clear each state's goods markets given dollar wealth y."
    out = np.zeros((S, C))
    for s in range(S):
        def excess(p_s):
            return sum(spot_demand(p_s, y[i, s], A[i])
                       for i in range(len(A))) - Ω[s]
        out[s] = root(excess, np.ones(C), tol=1e-13).x
    return out

def portfolios(forecast):
    "Optimal portfolios when consumers forecast the spot prices in forecast."
    y = np.zeros((len(A), S))
    for i in range(len(A)):
        g = np.array([π[s] * G(forecast[s], A[i]) for s in range(S)])
        y[i] = M[i] * (g / q)**2 / ((g / q)**2 * q).sum()
    return y

print("correct forecast reproduces itself:",
      np.allclose(realized_spot_prices(y_star), p_bar))

p_wrong = p_bar.copy()
p_wrong[0] = p_bar[0] * np.array([1.5, 1.0])     # misforecast relative prices
p_realized = realized_spot_prices(portfolios(p_wrong))

print(f"\nforecast relative price in state 1: "
      f"{p_wrong[0, 0] / p_wrong[0, 1]:.4f}")
print(f"realized relative price in state 1: "
      f"{p_realized[0, 0] / p_realized[0, 1]:.4f}")
```

```
correct forecast reproduces itself: True

forecast relative price in state 1: 1.0537
realized relative price in state 1: 0.6496
```

With the correct forecast, the prices that clear the spot markets are the prices that consumers used when choosing portfolios, so the forecast is confirmed and the sequential allocation is the contingent claims allocation.

With the incorrect forecast, consumers carry the wrong dollar wealth into each state, and the prices that then clear the spot markets are not the prices they forecast.

The plans are feasible and the securities market clears, but the economy is not in a rational expectations equilibrium.

Note that only *relative* prices within a state need to be forecast correctly.

Scaling all of state s's spot prices by a constant is absorbed by the security price q_s, which is why the normalization of $\bar{p}_s$ is harmless.

Exercise 134.13.5

In Exercise 134.13.4 the securities pay dollars, so the set of state-contingent wealth distributions that a portfolio can deliver does not depend on what consumers forecast.

This exercise, which follows Hart [1975], asks what happens when securities pay *goods* instead.

There are two dates.

Securities trade at the first date; at the second date a state $s \in \{1, 2\}$ is realized and $C = 2$ goods trade on spot markets.

Two consumers care only about consumption at the second date and assign probability $1/2$ to each state, with

$$u^1(x) = 2^{2.5}\sqrt{x_1} + 2\sqrt{x_2}, \qquad u^2(x) = 2\sqrt{x_1} + 2^{2.5}\sqrt{x_2}$$

and endowments

$$\omega_{11} = \left(\tfrac{5}{2}, \tfrac{50}{21}\right), \quad \omega_{21} = \left(\tfrac{1}{2}, \tfrac{13}{21}\right), \quad \omega_{12} = \left(\tfrac{13}{21}, \tfrac{1}{2}\right), \quad \omega_{22} = \left(\tfrac{50}{21}, \tfrac{5}{2}\right)$$

where ω_{is} is consumer i's endowment in state s.

The aggregate endowment of each good is 3 in each state, so all risk is idiosyncratic, and the two consumers and the two states are mirror images of each other.

1. Suppose first that no securities are available. Compute the spot market equilibrium in each state.
2. Now let two securities trade at the first date. A unit of security 1 delivers one unit of good 1 if state 1 occurs and two units of good 1 if state 2 occurs; a unit of security 2 delivers two units of good 2 if state 1 occurs and one unit of good 2 if state 2 occurs. Show that if consumers forecast the spot prices of part 1, the two securities become perfect substitutes, so that no portfolio of them can move wealth between states. Conclude that this is a rational expectations equilibrium, and compute each consumer's expected utility.
3. Show that the forecast $\hat{p}_1 = \hat{p}_2 = (1, 1)$ instead makes the two securities span, so that the equilibrium allocation is the contingent claims allocation $x_{1s} = (8/3, 1/3)$, $x_{2s} = (1/3, 8/3)$. Verify that this forecast is also confirmed, compute expected utilities, and find the portfolio that consumer 1 uses.
4. Compare the two equilibria, and explain why no market force selects the better one.
5. Change only the dividends: let security 1 deliver one unit of good 1 in *both* states and security 2 one unit of good 2 in both states. Show that now no rational expectations equilibrium exists.

Solution

Part 1.

```
a = np.array([[2**2.5, 2.0],       # consumer 1
              [2.0, 2**2.5]])      # consumer 2
ω_h = np.array([[[5 / 2, 50 / 21], [13 / 21, 1 / 2]],      # consumer 1, by state
                [[1 / 2, 13 / 21], [50 / 21, 5 / 2]]])     # consumer 2, by state
Ω_h = ω_h.sum(axis=0)
def u(x, a_i):
    "Expected utility of the bundle x for a consumer with tastes a_i."
    return a_i @ np.sqrt(x)

print("aggregate endowment by state:\n", Ω_h.round(6))

p_hat = np.array([[2.0, 1.0], [1.0, 2.0]])      # conjectured spot prices
x_auto = np.array([[spot_demand(p_hat[s], p_hat[s] @ ω_h[i, s], a[i])
                    for s in range(2)] for i in range(2)])

print("\nno securities, state by state:")
for s in range(2):
    print(f"  state {s+1}: p = {p_hat[s]}, "
          f"x_1 = {(21 * x_auto[0, s]).round(4)}/21, "
          f"x_2 = {(21 * x_auto[1, s]).round(4)}/21, "
          f"markets clear: {np.allclose(x_auto[:, s].sum(axis=0), Ω_h[s])}")
```

```
aggregate endowment by state:
 [[3. 3.]
 [3. 3.]]

no securities, state by state:
  state 1: p = [2. 1.], x_1 = [62. 31.]/21, x_2 = [ 1. 32.]/21, markets clear:␣
↪True
  state 2: p = [1. 2.], x_1 = [32.  1.]/21, x_2 = [31. 62.]/21, markets clear:␣
↪True
```

The spot equilibrium in state 1 has prices proportional to $(2, 1)$ and in state 2 proportional to $(1, 2)$.

Each consumer eats mostly the good that it likes more.

In each state the good favored by the richer consumer is the expensive one, because the aggregate endowment is the same in both states and only the distribution of wealth differs.

Part 2. The dollar payoff of a security is the forecast price of the good it delivers, times the quantity delivered.

```
def payoff_matrix(forecast, dividends):
    """
    Dollar payoffs of securities, given forecast spot prices.

    dividends[f, s] is the bundle that security f delivers in state s.
    """
    return np.array([[forecast[s] @ dividends[f, s] for s in range(2)]
                     for f in range(2)])

div_b = np.array([[[1.0, 0.0], [2.0, 0.0]],     # security 1: good 1, 1 then 2␣
↪units
                  [[0.0, 2.0], [0.0, 1.0]]])    # security 2: good 2, 2 then 1 unit

Z = payoff_matrix(p_hat, div_b)
print("payoff matrix with the part 1 forecast:\n", Z)
print("rank:", np.linalg.matrix_rank(Z))

EU_auto = [sum(0.5 * u(x_auto[i, s], a[i]) for s in range(2)) for i in range(2)]
```

```
print(f"\nexpected utilities with no risk sharing: {EU_auto[0]:.4f}, {EU_auto[1]:.
 ↪4f}")
```

```
payoff matrix with the part 1 forecast:
 [[2. 2.]
 [2. 2.]]
rank: 1

expected utilities with no risk sharing: 9.7846, 9.7846
```

Both securities pay $(2, 2)$, so they are perfect substitutes and must bear the same price.

No portfolio of them can move wealth between the two states, so no securities trade, the realized spot prices are those of part 1, and the forecast is confirmed.

This is a rational expectations equilibrium in which the asset market is effectively incomplete.

Part 3.

```
p_span = np.array([[1.0, 1.0], [1.0, 1.0]])
Z_span = payoff_matrix(p_span, div_b)
print("payoff matrix with the forecast (1,1):\n", Z_span)
print("rank:", np.linalg.matrix_rank(Z_span))

x_cm = np.array([[8 / 3, 1 / 3], [1 / 3, 8 / 3]])    # complete markets allocation
print("\nallocation clears:", np.allclose(x_cm.sum(axis=0), ω_h[0]))
for i in range(2):
    ratio = x_cm[i, 1] / x_cm[i, 0]
    print(f"  consumer {i+1}: x_2/x_1 = {ratio:.4f}, "
          f"(a_2/a_1)^2 = {(a[i, 1] / a[i, 0])**2:.4f}")

EU_cm = [u(x_cm[i], a[i]) for i in range(2)]
print(f"\nexpected utilities with complete markets: {EU_cm[0]:.4f}, {EU_cm[1]:.4f}
 ↪")

wealth = np.array([p_span[s] @ ω_h[0, s] for s in range(2)])
z1 = np.linalg.solve(Z_span.T, p_span[0] @ x_cm[0] - wealth)
print(f"\nconsumer 1 state-contingent wealth: {(42 * wealth).round(3)}/42")
print(f"consumer 1 portfolio: {(42 * z1).round(3)}/42")
```

```
payoff matrix with the forecast (1,1):
 [[1. 2.]
 [2. 1.]]
rank: 2

allocation clears: True
  consumer 1: x_2/x_1 = 0.1250, (a_2/a_1)^2 = 0.1250
  consumer 2: x_2/x_1 = 8.0000, (a_2/a_1)^2 = 8.0000

expected utilities with complete markets: 10.3923, 10.3923

consumer 1 state-contingent wealth: [205.  47.]/42
consumer 1 portfolio: [ 79. -79.]/42
```

With this forecast the payoff vectors $(1, 2)$ and $(2, 1)$ are linearly independent, so the securities span and the allocation must be the contingent claims allocation.

Because each consumer's demands satisfy $x_2/x_1 = (a_2/a_1)^2$ at prices proportional to $(1, 1)$, those are indeed the spot prices that clear the markets, so this forecast is confirmed too.

Consumer 1 reaches the allocation by buying $79/42$ units of security 1 and selling $79/42$ units of security 2, a swap that costs nothing because the two securities bear the same price, and that converts its state-contingent wealth of

$(205/42, 47/42)$ into $(3, 3)$.

Consumer 2 holds the mirror-image portfolio.

Part 4.

```
print(f"{'':12}{'no risk sharing':>18}{'complete markets':>19}")
for i in range(2):
    print(f"consumer {i+1}: {EU_auto[i]:>17.4f}{EU_cm[i]:>19.4f}")
```

```
                no risk sharing   complete markets
consumer 1:             9.7846            10.3923
consumer 2:             9.7846            10.3923
```

Both consumers strictly prefer the second equilibrium, so it Pareto dominates the first.

Yet both are rational expectations equilibria: in each, consumers forecast correctly and markets clear.

Hart's point is that no market force selects the better one, because the trading opportunities available to a consumer depend on forecasts of spot prices, and a competitive consumer takes those forecasts as given.

A consumer cannot unilaterally make the two securities span, since spanning is a property of prices that no individual chooses.

Part 5.

```
div_e = np.array([[[1.0, 0.0], [1.0, 0.0]],     # security 1: one unit of good 1
                  [[0.0, 1.0], [0.0, 1.0]]])    # security 2: one unit of good 2

# each candidate forecast implies a span, the span implies an allocation,
# and the allocation implies the spot prices that actually clear markets
for name, forecast, implied in [
    ("(2,1) and (1,2)", p_hat,  p_span),
    ("(1,1) and (1,1)", p_span, p_hat),
]:
    Z_e = payoff_matrix(forecast, div_e)
    rank = np.linalg.matrix_rank(Z_e)
    spans = "spans" if rank == 2 else "does not span"
    allocation = "complete markets" if rank == 2 else "no risk sharing"
    print(f"forecast {name}: rank {rank}, {spans}")
    print(f"  implied allocation: {allocation}")
    print(f"  spot prices it implies: {implied[0]} and {implied[1]}")
    print(f"  forecast confirmed: {np.allclose(forecast, implied)}\n")
```

```
forecast (2,1) and (1,2): rank 2, spans
  implied allocation: complete markets
  spot prices it implies: [1. 1.] and [1. 1.]
  forecast confirmed: False

forecast (1,1) and (1,1): rank 1, does not span
  implied allocation: no risk sharing
  spot prices it implies: [2. 1.] and [1. 2.]
  forecast confirmed: False
```

With these dividends the payoff vectors are $(\hat{p}_{11}, \hat{p}_{21})$ and $(\hat{p}_{12}, \hat{p}_{22})$, so the securities span if and only if the two states' price vectors are not proportional.

Suppose a forecast makes them proportional.

Then no risk can be traded, so the realized spot prices are those of part 1, which are *not* proportional, and the forecast is wrong.

Suppose instead that a forecast makes them non-proportional.

Then the securities span, so the allocation is the complete markets allocation of part 3, whose spot prices *are* proportional, and again the forecast is wrong.

So no rational expectations equilibrium exists, even though preferences are strictly concave, endowments are strictly positive, and there are no transaction costs.

Nothing like parts 2 to 5 can happen in the sequential economy of this lecture, where there is one good per date and history and the one-period Arrow securities are denominated in that good.

There the span of the asset market is n-dimensional whatever consumers forecast, and the only forecast that matters is of the pricing kernel Q itself.

CHAPTER

FIVE

HETEROGENEOUS BELIEFS AND BUBBLES

Contents

- *Heterogeneous Beliefs and Bubbles*
  - *Overview*
  - *Structure of the model*
  - *Solving the model*
  - *Exercises*

In addition to what's in Anaconda, this lecture uses following libraries:

```
!pip install quantecon
```

135.1 Overview

This lecture describes a version of a model of Harrison and Kreps [Harrison and Kreps, 1978].

The model determines the price of a dividend-yielding asset that is traded by two types of self-interested investors.

The model features

- heterogeneous beliefs
- incomplete markets
- short sales constraints, and possibly ...
- (leverage) limits on an investor's ability to borrow in order to finance purchases of a risky asset

Let's start with some standard imports:

```
import numpy as np
import quantecon as qe
import scipy.linalg as la
```

135.1.1 References

Prior to reading the following, you might like to review our lectures on

- *Markov chains*
- *Asset pricing with finite state space*

135.1.2 Bubbles

Economists differ in how they define a *bubble*.

The Harrison-Kreps model illustrates the following notion of a bubble that attracts many economists:

> *A component of an asset price can be interpreted as a bubble when all investors agree that the current price of the asset exceeds what they believe the asset's underlying dividend stream justifies*.

135.2 Structure of the model

The model simplifies things by ignoring alterations in the distribution of wealth among investors who have hard-wired different beliefs about the fundamentals that determine asset payouts.

There is a fixed number A of shares of an asset.

Each share entitles its owner to a stream of dividends $\{d_t\}$ governed by a Markov chain defined on a state space $S \in \{0, 1\}$.

The dividend obeys

$$d_t = \begin{cases} 0 & \text{if } s_t = 0 \\ 1 & \text{if } s_t = 1 \end{cases}$$

An owner of a share at the end of time t and the beginning of time $t+1$ is entitled to the dividend paid at time $t+1$.

Thus, the stock is traded **ex dividend**.

An owner of a share at the beginning of time $t+1$ is also entitled to sell the share to another investor during time $t+1$.

Two types $h = a, b$ of investors differ only in their beliefs about a Markov transition matrix P with typical element

$$P(i,j) = \mathbb{P}\{s_{t+1} = j \mid s_t = i\}$$

Investors of type a believe the transition matrix

$$P_a = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}$$

Investors of type b think the transition matrix is

$$P_b = \begin{bmatrix} \frac{2}{3} & \frac{1}{3} \\ \frac{1}{4} & \frac{3}{4} \end{bmatrix}$$

Thus, in state 0, a type a investor is more optimistic about next period's dividend than is investor b.

But in state 1, a type a investor is more pessimistic about next period's dividend than is investor b.

The stationary (i.e., invariant) distributions of these two matrices can be calculated as follows:

```
qa = np.array([[1/2, 1/2], [2/3, 1/3]])
qb = np.array([[2/3, 1/3], [1/4, 3/4]])
mca = qe.MarkovChain(qa)
mcb = qe.MarkovChain(qb)
mca.stationary_distributions
```

```
array([[0.57142857, 0.42857143]])
```

```
mcb.stationary_distributions
```

```
array([[0.42857143, 0.57142857]])
```

The stationary distribution of P_a is approximately $\pi_a = \begin{bmatrix} .57 & .43 \end{bmatrix}$.

The stationary distribution of P_b is approximately $\pi_b = \begin{bmatrix} .43 & .57 \end{bmatrix}$.

Thus, a type a investor is more pessimistic on average.

135.2.1 Ownership rights

An owner of the asset at the end of time t is entitled to the dividend at time $t+1$ and also has the right to sell the asset at time $t+1$.

Both types of investors are risk-neutral and both have the same fixed discount factor $\beta \in (0, 1)$.

In our numerical example, we'll set $\beta = .75$, just as Harrison and Kreps [Harrison and Kreps, 1978] did.

We'll eventually study the consequences of two alternative assumptions about the number of shares A relative to the resources that our two types of investors can invest in the stock.

1. Both types of investors have enough resources (either wealth or the capacity to borrow) so that they can purchase the entire available stock of the asset[1].
2. No single type of investor has sufficient resources to purchase the entire stock.

Case 1 is the case studied in Harrison and Kreps.

In case 2, both types of investors always hold at least some of the asset.

135.2.2 Short sales prohibited

No short sales are allowed.

This matters because it limits how pessimists can express their opinions.

- They *can* express themselves by selling their shares.
- They *cannot* express themselves more loudly by artificially "manufacturing shares" – that is, they cannot borrow shares from more optimistic investors and then immediately sell them.

[1] By assuming that both types of agents always have "deep enough pockets" to purchase all of the asset, the model takes wealth dynamics off the table. The Harrison-Kreps model generates high trading volume when the state changes either from 0 to 1 or from 1 to 0.

135.2.3 Optimism and pessimism

The above specifications of the perceived transition matrices P_a and P_b, taken directly from Harrison and Kreps, build in stochastically alternating temporary optimism and pessimism.

Remember that state 1 is the high dividend state.

- In state 0, a type a agent is more optimistic about next period's dividend than a type b agent.
- In state 1, a type b agent is more optimistic about next period's dividend than a type a agent is.

However, the stationary distributions $\pi_a = \begin{bmatrix} .57 & .43 \end{bmatrix}$ and $\pi_b = \begin{bmatrix} .43 & .57 \end{bmatrix}$ tell us that a type b person is more optimistic about the dividend process in the long run than is a type a person.

135.2.4 Information

Investors know a price function mapping the state s_t at t into the equilibrium price $p(s_t)$ that prevails in that state.

This price function is endogenous and to be determined below.

When investors choose whether to purchase or sell the asset at t, they also know s_t.

135.3 Solving the model

Now let's turn to solving the model.

We'll determine equilibrium prices under a particular specification of beliefs and constraints on trading selected from one of the specifications described above.

We shall compare equilibrium price functions under the following alternative assumptions about beliefs:

1. There is only one type of agent, either a or b.
2. There are two types of agents differentiated only by their beliefs. Each type of agent has sufficient resources to purchase all of the asset (Harrison and Kreps's setting).
3. There are two types of agents with different beliefs, but because of limited wealth and/or limited leverage, both types of investors hold the asset each period.

135.3.1 Summary table

The following table gives a summary of the findings obtained in the remainder of the lecture (in an exercise you will be asked to recreate the table and also reinterpret parts of it).

The table reports implications of Harrison and Kreps's specifications of P_a, P_b, β.

s_t	0	1
p_a	1.33	1.22
p_b	1.45	1.91
p_o	1.85	2.08
p_p	1	1
$\hat{p}_a$	1.85	1.69
$\hat{p}_b$	1.69	2.08

Here

- p_a is the equilibrium price function under homogeneous beliefs P_a
- p_b is the equilibrium price function under homogeneous beliefs P_b
- p_o is the equilibrium price function under heterogeneous beliefs with optimistic marginal investors
- p_p is the equilibrium price function under heterogeneous beliefs with pessimistic marginal investors
- $\hat{p}_a$ is the amount type a investors are willing to pay for the asset
- $\hat{p}_b$ is the amount type b investors are willing to pay for the asset

We'll explain these values and how they are calculated one row at a time.

The row corresponding to p_o applies when both types of investor have enough resources to purchase the entire stock of the asset and strict short sales constraints prevail so that temporarily optimistic investors always price the asset.

The row corresponding to p_p would apply if neither type of investor has enough resources to purchase the entire stock of the asset and both types must hold the asset.

The row corresponding to p_p would also apply if both types have enough resources to buy the entire stock of the asset but short sales are also possible so that temporarily pessimistic investors price the asset.

135.3.2 Single belief prices

We'll start by pricing the asset under homogeneous beliefs.

(This is the case treated in *the lecture* on asset pricing with finite Markov states)

Suppose that there is only one type of investor, either of type a or b, and that this investor always "prices the asset".

Let $p_h = \begin{bmatrix} p_h(0) \\ p_h(1) \end{bmatrix}$ be the equilibrium price vector when all investors are of type h.

The price today equals the expected discounted value of tomorrow's dividend and tomorrow's price of the asset:

$$p_h(s) = \beta \left(P_h(s,0)(0 + p_h(0)) + P_h(s,1)(1 + p_h(1)) \right), \quad s = 0, 1 \tag{135.1}$$

These equations imply that the equilibrium price vector is

$$\begin{bmatrix} p_h(0) \\ p_h(1) \end{bmatrix} = \beta [I - \beta P_h]^{-1} P_h \begin{bmatrix} 0 \\ 1 \end{bmatrix} \tag{135.2}$$

The first two rows of the table report $p_a(s)$ and $p_b(s)$.

Here's a function that can be used to compute these values

```
def price_single_beliefs(transition, dividend_payoff, β=.75):
    """
    Function to Solve Single Beliefs
    """
    # First compute (I - βQ)
    A = np.eye(transition.shape[0]) - β * transition

    # Next compute βQd
    b = β * (transition @ dividend_payoff)

    # Solve linear system
    prices = la.solve(A, b)

    return prices
```

Single belief prices as benchmarks

These equilibrium prices under homogeneous beliefs are important benchmarks for the subsequent analysis.

- $p_h(s)$ tells what a type h investor thinks is the "fundamental value" of the asset.
- Here "fundamental value" means the expected discounted present value of future dividends.

We will compare these fundamental values of the asset with equilibrium values when traders have different beliefs.

135.3.3 Pricing under heterogeneous beliefs

There are several cases to consider.

The first is when both types of agents have sufficient wealth to purchase all of the asset themselves.

In this case, the marginal investor who prices the asset is the more optimistic type so that the equilibrium price $\bar{p}$ satisfies Harrison and Kreps's key equation:

$$\bar{p}(s) = \beta \max \left\{ P_a(s,0)\bar{p}(0) + P_a(s,1)(1+\bar{p}(1)),\ P_b(s,0)\bar{p}(0) + P_b(s,1)(1+\bar{p}(1)) \right\} \tag{135.3}$$

for $s = 0, 1$.

In the above equation, the max on the right side is over the two prospective values of next period's payout from owning the asset.

The marginal investor who prices the asset in state s is of type a if

$$P_a(s,0)\bar{p}(0) + P_a(s,1)(1+\bar{p}(1)) > P_b(s,0)\bar{p}(0) + P_b(s,1)(1+\bar{p}(1))$$

The marginal investor is of type b if

$$P_a(s,0)\bar{p}(0) + P_a(s,1)(1+\bar{p}(1)) < P_b(s,0)\bar{p}(0) + P_b(s,1)(1+\bar{p}(1))$$

Thus the marginal investor is the (temporarily) optimistic type.

Equation (135.3) is a functional equation that, like a Bellman equation, can be solved by

- starting with a guess for the price vector $\bar{p}$ and
- iterating to convergence on the operator that maps a guess $\bar{p}^j$ into an updated guess $\bar{p}^{j+1}$ defined by the right side of (135.3), namely

$$\bar{p}^{j+1}(s) = \beta \max \left\{ P_a(s,0)\bar{p}^j(0) + P_a(s,1)(1+\bar{p}^j(1)),\ P_b(s,0)\bar{p}^j(0) + P_b(s,1)(1+\bar{p}^j(1)) \right\} \tag{135.4}$$

for $s = 0, 1$.

The third row of the table labeled p_o reports equilibrium prices that solve the functional equation when $\beta = .75$.

Here the type that is optimistic about s_{t+1} prices the asset in state s_t.

It is instructive to compare these prices with the equilibrium prices for the homogeneous belief economies that solve under beliefs P_a and P_b reported in the rows labeled p_a and p_b, respectively.

Equilibrium prices p_o in the heterogeneous beliefs economy evidently exceed what any prospective investor regards as the fundamental value of the asset in each possible state.

Nevertheless, the economy recurrently visits a state that makes each investor want to purchase the asset for more than he believes its future dividends are worth.

An investor is willing to pay more than what he believes is warranted by fundamental value of the prospective dividend stream because he expects to have the option later to sell the asset to another investor who will value the asset more highly than he will then.

- Investors of type a are willing to pay the following price for the asset

$$\hat{p}_a(s) = \begin{cases} \bar{p}(0) & \text{if } s_t = 0 \\ \beta(P_a(1,0)\bar{p}(0) + P_a(1,1)(1+\bar{p}(1))) & \text{if } s_t = 1 \end{cases}$$

- Investors of type b are willing to pay the following price for the asset

$$\hat{p}_b(s) = \begin{cases} \beta(P_b(0,0)\bar{p}(0) + P_b(0,1)(1+\bar{p}(1))) & \text{if } s_t = 0 \\ \bar{p}(1) & \text{if } s_t = 1 \end{cases}$$

Evidently, $\hat{p}_a(1) < \bar{p}(1)$ and $\hat{p}_b(0) < \bar{p}(0)$.

Investors of type a want to sell the asset in state 1 while investors of type b want to sell it in state 0.

- The asset changes hands whenever the state changes from 0 to 1 or from 1 to 0.
- The valuations $\hat{p}_a(s)$ and $\hat{p}_b(s)$ are displayed in the fourth and fifth rows of the table.
- Even pessimistic investors who don't buy the asset think that it is worth more than they think future dividends are worth.

Here's code to solve for $\bar{p}$, $\hat{p}_a$ and $\hat{p}_b$ using the iterative method described above

```
def price_optimistic_beliefs(transitions, dividend_payoff, β=.75,
                             max_iter=50000, tol=1e-16):
    """
    Function to Solve Optimistic Beliefs
    """
    # We will guess an initial price vector of [0, 0]
    p_new = np.array([[0], [0]])
    p_old = np.array([[10.], [10.]])

    # We know this is a contraction mapping, so we can iterate to conv
    for i in range(max_iter):
        p_old = p_new
        p_new = β * np.max([q @ p_old
                            + q @ dividend_payoff for q in transitions],
                            axis=0)

        # If we succeed in converging, break out of for loop
        if np.max(np.sqrt((p_new - p_old)**2)) < tol:
            break

    ptwiddle = β * np.min([q @ p_old
                           + q @ dividend_payoff for q in transitions],
                           axis=0)

    phat_a = np.array([p_new[0], ptwiddle[1]])
    phat_b = np.array([ptwiddle[0], p_new[1]])

    return p_new, phat_a, phat_b
```

135.3.4 Insufficient funds

Outcomes differ when the more optimistic type of investor has insufficient wealth — or insufficient ability to borrow enough — to hold the entire stock of the asset.

In this case, the asset price must adjust to attract pessimistic investors.

Instead of equation (135.3), the equilibrium price satisfies

$$\check{p}(s) = \beta \min \{P_a(s,0)\check{p}(0) + P_a(s,1)(1+\check{p}(1)),\ P_b(s,0)\check{p}(0) + P_b(s,1)(1+\check{p}(1))\} \tag{135.5}$$

and the marginal investor who prices the asset is always the one that values it *less* highly than does the other type.

Now the marginal investor is always the (temporarily) pessimistic type.

Notice from the sixth row of that the pessimistic price p_o is lower than the homogeneous belief prices p_a and p_b in both states.

When pessimistic investors price the asset according to (135.5), optimistic investors think that the asset is underpriced.

If they could, optimistic investors would willingly borrow at a one-period risk-free gross interest rate β^{-1} to purchase more of the asset.

Implicit constraints on leverage prohibit them from doing so.

When optimistic investors price the asset as in equation (135.3), pessimistic investors think that the asset is overpriced and would like to sell the asset short.

Constraints on short sales prevent that.

Here's code to solve for $\check{p}$ using iteration

```
def price_pessimistic_beliefs(transitions, dividend_payoff, β=.75,
                              max_iter=50000, tol=1e-16):
    """
    Function to Solve Pessimistic Beliefs
    """
    # We will guess an initial price vector of [0, 0]
    p_new = np.array([[0], [0]])
    p_old = np.array([[10.], [10.]])

    # We know this is a contraction mapping, so we can iterate to conv
    for i in range(max_iter):
        p_old = p_new
        p_new = β * np.min([q @ p_old
                            + q @ dividend_payoff for q in transitions],
                           axis=0)

        # If we succeed in converging, break out of for loop
        if np.max(np.sqrt((p_new - p_old)**2)) < tol:
            break

    return p_new
```

135.3.5 Further interpretation

Jose Scheinkman [Scheinkman, 2014] interprets the Harrison-Kreps model as a model of a bubble — a situation in which an asset price exceeds what every investor thinks is merited by his or her beliefs about the value of the asset's underlying dividend stream.

Scheinkman stresses these features of the Harrison-Kreps model:

- High volume occurs when the Harrison-Kreps pricing formula (135.3) prevails.
- Type a investors sell the entire stock of the asset to type b investors every time the state switches from $s_t = 0$ to $s_t = 1$.
- Type b investors sell the asset to type a investors every time the state switches from $s_t = 1$ to $s_t = 0$.

Scheinkman takes this as a strength of the model because he observes high volume during *famous bubbles*.

- If the *supply* of the asset is increased sufficiently either physically (more "houses" are built) or artificially (ways are invented to short sell "houses"), bubbles end when the asset supply has grown enough to outstrip optimistic investors' resources for purchasing the asset.
- If optimistic investors finance their purchases by borrowing, tightening leverage constraints can extinguish a bubble.

Scheinkman extracts insights about the effects of financial regulations on bubbles.

He emphasizes how limiting short sales and limiting leverage have opposite effects.

135.4 Exercises

Exercise 135.4.1

This exercise invites you to recreate the summary table using the functions we have built above.

s_t	0	1
p_a	1.33	1.22
p_b	1.45	1.91
p_o	1.85	2.08
p_p	1	1
$\hat{p}_a$	1.85	1.69
$\hat{p}_b$	1.69	2.08

You will want first to define the transition matrices and dividend payoff vector.

In addition, below we'll add an interpretation of the row corresponding to p_o by inventing two additional types of agents, one of whom is **permanently optimistic**, the other who is **permanently pessimistic**.

We construct subjective transition probability matrices for our permanently optimistic and permanently pessimistic investors as follows.

The permanently optimistic investors(i.e., the investor with the most optimistic beliefs in each state) believes the transition matrix

$$P_o = \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{4} & \frac{3}{4} \end{bmatrix}$$

The permanently pessimistic investor believes the transition matrix

$$P_p = \begin{bmatrix} \frac{2}{3} & \frac{1}{3} \\ \frac{2}{3} & \frac{1}{3} \end{bmatrix}$$

We'll use these transition matrices when we present our solution of exercise 1 below.

Solution

First, we will obtain equilibrium price vectors with homogeneous beliefs, including when all investors are optimistic or pessimistic.

```
qa = np.array([[1/2, 1/2], [2/3, 1/3]])     # Type a transition matrix
qb = np.array([[2/3, 1/3], [1/4, 3/4]])     # Type b transition matrix
# Optimistic investor transition matrix
qopt = np.array([[1/2, 1/2], [1/4, 3/4]])
# Pessimistic investor transition matrix
qpess = np.array([[2/3, 1/3], [2/3, 1/3]])

dividendreturn = np.array([[0], [1]])

transitions = [qa, qb, qopt, qpess]
labels = ['p_a', 'p_b', 'p_optimistic', 'p_pessimistic']

for transition, label in zip(transitions, labels):
    print(label)
    print("=" * 20)
    s0, s1 = np.round(price_single_beliefs(transition, dividendreturn), 2)
    print(f"State 0: {s0}")
    print(f"State 1: {s1}")
    print("-" * 20)
```

```
p_a
====================
State 0: [1.33]
State 1: [1.22]
--------------------
p_b
====================
State 0: [1.45]
State 1: [1.91]
--------------------
p_optimistic
====================
State 0: [1.85]
State 1: [2.08]
--------------------
p_pessimistic
====================
State 0: [1.]
State 1: [1.]
--------------------
```

We will use the price_optimistic_beliefs function to find the price under heterogeneous beliefs.

```
opt_beliefs = price_optimistic_beliefs([qa, qb], dividendreturn)
labels = ['p_optimistic', 'p_hat_a', 'p_hat_b']

for p, label in zip(opt_beliefs, labels):
    print(label)
    print("=" * 20)
    s0, s1 = np.round(p, 2)
    print(f"State 0: {s0}")
    print(f"State 1: {s1}")
    print("-" * 20)
```

```
p_optimistic
====================
State 0: [1.85]
State 1: [2.08]
--------------------
p_hat_a
====================
State 0: [1.85]
State 1: [1.69]
--------------------
p_hat_b
====================
State 0: [1.69]
State 1: [2.08]
--------------------
```

Notice that the equilibrium price with heterogeneous beliefs is equal to the price under single beliefs with **permanently optimistic** investors - this is due to the marginal investor in the heterogeneous beliefs equilibrium always being the type who is temporarily optimistic.

CHAPTER

SIX

SPECULATIVE BEHAVIOR WITH BAYESIAN LEARNING

Contents

- *Speculative Behavior with Bayesian Learning*
  - *Overview*
  - *Structure of the model*
  - *Information and beliefs*
  - *Source of heterogeneous priors*
  - *Beta priors*
  - *Market prices with learning*
  - *Two Traders*
  - *Concluding remarks*
  - *Exercise*

136.1 Overview

This lecture describes how Morris [1996] extended the Harrison–Kreps model [Harrison and Kreps, 1978] of speculative asset pricing.

Like Harrison and Kreps's model, Morris's model determines the price of a dividend-yielding asset that is traded by risk-neutral investors who have heterogeneous beliefs.

The Harrison-Kreps model assumes that the traders have dogmatic, hard-wired beliefs about the asset's dividend stream.

Morris replaced Harrison and Kreps's traders with hard-wired beliefs about the dividend stream with traders who use Bayes' Law to update their beliefs about prospective dividends as new dividend data arrive.

Note

Morris's traders don't use data on past prices of the asset to update their beliefs about the dividend process.

Key features of the environment in Morris's model include:

- All traders share a set of statistical models for prospective dividends

- A single parameter indexes the set of statistical models
- All traders observe the same dividend history
- All traders use Bayes' Law to update beliefs
- Traders have different initial *prior distributions* over the parameter
- Traders' *posterior distributions* over the parameter eventually merge
- Before their posterior distributions merge, traders disagree about the predictive density over prospective dividends
  - therefore they disagree about the value of the asset

Just as in the hard-wired beliefs model of Harrison and Kreps, those differences of opinion induce investors to engage in *speculative behavior* in the following sense:

- sometimes they are willing to pay more for the asset than what they think is its "fundamental" value, i.e., the expected discounted value of its prospective dividend stream

Prior to reading this lecture, you might want to review the following quantecon lectures:

- *Harrison-Kreps model*
- *Likelihood ratio processes*
- *Bayesian versus frequentist statistics*

Let's start with some standard imports:

```
import numpy as np
import matplotlib.pyplot as plt
```

136.2 Structure of the model

There is a fixed supply of shares of an asset.

Each share entitles its owner to a stream of *binary* i.i.d. dividends $\{d_t\}$ where

$$d_{t+1} \in \{0, 1\}$$

The dividend at time t equals 1 with unknown probability $\theta \in (0, 1)$ and equals 0 with probability $1 - \theta$.

Unlike [Harrison and Kreps, 1978] where traders have hard-wired beliefs about a Markov transition matrix, in Morris's model:

- The true dividend probability θ is unknown
- Traders have *prior beliefs* about θ
- Traders observe dividend realizations and update beliefs via Bayes' Law

There is a finite set $\mathcal{I}$ of *risk-neutral* traders.

All traders have the same discount factor $\beta \in (0, 1)$.

- You can think of β as being related to a net risk-free interest rate r by $\beta = 1/(1 + r)$.

Owning the asset at the end of period t entitles the owner to dividends at time $t + 1, t + 2, \ldots$.

Because the dividend process is i.i.d., trader i thinks that the fundamental value of the asset is the capitalized value of the dividend stream, namely, $\sum_{j=1}^{\infty} \beta^j \hat{\theta}_i = \frac{\hat{\theta}_i}{r}$, where $\hat{\theta}_i$ is the mean of the trader's posterior distribution over θ.

136.2.1 Possible trades

Traders buy and sell the risky asset in competitive markets each period $t = 0, 1, 2, ...$ after dividends are paid.

As in Harrison-Kreps:

- The asset is traded *ex dividend*
- An owner of a share at the end of time t is entitled to the dividend at time $t + 1$
- An owner of a share at the end of period t also has the right to sell the share at time $t + 1$ after having received the dividend at time $t + 1$.

Short sales are prohibited.

This matters because it limits how pessimists can express their opinions:

- They *can* express themselves by selling their shares
- They *cannot* express themselves more emphatically by borrowing shares and immediately selling them

All traders have sufficient wealth to purchase the risky asset.

136.3 Information and beliefs

At time $t \geq 1$, all traders observe $(d_1, d_2, ..., d_t)$.

All traders update their subjective distribution over θ by applying Bayes' rule.

Traders have *heterogeneous priors* over the unknown dividend probability θ.

This heterogeneity in priors produces heterogeneous posterior beliefs.

136.4 Source of heterogeneous priors

Imputing different statistical models to agents inside a model is controversial.

Many game theorists and rational expectations applied economists think it is a bad idea.

While these economists often construct models in which agents have different *information*, they prefer to assume that all of the agents inside their model always share the same statistical model – i.e., the same joint probability distribution over the random process being modeled.

For a statistician or an economic theorist, a statistical model is a joint probability distribution that is characterized by a known parameter vector.

When working with a *set* of statistical models swept out by parameters, say θ in a known set Θ, economic theorists reduce the set of models to a single model by imputing to all agents inside the model the same prior probability distribution over θ.

Note

A set of statistical models that has a particular geometric structure is called a manifold of statistical models. Morris endows traders with a shared manifold of statistical models.

Proceeding in this way adheres to the *Harsanyi Common Priors Doctrine*.

[Harsanyi, 1967], [Harsanyi, 1968], [Harsanyi, 1968] argued that if two rational agents have the same information and the same reasoning capabilities, they will have the same joint probability distribution over outcomes of interest.

Harsanyi interpreted disagreements about prospective outcomes as arising from differences in agents' information sets, not differences in their statistical models.

Evidently, [Harrison and Kreps, 1978] departed from the Harsanyi common statistical model assumption when they hard-wired dogmatic disparate beliefs.

Morris [1996] abandons the Harsanyi doctrine less completely than Harrison and Kreps had.

- Morris does assume that agents share the same set of statistical models, but ...
- Morris assumes that they have different initial prior distributions over the parameter that indexes the models

Morris's agents express their initial ignorance about the parameter differently – they have different priors.

Morris defends his assumption by alluding to the apparent "mispricing" of initial public offerings presented by [Miller, 1977].

Miller described a situation in which agents have access to little or no data about a new enterprise.

Morris wanted his traders to be open to changing their opinions as information about the parameter arrives.

Knowledgeable statisticians have been known to disagree about an appropriate prior.

For example, Morris described *different* respectable ways to express "maximal ignorance" about the parameter of a Bernoulli distribution

- a uniform distribution on $[0, 1]$
- a Jeffreys prior [Jeffreys, 1946] that is invariant to reparameterization; in the present situation, the Jeffreys prior takes the form of a Beta distribution with parameters $.5, .5$

Is one of these priors more "rational" than the other?

Morris thinks not.

136.5 Beta priors

For tractability, assume trader i has a Beta prior over the dividend probability

$$\theta \sim \text{Beta}(a_i, b_i)$$

where $a_i, b_i > 0$ are the prior parameters.

Note

The Beta distribution also appears in these quantecon lectures *Statistical Divergence Measures*, *Likelihood Ratio Processes*, *Job Search IX: Search with Learning*.

Suppose trader i observes a history of t periods in which a total of s dividends are paid (i.e., s successes with dividend and $t - s$ failures without dividend).

By Bayes' rule, the posterior density over θ is:

$$\pi_i(\theta \mid s, t) = \frac{\theta^s (1-\theta)^{t-s} \pi_i(\theta)}{\int_0^1 \theta^s (1-\theta)^{t-s} \pi_i(\theta) d\theta}$$

where $\pi_i(\theta)$ is trader i's prior density.

Note

The Beta distribution is the conjugate prior for the Binomial likelihood. This means that when the prior is Beta(a_i, b_i) and we observe s successes in t trials, the posterior is Beta$(a_i + s, b_i + t - s)$.

The posterior mean (or expected dividend probability) is:

$$\mu_i(s,t) = \int_0^1 \theta \pi_i(\theta \mid s,t) d\theta = \mathbb{E}[\text{Beta}(a_i + s, b_i + t - s)] = \frac{a_i + s}{a_i + b_i + t}$$

Morris refers to $\mu_i(s,t)$ as trader i's **fundamental valuation** of the asset after history (s,t).

This is the probability trader i assigns to receiving a dividend next period.

It embeds trader i's updated belief about θ.

136.6 Market prices with learning

Fundamental valuations equal expected present values of dividends that our heterogeneous traders attach to the option of holding the asset *forever*.

The equilibrium price process is determined by the condition that the asset is held at time t by the trader who attaches the highest valuation to the asset at time t.

An owner of the asset has the option to sell it after receiving that period's dividend.

Traders take that into account.

That opens the possibility that a trader will be willing to pay more for the asset than that trader's fundamental valuation.

Definition 136.6.1 (Most Optimistic Valuation)

After history (s,t), the *most optimistic fundamental valuation* is:

$$\mu^*(s,t) = \max_{i \in \mathcal{I}} \mu_i(s,t)$$

Definition 136.6.2 (Equilibrium Asset Price)

Write $\tilde{p}(s,t,r)$ for the competitive equilibrium price of the risky asset (in current dollars) after history (s,t) when the interest rate is r.

The equilibrium price satisfies:

$$\tilde{p}(s,t,r) = \frac{1}{1+r}\Big[\mu^*(s,t)\{1 + \tilde{p}(s+1,t+1,r)\} + (1 - \mu^*(s,t))\tilde{p}(s,t+1,r)\Big]$$

The equilibrium price equals the highest expected discounted return among all traders from holding the asset to the next period.

Definition 136.6.3 (Normalized Price)

Define the normalized price as:

$$p(s,t,r) = r\tilde{p}(s,t,r)$$

Since the current "dollar" price of the riskless asset is $1/r$, this represents the price of the risky asset in terms of the riskless asset.

Substituting the preceding formula into the equilibrium condition gives:

$$p(s,t,r) = \frac{r}{1+r}\mu^*(s,t) + \frac{1}{1+r}\Big[\mu^*(s,t)p(s+1,t+1,r) + (1-\mu^*(s,t))p(s,t+1,r)\Big]$$

or equivalently:

$$p(s,t,r) = \mu^*(s,t) + \frac{r}{1+r}\Big[\mu^*(s,t)p(s+1,t+1,r) + (1-\mu^*(s,t))p(s,t+1,r) - \mu^*(s,t)\Big]$$

A price function that satisfies the equilibrium condition can be computed recursively.

Set $p^0(s,t,r) = 0$ for all (s,t,r), and define $p^{n+1}(s,t,r)$ by:

$$p^{n+1}(s,t,r) = \frac{r}{1+r}\mu^*(s,t) + \frac{1}{1+r}\Big[\mu^*(s,t)p^n(s+1,t+1,r) + (1-\mu^*(s,t))p^n(s,t+1,r)\Big]$$

The sequence $\{p^n(s,t,r)\}$ converges to the equilibrium price $p(s,t,r)$.

Definition 136.6.4 (Speculative Premium)

When the identity of the most optimistic trader can switch with future dividend realizations, the market price exceeds *every* trader's fundamental valuation.

In normalized units:

$$p(s,t,r) > \mu_i(s,t) \quad \text{for all } i \in \mathcal{J}$$

Define the **speculative premium** as:

$$p(s,t,r) - \mu^*(s,t) > 0$$

136.7 Two Traders

We now focus on an example with two traders with Beta priors with parameters (a_1, b_1) and (a_2, b_2).

Definition 136.7.1 (Rate Dominance (Beta Priors))

Trader 1 **rate-dominates** trader 2 if:

$$a_1 \geq a_2 \quad \text{and} \quad b_1 \leq b_2$$

Theorem 136.7.1 (Global Optimist (Two Traders))

For two traders with Beta priors:

1. If trader 1 rate-dominates trader 2, then trader 1 is a **global optimist**: $\mu_1(s,t) \geq \mu_2(s,t)$ for all histories (s,t)
2. In this case where $p(s,t,r) = \mu_1(s,t)$ for all (s,t,r), there is *no speculative premium*.

When neither trader rate-dominates the other, the identity of the most optimistic trader can switch as dividends accrue.

Along a history in which perpetual switching occurs, the price of the asset strictly exceeds both traders' fundamental valuations so long as traders continue to disagree:

$$p(s,t,r) > \max\{\mu_1(s,t), \mu_2(s,t)\}$$

Thus, along such a history, there is a persistent speculative premium.

136.7.1 Implementation

For computational tractability, let's work with a finite horizon T and solve by backward induction.

Note

On page 1122, Morris [1996] provides an argument that the limit as $T \to +\infty$ of such finite-horizon economies provides a useful selection algorithm that excludes additional equilibria that involve a Ponzi-scheme price component that Morris dismisses as fragile.

Following *Definition 136.6.2*, we use the discount factor parameterization $\beta = 1/(1+r)$ and compute dollar prices $\tilde{p}(s,t)$ via:

$$\tilde{p}(s,t) = \beta \max_{i \in \{1,2\}} \left[\mu_i(s,t)\{1 + \tilde{p}(s+1,t+1)\} + (1-\mu_i(s,t))\tilde{p}(s,t+1)\right]$$

We set the terminal price $\tilde{p}(s,T)$ equal to the perpetuity value under the most optimistic belief.

```
def posterior_mean(a, b, s, t):
    """
    Compute posterior mean μ_i(s,t) for Beta(a, b) prior.
    """
    return (a + s) / (a + b + t)

def perpetuity_value(a, b, s, t, β=.75):
    """
    Compute perpetuity value (β/(1-β)) * μ_i(s,t).
    """
    return (β / (1 - β)) * posterior_mean(a, b, s, t)

def price_learning_two_agents(prior1, prior2, β=.75, T=200):
    """
    Compute \tilde p(s,t) for two Beta-prior traders via backward induction.
    """
    a1, b1 = prior1
    a2, b2 = prior2
    price_array = np.zeros((T+1, T+1))
```

(continues on next page)

(continued from previous page)

```
    # Terminal condition: set to perpetuity value under max belief
    for s in range(T+1):
        perp1 = perpetuity_value(a1, b1, s, T, β)
        perp2 = perpetuity_value(a2, b2, s, T, β)
        price_array[s, T] = max(perp1, perp2)

    # Backward induction
    for t in range(T-1, -1, -1):
        for s in range(t, -1, -1):
            μ1 = posterior_mean(a1, b1, s, t)
            μ2 = posterior_mean(a2, b2, s, t)

            # One-step continuation values under each trader's beliefs
            cont1 = μ1 * (1.0 + price_array[s+1, t+1]) \
                    + (1.0 - μ1) * price_array[s, t+1]
            cont2 = μ2 * (1.0 + price_array[s+1, t+1]) \
                    + (1.0 - μ2) * price_array[s, t+1]
            price_array[s, t] = β * max(cont1, cont2)

    def μ1_fun(s, t):
        return posterior_mean(a1, b1, s, t)
    def μ2_fun(s, t):
        return posterior_mean(a2, b2, s, t)

    return price_array, μ1_fun, μ2_fun
```

136.7.2 Case A: global optimist (no premium)

Pick priors with rate dominance, e.g., trader 1: Beta$(a_1, b_1) = (2, 1)$ and trader 2: $(a_2, b_2) = (1, 2)$.

Trader 1 is the global optimist, so the normalized price equals trader 1's fundamental valuation: $p(s, t, r) = \mu_1(s, t)$.

```
β = 0.75
price_go, μ1_go, μ2_go = price_learning_two_agents(
        (2,1), (1,2), β=β, T=200)

perpetuity_1 = (β / (1 - β)) * μ1_go(0, 0)
perpetuity_2 = (β / (1 - β)) * μ2_go(0, 0)

print("Price at (0, 0) =", price_go[0,0])
print("Valuation of trader 1 at (0, 0) =", perpetuity_1)
print("Valuation of trader 2 at (0, 0) =", perpetuity_2)
```

```
Price at (0, 0) = 2.0
Valuation of trader 1 at (0, 0) = 2.0
Valuation of trader 2 at (0, 0) = 1.0
```

The price equals trader 1's perpetuity value.

136.7.3 Case B: perpetual switching (positive premium)

Now assume trader 1 has Beta$(1, 1)$, trader 2 has Beta$(1/2, 1/2)$.

These produce crossing posteriors, so there is no global optimist and the price exceeds both fundamentals early on.

```
price_ps, μ1_ps, μ2_ps = price_learning_two_agents(
                                (1,1), (0.5,0.5), β=β, T=200)

price_00 = price_ps[0,0]
μ1_00 = μ1_ps(0,0)
μ2_00 = μ2_ps(0,0)

perpetuity_1 = (β / (1 - β)) * μ1_ps(0, 0)
perpetuity_2 = (β / (1 - β)) * μ2_ps(0, 0)

print("Price at (0, 0) =", np.round(price_00, 6))
print("Valuation of trader 1 at (0, 0) =", perpetuity_1)
print("Valuation of trader 2 at (0, 0) =", perpetuity_2)
```

```
Price at (0, 0) = 1.599322
Valuation of trader 1 at (0, 0) = 1.5
Valuation of trader 2 at (0, 0) = 1.5
```

The resulting premium reflects the option value of reselling to whichever trader becomes temporarily more optimistic as dividends arrive sequentially.

Within this setting, we can reproduce two key figures reported in Morris [1996]

```
def normalized_price_two_agents(prior1, prior2, r, T=250):
    """Return p(s,t,r) = r \tilde p(s,t,r) for two traders."""
    β = 1.0 / (1.0 + r)
    price_array, *_ = price_learning_two_agents(prior1, prior2, β=β, T=T)
    return r * price_array

# Figure I: p*(0,0,r) as a function of r
r_grid = np.linspace(1e-3, 5.0, 200)
priors = ((1,1), (0.5,0.5))
p00 = np.array([normalized_price_two_agents(
                priors[0], priors[1], r, T=300)[0,0]
                for r in r_grid])

fig, ax = plt.subplots()
ax.plot(r_grid, p00, lw=2)
ax.set_xlabel(r'$r$')
ax.set_ylabel(r'$p^*(0,0,r)$')
ax.axhline(0.5, color='C1', linestyle='--')
plt.show()
```

In the first figure, notice that:

- The resale option pushes the normalized price $p^*(0, 0, r)$ above fundamentals (0.5) for any finite r.
- As r increases (β decreases), the option value fades and $p^*(0, 0, r) \to 0.5$.
- At $r = 0.05$ the premium is about 8–9%, consistent with Morris (1996, Section IV).

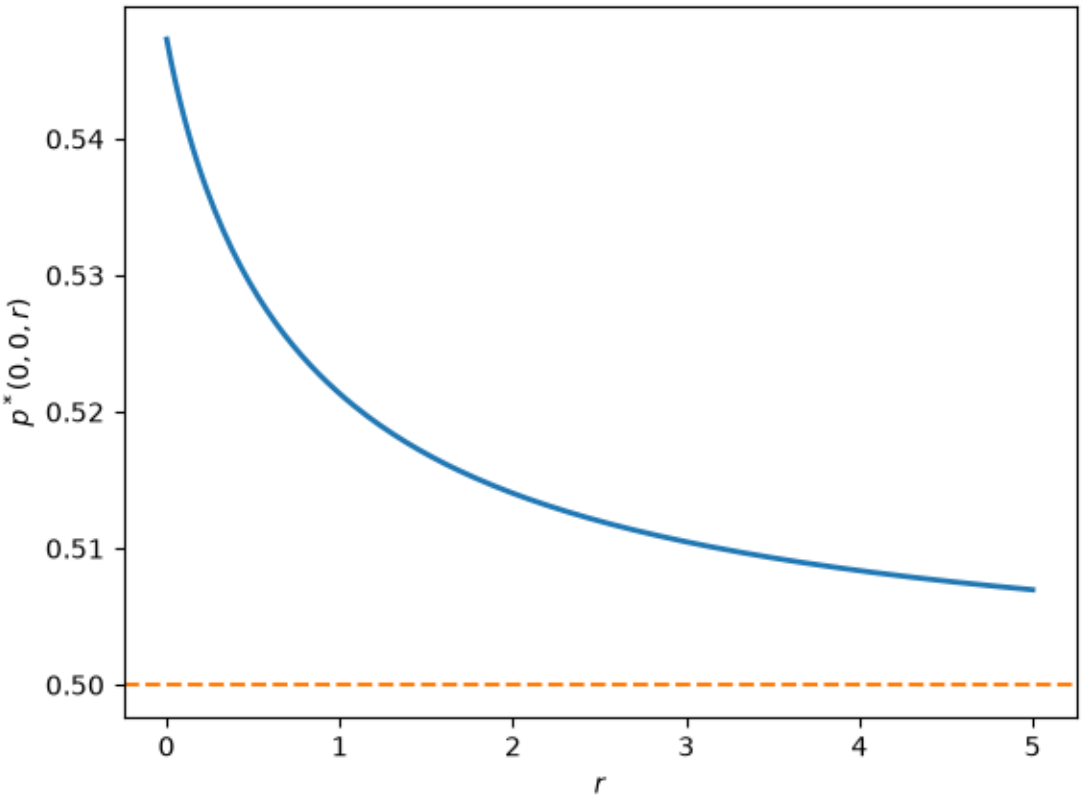


Fig. 136.1: Normalized price against interest rate

```
# Figure II: p*(t/2,t,0.05) as a function of t
r = 0.05
T = 60
p_mat = normalized_price_two_agents(priors[0], priors[1], r, T=T)
t_vals = np.arange(0, 54, 2)
s_vals = t_vals // 2
y = np.array([p_mat[s, t] for s, t in zip(s_vals, t_vals)])

fig, ax = plt.subplots()
ax.plot(t_vals, y, lw=2)
ax.set_xlabel(r'$t$')
ax.set_ylabel(r'$p^*(t/2,t,0.05)$')
ax.axhline(0.5, color='C1', linestyle='--')
plt.show()

p0 = p_mat[0,0]
μ0 = 0.5
print("Initial normalized premium at r=0.05 (%):",
      np.round(100 * (p0 / μ0 - 1.0), 2))
```

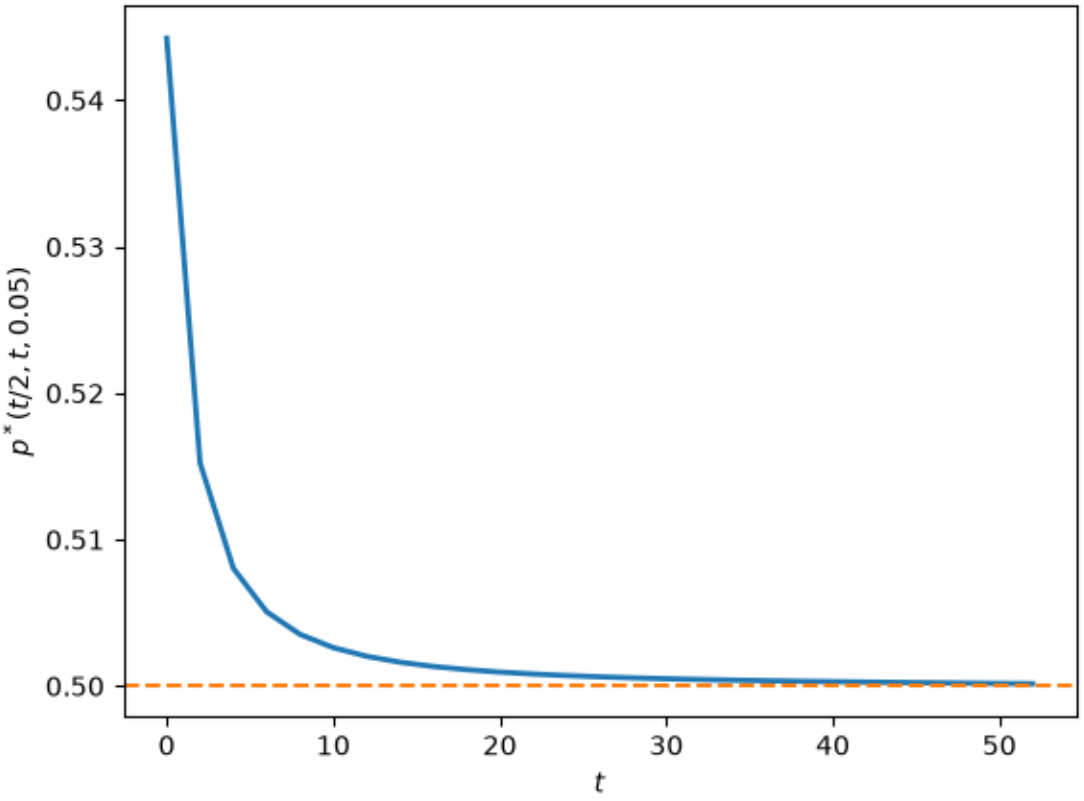


Fig. 136.2: Normalized price against time

```
Initial normalized premium at r=0.05 (%): 8.85
```

In the second figure, notice that:

- Along the symmetric path $s = t/2$, both traders' fundamental valuations equal 0.5 at every t, yet the price starts above 0.5 and declines toward 0.5 as learning reduces disagreement and the resale option loses value.

136.7.4 General N-trader extension

The same recursion extends to any finite set of Beta priors $\{(a_i, b_i)\}_{i=1}^N$ by taking a max over i each period.

```
def price_learning(priors, β=0.75, T=200):
    """
    N-trader version with heterogeneous Beta priors.
    """
    price_array = np.zeros((T+1, T+1))

    def perp_i(i, s, t):
        a, b = priors[i]
        return perpetuity_value(a, b, s, t, β)

    # Terminal condition
    for s in range(T+1):
        price_array[s, T] = max(
            perp_i(i, s, T) for i in range(len(priors)))

    # Backward induction
    for t in range(T-1, -1, -1):
        for s in range(t, -1, -1):
            conts = []
            for (a, b) in priors:
                μ = posterior_mean(a, b, s, t)
                conts.append(μ *
                (1.0 + price_array[s+1, t+1])
                     + (1.0 - μ) * price_array[s, t+1])
            price_array[s, t] = β * max(conts)

    return price_array

β = 0.75
priors = [(1,1), (0.5,0.5), (3,2)]
price_N = price_learning(priors, β=β, T=150)

# Compute valuations for each trader at (0,0)
μ_vals = [posterior_mean(a, b, 0, 0) for a, b in priors]
perp_vals = [(β / (1 - β)) * μ for μ in μ_vals]

print("Three-trader example at (s,t)=(0,0):")
print(f"Price at (0,0) = {np.round(price_N[0,0], 6)}")
print(f"\nTrader valuations:")
for i, (μ, perp) in enumerate(zip(μ_vals, perp_vals), 1):
    print(f"  Trader {i} = {np.round(perp, 6)}")
```

```
Three-trader example at (s,t)=(0,0):
Price at (0,0) = 1.937972

Trader valuations:
  Trader 1 = 1.5
  Trader 2 = 1.5
  Trader 3 = 1.8
```

Note that the asset price is above all traders' valuations.

Morris tells us that no rate dominance exists in this case.

Let's verify this using the code below

```
dominant = None
for i in range(len(priors)):
    is_dom = all(
        priors[i][0] >= priors[j][0] and priors[i][1] <= priors[j][1]
                for j in range(len(priors)) if i != j)
    if is_dom:
        dominant = i
        break

if dominant is not None:
    print(f"\nTrader {dominant+1} is the global optimist (rate-dominant)")
else:
    print(f"\nNo global optimist and speculative premium exists")
```

```
No global optimist and speculative premium exists
```

Indeed, there is no global optimist and a speculative premium exists.

136.8 Concluding remarks

Morris [1996] uses his model to interpret a "hot issue" anomaly described by [Miller, 1977] according to which opening market prices of initial public offerings seem higher than values prices that emerge later.

136.9 Exercise

Exercise 136.9.1

Morris [Morris, 1996] provides a sharp characterization of when speculative bubbles arise.

The key condition is that there is no *global optimist*.

In this exercise, you will verify this condition for the following sets of traders with Beta priors:

1. Trader 1: Beta$(2,1)$, Trader 2: Beta$(1,2)$
2. Trader 1: Beta$(1,1)$, Trader 2: Beta$(1/2,1/2)$
3. Trader 1: Beta$(3,1)$, Trader 2: Beta$(2,1)$, Trader 3: Beta$(1,2)$
4. Trader 1: Beta$(1,1)$, Trader 2: Beta$(1/2,1/2)$, Trader 3: Beta$(3/2,3/2)$

Solution

Here is one solution:

```
def check_rate_dominance(priors):
    """
    Check if any trader rate-dominates all others.
    """
    N = len(priors)

    for i in range(N):
        a_i, b_i = priors[i]
```

```
        is_dominant = True

        for j in range(N):
            if i == j:
                continue
            a_j, b_j = priors[j]

            # Check rate dominance condition
            if not (a_i >= a_j and b_i <= b_j):
                is_dominant = False
                break

        if is_dominant:
            return i

    return None

# Test cases
test_cases = [
    ([(2, 1), (1, 2)], "Global optimist exists"),
    ([(1, 1), (0.5, 0.5)], "Perpetual switching"),
    ([(3, 1), (2, 1), (1, 2)], "Three traders with dominant"),
    ([(1, 1), (0.5, 0.5), (1.5, 1.5)], "Three traders, no dominant")
]

for priors, description in test_cases:
    dominant = check_rate_dominance(priors)

    print(f"\n{description}")
    print(f"Priors: {priors}")
    print("=="*8)
    if dominant is not None:
        print(f"Trader {dominant+1} is the global optimist (rate-dominant)")
    else:
        print(f"No global optimist exists")
    print("=="*8 + "\n")
```

```
Global optimist exists
Priors: [(2, 1), (1, 2)]
================
Trader 1 is the global optimist (rate-dominant)
================


Perpetual switching
Priors: [(1, 1), (0.5, 0.5)]
================
No global optimist exists
================


Three traders with dominant
Priors: [(3, 1), (2, 1), (1, 2)]
================
Trader 1 is the global optimist (rate-dominant)
================
```

```
Three traders, no dominant
Priors: [(1, 1), (0.5, 0.5), (1.5, 1.5)]
================
No global optimist exists
================
```

CHAPTER

SEVEN

RATIONAL LEARNING AND RATIONAL EXPECTATIONS

Contents

- *Rational Learning and Rational Expectations*
  - *Overview*
  - *The economy*
  - *The rational learning equilibrium*
  - *Computing the equilibrium*
  - *Posterior concentration*
  - *Equilibrium trades and prices*
  - *Convergence of posterior assessments*
  - *Sharpening the convergence result*
  - *Obstacles to convergence*
  - *Learning within versus learning about a rational expectations equilibrium*
  - *Summary*
  - *Exercises*

137.1 Overview

This lecture explores an important question in economic theory: what can agents learn inside a rational expectations equilibrium?

This question is related to, but distinct from, the adaptive-learning question of whether a rational expectations equilibrium can be justified as a dynamic attractor.

The starting point is Bray and Kreps [1987], which gives a rigorous model of Bayesian learning inside a rational expectations equilibrium.

In a rational expectations equilibrium, agents use market prices to make inferences about other agents' private information.

Each agent knows the *statistical relationship* between prices and the underlying payoff-relevant variables and that relationship is *correct* given the equilibrium.

But this raises a question: where does that knowledge come from?

Their **rational learning** approach answers a narrower question.

Agents are uncertain about structural parameters, and the state space is enlarged to include, for each candidate value, the associated equilibrium price and allocation maps.

Agents then use Bayes' rule to update over those possibilities as prices and returns are observed.

This lecture develops that Bayesian formulation through an asset-market model.

The aim is to see what rational learning can explain, and where its limits appear, before turning to the computational illustration.

The discussion also connects to earlier work by Bray [1982], Bray and Savin [1986], and the rational expectations literature of Radner [1979], Grossman [1976], and Jordan [1982].

Let's start with the following imports

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.optimize import brentq
```

137.2 The economy

The model is a simple asset market built around the recurring information model of Grossman and Stiglitz [1980], repeated across infinitely many dates.

137.2.1 Agents and assets

Each date is economically disconnected from the others, so agents start each period afresh.

There are two types of agents and two assets:

- A **safe asset** with net return normalized to zero.
- A **risky asset** endowed one unit per agent and traded at date t at spot price p_t.

At each date $t = 0, 1, 2, \ldots$ the risky asset yields a gross return r_t at date $t + 1/2$.

An informed signal s_t satisfies

$$r_t = s_t + \epsilon_t, \qquad s_t \sim N(\mu_s, \tau^2), \qquad \epsilon_t \sim N(0, \sigma^2),$$

where $\{s_t\}$ and $\{\epsilon_t\}$ are IID normal sequences and are mutually independent.

Common knowledge of the prior moments (μ_s, τ^2) is what makes the price observation informative about θ^I, as we will see.

There are two representative agents:

- **Agent I (informed)** observes s_t before trading at date t.
- **Agent U (uninformed)** observes the equilibrium price p_t but not s_t before trading.

Both agents observe the previous return before current trading.

137.2.2 Preferences

Agent $n \in \{I, U\}$ has constant absolute risk tolerance θ^n.

If agent n holds x^n units of the risky asset and y^n units of the safe asset between dates t and $t + 1/2$, period utility is

$$-\exp\left[-\frac{x^n r_t + y^n}{\theta^n}\right].$$

Thus $1/\theta^n$ is the coefficient of absolute risk aversion.

Given the signal s_t, the informed agent's demand is

$$x_t^I = \frac{\theta^I}{\sigma^2}(s_t - p_t). \tag{137.1}$$

Because each of the two agents is endowed with one unit of the risky asset, market clearing is

$$x_t^I + x_t^U = 2.$$

137.2.3 Rational expectations equilibrium

If all agents knew s_t, agent n would demand

$$x_t^n = \frac{\theta^n}{\sigma^2}(s_t - p_t).$$

With N agents and total risky-asset supply N, market clearing gives the **full communication equilibrium price**

$$p_t = s_t - \frac{N\sigma^2}{\sum_{n=1}^N \theta^n}. \tag{137.2}$$

Thus if $\sum_n \theta^n$ is known, the price fully reveals s_t.

Following Radner [1979], this is called a full communication rational expectations equilibrium.

Suppose now that θ^I is unknown to agent U.

Following Bray and Kreps [1987], we treat this uncertainty by *expanding the state space*: we let the unknown parameter θ^I become a coordinate of the state, alongside the per-period shocks (s_t, ϵ_t).

Formally, the state space is $\Omega = \Theta \times \Phi^\infty$, where $\Theta = [a, b]$ supports the unknown θ^I and Φ supports each (s_t, ϵ_t).

Agent U knows σ^2 and θ^U, and starts with a prior density over θ^I on $[a, b]$.

This expansion is what turns the learning problem into Bayesian inference inside a single rational expectations equilibrium on Ω.

At a date when agent U has posterior density f over θ^I, his own trade reveals $x_t^I = 2 - x_t^U$ through market clearing.

Combining this inferred x_t^I with (137.1), each candidate θ^I implies

$$s_t = p_t + \frac{\sigma^2 x_t^I}{\theta^I}. \tag{137.3}$$

After trading, agent U observes r_t.

Bayes' rule then updates the posterior over θ^I using the normal density of the signal implied by (137.3) conditional on the realized return.

This is the main object learned in the two-agent example.

Even in this simple case, the equilibrium can be defined recursively but closed-form prices are unavailable.

137.3 The rational learning equilibrium

The model has two pieces that interact at each date.

The first is the within-period equilibrium given the uninformed agent's current posterior on θ^I.

The second is the Bayesian update of that posterior after the period closes.

137.3.1 Uninformed demand given beliefs

Suppose at date t agent U has posterior density f_t on θ^I, supported on $[a, b]$.

Suppose the equilibrium price is p and the equilibrium informed trade $X^I = 2 - x^U$ has been inferred from market clearing.

Conditional on θ, equation (137.3) pins down the signal as $s_t(\theta) = \sigma^2 X^I/\theta + p$.

Two sources of information about θ are therefore present at the start of date t trading: the carried-over posterior $f_t(\theta)$ and the Gaussian prior $\phi_s(\cdot; \mu_s, \tau^2)$ on s_t that values some implied signals as more plausible than others.

Bayes' rule combines them into the *intra-period* posterior

$$f_t^{(p,X^I)}(\theta) \propto f_t(\theta)\, \phi_s\!\left(\frac{\sigma^2 X^I}{\theta} + p\,;\, \mu_s, \tau^2\right), \tag{137.4}$$

which is the posterior on θ^I that the agent actually uses to forecast r_t before r_t is observed.

Conditional on a candidate value θ, the excess payoff on one unit of the risky asset is

$$r_t - p = \frac{\sigma^2 X^I}{\theta} + \epsilon_t, \qquad \epsilon_t \sim N(0, \sigma^2).$$

Because CARA preferences have no wealth effects, agent U's problem reduces to

$$\max_{x^U} \mathbb{E}[u^U(x^U, r_t, p)], \qquad u^U(x^U, r_t, p) = -\exp\!\left(-\frac{x^U(r_t - p)}{\theta^U}\right),$$

where the expectation integrates over $\theta^I \sim f_t^{(p,X^I)}$ and ϵ_t.

Substituting the conditional excess payoff and using the normal moment-generating formula gives

$$\mathbb{E}[u^U] = -\exp\!\left(\frac{(x^U)^2\sigma^2}{2(\theta^U)^2}\right) \int_a^b f_t(\theta)\, \phi_s\!\left(\tfrac{\sigma^2 X^I}{\theta} + p; \mu_s, \tau^2\right) \exp\!\left(-\tfrac{x^U \sigma^2 X^I}{\theta\, \theta^U}\right) d\theta,$$

up to a θ-independent constant absorbed in normalisation.

Define the tilted weight

$$w(\theta;\, p, X^I, x^U) = f_t(\theta)\, \phi_s\!\left(\tfrac{\sigma^2 X^I}{\theta} + p; \mu_s, \tau^2\right) \exp\!\left(-\tfrac{x^U \sigma^2 X^I}{\theta\, \theta^U}\right). \tag{137.5}$$

The first-order condition rearranges to

$$\frac{x^U}{\theta^U} = X^I\, \frac{\int_a^b \theta^{-1}\, w(\theta;\, p, X^I, x^U)\, d\theta}{\int_a^b w(\theta;\, p, X^I, x^U)\, d\theta}. \tag{137.6}$$

The right-hand side is X^I multiplied by a tilted expectation of $1/\theta^I$ under the weighting in (137.5).

Equation (137.6) implicitly defines $x^U(p, X^I; f_t)$, the uninformed agent's optimal demand at observed price p, conjectured informed trade X^I, and prior posterior f_t.

Dependence on p enters through the prior weight ϕ_s: at higher prices, candidate values of θ that imply s_t above the prior mean become less plausible, so the agent's demand schedule slopes downward in p as expected.

137.3.2 Market clearing

Equilibrium requires that the informed and uninformed demands sum to the total endowment.

Substituting (137.1) and the implicit function $x^U(p, X^I; f_t)$, the equilibrium (p_t, X_t^I) satisfies the two equations

$$X_t^I = \frac{\theta^I}{\sigma^2}(s_t - p_t), \qquad X_t^I + x^U(p_t, X_t^I; f_t) = 2. \tag{137.7}$$

Eliminating X_t^I between the two leaves a single root-finding problem for p_t.

Combining the two equations, the equilibrium price has the form

$$p_t = s_t - \frac{\sigma^2 X_t^I}{\theta^I}. \tag{137.8}$$

When f_t collapses to a point mass at the true θ^I, equation (137.6) simplifies to $x^U/\theta^U = X^I/\theta^I$, and market clearing gives the full-communication allocation

$$X_t^I = \frac{2\theta^I}{\theta^I + \theta^U}, \qquad x_t^U = \frac{2\theta^U}{\theta^I + \theta^U}. \tag{137.9}$$

This is the full-communication allocation we will use to check the simulation.

137.3.3 Bayesian update

After trading, agent U observes (p_t, x_t^U, r_t).

Market clearing gives $X_t^I = 2 - x_t^U$, and equation (137.3) assigns a candidate $s_t(\theta) = \sigma^2 X_t^I/\theta + p_t$ to each θ.

Since $s_t \sim N(\mu_s, \tau^2)$ independently of $\epsilon_t \sim N(0, \sigma^2)$, the conditional density of s_t given r_t is Gaussian:

$$g(s \mid r) = \phi\left(s;\, \frac{\sigma^2\mu_s + \tau^2 r}{\sigma^2 + \tau^2},\, \frac{\sigma^2\tau^2}{\sigma^2 + \tau^2}\right),$$

where $\phi(\cdot; m, v)$ denotes the Normal density with mean m and variance v.

Bayes' rule then produces the posterior update

$$f_{t+1}(\theta) \propto f_t(\theta)\; g\left(\frac{\sigma^2 X_t^I}{\theta} + p_t \,\middle|\, r_t\right). \tag{137.10}$$

This is the rule we simulate below.

137.4 Computing the equilibrium

We discretise the support $[a, b]$ of θ^I on a fine grid and represent f_t as a vector of density values.

The three computational primitives are:

- `uninformed_demand` solves the FOC in (137.6) for $x^U(p, X^I; f)$ by root-finding.
- `equilibrium_price` solves the market-clearing system (137.7) for p_t.
- `bayes_update` applies (137.10) and renormalises.

```
def uninformed_demand(p, XI, f, θ_grid, θ_U, σ2, μ_s, τ2):
    """
    Solve the FOC for x^U(p, X^I; f), the uninformed
    agent's optimal demand given observed price p, conjectured
    informed trade XI, and carried-over posterior density f.
    """
    with np.errstate(divide='ignore'):
        log_f = np.log(f)
    s_implied = σ2 * XI / θ_grid + p
    log_phi_s = -0.5 * (s_implied - μ_s)**2 / τ2  # prior weight on s_t

    def foc(xU):
        z = xU * σ2 * XI / (θ_grid * θ_U)
        log_w = log_f + log_phi_s - z
        M = log_w.max()
        w = np.exp(log_w - M)
        num = np.sum(w / θ_grid)
        den = np.sum(w)
        return xU / θ_U - XI * num / den

    return brentq(foc, -50.0, 50.0, xtol=1e-10)
```

```
def equilibrium_price(s_t, θ_I_true, f, θ_grid, θ_U, σ2, μ_s, τ2):
    """
    Solve the market-clearing system for the equilibrium
    price p_t given signal s_t, true informed risk tolerance
    θ_I_true, and posterior f.
    """
    def mc_residual(p):
        XI = θ_I_true * (s_t - p) / σ2
        xU = uninformed_demand(p, XI, f, θ_grid, θ_U, σ2, μ_s, τ2)
        return XI + xU - 2.0

    return brentq(mc_residual, s_t - 10.0, s_t, xtol=1e-8)
```

```
def bayes_update(f, θ_grid, p_t, xU_t, r_t, σ2, τ2, μ_s):
    """
    Bayesian update of the posterior on θ^I given the date-t
    observations (p_t, x^U_t, r_t).
    """
    XI = 2.0 - xU_t
    s_mean = (σ2 * μ_s + τ2 * r_t) / (σ2 + τ2)
    s_var = σ2 * τ2 / (σ2 + τ2)
    s_implied = σ2 * XI / θ_grid + p_t

    log_like = -0.5 * (s_implied - s_mean)**2 / s_var
    log_like -= log_like.max()  # log-shift for stability
    f_new = f * np.exp(log_like)
    dθ = θ_grid[1] - θ_grid[0]
    f_new /= np.sum(f_new) * dθ
    return f_new
```

The simulation loop chains (s_t, ϵ_t) shocks through these three functions.

```
def simulate(θ_I_true, θ_U, σ2, μ_s, τ2,
             a, b, n_grid, T, prior=None, seed=42):
    """
```

(continues on next page)

(continued from previous page)

```
    Simulate T periods of the Bray-Kreps rational-learning equilibrium.
    """
    rng = np.random.default_rng(seed)
    θ_grid = np.linspace(a, b, n_grid)
    dθ = θ_grid[1] - θ_grid[0]

    if prior is None:
        f = np.ones(n_grid) / (b - a)
    else:
        f = prior(θ_grid)
        f /= np.sum(f) * dθ

    s_seq = rng.normal(μ_s, np.sqrt(τ2), T)
    eps_seq = rng.normal(0.0, np.sqrt(σ2), T)

    XI_path = np.empty(T)
    p_path = np.empty(T)
    r_path = np.empty(T)
    post_mean = np.empty(T + 1)
    post_var = np.empty(T + 1)
    post_mean[0] = np.sum(θ_grid * f) * dθ
    post_var[0] = np.sum((θ_grid - post_mean[0])**2 * f) * dθ

    snap_times = {0, 5, 20, 50, 100, T}
    snapshots = {0: f.copy()}

    for t in range(T):
        p_t = equilibrium_price(
            s_seq[t], θ_I_true, f, θ_grid, θ_U, σ2, μ_s, τ2
        )
        XI = θ_I_true * (s_seq[t] - p_t) / σ2
        xU = 2.0 - XI
        r_t = s_seq[t] + eps_seq[t]
        f = bayes_update(f, θ_grid, p_t, xU, r_t, σ2, τ2, μ_s)

        XI_path[t] = XI
        p_path[t] = p_t
        r_path[t] = r_t
        post_mean[t + 1] = np.sum(θ_grid * f) * dθ
        post_var[t + 1] = np.sum(
            (θ_grid - post_mean[t + 1])**2 * f
        ) * dθ
        if (t + 1) in snap_times:
            snapshots[t + 1] = f.copy()

    return dict(
        θ_grid=θ_grid,
        snapshots=snapshots,
        XI_path=XI_path,
        p_path=p_path,
        r_path=r_path,
        post_mean=post_mean,
        post_var=post_var,
    )
```

137.5 Posterior concentration

We run the simulation with a uniform prior on $[0.5, 4]$ and true $\theta^I = 2$.

```
params = dict(
    θ_I_true=2.0,
    θ_U=1.0,
    σ2=1.0,
    μ_s=1.0,
    τ2=1.0,
    a=0.5,
    b=4.0,
    n_grid=300,
    T=200,
    seed=42,
)

res = simulate(**params)
```

The first picture shows snapshots of the posterior density at selected dates.

```
fig, ax = plt.subplots(figsize=(10, 5))
for t, ft in sorted(res['snapshots'].items()):
    ax.plot(res['θ_grid'], ft, lw=2, label=f't = {t}')
ax.axvline(params['θ_I_true'], color='black', ls='--', lw=1.5,
           label=r'$\theta^I_{\rm true}$')
ax.set_xlabel(r'$\theta^I$')
ax.set_ylabel('posterior density')
ax.legend()
plt.tight_layout()
plt.show()
```

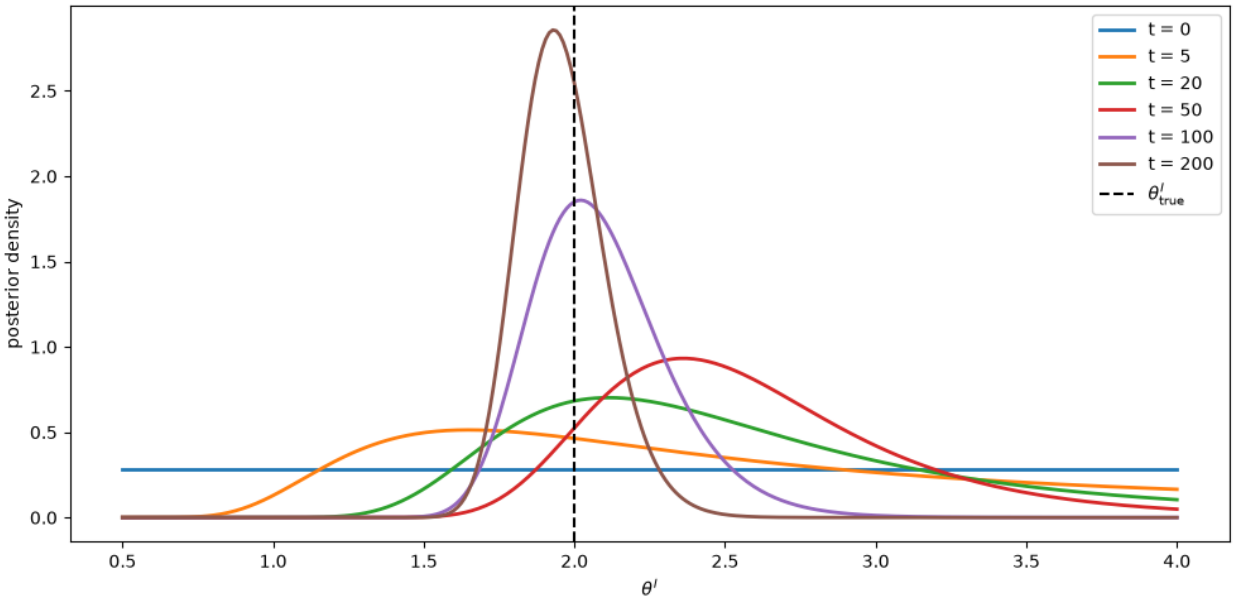


Fig. 137.1: posterior density over θ^I at selected dates

The posterior tightens around $\theta^I_{\text{true}} = 2$ as price and return data accumulate.

The next picture tracks the posterior mean and variance.

```
fig, axes = plt.subplots(1, 2, figsize=(13, 5))

ax = axes[0]
ax.plot(np.arange(params['T'] + 1), res['post_mean'], lw=2)
ax.axhline(params['θ_I_true'], color='red', ls='--', lw=2,
           label=r'$\theta^I_{\rm true}$')
ax.set_xlabel('$t$')
ax.set_ylabel(r'$E_t[\theta^I]$')
ax.legend()

ax = axes[1]
ax.plot(np.arange(params['T'] + 1), res['post_var'], lw=2)
ax.set_xlabel('$t$')
ax.set_ylabel(r'${\rm Var}_t[\theta^I]$')

plt.tight_layout()
plt.show()
```

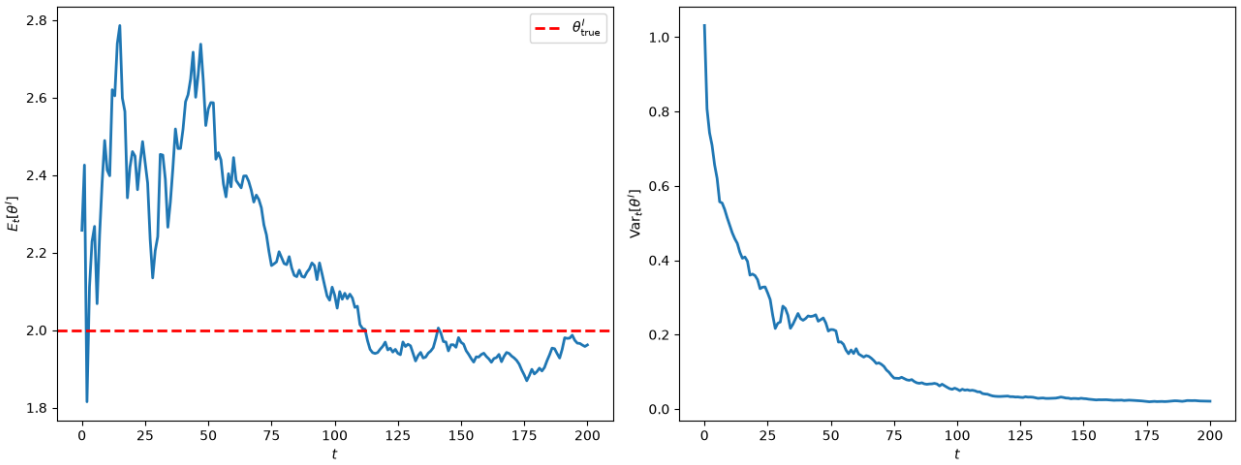


Fig. 137.2: posterior mean and variance over time

The posterior mean converges to the truth and the posterior variance vanishes.

This is the concrete manifestation of weak convergence of posteriors to a point mass at θ^I_{true}, which we describe in general terms below.

137.6 Equilibrium trades and prices

The equilibrium informed trade X_t^I depends on the current signal s_t, on θ_{true}^I, and on the carried-over posterior f_t, all through the market-clearing system (137.7).

As f_t tightens around θ_{true}^I, the average X_t^I approaches the full-communication allocation in (137.9).

```
XI_full = 2 * params['θ_I_true'] / (params['θ_I_true'] + params['θ_U'])
p_mean_full = params['μ_s'] - params['σ2'] * XI_full / params['θ_I_true']

fig, axes = plt.subplots(1, 2, figsize=(13, 5))

ax = axes[0]
ax.plot(res['XI_path'], lw=2, label='$X^I_t$ (learning)')
ax.axhline(XI_full, color='red', ls='--', lw=2,
           label='$X^I$ (full info)')
ax.set_xlabel('$t$')
ax.set_ylabel('$X^I_t$')
ax.legend()

ax = axes[1]
ax.plot(res['p_path'], lw=1.5, alpha=0.7, label='$p_t$ (learning)')
ax.axhline(p_mean_full, color='red', ls='--', lw=2,
           label='$E[p_t]$ (full info)')
ax.set_xlabel('$t$')
ax.set_ylabel('$p_t$')
ax.legend()

plt.tight_layout()
plt.show()
```

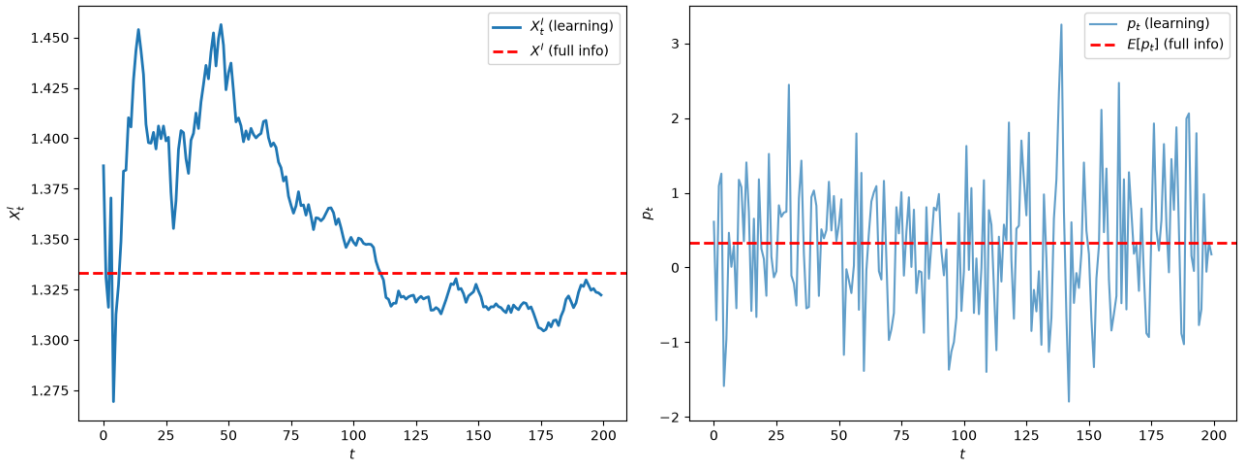


Fig. 137.3: equilibrium trade and prices over time

The left panel shows X_t^I approaching the full-information allocation as beliefs concentrate.

The right panel shows the price path, which fluctuates because p_t inherits the variation in s_t.

The simulation illustrates three features of this equilibrium.

The posterior density on θ^I concentrates around the true value, the posterior variance vanishes, and the equilibrium informed trade X_t^I converges to its full-information benchmark.

The next sections identify the theorems behind these outcomes and the assumptions they require.

We first state the two convergence results of Bray and Kreps [1987] for the abstract model, then specialize to the two-agent example to pin down the hypotheses that drive concentration on the true θ^I, and finally explain when those hypotheses can fail.

137.7 Convergence of posterior assessments

Let $(\Omega, \mathcal{F})$ be a measurable space carrying the equilibrium.

In the two-agent example, $\Omega = \Theta \times \Phi^\infty$, where $\Theta = [a, b] \times \{\theta^U\}$ collects the structural parameters, Φ collects the per-period shocks (s_t, ϵ_t), and $\mathcal{F}$ is the product Borel σ-field.

Agent n enters date 0 with a prior probability measure P^n on $(\Omega, \mathcal{F})$.

Let $G_t^n \subseteq \mathcal{F}$ denote the σ-field generated by agent n's private information through date t, and let

$$H_t^n(p) = G_t^n \vee \sigma(p_0, p_1, \dots, p_t)$$

be the σ-field that adds observation of equilibrium prices through date t.

The tail σ-field is

$$H_\infty^n(p) = \bigvee_{t \geq 0} H_t^n(p).$$

The first result, due to Bray and Kreps [1987], states that the conditional probability of any event converges almost surely.

Proposition 137.7.1

Fix an agent n and an event $A \in \mathcal{F}$.

The process $M_t = \mathbb{E}^n[\mathbf{1}_A \mid H_t^n(p)]$ is a P^n-bounded martingale with respect to $(H_t^n(p))_{t \geq 0}$, and

$$\lim_{t \to \infty} \mathbb{E}^n[\mathbf{1}_A \mid H_t^n(p)] = \mathbb{E}^n[\mathbf{1}_A \mid H_\infty^n(p)], \qquad P^n\text{-a.s.}$$

The proof is the bounded martingale convergence theorem, with $M_t \in [0, 1]$ supplying the uniform integrability needed for the limit identification.

Proposition 137.7.1 is convergence of posterior assessments, not convergence to "correct" beliefs.

Two qualifications are worth stating.

The "a.s." statement is relative to agent n's own prior P^n, so if two priors P^n and $P^{n'}$ are mutually singular, the conclusion need not hold simultaneously on a common P-positive event.

If the priors share a common null collection, simultaneous convergence holds outside a common null set.

The second result in Bray and Kreps [1987] sharpens convergence from individual events to the entire posterior measure on the parameter space, given a topological assumption on Θ.

Assumption 137.7.1

The parameter space Θ is a complete separable metric (Polish) space, and the Borel σ-field on Θ generated by its open sets makes $(\Theta, \mathcal{B}(\Theta))$ a Borel space.

In the two-agent example $\Theta = [a, b]$ trivially satisfies this assumption.

Under *Assumption 137.7.1* one can fix regular versions of the conditional probabilities: maps

$$P_t^n : \Omega \to \mathcal{P}(\Theta), \qquad \omega \mapsto P_t^n(\omega),$$

such that for each measurable $A \subseteq \Theta$, $\omega \mapsto P_t^n(\omega)(A)$ is a version of $\mathbb{E}^n[\mathbf{1}_{A \times \Phi^\infty} \mid H_t^n(p)](\omega)$, and $P_t^n(\omega) \in \mathcal{P}(\Theta)$ is a probability measure P^n-a.s.

The sharpened convergence result says these regular versions converge weakly almost surely.

Proposition 137.7.2

Under *Assumption 137.7.1*, the regular versions P_t^n converge weakly to a regular version P_∞^n, P^n-a.s.

Equivalently, for P^n-a.e. ω and every bounded continuous $f : \Theta \to \mathbb{R}$,

$$\int_\Theta f \, dP_t^n(\omega) \xrightarrow[t \to \infty]{} \int_\Theta f \, dP_\infty^n(\omega).$$

The proof in Bray and Kreps [1987] applies *Proposition 137.7.1* to a countable disjoint partition of Θ by $1/k$-balls, which exists because Θ is Polish, and then invokes the Portmanteau characterisation of weak convergence on bounded continuous functions.

Rational Bayesian learning therefore always produces a limiting posterior measure.

But *Proposition 137.7.2* alone does not pin down what that limit is, and additional structure is needed before the limit assesses the truth correctly.

137.8 Sharpening the convergence result

We now return to the two-agent example and identify hypotheses under which P_∞^U is a point mass at the true θ^I.

Write F_t for the CDF of agent U's posterior on θ^I at date t after observing $(r_{t-1}, p_{t-1}, x_{t-1}^U)$ and all earlier data.

Proposition 137.7.2 yields a random CDF F_∞ such that F_t converges weakly to F_∞, P^U-a.s.

Three ingredients sharpen this to concentration on the truth, corresponding to the three steps in Bray and Kreps [1987].

Assumption 137.8.1

The equilibrium uninformed demand $x^U(p, F)$ is continuous in F with respect to weak convergence, uniformly in p on a P^U-full-measure set of prices.

Assumption 137.8.2

For fixed θ^U and limiting posterior F_∞, the map

$$\theta^I \mapsto \mathcal{L}\{p_\infty(s; F_\infty, \theta^I, \theta^U)\}$$

is injective, where s is drawn from its marginal distribution.

A sufficient condition is strict stochastic monotonicity of this marginal price distribution in θ^I.

In the lecture's CARA-Normal setup, these assumptions are plausible under the maintained uniqueness and regularity of the implicit demand and price equations.

Continuity requires continuous dependence of the FOC root in (137.6) on F, not just pointwise continuity of the integrals.

Identification also uses the full equilibrium map: the equation $p_t = s_t - \sigma^2 X_t^I/\theta^I$ is informative, but X_t^I itself depends on θ^I, F_t, and the equilibrium fixed point.

The IID assumption on $\{s_t\}$, already part of the model, supplies the ergodicity used in step 2 below.

Under these three assumptions and the IID signal sequence, the limiting posterior in the two-agent example concentrates on the truth.

Proposition 137.8.1

Suppose $\theta^I_{\text{true}} \in [a, b]$ and the prior f_0 puts positive density in every neighbourhood of θ^I_{true}.

Under *Assumption 137.7.1*, *Assumption 137.8.1*, and *Assumption 137.8.2*, and given the IID signal sequence $\{s_t\}$, the limiting posterior on θ^I satisfies

$$F_\infty = \delta_{\theta^I_{\text{true}}} \qquad P^U\text{-a.s.}$$

The proof has three steps.

Step 1: price functional convergence.

Assumption 137.8.1 and the weak convergence $F_t \Rightarrow F_\infty$ from *Proposition 137.7.2* imply that equilibrium demands $x^U(p, F_t)$ converge to $x^U(p, F_\infty)$.

Combining with market clearing and the price equation (137.8) gives $p_t - p_\infty(s_t; F_\infty, \theta^I, \theta^U) \to 0$ on a P^U-full-measure set.

Step 2: the limit price distribution is observable.

Since the deviation $p_t - p_\infty(s_t; F_\infty, \theta^I, \theta^U) \to 0$ almost surely and $\{s_t\}$ is IID, the empirical distribution of observed prices has the same limit as the empirical distribution of the limiting price functional.

The latter equals the distribution of $p_\infty(s; F_\infty, \theta^I, \theta^U)$ for $s \sim N(\mu_s, \tau^2)$, and that limit is $H^U_\infty(p)$-measurable as a long-run frequency of an observable sequence.

Step 3: identification.

Assumption 137.8.2 makes the marginal distribution of p_∞ identify θ^I given (F_∞, θ^U).

Combined with step 2, this means θ^I is itself $H^U_\infty(p)$-measurable, so for any subinterval $[c, d] \subseteq [a, b]$ the limiting posterior satisfies $P^U_\infty(\theta^I \in [c, d]) = \mathbf{1}_{\{\theta^I_{\text{true}} \in [c,d]\}}$.

Combining steps 1, 2, and 3 yields $F_\infty = \delta_{\theta^I_{\text{true}}}$.

The numerical simulation above is consistent with this result.

The posterior density on θ^I collapses to a spike at $\theta^I_{\text{true}} = 2$, and the equilibrium informed trade X^I_t converges to the full-information value $2\theta^I_{\text{true}}/(\theta^I_{\text{true}} + \theta^U)$.

Hence, the path connecting *Proposition 137.7.1* (martingale convergence) to *Proposition 137.8.1* (concentration on the truth) depends on three model-specific ingredients: continuity, ergodicity, and identification.

137.9 Obstacles to convergence

It is natural to ask when these ingredients can fail, and what the consequences are for learning.

137.9.1 Obstacle 1: failure of continuity

If *Assumption 137.8.1* fails, step 1 of the proof breaks.

When the equilibrium price functional is discontinuous in F, small changes in beliefs can produce large changes in the information content of prices, and weak convergence of beliefs need not imply convergence of prices.

Bray and Kreps [1987] flag this as the most delicate step in their argument.

Continuity of $x^U(p, F)$ in F is plausible in this lecture's regular CARA-Normal case, but it also requires uniqueness of the FOC solution and continuous dependence of that solution on F.

Verifying the same property in richer market structures often requires non-trivial regularity arguments.

137.9.2 Obstacle 2: failure of identification

If *Assumption 137.8.2* fails, step 3 breaks even when steps 1 and 2 succeed.

Consider a variant with two informed agents and risk tolerances θ^{I1}, θ^{I2} both unknown to the uninformed agent.

With three agents each endowed with one unit of the risky asset, the full-communication formula (137.2) gives

$$p_t = s_t - \frac{3\sigma^2}{\theta^{I1} + \theta^{I2} + \theta^U},$$

which depends on $(\theta^{I1}, \theta^{I2})$ only through the sum $\theta^{I1} + \theta^{I2}$.

Proposition 137.7.2 still applies, but F_∞ is supported on the level set

$$\{(\theta_1, \theta_2) \in [a, b]^2 : \theta_1 + \theta_2 = \theta^{I1}_{\text{true}} + \theta^{I2}_{\text{true}}\},$$

not on the singleton $\{(\theta^{I1}_{\text{true}}, \theta^{I2}_{\text{true}})\}$.

Convergence occurs, but to a manifold of observationally equivalent parameter values rather than to the truth.

137.9.3 Obstacle 3: misspecification

A separate obstacle arises if the true pricing relation lies outside the agent's prior support.

Blume and Easley [1982] give a stylised version of this obstacle, and *Heterogeneous Beliefs and Financial Markets* develops the Blume-Easley heterogeneous-beliefs model in this lecture series.

Each agent entertains competing conditional likelihoods for other agents' information given his own information and the price.

An equilibrium can exist in which agents assign asymptotic probability one to an incorrect model that gives the observed events positive likelihood, while the true stable price relation receives zero posterior probability because that relation was absent from the prior model class.

In strict rational learning the prior already covers the truth: it is supported on Bayesian-consistent models in the expanded state space, and positive weight is placed on the true parameter.

Bayes' rule can only reweight that initial model class, so a pricing relation with zero prior probability stays at zero no matter what data accumulate.

137.10 Learning within versus learning about a rational expectations equilibrium

The framework above points to an important conceptual distinction in Bray and Kreps [1987].

137.10.1 The distinction

Learning *within* a rational expectations equilibrium is the topic of this lecture.

It is Bayesian inference inside a correctly specified model: *Assumption 137.7.1*, *Assumption 137.8.1*, and *Assumption 137.8.2* all hold, and the prior puts positive weight on the truth.

Agent U is uncertain about θ^I, but for every candidate value he already knows the equilibrium price and allocation maps.

The expanded-state-space formulation $\Omega = \Theta \times \Phi^\infty$ embeds a rational expectations equilibrium on the larger space, and inference reduces to conditional probability over Θ.

Learning *about* a rational expectations equilibrium is a fundamentally different exercise.

The agent does not begin with the equilibrium map embedded in his probability model.

Instead he must infer the price-state relation from data generated while his own beliefs and behavior co-evolve with the data.

137.10.2 The trade-off

The two notions sit on opposite sides of a precise trade-off.

A correctly-specified Bayesian learner enjoys the convergence guarantees in *Proposition 137.7.1* and *Proposition 137.7.2*, but only because the equilibrium has been built into the prior from date zero.

An adaptive learner who treats the price-state relation as something to be estimated can hope to discover it from data, but the estimator he uses cannot be derived from Bayes' rule applied to a correctly specified model.

Bayesian rational learning can update among equilibrium maps already included in the agent's prior, but it does not explain how agents come to obtain those maps in the first place.

The literature on learning *about* rational expectations equilibria, beginning with Bray [1982] and Bray and Savin [1986] and extended by Marcet and Sargent [1989], takes the second side of the trade-off and replaces Bayes' rule with **ordinary least squares** or related recursive estimators.

The companion lecture *Least Squares Learning in Self-Referential Models* develops this least-squares-learning framework in self-referential models and traces the resulting dynamics through the associated ordinary differential equation.

Those rules are computationally tractable and converge in important examples, but they are not the Bayesian update implied by the fully specified rational-learning equilibrium prior.

137.11 Summary

This lecture implemented the rational-learning equilibrium of Bray and Kreps [1987].

Posterior assessments converge by bounded martingale convergence (*Proposition 137.7.1*), and posterior measures converge weakly under a Polish-Borel assumption (*Proposition 137.7.2*).

Concentration on the truth additionally requires continuity (*Assumption 137.8.1*), ergodicity, and identification (*Assumption 137.8.2*); each obstacle above is a failure of one of these.

The simulation illustrates both conclusions: the posterior on θ^I collapses toward θ^I_{true} and the equilibrium informed trade approaches its full-information value.

Rational learning describes the limits of Bayesian inference *given* the equilibrium structure; adaptive learning, in *Least Squares Learning in Self-Referential Models*, describes how that structure can be learned in the first place.

137.12 Exercises

Exercise 137.12.1

Off-center prior

The baseline simulation uses a uniform prior on $\theta^I \in [0.5, 4]$.

1. Re-run the simulation with a prior whose mass sits *above* the true value, for example

```
prior = lambda θ: (θ - 0.5)**3 * (4 - θ)
```

which peaks near $\theta = 3.1$.

2. Plot the posterior mean over time alongside the uniform-prior baseline.
3. Does the posterior eventually concentrate on θ^I_{true}, and how does the speed compare?

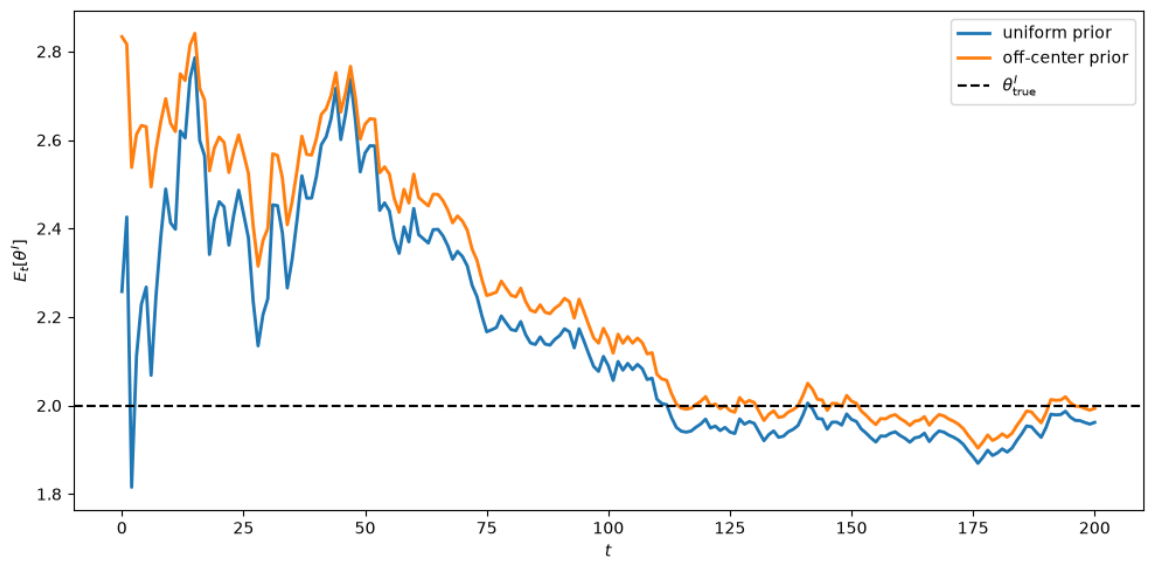
Solution

Here is one solution:

```
res_uniform = simulate(**params)

params_biased = dict(params)
params_biased['prior'] = lambda θ: (θ - 0.5)**3 * (4 - θ)
res_biased = simulate(**params_biased)

fig, ax = plt.subplots(figsize=(10, 5))
ax.plot(res_uniform['post_mean'], lw=2, label='uniform prior')
ax.plot(res_biased['post_mean'], lw=2, label='off-center prior')
ax.axhline(params['θ_I_true'], color='black', ls='--',
           label=r'$\theta^I_{\rm true}$')
ax.set_xlabel('$t$')
ax.set_ylabel(r'$E_t[\theta^I]$')
ax.legend()
plt.tight_layout()
plt.show()
```



The off-center prior starts the posterior mean well above $\theta^I_{\text{true}} = 2$, but Bayesian updating drives it down to the truth.

This is the rational-learning convergence result in action: any prior that puts positive density on θ^I_{true} eventually concentrates around it.

Exercise 137.12.2

Speed of learning across θ^I

Information from one period about θ^I comes through the implied signal

$$s_t(\theta) = \frac{\sigma^2 X_t^I}{\theta} + p_t.$$

The sensitivity $|\partial s_t / \partial \theta| = \sigma^2 X_t^I / \theta^2$ depends on the level of θ^I_{true} through X_t^I and θ^{-2}.

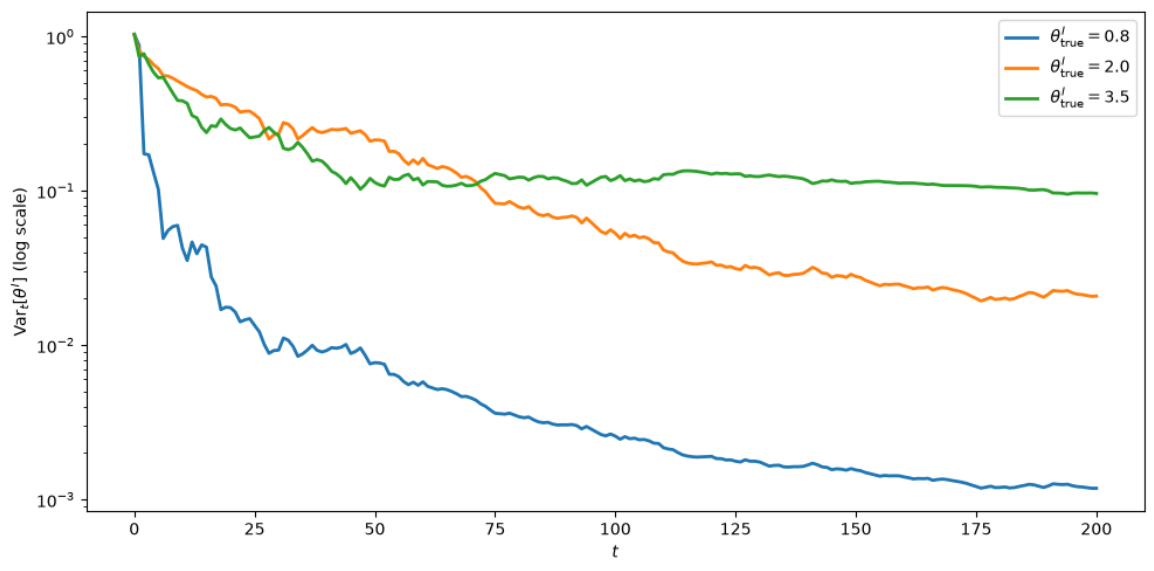
1. Run the simulation for $\theta^I_{\text{true}} \in \{0.8, 2.0, 3.5\}$, holding everything else at the baseline.
2. Plot the posterior variance on a log scale for each case.
3. Which value of θ^I_{true} yields the fastest concentration, and does the result match the sensitivity formula above?

Solution

Here is one solution:

```
fig, ax = plt.subplots(figsize=(10, 5))
for θ_val in [0.8, 2.0, 3.5]:
    params_θ = dict(params)
    params_θ['θ_I_true'] = θ_val
    res_θ = simulate(**params_θ)
    ax.semilogy(res_θ['post_var'], lw=2,
                label=fr'$\theta^I_{{\rm true}} = {θ_val}$')
ax.set_xlabel('$t$')
ax.set_ylabel(r'${\rm Var}_t[\theta^I]$ (log scale)')
ax.legend()
```

```
plt.tight_layout()
plt.show()
```



The smallest θ^I_{true} produces the steepest decline in posterior variance.

The reason is that the sensitivity $\sigma^2 X_t^I / \theta^2$ scales as θ^{-2} for fixed X_t^I, so the same noise level conveys much more information about θ^I when θ^I is small.

Exercise 137.12.3

Effect of return noise

Larger σ^2 widens the conditional density of s_t given r_t, which one might guess slows learning.

But σ^2 also scales the price intercept in (137.8), so price dispersion across candidate θ grows with σ^2.

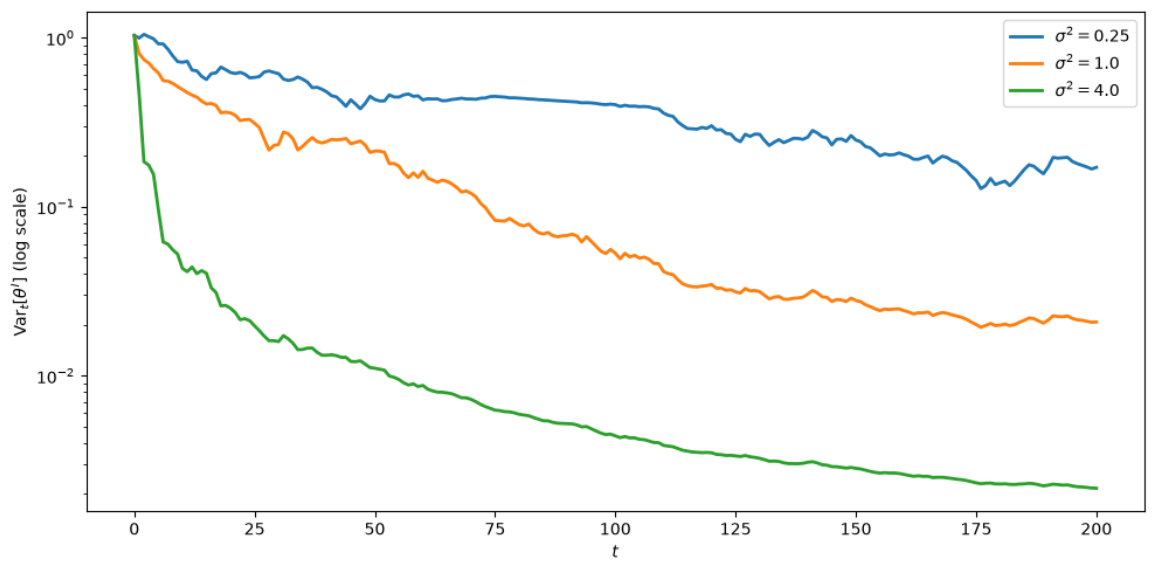
1. Run the simulation with $\sigma^2 \in \{0.25, 1.0, 4.0\}$, keeping $\tau^2 = 1$ fixed.
2. Plot the posterior variance on a log scale for each σ^2.
3. Explain which effect dominates in terms of the signal-to-noise ratio for inferring θ^I from the price.

Solution

Here is one solution:

```
fig, ax = plt.subplots(figsize=(10, 5))
for σ2_val in [0.25, 1.0, 4.0]:
    params_σ = dict(params)
    params_σ['σ2'] = σ2_val
    res_σ = simulate(**params_σ)
    ax.semilogy(res_σ['post_var'], lw=2,
                label=fr'$\sigma^2 = {σ2_val}$')
ax.set_xlabel('$t$')
ax.set_ylabel(r'${\rm Var}_t[\theta^I]$ (log scale)')
ax.legend()
plt.tight_layout()
```

```
plt.show()
```



The posterior variance falls *faster* for larger σ^2.

To see why, write $\nu = \sigma^2$ and hold the realized trade X_t^I fixed.

Two nearby values of θ imply signals separated by approximately

$$\left| \frac{\partial s_t(\theta)}{\partial \theta} \right| = \frac{\nu |X_t^I|}{\theta^2}.$$

The likelihood compares these implied signals using $g(s \mid r_t)$, whose conditional variance is

$$\mathrm{Var}(s_t \mid r_t) = \frac{\nu \tau^2}{\nu + \tau^2}.$$

Thus the local signal-to-noise ratio for distinguishing nearby θ values is proportional to

$$\frac{\nu |X_t^I| / \theta^2}{\sqrt{\nu \tau^2 / (\nu + \tau^2)}} = \frac{|X_t^I|}{\theta^2} \sqrt{\frac{\nu(\nu + \tau^2)}{\tau^2}},$$

This ratio rises with ν when τ^2 is fixed.

The price-revelation effect therefore dominates the extra return noise in this experiment.

CHAPTER

EIGHT

LEAST SQUARES LEARNING IN SELF-REFERENTIAL MODELS

Contents

- *Least Squares Learning in Self-Referential Models*
  - *Overview*
  - *The self-referential structure*
  - *Why a differential equation governs the limit*
  - *The governing ODE*
  - *Five examples*
  - *Example 1: Ordinary linear stochastic difference equations*
  - *Example 2: Bray's cobweb model*
  - *Counterexample: Unstable Bray feedback*
  - *Example 3: Bray-Savin supply-shifter model*
  - *Example 4: Hyperinflation / asset prices*
  - *Example 5: Investment under uncertainty*
  - *Connection to rational learning*
  - *Summary*
  - *Exercises*

138.1 Overview

This lecture is a companion to *Rational Learning and Rational Expectations*, which presents the Bray-Kreps perspective on rational learning.

We examine the related but distinct question of whether *least squares* learning converges to a rational expectations equilibrium in self-referential models.

We present the framework of Marcet and Sargent [1989] for studying **least squares learning** in a class of **self-referential** linear stochastic models.

A self-referential model is one where the *actual* law of motion for the economy depends on the *perceived* law of motion of agents within it.

In a rational expectations equilibrium (REE) the two coincide: the perceived and actual laws of motion are the same.

But if agents start away from equilibrium and update their beliefs by running least squares regressions, will they converge to the REE?

Marcet and Sargent [1989] answer this question by using a powerful technique from systems-control engineering: the differential equation approach of Ljung [1977].

The key insight is that the stochastic difference equation of belief evolution can be approximated by a deterministic ordinary differential equation (ODE) in the limit.

Almost-sure convergence can be studied through the stability of that ODE.

Local stability of the small ODE is closely related to local convergence, while global almost-sure convergence also requires boundedness and domain-of-attraction conditions.

The framework unifies and extends earlier work by Bray [1982] and Bray and Savin [1986] and connects naturally to the distinction between learning *within* a rational expectations equilibrium (Bayesian updating inside a correctly specified model) and learning *about* one (adapting an OLS estimator whose data-generating process shifts with beliefs) discussed in Bray and Kreps [1987].

Let's begin with the imports we'll use throughout.

```
import numpy as np
import matplotlib.pyplot as plt
from numpy import linalg as la
from scipy.integrate import solve_ivp
from scipy.optimize import fsolve
```

Before simulating anything, we describe the object being learned.

138.2 The self-referential structure

138.2.1 Perceived and actual laws of motion

At each date t, agents hold a **perceived law of motion** summarized by a parameter matrix β_t.

They believe that the variable z_{1t} they care about evolves according to

$$z_{1t} = \beta_t z_{2,t-1} + \eta_t, \tag{138.1}$$

where z_{2t} is a vector of variables agents use to forecast $z_{1,t+1}$, and η_t is orthogonal to all past z_2's.

Because agents optimize (or behave) on the basis of this belief, their actions feed back into the economy.

To write the actual law of motion for the full state vector $z_t = (z_{1t}, z_{1t}^c)^\top$, partition $A(\beta_t)$ conformably with $(z_{2,t-1}^c, z_{2,t-1})$ as $A(\beta_t) = [A_c(\beta_t)\ A_2(\beta_t)]$.

$$z_t = \begin{bmatrix} 0 & T(\beta_t) \\ A_c(\beta_t) & A_2(\beta_t) \end{bmatrix} \begin{bmatrix} z_{2,t-1}^c \\ z_{2,t-1} \end{bmatrix} + \begin{bmatrix} V(\beta_t) \\ B(\beta_t) \end{bmatrix} u_t, \tag{138.2}$$

where u_t is IID white noise with covariance Σ.

The mapping T is the key object, as it maps the *perceived* coefficient β to the coefficient that *actually* governs z_{1t} in equilibrium.

A **rational expectations equilibrium** is a fixed point $\beta_f = T(\beta_f)$.

138.2.2 The learning scheme

Agents update β_t each period using **recursive least squares** (RLS).

Define R_t as a running estimate of the second-moment matrix $Ez_{2t}z_{2t}^\top$.

Updating equations are

$$\beta_t^\top = \beta_{t-1}^\top + \frac{\alpha_t}{t} R_{t-1}^{-1} \left\{ z_{2,t-2} z_{2,t-2}^\top [T(\beta_{t-1}) - \beta_{t-1}]^\top + z_{2,t-2} u_{t-1}^\top V(\beta_{t-1})^\top \right\}, \tag{138.3}$$

$$R_t = R_{t-1} + \frac{\alpha_t}{t} [z_{2,t-1} z_{2,t-1}^\top - R_{t-1}/\alpha_t], \tag{138.4}$$

where $\{\alpha_t\}$ is a positive, non-decreasing sequence with $\alpha_t \to 1$ as $t \to \infty$.

When $\alpha_t = 1$ for all t, equations (138.3)-(138.4) reduce to **ordinary least squares** updated recursively.

138.2.3 Lagged and contemporaneous data

The recursion above is written with *lagged information*, so the estimate β_t uses observations available through date $t-1$.

Marcet and Sargent [1989] also treat a *contemporaneous-data* version in which agents update using z_{1t} and $z_{2,t-1}$ at date t.

That timing creates simultaneous determination, because z_t depends on the same estimate β_t that is being updated from z_t.

An extra requirement is that the date-t system must have a unique solution (β_t, R_t, z_t) for each history.

If this uniqueness condition holds, and if the regularity and boundedness conditions used below also hold, convergence is still governed by the same full ODE (138.6) and small ODE (138.5).

Note

As Bray and Savin [1986] and Bray and Kreps [1987] emphasize, the RLS algorithm cannot be derived from Bayes' rule applied to a correctly specified model, because during the learning transition the data-generating process is non-stationary — beliefs shift the equilibrium, which shifts the data.

The algorithm is *"irrational"* in the sense that it acts as if the environment were stationary, when it is not.

Thus, any state variable that matters for the actual law of motion must appear among the regressors in the perceived law.

138.3 Why a differential equation governs the limit

The RLS recursion (138.3)-(138.4) is a *stochastic difference equation* with two key features.

First, the *step size* in front of each update is α_t/t, which shrinks to zero as t grows.

Second, the bracketed expression on the right-hand side of (138.3) is, at the long-run average values of the regressors and noise, just the discrepancy $T(\beta_{t-1}) - \beta_{t-1}$ that measures how far perceived beliefs are from the law of motion they actually generate.

Combine these two facts and a clear picture emerges.

The recursion makes only small adjustments to β_t each period and those adjustments average out, by the law of large numbers, to the deterministic drift $T(\beta) - \beta$.

As the gain α_t/t vanishes, the time-rescaled trajectory β_t looks more and more like a solution to the continuous-time differential equation

$$\frac{d\beta}{dt} = T(\beta) - \beta. \tag{138.5}$$

To compare the ODE with RLS after t observations, evaluate the ODE at $\tau_t = \sum_{s=1}^{t} \alpha_s/s$, the total step size accumulated by the learning algorithm.

When $\alpha_t = 1$, this accumulated step size is $\tau_t = \sum_{s=1}^{t} 1/s \approx \log t$.

This idea, due to Ljung [1977], is what lets Marcet and Sargent [1989] reduce the analysis of a noisy adaptive learning rule to the much easier study of a deterministic ODE.

The rest of this section makes the connection precise.

138.4 The governing ODE

138.4.1 Small and full ODEs

The *small ODE* (138.5) keeps only the drift in β, holding the variance estimator R at its long-run value.

Its rest points are precisely the rational expectations equilibria of the model: $\dot\beta = 0$ if and only if $T(\beta) = \beta$.

The *full ODE* tracks both β and R jointly:

$$\frac{d}{dt}\begin{bmatrix}\beta^\top \\ R\end{bmatrix} = \begin{bmatrix} R^{-1} M_{z_2}(\beta)\,[T(\beta) - \beta]^\top \\ M_{z_2}(\beta) - R \end{bmatrix}, \tag{138.6}$$

where $M_{z_2}(\beta) = \mathbb{E}[z_{2t} z_{2t}^\top]$ is computed at the stationary distribution of z_{2t} that prevails when agents believe the perceived law has constant parameter β.

The fixed point of (138.6) is (β_f, R_f) with $R_f = M_{z_2}(\beta_f)$ — the same β_f as the small ODE, paired with the second-moment matrix consistent with it.

138.4.2 Regularity assumptions

Marcet and Sargent [1989] distinguish two groups of assumptions, with quite different roles.

The first five are standard regularity conditions on the operator T, the shocks u_t, and the gain sequence; they are usually easy to check in applications.

Let $D_s \subset \mathbb{R}^{n_1 \times n_2}$ be the set on which $T(\beta)$, $A(\beta)$, $B(\beta)$, $V(\beta)$ are well defined and the eigenvalues of

$$\begin{bmatrix} 0 & T(\beta) \\ A_c(\beta) & A_2(\beta) \end{bmatrix}$$

are less than one in modulus, so that the state process induced by belief β is covariance-stationary.

Assumption 138.4.1 (Unique REE)

The operator T has a unique fixed point $\beta_f = T(\beta_f)$ with $\beta_f \in D_s$.

Assumption 138.4.2 (Smoothness of T, A, B, V)

T is twice differentiable and A, B, V each have one derivative in D_s.

Assumption 138.4.3 (Nonsingular limit covariance)

The covariance matrix $M_{z_2}(\beta_f)$ is nonsingular.

Assumption 138.4.4 (Gain sequence)

For all t, $\alpha_t > 0$; α_t is non-decreasing in t; $\alpha_t \to 1$ as $t \to \infty$; and $\limsup_{t\to\infty} t\,|\alpha_t - \alpha_{t-1}| < \infty$.

Assumption 138.4.5 (Shock moments)

The vector u_t is serially independent, and $E|u_{it}|^p < \infty$ for all $p > 1$ and all $i = 1, \ldots, m$.

These five say: there is a unique target, the model is smooth around it, the regression has a well-defined precision matrix in the limit, the gain shrinks at the right speed, and the shocks have enough moments for laws of large numbers to work.

The remaining two are *boundedness* conditions.

They are needed because Ljung's theorem requires the relevant sample path to return to a bounded region infinitely often.

In this model that includes the regressor process z_{2t} and the covariance estimate R_t.

Assumption 138.4.6 (Boundedness along a subsequence)

There exist a set Ω_0 with $\mathbb{P}(\Omega_0) = 1$, random variables $C_1(\omega)$ and $C_2(\omega)$, and a subsequence $\{t_k(\omega)\}$ such that

$$|z_{2t_k}(\omega)| < C_1(\omega) \quad \text{and} \quad |R_{t_k}(\omega)| < C_2(\omega)$$

for all $\omega \in \Omega_0$ and all $k = 1, 2, \ldots$.

Assumption 138.4.7 (Sample path stays in a workable domain)

Either the compact-state version or the projection-set version holds.

- *Compact-state version.*
  - $D_1 = D_2 = \mathbb{R}^{n_1 \times (n_2)^3}$.
  - There exists a compact $D' \subset D_s$ that contains $\beta_{t_k}(\omega)$ for all k and all $\omega \in \Omega_0$.
  - For every initial condition $(\beta(0), R(0))$ with $\beta(0) \in D'$ and $R(0)$ bounded as in *Assumption 138.4.6*, the β-component of the trajectory of (138.6) remains in a closed subset of D_s.
- *Projection-set version.*
  - D_2 is closed.

- D_1 is open and bounded.
- $\beta \in D_s$ for every $(\beta, R) \in D_1$.
- Trajectories of (138.6) with initial conditions in D_2 never leave a closed subset of D_1.

Assumption 138.4.6 is automatic when the regressors z_{2t} are *exogenous* and ergodic, but it can be delicate when z_{2t} contains endogenous variables.

Assumption 138.4.7 is satisfied in the compact-state version when the model has a natural bounded domain on which T is well-defined; otherwise the projection-set version corresponds to using the projection facility described below.

Let D_A denote the domain of attraction of the unique equilibrium (β_f, R_f) of (138.6).

138.4.3 Convergence of least squares

Proposition 138.4.1 (Least-squares convergence)

Assume *Assumption 138.4.1*, *Assumption 138.4.2*, *Assumption 138.4.3*, *Assumption 138.4.4*, *Assumption 138.4.5*, and *Assumption 138.4.6*.

If either

- the compact-state version of *Assumption 138.4.7* is satisfied and $D' \subset D_A$, or
- the projection-set version of *Assumption 138.4.7* is satisfied and $D_1 \subset D_A$,

then $\beta_t \to \beta_f$ almost surely as $t \to \infty$.

Proposition 138.4.1 reduces the problem of almost-sure convergence of recursive least squares to checking the *stability* of the ODE (138.6) at (β_f, R_f) and guaranteeing the boundedness of the sample path.

138.4.4 Stability governs convergence

Let $\mathcal{M}$ be the Jacobian matrix of $T(\beta) - \beta$ evaluated at the REE β_f:

$$\mathcal{M} = \frac{d\ \mathrm{col}(T(\beta) - \beta)}{d\ \mathrm{col}(\beta)^\top}\bigg|_{\beta=\beta_f}, \tag{138.7}$$

and let $h(\beta, R)$ denote the Jacobian of the right-hand side of the full ODE (138.6) after stacking $(\beta^\top, R)$ into a column vector.

Proposition 138.4.2 (Jacobian reduction)

The matrix $h(\beta_f, R_f)$ has $(n_2)^2$ repeated eigenvalues equal to -1; its remaining $n_1 \times n_2$ eigenvalues coincide with the eigenvalues of $\mathcal{M}$.

Consequently:

- If all eigenvalues of $\mathcal{M}$ have *strictly negative real parts*, both (138.5) and (138.6) are locally stable.
- Under the boundedness and domain-of-attraction conditions in *Proposition 138.4.1*, this local stability can be used to obtain almost-sure convergence of β_t to β_f.

- If any eigenvalue of $\mathcal{M}$ has *positive real part*, then the next proposition shows that convergence is impossible.

Proposition 138.4.3 (Necessity)

Assume *Assumption 138.4.1*, *Assumption 138.4.2*, *Assumption 138.4.3*, *Assumption 138.4.4*, and *Assumption 138.4.5*.

1. Let $\hat{\beta} \neq \beta_f$ and suppose $M_{z_2}(\hat{\beta})$ is positive definite and $\hat{\beta} \in \text{int}(D_2)$.

   Then $\mathbb{P}(\beta_t \to \hat{\beta}) = 0$.

2. If $h(\beta_f, R_f)$ has at least one eigenvalue with strictly positive real part, then $\mathbb{P}(\beta_t \to \beta_f) = 0$.

The first part says that recursive least squares cannot settle on a non-equilibrium belief in the interior of the admissible domain.

If learning has a limit, that limit must be a fixed point of T.

The second part says that local instability is not just slow convergence.

If the REE is unstable for the learning ODE, convergence to that REE has probability zero.

The stability condition $\text{Re}(\lambda_i(\mathcal{M})) < 0$ for all i is what is called **E-stability** (See Evans [1985]): the REE is a stable rest point of the expectational dynamics $\dot{\beta} = T(\beta) - \beta$.

E-stability plays the role here that prior support plays in Bayesian learning: it tells us when the learning process can find its way to the equilibrium.

138.4.5 The projection facility

E-stability is the local condition that makes the REE attractive for the learning ODE.

By itself, it is not an almost-sure convergence theorem.

The stochastic recursion must also remain in a bounded region where the ODE approximation is valid.

Conversely, a strictly positive eigenvalue rules out convergence to the REE, while boundary cases with zero real parts are not covered by the theorem.

Ljung's theorem requires the learning process and the relevant regressors to return to bounded regions with probability one.

The **projection facility** is a procedure that enforces boundedness of the estimated coefficients and covariance matrix.

What the projection facility does

The full learning algorithm augments the plain RLS update with a constraint set $D_1 \supset D_2$ in (β, R)-space.

After each unconstrained RLS step produces a candidate $(\tilde{\beta}_t, \tilde{R}_t)$, the projection facility enforces:

$$(\beta_t, R_t) = \begin{cases} (\tilde{\beta}_t,\ \tilde{R}_t) & \text{if } (\tilde{\beta}_t, \tilde{R}_t) \in D_1, \\ \text{some point in } D_2 & \text{otherwise.} \end{cases} \tag{138.8}$$

The set D_1 is chosen so that the model remains well-defined (e.g., R_t stays positive definite; β_t stays in a region where $T(\beta)$ is well-defined and the state process is covariance-stationary).

The set $D_2 \subset D_1$ is a slightly smaller "safe" region to which the algorithm is retracted whenever it threatens to leave D_1.

Without the projection facility, the stochastic path (β_t, R_t) might temporarily wander to regions where the system (138.2) is non-stationary (e.g., an explosive VAR).

The projection facility keeps (β_t, R_t) inside a chosen admissible region.

When regressors include endogenous variables, the separate boundedness condition *Assumption 138.4.6* still has to control z_{2t} along a subsequence.

Formally, Marcet and Sargent [1989] require that trajectories of the full ODE (138.6) originating in D_1 do not leave D_1.

This is often checked by verifying that the full vector field points toward the interior at the boundary ∂D_1.

Corollary 138.4.1

Assume *Assumption 138.4.1*, *Assumption 138.4.2*, *Assumption 138.4.3*, *Assumption 138.4.4*, *Assumption 138.4.5*, and *Assumption 138.4.6*.

Suppose also that $(\beta, R) \in D_1$ implies $\beta \in D_s$, and that D_1 is open and bounded with $D_1 \subset D_A$.

Then for some subsequence $\{t_k(\omega)\}$,

$$\mathbb{P}(\beta_t \to \beta_f) + \mathbb{P}(\beta_{t_k} \to (D_1 \setminus D_2)) = 1.$$

The second event has probability zero whenever the ODE trajectories point inward at ∂D_1, in which case *Corollary 138.4.1* reduces to $\beta_t \to \beta_f$ almost surely.

The exogenous-regressor case

When the regressors z_{2t} are *exogenous*, so that $E(z_{2t} z_{2t}^\top) = M_{z_2}(\beta) \equiv M$ does not depend on β, the verification of the boundary condition becomes routine.

Let $H(\beta)$ be the mean-value slope of T, i.e. the matrix satisfying

$$\operatorname{col}\{T(\beta) - T(\beta_f)\} = H(\beta) \operatorname{col}(\beta - \beta_f). \tag{138.9}$$

Corollary 138.4.2

Consider the algorithm defined by (138.3)-(138.4) with projection rule (138.8).

Choose $0 < K' < K < \infty$ and assume

1. *Assumption 138.4.1*, *Assumption 138.4.2*, *Assumption 138.4.3*, *Assumption 138.4.4*, and *Assumption 138.4.5* hold;
2. z_{2t} is exogenous, so that $E(z_{2t} z_{2t}^\top) = M_{z_2}(\beta) \equiv M$;
3. the small ODE $\dot\beta = T(\beta) - \beta$ is globally stable in $\mathbb{R}^{n_1 \times n_2}$;
4. there exists $\bar\varepsilon > 0$ such that for all $0 < \varepsilon \le \bar\varepsilon$ and all β with $|\beta - \beta_f| = K$, every eigenvalue of $[I(1-\varepsilon) + \varepsilon H(\beta)]^\top [I(1-\varepsilon) + \varepsilon H(\beta)]$ has modulus less than α^2 with $\alpha < 1$.

Take $D_1 = \{(\beta, R) : |\beta - \beta_f| < K\}$ and $D_2 = \{(\beta, R) : |\beta - \beta_f| \le K'\}$, and let the projection rule retract β_t to any value with $|\beta - \beta_f| \le K'$, while leaving $R_t = \tilde R_t$.

Then $\beta_t \to \beta_f$ almost surely.

For the scalar linear examples below, condition (4) reduces to checking that the feedback slope of T is not too strong.

For the first four examples below, T is linear and M_{z_2} is independent of β, so *Corollary 138.4.2* uses D_1 and D_2 only to justify the projection argument and leaves the examples with the scalar ODE stability checks.

D_1 and D_2 are therefore not economic primitives in the examples below.

They are proof devices for the projected learning algorithm: D_1 is the outer admissible region where the model and ODE remain well defined, and D_2 is the smaller region used to reset the estimates if an update tries to leave D_1.

Once the corollary says suitable sets can be chosen, the examples no longer need to display them and can focus on T, β_f, and the local stability slope.

138.5 Five examples

We now work through five examples from Marcet and Sargent [1989].

Each substantive learning example follows the same template.

1. Write down the economic equations that determine the equilibrium.
2. Substitute the perceived law of motion for any expectational variables to obtain the actual law of motion as a function of β.
3. Read off the operator T and the REE $\beta_f = T(\beta_f)$.
4. Check E-stability by computing $\mathcal{M} = dT/d\beta - I$ at β_f.
5. Plot the RLS paths, the small ODE evaluated at cumulative-gain time, and the drift $T(\beta) - \beta$.

In each case, β_t is the law of motion agents currently use when making decisions, while $T(\beta_t)$ is the law of motion their decisions actually produce.

Least-squares learning asks whether repeated observations move β_t toward a fixed point of this map.

The first example is the control case with no self-referential feedback.

For the remaining scalar examples, the simulation figure appears immediately after the model description.

138.6 Example 1: Ordinary linear stochastic difference equations

The first example has no self-referential component.

Let the actual law of motion be fixed, with $T(\beta) = \Gamma$ for a stable matrix Γ and with $V(\beta) = I$.

Economically, this is the control case.

Agents are estimating a stable, exogenous data-generating process.

Their beliefs do not feed back into prices, quantities, or future data.

The REE is $\beta_f = \Gamma$.

Since T is constant, $H(\beta) = 0$ in (138.9).

The small-ODE Jacobian is $\mathcal{M} = -I$, so the small ODE is globally stable.

Corollary 138.4.2 then implies that recursive least squares converges almost surely to the true law of motion.

In this no-feedback benchmark, the Marcet-Sargent result reduces to the standard consistency of least squares for stable linear stochastic difference equations.

138.7 Example 2: Bray's cobweb model

Bray [1982] studied a simple cobweb economy in which the equilibrium price satisfies

$$p_t = a + b\beta_t + \tilde{u}_t, \tag{138.10}$$

where β_t is agents' OLS estimate of the price (their point forecast of p_t), and $\tilde{u}_t$ is IID noise with mean zero and variance σ_u^2.

Here the forecast itself is a state variable for the economy.

If producers expect a high price, their current supply decisions alter the market-clearing price.

The parameter b measures the strength of this expectational feedback.

When $b < 1$, the actual price response is weaker than the forecast error, so least squares has a force pushing beliefs back toward the fixed point.

When $b > 1$, the feedback is too strong and the same learning rule moves beliefs away from the REE.

The mapping T is simply $T(\beta) = a + b\beta$.

The REE is

$$\beta_f = \frac{a}{1-b}, \quad b \neq 1. \tag{138.11}$$

The small ODE is

$$\dot{\beta} = T(\beta) - \beta = a + b\beta - \beta = a - (1-b)\beta, \tag{138.12}$$

which has the unique fixed point $\beta_f = a/(1-b)$.

Its Jacobian is $\mathcal{M} = b - 1 < 0$ when $b < 1$.

Under the boundedness and domain conditions above, RLS then converges almost surely.

When $b > 1$, $\mathcal{M} > 0$ and convergence fails.

The code below sets $a = 1$ and $b = 0.6$.

The rational expectations price forecast is then $\beta_f = 2.5$.

The three panels show the noisy RLS paths, the small-ODE approximation, and the learning drift $T(\beta) - \beta$.

The next cell sets up the simulator and plotting helpers used in the Bray figures.

```
def simulate_rls_scalar(T_map, σ_u, β0, T_periods=500, N_paths=100,
                        a_seq=None, seed=0):
    """
    Simulate the RLS recursion for z1_t = T(β_t) + u_t, z2_t = 1.
    Returns an (N_paths, T_periods) array of belief paths.
    """
    rng = np.random.default_rng(seed)
    if a_seq is None:
        a_seq = np.ones(T_periods)

    β_paths = np.empty((N_paths, T_periods))

    for i in range(N_paths):
        β = β0
        R = 1.0
```

(continues on next page)

(continued from previous page)

```
        for t in range(T_periods):
            α_t = a_seq[t]
            z2 = 1.0
            u_t = rng.normal(0, σ_u)
            z1 = T_map(β) * z2 + u_t

            R_old = max(R, 1e-8)
            β = β + (α_t / (t + 1)) / R_old * z2 * (z1 - β * z2)
            R = R + (α_t / (t + 1)) * (z2**2 - R / α_t)
            R = max(R, 1e-8)

            β_paths[i, t] = β

    return β_paths


def solve_ode_calendar(f_ode, β0, T_periods, a_seq=None):
    """
    Solve dβ/dτ = f_ode(β) and evaluate it at RLS cumulative-gain time.
    """
    if a_seq is None:
        a_seq = np.ones(T_periods)

    periods = np.arange(T_periods + 1)
    gains = a_seq / np.arange(1, T_periods + 1)
    ode_time = np.concatenate(([0.0], np.cumsum(gains)))
    sol = solve_ivp(
        lambda τ, y: [f_ode(y[0])],
        (0.0, ode_time[-1]),
        [β0],
        t_eval=ode_time,
        method='RK45',
        max_step=0.05
    )
    return periods, sol.y[0]


def plot_scalar_drift(ax, β_grid, drift, β_f, color):
    """Plot the one-dimensional learning drift T(β) - β."""
    ax.plot(β_grid, drift, color=color, lw=2)
    ax.axhline(0, color='black', lw=1.5)
    ax.axvline(β_f, color='red', ls='--', lw=2,
               label=f'$\\beta_f = {β_f:.2f}$')
    ax.fill_between(β_grid, drift, 0, where=(drift > 0),
                    color=color, alpha=0.12)
    ax.fill_between(β_grid, drift, 0, where=(drift < 0),
                    color=color, alpha=0.12)

    for β_arrow in np.linspace(β_grid[20], β_grid[-20], 7):
        dβ = np.interp(β_arrow, β_grid, drift)
        if abs(dβ) > 1e-10:
            ax.annotate(
                '', xy=(β_arrow + 0.25 * np.sign(dβ), 0),
                xytext=(β_arrow, 0),
                arrowprops=dict(arrowstyle='->', color=color, lw=1.8)
            )
```

(continues on next page)

(continued from previous page)

```
    ax.set_xlabel('$\\beta$')
    ax.set_ylabel('$T(\\beta) - \\beta$')
    ax.legend(fontsize=9)


T_sim = 400
N_sim = 80
```

```
a_bray, b_bray, σ_bray = 1.0, 0.6, 1.0


def T_bray(β):
    return a_bray + b_bray * β


β_f_bray = a_bray / (1 - b_bray)

β0_bray = 0.0

β_paths_bray = simulate_rls_scalar(
    T_bray, σ_bray, β0_bray,
    T_periods=T_sim, N_paths=N_sim
)


def ode_bray(β):
    return a_bray + b_bray * β - β


t_ode, sol_low = solve_ode_calendar(ode_bray, 0.0, T_sim)
_, sol_high = solve_ode_calendar(ode_bray, 4.5, T_sim)

β_grid_bray = np.linspace(-0.5, 5.0, 300)
drift_bray = np.array([ode_bray(b) for b in β_grid_bray])

fig, axes = plt.subplots(1, 3, figsize=(15, 4.8))

ax = axes[0]
for i in range(min(30, N_sim)):
    ax.plot(β_paths_bray[i], color='steelblue', alpha=0.25, lw=2)
ax.plot(np.mean(β_paths_bray, axis=0), color='navy', lw=2,
        label='cross-path average')
ax.axhline(β_f_bray, color='red', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_bray:.2f}$')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta_t$')
ax.legend()

ax = axes[1]
ax.plot(
    t_ode, sol_low, color='steelblue', lw=2,
    label='ODE at $\\tau_t$, $\\beta_0=0$'
)
ax.plot(
    t_ode, sol_high, color='darkorange', lw=2,
    label='ODE at $\\tau_t$, $\\beta_0=4.5$'
```

(continues on next page)

(continued from previous page)

```
)
ax.axhline(β_f_bray, color='red', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_bray:.2f}$')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta(\\tau_t)$')
ax.legend()

plot_scalar_drift(
    axes[2], β_grid_bray, drift_bray, β_f_bray, 'steelblue'
)

plt.tight_layout()
plt.show()
print(f"REE: β_f = a/(1-b) = {β_f_bray:.4f}")
print(f"Jacobian M = b - 1 = {b_bray - 1:.4f}  (< 0: E-stable)")
```

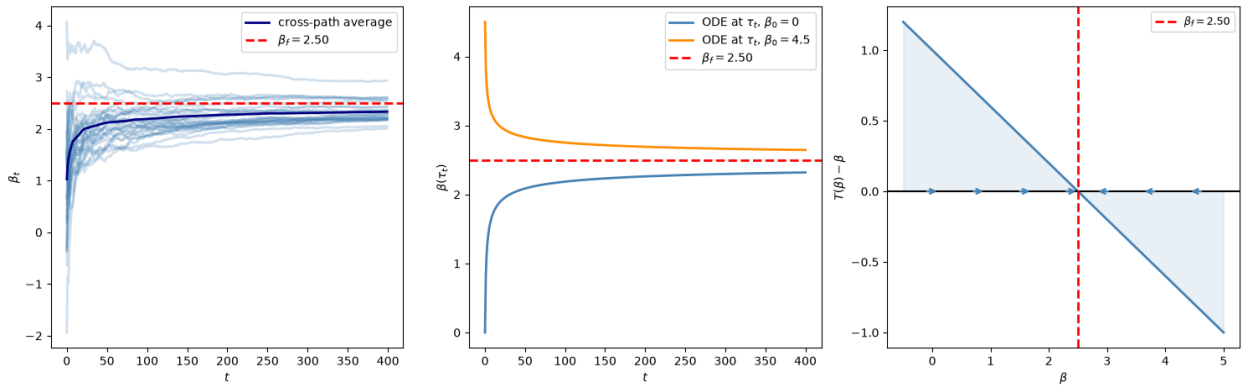


Fig. 138.1: Bray learning dynamics

```
REE: β_f = a/(1-b) = 2.5000
Jacobian M = b - 1 = -0.4000  (< 0: E-stable)
```

The stochastic paths are noisy because prices contain $\tilde{u}_t$.

The average path and the ODE paths remove most of that noise and reveal the mean learning direction.

138.8 Counterexample: Unstable Bray feedback

To see what happens when E-stability is violated, repeat Bray's model with $b > 1$.

It represents an economy where expectations affect actual prices more than one-for-one.

If agents start with a forecast that is too high, the market outcome induced by that forecast is higher still, so least squares raises the forecast.

The same logic works in reverse below the REE.

The REE exists, but it is not learnable by this adaptive rule.

The same three diagnostic panels now show how beliefs are pushed away from the REE rather than toward it, and the drift panel shows that the learning force points away from the fixed point on both sides.

```
b_unstable = 1.4


def T_unstable(β):
    return a_bray + b_unstable * β


β_f_unstable = a_bray / (1 - b_unstable)

β_paths_unstable = simulate_rls_scalar(
    T_unstable, σ_bray, β0=0.0,
    T_periods=200, N_paths=50
)


def ode_unstable(β):
    return T_unstable(β) - β


β_grid = np.linspace(-5, 5, 300)
drift = np.array([ode_unstable(b) for b in β_grid])

t_ode_un, sol_un_low = solve_ode_calendar(ode_unstable, -4.0, 200)
_, sol_un_high = solve_ode_calendar(ode_unstable, 0.0, 200)

fig, axes = plt.subplots(1, 3, figsize=(15, 4.8))

ax = axes[0]
for i in range(min(30, 50)):
    ax.plot(β_paths_unstable[i], color='crimson', alpha=0.3, lw=2)
ax.axhline(β_f_unstable, color='black', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_unstable:.2f}$ (unstable)')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta_t$')
ax.legend()

ax = axes[1]
ax.plot(
    t_ode_un, sol_un_low, color='crimson', lw=2,
    label='ODE at $\\tau_t$, $\\beta_0=-4$'
)
ax.plot(
    t_ode_un, sol_un_high, color='steelblue', lw=2,
    label='ODE at $\\tau_t$, $\\beta_0=0$'
)
ax.axhline(β_f_unstable, color='black', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_unstable:.2f}$')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta(\\tau_t)$')
ax.legend()

plot_scalar_drift(
    axes[2], β_grid, drift, β_f_unstable, 'crimson'
)

plt.tight_layout()
plt.show()
print(f"Jacobian M = b - 1 = {b_unstable - 1:.2f}  (> 0: NOT E-stable)")
```

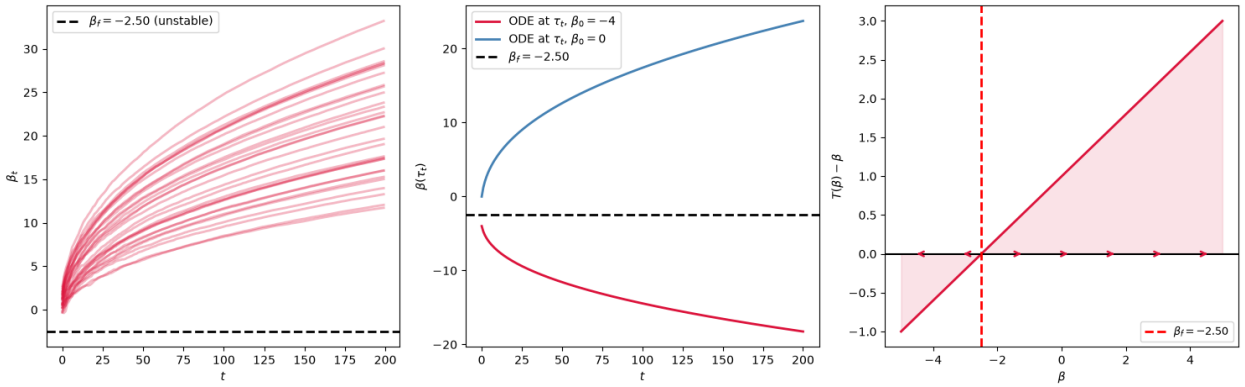


Fig. 138.2: Unstable Bray dynamics

```
Jacobian M = b - 1 = 0.40  (> 0: NOT E-stable)
```

138.9 Example 3: Bray-Savin supply-shifter model

Bray and Savin [1986] studied a model where

$$p_t = x_t^\top (m + a\beta_{t-1}) + \tilde{u}_t, \quad p_t^e = x_t^\top \beta_{t-1}, \tag{138.13}$$

with x_t an exogenous supply-shifter, a a scalar feedback parameter, and agents running an OLS regression of p on x.

This is the same learning problem with one extra economic ingredient: expectations are conditional on an observed shifter.

Agents do not merely forecast the average price.

They estimate how price responds to x_t.

The actual coefficient on x_t is $m + a\beta$, so the perceived slope changes the true slope generated by the market.

If $a < 1$, a mistaken slope feeds back into the actual price equation with attenuation, and the least-squares estimate is pulled toward the REE slope.

The mapping is $T(\beta) = m + a\beta$ (scalar case), giving

$$\dot{\beta} = (a - 1)\beta + m, \quad \beta_f = \frac{m}{1 - a}, \tag{138.14}$$

with Jacobian $\mathcal{M} = a - 1 < 0$ iff $a < 1$.

The simulation below draws the exogenous shifter x_t and lets agents update the slope in the regression of p_t on x_t.

It uses $a = 0.7$.

The drift panel shows that mistaken slopes are pushed back toward the REE slope rather than amplified.

The next cell adapts the RLS simulator to the supply-shifter regression.

```
def simulate_bray_savin(m, a, σ_x, σ_u, β0, T_periods=500,
                        N_paths=100, seed=0):
    """
    Simulate RLS for p_t = x_t (m + a β_t) + u_t.
    Agents regress p_t on x_t.
```

(continues on next page)

(continued from previous page)

```
    """
    rng = np.random.default_rng(seed)
    β_paths = np.empty((N_paths, T_periods))

    for i in range(N_paths):
        β = β0
        R = 1.0

        for t in range(T_periods):
            x_t = rng.normal(0.0, σ_x)
            u_t = rng.normal(0.0, σ_u)
            p_t = x_t * (m + a * β) + u_t
            step = 1.0 / (t + 1)

            R_old = max(R, 1e-8)
            β = β + step / R_old * x_t * (p_t - β * x_t)
            R = R + step * (x_t**2 - R)
            R = max(R, 1e-8)
            β_paths[i, t] = β

    return β_paths
```

```
m_bs, a_bs = 0.5, 0.7
σ_x_bs, σ_u_bs = 1.0, 1.0


def T_bs(β):
    return m_bs + a_bs * β


β_f_bs = m_bs / (1 - a_bs)

β_paths_bs = simulate_bray_savin(
    m_bs, a_bs, σ_x_bs, σ_u_bs, 0.0,
    T_periods=T_sim, N_paths=N_sim
)


def ode_bs(β):
    return T_bs(β) - β


t_ode_bs, sol_bs_low = solve_ode_calendar(ode_bs, 0.0, T_sim)
_, sol_bs_high = solve_ode_calendar(ode_bs, 4.0, T_sim)

β_grid_bs = np.linspace(-0.5, 4.0, 300)
drift_bs = np.array([ode_bs(b) for b in β_grid_bs])

fig, axes = plt.subplots(1, 3, figsize=(15, 4.8))

ax = axes[0]
for i in range(min(30, N_sim)):
    ax.plot(β_paths_bs[i], color='darkorange', alpha=0.25, lw=2)
ax.plot(np.mean(β_paths_bs, axis=0), color='saddlebrown', lw=2,
        label='cross-path average')
ax.axhline(β_f_bs, color='red', ls='--', lw=2,
```

(continues on next page)

(continued from previous page)

```
            label=f'$\\beta_f = {β_f_bs:.2f}$')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta_t$')
ax.legend()

ax = axes[1]
ax.plot(
    t_ode_bs, sol_bs_low, color='darkorange', lw=2,
    label='ODE at $\\tau_t$, $\\beta_0=0$'
)
ax.plot(
    t_ode_bs, sol_bs_high, color='steelblue', lw=2,
    label='ODE at $\\tau_t$, $\\beta_0=4$'
)
ax.axhline(β_f_bs, color='red', ls='--', lw=2,
            label=f'$\\beta_f = {β_f_bs:.2f}$')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta(\\tau_t)$')
ax.legend()

plot_scalar_drift(
    axes[2], β_grid_bs, drift_bs, β_f_bs, 'darkorange'
)

plt.tight_layout()
plt.show()
print(f"REE: β_f = m/(1-a) = {β_f_bs:.4f}")
print(f"Jacobian M = a - 1 = {a_bs - 1:.4f}  (< 0: E-stable)")
```

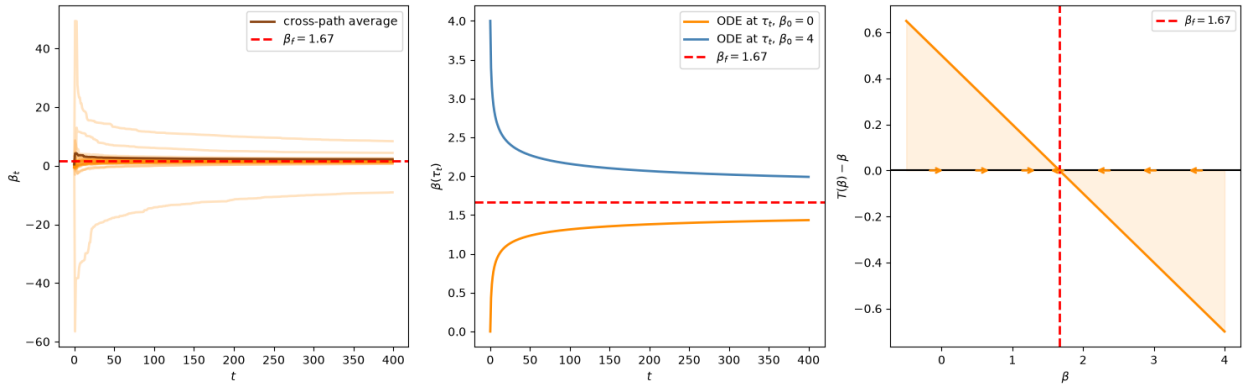


Fig. 138.3: Bray-Savin learning dynamics

```
REE: β_f = m/(1-a) = 1.6667
Jacobian M = a - 1 = -0.3000  (< 0: E-stable)
```

The interpretation is parallel to Bray's cobweb model.

The difference is that agents are now learning how prices co-move with an exogenous variable, not just learning an unconditional mean.

138.10 Example 4: Hyperinflation / asset prices

Consider the present-value asset pricing model

$$y_t = \lambda \mathbb{E}_t y_{t+1} + x_t, \quad x_t = \rho x_{t-1} + \varepsilon_t, \tag{138.15}$$

where $|\lambda| < 1$, $|\rho| < 1$, and agents perceive $y_t = \beta_t x_{t-1} + v_t$.

Under the perceived law, $\mathbb{E}_t y_{t+1} = \beta_t x_t$, so (138.15) becomes

$$y_t = (\lambda\beta_t + 1)x_t = (\lambda\beta_t + 1)\rho x_{t-1} + (\lambda\beta_t + 1)\varepsilon_t.$$

The coefficient on x_{t-1} in the actual law of motion is therefore $(\lambda\beta_t + 1)\rho$, which gives the operator T below.

If y_t is a price level, the equation resembles a Cagan-style hyperinflation model.

If y_t is an asset price and x_t is a dividend or fundamental, it is a present-value relation.

In both cases, beliefs about the future value of y enter the current value of y.

The product $\lambda\rho$ measures the strength of the forward-looking feedback: λ discounts the future and ρ measures how persistent the fundamental is.

When $|\lambda\rho| < 1$, a mistaken perceived coefficient does not amplify without bound.

The mapping is $T(\beta) = (\lambda\beta + 1)\rho$ and the REE is

$$\beta_f = \frac{\rho}{1 - \lambda\rho}. \tag{138.16}$$

The small ODE is

$$\dot{\beta} = (\lambda\rho - 1)\beta + \rho, \tag{138.17}$$

with Jacobian $\mathcal{M} = \lambda\rho - 1 < 0$ for $|\lambda\rho| < 1$, so the small ODE is stable.

The convergence theorem then applies under the exogenous-regressor conditions above.

The simulation uses $\lambda = 0.8$ and $\rho = 0.9$, so $\lambda\rho = 0.72$.

This means the feedback from expectations is strong enough to matter but still less than one.

The ODE panel below uses τ itself on the horizontal axis, where τ is the total step size accumulated by least squares.

The code simulates the autoregressive fundamental x_t and updates an OLS regression of y_t on x_{t-1}.

The next cell simulates the autoregressive fundamental and the matching RLS update.

It keeps the covariance estimate R_t bounded away from zero, mirroring the projection idea above and avoiding misleading early-sample explosions.

```
def simulate_present_value_rls(λ, ρ, σ_ε, β0, T_periods=500,
                               N_paths=100, seed=0):
    """
    Simulate RLS for y_t = (λ β_t + 1) x_t,
    x_t = ρ x_{t-1} + ε_t, with regression of y_t on x_{t-1}.
    """
    rng = np.random.default_rng(seed)
    β_paths = np.empty((N_paths, T_periods))
    σ_x = σ_ε / np.sqrt(1 - ρ**2)
    R_min = 0.05 * σ_x**2
```

(continues on next page)

(continued from previous page)

```
    for i in range(N_paths):
        β = β0
        R = σ_x**2
        x_lag = rng.normal(0.0, σ_x)

        for t in range(T_periods):
            ε_t = rng.normal(0.0, σ_ε)
            x_t = ρ * x_lag + ε_t
            y_t = (λ * β + 1.0) * x_t
            step = 1.0 / (t + 1)

            R_old = max(R, R_min)
            β = β + step / R_old * x_lag * (y_t - β * x_lag)
            R = R + step * (x_lag**2 - R)
            R = max(R, R_min)
            β_paths[i, t] = β
            x_lag = x_t

    return β_paths
```

```
λ, ρ_pv, σ_pv = 0.8, 0.9, 1.0
T_pv_sim = T_sim
N_pv_sim = N_sim


def T_pv(β):
    return (λ * β + 1) * ρ_pv


β_f_pv = ρ_pv / (1 - λ * ρ_pv)

β_paths_pv = simulate_present_value_rls(
    λ, ρ_pv, σ_pv, 0.0,
    T_periods=T_pv_sim, N_paths=N_pv_sim
)


def ode_pv(β):
    return T_pv(β) - β


τ_grid_pv = np.linspace(0.0, 30.0, 300)
sol_pv_low = solve_ivp(
    lambda τ, y: [ode_pv(y[0])],
    (0.0, τ_grid_pv[-1]),
    [0.0],
    t_eval=τ_grid_pv,
    max_step=0.05
).y[0]
sol_pv_high = solve_ivp(
    lambda τ, y: [ode_pv(y[0])],
    (0.0, τ_grid_pv[-1]),
    [10.0],
    t_eval=τ_grid_pv,
    max_step=0.05
).y[0]
```

(continues on next page)

(continued from previous page)

```
β_grid_pv = np.linspace(-1.0, 7.0, 300)
drift_pv = np.array([ode_pv(b) for b in β_grid_pv])

fig, axes = plt.subplots(1, 3, figsize=(15, 4.8))

ax = axes[0]
for i in range(min(30, N_pv_sim)):
    ax.plot(β_paths_pv[i], color='seagreen', alpha=0.25, lw=2)
ax.plot(np.mean(β_paths_pv, axis=0), color='darkgreen', lw=2,
        label='cross-path average')
ax.axhline(β_f_pv, color='red', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_pv:.2f}$')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta_t$')
ax.legend()

ax = axes[1]
ax.plot(
    τ_grid_pv, sol_pv_low, color='seagreen', lw=2,
    label='ODE, $\\beta_0=0$'
)
ax.plot(
    τ_grid_pv, sol_pv_high, color='steelblue', lw=2,
    label='ODE, $\\beta_0=10$'
)
ax.axhline(β_f_pv, color='red', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_pv:.2f}$')
ax.set_xlabel('Accumulated step size $\\tau$')
ax.set_ylabel('$\\beta(\\tau)$')
ax.legend()

plot_scalar_drift(
    axes[2], β_grid_pv, drift_pv, β_f_pv, 'seagreen'
)

plt.tight_layout()
plt.show()
print(f"REE: β_f = ρ/(1 - lambda*ρ) = {β_f_pv:.4f}")
print(f"Jacobian M = lambda*ρ - 1 = {λ*ρ_pv - 1:.4f}  (< 0: E-stable)")
```

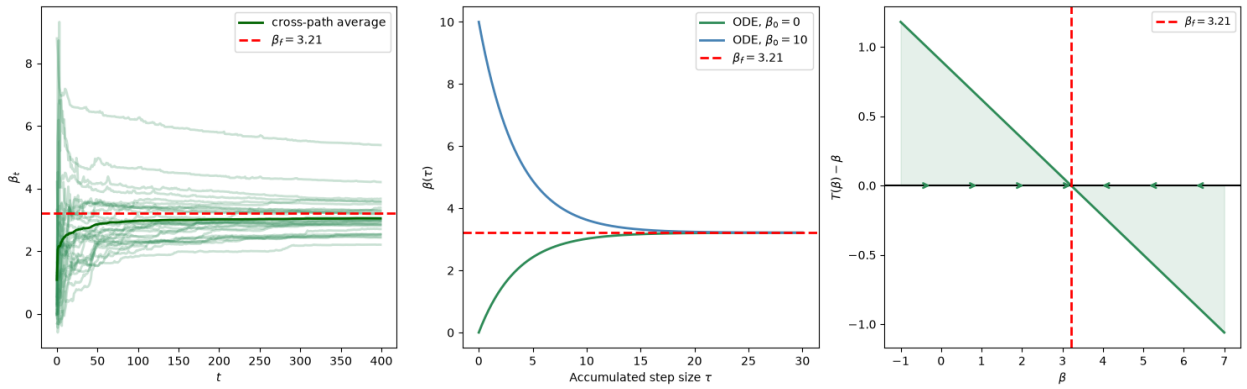


Fig. 138.4: Present-value learning dynamics

```
REE: β_f = ρ/(1 - lambda*ρ) = 3.2143
Jacobian M = lambda*ρ - 1 = -0.2800  (< 0: E-stable)
```

The REE coefficient is larger than ρ because current prices capitalize expected future fundamentals.

Learning converges because each revision in β_t changes expectations, but the discounted persistence term $\lambda\rho$ prevents revisions from feeding back explosively.

138.11 Example 5: Investment under uncertainty

Let's now consider a version of the Lucas-Prescott investment model where agents learn about the aggregate capital stock K_t by regressing on (K_{t-1}, w_{t-1}) where w_t is an exogenous cost shock.

Firms choose investment using a perceived law of motion for aggregate capital.

Those investment choices then determine the actual law of motion for aggregate capital through market clearing.

Learning is therefore about a state variable that agents themselves help create through their investment decisions.

The perceived law of motion is

$$K_t = \beta_1 K_{t-1} + \beta_2 w_{t-1} + \eta_t,$$

while the actual law (from firms' optimal investment decisions and market clearing) is

$$K_t = T_1(\beta) K_{t-1} + T_2(\beta) w_{t-1} + V(\beta) u_t, \tag{138.18}$$

where the nonlinear mappings T_1, T_2 come from solving the firms' linear quadratic control problems.

The two coefficients have different economic roles.

β_1 describes persistence in aggregate capital, while β_2 describes how the exogenous shock w_t affects next period's capital.

Because the optimal investment rule depends on the whole perceived law of motion, the map T is nonlinear.

This is why the larger ODE is harder to analyze here than in the scalar examples with exogenous regressors.

The small ODE decomposes as:

$$\dot{\beta}_1 = T_1(\beta_1) - \beta_1, \quad \dot{\beta}_2 = T_2(\beta_1, \beta_2) - \beta_2, \tag{138.19}$$

and E-stability can be verified analytically for $|\beta_1| < b^{-1/2}$ (where b is the discount factor).

The phase portrait below plots the vector field $T(\beta) - \beta$ over (β_1, β_2) space.

Each arrow shows how the perceived law of motion would be revised if the economy were currently operating under that belief.

The plotted trajectories are deterministic ODE paths, not noisy least-squares sample paths.

```
def T_invest(β, b=0.95, d=1.0, f=1.0, A1=1.0, N=1.0, ρ_w=0.5):
    """Mapping T for the investment model with β = [β1, β2]."""
    b1, b2 = β
    denom1 = 1 - b1 * b + (1 / d) * f**2 * A1 * N
    T1 = (1 - b1 * b) / denom1
    numer2 = 1 - b1 * b + f**2 * A1 * b2 * b * ρ_w
    T2 = (-N / (d * (1 - ρ_w * b))) * (numer2 / denom1) * ρ_w
    return np.array([T1, T2])
```

(continues on next page)

(continued from previous page)

```
def ode_invest(t, β, **kwargs):
    Tb = T_invest(β, **kwargs)
    return Tb - β


params = dict(b=0.95, d=1.0, f=1.0, A1=1.0, N=1.0, ρ_w=0.5)
β1_bound = params["b"] ** (-0.5)
β_f_inv = fsolve(lambda b: T_invest(b, **params) - b, [0.5, 0.1])
print(f"REE: β_f = {β_f_inv}")

eps = 1e-6
J = np.zeros((2, 2))
for j in range(2):
    e = np.zeros(2)
    e[j] = eps
    J[:, j] = (
        T_invest(β_f_inv + e, **params)
        - T_invest(β_f_inv - e, **params)
    ) / (2 * eps)
M = J - np.eye(2)
eigs = la.eigvals(M)
print(f"Jacobian M eigenvalues: {eigs}")
print(f"E-stable: {np.all(eigs.real < 0)}")

fig, ax = plt.subplots(figsize=(8, 6))

b1_grid = np.linspace(-0.1, min(1.0, β1_bound - 1e-3), 20)
b2_grid = np.linspace(-0.8, 0.5, 20)
B1, B2 = np.meshgrid(b1_grid, b2_grid)
U = np.zeros_like(B1)
V_field = np.zeros_like(B2)
for i in range(B1.shape[0]):
    for j in range(B1.shape[1]):
        β_ij = np.array([B1[i, j], B2[i, j]])
        drift = T_invest(β_ij, **params) - β_ij
        U[i, j] = drift[0]
        V_field[i, j] = drift[1]

speed = np.sqrt(U**2 + V_field**2)
speed[speed == 0] = 1e-8
ax.streamplot(b1_grid, b2_grid, U, V_field, color=speed,
              cmap='Blues', density=1.3, linewidth=1)

starts = [(0.1, 0.0), (0.9, 0.4), (0.98, -0.6), (0.3, -0.7)]
colors_traj = ['red', 'darkorange', 'green', 'purple']
for (b10, b20), col in zip(starts, colors_traj):
    sol = solve_ivp(lambda t, β: ode_invest(t, β, **params),
                    [0, 30], [b10, b20],
                    t_eval=np.linspace(0, 30, 300),
                    method='RK45')
    ax.plot(sol.y[0], sol.y[1], color=col, lw=2)
    ax.plot(b10, b20, 'o', color=col, ms=7)

ax.plot(*β_f_inv, 'k*', ms=14, label='REE $\\beta_f$')
ax.set_xlabel('$\\beta_1$', fontsize=12)
ax.set_ylabel('$\\beta_2$', fontsize=12)
```

(continues on next page)

(continued from previous page)

```
ax.legend()
plt.tight_layout()
plt.show()
```

```
REE: β_f = [ 0.38728475 -0.28879438]
Jacobian M eigenvalues: [-1.27718071 -1.35664898]
E-stable: True
```

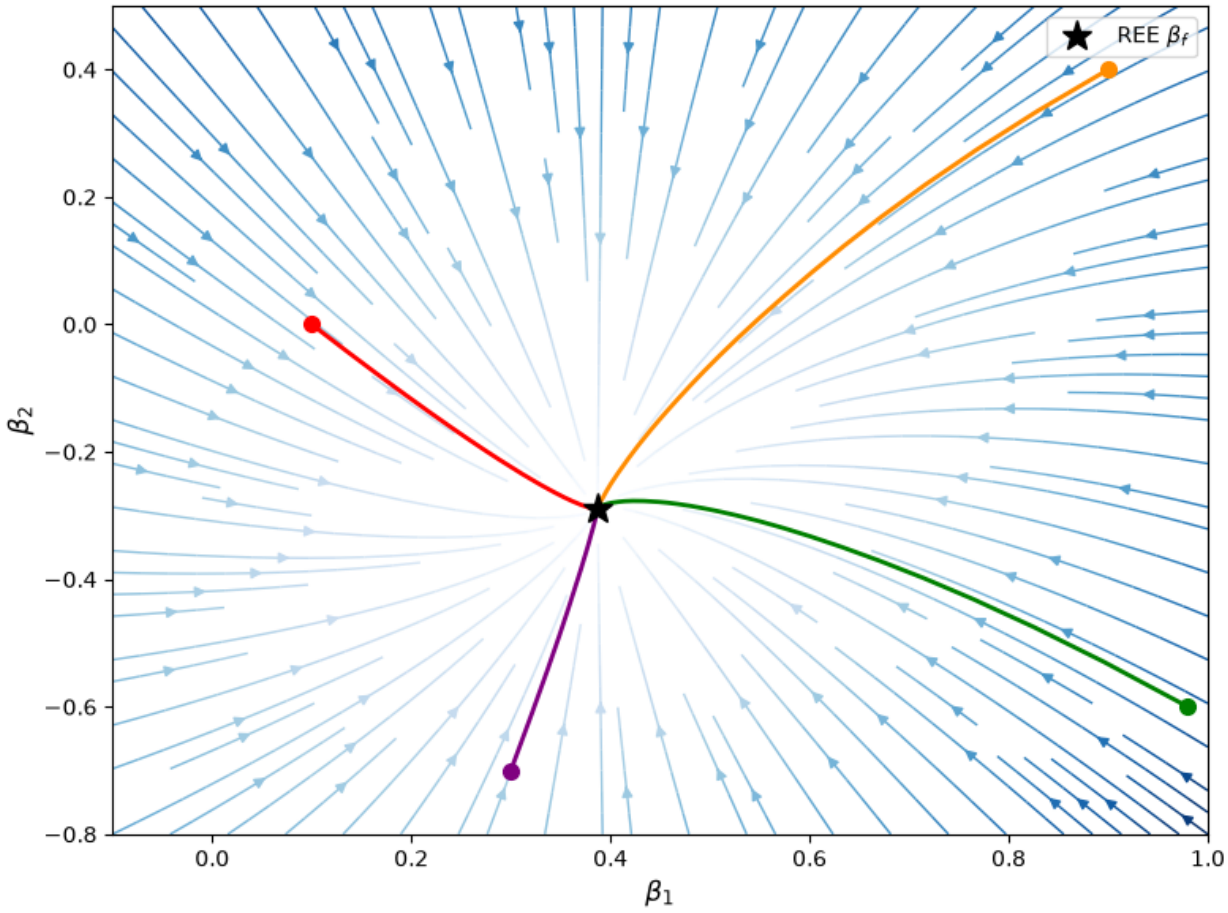


Fig. 138.5: Investment phase portrait

The star marks the REE.

The paths converge because the capital-persistence coefficient β_1 is stabilized first, and then the shock coefficient β_2 adjusts given the limiting value of β_1.

This recursive structure is why the small ODE is tractable even though the full least-squares system has endogenous regressors.

138.12 Connection to rational learning

The framework of Marcet and Sargent [1989] belongs to the program of learning *about* a rational expectations equilibrium, as distinct from learning *within* one — a distinction emphasized by Bray and Kreps [1987].

Learning *within* an REE (the subject of the companion lecture *Rational Learning and Rational Expectations*) refers to Bayesian inference inside a correctly specified model.

In that setting the data-generating process is stationary from the agent's perspective, and Bayes' rule is fully rationalized.

Learning *about* an REE involves an agent who does not know the equilibrium price function.

Because the agent's beliefs shift the equilibrium price, the data the agent uses to update beliefs are themselves generated by a non-stationary process.

As Marcet and Sargent [1989] put it,

> The models do not incorporate fully optimal behavior or rational expectations, because agents operate under the continually falsified assumption that the law of motion is time invariant and known for sure.

Yet it is a compelling rule.

Convergence holds despite the misspecification whenever E-stability combines with the required boundedness and domain-of-attraction conditions.

Unlike Bayesian learning, it does not require strong prior assumptions about the statistical structure of the economy.

138.13 Summary

This lecture studied least-squares learning in the self-referential models of Marcet and Sargent [1989].

The central object is the map T from a perceived law of motion to the actual law of motion generated when agents act on that perception.

A rational expectations equilibrium is a fixed point of this map.

Recursive least squares converges to that fixed point when the associated ODE $\dot{\beta} = T(\beta) - \beta$ is locally stable and the learning process remains in a region where the model is well defined.

It complements rational learning, studied in *Rational Learning and Rational Expectations*, where Bayesian agents learn *within* an equilibrium structure that is already specified.

The Phillips curve lectures apply this machinery to a government that learns: a scalar case in *The Credibility Problem*, the self-confirming equilibrium as a rest point in *Self-Confirming Equilibria*, constant-gain mean dynamics and escapes in *Adaptive Learning and Escape Dynamics* and *Escaping Nash Inflation*, and priors under which E-stability no longer guarantees convergence in *Priors, Escapes, and Learning Cycles*.

138.14 Exercises

Exercise 138.14.1

Consider the scalar model with $T(\beta) = a + b\beta$.

1. Derive a formula for the unique REE β_f in terms of a and b.
2. Show that the small ODE $\dot{\beta} = T(\beta) - \beta$ is globally stable if and only if $b < 1$.

3. Simulate $N = 200$ paths of length $T = 500$ for $a = 1$ and each of $b \in \{0.3, 0.7, 0.9, 0.99\}$ (all less than 1).

Plot the cross-path average of β_t for each b value on the same figure and comment on how the rate of convergence changes as $b \to 1$.

Solution

Here is one solution:

Part 1. The REE satisfies $\beta_f = T(\beta_f) = a + b\beta_f$, so

$$\beta_f(1-b) = a \implies \beta_f = \frac{a}{1-b}.$$

Part 2. The small ODE is $\dot{\beta} = a + b\beta - \beta = a - (1-b)\beta$.

This is linear with slope $-(1-b)$, so the unique fixed point $\beta_f = a/(1-b)$ is globally stable iff $1 - b > 0$, i.e., $b < 1$.

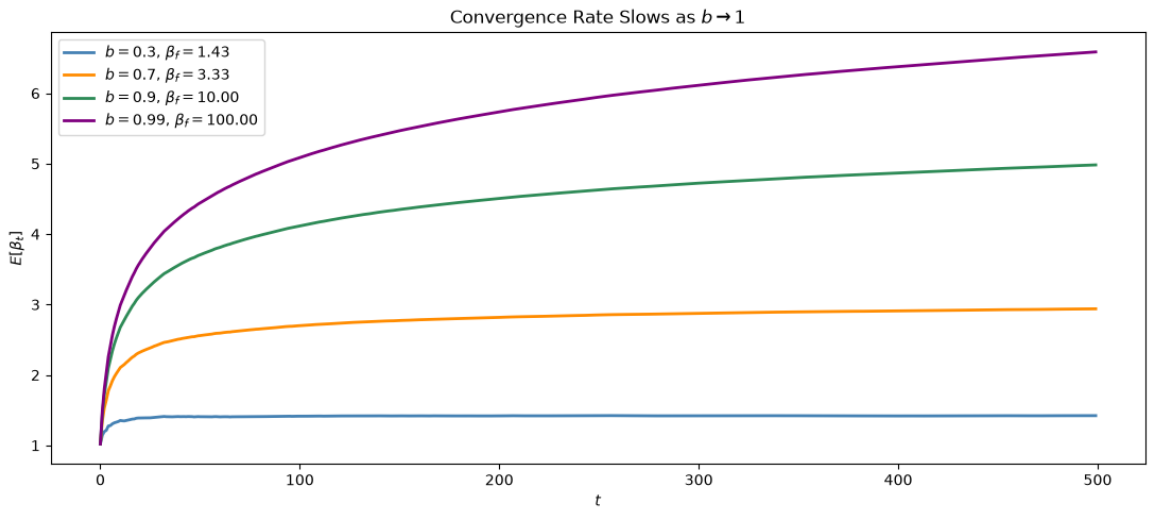
Part 3.

```
a_ex, T_ex, N_ex = 1.0, 500, 200
b_values = [0.3, 0.7, 0.9, 0.99]
colors_ex = ['steelblue', 'darkorange', 'seagreen', 'purple']

fig, ax = plt.subplots(figsize=(11, 5))
for b_val, col in zip(b_values, colors_ex):
    def T_fn(β, b_val=b_val):
        return a_ex + b_val * β

    paths = simulate_rls_scalar(T_fn, σ_u=1.0, β0=0.0,
                                T_periods=T_ex, N_paths=N_ex, seed=0)
    bf = a_ex / (1 - b_val)
    ax.plot(np.mean(paths, axis=0), color=col, lw=2,
            label=f'$b={b_val}$, $\\beta_f={bf:.2f}$')

ax.set_xlabel('$t$')
ax.set_ylabel('$E[\\beta_t]$')
ax.set_title('Convergence Rate Slows as $b \\to 1$')
ax.legend()
plt.tight_layout()
plt.show()
```



As $b \to 1$, the Jacobian $\mathcal{M} = b - 1$ approaches zero, so the ODE becomes slow to return to the fixed point.

Convergence still occurs, but it takes longer.

Exercise 138.14.2

Proposition 138.4.3 states that $\mathbb{P}(\beta_t \to \hat{\beta}) = 0$ for any $\hat{\beta} \neq \beta_f$ in the interior.

1. Using the Bray model with $a = 1$, $b = 0.6$, simulate 100 paths of length $T = 600$ starting from $\beta_0 = 6$ (far from $\beta_f = 2.5$) and show that paths still converge to β_f.
2. Now consider the *unstable* case $b = 1.5$, simulate 50 paths of length $T = 200$ starting from $\beta_0 = -1.9$ (close to the REE $\beta_f = -2$), and describe what happens.

Solution

Here is one solution:

Parts 1 and 2.

```
fig, axes = plt.subplots(1, 2, figsize=(14, 5))

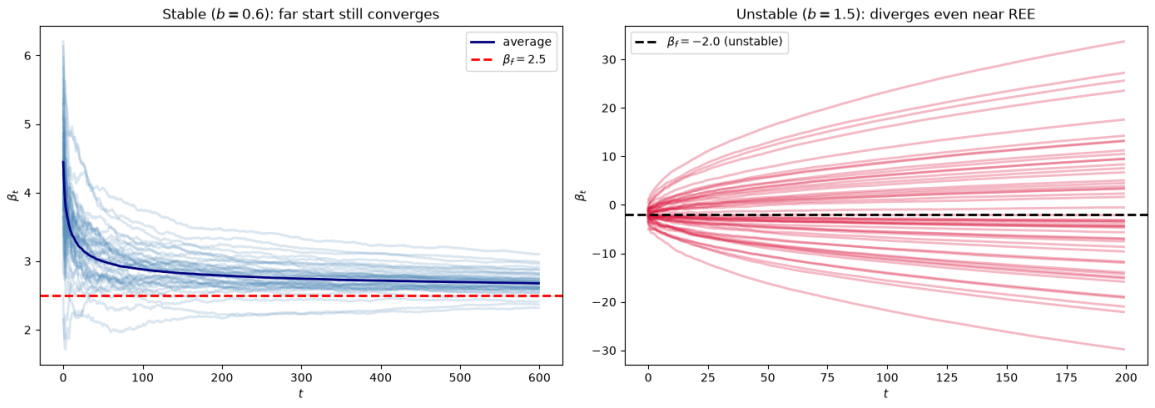
def T_st(β):
    return 1.0 + 0.6 * β


paths_far = simulate_rls_scalar(T_st, 1.0, β0=6.0,
                                T_periods=600, N_paths=100, seed=1)
ax = axes[0]
for i in range(40):
    ax.plot(paths_far[i], color='steelblue', alpha=0.2, lw=2)
ax.plot(np.mean(paths_far, axis=0), color='navy', lw=2, label='average')
ax.axhline(2.5, color='red', ls='--', lw=2, label='$\\beta_f = 2.5$')
ax.set_title('Stable ($b=0.6$): far start still converges')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta_t$')
ax.legend()
```

```
def T_un(β):
    return 1.0 + 1.5 * β


β_f_un = 1.0 / (1 - 1.5)
paths_un = simulate_rls_scalar(T_un, 1.0, β0=-1.9,
                               T_periods=200, N_paths=50, seed=2)
ax = axes[1]
for i in range(50):
    ax.plot(paths_un[i], color='crimson', alpha=0.3, lw=2)
ax.axhline(β_f_un, color='black', ls='--', lw=2,
           label=f'$\\beta_f = {β_f_un}$ (unstable)')
ax.set_title('Unstable ($b=1.5$): diverges even near REE')
ax.set_xlabel('$t$')
ax.set_ylabel('$\\beta_t$')
ax.legend()

plt.tight_layout()
plt.show()
```



Geometrically, the slope $dT/d\beta - 1 = b - 1 = 0.5$ is positive at the REE.

The ODE therefore pushes β away from β_f in both directions.

Exercise 138.14.3

The present-value model: effect of λ on E-stability

In the present-value model (138.15), $T(\beta) = (\lambda\beta + 1)\rho$ and the Jacobian is $\mathcal{M} = \lambda\rho - 1$.

1. For $\rho = 0.9$ and each of $\lambda \in \{0.5, 0.8, 0.95, 0.99\}$, compute β_f and $\mathcal{M}$ and determine whether the REE is E-stable.
2. For the E-stable cases, simulate 100 paths of length $T = 400$ and plot representative noisy paths in calendar time, then plot the limiting ODE paths in accumulated-gain time.
3. At $\lambda = 0.99$, $\mathcal{M} = 0.99\rho - 1 < 0$.

   Simulate paths for this near-boundary case and compare the convergence speed with the $\lambda = 0.5$ case, providing an intuitive explanation.

Solution

Here is one solution:

Part 1.

```
ρ_ex = 0.9
λ_values = [0.5, 0.8, 0.95, 0.99]

print(f"{'lambda':>8}  {'β_f':>10}  {'M = λ*ρ-1':>15}  {'E-stable':>10}")
print("-" * 50)
for lv in λ_values:
    bf = ρ_ex / (1 - lv * ρ_ex) if abs(lv * ρ_ex) < 1 else float('inf')
    M_jac = lv * ρ_ex - 1
    estab = "YES" if M_jac < 0 else "NO"
    print(f"{lv:>8.2f}  {bf:>10.4f}  {M_jac:>15.4f}  {estab:>10}")
```

```
  lambda          β_f        M = λ*ρ-1    E-stable
--------------------------------------------------
    0.50       1.6364          -0.5500         YES
    0.80       3.2143          -0.2800         YES
    0.95       6.2069          -0.1450         YES
    0.99       8.2569          -0.1090         YES
```

Parts 2 and 3.

```
fig, axes = plt.subplots(len(λ_values), 2, figsize=(13, 14))
colors_λ = ['steelblue', 'darkorange', 'seagreen', 'purple']

for row, (lv, col) in enumerate(zip(λ_values, colors_λ)):
    def ode_fn(β, λ_val=lv):
        return (λ_val * β + 1) * ρ_ex - β

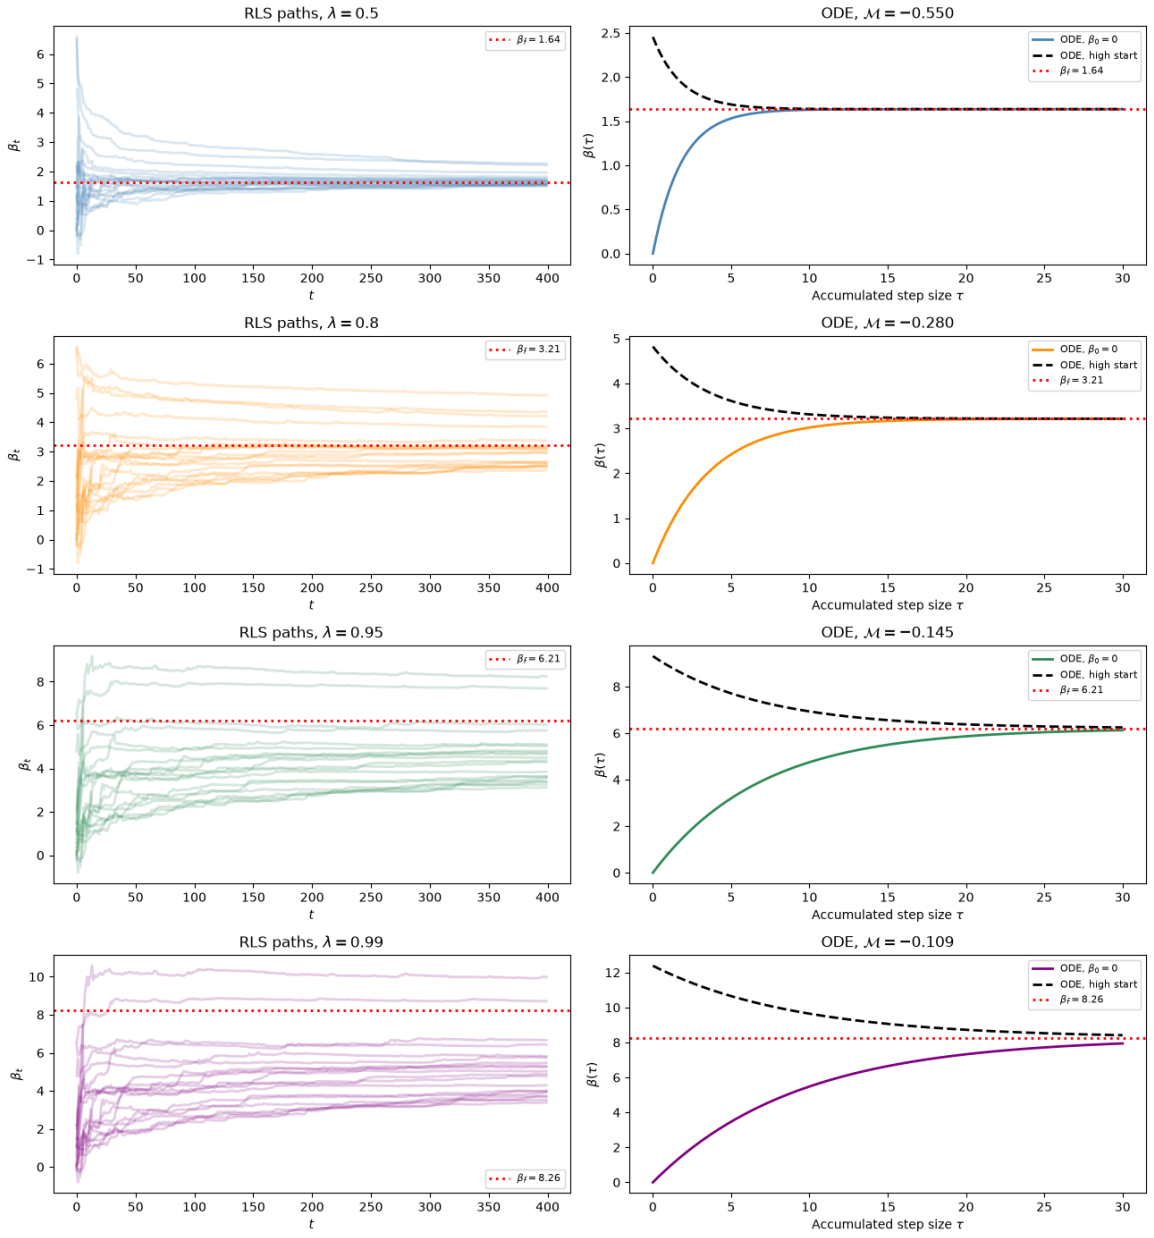
    bf = ρ_ex / (1 - lv * ρ_ex)

    paths_λ = simulate_present_value_rls(
        lv, ρ_ex, 1.0, β0=0.0,
        T_periods=400, N_paths=100, seed=3
    )

    ax = axes[row, 0]
    for i in range(20):
        ax.plot(paths_λ[i], color=col, alpha=0.2, lw=2)
    ax.axhline(bf, color='red', ls=':', lw=2,
               label=f'$\\beta_f={bf:.2f}$')
    ax.set_title(f'RLS paths, $\\lambda={lv}$')
    ax.set_xlabel('$t$')
    ax.set_ylabel('$\\beta_t$')
    ax.legend(fontsize=8)

    τ_grid = np.linspace(0.0, 30.0, 300)
    sol_low = solve_ivp(
        lambda τ, y: [ode_fn(y[0])],
        (0.0, τ_grid[-1]),
        [0.0],
        t_eval=τ_grid,
        max_step=0.05
    ).y[0]
    sol_high = solve_ivp(
        lambda τ, y: [ode_fn(y[0])],
        (0.0, τ_grid[-1]),
        [1.5 * bf],
        t_eval=τ_grid,
        max_step=0.05
    ).y[0]

    ax = axes[row, 1]
```



In each row, the left panel shows finite-sample RLS paths in calendar time and the right panel shows deterministic ODE paths in accumulated-gain time.

When $\lambda = 0.99$, $\mathcal{M} = 0.99\rho - 1 \approx -0.109$ is small in absolute value.

The ODE is nearly flat near β_f, so the restoring force is weak and convergence is slow.

When $\lambda = 0.5$, $\mathcal{M} = -0.55$, which gives a stronger restoring force and faster convergence.

CHAPTER

NINE

AFFINE MODELS OF ASSET PRICES

139.1 Overview

This lecture describes a class of **affine** or **exponential quadratic** models of the stochastic discount factor that have become widely used in empirical finance.

These models are presented in chapter 15 of Ljungqvist and Sargent [2018].

The models discussed here take a different approach from the time-separable CRRA stochastic discount factor of Hansen and Singleton [1983], studied in our *companion lecture*.

The CRRA stochastic discount factor is

$$m_{t+1} = \exp\left(-r_t - \frac{1}{2}\sigma_c^2\gamma^2 - \gamma\sigma_c\varepsilon_{t+1}\right)$$

where ρ is the rate of time preference, γ is the coefficient of relative risk aversion, log consumption growth is $g + \sigma_c\varepsilon_{t+1}$, and $r_t = \rho + \gamma g - \frac{1}{2}\sigma_c^2\gamma^2$.

This model asserts that exposure to the random part of aggregate consumption growth, $\sigma_c\varepsilon_{t+1}$, is the *only* priced risk, the sole source of discrepancies among expected returns across assets.

Empirical difficulties with this specification (the equity premium puzzle, the risk-free rate puzzle, and the Hansen-Jagannathan bounds discussed in Doubts or Variability?) motivate the alternative approach described in this lecture.

Put bluntly, the model to be studied in this lecture declares the Lucas asset pricing model's stochastic discount factor to be a failure.

The **affine model** maintains $\mathbb{E}(m_{t+1}R_{j,t+1}) = 1$ but *divorces* the stochastic discount factor from consumption risk, and consequently, from much of macroeconomics too.

Instead, it

- specifies an analytically tractable stochastic process for m_{t+1}, and
- uses overidentifying restrictions from $\mathbb{E}(m_{t+1}R_{j,t+1}) = 1$ applied to N assets to let the data reveal risks and their prices.

Note

Researchers including [Bansal and Yaron, 2004] and [Hansen *et al.*, 2008] have been less willing to give up on consumption-based models of the stochastic discount factor.

Key applications we study include:

1. *Pricing risky assets*: how risk prices and exposures determine excess returns.

2. *Affine term structure models*: bond yields as affine functions of a state vector (Ang and Piazzesi [2003]).
3. *Risk-neutral probabilities*: a change-of-measure representation of the pricing equation.
4. *Distorted beliefs*: reinterpreting risk price estimates when agents hold systematically biased forecasts (Piazzesi *et al.* [2015]); see also Risk Aversion or Mistaken Beliefs?.

The lecture uses lognormal pricing tools that also appear in *Asset Pricing: Finite State Models* and *Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns*.

The change-of-measure arguments rely on likelihood ratios of the kind studied in *Likelihood Ratio Processes* and *Statistical Divergence Measures*.

We start with the following imports:

```
import numpy as np
import matplotlib.pyplot as plt
from collections import namedtuple
from numpy.linalg import eigvals
from scipy.linalg import solve_discrete_lyapunov
from scipy.stats import norm
```

139.2 The model

139.2.1 State dynamics and short rate

The model has two components.

Component 1 is a vector autoregression that describes the state of the economy and the evolution of the short rate:

$$z_{t+1} = \mu + \phi z_t + C\varepsilon_{t+1} \tag{139.1}$$

$$r_t = \delta_0 + \delta_1^\top z_t \tag{139.2}$$

Here

- ϕ is a stable $m \times m$ matrix,
- C is an $m \times m$ matrix,
- $\varepsilon_{t+1} \sim \mathcal{N}(0, I)$ is an i.i.d. $m \times 1$ random vector,
- z_t is an $m \times 1$ state vector.

Equation (139.2) says that the **short rate** r_t, the net yield on a one-period risk-free claim, is an affine function of the state z_t.

Component 2 is a vector of **risk prices** λ_t and an associated stochastic discount factor m_{t+1}:

$$\lambda_t = \lambda_0 + \lambda_z z_t \tag{139.3}$$

$$\log(m_{t+1}) = -r_t - \frac{1}{2}\lambda_t^\top \lambda_t - \lambda_t^\top \varepsilon_{t+1} \tag{139.4}$$

Here λ_0 is $m \times 1$ and λ_z is $m \times m$.

The entries of λ_t that multiply corresponding entries of the risks ε_{t+1} are called **risk prices** because they determine how exposures to each risk component affect expected returns (as we show below).

Because λ_t is affine in z_t, the stochastic discount factor m_{t+1} is **exponential quadratic** in the state z_t.

We implement the model components as follows.

```
AffineModel = namedtuple('AffineModel',
    ('μ', 'φ', 'C', 'δ_0', 'δ_1', 'λ_0', 'λ_z', 'm', 'φ_rn', 'μ_rn'))

def create_affine_model(μ, φ, C, δ_0, δ_1, λ_0, λ_z):
    """Create an affine term structure model."""
    μ = np.asarray(μ, float)
    φ = np.asarray(φ, float)
    C = np.asarray(C, float)
    δ_1 = np.asarray(δ_1, float)
    λ_0, λ_z = np.asarray(λ_0, float), np.asarray(λ_z, float)
    return AffineModel(μ=μ, φ=φ, C=C, δ_0=float(δ_0), δ_1=δ_1,
                       λ_0=λ_0, λ_z=λ_z, m=len(μ),
                       φ_rn=φ - C @ λ_z, μ_rn=μ - C @ λ_0)

def simulate(model, z0, T, rng=None):
    """Simulate z_{t+1} = μ + φ z_t + C ε_{t+1} for T periods."""
    if rng is None:
        rng = np.random.default_rng(42)
    Z = np.zeros((T + 1, model.m))
    Z[0] = z0
    for t in range(T):
        ε = rng.standard_normal(model.m)
        Z[t + 1] = model.μ + model.φ @ Z[t] + model.C @ ε
    return Z

def short_rate(model, z):
    """Compute r_t = δ_0 + δ_1^⊤ z_t."""
    return model.δ_0 + model.δ_1 @ z

def risk_prices(model, z):
    """Compute λ_t = λ_0 + λ_z z_t."""
    return model.λ_0 + model.λ_z @ z
```

139.2.2 Properties of the SDF

Since $\lambda_t^\top \varepsilon_{t+1}$ is conditionally normal, it follows that

$$\mathbb{E}_t(m_{t+1}) = \exp(-r_t)$$

and

$$\operatorname{std}_t(m_{t+1}) \approx \|\lambda_t\|.$$

Exercise 139.2.1

Show that the SDF defined in (139.4) satisfies

$$\mathbb{E}_t(m_{t+1}) = \exp(-r_t)$$

and

$$\operatorname{std}_t(m_{t+1}) \approx \|\lambda_t\|$$

where $\|\lambda_t\| = \sqrt{\lambda_t^\top \lambda_t}$ denotes the Euclidean norm of the risk price vector.

For the second result, use the lognormal variance formula and the approximations $\exp(x) \approx 1 + x$ and $\exp(-r_t) \approx 1$ for small x and r_t.

Solution

From (139.4), we have

$$m_{t+1} = \exp\left(-r_t - \frac{1}{2}\lambda_t^\top \lambda_t - \lambda_t^\top \varepsilon_{t+1}\right)$$

Since $-\lambda_t^\top \varepsilon_{t+1} \sim \mathcal{N}(0, \lambda_t^\top \lambda_t)$, we have $\mathbb{E}_t[\exp(-\lambda_t^\top \varepsilon_{t+1})] = \exp\left(\frac{1}{2}\lambda_t^\top \lambda_t\right)$.

Therefore,

$$\mathbb{E}_t(m_{t+1}) = \exp(-r_t - \frac{1}{2}\lambda_t^\top \lambda_t)\mathbb{E}_t[\exp(-\lambda_t^\top \varepsilon_{t+1})] = \exp(-r_t)$$

m_{t+1} is conditionally lognormal with $\log m_{t+1} \sim \mathcal{N}(-r_t - \frac{1}{2}\lambda_t^\top \lambda_t, \lambda_t^\top \lambda_t)$.

By the lognormal variance formula $\mathrm{Var}(\exp(X)) = (\exp(\sigma^2) - 1)\exp(2\mu + \sigma^2)$ for $X \sim \mathcal{N}(\mu, \sigma^2)$, we have

$$\begin{aligned}\mathrm{Var}_t(m_{t+1}) &= (\exp(\lambda_t^\top \lambda_t) - 1)\exp(-2r_t) \\ &\approx \lambda_t^\top \lambda_t \exp(-2r_t)\end{aligned}$$

by the approximation $\exp(x) \approx 1 + x$ for small x.

Hence,

$$\mathrm{std}_t(m_{t+1}) \approx \|\lambda_t\| \exp(-r_t)$$

With $\exp(-r_t) \approx 1$ for small r_t, we obtain

$$\mathrm{std}_t(m_{t+1}) \approx \|\lambda_t\|$$

The first equation confirms that r_t is the net yield on a risk-free one-period bond.

That is why r_t is called **the short rate** in the exponential quadratic literature.

The second equation says that the conditional standard deviation of the SDF is approximately the magnitude of the vector of risk prices, a measure of overall **market price of risk**.

The approximation comes from an exact formula: because m_{t+1} is conditionally lognormal,

$$\frac{\mathrm{std}_t(m_{t+1})}{\mathbb{E}_t(m_{t+1})} = \sqrt{\exp(\lambda_t^\top \lambda_t) - 1}$$

By the Hansen–Jagannathan bound [Hansen and Jagannathan, 1991], this ratio bounds the conditional Sharpe ratio of every excess return, so $\|\lambda_t\|$ tells us how large Sharpe ratios can be.

139.3 Pricing risky assets

139.3.1 Lognormal returns

Consider a risky asset j whose gross return has a lognormal conditional distribution:

$$R_{j,t+1} = \exp\left(\nu_t(j) - \frac{1}{2}\alpha_t(j)^\top \alpha_t(j) + \alpha_t(j)^\top \varepsilon_{t+1}\right) \tag{139.5}$$

where the **exposure vector**

$$\alpha_t(j) = \alpha_0(j) + \alpha_z(j)\, z_t \tag{139.6}$$

Here $\alpha_0(j)$ is $m \times 1$ and $\alpha_z(j)$ is $m \times m$.

The components of $\alpha_t(j)$ express the **exposures** of $\log R_{j,t+1}$ to corresponding components of the risk vector ε_{t+1}.

The specification (139.5) implies $\mathbb{E}_t R_{j,t+1} = \exp(\nu_t(j))$, so $\nu_t(j)$ is the log of the expected gross return.

139.3.2 Expected excess returns

Applying the pricing equation $\mathbb{E}_t(m_{t+1} R_{j,t+1}) = 1$ together with the formula for the mean of a lognormal random variable gives

$$\nu_t(j) = r_t + \alpha_t(j)^\top \lambda_t \tag{139.7}$$

Exercise 139.3.1

Using the SDF (139.4) and the return specification (139.5), derive the expected excess return formula (139.7):

$$\nu_t(j) = r_t + \alpha_t(j)^\top \lambda_t$$

Hint: Start by computing $\log(m_{t+1} R_{j,t+1})$, identify its conditional distribution, and apply the pricing condition $\mathbb{E}_t(m_{t+1} R_{j,t+1}) = 1$.

Solution

Combining (139.4) and (139.5), we get

$$\log(m_{t+1} R_{j,t+1}) = -r_t + \nu_t(j) - \frac{1}{2}\lambda_t^\top \lambda_t - \frac{1}{2}\alpha_t(j)^\top \alpha_t(j) + (\alpha_t(j) - \lambda_t)^\top \varepsilon_{t+1}$$

This is conditionally normal with mean $\mu = -r_t + \nu_t(j) - \frac{1}{2}\lambda_t^\top \lambda_t - \frac{1}{2}\alpha_t(j)^\top \alpha_t(j)$ and variance $\sigma^2 = (\alpha_t(j) - \lambda_t)^\top (\alpha_t(j) - \lambda_t)$.

Since $\mathbb{E}_t[\exp(X)] = \exp(\mu + \frac{1}{2}\sigma^2)$ for $X \sim \mathcal{N}(\mu, \sigma^2)$, the pricing condition $\mathbb{E}_t(m_{t+1} R_{j,t+1}) = 1$ requires $\mu + \frac{1}{2}\sigma^2 = 0$.

Expanding $\frac{1}{2}\sigma^2 = \frac{1}{2}\alpha_t(j)^\top \alpha_t(j) - \alpha_t(j)^\top \lambda_t + \frac{1}{2}\lambda_t^\top \lambda_t$ and adding to μ, the $\frac{1}{2}\lambda_t^\top \lambda_t$ and $\frac{1}{2}\alpha_t(j)^\top \alpha_t(j)$ terms cancel, leaving

$$-r_t + \nu_t(j) - \alpha_t(j)^\top \lambda_t = 0$$

which gives (139.7).

This is a central result.

It says:

> The log expected gross return on asset j equals the short rate plus the inner product of the asset's exposure vector $\alpha_t(j)$ with the risk price vector λ_t.

Each component of λ_t prices the corresponding component of ε_{t+1}.

An asset that loads heavily on a risk component with a large risk price earns a correspondingly high expected return.

Exercise 139.3.2

Suppose that a representative agent has CRRA utility with risk aversion γ and discount factor $\beta = e^{-\rho}$, and that log consumption growth is $\log(c_{t+1}/c_t) = g + \sigma_c^\top \varepsilon_{t+1}$, where σ_c is an $m \times 1$ vector.

1. Show that $m_{t+1} = \beta(c_{t+1}/c_t)^{-\gamma}$ has the form (139.4) with a constant risk price vector $\lambda_t = \gamma\sigma_c$, and find r_t.
2. Define the conditional Sharpe ratio of the log return on asset j as $(\nu_t(j) - r_t)/\|\alpha_t(j)\|$. Use (139.7) and the Cauchy–Schwarz inequality to show that the largest such Sharpe ratio is $\|\lambda_t\|$, and that it is attained by returns whose exposure vector is proportional to λ_t.
3. Suppose that $\|\sigma_c\| = 0.02$ per year and that some asset has an annual Sharpe ratio of 0.4. Using the exact bound $\sqrt{\exp(\lambda^\top\lambda) - 1} \geq 0.4$, compute the smallest γ consistent with the CRRA model, and compare it with the first-order answer $\gamma \geq 0.4/0.02$.

Solution

Part 1. Taking logs,

$$\log m_{t+1} = -\rho - \gamma g - \gamma\sigma_c^\top \varepsilon_{t+1}$$

Matching the shock term with (139.4) gives $\lambda_t = \gamma\sigma_c$.

Matching the constant gives $-r_t - \frac{1}{2}\gamma^2\sigma_c^\top\sigma_c = -\rho - \gamma g$, so

$$r_t = \rho + \gamma g - \frac{1}{2}\gamma^2\sigma_c^\top\sigma_c$$

This is the CRRA model from the overview, now written as a special case of the affine model with $\lambda_z = 0$.

Part 2. By (139.7), the Sharpe ratio is $\alpha_t(j)^\top\lambda_t/\|\alpha_t(j)\|$.

The Cauchy–Schwarz inequality gives $\alpha_t(j)^\top\lambda_t \leq \|\alpha_t(j)\|\,\|\lambda_t\|$, with equality exactly when $\alpha_t(j)$ is a positive multiple of λ_t.

The following check maximizes the Sharpe ratio over many random exposure vectors.

```
rng = np.random.default_rng(1234)
λ_example = np.array([0.3, 0.1])
α_draws = rng.standard_normal((20_000, 2))
sharpe = α_draws @ λ_example / np.linalg.norm(α_draws, axis=1)
print(f"largest Sharpe ratio over draws: {sharpe.max():.5f}")
print(f"||λ||:                           {np.linalg.norm(λ_example):.5f}")
```

```
largest Sharpe ratio over draws: 0.31623
||λ||:                           0.31623
```

Part 3.

```
σ_c_norm, target_sr = 0.02, 0.4
λ_norm_required = np.sqrt(np.log(1 + target_sr**2))
print(f"required ||λ|| (exact bound): {λ_norm_required:.4f}")
print(f"implied γ (exact bound):      {λ_norm_required / σ_c_norm:.1f}")
print(f"implied γ (first order):      {target_sr / σ_c_norm:.1f}")
```

```
required ||λ|| (exact bound): 0.3853
implied γ (exact bound):      19.3
implied γ (first order):      20.0
```

Either way, the CRRA model needs a coefficient of relative risk aversion near 20.

This is the equity premium puzzle expressed in the language of this lecture.

The affine model sidesteps the puzzle by treating λ_t as free parameters to be estimated from asset returns rather than tying them to consumption growth.

139.4 Affine term structure of yields

One of the most important applications is the **affine term structure model** studied by Ang and Piazzesi [2003].

139.4.1 Bond prices

Let $p_t(n)$ be the price at time t of a risk-free pure discount bond maturing at $t+n$ (paying one unit of consumption).

The one-period gross return on holding an $(n+1)$-period bond from t to $t+1$ is

$$R_{t+1} = \frac{p_{t+1}(n)}{p_t(n+1)}$$

The pricing equation $\mathbb{E}_t(m_{t+1} R_{t+1}) = 1$ implies

$$p_t(n+1) = \mathbb{E}_t(m_{t+1}\, p_{t+1}(n)) \tag{139.8}$$

with the initial condition

$$p_t(1) = \mathbb{E}_t(m_{t+1}) = \exp(-r_t) = \exp(-\delta_0 - \delta_1^\top z_t).$$

139.4.2 Exponential affine prices

The recursion (139.8) has an **exponential affine** solution:

$$p_t(n) = \exp(\bar{A}_n + \bar{B}_n^\top z_t) \tag{139.9}$$

where the scalar $\bar{A}_n$ and the $m \times 1$ vector $\bar{B}_n$ satisfy the **Riccati difference equations**

$$\bar{A}_{n+1} = \bar{A}_n + \bar{B}_n^\top(\mu - C\lambda_0) + \frac{1}{2}\bar{B}_n^\top CC^\top \bar{B}_n - \delta_0 \tag{139.10}$$

$$\bar{B}_{n+1}^\top = \bar{B}_n^\top(\phi - C\lambda_z) - \delta_1^\top \tag{139.11}$$

with initial conditions $\bar{A}_1 = -\delta_0$ and $\bar{B}_1 = -\delta_1$.

Exercise 139.4.1

Derive the Riccati difference equations (139.10) and (139.11) by substituting the conjectured bond price (139.9) into the pricing recursion (139.8) and matching coefficients.

Hint: Substitute $p_{t+1}(n) = \exp(\bar{A}_n + \bar{B}_n^\top z_{t+1})$ and $\log m_{t+1}$ from (139.4) into (139.8).

Use the state dynamics (139.1) to express z_{t+1} in terms of z_t and ε_{t+1}, then evaluate the conditional expectation using the lognormal moment generating function.

Solution

We want to show that if $p_t(n) = \exp(\bar{A}_n + \bar{B}_n^\top z_t)$, then the recursion $p_t(n+1) = \mathbb{E}_t(m_{t+1}\, p_{t+1}(n))$ yields $p_t(n+1) = \exp(\bar{A}_{n+1} + \bar{B}_{n+1}^\top z_t)$ with $\bar{A}_{n+1}$ and $\bar{B}_{n+1}$ given by (139.10) and (139.11).

From (139.4) and (139.9),

$$\log(m_{t+1}\, p_{t+1}(n)) = -r_t - \frac{1}{2}\lambda_t^\top \lambda_t - \lambda_t^\top \varepsilon_{t+1} + \bar{A}_n + \bar{B}_n^\top z_{t+1}$$

Substituting $z_{t+1} = \mu + \phi z_t + C\varepsilon_{t+1}$ from (139.1) and $r_t = \delta_0 + \delta_1^\top z_t$ from (139.2) gives

$$\log(m_{t+1}\, p_{t+1}(n)) = \bar{A}_n + \bar{B}_n^\top \mu - \delta_0 + (\bar{B}_n^\top \phi - \delta_1^\top) z_t - \frac{1}{2}\lambda_t^\top \lambda_t + (\bar{B}_n^\top C - \lambda_t^\top)\varepsilon_{t+1}$$

Since $\varepsilon_{t+1} \sim \mathcal{N}(0, I)$, and writing the exponent as $a + b^\top \varepsilon_{t+1}$ where $b = C^\top \bar{B}_n - \lambda_t$, we have

$$\mathbb{E}_t[\exp(a + b^\top \varepsilon_{t+1})] = \exp\left(a + \frac{1}{2} b^\top b\right)$$

Computing $\frac{1}{2} b^\top b$:

$$\frac{1}{2}(\bar{B}_n^\top C - \lambda_t^\top)(\bar{B}_n^\top C - \lambda_t^\top)^\top = \frac{1}{2}\bar{B}_n^\top C C^\top \bar{B}_n - \bar{B}_n^\top C \lambda_t + \frac{1}{2}\lambda_t^\top \lambda_t$$

The $\frac{1}{2}\lambda_t^\top \lambda_t$ cancels with the $-\frac{1}{2}\lambda_t^\top \lambda_t$ already in a, and $-\bar{B}_n^\top C \lambda_t = -\bar{B}_n^\top C(\lambda_0 + \lambda_z z_t)$.

$$\log p_t(n+1) = \underbrace{\bar{A}_n + \bar{B}_n^\top(\mu - C\lambda_0) + \frac{1}{2}\bar{B}_n^\top C C^\top \bar{B}_n - \delta_0}_{\bar{A}_{n+1}} + \underbrace{(\bar{B}_n^\top(\phi - C\lambda_z) - \delta_1^\top)}_{\bar{B}_{n+1}^\top} z_t$$

Matching the constant and the coefficient on z_t gives the Riccati equations (139.10) and (139.11).

Setting $n = 0$ with $p_t(1) = \exp(-r_t) = \exp(-\delta_0 - \delta_1^\top z_t)$ gives $\bar{A}_1 = -\delta_0$ and $\bar{B}_1 = -\delta_1$.

139.4.3 Yields

The **yield to maturity** on an n-period bond is the constant rate y at which one would discount the face value to obtain the observed price, i.e., $p_t(n) = e^{-n\,y}$.

Solving for y gives

$$y_t(n) = -\frac{\log p_t(n)}{n}$$

Substituting (139.9) gives

$$y_t(n) = A_n + B_n^\top z_t \tag{139.12}$$

where $A_n = -\bar{A}_n/n$ and $B_n = -\bar{B}_n/n$.

Yields are affine functions of the state vector z_t.

This is the defining property of affine term structure models.

We now implement the bond pricing formulas (139.10), (139.11), and (139.12).

```
def bond_coefficients(model, n_max):
    """Compute (A_bar_n, B_bar_n) for n = 1, ..., n_max."""
    A_bar = np.zeros(n_max + 1)
    B_bar = np.zeros((n_max + 1, model.m))
    A_bar[1], B_bar[1] = -model.δ_0, -model.δ_1
    CC = model.C @ model.C.T
    for n in range(1, n_max):
        Bn = B_bar[n]
        A_bar[n + 1] = (A_bar[n] + Bn @ model.μ_rn
                        + 0.5 * Bn @ CC @ Bn
                        - model.δ_0)
        B_bar[n + 1] = model.ψ_rn.T @ Bn - model.δ_1
    return A_bar, B_bar

def compute_yields(model, z, n_max):
    """Compute yield curve y_t(n) for n = 1, ..., n_max."""
    A_bar, B_bar = bond_coefficients(model, n_max)
    return np.array([(-A_bar[n] - B_bar[n] @ z) / n
                     for n in range(1, n_max + 1)])

def bond_prices(model, z, n_max):
    """Compute bond prices p_t(n) for n = 1, ..., n_max."""
    A_bar, B_bar = bond_coefficients(model, n_max)
    return np.array([np.exp(A_bar[n] + B_bar[n] @ z)
                     for n in range(1, n_max + 1)])
```

139.4.4 A one-factor Gaussian example

To build intuition, we start with a single-factor ($m = 1$) Gaussian model.

With $m = 1$, the state z_t follows an AR(1) process $z_{t+1} = \mu + \phi z_t + C\varepsilon_{t+1}$.

The unconditional standard deviation of z_t is $\sigma_z = C/\sqrt{1-\phi^2}$, which determines the range of short rates the model generates via $r_t = \delta_0 + \delta_1 z_t$.

```
# One-factor Gaussian model (quarterly)
μ      = np.array([0.0])
φ      = np.array([[0.95]])
C      = np.array([[1.0]])
δ_0    = 0.01                  # 1%/quarter ≈ 4% p.a.
δ_1    = np.array([0.001])
λ_0    = np.array([-0.05])
λ_z    = np.array([[-0.01]])

model_1f = create_affine_model(μ, φ, C, δ_0, δ_1, λ_0, λ_z)
```

139.4.5 Yield curve shapes

We compute yield curves $y_t(n)$ across a range of short-rate states z_t.

```
n_max_1f = 200
maturities_1f = np.arange(1, n_max_1f + 1)

z_low  = np.array([-5.0])
z_mid  = np.array([0.0])
z_high = np.array([5.0])

fig, ax = plt.subplots(figsize=(9, 5.5))

r_low = short_rate(model_1f, z_low) * 4 * 100
r_mid = short_rate(model_1f, z_mid) * 4 * 100
r_high = short_rate(model_1f, z_high) * 4 * 100

for z, label in [
    (z_low,  f"Low state  ($y_t(1) = ${r_low:.1f}%)"),
    (z_mid,  f"Median state ($y_t(1) = ${r_mid:.1f}%)"),
    (z_high, f"High state ($y_t(1) = ${r_high:.1f}%)"),
]:
    y = compute_yields(model_1f, z, n_max_1f) * 4 * 100
    line, = ax.plot(maturities_1f, y, lw=2.2, label=label)
    ax.plot(1, y[0], 'o', color=line.get_color(), ms=7, zorder=5)

r_bar = short_rate(model_1f, np.array([0.0])) * 4 * 100
ax.axhline(r_bar, color='grey', ls=':', lw=1.2, alpha=0.7,
           label=f"Mean short rate ({r_bar:.1f}%)")

# Long-run yield: B_bar_n converges, so y_inf = lim -A_bar_n / n
φ_Cλ = (model_1f.φ_rn)[0, 0]          # φ - Cλ_z (scalar)
B_inf = -model_1f.δ_1[0] / (1 - φ_Cλ) # fixed point of B recursion
A_increment = (B_inf * model_1f.μ_rn[0]
               + 0.5 * B_inf**2 * (model_1f.C @ model_1f.C.T)[0, 0]
               - model_1f.δ_0)
y_inf = -A_increment * 4 * 100         # annualised %
ax.axhline(y_inf, color='black', ls='--', lw=1.2, alpha=0.7,
           label=f"Long-run yield ({y_inf:.1f}%)")

ax.set_xlabel("Maturity (quarters)")
ax.set_ylabel("Yield (% per annum)")
ax.legend(fontsize=10, loc='best')
ax.set_xlim(1, n_max_1f)
```

(continues on next page)

(continued from previous page)

```
ax2 = ax.twiny()
ax2.set_xlim(ax.get_xlim())
year_ticks = [4, 40, 80, 120, 160, 200]
ax2.set_xticks(year_ticks)
ax2.set_xticklabels([f"{t/4:.0f}y" for t in year_ticks])
ax2.set_xlabel("Maturity (years)")

plt.tight_layout()
plt.show()
```

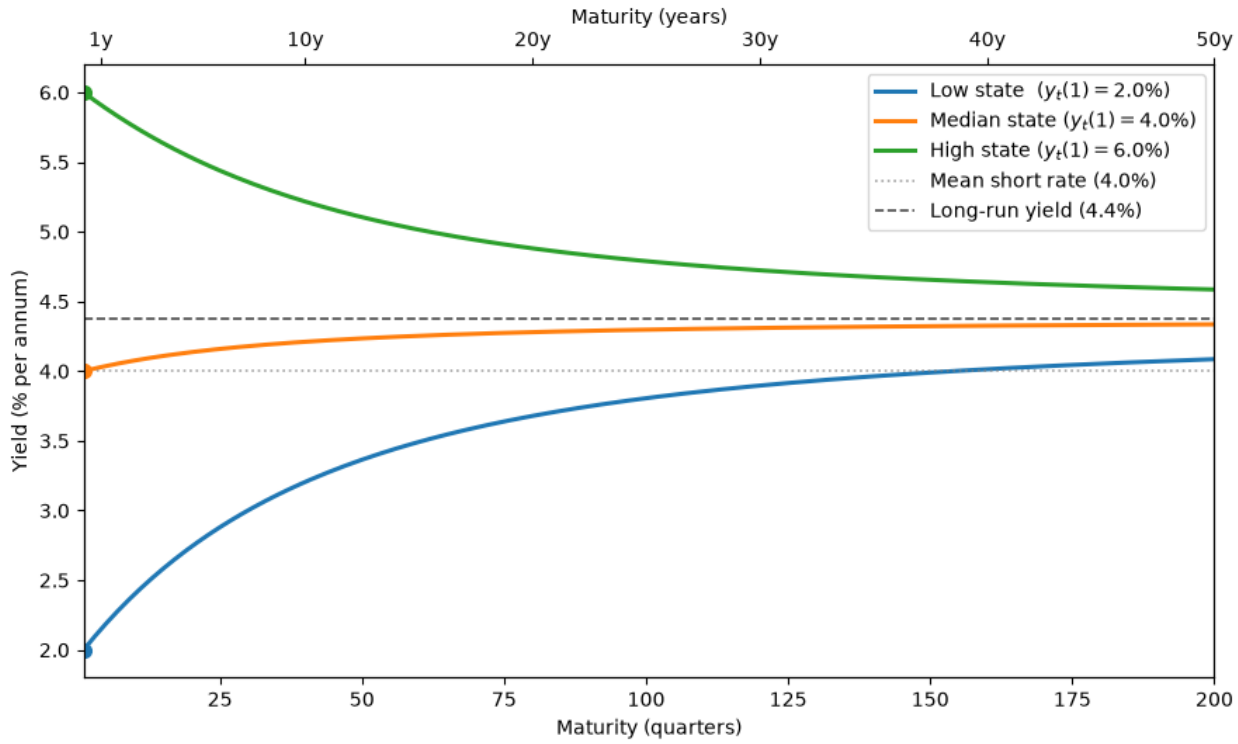


Fig. 139.1: Yield curves under the one-factor affine model

When the short rate is low, the yield curve is upward-sloping, while when the short rate is high, it is downward-sloping.

All three curves approach the same long-run yield y_∞, which lies above the mean short rate δ_0.

The approach is slow.

Because $\bar{B}_n$ converges to a finite limit, the state enters $y_t(n)$ through $\bar{B}_n^\top z_t / n$, so gaps between curves shrink only at rate $1/n$.

That is why we plot maturities out to 50 years: at 15 years the three curves are still spread between about 3.5% and 5.0%, while at 50 years they lie between about 4.1% and 4.6%.

Exercise 139.4.2

Show that the long-run yield satisfies

$$y_\infty = \delta_0 - \bar{B}_\infty^\top(\mu - C\lambda_0) - \tfrac{1}{2}\bar{B}_\infty^\top CC^\top \bar{B}_\infty \tag{139.13}$$

where $\bar{B}_\infty = -(I - (\phi - C\lambda_z)^\top)^{-1}\delta_1$ is the fixed point of the recursion (139.11).

Then explain why $y_\infty > \delta_0$ under this parameterization.

Hint: Use (139.12) and the Riccati equations (139.10)–(139.11).

For the inequality, consider each subtracted term separately.

Solution

The recursion (139.11) is a linear difference equation $\bar{B}_{n+1} = (\phi - C\lambda_z)^\top \bar{B}_n - \delta_1$.

When $\phi - C\lambda_z$ has eigenvalues inside the unit circle, $\bar{B}_n$ converges to $\bar{B}_\infty = -(I - (\phi - C\lambda_z)^\top)^{-1}\delta_1$.

Since $\bar{B}_\infty$ is finite, $\bar{B}_n^\top z_t/n \to 0$ in (139.12), so $y_t(n) \to \lim_{n\to\infty} -\bar{A}_n/n$ regardless of z_t.

To find this limit, write $\bar{A}_n = \bar{A}_1 + \sum_{k=1}^{n-1}(\bar{A}_{k+1} - \bar{A}_k)$.

By (139.10), each increment depends on $\bar{B}_k$, which converges to $\bar{B}_\infty$, so the increment converges to $L \equiv \bar{B}_\infty^\top(\mu - C\lambda_0) + \frac{1}{2}\bar{B}_\infty^\top CC^\top \bar{B}_\infty - \delta_0$.

Therefore $\bar{A}_n/n \to L$ and $y_\infty = -L$, giving (139.13).

To see why $y_\infty > \delta_0$, note that the two subtracted terms in (139.13) have opposite signs under this parameterization.

The quadratic term $\frac{1}{2}\bar{B}_\infty^\top CC^\top \bar{B}_\infty = \frac{1}{2}\|C^\top \bar{B}_\infty\|^2 \geq 0$ always.

This is a **convexity effect** from Jensen's inequality that pushes y_∞ below δ_0.

The linear term $\bar{B}_\infty^\top(\mu - C\lambda_0)$ is negative because $\bar{B}_\infty < 0$ (since $\delta_1 > 0$) while $\mu - C\lambda_0 > 0$ (since $\lambda_0 < 0$).

Subtracting this negative quantity raises y_∞ above δ_0.

This is a **risk-premium effect**: positive term premiums tilt the average yield curve upward.

Under this parameterization the risk-premium effect dominates the convexity effect, so $y_\infty > \delta_0$.

Let's also simulate the short rate path:

```
T = 200
Z = simulate(model_1f, np.array([0.0]), T)
short_rates = np.array([short_rate(model_1f, Z[t]) * 4 * 100
                        for t in range(T + 1)])
r_bar_pct = short_rate(model_1f, np.array([0.0])) * 4 * 100

fig, ax = plt.subplots(figsize=(10, 4))
quarters = np.arange(T + 1)
line, = ax.plot(quarters, short_rates, lw=1.3)
ax.axhline(r_bar_pct, ls="--", lw=1.3,
           label=f"Unconditional mean ({r_bar_pct:.1f}%)")
ax.fill_between(quarters, short_rates, r_bar_pct,
                alpha=0.08, color=line.get_color())
ax.set_xlabel("Quarter")
ax.set_ylabel("Short rate (% p.a.)")
ax.set_xlim(0, T)
```

(continues on next page)

(continued from previous page)

```
ax.legend(fontsize=11)
plt.tight_layout()
plt.show()
```

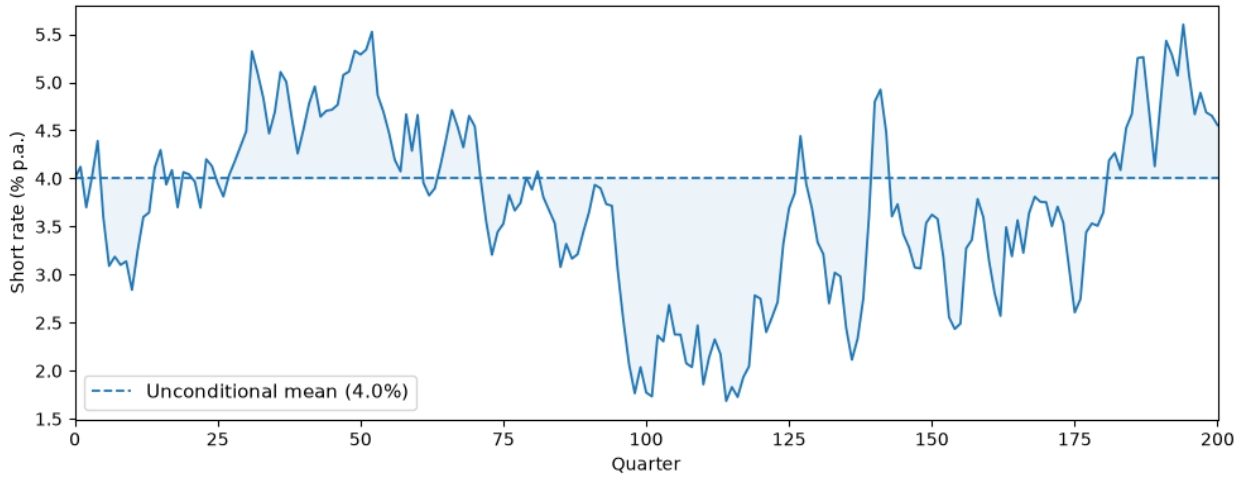


Fig. 139.2: Simulated short rate path

Because r_t is an affine function of a Gaussian state, a Gaussian affine model assigns positive probability to negative short rates.

Under the stationary distribution of the one-factor model, r_t is normal with mean δ_0 and standard deviation $\delta_1 \sigma_z$, which lets us compute that probability.

```
σ_z = model_1f.C[0, 0] / np.sqrt(1 - model_1f.ψ[0, 0]**2)
σ_r = model_1f.δ_1[0] * σ_z
print(f"Stationary std of r_t: {σ_r * 4 * 100:.2f}% p.a.")
print(f"P(r_t < 0):            {norm.cdf(-model_1f.δ_0 / σ_r):.2e}")
```

```
Stationary std of r_t: 1.28% p.a.
P(r_t < 0):            8.97e-04
```

The probability is small in this calibration, but it can be substantial in calibrations fit to periods of low interest rates.

Square-root models such as Cox *et al.* [1985] and shadow-rate models such as Black [1995] are two ways of keeping nominal rates non-negative, at the cost of some of the tractability that the Gaussian specification delivers.

139.4.6 A two-factor model

To match richer yield-curve dynamics, practitioners routinely use $m \geq 2$ factors.

We now introduce a two-factor specification with state $z_t = (z_{1t},\, z_{2t})^\top$, where

$$z_{t+1} = \mu + \phi\, z_t + C\, \varepsilon_{t+1}, \qquad \phi = \begin{pmatrix} 0.97 & -0.03 \\ 0 & 0.90 \end{pmatrix}, \qquad C = I_2$$

The first factor z_{1t} is highly persistent ($\phi_{11} = 0.97$) and drives most of the variation in the short rate through δ_1, so we interpret it as a **level** factor.

The second factor z_{2t} mean-reverts faster ($\phi_{22} = 0.90$) and affects the short rate with a smaller loading, capturing the **slope** of the yield curve.

The off-diagonal entry $\phi_{12} = -0.03$ allows the level factor to respond to the current slope state z_{2t}.

The short rate is $r_t = \delta_0 + \delta_1^\top z_t$ with $\delta_1 = (0.002,\ 0.001)^\top$, so both factors raise the short rate when positive, but the level factor has twice the impact.

Risk prices are $\lambda_t = \lambda_0 + \lambda_z z_t$ with $\lambda_0 = (-0.01,\ -0.005)^\top$ and $\lambda_z = \text{diag}(-0.005,\ -0.003)$.

The negative diagonal entries of λ_z make $\phi - C\lambda_z$ have larger eigenvalues than ϕ, so the state is more persistent under the risk-neutral measure and the yield curve is more sensitive to the current state at long horizons.

```
# Two-factor model: z = [level, slope]
μ_2  = np.array([0.0,  0.0])
φ_2  = np.array([[0.97, -0.03],
                 [0.00,  0.90]])
C_2  = np.eye(2)
δ_0_2 = 0.01
δ_1_2 = np.array([0.002, 0.001])
λ_0_2 = np.array([-0.01, -0.005])
λ_z_2 = np.array([[-0.005, 0.0],
                  [ 0.0, -0.003]])

model_2f = create_affine_model(μ_2, φ_2, C_2, δ_0_2, δ_1_2, λ_0_2, λ_z_2)

print(f"Eigenvalues of φ:       {eigvals(φ_2).real.round(4)}")
print(f"Eigenvalues of φ - Cλ_z: {eigvals(model_2f.φ_rn).real.round(4)}")
```

```
Eigenvalues of φ:       [0.97 0.9 ]
Eigenvalues of φ - Cλ_z: [0.975 0.903]
```

This confirms that the eigenvalues of $\phi - C\lambda_z$ are larger than those of ϕ, so the state is more persistent under the risk-neutral measure.

The following figure shows yield curves across different states of the world, as well as the factor loadings $B_{n,1}$ and $B_{n,2}$ that determine how yields load on the level and slope factors at each maturity

```
n_max_2f = 60
maturities_2f = np.arange(1, n_max_2f + 1)

states = {
    "Normal":              np.array([0.0,   0.0]),
    "Low short rate":      np.array([-4.0,  3.0]),
    "High short rate":     np.array([4.0,  -3.0]),
}

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(14, 5.5))

for label, z in states.items():
    r_now = short_rate(model_2f, z) * 4 * 100
    y = compute_yields(model_2f, z, n_max_2f) * 4 * 100
    line, = ax1.plot(maturities_2f, y, lw=2.2,
                     label=f"{label} (r₁ = {r_now:.1f}%)")
    ax1.plot(1, y[0], 'o', color=line.get_color(), ms=7, zorder=5)

ax1.set_xlabel("Maturity (quarters)")
ax1.set_ylabel("Yield (% p.a.)")
ax1.legend(fontsize=10)
```

(continues on next page)

(continued from previous page)

```
ax1.set_xlim(1, n_max_2f)

A_bar, B_bar = bond_coefficients(model_2f, n_max_2f)
ns = np.arange(1, n_max_2f + 1)
B_n = np.array([-B_bar[n] / n for n in ns])

ax2.plot(ns, B_n[:, 0], lw=2.2,
         label=r"Level loading $B_{n,1}$")
ax2.plot(ns, B_n[:, 1], lw=2.2,
         label=r"Slope loading $B_{n,2}$")
ax2.axhline(0, color='black', lw=0.6)
ax2.set_xlabel("Maturity (quarters)")
ax2.set_ylabel(r"Yield loading $B_{n,k}$")
ax2.legend(fontsize=11)
ax2.set_xlim(1, n_max_2f)

for ax in (ax1, ax2):
    ax_top = ax.twiny()
    ax_top.set_xlim(ax.get_xlim())
    year_ticks = [4, 12, 20, 40, 60]
    ax_top.set_xticks(year_ticks)
    ax_top.set_xticklabels([f"{t/4:.0f}y" for t in year_ticks])

plt.tight_layout()
plt.show()
```

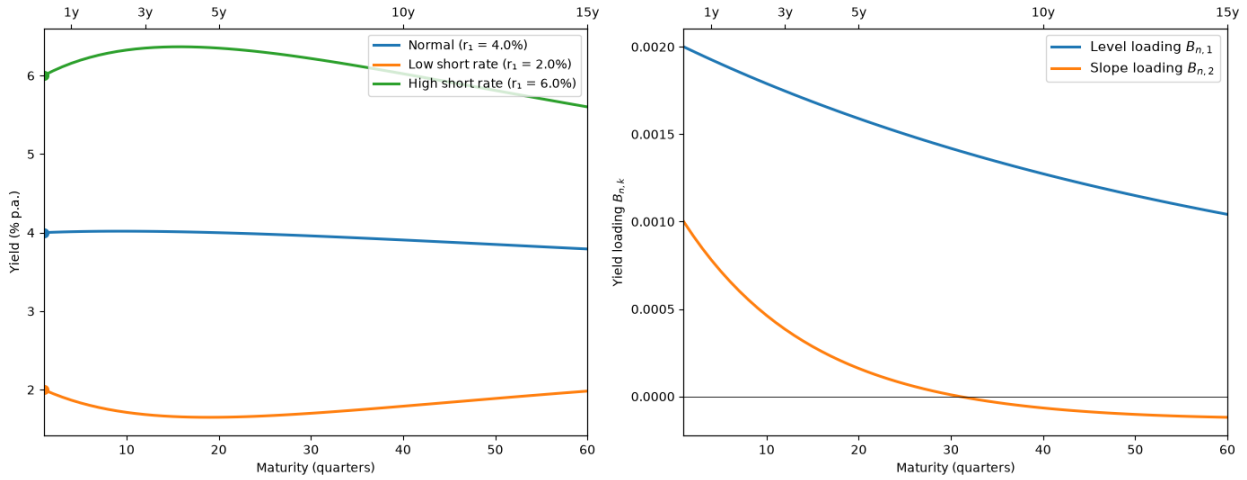


Fig. 139.3: Yield curves and factor loadings under the two-factor model

The right panel shows how yields at different maturities respond to the two factors.

The level loading $B_{n,1}$ declines gradually with maturity, from 0.002 at one quarter to about half that at 15 years, so the persistent factor moves yields at all maturities.

The slope loading $B_{n,2}$ starts at 0.001, decays quickly, and turns slightly negative beyond about 30 quarters.

The sign change reflects the off-diagonal entry $\phi_{12} = -0.03$: a high z_{2t} raises the short rate today but pushes the level factor, and hence future short rates, down.

As a result, the level factor dominates yields at long maturities, while the slope factor matters mainly at the short end.

As in the one-factor case, the curves in the left panel are still approaching their common long-run yield at 60 quarters.

139.5 Risk premiums

A key object in the affine term structure model is the **term premium**, the expected excess return on a long-term bond relative to rolling over short-term bonds.

For an $(n+1)$-period bond held for one period, the shock loading is $\alpha_n = C^\top \bar{B}_n$, so (139.7) gives

$$\log \mathbb{E}_t R_{t+1}^{(n+1)} - r_t \;=\; \bar{B}_n^\top C \lambda_t$$

The term premium equals the inner product of the bond's shock exposure $\bar{B}_n^\top C$ with the risk price vector λ_t.

This formula measures the premium as the log of an expected gross return.

The expected log excess return is smaller by a Jensen's inequality term:

$$\mathbb{E}_t \log R_{t+1}^{(n+1)} - r_t = \bar{B}_n^\top C \lambda_t - \tfrac{1}{2} \bar{B}_n^\top C C^\top \bar{B}_n$$

Only the first term depends on the state, so the two measures differ by a maturity-specific constant that does not affect how term premiums vary over time.

Because the term premium equals $\bar{B}_n^\top C \lambda_t$, its sign depends on the *current* risk-price vector λ_t, which is state-dependent whenever $\lambda_z \neq 0$.

To see this more concretely, consider a state where $C\lambda_t$ is negative componentwise (for example, $z_t = 0$ in our calibration below).

When $\delta_1 > 0$, a positive shock ε_{t+1} raises the short rate and lowers long-bond prices, so the bond shock loading $\alpha_n = C^\top \bar{B}_n$ is negative.

A negative $C\lambda_t$ then means the stochastic discount factor m_{t+1} loads positively on ε_{t+1}, i.e. the SDF is high in states where interest rates rise and bond prices fall.

This makes $\mathrm{Cov}_t(m_{t+1}, R_{t+1}^{(n+1)}) < 0$, so long bonds are risky and carry a positive term premium.

Algebraically, $\bar{B}_n < 0$ and $C\lambda_t < 0$ combine to give $\bar{B}_n^\top C \lambda_t > 0$.

In other states, however, components of λ_t can change sign.

With $\lambda_z < 0$, a low value of the level factor z_{1t} pushes the first component of λ_t above zero.

Since long bonds load mainly on the level shock, their term premiums then turn negative.

Exercise 139.5.1

Derive the term premium formula above by computing the one-period holding return on an $(n+1)$-period bond and identifying its shock loading.

Hint: Use $R_{t+1}^{(n+1)} = p_{t+1}(n)/p_t(n+1)$ with $\log p_t(n) = \bar{A}_n + \bar{B}_n^\top z_t$, substitute the state dynamics (139.1), and apply the Riccati equations (139.10)–(139.11) to simplify.

Solution

The one-period holding return on an $(n+1)$-period bond is $R_{t+1}^{(n+1)} = p_{t+1}(n)/p_t(n+1)$, so

$$\log R_{t+1}^{(n+1)} = \bar{A}_n + \bar{B}_n^\top z_{t+1} - \bar{A}_{n+1} - \bar{B}_{n+1}^\top z_t$$

Substituting $z_{t+1} = \mu + \phi z_t + C\varepsilon_{t+1}$ from (139.1):

$$= \underbrace{(\bar{A}_n + \bar{B}_n^\top \mu - \bar{A}_{n+1})}_{\text{constant}} + \underbrace{(\bar{B}_n^\top \phi - \bar{B}_{n+1}^\top)}_{\text{loading on } z_t} z_t + \underbrace{\bar{B}_n^\top C}_{\text{shock loading}} \varepsilon_{t+1}$$

We now use the Riccati equations to simplify each piece.

For the constant piece, (139.10) gives $\bar{A}_{n+1} = \bar{A}_n + \bar{B}_n^\top(\mu - C\lambda_0) + \frac{1}{2}\bar{B}_n^\top CC^\top \bar{B}_n - \delta_0$, so

$$\bar{A}_n + \bar{B}_n^\top \mu - \bar{A}_{n+1} = \bar{B}_n^\top C\lambda_0 - \tfrac{1}{2}\bar{B}_n^\top CC^\top \bar{B}_n + \delta_0$$

For the z_t coefficient, (139.11) gives $\bar{B}_{n+1}^\top = \bar{B}_n^\top(\phi - C\lambda_z) - \delta_1^\top$, so

$$\bar{B}_n^\top \phi - \bar{B}_{n+1}^\top = \bar{B}_n^\top C\lambda_z + \delta_1^\top$$

Combining the pieces:

$$\log R_{t+1}^{(n+1)} = \underbrace{(\delta_0 + \delta_1^\top z_t)}_{r_t} + \bar{B}_n^\top C \underbrace{(\lambda_0 + \lambda_z z_t)}_{\lambda_t} - \tfrac{1}{2}\bar{B}_n^\top CC^\top \bar{B}_n + \bar{B}_n^\top C\,\varepsilon_{t+1}$$

Writing $\alpha_n = C^\top \bar{B}_n$, this takes the generic return form (139.5):

$$\log R_{t+1}^{(n+1)} = \underbrace{(r_t + \alpha_n^\top \lambda_t)}_{\nu_t} - \tfrac{1}{2}\alpha_n^\top \alpha_n + \alpha_n^\top \varepsilon_{t+1}$$

Since $\mathbb{E}_t R_{t+1}^{(n+1)} = \exp(\nu_t)$, we obtain

$$\log \mathbb{E}_t R_{t+1}^{(n+1)} - r_t = \alpha_n^\top \lambda_t = \bar{B}_n^\top C\lambda_t$$

The following figure plots term premiums across maturities for different states of the world, as well as the level and slope factor contributions to the term premium in the normal state

```
def term_premiums(model, z, n_max):
    """
    Compute one-period term premiums on bonds of maturity 1, ..., n_max.

    An n-period bond held for one period becomes an (n-1)-period bond,
    so its premium is B_bar_{n-1}^⊤ C λ_t (and zero for n = 1).
    """
    A_bar, B_bar = bond_coefficients(model, n_max + 1)
    λ_t = risk_prices(model, z)
    return np.array([B_bar[n-1] @ model.C @ λ_t
                     for n in range(1, n_max + 1)])

n_max_tp = 60
maturities_tp = np.arange(1, n_max_tp + 1)

z_states_tp = {
    "Low rate ($z_1 < 0$)":  np.array([-3.0, 2.0]),
    "High rate ($z_1 > 0$)": np.array([3.0, -2.0]),
}

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(14, 5.5))

for label, z in z_states_tp.items():
```

(continues on next page)

(continued from previous page)

```
    tp = term_premiums(model_2f, z, n_max_tp) * 4 * 100
    r_now = short_rate(model_2f, z) * 4 * 100
    lam = risk_prices(model_2f, z)
    ax1.plot(maturities_tp, tp, lw=2.2,
             label=(f"{label}\n  r={r_now:.1f}%,"
                    f" λ=[{lam[0]:.3f}, {lam[1]:.3f}]"))

ax1.axhline(0, color="black", lw=0.8, ls="--")
ax1.set_xlabel("Maturity (quarters)")
ax1.set_ylabel("Term premium (% p.a.)")
ax1.legend(fontsize=9)
ax1.set_xlim(1, n_max_tp)

z_decomp = np.array([0.0, 0.0])
A_bar_d, B_bar_d = bond_coefficients(model_2f, n_max_tp + 1)
λ_t = risk_prices(model_2f, z_decomp)
C_lam = model_2f.C @ λ_t

tp_level = np.array([B_bar_d[n-1, 0] * C_lam[0]
                     for n in range(1, n_max_tp + 1)]) * 4 * 100
tp_slope = np.array([B_bar_d[n-1, 1] * C_lam[1]
                     for n in range(1, n_max_tp + 1)]) * 4 * 100
tp_total = tp_level + tp_slope

ax2.plot(maturities_tp, tp_total, 'k-', lw=2.2, label="Total")
ax2.plot(maturities_tp, tp_level, lw=1.8, ls="--",
         label="Level factor")
ax2.plot(maturities_tp, tp_slope, lw=1.8, ls="--",
         label="Slope factor")
ax2.axhline(0, color="black", lw=0.6, ls=":")
ax2.set_xlabel("Maturity (quarters)")
ax2.set_ylabel("Term premium (% p.a.)")
ax2.legend(fontsize=10)
ax2.set_xlim(1, n_max_tp)

for ax in (ax1, ax2):
    ax_top = ax.twiny()
    ax_top.set_xlim(ax.get_xlim())
    year_ticks = [4, 12, 20, 40, 60]
    ax_top.set_xticks(year_ticks)
    ax_top.set_xticklabels([f"{t/4:.0f}y" for t in year_ticks])

plt.tight_layout()
plt.show()
```

The left panel shows that the sign of the term premium depends on the state.

In the high-rate state, $\lambda_t \approx (-0.025, 0.001)$, and the term premium is positive and rises with maturity, reaching about 0.6% per year at 15 years.

In the low-rate state, the first component of λ_t has turned positive, $\lambda_t \approx (0.005, -0.011)$, and the term premium is negative at every maturity, falling to about -0.15% per year at 15 years.

Investors in the low-rate state accept a lower expected return on long bonds than on rolling over short bonds because long bonds pay off well in the states that they value most.

The right panel decomposes the term premium at $z_t = 0$, where both components of λ_t are negative.

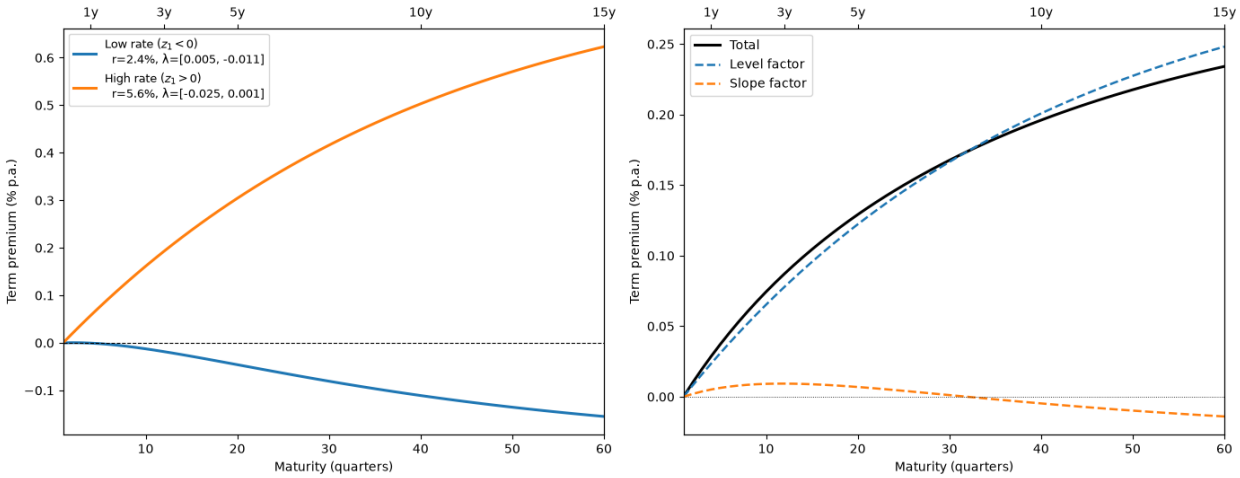


Fig. 139.4: Term premiums and factor decomposition under the two-factor model

The level factor accounts for almost all of the premium at long maturities.

The slope contribution is small and turns negative beyond about 30 quarters, mirroring the sign change in the slope loading shown earlier.

Exercise 139.5.2

The expectations hypothesis says that expected excess returns on long bonds are constant over time, so that they cannot be forecast by the yield spread.

1. Simulate the two-factor model `model_2f` for $T = 200{,}000$ quarters. Regress the one-quarter excess holding return on a 20-quarter bond, $\log p_{t+1}(19) - \log p_t(20) - r_t$, on a constant and the spread $y_t(20) - r_t$.
2. Compute the population regression slope implied by the model, using $\bar{B}_n$ and the stationary covariance matrix Σ_z of z_t, and compare it with your estimate.
3. Repeat both steps with $\lambda_z = 0$ and explain the result.
4. In U.S. data, Fama and Bliss [1987] and Campbell and Shiller [1991] find that high spreads forecast *high* excess returns on long bonds. Does `model_2f` reproduce this pattern? What happens if you flip the sign of λ_z?

Solution

From (139.7) and the term-premium formula, the log excess holding return on a 20-quarter bond is

$$x_{t+1} = \bar{B}_{19}^\top C \lambda_t - \tfrac{1}{2} \bar{B}_{19}^\top C C^\top \bar{B}_{19} + \bar{B}_{19}^\top C \varepsilon_{t+1}$$

Its conditional mean is a constant plus $a^\top z_t$ with $a = \lambda_z^\top C^\top \bar{B}_{19}$.

The spread is $s_t = y_t(20) - r_t = \text{constant} + b^\top z_t$ with $b = -\bar{B}_{20}/20 - \delta_1$.

Because ε_{t+1} is orthogonal to z_t, the population slope is

$$\beta = \frac{a^\top \Sigma_z b}{b^\top \Sigma_z b}$$

where Σ_z solves $\Sigma_z = \phi \Sigma_z \phi^\top + C C^\top$.

```
def eh_regression(model, n=20, T=200_000, seed=1):
    """OLS and population slopes of excess returns on the yield spread."""
    A_bar, B_bar = bond_coefficients(model, n)
    z_bar = np.linalg.solve(np.eye(model.m) - model.ψ, model.μ)
    Z = simulate(model, z_bar, T, rng=np.random.default_rng(seed))

    r = model.δ_0 + Z[:-1] @ model.δ_1
    log_p = lambda k, z: A_bar[k] + z @ B_bar[k]
    excess = log_p(n - 1, Z[1:]) - log_p(n, Z[:-1]) - r
    spread = -log_p(n, Z[:-1]) / n - r

    X = np.column_stack([np.ones(T), spread])
    coef = np.linalg.lstsq(X, excess, rcond=None)[0]
    resid = excess - X @ coef
    se = np.sqrt(resid.var() / (T * spread.var()))

    Σ_z = solve_discrete_lyapunov(model.ψ, model.C @ model.C.T)
    a = model.λ_z.T @ model.C.T @ B_bar[n - 1]
    b = -B_bar[n] / n - model.δ_1
    β_pop = (a @ Σ_z @ b) / (b @ Σ_z @ b)
    return coef[1], se, β_pop

cases = {
    "λ_z as calibrated": λ_z_2,
    "λ_z = 0":           np.zeros((2, 2)),
    "λ_z sign flipped":  -λ_z_2,
}
print(f"{'case':>20}  {'OLS slope':>10}  {'s.e.':>7}  {'population':>10}")
for label, lz in cases.items():
    mod = create_affine_model(μ_2, ψ_2, C_2, δ_0_2, δ_1_2, λ_0_2, lz)
    b_ols, se, b_pop = eh_regression(mod)
    print(f"{label:>20}  {b_ols:>10.4f}  {se:>7.4f}  {b_pop:>10.4f}")
```

```
                case   OLS slope      s.e.  population
   λ_z as calibrated     -0.1614    0.0280     -0.1723
             λ_z = 0      0.0112    0.0250      0.0000
    λ_z sign flipped      0.1728    0.0224      0.1618
```

With the calibrated λ_z, the OLS slope is close to the population slope of about -0.17 and is many standard errors from zero, so the expectations hypothesis fails in simulated data.

With $\lambda_z = 0$, the vector a is zero, so expected excess returns are constant and the population slope is exactly zero.

The OLS estimate is then within sampling error of zero.

The expectations hypothesis holds in an affine model exactly when risk prices do not vary with the state.

The calibrated model gets the *sign* of the Fama–Bliss and Campbell–Shiller finding wrong: a high spread forecasts a low excess return.

Flipping the sign of λ_z reverses the sign of the slope, which shows that the regression evidence is informative about λ_z and not just about the size of average term premiums.

139.6 Risk-neutral probabilities

We return to the VAR and short-rate equations (139.1)–(139.2), which for convenience we repeat here:

$$z_{t+1} = \mu + \phi z_t + C\varepsilon_{t+1}, \qquad r_t = \delta_0 + \delta_1^\top z_t$$

where $\varepsilon_{t+1} \sim \mathcal{N}(0, I)$.

We suppose that this structure describes the data-generating mechanism.

Finance economists call this the **physical measure** P, to distinguish it from the **risk-neutral measure** Q that we now describe.

Under the physical measure, the conditional distribution of z_{t+1} given z_t is $\mathcal{N}(\mu + \phi z_t,\ CC^\top)$.

139.6.1 Change of measure

With the risk-price vector $\lambda_t = \lambda_0 + \lambda_z z_t$ from (139.3), define the non-negative random variable

$$\frac{\xi_{t+1}^Q}{\xi_t^Q} = \exp\!\left(-\tfrac{1}{2}\lambda_t^\top \lambda_t - \lambda_t^\top \varepsilon_{t+1}\right) \tag{139.14}$$

This is a log-normal random variable with mean 1, so it is a valid likelihood ratio that can be used to twist the conditional distribution of z_{t+1}.

To see what the twist does to the shocks, multiply the standard normal density of ε_{t+1} by (139.14):

$$(2\pi)^{-m/2}\exp\!\left(-\tfrac{1}{2}\varepsilon^\top\varepsilon\right)\exp\!\left(-\tfrac{1}{2}\lambda_t^\top\lambda_t - \lambda_t^\top\varepsilon\right) = (2\pi)^{-m/2}\exp\!\left(-\tfrac{1}{2}(\varepsilon+\lambda_t)^\top(\varepsilon+\lambda_t)\right)$$

So under Q the shock is $\varepsilon_{t+1} \sim \mathcal{N}(-\lambda_t, I)$: the twist shifts its mean by $-\lambda_t$ and leaves its covariance matrix unchanged.

Substituting $\varepsilon_{t+1} = -\lambda_t + \varepsilon_{t+1}^Q$ into (139.1) shows that multiplying the physical conditional distribution by this likelihood ratio transforms it into the **risk-neutral conditional distribution**

$$z_{t+1} \mid z_t \overset{Q}{\sim} \mathcal{N}(\mu - C\lambda_0 + (\phi - C\lambda_z)z_t,\ CC^\top)$$

In other words, under Q the state follows

$$z_{t+1} = (\mu - C\lambda_0) + (\phi - C\lambda_z)\, z_t + C\varepsilon_{t+1}^Q$$

where $\varepsilon_{t+1}^Q \sim \mathcal{N}(0, I)$ under Q.

The risk-neutral distribution twists the conditional mean from $\mu + \phi z_t$ to $\mu - C\lambda_0 + (\phi - C\lambda_z)z_t$.

The adjustments $-C\lambda_0$ (constant) and $-C\lambda_z$ (state-dependent) encode how the pricing equation $\mathbb{E}_t^P m_{t+1} R_{j,t+1} = 1$ adjusts expected returns for exposure to the risks ε_{t+1}.

Exercise 139.6.1

How different are the risk-neutral and physical measures?

A natural measure is relative entropy (Kullback–Leibler divergence), discussed in *Statistical Divergence Measures*.

1. Show that the conditional relative entropy of Q with respect to P, $\mathbb{E}_t^Q\left[\log(\xi_{t+1}^Q/\xi_t^Q)\right]$, and the conditional relative entropy of P with respect to Q, $\mathbb{E}_t^P\left[-\log(\xi_{t+1}^Q/\xi_t^Q)\right]$, both equal $\frac{1}{2}\lambda_t^\top\lambda_t$.

2. Suppose that z_0 is drawn from the stationary distribution of z_t under P. Show that the relative entropy of P with respect to Q for a sample $z_1, \ldots, z_T$ equals T times $\$\frac{1}{2}\left(\bar{\lambda}^\top \bar{\lambda} + \text{tr}(\lambda_z \Sigma_z \lambda_z^\top)\right) where \bar z = (I - \phi)^{-1}\mu, \bar\lambda = \lambda_0 + \lambda_z \bar z, and \Sigma_z solves \Sigma_z = \phi\Sigma_z\phi^\top + CC^\top\$$.
3. Evaluate this formula for `model_2f`, verify it by simulation, and use Pinsker's inequality to bound how well any test based on 100 years of quarterly data on z_t could distinguish P from Q.

Solution

Part 1. We showed above that $\varepsilon_{t+1} \sim \mathcal{N}(-\lambda_t, I)$ under Q, while $\varepsilon_{t+1} \sim \mathcal{N}(0, I)$ under P.

Hence

$$\mathbb{E}_t^Q \left[-\tfrac{1}{2}\lambda_t^\top \lambda_t - \lambda_t^\top \varepsilon_{t+1}\right] = -\tfrac{1}{2}\lambda_t^\top \lambda_t + \lambda_t^\top \lambda_t = \tfrac{1}{2}\lambda_t^\top \lambda_t$$

and

$$\mathbb{E}_t^P \left[\tfrac{1}{2}\lambda_t^\top \lambda_t + \lambda_t^\top \varepsilon_{t+1}\right] = \tfrac{1}{2}\lambda_t^\top \lambda_t$$

The two divergences coincide because the two conditional distributions are normal with the same covariance matrix.

Part 2. When C is invertible, the path $z_1, \ldots, z_T$ and the shocks $\varepsilon_1, \ldots, \varepsilon_T$ determine each other given z_0.

The log likelihood ratio of the path is therefore the sum of the one-period log likelihood ratios, and by the law of iterated expectations its expectation under P is $\sum_{t=0}^{T-1} \mathbb{E}^P \left[\frac{1}{2}\lambda_t^\top \lambda_t\right]$.

Under the stationary distribution, $\lambda_t = \bar\lambda + \lambda_z(z_t - \bar z)$ with $\mathbb{E}(z_t - \bar z) = 0$ and $\text{Var}(z_t) = \Sigma_z$.

Therefore $\mathbb{E}(\lambda_t^\top \lambda_t) = \bar\lambda^\top \bar\lambda + \text{tr}(\lambda_z \Sigma_z \lambda_z^\top)$ for every t, which gives the formula.

Part 3. Pinsker's inequality says that the total variation distance between two distributions is at most $\sqrt{D/2}$, where D is their relative entropy.

With equal prior probabilities on P and Q, the smallest achievable probability of choosing the wrong model is $(1 - \text{TV})/2$, so it is at least $(1 - \sqrt{D/2})/2$.

```
z_bar_2 = np.linalg.solve(np.eye(2) - model_2f.φ, model_2f.μ)
Σ_z_2 = solve_discrete_lyapunov(model_2f.φ, model_2f.C @ model_2f.C.T)
λ_bar_2 = risk_prices(model_2f, z_bar_2)
entropy = 0.5 * (λ_bar_2 @ λ_bar_2
                 + np.trace(model_2f.λ_z @ Σ_z_2 @ model_2f.λ_z.T))

Z_sim = simulate(model_2f, z_bar_2, 200_000, rng=np.random.default_rng(7))
Λ_sim = model_2f.λ_0 + Z_sim @ model_2f.λ_z.T
entropy_sim = 0.5 * np.mean(np.sum(Λ_sim**2, axis=1))

T_years = 100
D = 4 * T_years * entropy
print(f"relative entropy per quarter (formula):    {entropy:.6f}")
print(f"relative entropy per quarter (simulation): {entropy_sim:.6f}")
print(f"relative entropy over {T_years} years:           {D:.4f}")
print(f"lower bound on error probability:          {(1 - np.sqrt(D / 2)) / 2:.3f}")
```

```
relative entropy per quarter (formula):    0.000312
relative entropy per quarter (simulation): 0.000310
relative entropy over 100 years:           0.1250
lower bound on error probability:          0.375
```

Even with a century of quarterly data, no test can push the average probability of choosing the wrong model below about 37%.

Yet the gap between these two hard-to-distinguish measures is exactly what generates the term premiums plotted above.

Time-series data on the state alone therefore say little about risk prices, which is why estimates of λ_0 and λ_z rely heavily on the cross-section of bond yields.

This echoes the detection-error calculations of Doubts or Variability?, where plausible amounts of model uncertainty are calibrated in the same way.

139.6.2 Asset pricing in a nutshell

Let $\mathbb{E}^P$ denote an expectation under the physical measure that nature uses to generate the data.

Our key asset pricing equation is $\mathbb{E}_t^P m_{t+1} R_{j,t+1} = 1$ for all returns $R_{j,t+1}$.

Using (139.14), we can express the SDF (139.4) as

$$m_{t+1} = \frac{\xi_{t+1}^Q}{\xi_t^Q} \exp(-r_t)$$

Then the condition $\mathbb{E}_t^P(\exp(-r_t)\,\frac{\xi_{t+1}^Q}{\xi_t^Q}\,R_{j,t+1}) = 1$ is equivalent to

$$\mathbb{E}_t^Q R_{j,t+1} = \exp(r_t) \tag{139.15}$$

Under the risk-neutral measure, expected returns on all assets equal the risk-free return.

139.6.3 Verification via risk-neutral pricing

Bond prices can be computed by discounting at r_t under Q:

$$p_t(n) = \mathbb{E}_t^Q \left[\exp\left(-\sum_{s=0}^{n-1} r_{t+s} \right) \right]$$

We can verify that this agrees with (139.9) by iterating the affine recursion under the risk-neutral VAR.

Below we confirm this numerically

```
def bond_price_mc_Q(model, z0, n, n_sims=50_000, rng=None):
    """Estimate p_t(n) by Monte Carlo under Q; return estimate and std. error."""
    if rng is None:
        rng = np.random.default_rng(0)
    m = len(z0)
    Z = np.tile(z0, (n_sims, 1))
    disc = np.zeros(n_sims)
    for _ in range(n):
        disc += model.δ_0 + Z @ model.δ_1
        ε = rng.standard_normal((n_sims, m))
        Z = model.μ_rn + Z @ model.ψ_rn.T + ε @ model.C.T
    payoffs = np.exp(-disc)
    return payoffs.mean(), payoffs.std() / np.sqrt(n_sims)
```

(continues on next page)

(continued from previous page)

```
z_test = np.array([0.01, 0.005])
p_analytic = bond_prices(model_2f, z_test, 40)

rng = np.random.default_rng(0)
maturities_check = [4, 12, 24, 40]
mc_results = [bond_price_mc_Q(model_2f, z_test, n, n_sims=100_000, rng=rng)
              for n in maturities_check]

header = (f"{'Maturity':>10}  {'Analytic':>10}  {'Monte Carlo':>11}"
          f"  {'Error (bps)':>11}  {'s.e. (bps)':>10}  {'z-score':>8}")
print(header)
print("-" * len(header))
for n, (mc, se) in zip(maturities_check, mc_results):
    analytic = p_analytic[n - 1]
    error_bp = (mc - analytic) / analytic * 10_000
    se_bp = se / analytic * 10_000
    print(f"{n:>10}  {analytic:>10.6f}  {mc:>11.6f}"
          f"  {error_bp:>11.2f}  {se_bp:>10.2f}  {error_bp / se_bp:>8.2f}")
```

```
  Maturity    Analytic  Monte Carlo  Error (bps)  s.e. (bps)   z-score
------------------------------------------------------------------------
         4    0.960592     0.960571        -0.23        0.25     -0.89
        12    0.886276     0.886200        -0.86        1.35     -0.64
        24    0.787066     0.787580         6.53        3.47      1.88
        40    0.676314     0.676398         1.24        6.61      0.19
```

The table reports each Monte Carlo error alongside its standard error, both in basis points of the analytical price.

All the z-scores are within ± 2, so the differences between analytical and simulated prices are consistent with pure simulation noise.

This validates the Riccati recursion (139.10)–(139.11).

139.7 Distorted beliefs

Piazzesi *et al.* [2015] assemble survey evidence suggesting that economic experts' forecasts are systematically biased relative to the physical measure.

139.7.1 The subjective measure

Let $\{z_t\}_{t=1}^T$ be a record of observations on the state and let $\{\check{z}_{t+1}\}_{t=1}^T$ be a record of one-period-ahead expert forecasts.

Let $\check{\mu}, \check{\phi}$ be the regression coefficients in

$$\check{z}_{t+1} = \check{\mu} + \check{\phi}\, z_t + e_{t+1}$$

where the residual e_{t+1} has mean zero, is orthogonal to z_t, and satisfies $\mathbb{E}\, e_{t+1} e_{t+1}^\top = CC^\top$.

By comparing estimates of μ, ϕ from (139.1) with estimates of $\check{\mu}, \check{\phi}$ from the experts' forecasts, Piazzesi *et al.* [2015] deduce that the experts' beliefs are systematically distorted.

To organize this evidence, let $\kappa_t = \kappa_0 + \kappa_z z_t$ and define the likelihood ratio

$$\frac{\xi_{t+1}^S}{\xi_t^S} = \exp\left(-\tfrac{1}{2}\kappa_t^\top \kappa_t - \kappa_t^\top \varepsilon_{t+1}\right) \tag{139.16}$$

This is log-normal with mean 1, so it is a valid likelihood ratio.

Multiplying the physical conditional distribution of z_{t+1} by this likelihood ratio transforms it to the experts' **subjective conditional distribution**

$$z_{t+1} \mid z_t \overset{S}{\sim} \mathcal{N}(\mu - C\kappa_0 + (\phi - C\kappa_z)\, z_t,\; CC^\top)$$

In the experts' forecast regression, $\breve{\mu}$ estimates $\mu - C\kappa_0$ and $\breve{\phi}$ estimates $\phi - C\kappa_z$.

Piazzesi *et al.* [2015] find that the experts behave as if the level and slope of the yield curve are more persistent than under the physical measure: $\breve{\phi}$ has larger eigenvalues than ϕ.

139.7.2 Pricing under distorted beliefs

Piazzesi *et al.* [2015] explore the hypothesis that a representative agent with these distorted beliefs prices assets and makes returns satisfy

$$\mathbb{E}_t^S(m_{t+1}^\star\, R_{j,t+1}) = 1$$

where $\mathbb{E}_t^S$ is the conditional expectation under the subjective S measure and $m_{t+1}^\star$ is the SDF of an agent with these beliefs.

To express the agent's SDF, define the **subjective shocks**

$$\varepsilon_{t+1}^S = \varepsilon_{t+1} + \kappa_t$$

By the same argument we used for the risk-neutral measure, $\varepsilon_{t+1}^S \sim \mathcal{N}(0, I)$ under S, and the state evolves as

$$z_{t+1} = (\mu - C\kappa_0) + (\phi - C\kappa_z)z_t + C\varepsilon_{t+1}^S$$

The agent's SDF is exponential quadratic in these subjective shocks:

$$m_{t+1}^\star = \exp\!\big(-r_t - \tfrac{1}{2}\lambda_t^{\star\top}\lambda_t^\star - \lambda_t^{\star\top}\varepsilon_{t+1}^S\big)$$

where $\lambda_t^\star = \lambda_0^\star + \lambda_z^\star z_t$ is the agent's vector of risk prices.

Because $\mathbb{E}_t^S m_{t+1}^\star = \exp(-r_t)$, the short rate in this SDF is the market short rate r_t that we observe.

Using (139.16) to convert to the physical measure, the subjective pricing equation becomes

$$\mathbb{E}_t^P\big[\exp\!\big(-r_t - \tfrac{1}{2}\lambda_t^{\star\top}\lambda_t^\star - \lambda_t^{\star\top}(\varepsilon_{t+1} + \kappa_t)\big)\exp\!\big(-\tfrac{1}{2}\kappa_t^\top\kappa_t - \kappa_t^\top\varepsilon_{t+1}\big)\, R_{j,t+1}\big] = 1$$

The constant terms in the combined exponent are $-r_t - \frac{1}{2}\lambda_t^{\star\top}\lambda_t^\star - \lambda_t^{\star\top}\kappa_t - \frac{1}{2}\kappa_t^\top\kappa_t = -r_t - \frac{1}{2}(\lambda_t^\star + \kappa_t)^\top(\lambda_t^\star + \kappa_t)$, so the two exponentials combine exactly into

$$\mathbb{E}_t^P\big[\exp\!\big(-r_t - \tfrac{1}{2}(\lambda_t^\star + \kappa_t)^\top(\lambda_t^\star + \kappa_t) - (\lambda_t^\star + \kappa_t)^\top\varepsilon_{t+1}\big)\, R_{j,t+1}\big] = 1$$

Comparing this with the rational-expectations econometrician's pricing equation

$$\mathbb{E}_t^P\big[\exp\!\big(-r_t - \tfrac{1}{2}\lambda_t^\top\lambda_t - \lambda_t^\top\varepsilon_{t+1}\big)\, R_{j,t+1}\big] = 1$$

we see that what the econometrician interprets as λ_t is actually

$$\widehat{\lambda}_t = \lambda_t^\star + \kappa_t$$

139.7.3 What bond prices can and cannot reveal

The decomposition has a simple interpretation in terms of the risk-neutral measure.

Starting from the physical measure, the econometrician reaches Q by twisting with $\hat{\lambda}_t$:

$$\mu - C\hat{\lambda}_0 = (\mu - C\kappa_0) - C\lambda_0^\star, \qquad \phi - C\hat{\lambda}_z = (\phi - C\kappa_z) - C\lambda_z^\star$$

The right sides show that the agent reaches the *same* Q by twisting the subjective measure with $\lambda_t^\star$.

Since bond prices depend only on Q, the econometrician and the agent agree about every bond price.

They disagree about expected returns.

The term premium measured by the econometrician is $\bar{B}_n^\top C\hat{\lambda}_t$, while the premium that the agent expects is $\bar{B}_n^\top C\lambda_t^\star$.

The difference, $\bar{B}_n^\top C\kappa_t$, is the part of measured excess returns that the agent does not expect and that therefore shows up as predictable forecast errors.

Bond prices and data on z_t identify $\hat{\lambda}_t$, but they cannot split it into $\lambda_t^\star$ and κ_t.

Splitting it requires direct evidence on beliefs, such as the survey forecasts used by Piazzesi *et al.* [2015].

The same identification problem is central to *The Recovery Theorem* and *Misspecified Recovery*, which ask when risk-neutral prices alone can reveal subjective beliefs.

139.7.4 A numerical illustration

We keep the same physical state dynamics and short-rate specification as above, choose the agent's risk prices $\lambda_t^\star$, and deduce the econometrician's risk prices $\hat{\lambda}_t = \lambda_t^\star + \kappa_t$.

We set the subjective parameters $\check{\mu}, \check{\phi}$ to match the evidence in Piazzesi *et al.* [2015] that experts behave as if the level and slope of the yield curve are more persistent than under the physical measure.

In particular, we use

$$\check{\phi} = \begin{pmatrix} 0.985 & -0.025 \\ 0.00 & 0.94 \end{pmatrix}$$

```
ψ_P = ψ_2.copy()
μ_P = μ_2.copy()

# Subjective parameters: experts believe factors are MORE persistent
ψ_S = np.array([[0.985, -0.025], [0.00, 0.94]])
μ_S = np.array([0.005, 0.0])

# κ_t = κ_0 + κ_z z_t twists P into S
κ_z = np.linalg.solve(C_2, ψ_P - ψ_S)
κ_0 = np.linalg.solve(C_2, μ_P - μ_S)

# Agent's risk prices, which price subjective shocks ε^S
λ_star_0 = np.array([-0.03, -0.015])
λ_star_z = np.array([[-0.006, 0.0], [0.0, -0.004]])

# Econometrician's risk prices, which price physical shocks ε
λ_hat_0 = λ_star_0 + κ_0
λ_hat_z = λ_star_z + κ_z
```

The agent's model pairs the subjective dynamics with $\lambda_t^\star$, and the econometrician's model pairs the physical dynamics with $\widehat{\lambda}_t$.

We first confirm that the two models imply the same risk-neutral dynamics and hence the same bond prices.

```
# Agent: subjective dynamics, risk prices λ*
model_subj = create_affine_model(
    μ_S, ψ_S, C_2, δ_0_2, δ_1_2, λ_star_0, λ_star_z)
# Econometrician: physical dynamics, risk prices λ̂ = λ* + κ
model_econ = create_affine_model(
    μ_P, ψ_P, C_2, δ_0_2, δ_1_2, λ_hat_0, λ_hat_z)

print("Same risk-neutral dynamics:",
      np.allclose(model_subj.ψ_rn, model_econ.ψ_rn)
      and np.allclose(model_subj.μ_rn, model_econ.μ_rn))
print("Same yields at z = (1, -1):",
      np.allclose(compute_yields(model_subj, np.array([1.0, -1.0]), 60),
                  compute_yields(model_econ, np.array([1.0, -1.0]), 60)))
```

```
Same risk-neutral dynamics: True
Same yields at z = (1, -1): True
```

Now we compare the term premium the agent expects with the term premium the econometrician measures.

```
z_ref = np.array([0.0, 0.0])
n_max_db = 60
maturities_db = np.arange(1, n_max_db + 1)

tp_subj = term_premiums(model_subj, z_ref, n_max_db) * 4 * 100
tp_econ = term_premiums(model_econ, z_ref, n_max_db) * 4 * 100

fig, (ax1, ax2) = plt.subplots(1, 2, figsize=(14, 5.5))

ax1.plot(maturities_db, tp_subj, lw=2.2,
         label=r"Agent's expected premium, $\lambda^\star_t$")
line_econ, = ax1.plot(maturities_db, tp_econ, lw=2.2, ls="--",
         label=(r"Measured by RE econometrician,"
                r" $\hat\lambda_t = \lambda^\star_t + \kappa_t$"))
ax1.fill_between(maturities_db, tp_subj, tp_econ,
                 alpha=0.15, color=line_econ.get_color(),
                 label="Belief distortion component")
ax1.axhline(0, color="black", lw=0.8, ls=":")
ax1.set_xlabel("Maturity (quarters)")
ax1.set_ylabel("Term premium (% p.a.)")
ax1.legend(fontsize=9.5)
ax1.set_xlim(1, n_max_db)

mask = np.abs(tp_subj) > 1e-8
ratio = np.full_like(tp_subj, np.nan)
ratio[mask] = tp_econ[mask] / tp_subj[mask]

ax2.plot(maturities_db[mask], ratio[mask], lw=2.2)
ax2.axhline(1, color="black", lw=0.8, ls="--",
            label="No distortion (ratio = 1)")
ax2.set_xlabel("Maturity (quarters)")
ax2.set_ylabel(r"$\hat{tp}\, /\, tp^\star$")
ax2.legend(fontsize=11)
ax2.set_xlim(1, n_max_db)
```

(continues on next page)

(continued from previous page)

```
for ax in (ax1, ax2):
    ax_top = ax.twiny()
    ax_top.set_xlim(ax.get_xlim())
    year_ticks = [4, 12, 20, 40, 60]
    ax_top.set_xticks(year_ticks)
    ax_top.set_xticklabels([f"{t/4:.0f}y" for t in year_ticks])

plt.tight_layout()
plt.show()

for n in [4, 20, 40, 60]:
    print(f"n = {n:>2}: agent {tp_subj[n-1]:.3f}%, "
          f"econometrician {tp_econ[n-1]:.3f}%, "
          f"ratio {tp_econ[n-1] / tp_subj[n-1]:.2f}")
```

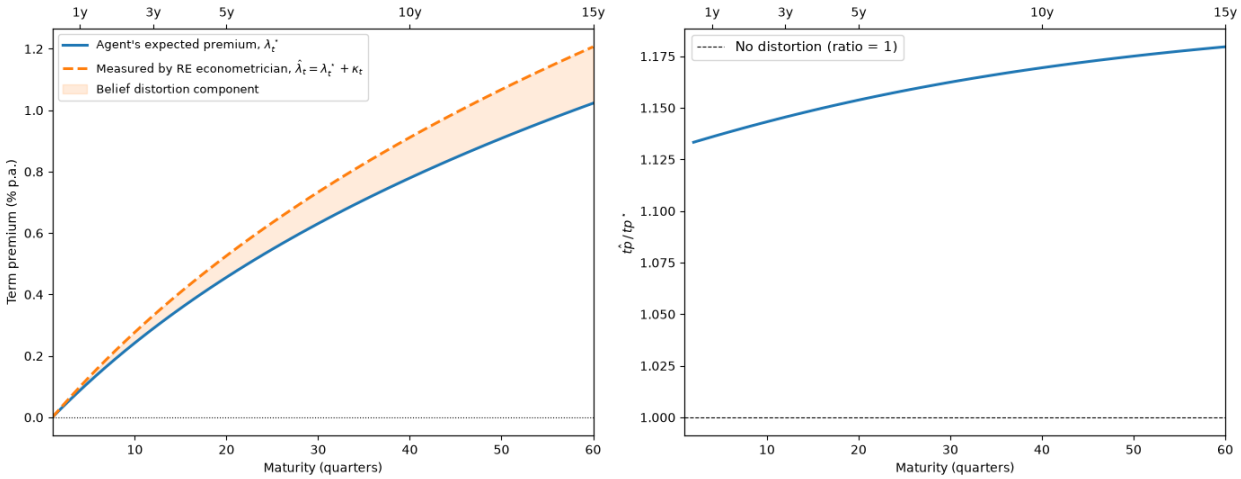


Fig. 139.5: Subjective vs. measured term premiums and overstatement ratio

```
n =  4: agent 0.087%, econometrician 0.099%, ratio 1.14
n = 20: agent 0.456%, econometrician 0.526%, ratio 1.15
n = 40: agent 0.779%, econometrician 0.911%, ratio 1.17
n = 60: agent 1.023%, econometrician 1.207%, ratio 1.18
```

At $z_t = 0$, the distortion is $\kappa_t = \kappa_0 = (-0.005, 0)$, which comes from the experts' upward bias $\check{\mu}_1 = 0.005$ in forecasting the level factor.

Experts who expect higher future interest rates expect lower returns on long bonds than the returns that actually materialize on average.

The econometrician therefore overstates the term premium that the agent requires, by 14 to 18 percent in this calibration.

The persistence distortion κ_z makes the size and even the sign of the gap depend on the state.

```
for label, z in [("High rate", np.array([3.0, -2.0])),
                 ("Low rate",  np.array([-3.0, 2.0]))]:
    gap = (term_premiums(model_econ, z, n_max_db)
           - term_premiums(model_subj, z, n_max_db)) * 4 * 100
    print(f"{label}: belief-distortion gap at 5 and 15 years = "
          f"{gap[19]:.2f}%, {gap[59]:.2f}%")
```

```
High rate: belief-distortion gap at 5 and 15 years = 0.37%, 1.89%
Low rate: belief-distortion gap at 5 and 15 years = -0.23%, -1.53%
```

When rates are high, experts extrapolate them too far into the future, so the econometrician's measured premium exceeds the agent's premium by almost 2 percentage points at 15 years.

When rates are low, the gap reverses sign.

A rational-expectations econometrician therefore attributes to time-varying risk prices movements in measured excess returns that actually come from the agent's systematic forecast errors.

Our Risk Aversion or Mistaken Beliefs? lecture explores this confounding in greater depth.

139.8 Concluding remarks

The affine model of the stochastic discount factor provides a flexible and tractable framework for studying asset prices.

Key features are:

1. **Analytical tractability:** Bond prices are exponential affine in z_t; expected returns decompose cleanly into a short rate plus a risk-price×exposure inner product.
2. **Empirical flexibility:** The free parameters $(\mu, \phi, C, \delta_0, \delta_1, \lambda_0, \lambda_z)$ can be estimated by maximum likelihood (the *Kalman filter* chapter describes the relevant methods) without imposing restrictions from a full general equilibrium model.
3. **Multiple risks:** The vector structure accommodates many sources of risk (monetary policy, real activity, volatility, etc.).
4. **Belief distortions:** The framework naturally accommodates non-rational beliefs via likelihood-ratio twists of the physical measure, as in Piazzesi *et al.* [2015].
5. **Limits to identification:** Bond prices reveal only the risk-neutral measure, so separating risk prices from belief distortions requires additional evidence, such as survey forecasts; see also *The Recovery Theorem* and *Misspecified Recovery*.

The model also connects directly to the Hansen–Jagannathan bounds studied in Doubts or Variability? and to robust control interpretations of the stochastic discount factor described in other chapters of Ljungqvist and Sargent [2018].

For finite-state approaches to asset pricing that complement the continuous-state framework here, see *Asset Pricing: Finite State Models*.

CHAPTER

THE RECOVERY THEOREM

Contents

- *The Recovery Theorem*
  - *Overview*
  - *Model setup*
  - *The recovery theorem*
  - *Numerical example*
  - *Natural vs. risk-neutral distributions*
  - *Effect of risk aversion*
  - *Recovering the discount rate*
  - *Tail risk: natural vs. risk-neutral probabilities of catastrophe*
  - *From option prices to transition prices*
  - *Testing efficient markets*
  - *Limitations and extensions*
  - *Exercises*

140.1 Overview

Asset prices are forward-looking: they encode investors' expectations about future economic states and their valuations of different risks.

A long-standing question in finance is whether one can *recover* the probability distribution used by investors – their subjective beliefs – from observed asset prices alone.

Option prices reveal **state prices**; once these are normalized by the riskless discount factor, the resulting probabilities are the **risk-neutral probabilities** implied by asset prices after risk adjustments have been folded in.

These are not the **natural probabilities** that investors actually assign to future states of the world.

The two differ because risk-neutral probabilities blend together two distinct objects: the market's true beliefs about the future, and investors' aversion to risk.

The link between them is the **pricing kernel**, which reweights natural probabilities to deliver state prices.

Separating beliefs from risk aversion has traditionally required parametric assumptions about the preferences of a representative investor.

Ross [2015] showed otherwise.

Ross's theorem says that, in a finite-state Markov economy, state prices can be enough.

Suppose the Arrow–Debreu state-price transition matrix is arbitrage-free and irreducible.

If the pricing kernel also satisfies a structural restriction called **transition independence**, then state prices uniquely determine both the natural probability transition matrix and the transition pricing kernel.

No historical return data or assumed utility function is needed if some assumptions about the structure of the pricing kernel hold.

This is the **Recovery Theorem**.

It has several important implications:

- It shows how state-price transition data can identify the market's forward-looking natural distribution when the assumption holds.
- It provides tests of the efficient market hypothesis.
- It sheds light on the "dark matter" of finance: the probability of rare catastrophic events embedded in market prices.

This lecture covers

- the Arrow–Debreu framework linking state prices, risk-neutral probabilities, the pricing kernel, and natural probabilities,
- Ross's Recovery Theorem and its proof via the Perron–Frobenius theorem,
- an implementation that recovers the natural distribution from a simulated state-price matrix,
- how option prices and forward equations can be used to estimate transition state prices, and
- comparisons between risk-neutral and recovered natural densities.

Let's import the packages we'll need.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.linalg import eig
from scipy.stats import norm
import matplotlib.cm as cm
```

140.2 Model setup

140.2.1 Arrow–Debreu state prices

Consider a discrete-time, discrete-state economy.

At each date the economy occupies one of m states $\theta_1, \ldots, \theta_m$.

An **Arrow–Debreu security** pays \$1 if the economy is in state θ_j next period and nothing otherwise.

Denote by $p(\theta_i, \theta_j)$ the price today, when the current state is θ_i, of the Arrow–Debreu security paying in state θ_j next period.

Collect these into an $m \times m$ **state price transition matrix**

$$P = [p(\theta_i, \theta_j)]_{i,j=1}^{m}.$$

As in *Competitive Equilibria with Arrow Securities*, the row sums give the state-dependent riskless discount factor: $\sum_j p(\theta_i, \theta_j) = e^{-r(\theta_i)}$.

Here $r(\theta_i)$ is the one-period continuously compounded riskless rate in current state θ_i.

More generally, if an asset pays $g(\theta_j)$ next period, then its price in state θ_i is

$$p_g(\theta_i) = \sum_j p(\theta_i, \theta_j) g(\theta_j).$$

Let

$$b(\theta_i) \equiv \sum_j p(\theta_i, \theta_j) = e^{-r(\theta_i)}$$

be the price of a one-period riskless bond in state θ_i.

Normalizing Arrow prices by this bond price gives the **risk-neutral transition probabilities**

$$q^*(\theta_i, \theta_j) = \frac{p(\theta_i, \theta_j)}{b(\theta_i)} = e^{r(\theta_i)} p(\theta_i, \theta_j).$$

Thus the same asset price can be written as

$$p_g(\theta_i) = b(\theta_i) \sum_j q^*(\theta_i, \theta_j) g(\theta_j) = e^{-r(\theta_i)} E_i^*[g(\theta_{t+1})].$$

Here E_i^* denotes conditional expectation under $q^*(\theta_i, \cdot)$.

The asterisk marks the risk-neutral, or martingale, probability measure.

It is useful to separate this one-period normalization from the dynamic transition structure.

If $Q(\theta_i, \theta_j, T)$ denotes the risk-neutral probability of moving from θ_i to θ_j over T periods, and $0 < t < T$ is an intermediate horizon, then the Markov forward equation is

$$Q(\theta_i, \theta_j, T) = \sum_k Q(\theta_i, \theta_k, t) Q(\theta_k, \theta_j, T - t).$$

In matrix notation, multiperiod risk-neutral transition matrices compose by matrix multiplication.

140.2.2 The pricing kernel

Using the stochastic-discount-factor notation studied in *Asset Pricing: Finite State Models* and the Arrow-security notation used in *Competitive Equilibria with Arrow Securities*, the pricing kernel $\phi(\theta_i, \theta_j)$ relates state prices to natural probabilities via

$$p(\theta_i, \theta_j) = \phi(\theta_i, \theta_j)\, f(\theta_i, \theta_j),$$

where $f(\theta_i, \theta_j)$ is the natural (conditional) probability of transitioning from state θ_i to θ_j.

As in the representative-agent equilibrium calculation in *Competitive Equilibria with Arrow Securities*, the canonical additively separable model with discount factor β gives

$$\phi(\theta_i, \theta_j) = \frac{p(\theta_i, \theta_j)}{f(\theta_i, \theta_j)} = \frac{\beta U'(c(\theta_j))}{U'(c(\theta_i))}. \tag{140.1}$$

This formula has a special structure: the kernel can be written as a ratio of two state-specific terms.

Ross calls this property **transition independence**.

We will say more about it soon.

140.2.3 The identification challenge

Before stating the restriction, it helps to see why one is needed at all.

Given P, any pair (ϕ, f) satisfying $p_{ij} = \phi_{ij} f_{ij}$ for every (i, j) is consistent with observed state prices.

The state-price matrix P supplies m^2 equations.

A natural transition matrix F contributes $m(m-1)$ free entries (rows sum to one), and an arbitrary kernel ϕ contributes another m^2 – a total of $2m^2 - m$ unknowns against only m^2 equations.

The system is under-identified by $m^2 - m$ parameters, so some structural restriction on the kernel is needed to pin down ϕ and f separately.

The transition independence restriction does the job, as we will see in the next section.

140.2.4 Transition independence

Definition (Transition Independence)

A pricing kernel is **transition independent** if there exists a positive function h on the state space and a positive scalar β such that for every transition from state θ_i to θ_j,

$$\phi(\theta_i, \theta_j) = \beta \frac{h(\theta_j)}{h(\theta_i)}.$$

Transition independence says the kernel depends on the *ending* state and normalizes by the *beginning* state.

In the representative-agent complete-markets environment above, it holds under intertemporally additive separable utility (where $h = U'$).

In particular, this holds for (140.1).

Transition independence helps because it ties all m^2 entries of ϕ together: once the m state-specific values are known, the whole kernel is pinned down.

It therefore cuts ϕ from m^2 free entries down to m, so the system becomes exactly identified.

Under transition independence, the state-price equation becomes

$$p(\theta_i, \theta_j) = \beta \frac{h(\theta_j)}{h(\theta_i)} f(\theta_i, \theta_j).$$

In matrix notation, defining the diagonal matrix D with $D_{ii} = h(\theta_i)/\beta$,

$$DP = \beta F D,$$

or equivalently,

$$F = \frac{1}{\beta} D P D^{-1}.$$

140.3 The recovery theorem

140.3.1 Reduction to an eigenvalue problem

Since F is a stochastic matrix, its rows sum to one: $Fe = e$ where e is the vector of ones.

Substituting the expression for F:

$$\frac{1}{\beta} DPD^{-1}e = e \quad \Longrightarrow \quad Pz = \beta z, \quad z \equiv D^{-1}e.$$

This is an **eigenvalue problem** where we seek a positive vector z and scalar β satisfying $Pz = \beta z$.

In principle every eigenvalue-eigenvector pair of P is a formal solution, but only the one with a strictly positive eigenvector is economically valid: $D_{ii} = 1/z_i$ must be positive (so $z_i > 0$), and F must have nonnegative entries.

The Perron–Frobenius theorem guarantees that exactly one such pair exists.

Theorem (Perron–Frobenius)

If A is a nonnegative irreducible matrix, then

1. A has a positive real eigenvalue r equal to its spectral radius (the Perron root).
2. There exists a strictly positive eigenvector $z \gg 0$ with $Az = rz$, unique up to scaling.
3. No other eigenvector is strictly positive.

Other eigenvalues can have the same modulus when the matrix is imprimitive, but the strictly positive eigenvector is unique up to scale.

See Section 1.2.3 of Sargent and Stachurski [2024] for details.

See also the full statement in The Perron-Frobenius Theorem.

Applied to the recovery problem: the Perron root is β (the subjective discount factor) and the Perron vector z determines D via $D_{ii} = 1/z_i$.

140.3.2 Ross's recovery theorem

The three assumptions in the theorem each carry a specific role.

Assuming the Arrow–Debreu state prices are identified, no-arbitrage guarantees that P has nonnegative entries and that the state prices encode a well-defined pricing measure.

Irreducibility ensures the economy is not divided into disconnected sub-economies – without it, the Perron–Frobenius theorem gives multiple candidate eigenvectors and recovery breaks down.

Transition independence is the key economic restriction.

It says the pricing kernel factors as $\beta h(\theta_j)/h(\theta_i)$, so the entire kernel is pinned down by a single vector h (or equivalently z).

With these in mind, the Recovery Theorem follows from the Perron–Frobenius theorem.

Theorem (Recovery Theorem)

Suppose prices provide no arbitrage opportunities, that the state price transition matrix P is irreducible, and that the pricing kernel is transition independent.

Then there exists a positive solution (β, z, F) to the recovery problem in which z is unique up to normalization, and the implied natural probability transition matrix F and transition pricing kernel are unique.

Proof. Because P is nonnegative and irreducible, the Perron–Frobenius theorem gives a unique positive eigenvector $z \gg 0$ with positive eigenvalue $\lambda > 0$ satisfying $Pz = \lambda z$.

Setting

$$\beta = \lambda, \qquad D_{ii} = \frac{1}{z_i},$$

the natural probability transition matrix is uniquely recovered as

$$f_{ij} = \frac{1}{\beta} \frac{z_j}{z_i} p_{ij}.$$

To confirm F is stochastic, note that all entries are nonnegative (since $p_{ij} \geq 0$ and $z_i, z_j > 0$) and

$$\sum_j f_{ij} = \frac{1}{\beta z_i} \sum_j z_j \, p_{ij} = \frac{[Pz]_i}{\beta z_i} = \frac{\beta z_i}{\beta z_i} = 1.$$

Uniqueness follows from the uniqueness of the Perron–Frobenius eigenvector.

140.3.3 Pricing kernel from the eigenvector

The recovered transition-kernel values are

$$\phi(\theta_i, \theta_j) = \beta \frac{z_i}{z_j}, \qquad h(\theta_i) = \frac{\beta}{z_i},$$

where $h(\theta_i) = \beta / z_i$ follows from $D_{ii} = h(\theta_i)/\beta = 1/z_i$.

It is useful to distinguish the **full transition kernel** $\phi_{ij} = \beta z_i / z_j$, which depends on both origin and destination states, from the **relative kernel component** $1/z_j$, which depends only on the destination state.

Ross's Table I reports the destination-state shape $1/z_j$, normalized so that the middle state equals one.

Destination states with high z_j have *low* kernel values: for a fixed origin i, the kernel $\beta z_i / z_j$ is decreasing in z_j.

When h represents marginal utility and states are ordered by consumption or payoff, larger z_j corresponds to lower marginal utility – "good times" that require less insurance and so receive less pricing weight per unit of natural probability.

The same eigenvector argument also yields a useful limiting case.

If the one-period bond price is identical in every current state, then the vector of ones is already the Perron vector, so recovery has no state-dependent change of measure left to perform.

Corollary

If the riskless rate is the same in all states ($Pe = be$ for some scalar b), then the unique natural distribution consistent with recovery is the risk-neutral (martingale) distribution itself: $F = (1/b)P$.

Proof. When $Pe = be$, the vector of ones e is the Perron eigenvector with eigenvalue b.

By the uniqueness part of the Perron–Frobenius theorem, $z = e$ (up to scaling) and $\beta = b$.

Setting $z = e$ gives $D = I$, so

$$F = \frac{1}{\beta} DPD^{-1} = \frac{1}{b} P. \qquad \square$$

140.3.4 Single crossing and the risk premium

Ross also uses the representative-agent formula to compare the natural and risk-neutral densities directly.

For a fixed current state θ_i,

$$\frac{q^*(\theta_i, \theta_j)}{f(\theta_i, \theta_j)} = e^{r(\theta_i)} \phi(\theta_i, \theta_j) = e^{r(\theta_i)} \beta \frac{U'(c(\theta_j))}{U'(c(\theta_i))}. \tag{140.2}$$

If U' is decreasing in consumption, then the ratio in (140.2) is decreasing in next-period consumption $c(\theta_j)$.

Since both $q^*(\theta_i, \cdot)$ and $f(\theta_i, \cdot)$ integrate to one, there is a crossing point v defined by

$$e^{r(\theta_i)} \beta U'(v) = U'(c(\theta_i)).$$

Below v, risk-neutral probability exceeds natural probability; above v, the natural probability exceeds the risk-neutral probability.

Hence the natural consumption distribution first-order stochastically dominates the risk-neutral one.

In a one-period model where terminal consumption is the market payoff, this also gives a positive market risk premium.

Let R denote the market return under the natural law, let R^* denote the same return under the risk-neutral law, and let R_f denote the riskless return in the same one-period units.

The stochastic-dominance result can be represented as

$$R^* \sim R - Z + \epsilon,$$

where $Z \geq 0$ captures the downward shift induced by risk adjustment and ϵ is a residual satisfying $E[\epsilon \mid R - Z] = 0$.

Because R^* is the market return under the risk-neutral measure, its expectation is the riskless return: $E[R^*] = R_f$.

Taking expectations of the representation above therefore gives

$$E[R] = E[R^*] + E[Z] = R_f + E[Z],$$

which strictly exceeds R_f whenever Z is positive with positive probability – that is, whenever the two densities genuinely differ.

140.4 Numerical example

We now demonstrate the Recovery Theorem numerically.

140.4.1 Building a finite-state example

We build the economy directly on a finite grid of log payoff states $s_1, \dots, s_m$.

On this grid we choose three primitives:

1. a row-stochastic irreducible natural transition matrix F,
2. a subjective discount factor $\beta = e^{-\rho T}$, and
3. a CRRA transition pricing kernel $\phi_{ij} = \beta e^{-\gamma(s_j - s_i)}$.

The state-price matrix is then constructed from

$$p_{ij} = \phi_{ij} f_{ij}.$$

This means the Recovery Theorem assumptions hold by construction: P is nonnegative, F is a Markov transition matrix, and the kernel is transition independent with $z_i \propto e^{\gamma s_i}$.

To keep the example close to Ross's Section IV, we choose F to have lognormal-shaped rows.

The continuous benchmark is a lognormal payoff with CRRA utility:

$$U(S_T) = \frac{S_T^{1-\gamma}}{1-\gamma}, \qquad S_T = S_0 \exp\Big((\mu - \tfrac{1}{2}\sigma^2)T + \sigma\sqrt{T}\xi\Big),$$

where $\xi \sim N(0,1)$, μ is the expected growth-rate parameter, σ is volatility, T is the horizon, γ is the CRRA coefficient, and ρ is the continuously compounded subjective discount rate.

The T-period pricing kernel is

$$\phi_T = e^{-\rho T} \left(\frac{S_T}{S_0}\right)^{-\gamma}.$$

Equivalently, if $s = \log S_0$ and $s_T = \log S_T$, then the state-price density with respect to the future log state s_T is

$$p_T(s, s_T) = e^{-\rho T} e^{-\gamma(s_T - s)} \frac{1}{\sigma\sqrt{T}} n\left(\frac{s_T - s - (\mu - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right),$$

where n is the standard normal density.

Thus the natural log return satisfies

$$\log(S_T/S_0) \sim \mathcal{N}\big((\mu - \tfrac{1}{2}\sigma^2)T, \sigma^2 T\big).$$

Following Ross's Table I, we represent the distribution on a finite grid of states.

This example is Ross-inspired rather than an exact reproduction of Ross's Table I.

Ross's Table I uses a fixed future payoff distribution, so its rows of F are identical.

Here the same CRRA/lognormal pricing logic is embedded in a finite Markov transition matrix whose rows shift with the current state.

Ross uses states from -5 to $+5$ standard deviations; we use the same range below.

The truncation is an essential part of the finite-state model: it is what brings the example into the Perron–Frobenius setting.

In the unbounded continuous lognormal growth model, Ross shows that recovery is not unique.

On the finite grid, the natural transition probabilities and state prices are

$$f_{ij} \propto n\left(\frac{s_j - s_i - (\mu - \frac{1}{2}\sigma^2)T}{\sigma\sqrt{T}}\right)\Delta s, \qquad p_{ij} = e^{-\rho T}e^{-\gamma(s_j - s_i)}f_{ij},$$

where $s_i = \ln(S_i/S_0)$ and $s_j = \ln(S_j/S_0)$ are log states measured relative to today's index level, $n(\cdot)$ is the standard normal density, and the discretized probabilities f_{ij} are normalized row by row.

The next cell constructs this finite grid and builds P.

```
def build_state_price_matrix(μ, σ, γ, ρ, T=1.0, n_states=11, n_σ=5):
    """Build a discretized lognormal/CRRA state-price matrix."""
    states = np.linspace(-n_σ * σ * np.sqrt(T),
                         n_σ * σ * np.sqrt(T),
                         n_states)
    ds = states[1] - states[0]

    m = n_states
    P = np.zeros((m, m))
    F = np.zeros((m, m))

    drift = (μ - 0.5 * σ**2) * T

    # First build a row-stochastic natural transition matrix on the bounded grid
    for i in range(m):
        s_i = states[i]
        for j in range(m):
            s_j = states[j]
            log_return = s_j - s_i
            F[i, j] = norm.pdf(log_return, loc=drift,
                               scale=σ * np.sqrt(T)) * ds

        F[i] = F[i] / F[i].sum()

        # Price each Arrow claim as natural probability times the CRRA kernel
        for j in range(m):
            log_return = states[j] - s_i
            kernel = np.exp(-ρ * T) * np.exp(-γ * log_return)
            P[i, j] = kernel * F[i, j]

    return P, states
```

Now choose a calibration and build the state-price matrix.

```
μ = 0.08    # 8% annual expected return
σ = 0.20    # 20% annual volatility
γ = 3.0     # CRRA coefficient
ρ = 0.02    # 2% annual continuous discount rate
T = 1.0     # one-year horizon

P, states = build_state_price_matrix(μ, σ, γ, ρ, T,
                                     n_states=11, n_σ=5)

print("State-price row sums:")
```

(continues on next page)

(continued from previous page)

```
print(np.round(P.sum(axis=1), 4))
print(f"Middle-state risk-free rate: {-np.log(P[5].sum()):.4f}")
```

```
State-price row sums:
[0.7154 0.9046 0.9713 0.9798 0.9802 0.9802 0.9802 0.9806 0.9892 1.0621
 1.3429]
Middle-state risk-free rate: 0.0200
```

The row sums are the model-implied one-period bond prices in each current state.

They vary near the boundaries because the finite grid truncates and renormalizes the conditional transition probabilities.

The middle-state rate printed above happens to equal ρ only because this calibration satisfies $\mu = \frac{1}{2}\sigma^2(1+\gamma)$, which holds at $\mu = 0.08$, $\sigma = 0.2$ and $\gamma = 3$.

It is a coincidence of the numbers chosen here, not a check on the code or a general property.

140.4.2 Applying the recovery theorem

The Recovery Theorem requires computing the **Perron eigenvector** of P.

```
def recover_natural_distribution(P, tol=1e-10):
    """
    Recover natural probabilities and the relative pricing kernel
    from state prices.
    """

    m = P.shape[0]

    eigenvalues, eigenvectors = eig(P)
    eigenvalues = np.real_if_close(eigenvalues, tol=1000)
    eigenvectors = np.real_if_close(eigenvectors, tol=1000)

    # Ross recovery uses the Perron root and its strictly positive eigenvector
    real_mask = np.isreal(eigenvalues)
    real_eigenvalues = np.asarray(
        eigenvalues[real_mask].real, dtype=float)
    real_eigenvectors = np.asarray(
        eigenvectors[:, real_mask].real, dtype=float)

    order = np.argsort(real_eigenvalues)[::-1]

    for idx in order:
        β_candidate = real_eigenvalues[idx]
        z_candidate = real_eigenvectors[:, idx]

        if np.mean(z_candidate) < 0:
            z_candidate = -z_candidate

        # eig returns eigenvectors of unit Euclidean norm, so the smallest entry
        # of a legitimately positive z can be far below any absolute threshold
        # (it is of order e^{-2γ n_σ σ√T} here). We therefore rescale and test
        # the sign, which makes the test invariant to the eigenvector's scale.
        z_candidate = z_candidate / np.max(np.abs(z_candidate))

        if β_candidate > 0 and np.all(z_candidate > 0):
```

(continues on next page)

(continued from previous page)

```
            β_recovered = β_candidate
            z = z_candidate
            break
    else:
        raise ValueError("No strictly positive real eigenvector found")

    z = z / z[m // 2]

    D = np.diag(1.0 / z)
    D_inv = np.diag(z)

    # Converts state prices into probabilities
    F = (1.0 / β_recovered) * D @ P @ D_inv

    min_entry = F.min()
    row_sum_error = np.max(np.abs(F.sum(axis=1) - 1.0))

    if min_entry < -tol:
        raise ValueError(f"Recovered F has negative entries: min={min_entry}")

    if row_sum_error > 1e-8:
        raise ValueError(
            f"Recovered F row sums are not one: max error={row_sum_error}"
        )

    # The kernel relative to the middle state normalization
    ψ_relative = 1.0 / z

    return F, z, β_recovered, ψ_relative
```

The Perron vector also recovers the shape of the pricing kernel.

Ross's Table I reports this shape with the middle state normalized to one, which is $1/z_j$ under our normalization $z_{\text{mid}} = 1$.

```
F, z, β_rec, ψ_relative = recover_natural_distribution(P)

print("Ross-normalized kernel 1/z (middle state = 1):")
print(np.round(ψ_relative, 4))
```

```
Ross-normalized kernel 1/z (middle state = 1):
[20.0855 11.0232  6.0496  3.3201  1.8221  1.      0.5488  0.3012  0.1653
  0.0907  0.0498]
```

Because we know the data-generating natural transition matrix used to construct P, we can verify that recovery works in this simulation.

```
def true_lognormal_transition_matrix(states, μ, σ, T):
    """
    Construct the bounded-grid natural transition matrix used in the simulation.
    """
    m = len(states)
    ds = states[1] - states[0]
    drift = (μ - 0.5 * σ**2) * T
    F_true = np.zeros((m, m))

    for i in range(m):
        log_returns = states - states[i]
```

(continues on next page)

(continued from previous page)

```
        F_true[i] = norm.pdf(log_returns, loc=drift,
                             scale=σ * np.sqrt(T)) * ds
        F_true[i] = F_true[i] / F_true[i].sum()

    return F_true


F_true = true_lognormal_transition_matrix(states, μ, σ, T)
P_reconstructed = β_rec * (z[:, None] / z[None, :]) * F

print("Recovery numerical checks")
print(f"max |F - true F| = {np.max(np.abs(F - F_true)):.2e}")
print(f"max |P - recovered kernel times F| = "
      f"{np.max(np.abs(P - P_reconstructed)):.2e}")
```

```
Recovery numerical checks
max |F - true F| = 3.46e-14
max |P - recovered kernel times F| = 1.11e-16
```

Indeed, the discrepancies are at the level of numerical roundoff.

140.5 Natural vs. risk-neutral distributions

A key insight of Ross [2015] is that the natural distribution can differ systematically from the risk-neutral one.

In this CRRA example, where states are ordered from low to high payoff, the single-crossing argument in *Single crossing and the risk premium* implies that the natural marginal density **first-order stochastically dominates** the risk-neutral density: the CDF of the natural distribution lies *below* that of the risk-neutral distribution.

Because the pricing kernel is declining (investors fear bad outcomes), risk-neutral probabilities overweight bad states and underweight good states relative to the natural measure.

We first plot the natural distribution against the risk-neutral one and the recovered relative pricing kernel.

```
mid = len(states) // 2

row_sums = P.sum(axis=1, keepdims=True)

# Normalize Arrow prices by the one-period riskless bond price in each state
Q_rn = P / row_sums

f_nat = F[mid, :]
f_rn = Q_rn[mid, :]

gross_returns = np.exp(states)

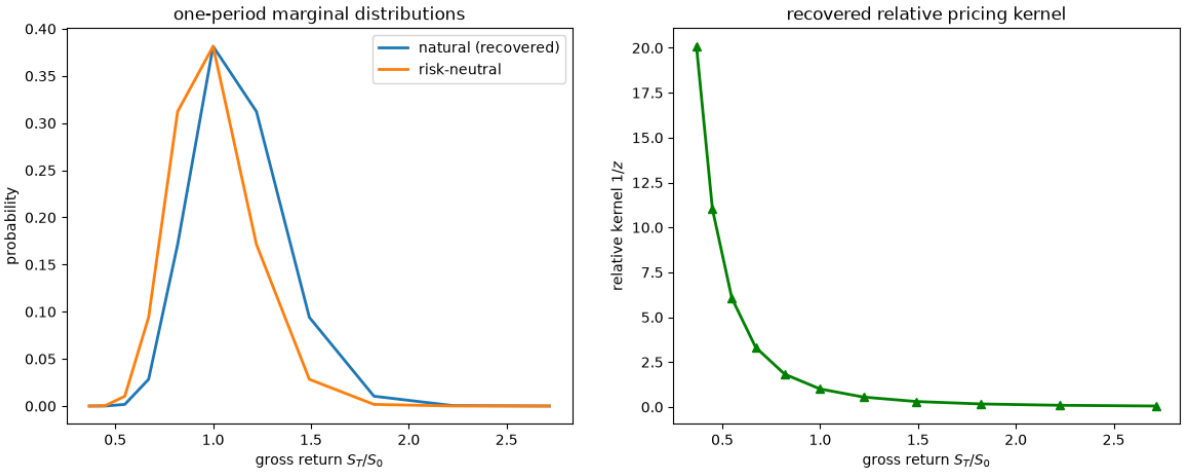
fig, axes = plt.subplots(1, 2, figsize=(14, 5))

axes[0].plot(gross_returns, f_nat, label='natural (recovered)', lw=2)
axes[0].plot(gross_returns, f_rn, label='risk-neutral', lw=2)
axes[0].set_xlabel('gross return $S_T / S_0$')
axes[0].set_ylabel('probability')
axes[0].set_title('one-period marginal distributions')
axes[0].legend()
```

(continues on next page)

(continued from previous page)

```
axes[1].plot(gross_returns, ψ_relative, 'g-^', lw=2)
axes[1].set_xlabel('gross return $S_T / S_0$')
axes[1].set_ylabel('relative kernel $1/z$')
axes[1].set_title('recovered relative pricing kernel')
plt.show()
```

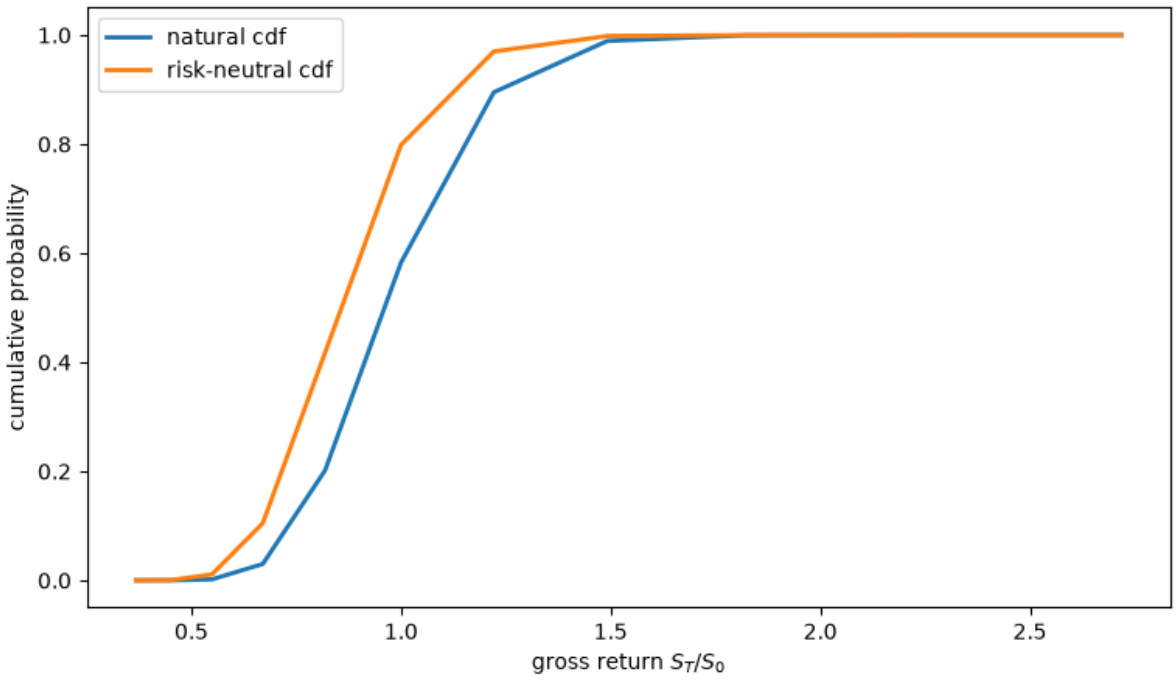


The CDF clearly shows the first-order stochastic dominance.

```
cdf_nat = np.cumsum(f_nat)
cdf_rn = np.cumsum(f_rn)

fig, ax = plt.subplots(figsize=(9, 5))
ax.plot(gross_returns, cdf_nat, lw=2, label='natural cdf')
ax.plot(gross_returns, cdf_rn, lw=2, label='risk-neutral cdf')
ax.set_xlabel('gross return $S_T / S_0$')
ax.set_ylabel('cumulative probability')
ax.legend()
plt.show()

print(f"Natural CDF <= Risk-neutral CDF at all states: "
      f"{np.all(cdf_nat <= cdf_rn + 1e-10)}")
```




```
Natural CDF <= Risk-neutral CDF at all states: True
```

The gap between the two CDFs is generated by the slope of the pricing kernel.

In the CRRA benchmark, this slope is controlled by the risk-aversion coefficient γ.

We next vary γ to see how the recovered kernel and the natural/risk-neutral wedge change.

140.6 Effect of risk aversion

The shape of the pricing kernel, and hence the gap between natural and risk-neutral probabilities, depends on the coefficient of risk aversion γ.

We illustrate this by plotting the relative pricing kernel $1/z$ and the gap between the natural and risk-neutral densities for a range of values of γ.

The left panel uses a logarithmic vertical axis because $1/z$ spans about seven orders of magnitude across these values of γ, so on a linear axis the curves for small γ would be flat against the horizontal axis and invisible.

On the log axis the economic content is immediate: since $1/z_j = e^{-\gamma s_j}$ exactly, $\log(1/z_j)$ is linear in the log state s_j with slope $-\gamma$, so the recovered kernel is log-linear and γ is simply its slope.

```
γs = [1.0, 2.0, 3.0, 5.0, 8.0]
colors = cm.viridis(np.linspace(0.1, 0.9, len(γs)))

fig, axes = plt.subplots(1, 2, figsize=(14, 5))

for γ_val, color in zip(γs, colors):
    P_g, states_g = build_state_price_matrix(μ, σ, γ_val, ρ, T)
    F_g, z_g, β_g, ψ_relative_g = recover_natural_distribution(P_g)
    mid_g = len(states_g) // 2
```

(continues on next page)

(continued from previous page)

```
    f_nat_g = F_g[mid_g, :]
    row_sum = P_g[mid_g].sum()
    f_rn_g = P_g[mid_g] / row_sum

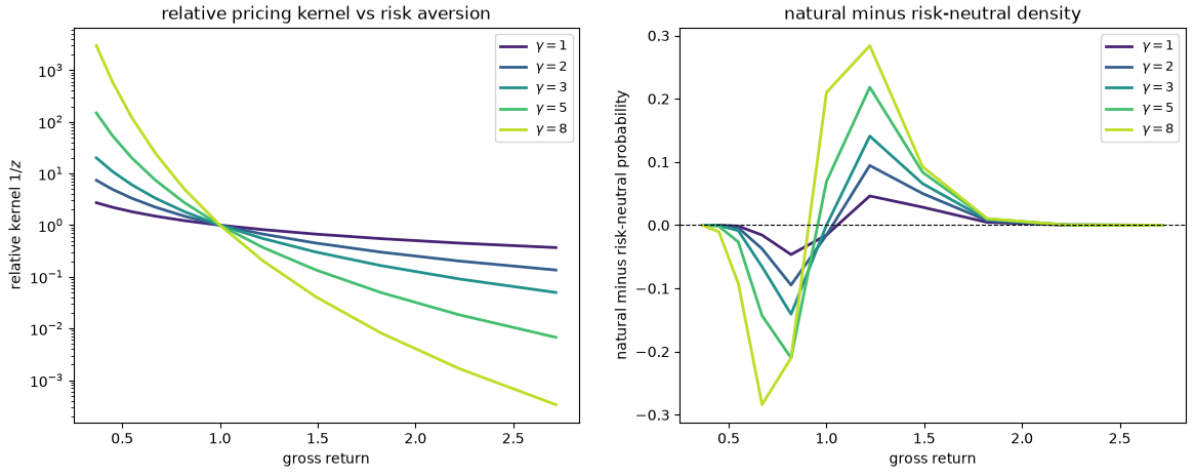
    gross = np.exp(states_g)

    axes[0].plot(gross, ψ_relative_g, color=color, lw=2,
                 label=f'$\\gamma={γ_val:.0f}$')
    axes[1].plot(gross, f_nat_g - f_rn_g, color=color, lw=2,
                 label=f'$\\gamma={γ_val:.0f}$')

axes[0].set_yscale('log')
axes[0].set_xlabel('gross return')
axes[0].set_ylabel('relative kernel $1/z$')
axes[0].set_title('relative pricing kernel vs risk aversion')
axes[0].legend(fontsize=9)

axes[1].axhline(0, color='k', lw=0.8, ls='--')
axes[1].set_xlabel('gross return')
axes[1].set_ylabel('natural minus risk-neutral probability')
axes[1].set_title('natural minus risk-neutral density')
axes[1].legend(fontsize=9)

plt.show()
```



Because the states are ordered from low to high payoff, the plots show the single-crossing property discussed in *Single crossing and the risk premium*: for returns below some threshold v, risk-neutral probability exceeds natural probability; above v the natural probability dominates.

A higher γ amplifies this wedge.

A caution about the large-γ cases: the one-period bond price in the middle state rises to 1.1972 at $\gamma = 5$ and 2.1815 at $\gamma = 8$, implied riskless rates of -18% and -78%, and the largest row sum of P reaches 3.63.

This is the usual CRRA precautionary-savings blow-up rather than a defect of the recovery calculation, whose Perron root remains $e^{-\rho T}$ for every γ.

We include these values to make the shape of the kernel visible across a wide range, not as plausible calibrations of an actual economy.

140.7 Recovering the discount rate

A useful by-product of the Recovery Theorem is the *recovered subjective discount factor* β, which equals the Perron–Frobenius eigenvalue of P.

If the horizon is T, the corresponding continuously compounded subjective discount rate is

$$\rho = -\frac{\log \beta}{T}.$$

In the numerical examples below, $T = 1$, so this reduces to $\rho = -\log \beta$.

Corollary 1 of Ross [2015] states that β is bounded above by the largest state-dependent one-period discount factor — equivalently, the maximum row sum of P:

$$\beta \leq \max_i \sum_j p(\theta_i, \theta_j).$$

Sweeping the true ρ over a grid and reporting the recovered values alongside the recovery error confirms that the eigenvalue calculation pins down β accurately:

```
true_ρs = np.linspace(0.00, 0.06, 13)
recovered_ρs = np.empty_like(true_ρs)

for k, rho in enumerate(true_ρs):
    P_d, _ = build_state_price_matrix(μ, σ, γ=3.0, ρ=rho, T=1.0)
    _, _, β_d, _ = recover_natural_distribution(P_d)
    recovered_ρs[k] = -np.log(β_d)

print(
    f"max |true ρ - recovered ρ| = {np.max(np.abs(true_ρs - recovered_ρs)):.2e}")
np.column_stack([true_ρs, recovered_ρs])
```

```
max |true ρ - recovered ρ| = 2.11e-15
```

```
array([[ 0.0000000e+00, -4.4408921e-16],
       [ 5.0000000e-03,  5.0000000e-03],
       [ 1.0000000e-02,  1.0000000e-02],
       [ 1.5000000e-02,  1.5000000e-02],
       [ 2.0000000e-02,  2.0000000e-02],
       [ 2.5000000e-02,  2.5000000e-02],
       [ 3.0000000e-02,  3.0000000e-02],
       [ 3.5000000e-02,  3.5000000e-02],
       [ 4.0000000e-02,  4.0000000e-02],
       [ 4.5000000e-02,  4.5000000e-02],
       [ 5.0000000e-02,  5.0000000e-02],
       [ 5.5000000e-02,  5.5000000e-02],
       [ 6.0000000e-02,  6.0000000e-02]])
```

140.8 Tail risk: natural vs. risk-neutral probabilities of catastrophe

One of the most striking applications of the Recovery Theorem is its ability to separate the market's recovered natural probability of catastrophes from the risk premium attached to them.

Mehra and Prescott [1985] posed the equity premium puzzle, and Barro [2006] argues that a small probability of rare macroeconomic disasters can help account for it.

The risk-neutral probability of a large decline is elevated both because (a) the market assigns a high natural probability to such events and (b) the pricing kernel upweights bad outcomes.

Ross's recovery machinery lets us decompose these two forces.

The next cell plots left-tail probabilities under the recovered natural and the risk-neutral measures from the middle state, so the gap between the curves isolates the pricing-kernel contribution to crash probabilities.

```
thresholds = np.linspace(-0.40, 0.10, 200)

def tail_prob(f_dist, states, threshold):
    """Left-tail probability for log returns."""
    return float(np.sum(f_dist[states <= threshold]))

P_base, states_base = build_state_price_matrix(
    μ, σ, γ=3.0, ρ=0.02, T=1.0,
    n_states=41, n_σ=5)
F_base, z_base, β_base, ψ_relative_base = recover_natural_distribution(P_base)

mid_b = len(states_base) // 2
f_nat_base = F_base[mid_b]
f_rn_base = P_base[mid_b] / P_base[mid_b].sum()

prob_nat = [tail_prob(f_nat_base, states_base, t) for t in thresholds]
prob_rn = [tail_prob(f_rn_base, states_base, t) for t in thresholds]

fig, ax = plt.subplots(figsize=(10, 5))
ax.plot(np.exp(thresholds), prob_nat, lw=2, label='natural (recovered)')
ax.plot(np.exp(thresholds), prob_rn, lw=2, label='risk-neutral')
ax.set_xlabel('gross return threshold')
ax.set_ylabel('probability of decline below threshold')
ax.axvline(x=0.75, color='gray', ls=':', lw=1.5, label='25% decline')
ax.axvline(x=0.70, color='silver', ls=':', lw=1.5, label='30% decline')
ax.legend()
plt.show()
```

This is a simulation illustrating Ross's decomposition.

The risk-neutral density assigns higher probability to large drops than the recovered natural density.

In this CRRA simulation, increasing risk aversion makes the risk-neutral crash probability rise faster than the recovered natural crash probability.

We will say more in Exercise 140.12.3.

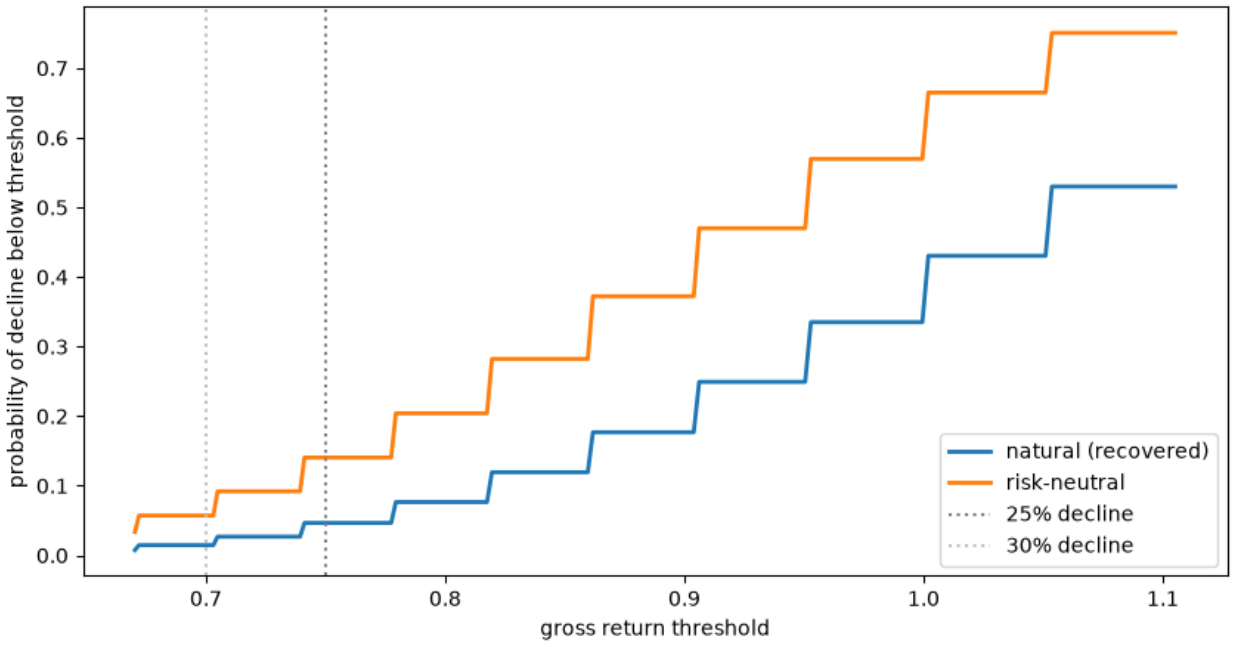


Fig. 140.1: Tail probabilities under the recovered natural and risk-neutral measures

140.9 From option prices to transition prices

The numerical example above starts from a known state-price transition matrix P.

Empirically, Ross starts one step earlier: option prices reveal state-price densities at different maturities from the current state, and the transition matrix must be inferred from those maturity-by-maturity state prices.

Let $C(K,T)$ be the price of a call option with strike K and maturity T.

If $p(S,T)$ is the state-price density for terminal index level S, then

$$C(K,T) = \int_K^\infty (S-K)p(S,T)\,dS.$$

Differentiating twice with respect to the strike gives the Breeden and Litzenberger [1978] formula

$$p(K,T) = \frac{\partial^2 C(K,T)}{\partial K^2}.$$

After discretizing strikes and maturities, let

$$p_t(c) = (p_t(c,1), \ldots, p_t(c,m))$$

be the vector of state prices at horizon t observed from today's state c.

Here c indexes the current state and t counts discrete maturity steps.

The first one-period vector $p_1(c)$ identifies the row of P corresponding to the current state c, supplying m equations.

If the one-period state-price transition matrix P is time homogeneous, these vectors satisfy the forward recursion

$$p_{t+1}(c) = p_t(c)P, \qquad t = 1, \ldots, m-1.$$

Componentwise,

$$p_{t+1}(c, j) = \sum_k p_t(c, k) p(k, j).$$

The remaining $m - 1$ forward equations $p_{t+1}(c) = p_t(c)P$, each with m components, supply the remaining $m(m - 1)$ equations.

Together these give m^2 equations for the m^2 transition prices $p(k, j)$.

In practice this step is numerically delicate because the second derivative in the option-price formula amplifies measurement error, and because additional shape restrictions such as positivity or unimodality may be needed to obtain a reasonable transition matrix.

140.10 Testing efficient markets

The recovered pricing kernel can also be used to test market efficiency, under the assumptions of the Recovery Theorem.

If a trading strategy has a very high Sharpe ratio, then some pricing kernel must be volatile enough to price that payoff.

The Hansen–Jagannathan bound [Hansen and Jagannathan, 1991] says that, for any excess return with mean μ_{excess} and standard deviation σ_{excess},

$$\frac{|\mu_{\text{excess}}|}{\sigma_{\text{excess}}} \leq e^{rT}\, \sigma(\phi),$$

where ϕ is the one-period pricing kernel of (140.1) and r is the continuously compounded riskless rate over horizon T.

Ross's point is that recovery gives an estimate of the relevant volatility $\sigma(\phi)$.

Hence it gives an upper bound on the Sharpe ratio of any strategy based on the same stock-market information used in recovery.

If such a strategy has a Sharpe ratio above the bound, then it is too profitable to be consistent with efficiency, under the assumptions of the Recovery Theorem.

The same logic gives a bound on return predictability.

Suppose excess returns are decomposed as

$$x_{t+1} = \mu(I_t) + \epsilon_{t+1},$$

where I_t is the stock-market information set and ϵ_{t+1} is unpredictable from I_t.

Then the R^2 of a forecasting regression based on I_t is bounded above by the variance of the recovered kernel:

$$R^2 \leq e^{2rT}\, \sigma^2(\phi).$$

Only the component of the kernel projected on this information set is relevant.

Adding unrelated noise to a candidate pricing kernel would raise its variance, but it would not justify stronger return predictability from stock-market information.

140.11 Limitations and extensions

The Recovery Theorem is a remarkable theoretical result, but several caveats apply in practice.

Finite state space:

Ross's theorem is proved for a finite-state irreducible Markov chain.

Recovery on a continuous state space requires a boundedness or compactness restriction: Carr and Yu [2012], discussed below, establish recovery for a bounded diffusion.

In continuous, unbounded state spaces (e.g., a lognormal diffusion), uniqueness fails because any exponential $e^{\alpha x}$ satisfies the characteristic equation.

To see the issue, consider the continuous lognormal growth state-price density above.

The natural continuous-space analogue of the Perron–Frobenius problem is

$$\int p_T(s, y) v(y)\, dy = \lambda v(s).$$

Here y is a possible future log state, v is a candidate positive eigenfunction, and λ is its eigenvalue.

For every real α, the exponential function $v_\alpha(s) = e^{\alpha s}$ solves this equation with eigenvalue

$$\lambda(\alpha) = \exp\!\left(-\rho T + (\alpha - \gamma)(\mu - \tfrac{1}{2}\sigma^2)T + \tfrac{1}{2}\sigma^2 T(\alpha - \gamma)^2\right).$$

The positive eigenfunction is therefore not unique.

This is why truncation or boundedness assumptions matter: they turn the continuous operator problem back into a Perron–Frobenius problem with a unique positive eigenvector.

Carr and Yu [2012] establish recovery with a bounded diffusion.

Transition independence:

If the kernel is not transition independent, recovery is not guaranteed.

Borovička *et al.* [2016] show that Ross recovery can confound the long-run risk component of the kernel with the natural probability distribution, yielding an incorrect decomposition.

We discuss this in the sequel lecture *Misspecified Recovery*.

See also

Long-term Risk: An Operator Approach develops the general operator treatment of this problem, in which the principal eigenfunction of the pricing operator and the resulting multiplicative martingale factorization of the stochastic discount factor are precisely what transition independence assumes away.

Be warned that the symbol ϕ denotes the pricing kernel in the present lecture but the principal eigenfunction there; it is this lecture's eigenvector z that plays the role of that lecture's ϕ.

Empirical estimation:

Extracting reliable state prices from observed option prices requires careful interpolation and extrapolation.

The mapping from implied volatilities to state prices via the Breeden and Litzenberger [1978] formula involves second derivatives, which amplify measurement error.

140.12 Exercises

Exercise 140.12.1

The Perron–Frobenius vector and the pricing kernel.

Consider the 3×3 state price matrix

$$P = \begin{pmatrix} 0.5950 & 0.1700 & 0.0272 \\ 0.159375 & 0.5525 & 0.1360 \\ 0.06640625 & 0.31875 & 0.5525 \end{pmatrix}.$$

1. Compute the Perron eigenvalue β and the corresponding eigenvector z of P.
2. Use z to recover the natural probability transition matrix F via
$$f_{ij} = \frac{1}{\beta} \frac{z_j}{z_i} p_{ij}.$$
3. Verify that each row of F sums to one and all entries are positive.
4. For destination state j, the relative kernel component is $1/z_j$; for a transition from state i to state j, the full pricing kernel is $\beta z_i / z_j$. Compute $1/z_j$ for each state.

Does the kernel decrease as we move from state 1 to state 3 (i.e., from bad to good states)?

Solution

Here is one solution:

```
P_ex = np.array([
    [0.5950, 0.1700, 0.0272],
    [0.159375, 0.5525, 0.1360],
    [0.06640625, 0.31875, 0.5525]
])

eigenvalues, eigenvectors = eig(P_ex)
real_mask = np.isreal(eigenvalues)
real_ev = eigenvalues[real_mask].real
real_evec = eigenvectors[:, real_mask].real

idx = np.argmax(real_ev)
β_ex = real_ev[idx]
z_ex = real_evec[:, idx]
if z_ex.min() < 0:
    z_ex = -z_ex
z_ex = z_ex / z_ex[1]

print(f"β = {β_ex:.6f}")
print(f"z = {z_ex}")

D_ex = np.diag(1.0 / z_ex)
D_inv_ex = np.diag(z_ex)
F_ex = (1.0 / β_ex) * D_ex @ P_ex @ D_inv_ex

print("\nRecovered F:")
print(np.round(F_ex, 4))
```

```
print(f"\nRow sums: {np.round(F_ex.sum(axis=1), 8)}")
print(f"Nonnegative: {(F_ex >= -1e-10).all()}")

ψ_relative_ex = 1.0 / z_ex
print(f"\nrelative kernel 1/z = {np.round(ψ_relative_ex, 4)}")
print(f"Decreasing: {ψ_relative_ex[0] > ψ_relative_ex[1] > ψ_relative_ex[2]}")
```

```
β = 0.850000
z = [0.8  1.   1.25]

Recovered F:
[[0.7  0.25 0.05]
 [0.15 0.65 0.2 ]
 [0.05 0.3  0.65]]

Row sums: [1. 1. 1.]
Nonnegative: True

relative kernel 1/z = [1.25 1.   0.8 ]
Decreasing: True
```

Exercise 140.12.2

Stochastic dominance.

Using the recovered F and the normalized risk-neutral matrix $Q = P/\text{row sums}$ from the exercise above:

1. Compute the one-step marginal distributions $f_j = F_{2,j}$ and $q_j = Q_{2,j}$ starting from state 2 (index 1 in Python).
2. Compute the CDFs $\hat{F}_k = \sum_{j \le k} f_j$ and $\hat{Q}_k = \sum_{j \le k} q_j$ for each state.
3. Verify numerically that $\hat{F}_k \le \hat{Q}_k$ for every k, confirming stochastic dominance in this ordered three-state example.

Solution

Here is one solution:

```
eigenvalues, eigenvectors = eig(P_ex)
real_mask = np.isreal(eigenvalues)
real_ev = eigenvalues[real_mask].real
real_evec = eigenvectors[:, real_mask].real
idx = np.argmax(real_ev)
β_ex = real_ev[idx]
z_ex = real_evec[:, idx]
if z_ex.min() < 0:
    z_ex = -z_ex
z_ex = z_ex / z_ex[1]

D_ex = np.diag(1.0 / z_ex)
D_inv_ex = np.diag(z_ex)
F_ex = (1.0 / β_ex) * D_ex @ P_ex @ D_inv_ex

start = 1
```

```
f_marg = F_ex[start]
q_marg = P_ex[start] / P_ex[start].sum()

print("One-step marginals from state 2:")
print(f"natural      = {np.round(f_marg, 4)}")
print(f"risk-neutral = {np.round(q_marg, 4)}")

cdf_nat = np.cumsum(f_marg)
cdf_rn = np.cumsum(q_marg)

print("\nCDFs:")
for k in range(3):
    print(f"state {k+1}: natural = {cdf_nat[k]:.4f}, risk-neutral = {cdf_rn[k]:.4f}
 ↪")

dominates = np.all(cdf_nat <= cdf_rn + 1e-10)
print(f"\nNatural CDF <= risk-neutral CDF: {dominates}")
```

```
One-step marginals from state 2:
natural      = [0.15 0.65 0.2 ]
risk-neutral = [0.188  0.6516 0.1604]

CDFs:
state 1: natural = 0.1500, risk-neutral = 0.1880
state 2: natural = 0.8000, risk-neutral = 0.8396
state 3: natural = 1.0000, risk-neutral = 1.0000

Natural CDF <= risk-neutral CDF: True
```

Exercise 140.12.3

Risk aversion and tail risk.

Write a function `tail_risk_ratio(γ, threshold, μ, σ, ρ, T)` that:

1. Constructs the state price matrix P using `build_state_price_matrix` with the given parameters and `n_states=41`.
2. Applies `recover_natural_distribution` to obtain F.
3. Computes $P(\text{log-return} \leq \text{threshold})$ under both the natural and risk-neutral distributions starting from the middle state.
4. Returns the ratio $p_{\text{risk-neutral}}/p_{\text{natural}}$.

Using this function, plot the ratio as a function of $\gamma \in [1, 10]$ for a 30 percent simple decline, i.e. `threshold = np.log(0.70)`.

Explain the economic interpretation: why does a higher γ raise the ratio?

Solution

Here is one solution:

```
def tail_risk_ratio(γ, threshold, μ=0.08, σ=0.20, ρ=0.02, T=1.0):
    """Risk-neutral / natural left-tail probability."""
    P_g, states_g = build_state_price_matrix(
```

```
        μ, σ, γ, ρ, T, n_states=41, n_σ=5)

    F_g, _, _, _ = recover_natural_distribution(P_g)

    mid_g = len(states_g) // 2

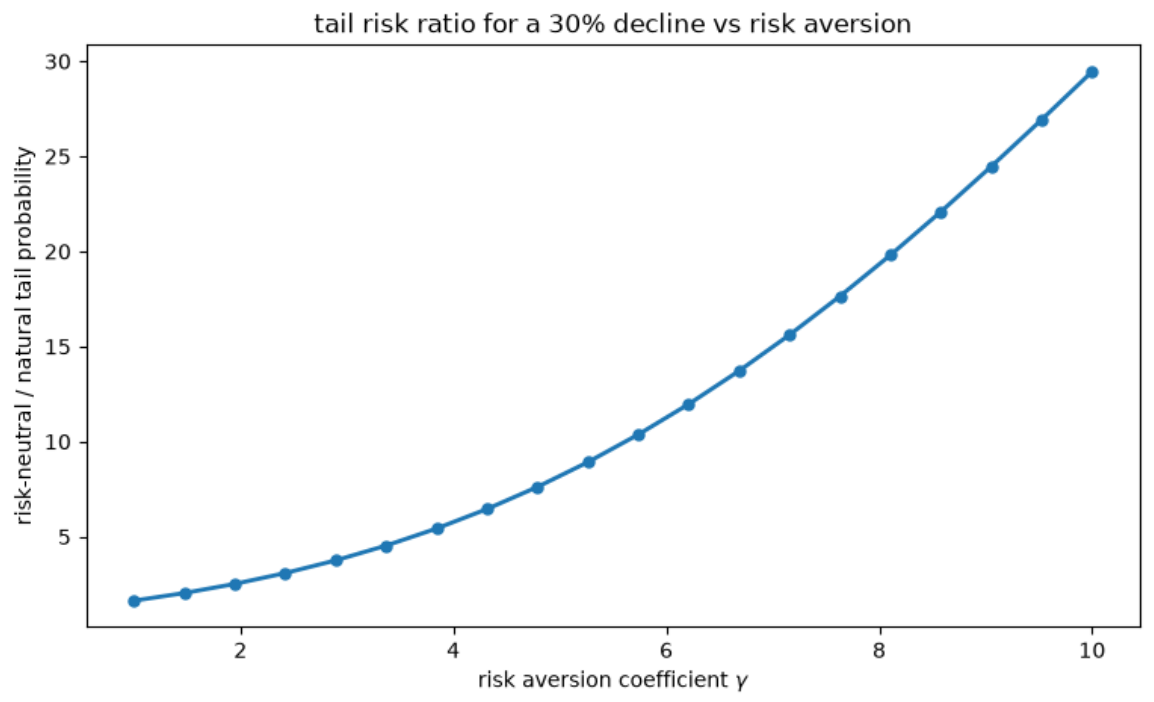
    f_nat_g = F_g[mid_g]
    f_rn_g = P_g[mid_g] / P_g[mid_g].sum()

    p_nat = float(np.sum(f_nat_g[states_g <= threshold]))
    p_rn = float(np.sum(f_rn_g[states_g <= threshold]))

    if p_nat < 1e-12:
        return np.nan
    return p_rn / p_nat


γs = np.linspace(1.0, 10.0, 20)
threshold_30 = np.log(0.70)
ratios = [tail_risk_ratio(g, threshold_30) for g in γs]

plt.figure(figsize=(9, 5))
plt.plot(γs, ratios, '-o', ms=5, lw=2)
plt.xlabel('risk aversion coefficient $\\gamma$')
plt.ylabel('risk-neutral / natural tail probability')
plt.title('tail risk ratio for a 30% decline vs risk aversion')
plt.show()
```



A higher coefficient of risk aversion γ makes the pricing kernel steeper: the market assigns a larger premium per unit of probability to bad-state payoffs.

Risk-neutral probabilities incorporate this premium, so in this CRRA simulation the risk-neutral crash probability rises faster with γ than the recovered natural crash probability.

Recovery separates the market's estimated natural crash probability from the pricing-kernel premium attached to crash states.

CHAPTER

ONE

LONG-TERM RISK: AN OPERATOR APPROACH

Contents

- *Long-term Risk: An Operator Approach*
  - *Overview*
  - *Multiplicative functionals*
  - *Semigroups*
  - *The generator*
  - *Principal eigenfunctions*
  - *A finite-state Markov chain*
  - *The affine diffusion example*
  - *Long-run risk prices*
  - *Assumptions behind the scenes*
  - *Summary*
  - *Exercises*

141.1 Overview

Standard short-horizon asset pricing tells us how equilibrium prices compensate investors for tiny, instantaneous exposures to shocks.

That is, they tell us about the *short end* of the term structure of risk prices.

But many of the most interesting asset pricing questions (e.g., the equity premium puzzle, the slope of the yield curve, the prices of long-dated options) concern the *long end* of the term structure of risk prices.

This lecture studies the long end using the operator approach of Hansen and Scheinkman [2009].

At center stage will be a stochastic discount factor process and a return process that are multiplicative across time in the sense that conditional expectations define a *semigroup* of valuation operators indexed by horizon t.

Long-horizon behavior of the semigroup is controlled by a single eigenvalue problem on the state space.

By solving that eigenvalue problem and selecting an appropriate eigenfunction, we can factor the multiplicative functional M_t into three economically meaningful pieces: a deterministic exponential trend, a non-negative martingale that changes

probability measure, and a transient state-dependent component.

In particular, the factorization takes the form

$$M_t = \exp(\rho t)\, \widehat{M}_t \, \frac{\phi(X_0)}{\phi(X_t)}, \tag{141.1}$$

where

- ρ is a scalar (the **principal eigenvalue**) giving the asymptotic growth or decay rate,
- ϕ is a strictly positive **principal eigenfunction** on the state space, capturing the long-run dependence on the current state,
- $\widehat{M}$ is a positive martingale used to change probability measure (the **martingale component**), and
- $\phi(X_0)/\phi(X_t)$ is a transient component that washes out as the twisted process settles into its stationary distribution.

In finite-state problems this is exactly the Perron-Frobenius decomposition of a positive matrix; in general state spaces it is a continuous-state counterpart.

We will refer to (141.1) as the **multiplicative factorization** associated with $(\rho, \phi, \widehat{M})$.

Alvarez and Jermann [2005] applied a related permanent-transitory decomposition to stochastic discount factors.

The operator approach links that decomposition to an explicit eigenvalue problem.

See also

This lecture is closely related to the lecture Additive and Multiplicative Functionals, which studies the same kind of permanent-transitory decomposition for additive and multiplicative functionals in a discrete-time linear-Gaussian setting.

Reading these two lectures together is a good way to learn about representations of long-run risks in both continuous and discrete time.

We will build up to (141.1) and use it to compute long-run risk prices in concrete models.

The plan of this lecture is to:

1. Set up positive multiplicative functionals M (discount factors, returns, stochastic growth) and the valuation semigroups they generate.
2. Introduce the **generator** of a semigroup, a local operator whose eigenvalue problem controls long-run behavior.
3. Find the principal eigenfunction ϕ and derive the factorization.
4. Work two examples in detail: a finite-state Markov chain (where everything reduces to Perron-Frobenius theory) and an affine diffusion model (where the eigenfunction is exponential-affine and we get closed-form formulas).
5. Use the factorization to compute long-run risk prices and compare them to the local risk prices appropriate for short-horizon asset pricing.

A recurring theme is that shocks to persistent state variables drive a wedge between local and long-run risk prices.

Generating that wedge is the mechanism through which long-run risk models like Bansal and Yaron [2004] generate large equity premia.

We start with the following imports

```
import numpy as np
import matplotlib.pyplot as plt
from scipy.linalg import eig, expm
```

141.2 Multiplicative functionals

141.2.1 Setting

Let $\{X_t : t \geq 0\}$ be a continuous-time Markov process with state space $\mathcal{D}_0$, and let $\mathcal{F}_t$ denote the filtration generated by its history.

We will work with a strong Markov process whose sample paths are càdlàg (defined below).

To arrive at the explicit formulas presented later we will specialize to a semimartingale that decomposes into a continuous component X^c and a pure-jump component X^j:

$$X_t = X_t^c + X_t^j.$$

We write the continuous-component dynamics as

$$dX_t^c = \xi(X_{t-})\,dt + \Gamma(X_{t-})\,dB_t,$$

and the pure-jump component as

$$dX_t^j = \int (y - X_{t-})\,\zeta(dy, dt),$$

where ζ is the random counting measure of jumps indexed by the *post-jump state* y, so that $\zeta(B, [0, t])$ counts the number of jumps in $[0, t]$ landing in the Borel set B.

Its compensator is $\eta(dy \mid X_{t-})\,dt$, the rate at which X jumps from X_{t-} into a region dy.

We will use this "y = post-jump state" convention throughout, matching the form $\phi(y) - \phi(x)$ that appears in the generator below.

We also impose two simplifying assumptions:

- **Finite jumps** on finite time intervals: only finitely many jumps occur on any bounded interval; this keeps integrals against the jump measure well-defined and finite.
- **Sufficient rank in** Γ so that the Brownian shocks relevant for pricing can be recovered from the state history; this makes the Markov state X "rich enough" to describe valuation.

These assumptions let us write the generator in closed form and use martingale-based changes of measure freely.

141.2.2 Functionals and càdlàg paths

We need a name for "a process that records something about the history of X".

This includes, for example, a stochastic discount factor or a cumulated return.

Definition 141.2.1 (Functional)

A **functional** is a real-valued adapted process $\{M_t : t \geq 0\}$ constructed from the history of X, so that M_t is $\mathcal{F}_t$-measurable for each t.

We will always work with the **càdlàg** version of a functional, the French acronym for "right-continuous with left limits".

Concretely, for almost every sample path ω,

$$\lim_{s \downarrow t} M_s(\omega) = M_t(\omega) \qquad \text{for all } t \geq 0,$$

and the left limit $M_{t-}(\omega) := \lim_{s \uparrow t} M_s(\omega)$ exists and is finite for all $t > 0$.

Thus, paths can jump, but each jump $\Delta M_t := M_t - M_{t-}$ occurs instantaneously.

At the jump time t, the value is the post-jump value, not the pre-jump value.

Note

Why is the càdlàg property worth insisting on?

Because we will later want to (i) integrate functionals against time, (ii) apply optional stopping arguments, and (iii) take limits like $\lim_{t \to \infty} t^{-1} \log M_t$.

All three operations need the joint measurability in (ω, t) that càdlàg paths give us automatically.

141.2.3 Multiplicativity

The central restriction we impose on the functional M is that it is *multiplicative across time*.

Definition 141.2.2 (Multiplicative Functional)

A functional $\{M_t : t \geq 0\}$ is **multiplicative** if $M_0 = 1$ and

$$M_{t+u} = M_u(\theta_t)\, M_t, \qquad t, u \geq 0, \tag{141.2}$$

where θ_t shifts the underlying Markov path forward by t units.

Why is this a useful condition to require?

Think of $M_t = S_t$, a stochastic discount factor.

The date-0 value of a date-t payoff Π_t is $\mathbb{E}[S_t \Pi_t \mid \mathcal{F}_0]$.

If we instead buy this payoff at intermediate date τ, its date-τ price must be

$$\mathbb{E}\left[\frac{S_t}{S_\tau}\Pi_t \,\middle|\, \mathcal{F}_\tau\right].$$

For the price to depend only on the current Markov state X_τ (and not on the entire history up to τ), the ratio S_t/S_τ must be a function only of the Markov path *after* τ.

Thus, $S_{\tau+u}/S_\tau = S_u(\theta_\tau)$, which is exactly (141.2).

So multiplicativity is the Markov form of the law of one price.

The same structural property is then carried over to stochastic growth and cumulated returns.

141.2.4 Additive functionals

It is often easier to work with the *log* of a strictly positive multiplicative functional.

If $M_t > 0$, we can write $M_t = \exp(A_t)$, and the multiplicative property (141.2) becomes the corresponding additive property of A.

Definition 141.2.3 (Additive Functional)

A functional $\{A_t : t \geq 0\}$ is **additive** if $A_0 = 0$ and

$$A_{t+u} = A_u(\theta_t) + A_t.$$

So exponentials of additive functionals are exactly the strictly positive multiplicative functionals.

In our jump-diffusion setting, a useful parameterization is

$$A_t = \int_0^t \beta(X_s)\, ds + \int_0^t \gamma(X_{s-})^\top dB_s + \sum_{0 \leq s \leq t} \kappa(X_s, X_{s-}), \tag{141.3}$$

where the three functions (β, γ, κ) play the following roles:

- $\beta(x)$ is a **state-dependent drift** rate (e.g., a pure discount factor with short rate $r(x)$ has $\beta(x) = -r(x)$).
- $\gamma(x)$ is a **Brownian loading**, controlling how Brownian shocks feed into A at state x.
- $\kappa(y, x)$ is a **jump amplitude** that fires whenever X jumps from x to y.

For everything to be well defined, we impose the natural integrability conditions

$$\int_0^t |\beta(X_s)|\, ds < \infty, \qquad \int_0^t \|\gamma(X_s)\|^2\, ds < \infty,$$

together with $\kappa(x, x) = 0$ (no jump if the state doesn't change) and

$$\int \exp[\kappa(y, x)]\, \eta(dy \mid x) < \infty.$$

This parameterization is rich enough for everything we do in this lecture, though additive functionals can also include things like occupation times.

141.3 Semigroups

A multiplicative functional M together with the Markov process X defines a **valuation operator** for each horizon t:

$$\mathbb{M}_t \psi(x) = \mathbb{E}\left[M_t \psi(X_t) \mid X_0 = x\right].$$

You should read $\mathbb{M}_t \psi(x)$ as "the date-0 value, starting from state x, of a date-t payoff $\psi(X_t)$", weighted by M_t.

The family of operators $\{\mathbb{M}_t\}_{t \geq 0}$ has a key compositional structure, the *semigroup property*.

Definition 141.3.1 (One-Parameter Semigroup)

A family of linear operators $\{T_t : t \geq 0\}$ is a **one-parameter semigroup** if $T_0 = I$ and $T_{t+s} = T_t T_s$ for all $s, t \geq 0$.

It is **positive** if $T_t \psi \geq 0$ whenever $\psi \geq 0$.

Definition 141.3.2 (Multiplicative Semigroup)

Given a multiplicative functional M, the associated **multiplicative semigroup** is the family of operators

$$\mathbb{M}_t\psi(x) = \mathbb{E}\left[M_t\psi(X_t) \mid X_0 = x\right]. \tag{141.4}$$

The semigroup identity $\mathbb{M}_{t+u} = \mathbb{M}_t\mathbb{M}_u$ follows from iterated expectations and the multiplicative property of M.

Economically, this is the *Markov law of iterated values*: to value a date-$(t+u)$ payoff today, we can either

- discount it directly from $t+u$ back to 0 in one step (apply $\mathbb{M}_{t+u}$), or
- first discount it from $t+u$ back to t (apply $\mathbb{M}_u$), then discount the resulting date-t price back to 0 (apply $\mathbb{M}_t$).

The semigroup identity says these two procedures give the same answer.

This is the operator-level version of the intertemporal consistency that rules out arbitrage across horizons.

Four positive multiplicative functionals will appear often below.

Symbol	Object	Semigroup
S	stochastic discount factor	$\{\mathbb{S}_t\}$
V	cumulated return on an asset	$\{\mathbb{V}_t\}$
G	stochastic growth in cash flows	$\{\mathbb{G}_t\}$
$Q = GS$	valuation of growing cash flows	$\{\mathbb{Q}_t\}$

The first three are primitives.

The fourth combines discounting and growth to value cash flows that grow stochastically over time.

Definition 141.3.3 (Stochastic Discount Factor)

A **stochastic discount factor** S is a positive multiplicative functional for which $\mathbb{E}[S_tZ_t \mid X_0 = x]$ is the date-0 value of an $\mathcal{F}_t$-measurable payoff Z_t.

Definition 141.3.4 (Valuation Functional)

Given a stochastic discount factor S, a **valuation functional** V is a positive multiplicative functional such that $\{V_tS_t : t \geq 0\}$ is a martingale.

Definition 141.3.5 (Stochastic Growth Functional)

A **stochastic growth functional** G is a positive multiplicative functional, interpreted as a multiplicative growth factor applied to a date-0 cash flow.

Consider a cash flow $D_t = D_0\, G_t\, \psi(X_t)$, where $D_0 > 0$ is the date-0 level and ψ is a Borel state-payoff function.

Discounting with S and growing with G, its date-0 value is

$$D_0\, \mathbb{Q}_t\psi(X_0), \qquad \mathbb{Q}_t\psi(x) = \mathbb{E}\left[G_tS_t\psi(X_t) \mid X_0 = x\right].$$

Definition 141.3.6 (Cash-Flow Valuation Semigroup)

The **cash-flow valuation semigroup** is the multiplicative semigroup generated by $Q = GS$.

The long-horizon behavior of $\mathbb{Q}_t$ is the central object of the lecture: it tells us how current prices value cash-flow growth risk that materializes far in the future.

Note

The factorization $D_t = D_0 G_t \psi(X_t)$ is not unique.

For any positive function φ,

$$D_t = D_0 \left[G_t \frac{\varphi(X_t)}{\varphi(X_0)} \right] \left[\frac{\psi(X_t)\varphi(X_0)}{\varphi(X_t)} \right],$$

so a transient state-dependent factor can be shuffled between G and ψ without changing D_t.

We resolve this indeterminacy by normalizing the growth component so that its permanent part is a martingale: $G_t = \exp(\delta t)\widehat{G}_t$, with $\widehat{G}$ a martingale and δ a constant trend.

The eigenfunction construction below will tell us exactly which martingale to pick.

141.3.1 Local pricing restriction

Before tackling long horizons, it is good to understand valuation at the *short* end.

That is the standard instantaneous risk-return relation.

This will give us a benchmark against which to compare long-run risk prices.

For a textbook discrete-time treatment of the same SDF-based asset-pricing ideas, see Elementary Asset Pricing Theory.

For an estimation perspective on Euler-equation-based asset pricing, see *Estimating Euler Equations by Generalized Method of Moments*.

The key starting point is that a valuation functional V must satisfy the no-arbitrage requirement that VS is a martingale (*Definition 141.3.4*).

We parameterize the stochastic discount factor S and valuation functional V as additive functionals with coefficients $(\beta^s, \gamma^s, \kappa^s)$ and $(\beta^v, \gamma^v, \kappa^v)$ respectively, in the notation of (141.3).

For a generic positive multiplicative functional $M = \exp(A)$ with parameters (β, γ, κ), applying Itô's formula to $\exp(A_t)$ and requiring the drift of dM_t/M_{t-} to vanish gives the **local martingale restriction**:

$$\beta(x) + \frac{\gamma(x)^\top \gamma(x)}{2} + \int \left(\exp[\kappa(y, x)] - 1 \right) \eta(dy \mid x) = 0. \tag{141.5}$$

The three terms correspond to:

- the drift of A itself,
- the Itô correction from the Brownian part (because $M = e^A$ picks up a quadratic-variation contribution), and
- the compensated jumps of M at the multiplier $\exp[\kappa(y, x)]$.

We ask you to verify this in Exercise 141.11.4.

Applying this to VS, whose parameters add: $(\beta^v + \beta^s, \gamma^v + \gamma^s, \kappa^v + \kappa^s)$, gives the **local pricing restriction**

$$\beta^v + \beta^s = -\frac{\|\gamma^v + \gamma^s\|^2}{2} - \int (\exp[\kappa^v(y,\cdot) + \kappa^s(y,\cdot)] - 1)\, \eta(dy \mid \cdot). \tag{141.6}$$

This determines the drift β^v of any candidate valuation functional in terms of its Brownian and jump exposures.

To turn this into an *expected return*, note that the expected net rate of return on V is

$$\epsilon^v = \beta^v + \frac{\|\gamma^v\|^2}{2} + \int (\exp[\kappa^v(y,\cdot)] - 1)\, \eta(dy \mid \cdot).$$

Combining with (141.6) gives

$$\begin{aligned} \epsilon^v = &-\beta^s - \gamma^v \cdot \gamma^s - \frac{\|\gamma^s\|^2}{2} \\ &- \int (\exp[\kappa^v(y,\cdot) + \kappa^s(y,\cdot)] - \exp[\kappa^v(y,\cdot)])\, \eta(dy \mid \cdot). \end{aligned} \tag{141.7}$$

Thus, the Brownian local risk-price vector is $-\gamma^s(x)$, expressed in the same exposure units as $\gamma^v(x)$.

Jump risk is priced through the function κ^s.

This local relation is one end of the term structure of risk prices.

The eigenvalue calculations below describe the other end.

141.4 The generator

So far we have a family of operators $\{\mathbb{M}_t\}_{t\geq 0}$, one for each horizon t.

That is more information than we can use directly.

What we care about is the long-run behavior: how $\mathbb{M}_t\psi$ grows as $t \to \infty$.

The **generator** $\mathbb{A}$ compresses the entire semigroup into one time-independent operator on the state space.

It records the *instantaneous* rate of change of $M_t\psi(X_t)$, and its eigenvalues drive the long-run growth rate of $\mathbb{M}_t$.

This lets us turn an asymptotic question about a family of operators into a single eigenvalue problem.

141.4.1 Discrete-time intuition

The role of the generator is easiest to see when time is discrete, so we build the picture there once and then carry it over.

Let X_n be Markov and M_n a multiplicative functional, and define the one-period valuation operator

$$K\psi(x) = \mathbb{E}\left[M_1\psi(X_1) \mid X_0 = x\right].$$

Iterating gives the n-period operator K^n.

This parallels the logic of a transition matrix: just as P^n gives n-step probabilities, K^n weights n-period payoffs, with M_1 folded in at each step.

So one local object, K, controls the entire horizon-indexed family.

It also controls long-run growth: if $K\phi = \lambda\phi$, then

$$K^n\phi = \lambda^n\phi,$$

so λ is the long-run geometric growth (or decay) rate and ϕ is the long-run dependence on the current state.

The same operator gives the predictable rate of change of the weighted payoff.

By multiplicativity,

$$\mathbb{E}\left[M_{n+1}\psi(X_{n+1}) - M_n\psi(X_n) \mid \mathcal{F}_n\right] = M_n(K\psi - \psi)(X_n),$$

so subtracting the cumulative predictable change from the total change leaves a martingale:

$$M_n\psi(X_n) - \psi(X_0) - \sum_{j=0}^{n-1} M_j(K\psi - \psi)(X_j).$$

Here $K - I$ is the *local* rate of change of $M_n\psi(X_n)$, and through K^n it also controls long-run growth.

141.4.2 From discrete to continuous time

Continuous time keeps the same logic.

The natural counterpart $K - I$ is the **infinitesimal generator** of the semigroup $\{\mathbb{M}_t\}$, the time derivative at zero:

$$\mathbb{A}\psi(x) := \lim_{h \downarrow 0} \frac{\mathbb{M}_h\psi(x) - \psi(x)}{h},$$

so that for small $h > 0$,

$$\mathbb{M}_h\psi(x) \approx \psi(x) + h\,\mathbb{A}\psi(x).$$

Note

When $M \equiv 1$, the multiplicative semigroup reduces to the standard Markov transition semigroup $\mathbb{M}_t\psi(x) = \mathbb{E}[\psi(X_t) \mid X_0 = x]$, and $\mathbb{A}$ becomes the familiar infinitesimal generator $\mathcal{L}$ of X from textbook stochastic calculus.

For general M, $\mathbb{A}$ is the same kind of object but it carries extra terms that encode the instantaneous "yield" of M: discount rates, Brownian risk prices, jump multipliers.

That is exactly what we will see when we write down its closed form in (141.8) below.

Why is this the right object?

If $\mathbb{A}\phi = \rho\phi$, then the candidate

$$\hat{M}_t := \exp(-\rho t)\, M_t\, \frac{\phi(X_t)}{\phi(X_0)}$$

is automatically a *local* martingale.

When $\hat{M}$ is in fact a martingale (the content of Assumption 6.1 in Hansen and Scheinkman [2009]), taking expectations gives the continuous-time analogue of $K^n\phi = \lambda^n\phi$:

$$\mathbb{M}_t\phi = \exp(\rho t)\,\phi.$$

Without that upgrade we get only the supermartingale inequality $\mathbb{M}_t\phi \le \exp(\rho t)\,\phi$, which we will revisit below.

So the long-run behavior of $\mathbb{M}_t$ is encoded in an eigenvalue problem for the local operator $\mathbb{A}$, together with the martingale property of $\hat{M}$.

141.4.3 Extended generator

The limit definition above needs a qualification.

To make the limit $h \downarrow 0$ rigorous, the textbook definition requires $(\mathbb{M}_h\psi - \psi)/h$ to converge to $\mathbb{A}\psi$ in a chosen norm, typically uniform convergence over bounded continuous functions, where the semigroup is a contraction.

But the functions we care about most are the principal eigenfunctions ϕ solving $\mathbb{A}\phi = \rho\phi$, and these typically *grow* with the state X (in the affine-Gaussian benchmark, ϕ is exponential-affine in x).

So they do not lie in this space, and the limit need not converge for them.

Hansen and Scheinkman [2009] sidestep this by characterizing $\mathbb{A}$ through a *Doob-Meyer style* semimartingale decomposition of $M_t\psi(X_t)$, a pathwise condition that does not require any norm.

This is the continuous-time counterpart to writing $K - I$ as the predictable rate of change of $M_n\psi(X_n)$ in discrete time.

The resulting **extended generator** admits unbounded ψ, has a strictly larger domain than the textbook generator, and agrees with it where both are defined.

Concretely:

Definition 141.4.1 (Extended Generator)

Fix a Borel function ψ, and look for a second Borel function χ that will play the role of "the instantaneous rate of change of $M_t\psi(X_t)$ at the current state".

We ask whether there exists χ such that

$$N_t = M_t\psi(X_t) - \psi(X_0) - \int_0^t M_s\chi(X_s)\,ds$$

is a local martingale.

We then say ψ lies in the domain of the **extended generator** $\mathbb{A}$ of M, and we *define* the operator by

$$\mathbb{A}\psi := \chi.$$

The three terms play the same roles as in discrete time:

- $M_t\psi(X_t) - \psi(X_0)$ is the total change in the weighted payoff over $[0, t]$,
- $\int_0^t M_s\chi(X_s)\,ds$ accumulates the expected rate of change along the path,
- N_t is the residual.

Requiring N_t to be a local martingale pins down $\chi = \mathbb{A}\psi$ as the instantaneous expected rate of change of $M_t\psi(X_t)$ at the current state.

- When $M \equiv 1$, the definition reduces to Dynkin's formula for the standard Markov generator $\mathcal{L}\psi(x) = \lim_{t\downarrow 0} t^{-1}[\mathbb{E}\psi(X_t) - \psi(x)]$.
- When X is a jump diffusion, Itô's formula applied to $M_t\psi(X_t)$ produces the closed-form expression for $\mathbb{A}\psi$ in (141.8) below.

141.4.4 A closed form for jump diffusions

For the jump-diffusion setting introduced above, we can compute $\mathbb{A}$ explicitly by applying Itô's formula to $M_t\phi(X_t)$.

Suppose the continuous part of X satisfies

$$dX_t^c = \xi(X_t)dt + \Gamma(X_t)dB_t$$

with diffusion matrix $\Sigma = \Gamma\Gamma^\top$, and let $\eta(dy \mid x)$ denote the jump compensator.

If $M = \exp(A)$ is parameterized by (β, γ, κ) as in (141.3), then for smooth ϕ,

$$\begin{aligned}\mathbb{A}\phi(x) = &\left[\xi(x) + \Gamma(x)\gamma(x)\right]^\top \frac{\partial\phi(x)}{\partial x} \\ &+ \frac{1}{2}\,\text{trace}\left[\Sigma(x)\frac{\partial^2\phi(x)}{\partial x\partial x^\top}\right] \\ &+ \int [\phi(y) - \phi(x)]\exp[\kappa(y,x)]\eta(dy \mid x) \\ &+ \left[\beta(x) + \frac{\gamma(x)^\top\gamma(x)}{2} + \int \left(\exp[\kappa(y,x)] - 1\right)\eta(dy \mid x)\right]\phi(x).\end{aligned} \tag{141.8}$$

The four terms have transparent interpretations:

1. The first term is the standard Markov drift, modified by $\Gamma\gamma$, a *covariance correction* between the Brownian shocks driving X and those driving M.
2. The second is the standard diffusion (Itô) term.
3. The third integrates ϕ against the jump-compensated transition rates, reweighted by the jump multiplier $\exp[\kappa(y,x)]$.
4. The fourth is a multiplicative *yield-like* term that multiplies $\phi(x)$ itself and combines the drift of M, the Brownian Itô correction, and the compensated jumps.

Note

When $M = S$ is a stochastic discount factor, the local risk prices are spread across the four pieces of the generator rather than concentrated in the level term.

- Brownian factor prices enter through the drift modification $\xi \to \xi + \Gamma\gamma$ in the first line.
- Jump-risk prices enter through the tilted jump measure $\eta \to \exp[\kappa]\,\eta$ in the third line.
- The level term in the fourth line contains the instantaneous risk-free rate component (the *short end* of the term structure) plus the jump-compensator contribution $\int(\exp[\kappa] - 1)\,\eta$.

We ask readers to verify the derivation of (141.8) in Exercise 141.11.5.

We will apply this formula directly in the affine-diffusion example below.

141.5 Principal eigenfunctions

We now arrive at the central technical question of the lecture:

> Which positive payoffs grow at a constant proportional rate under the valuation semigroup?

The answer, when it exists, is a positive eigenfunction of the generator.

Definition 141.5.1 (Eigenfunction of the Extended Generator)

A Borel function ϕ is an **eigenfunction** of $\mathbb{A}$ with eigenvalue ρ if

$$\mathbb{A}\phi = \rho\phi. \tag{141.9}$$

A **principal eigenfunction** is an eigenfunction ϕ that is strictly positive on the state space: $\phi(x) > 0$ for all $x \in \mathcal{D}_0$.

The strict-positivity requirement matters because ϕ will appear in denominators throughout: it has to be safe to divide by it.

141.5.1 From eigenfunction to factorization

Why does an eigenfunction of $\mathbb{A}$ give us the multiplicative factorization (141.1)?

The discrete-time analogy points the way.

If $K\phi = \lambda\phi$, then

$$\lambda^{-n} M_n \frac{\phi(X_n)}{\phi(X_0)}$$

is a martingale: the eigenvalue equation absorbs the one-step drift of $M_n\phi(X_n)$ exactly.

In continuous time, λ^n becomes $\exp(\rho t)$, and the candidate martingale is

$$\hat{M}_t = \exp(-\rho t)\, M_t \frac{\phi(X_t)}{\phi(X_0)}. \tag{141.10}$$

Let's check that $\hat{M}$ is indeed a local martingale.

By the definition of $\mathbb{A}$, the semimartingale $Z_t := M_t\phi(X_t)$ admits the decomposition

$$dZ_t = M_t\, \mathbb{A}\phi(X_t)\, dt + dN_t,$$

where N is a local martingale.

The eigenvalue equation $\mathbb{A}\phi = \rho\phi$ replaces the drift by $\rho Z_t\, dt$, and integration by parts gives

$$d(\exp(-\rho t)Z_t) = \exp(-\rho t)\left[dZ_t - \rho Z_t\, dt\right] = \exp(-\rho t)\, dN_t.$$

So $\exp(-\rho t)Z_t$ is a local martingale, and dividing by $\phi(X_0)$ shows the same for $\hat{M}$.

Rearranging (141.10) for M_t recovers the factorization (141.1), with $\hat{M}$ in the role of the martingale component.

Definition 141.5.2 (Martingale Component and Twisted Measure)

When $\hat{M}$ is a martingale (not merely a local martingale), it is the **martingale component** associated with (ρ, ϕ) and defines the **twisted probability measure** by weighting date-t events with $\hat{M}_t$:

$$\widehat{\Pr}(F \mid X_0 = x) = \mathbb{E}[\hat{M}_t \mathbf{1}_F \mid X_0 = x], \qquad F \in \mathcal{F}_t.$$

The verification establishes only that $\hat{M}$ is a *local* martingale, but the definition above (and the change-of-measure interpretation of $\hat{M}$) require it to be a martingale.

We close this gap by adopting Assumption 6.1 of Hansen and Scheinkman [2009]:

> The local martingale $\hat{M}$ defined in (141.10) is a martingale, i.e. $\mathbb{E}[\hat{M}_t \mid X_0 = x] = 1$ for every $t \geq 0$ and x.

We carry this assumption from here on.

Even without closing the gap we still get one-sided control.

Since $\hat{M}$ is nonnegative with $\hat{M}_0 = 1$, it is a supermartingale, so $\mathbb{E}[\hat{M}_t \mid X_0 = x] \leq 1$.

Taking expectations in (141.10),

$$\mathbb{M}_t \phi(x) \leq \exp(\rho t)\, \phi(x),$$

so ρ is at least an upper bound on the long-run growth rate of $\mathbb{M}_t \phi$.

When $\hat{M}$ is in fact a martingale, $\mathbb{E}[\hat{M}_t] = 1$, the inequality becomes an equality, and the local condition $\mathbb{A}\phi = \rho\phi$ lifts to the semigroup eigenvalue equation

$$\mathbb{M}_t \phi = \exp(\rho t)\, \phi, \qquad t \geq 0. \tag{141.11}$$

141.5.2 Stability of the twisted process

We now have a factorization (141.1) for *any* principal eigenfunction.

But for (ρ, ϕ) to describe **long-run** behavior of $\mathbb{M}_t$ the twisted process must settle into a stationary regime as $t \to \infty$.

If it doesn't, the transient factor $\phi(X_0)/\phi(X_t)$ will not vanish, and we cannot read off the asymptotics from ρ alone.

We need three conditions, each ruling out a specific way in which the twisted process can fail to settle down.

Let $\hat{E}$ and $\widehat{\Pr}$ denote expectation and probability under the twisted measure, and let $\hat{\mathbb{A}}$ be the generator of X under that measure.

Condition 1: a stationary distribution exists.

Definition 141.5.3 (Stationary Distribution of the Twisted Process)

A probability measure $\hat{\varsigma}$ on $\mathcal{D}_0$ is a **stationary distribution** for the twisted process if

$$\int \hat{\mathbb{A}}\psi \, d\hat{\varsigma} = 0$$

for every ψ in the L^∞ domain of $\hat{\mathbb{A}}$.

$\hat{\varsigma}$ is the candidate long-run distribution.

Without it, the twisted process has no steady state, and the long-run limit cannot be expressed as a state-space integral.

Condition 2: every important region is reachable.

Definition 141.5.4 (Irreducible Skeleton)

The discretely sampled skeleton $\{X_{\Delta j} : j = 0, 1, ...\}$ is **irreducible** relative to $\hat{\varsigma}$ if there exists $\Delta > 0$ such that, for every Borel set $\Lambda \subseteq \mathcal{D}_0$ with $\hat{\varsigma}(\Lambda) > 0$,

$$\hat{E}\left[\sum_{j=0}^{\infty} \mathbf{1}_{\{X_{\Delta j} \in \Lambda\}} \,\middle|\, X_0 = x\right] > 0 \qquad \text{for all } x \in \mathcal{D}_0.$$

Without it, the long-run distribution could depend on the starting state; different basins of attraction would give different limits.

Condition 3: every important region is visited infinitely often.

Definition 141.5.5 (Harris Recurrence)

The process X is **Harris recurrent** under the twisted measure if, for every Borel set $\Lambda \subseteq \mathcal{D}_0$ with $\hat{\varsigma}(\Lambda) > 0$,

$$\widehat{\Pr}\left\{\int_0^{\infty} \mathbf{1}_{\{X_t \in \Lambda\}}\, dt = \infty \,\middle|\, X_0 = x\right\} = 1 \qquad \text{for all } x \in \mathcal{D}_0.$$

Reachability (Condition 2) is not enough.

A set may be reachable but visited with vanishingly small probability, so time averages fail to converge to $\hat{\varsigma}$-averages.

Harris recurrence is the continuous-state counterpart to a "recurrent state" in a finite chain.

Collecting the three conditions:

Definition 141.5.6 (Stochastically Stable Twisted Process)

The $\hat{M}$-twisted Markov process is **stochastically stable** if it has a stationary distribution $\hat{\varsigma}$, the skeleton $\{X_{\Delta j}\}$ is irreducible relative to $\hat{\varsigma}$, and X is Harris recurrent under the twisted measure.

141.5.3 The long-run approximation

Under the martingale condition for $\hat{M}$, strict positivity of M, and stochastic stability, the long-run limit takes a clean form:

$$\lim_{t \to \infty} \exp(-\rho t)\, \mathbb{M}_t \psi = \phi \int \frac{\psi}{\phi}\, d\hat{\varsigma}. \tag{141.12}$$

Read this as follows:

- The factor $\exp(\rho t)$ captures the exponential growth or decay of the semigroup; once we divide it out, what remains has a finite limit.
- The state dependence in that limit is *entirely* captured by $\phi(x)$.

- The scalar $\int (\psi/\phi)\, d\hat{\varsigma}$ is the **long-run intensity** of the payoff ψ, weighted by $1/\phi$ and averaged against the twisted stationary distribution.

The two precise cases of Proposition 7.1 of Hansen and Scheinkman [2009] are:

- **Sampled grid, $\hat{\varsigma}$-almost every starting state.** Fix any $\Delta > 0$ and assume $\int |\psi|/\phi \, d\hat{\varsigma} < \infty$; then the limit holds along $t = \Delta j$ for $\hat{\varsigma}$-almost every $x \in \mathcal{D}_0$.
- **Continuous t, every starting state.** If ψ/ϕ is bounded then the limit holds for every $x \in \mathcal{D}_0$ and every $t \to \infty$, with no continuity assumption on ψ.

Note

There can be more than one positive eigenfunction of $\mathbb{A}$ for which the associated $\hat{M}$ is a martingale, possibly with different eigenvalues.

Stochastic stability selects a particular one.

By Proposition 7.2 of Hansen and Scheinkman [2009], the stable eigenfunction's eigenvalue ρ is the **smallest** among all eigenvalues of $\mathbb{A}$ that admit a strictly positive eigenfunction satisfying the maintained conditions.

Any other positive eigenfunction sharing this ρ is proportional to ϕ $\hat{\varsigma}$-almost surely.

The finite-state section below states the same selection in the more familiar Perron-Frobenius language, calling ρ the eigenvalue with the *largest* real part among all eigenvalues of A.

These two descriptions identify the same eigenvalue because in the irreducible finite-state case only one eigenvalue of A admits a strictly positive eigenvector.

In the affine example below the two sets pull apart, since the quadratic (141.18) has two roots c_f that both give strictly positive exponential-affine eigenfunctions with distinct eigenvalues.

Stochastic stability picks the smaller ρ by rejecting the root that produces an explosive twisted process.

141.6 A finite-state Markov chain

We now apply the framework to a concrete example.

We start with the simplest case: a finite-state Markov chain.

For background on finite Markov chains in discrete time, see *Finite Markov Chains*.

For the asset-pricing applications of finite-state chains that motivate the construction here, see *Asset Pricing: Finite State Models*.

Here, every abstract object collapses to a familiar one:

Abstract object	Finite-state version
Markov process X	continuous-time chain with intensity matrix U
Generator $\mathbb{A}$	a matrix A
Semigroup $\mathbb{M}_t$	matrix exponential $\exp(tA)$
Principal eigenfunction ϕ	Perron right eigenvector
Principal eigenvalue ρ	dominant real eigenvalue of A
Stationary distribution $\hat{\varsigma}$	left eigenvector of twisted generator

So the long-run analysis reduces to Perron-Frobenius theory.

(For an introduction to the Perron-Frobenius theorem, see The Perron-Frobenius Theorem.)

141.6.1 Setup

Let X take values in $\{x_1, \dots, x_N\}$ and let U be its intensity matrix.

Definition 141.6.1 (Intensity Matrix)

An **intensity matrix** U for a finite-state continuous-time Markov chain satisfies $u_{ij} \geq 0$ for $i \neq j$ and $\sum_j u_{ij} = 0$ for each state i.

Off-diagonal entries u_{ij} are the jump rates from state i to state j; the diagonal entry $u_{ii} = -\sum_{j \neq i} u_{ij}$ is minus the exit rate from state i.

Let the multiplicative functional have

- a **discount or decay rate** r_i in state i (the analogue of the drift $\beta(X_s)$ in the additive parameterization), and
- a **jump multiplier** $\exp[\kappa(x_j, x_i)]$ that fires whenever the state jumps from i to j.

In code below, `κ[j, i]` means $\kappa(x_j, x_i)$.

Then the generator matrix A for the multiplicative semigroup is

$$a_{ij} = \begin{cases} u_{ii} - r_i, & i = j, \\ u_{ij} \exp[\kappa(x_j, x_i)], & i \neq j. \end{cases} \tag{141.13}$$

The off-diagonal entries are the chain's jump rates *weighted* by the jump multipliers; the diagonal entries combine the exit rate with the in-state decay rate.

The semigroup is then just $\mathbb{M}_t = \exp(tA)$.

For an irreducible chain with strictly positive jump multipliers, the principal eigenvalue ρ is the unique real eigenvalue of A with largest real part, and the associated right eigenvector is strictly positive.

The twisted generator under the principal eigenpair (ρ, ϕ) is

$$\hat{A} = D_\phi^{-1} A D_\phi - \rho I,$$

where $D_\phi = \operatorname{diag}(\phi)$.

The row sums of $\hat{A}$ vanish, so $\hat{A}$ is itself a valid intensity matrix.

The stationary distribution $\hat{\varsigma}$ solves $\hat{\varsigma}^\top \hat{A} = 0$.

The helper functions below implement these three calculations.

```
def build_generator(U, r, κ):
    """Generator matrix for a finite-state multiplicative semigroup.

    κ[j, i] is the log jump multiplier for a transition i -> j.
    """
    U = np.asarray(U, dtype=float)
    r = np.asarray(r, dtype=float)
    κ = np.asarray(κ, dtype=float)

    A = U * np.exp(κ.T)
    np.fill_diagonal(A, np.diag(U) - r)
```

(continues on next page)

(continued from previous page)

```
    return A


def principal_eigenpair(A):
    """Perron eigenpair with a positive right eigenvector."""
    vals, vecs = eig(A)
    idx = np.argmax(vals.real)

    ρ = vals[idx].real
    ψ = vecs[:, idx].real

    if ψ.sum() < 0:
        ψ = -ψ

    if np.any(ψ <= 0):
        if np.min(ψ) > -1e-10:
            ψ = np.maximum(ψ, 0)
        else:
            raise ValueError(
                "Dominant eigenvector is not strictly positive.")

    ψ = ψ / ψ.mean()
    return ρ, ψ


def twisted_generator(A, ρ, ψ):
    """Generator under the twisted measure."""
    D = np.diag(ψ)
    D_inv = np.diag(1 / ψ)
    return D_inv @ A @ D - ρ * np.eye(len(ψ))


def stationary_distribution(Q):
    """Stationary distribution of a finite-state generator."""
    vals, vecs = eig(Q.T)
    idx = np.argmin(np.abs(vals))
    π = vecs[:, idx].real

    if π.sum() < 0:
        π = -π

    π = np.maximum(π, 0)
    return π / π.sum()
```

141.6.2 Two states

Consider a boom-recession economy.

State 1 is a *boom* (higher short rate $r_1 = 0.05$, switching to recession at rate $\lambda_1 = 0.30$).

State 2 is a *recession* (lower short rate $r_2 = 0.02$, switching to boom at rate $\lambda_2 = 0.50$).

For now we set the log jump multipliers κ to zero — equivalently, the jump multipliers $\exp[\kappa]$ to one — so the SDF only changes continuously through the in-state decay rates.

```
λ_1 = 0.30
λ_2 = 0.50

U = np.array([[-λ_1,  λ_1],
              [ λ_2, -λ_2]])

r = np.array([0.05, 0.02])
κ = np.zeros((2, 2))

A = build_generator(U, r, κ)
ρ, ψ = principal_eigenpair(A)

print("A =")
print(np.round(A, 4))
print(f"\nρ = {ρ:.6f}")
print(f"ψ = {ψ}")
print(f"long-run zero-coupon yield = {-ρ:.4f}")
```

```
A =
[[-0.35  0.3 ]
 [ 0.5  -0.52]]

ρ = -0.038484
ψ = [0.98116799 1.01883201]
long-run zero-coupon yield = 0.0385
```

Note that $-\rho$ is the asymptotic yield on a zero-coupon bond: from (141.1), the date-0 price of a long zero-coupon bond decays like $\exp(\rho t)$, so its yield is $-\rho$.

Let's verify the semigroup eigenvalue equation $\mathbb{M}_t \phi = \exp(\rho t)\phi$ numerically.

```
for t in [1.0, 5.0, 25.0]:
    lhs = expm(t * A) @ ψ
    rhs = np.exp(ρ * t) * ψ
    err = np.max(np.abs(lhs - rhs))
    print(f"t = {t:4.1f}, error = {err:.2e}")
```

```
t =  1.0, error = 1.11e-16
t =  5.0, error = 1.11e-16
t = 25.0, error = 3.33e-16
```

The error is at machine precision for every t, so the equation holds exactly up to floating-point noise from the eigendecomposition.

Next we compute the twisted generator $\hat{A}$ and the stationary distribution $\hat{\varsigma}$ of the chain under the twisted measure.

This is the candidate long-run distribution that appears in the long-run limit (141.12).

```
A_hat = twisted_generator(A, ρ, ψ)
ς_hat = stationary_distribution(A_hat)

print("twisted generator row sums:")
print(np.round(A_hat.sum(axis=1), 12))

print("\ntwisted stationary distribution:")
print(f"  boom      {ς_hat[0]:.4f}")
print(f"  recession {ς_hat[1]:.4f}")
```

```
twisted generator row sums:
[-0.  0.]

twisted stationary distribution:
  boom      0.6072
  recession 0.3928
```

This twisted stationary distribution (not the original chain's stationary distribution) is what determines long-horizon valuations.

It differs from the stationary distribution of the original chain because the eigenfunction ϕ reweights states by how persistently they affect the multiplicative functional.

For any payoff function ψ, the long-run limit (141.12) is the vector

$$\phi \sum_i \frac{\psi_i}{\phi_i} \hat{\varsigma}_i.$$

Let's check that the rescaled semigroup converges to this limit as t grows.

```
ψ = np.array([1.0, 2.0])
limit = ρ * np.sum((ψ / ϕ) * ς_hat)

for t in [1, 5, 20, 80]:
    approx = np.exp(-ρ * t) * expm(t * A) @ ψ
    print(f"t = {t:2d}, normalized value = {approx}")

print("\nlimit =", limit)
```

```
t =  1, normalized value = [1.19917652 1.68031101]
t =  5, normalized value = [1.35687434 1.42719756]
t = 20, normalized value = [1.36377325 1.41612444]
t = 80, normalized value = [1.3637733  1.41612436]

limit = [1.3637733  1.41612436]
```

The numerical values converge to the limit; the next plot shows the same convergence pictorially.

```
t_grid = np.linspace(0.01, 80, 400)
rescaled = np.array([np.exp(-ρ * t) * expm(t * A) @ ψ for t in t_grid])

fig, ax = plt.subplots()
ax.plot(t_grid, rescaled[:, 0], lw=2, label="starting from boom")
ax.plot(t_grid, rescaled[:, 1], lw=2, label="starting from recession")
ax.axhline(limit[0], color="C0", ls="--", lw=1, alpha=0.7)
ax.axhline(limit[1], color="C1", ls="--", lw=1, alpha=0.7)
ax.set_xlabel("$t$")
ax.set_ylabel(r"$e^{-\rho t}\,(\mathbb{M}_t \psi)(x)$")
ax.legend()
plt.show()
```

The two solid curves trace the rescaled values starting from each state, and the dashed horizontal lines mark the limits predicted by (141.12).

Both curves settle onto their predicted limits, confirming that the long-run behavior depends on the starting state only through ϕ.

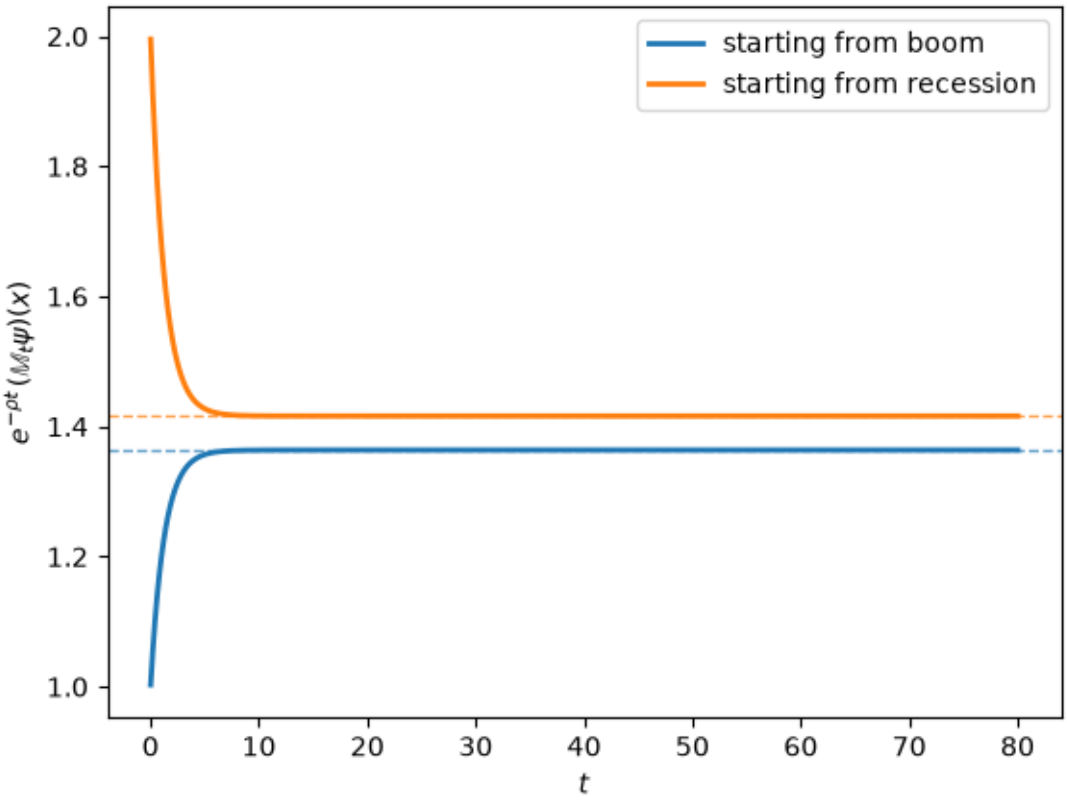


Fig. 141.1: Convergence of the rescaled semigroup to its long-run limit

Note

The asymptotic exponential rate of convergence is governed by the gap between the *real part* of the leading eigenvalue ρ and the largest real part among the remaining eigenvalues of A.

For an irreducible Metzler matrix the leading eigenvalue is real and its real part is strictly larger than the others, so this gap is well defined.

Exercise 141.11.3 works through a three-state example where the gap can be checked directly.

141.6.3 Adding jumps

State transitions in this chain are discontinuous, so it is natural to allow the multiplicative functional to jump at the transition times.

These jumps are the analogue of the κ function in the jump-diffusion parameterization.

A natural example arises with a stochastic discount factor that jumps *down* when the economy moves from recession into boom and *up* on the reverse transition.

The matrix `κ_jump` below encodes this.

We use the convention `κ[j, i]` = log jump multiplier of M for the transition $i \to j$, with state index 0 = boom and state index 1 = recession.

```
# recession (1) -> boom (0): SDF jumps down on good news (exp(-0.20))
# boom (0) -> recession (1): SDF jumps up on bad news (exp(+0.30))
κ_jump = np.array([[ 0.0, -0.20],
                   [ 0.30,  0.0]])

A_jump = build_generator(U, r, κ_jump)
ρ_jump, ψ_jump = principal_eigenpair(A_jump)

print(f"ρ without jumps = {ρ:.6f}")
print(f"ρ with jumps    = {ρ_jump:.6f}")
print("\nψ with jumps:")
print(ψ_jump)
```

```
ρ without jumps = -0.038484
ρ with jumps    = -0.019067

ψ with jumps:
[1.10059123 0.89940877]
```

To see how the long-run rate ρ responds to jump risk, we hold the boom-to-recession log multiplier fixed at its calibrated value and trace out ρ as the recession-to-boom log multiplier varies.

The right benchmark to compare against is therefore the eigenvalue at $\kappa(\text{rec} \to \text{boom}) = 0$ with the boom-to-recession jump still switched on, which the black dot marks.

```
κ_boom_to_rec = κ_jump[1, 0]

κ_grid = np.linspace(-0.5, 0.5, 100)
ρ_grid = np.empty_like(κ_grid)

for n, k in enumerate(κ_grid):
    κ_temp = np.array([[0.0, k],
```

(continues on next page)

(continued from previous page)

```
                        [κ_boom_to_rec, 0.0]])
    A_temp = build_generator(U, r, κ_temp)
    ρ_grid[n], _ = principal_eigenpair(A_temp)

ρ_no_up_jump, _ = principal_eigenpair(
    build_generator(U, r, np.array([[0.0, 0.0],
                                    [κ_boom_to_rec, 0.0]]))
)

fig, ax = plt.subplots()
ax.plot(κ_grid, ρ_grid, lw=2)
ax.plot([0.0], [ρ_no_up_jump], "o", color="black",
        label="no recession-to-boom jump")
ax.axvline(0, color="black", ls=":", lw=1)
ax.set_xlabel("jump log multiplier for recession to boom")
ax.set_ylabel("principal eigenvalue")
ax.legend()
plt.show()

print(f"ρ at κ(rec -> boom) = 0 = {ρ_no_up_jump:.6f}")
```

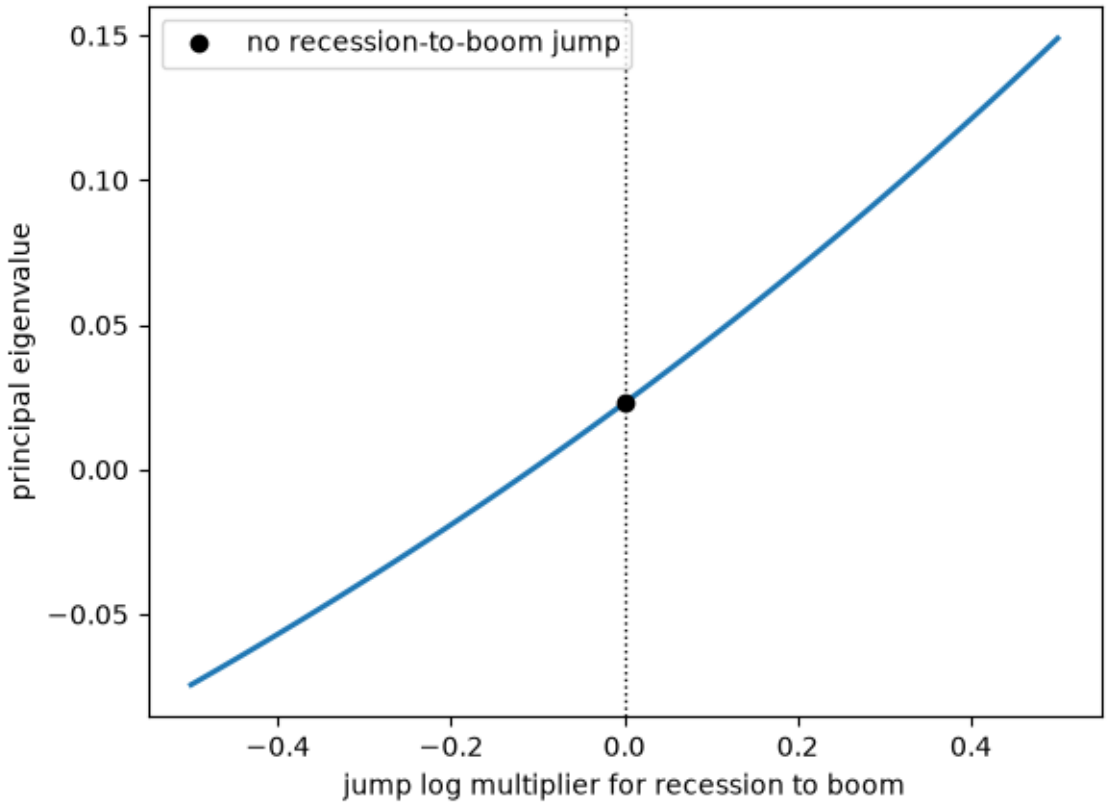


Fig. 141.2: Jumps and the long-run growth rate, holding the boom-to-recession log multiplier fixed

```
ρ at κ(rec -> boom) = 0 = 0.022934
```

The principal eigenvalue is monotonically increasing in the recession-to-boom log multiplier: as that multiplier rises, M jumps less downward (or more upward) on good news, which mechanically pushes ρ up.

The black dot sits at $\rho = 0.022934$.

Switching the recession-to-boom jump on at the calibrated $\kappa = -0.20$ moves ρ down to -0.019067, the value printed above.

The economically sensible region for a stochastic discount factor is to the left of zero, where the log multiplier is negative and M jumps down on good news.

141.7 The affine diffusion example

We now move to a continuous-state model.

We will use a two-factor affine specification that captures the two main empirical features of asset returns:

- **stochastic volatility**, where the dispersion of shocks is itself a state variable, and
- **predictable growth**, where a small, persistent state variable shifts expected growth rates.

This is the kind of state process used in long-run risk models like Bansal and Yaron [2004].

We work with two independent state components: a Feller square-root process X^f (stochastic volatility) and an Ornstein-Uhlenbeck process X^o (predictable growth):

$$\begin{aligned} dX_t^f &= \xi_f(\bar{x}_f - X_t^f)dt + \sqrt{X_t^f}\sigma_f dB_t^f, \\ dX_t^o &= \xi_o(\bar{x}_o - X_t^o)dt + \sigma_o dB_t^o. \end{aligned} \tag{141.14}$$

The parameters $\xi_f, \xi_o > 0$ are mean-reversion speeds, $\bar{x}_f, \bar{x}_o$ are the unconditional means, and σ_f, σ_o are diffusion coefficients.

The OU process X^o is the continuous-time analogue of the AR(1) process studied in AR(1) Processes, and continuous-time linear asset-pricing models in the same family are developed in *Affine Models of Asset Prices*.

We follow a sign convention with $\sigma_o > 0$ and $\sigma_f < 0$: a positive B^f shock then *reduces* volatility, in line with the empirical "leverage effect."

Now consider a multiplicative functional $M = \exp(A)$ with affine parameters:

$$\begin{aligned} A_t = \bar{\beta} t + \int_0^t \beta_f X_s^f \, ds + \int_0^t \beta_o X_s^o \, ds \\ + \int_0^t \sqrt{X_s^f}\, \gamma_f \, dB_s^f + \int_0^t \gamma_o \, dB_s^o. \end{aligned} \tag{141.15}$$

So the drift of A is affine in the state ($\bar{\beta} + \beta_f X^f + \beta_o X^o$), the Brownian loadings are constant in the B^o direction and proportional to $\sqrt{X^f}$ in the B^f direction.

141.7.1 Why exponential-affine eigenfunctions work

When the state is affine and the drift of A is affine, applying the generator to an exponential-affine function $\phi(x^f, x^o) = \exp(c_f x^f + c_o x^o)$ returns another exponential-affine function.

This closure property turns the eigenvalue equation $\mathbb{A}\phi = \rho\phi$ into a small system of algebraic equations in (c_f, c_o, ρ).

Definition 141.7.1 (Exponential-Affine Eigenfunction)

An eigenfunction ϕ of $\mathbb{A}$ on a state space $\mathcal{D}_0 \subseteq \mathbb{R}^n$ is **exponential-affine** if

$$\phi(x) = \exp(c_0 + c^\top x) \qquad \text{for all } x \in \mathcal{D}_0,$$

for some constant $c_0 \in \mathbb{R}$ and vector $c \in \mathbb{R}^n$.

Substituting $\phi(x^f, x^o) = \exp(c_f x^f + c_o x^o)$ into the generator formula (141.8) and matching coefficients of x^f, x^o, and the constant term gives

$$0 = \beta_f + \frac{\gamma_f^2}{2} + c_f(\gamma_f \sigma_f - \xi_f) + c_f^2 \frac{\sigma_f^2}{2}, \tag{141.16}$$

and

$$c_o = \frac{\beta_o}{\xi_o}. \tag{141.17}$$

The two candidate values for c_f are

$$c_f = \frac{\xi_f - \gamma_f \sigma_f \pm \sqrt{(\xi_f - \gamma_f \sigma_f)^2 - \sigma_f^2(2\beta_f + \gamma_f^2)}}{\sigma_f^2}. \tag{141.18}$$

The eigenvalue is then determined by matching the constant term:

$$\rho = \bar{\beta} + \frac{\gamma_o^2}{2} + c_f \xi_f \bar{x}_f + c_o(\xi_o \bar{x}_o + \gamma_o \sigma_o) + c_o^2 \frac{\sigma_o^2}{2}. \tag{141.19}$$

Equation (141.18) gives two candidate values of c_f, and we need to know which one is the principal eigenfunction.

This is where stochastic stability (*Definition 141.5.6*) does real work.

Under the twisted measure, the drift of X^f is

$$\xi_f(\bar{x}_f - x^f) + x^f \sigma_f(\gamma_f + c_f \sigma_f),$$

so the mean-reversion coefficient becomes

$$\xi_f - \sigma_f(\gamma_f + c_f \sigma_f).$$

If this is positive, the twisted square-root process stays stationary; if it is negative, the twisted process is explosive and the eigenfunction is not the long-run-relevant one.

So we **pick the root that keeps the twisted process mean-reverting**, exactly the way stochastic stability selects the principal eigenfunction in the abstract theory.

```
def solve_affine_eigenfunction(params):
    """Solve the exponential-affine eigenvalue problem."""
    ξ_f = params["ξ_f"]
    xbar_f = params["xbar_f"]
    σ_f = params["σ_f"]
    ξ_o = params["ξ_o"]
    xbar_o = params["xbar_o"]
    σ_o = params["σ_o"]
    β_bar = params["β_bar"]
```

(continues on next page)

(continued from previous page)

```
    β_f = params["β_f"]
    β_o = params["β_o"]
    γ_f = params["γ_f"]
    γ_o = params["γ_o"]

    co = β_o / ξ_o

    disc = ((ξ_f - γ_f * σ_f) ** 2
            - σ_f ** 2 * (2 * β_f + γ_f ** 2))

    if disc < 0:
        raise ValueError("No real affine eigenfunction for these parameters.")

    root = np.sqrt(disc)
    cf_roots = (
        ((ξ_f - γ_f * σ_f) - root) / σ_f ** 2,
        ((ξ_f - γ_f * σ_f) + root) / σ_f ** 2
    )

    def mean_reversion(cf):
        return ξ_f - σ_f * (γ_f + cf * σ_f)

    for cf in cf_roots:
        mr = mean_reversion(cf)
        if mr > 0:
            break
    else:
        raise ValueError("Neither root gives a stationary twisted process.")

    ρ = (β_bar
         + γ_o ** 2 / 2
         + cf * ξ_f * xbar_f
         + co * (ξ_o * xbar_o + γ_o * σ_o)
         + co ** 2 * σ_o ** 2 / 2)

    return cf, co, ρ, mr
```

141.7.2 A Breeden SDF

To make things concrete, we now plug in a specific stochastic discount factor: the Breeden CRRA consumption-based SDF.

Breeden [1979] derived the continuous-time SDF for an investor with time-separable CRRA preferences over a consumption stream.

We specify log consumption with the same affine state X^o driving expected growth and $\sqrt{X^f}$ driving volatility:

$$dc_t = X_t^o\,dt + \sqrt{X_t^f}\,\vartheta_f\,dB_t^f + \vartheta_o\,dB_t^o.$$

With time-separable CRRA utility (risk aversion a) and subjective discount rate b, the stochastic discount factor is

$$S_t = \exp(-bt - a(c_t - c_0)),$$

i.e. an exponential of $-b$ times time minus a times log consumption growth.

Reading off the additive functional coefficients gives

$$\bar{\beta}^s = -b, \quad \beta_f^s = 0, \quad \beta_o^s = -a, \quad \gamma_f^s = -a\vartheta_f, \quad \gamma_o^s = -a\vartheta_o. \tag{141.20}$$

We will use these parameters in the numerical example below.

141.7.3 Recursive preferences (optional)

A famous limitation of CRRA preferences is that they conflate risk aversion with the elasticity of intertemporal substitution.

Recursive preferences Epstein and Zin [1989], Kreps and Porteus [1978] separate the two, and add a forward-looking continuation-value term to the discount-factor expression.

This is what powers the equity-premium results in Bansal and Yaron [2004].

A QuantEcon lecture that studies long-run dynamics under recursive preferences in a different setting is *Survival and Long-Run Dynamics under Recursive Preferences*.

This section derives the SDF coefficients for the unit-elasticity recursive specification.

You can skip on a first read and come back later — the numerical example uses the simpler Breeden parameters above.

For the unit-elasticity recursive specification, conjecture a continuation value of the form

$$W_t = \frac{1}{1-a}\exp\left[(1-a)(w_f X_t^f + w_o X_t^o + c_t + \bar{w})\right], \qquad a > 1. \tag{141.21}$$

Matching the local mean of this continuation value gives

$$\begin{aligned} 0 &= -\xi_f w_f + \frac{(1-a)\sigma_f^2}{2} w_f^2 + (1-a)\vartheta_f \sigma_f w_f + \frac{(1-a)\vartheta_f^2}{2} - b w_f, \\ 0 &= -\xi_o w_o + 1 - b w_o, \\ b\bar{w} &= \xi_f \bar{x}_f w_f + \xi_o \bar{x}_o w_o + \frac{(1-a)\sigma_o^2}{2} w_o^2 + (1-a)\vartheta_o \sigma_o w_o + \frac{(1-a)\vartheta_o^2}{2}. \end{aligned} \tag{141.22}$$

The relevant solution has $w_o = 1/(\xi_o + b)$, selects the stable root of the quadratic for w_f, and then determines $\bar{w}$ from the constant equation.

The stochastic discount factor is the product $S_t = \exp(A_t^B)\exp(A_t^W)$, where the Breeden-like logarithmic component is

$$A_t^B = -bt - \int_0^t X_s^o ds - \int_0^t \sqrt{X_s^f}\vartheta_f dB_s^f - \int_0^t \vartheta_o dB_s^o$$

The continuation-value martingale component is

$$\begin{aligned} A_t^W = (1-a)\int_0^t \sqrt{X_s^f}(\vartheta_f + w_f\sigma_f) dB_s^f + (1-a)\int_0^t (\vartheta_o + w_o\sigma_o) dB_s^o \\ - \frac{(1-a)^2}{2}\int_0^t X_s^f (\vartheta_f + w_f\sigma_f)^2 ds - \frac{(1-a)^2}{2}(\vartheta_o + w_o\sigma_o)^2 t. \end{aligned} \tag{141.23}$$

Thus the recursive-utility SDF has affine parameters

$$\begin{aligned} \bar{\beta}^s &= -b - \frac{(1-a)^2}{2}(\vartheta_o + w_o\sigma_o)^2, \\ \beta_f^s &= -\frac{(1-a)^2}{2}(\vartheta_f + w_f\sigma_f)^2, \\ \beta_o^s &= -1, \\ \gamma_f^s &= -a\vartheta_f - (a-1)w_f\sigma_f, \\ \gamma_o^s &= -a\vartheta_o - (a-1)w_o\sigma_o. \end{aligned} \tag{141.24}$$

The local Brownian risk prices are therefore $\sqrt{x^f}[a\vartheta_f + (a-1)w_f\sigma_f]$ for B^f exposure and $a\vartheta_o + (a-1)w_o\sigma_o$ for B^o exposure.

The numerical example below uses the simpler Breeden specification, but the same operator calculation applies once the SDF parameters are replaced by (141.24).

141.7.4 Numerical example

Let's set up parameters and solve for the principal eigenpair.

We use illustrative parameters loosely in the long-run-risk family: a mean-reverting volatility factor X^f with mean 0.04, a more persistent predictable-growth factor X^o with mean 0.02, risk aversion $a = 4$, and a time discount rate $b = 0.03$.

These are chosen to make the operator calculations transparent, not to match asset-price moments.

With time-separable CRRA preferences and $a = 4$ they imply an instantaneous riskless rate of about 10.6% at the state means and a long zero-coupon yield of 9.6% (printed below), which is the familiar risk-free rate puzzle rather than a resolution of it.

```
params_state = {
    "ξ_f": 0.70,
    "xbar_f": 0.04,
    "σ_f": -0.20,
    "ξ_o": 0.50,
    "xbar_o": 0.02,
    "σ_o": 0.01,
}

a = 4.0
b = 0.03
ϑ_f = 0.06
ϑ_o = 0.02

params_sdf = {
    **params_state,
    "β_bar": -b,
    "β_f": 0.0,
    "β_o": -a,
    "γ_f": -a * ϑ_f,
    "γ_o": -a * ϑ_o,
}

cf_s, co_s, ρ_s, mr_s = solve_affine_eigenfunction(params_sdf)

print("principal eigenfunction ψ(xf, xo) = exp(cf xf + co xo)")
print(f"cf = {cf_s:.6f}")
print(f"co = {co_s:.6f}")
print(f"ρ = {ρ_s:.6f}")
print(f"twisted mean-reversion coefficient for Xf = {mr_s:.6f}")
print(f"long-run zero-coupon yield = {-ρ_s:.4f}")
```

```
principal eigenfunction ψ(xf, xo) = exp(cf xf + co xo)
cf = 0.044232
co = -8.000000
ρ = -0.095962
twisted mean-reversion coefficient for Xf = 0.650231
long-run zero-coupon yield = 0.0960
```

The long-run zero-coupon yield $-\rho_s$ represents the asymptotic decay rate in the SDF expectation $\mathbb{E}[S_t]$.

We can also check that the rejected root for c_f would have produced a non-stationary twisted process, a clear example of stochastic stability selecting one of two algebraically valid eigenfunctions.

```
ξ_f = params_sdf["ξ_f"]
σ_f = params_sdf["σ_f"]
γ_f = params_sdf["γ_f"]
β_f = params_sdf["β_f"]

disc = ((ξ_f - γ_f * σ_f) ** 2
        - σ_f ** 2 * (2 * β_f + γ_f ** 2))
root = np.sqrt(disc)

cf_candidates = np.array([
    ((ξ_f - γ_f * σ_f) - root) / σ_f ** 2,
    ((ξ_f - γ_f * σ_f) + root) / σ_f ** 2
])

for cf in cf_candidates:
    mr = ξ_f - σ_f * (γ_f + cf * σ_f)
    print(f"cf = {cf:8.4f}, twisted mean reversion = {mr:8.4f}")
```

```
cf =   0.0442, twisted mean reversion =   0.6502
cf =  32.5558, twisted mean reversion =  -0.6502
```

141.7.5 The martingale component

Now that we have the eigenpair (ρ, ϕ), we can write down all three pieces of the factorization (141.1) explicitly.

The martingale component $\hat{M}_t = \exp(\hat{A}_t)$ defined in (141.10) has log

$$\begin{aligned}\hat{A}_t = &\int_0^t \sqrt{X_s^f}(\gamma_f + c_f\sigma_f)\, dB_s^f + \int_0^t (\gamma_o + c_o\sigma_o)\, dB_s^o \\ &- \frac{1}{2}\int_0^t X_s^f(\gamma_f + c_f\sigma_f)^2\, ds - \frac{1}{2}\int_0^t (\gamma_o + c_o\sigma_o)^2\, ds.\end{aligned}$$

The first line is the Brownian integral (the "exponential martingale" piece); the second is the Itô correction needed to make it a martingale.

Under the twisted measure induced by $\hat{M}$, the drifts of the state variables shift to

$$\begin{aligned}dX_t^f &: \quad \xi_f(\bar{x}_f - X_t^f) + X_t^f\sigma_f(\gamma_f + c_f\sigma_f), \\ dX_t^o &: \quad \xi_o(\bar{x}_o - X_t^o) + \sigma_o(\gamma_o + c_o\sigma_o).\end{aligned}$$

The drift distortions are exactly the Girsanov shifts induced by the Brownian loadings of $\hat{M}$.

Let's now simulate the state and check the factorization numerically.

The first check is the *algebraic identity*: once we define $\hat{M}$ through (141.10), the equation $M_t = \exp(\rho t)\hat{M}_t\phi(X_0)/\phi(X_t)$ is automatic for any choice of (ρ, ϕ), and the error below is just floating-point round-off.

The second, substantive, check is whether the eigenpair (ρ, ϕ) we solved for really makes $\hat{M}$ a martingale, which we approximate by computing $\mathbb{E}[\hat{M}_t]$ across many simulated paths.

```
def brownian_increments(n, dt, seed=1234):
    """Independent Brownian increments for the simulation."""
    rng = np.random.default_rng(seed)
    draws = rng.normal(scale=np.sqrt(dt), size=(n, 2))
    return draws[:, 0], draws[:, 1]


def simulate_states(params, T=40.0, dt=0.01, seed=1234):
    """Euler simulation of the affine state process."""
    n = int(T / dt)
    t = np.linspace(0, T, n + 1)
    Xf = np.empty(n + 1)
    Xo = np.empty(n + 1)
    dBf, dBo = brownian_increments(n, dt, seed)

    Xf[0] = params["xbar_f"]
    Xo[0] = params["xbar_o"]

    for k, (dBf_k, dBo_k) in enumerate(zip(dBf, dBo)):
        xf = max(Xf[k], 1e-10)
        xo = Xo[k]

        Xf[k + 1] = (xf
                     + params["ξ_f"] * (params["xbar_f"] - xf) * dt
                     + np.sqrt(xf) * params["σ_f"] * dBf_k)
        Xf[k + 1] = max(Xf[k + 1], 1e-10)

        Xo[k + 1] = (xo
                     + params["ξ_o"] * (params["xbar_o"] - xo) * dt
                     + params["σ_o"] * dBo_k)

    return t, Xf, Xo, dBf, dBo


def additive_log_M(params, t, Xf, Xo, dBf, dBo):
    """Additive log functional along a simulated path."""
    dt = t[1] - t[0]
    A = np.zeros_like(t)

    for k, (dBf_k, dBo_k) in enumerate(zip(dBf, dBo)):
        xf = max(Xf[k], 1e-10)
        xo = Xo[k]

        drift = (params["β_bar"]
                 + params["β_f"] * xf
                 + params["β_o"] * xo)

        shock = (np.sqrt(xf) * params["γ_f"] * dBf_k
                 + params["γ_o"] * dBo_k)

        A[k + 1] = A[k] + drift * dt + shock

    return A


t, Xf, Xo, dBf, dBo = simulate_states(params_sdf)
A_log = additive_log_M(params_sdf, t, Xf, Xo, dBf, dBo)
```

(continues on next page)

(continued from previous page)

```
ψ_0 = np.exp(cf_s * Xf[0] + co_s * Xo[0])
ψ_t = np.exp(cf_s * Xf + co_s * Xo)

M = np.exp(A_log)
M_hat = np.exp(-ρ_s * t) * M * ψ_t / ψ_0
transient = ψ_0 / ψ_t

identity_error = np.max(np.abs(M - np.exp(ρ_s * t) * M_hat * transient))
print(f"algebraic identity error = {identity_error:.2e}")
```

```
algebraic identity error = 2.22e-16
```

The error above is up to machine precision, as expected.

Next we estimate $\mathbb{E}[\hat{M}_t \mid X_0 = \bar{x}]$ over a Monte Carlo sample of paths.

If $\hat{M}$ is a martingale, the population mean is exactly 1 at every t, and the sample mean should lie within a few standard errors of 1.

```
def simulate_M_hat(params, ρ, cf, co, n_paths=2000, T=20.0, dt=0.01, seed=2024):
    """Monte Carlo paths of hat M_t along an ensemble of trajectories."""
    rng = np.random.default_rng(seed)
    n = int(T / dt)
    t = np.linspace(0, T, n + 1)
    M_hat_paths = np.empty((n_paths, n + 1))

    for k in range(n_paths):
        seed_k = rng.integers(1, 10**9)
        t_k, Xf_k, Xo_k, dBf_k, dBo_k = simulate_states(
            params, T=T, dt=dt, seed=int(seed_k)
        )
        A_k = additive_log_M(params, t_k, Xf_k, Xo_k, dBf_k, dBo_k)
        ψ_t_k = np.exp(cf * Xf_k + co * Xo_k)
        ψ_0_k = np.exp(cf * Xf_k[0] + co * Xo_k[0])
        M_hat_paths[k] = np.exp(-ρ * t_k) * np.exp(A_k) * ψ_t_k / ψ_0_k

    return t, M_hat_paths


t_mc, M_hat_paths = simulate_M_hat(params_sdf, ρ_s, cf_s, co_s)
M_hat_mean = M_hat_paths.mean(axis=0)
M_hat_se = M_hat_paths.std(axis=0, ddof=1) / np.sqrt(M_hat_paths.shape[0])

print("   t      mean       se      (mean - 1) / se")
for t_check in [1.0, 5.0, 10.0, 20.0]:
    idx = np.argmin(np.abs(t_mc - t_check))
    mean = M_hat_mean[idx]
    se = M_hat_se[idx]
    z = (mean - 1.0) / se
    print(f"{t_mc[idx]:5.2f}   {mean:7.4f}  {se:7.4f}   {z:+6.2f}")
```

```
   t      mean       se      (mean - 1) / se
 1.00    0.9974   0.0038    -0.68
 5.00    0.9928   0.0087    -0.83
10.00    1.0005   0.0129    +0.04
20.00    1.0020   0.0201    +0.10
```

```
fig, axes = plt.subplots(2, 2, figsize=(12, 8))

axes[0, 0].plot(t, Xf)
axes[0, 0].set_ylabel("$X_t^f$")
axes[0, 0].set_xlabel("$t$")

axes[0, 1].plot(t, Xo)
axes[0, 1].set_ylabel("$X_t^o$")
axes[0, 1].set_xlabel("$t$")

axes[1, 0].plot(t, M, label="$M_t$")
axes[1, 0].plot(t, np.exp(ρ_s * t) * M_hat * transient,
                "--", label="factorization")
axes[1, 0].set_ylabel("multiplicative factorization")
axes[1, 0].set_xlabel("$t$")
axes[1, 0].legend()

axes[1, 1].plot(t, np.exp(ρ_s * t), label="$\\exp(\\rho t)$")
axes[1, 1].plot(t, M_hat, label="$\\hat M_t$", alpha=0.8)
axes[1, 1].plot(t, transient, label="$\\phi(X_0)/\\phi(X_t)$", alpha=0.8)
axes[1, 1].set_ylabel("three components")
axes[1, 1].set_xlabel("$t$")
axes[1, 1].legend()

plt.tight_layout()
plt.show()
```

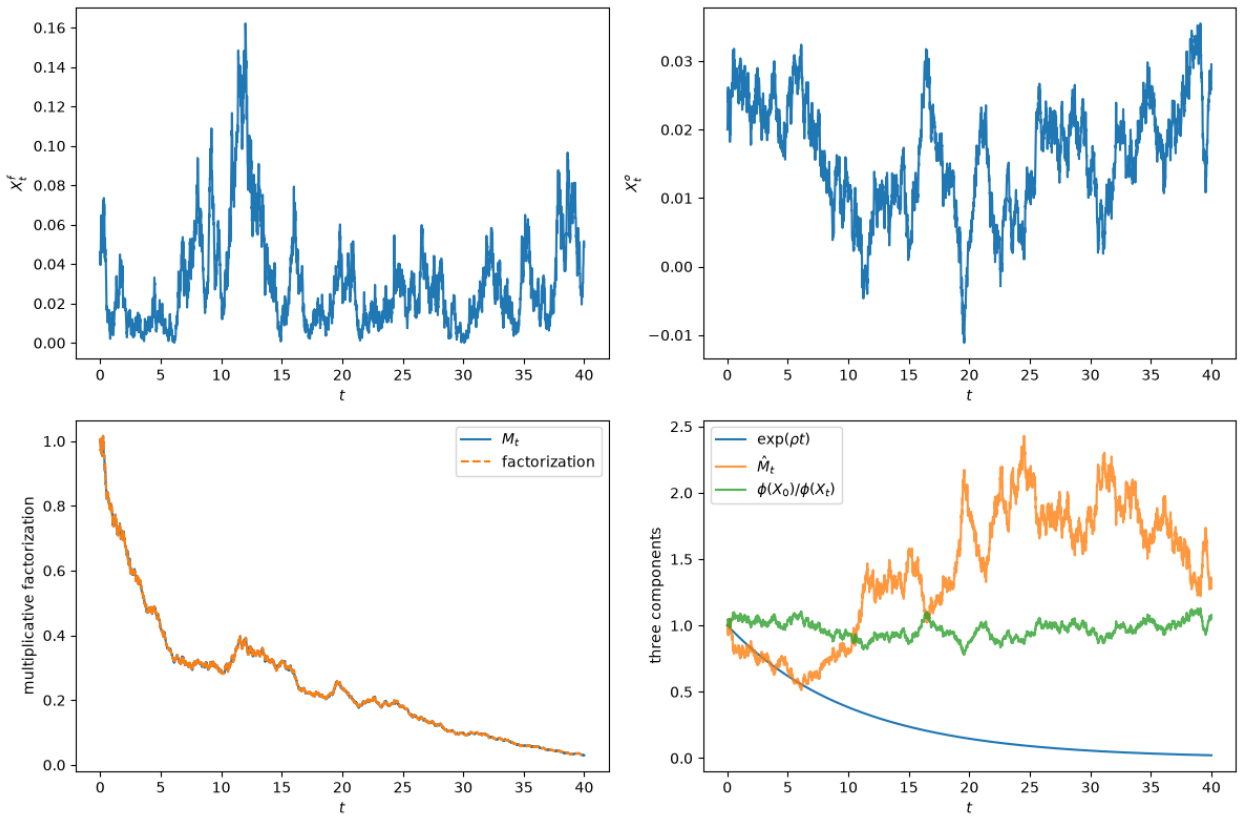


Fig. 141.3: Affine state paths and multiplicative factorization

The figure shows how the three components evolve over time.

141.8 Long-run risk prices

We can now use the factorization to compute long-run analogues of the instantaneous risk prices that arise in standard continuous-time asset pricing.

We can ask:

> If an investor takes on a small exposure to a shock today, how much extra expected return do they need when that return is measured as a long-horizon rate rather than an instantaneous one?

The two answers, local and long-run, need not agree.

A shock that moves a persistent state variable has a small *immediate* effect on the cash flow but a large *cumulative* effect on future expected growth and discounting, which is why the two answers come apart.

So the long-run risk price is the local price *plus a persistence correction.*

The size of the correction depends on the speed of mean reversion.

141.8.1 Defining the prices

Definition 141.8.1 (Local Brownian Risk Price)

The **local Brownian risk price** is the state-dependent vector $-\gamma^s(x)$.

A small Brownian exposure γ_i^v earns a required expected return of $-\gamma_i^v \gamma_i^s$ per unit time, so a unit of exposure γ_i^v is priced at $-\gamma_i^s$.

For the local price we just read off the SDF coefficients.

Definition 141.8.2 (Long-Run Risk Price)

The **long-run risk price** is the marginal change in the long-run required return on a cash flow with respect to a small change in its risk exposure.

When we work with growing cash flows, the long-run required return is $R_\infty = -\rho + \delta$, where ρ is the principal eigenvalue of the GS semigroup and δ is the trend growth rate, held fixed.

Computing the long-run price requires solving the principal eigenvalue problem, which captures how a shock propagates through the persistent state component.

141.8.2 Two frontiers

We will see *two* related ways to vary risk exposure, each leading to a slightly different long-run risk price:

1. **Valuation-functional frontier:**
   - Hold the SDF S fixed and vary the asset's Brownian exposures (γ_f^v, γ_o^v).
   - Use the local pricing restriction to determine the drift β^v, then compute ρ^v for the V-semigroup.
2. **Cash-flow frontier:**
   - Hold the SDF S fixed and vary the cash-flow's growth exposures (γ_f^g, γ_o^g).
   - Set $M = GS$ and compute the principal eigenvalue ρ of the cash-flow valuation semigroup.

In simple log-normal examples, these two frontiers coincide.

But they can differ with stochastic volatility, nonlinear dynamics, or jump risk.

We will work out both types of examples in the affine model below: in the B^o direction the two frontiers agree exactly, while in the B^f direction stochastic volatility pulls them apart.

141.8.3 Stochastic discount factor decomposition

Before getting into risk prices, a natural benchmark is the case $M = S$: applying the factorization (141.1) directly to the SDF itself.

The factorization becomes

$$S_t = \exp(\rho t)\widehat{M}_t \frac{\phi(X_0)}{\phi(X_t)}.$$

This is the **permanent-transitory decomposition** of Alvarez and Jermann [2005], linked now to a concrete eigenfunction construction.

The factor $\exp(\rho t)$ is the deterministic trend in the SDF and the martingale $\widehat{M}$ is its random *permanent* component; the state-dependent ratio is *transient* and washes out.

See also

The same spectral decomposition of the pricing operator is the central tool in *The Recovery Theorem*, which uses it to "recover" subjective beliefs from observed prices, and in *Misspecified Recovery*, which examines what goes wrong when the permanent martingale component is mistakenly assumed away.

For a long zero-coupon bond,

$$\exp(-\rho t)\,\mathbb{E}[S_t \mid X_0 = x] \to \phi(x) \int \frac{1}{\phi}\, d\hat{\varsigma},$$

so long-maturity discount bond prices depend on the current state primarily through the eigenfunction ϕ.

141.8.4 Comparison in the affine model

For the affine specification, we can write closed-form expressions for both local and long-run prices of a B^o shock.

The local price is just

$$\text{local price of } B^o = -\gamma_o^s.$$

The long-run price (which we will derive below using the valuation-functional frontier) is

$$\text{long-run price of } B^o = -\gamma_o^s - \frac{\beta_o^s}{\xi_o}\sigma_o. \tag{141.25}$$

The extra term $-(\beta_o^s/\xi_o)\sigma_o$ is the **persistence correction**.

It arises because:

- a B^o shock moves the persistent growth predictor X^o, and
- X^o mean reverts at rate ξ_o, so the cumulative effect of the shock on future SDF growth scales like $1/\xi_o$.

As ξ_o shrinks, persistence grows and the long-run price diverges from the local one, which is the central economic content of long-run risk models.

The local price of B^f exposure is state dependent (it scales with $\sqrt{X_t^f}$), and the long-run price of B^f exposure is nonlinear, since c_f is the root of a quadratic.

```
γ_s_o = params_sdf["γ_o"]
β_s_o = params_sdf["β_o"]
ξ_o = params_sdf["ξ_o"]
σ_o = params_sdf["σ_o"]

local_price_o = -γ_s_o
long_run_price_o = -γ_s_o - (β_s_o / ξ_o) * σ_o

print(f"local price of B^o exposure    = {local_price_o:.4f}")
print(f"long-run price of B^o exposure = {long_run_price_o:.4f}")
```

```
local price of B^o exposure    = 0.0800
long-run price of B^o exposure = 0.1600
```

The next cell illustrates how persistence changes the wedge between local and long-run prices.

```
ξ_o_grid = np.array([0.10, 0.20, 0.50, 1.00, 2.00, 5.00])
local_grid = np.full_like(ξ_o_grid, local_price_o)
long_grid = -γ_s_o - (β_s_o / ξ_o_grid) * σ_o

fig, ax = plt.subplots()
ax.plot(ξ_o_grid, local_grid, "--", lw=2, label="local")
ax.plot(ξ_o_grid, long_grid, "o-", lw=2, label="long-run")
ax.set_xscale("log")
ax.set_xlabel("mean-reversion speed $\\xi_o$")
ax.set_ylabel("risk price")
ax.legend()
plt.show()
```

As we can see, the long-run price of B^o exposure is sensitive to the mean-reversion speed ξ_o, but the local price is a constant.

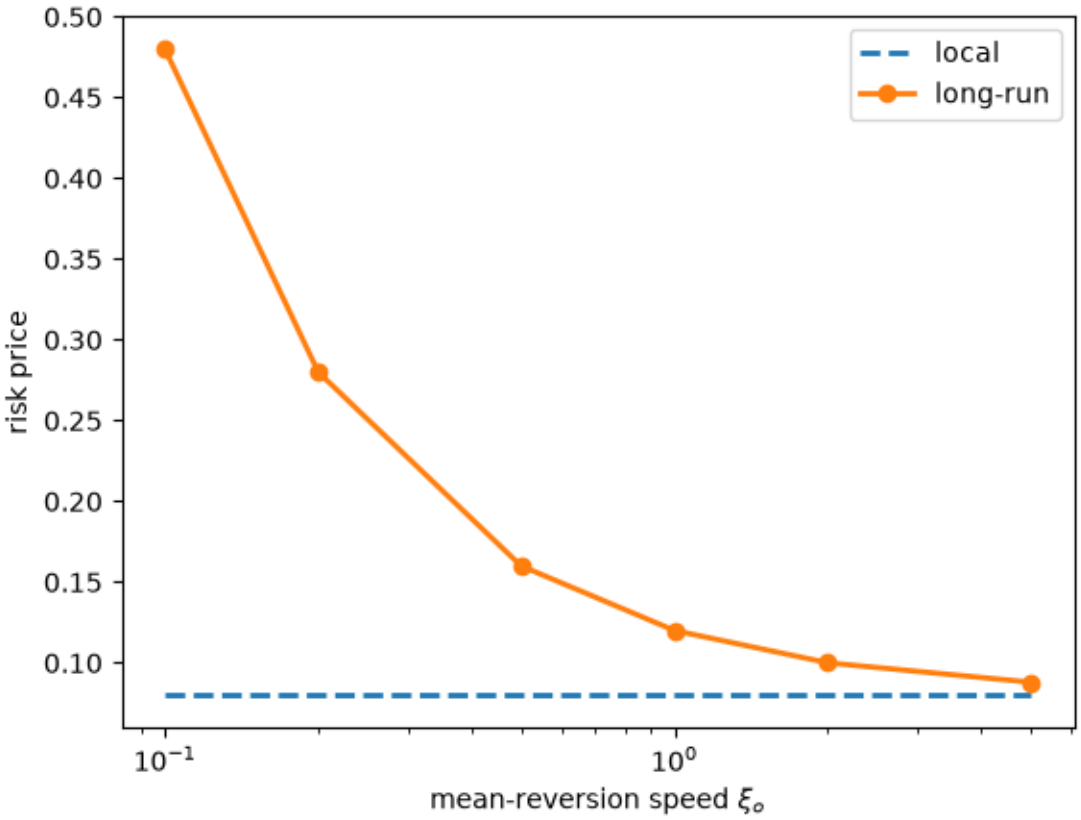


Fig. 141.4: Persistence and long-run risk prices

141.8.5 Changing valuation functionals

Now we work out the long-run risk price formula by varying the asset's exposure, the **valuation-functional frontier** introduced above.

We hold the SDF S fixed and pick Brownian exposures (γ_f^v, γ_o^v) for the asset return, parameterizing the valuation functional as

$$
\begin{aligned}
A_t^v = \bar{\beta}^v t + \int_0^t \beta_f^v X_s^f ds + \int_0^t \beta_o^v X_s^o ds \\
+ \int_0^t \sqrt{X_s^f} \gamma_f^v dB_s^f + \int_0^t \gamma_o^v dB_s^o.
\end{aligned}
$$

The martingale restriction on VS determines the drift coefficients from the chosen Brownian exposures (γ_f^v, γ_o^v):

$$
\begin{aligned}
\bar{\beta}^v &= -\bar{\beta}^s - \frac{1}{2}(\gamma_o^s + \gamma_o^v)^2, \\
\beta_f^v &= -\beta_f^s - \frac{1}{2}(\gamma_f^s + \gamma_f^v)^2, \\
\beta_o^v &= -\beta_o^s.
\end{aligned}
\tag{141.26}
$$

Applying the affine eigenvalue formula to $M = V$ gives

$$
\begin{aligned}
\rho^v &= \bar{\beta}^v + \frac{(\gamma_o^v)^2}{2} + c_f^v \xi_f \bar{x}_f + c_o^v(\xi_o \bar{x}_o + \gamma_o^v \sigma_o) + (c_o^v)^2 \frac{\sigma_o^2}{2} \\
&= -\bar{\beta}^s - \frac{(\gamma_o^s)^2}{2} - \gamma_o^s \gamma_o^v + c_f^v \xi_f \bar{x}_f + c_o^v(\xi_o \bar{x}_o + \gamma_o^v \sigma_o) + (c_o^v)^2 \frac{\sigma_o^2}{2},
\end{aligned}
\tag{141.27}
$$

where $c_o^v = \beta_o^v/\xi_o = -\beta_o^s/\xi_o$ and c_f^v solves the same quadratic equation as in (141.16) with (β_f, γ_f) replaced by (β_f^v, γ_f^v).

Holding γ_f^v fixed, differentiating (141.27) with respect to γ_o^v gives the long-run valuation-functional price of B^o exposure:

$$
\frac{\partial \rho^v}{\partial \gamma_o^v} = -\gamma_o^s + c_o^v \sigma_o = -\gamma_o^s - \frac{\beta_o^s}{\xi_o} \sigma_o. \tag{141.28}
$$

This matches the formula (141.25) we previewed above, namely the local price plus the persistence correction $-(\beta_o^s/\xi_o)\sigma_o$.

Let's verify the formula numerically by finite-differencing the eigenvalue computation.

```
def valuation_params_from_exposure(γ_v_o, γ_v_f=0.0):
    """Affine parameters for a valuation functional V."""
    p = dict(params_sdf)
    p.update({
        "β_bar": (-params_sdf["β_bar"]
                  - 0.5 * (params_sdf["γ_o"] + γ_v_o) ** 2),
        "β_f": (-params_sdf["β_f"]
                - 0.5 * (params_sdf["γ_f"] + γ_v_f) ** 2),
        "β_o": -params_sdf["β_o"],
        "γ_f": γ_v_f,
        "γ_o": γ_v_o,
    })
    return p


def valuation_eigenvalue_for_exposure(γ_v_o, γ_v_f=0.0):
```

(continues on next page)

(continued from previous page)

```
    """Principal eigenvalue for the valuation functional frontier."""
    p = valuation_params_from_exposure(γ_v_o, γ_v_f)
    _, _, ρ, _ = solve_affine_eigenfunction(p)
    return ρ
```

```
γ_v_o_grid = np.linspace(-0.5, 0.5, 101)
ρ_v_grid = np.array([
    valuation_eigenvalue_for_exposure(g) for g in γ_v_o_grid
])

fig, ax = plt.subplots()
ax.plot(γ_v_o_grid, ρ_v_grid, lw=2)
ax.set_xlabel("valuation exposure $\\gamma_o^v$")
ax.set_ylabel("principal eigenvalue $\\rho^v$")
plt.show()
```

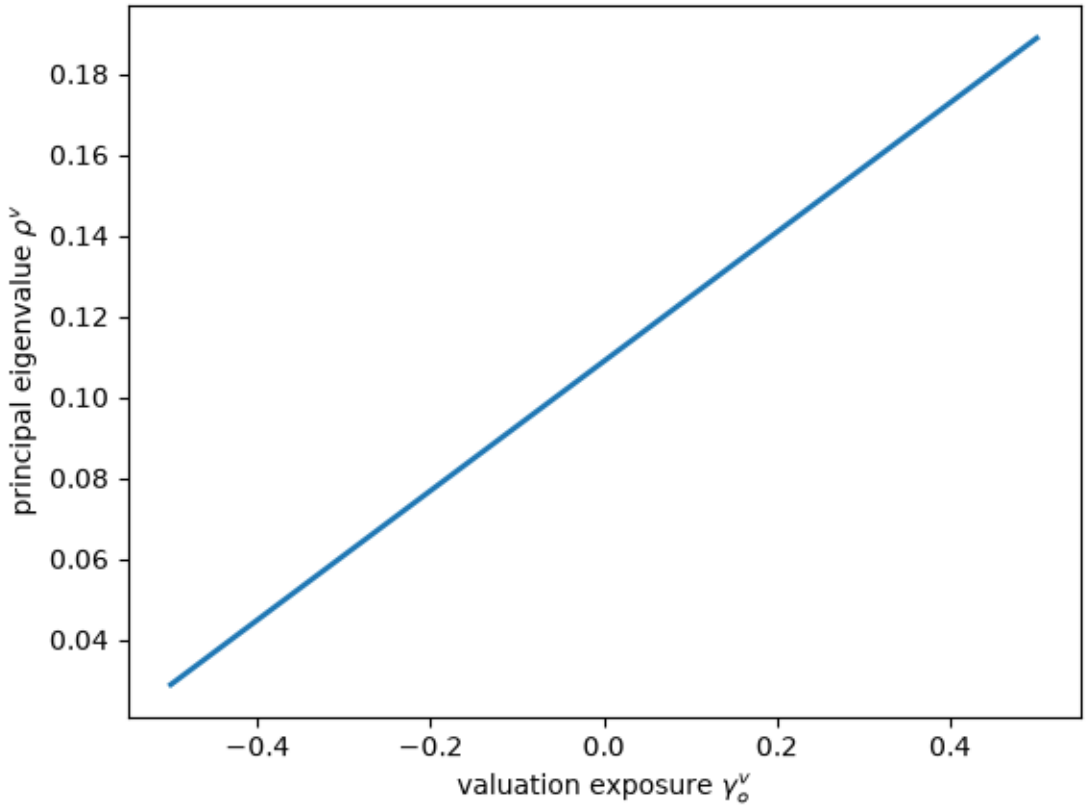


Fig. 141.5: Changing valuation functionals

```
valuation_slope = (
    valuation_eigenvalue_for_exposure(0.001)
    - valuation_eigenvalue_for_exposure(-0.001)
) / 0.002

print(f"finite-difference slope = {valuation_slope:.6f}")
```

(continues on next page)

(continued from previous page)

```
print(f"formula                  = {long_run_price_o:.6f}")
```

```
finite-difference slope = 0.160000
formula                 = 0.160000
```

141.8.6 Changing cash-flow risk

The second long-run frontier varies the cash-flow's exposure to risk rather than the return.

We pick a growth functional $G_t = \exp(A_t^g)$ with the affine parameterization

$$\begin{aligned} A_t^g = \delta t + \int_0^t \sqrt{X_s^f}\gamma_f^g dB_s^f + \int_0^t \gamma_o^g dB_s^o \\ - \frac{1}{2}\int_0^t \left[X_s^f(\gamma_f^g)^2 + (\gamma_o^g)^2\right] ds. \end{aligned} \tag{141.29}$$

The last line is the Itô compensator that makes $\exp(A_t^g - \delta t) = \hat{G}_t$ a *local* martingale, with δ the constant trend growth rate.

In this affine setting the three abstract conditions of *Definition 141.5.6* — a stationary distribution, an irreducible skeleton, and Harris recurrence — all follow once the twisted X^f is mean reverting and does not hit zero.

Alongside them sits the separate requirement that $\hat{G}$ itself be a martingale, the Assumption-6.1 analogue for the growth twist.

Take nonattainment first.

Under the twist the drift of X^f becomes

$$\xi_f \bar{x}_f - [\xi_f - \sigma_f(\gamma_f + c_f\sigma_f)]x^f, \qquad \gamma_f = \gamma_f^s + \gamma_f^g,$$

while Girsanov leaves the diffusion coefficient $\sqrt{x^f}\sigma_f$ unchanged.

Writing this drift in the canonical square-root form $\hat{\kappa}(\hat{\theta} - x^f)$, the twist rescales the mean-reversion speed $\hat{\kappa}$ and the long-run mean $\hat{\theta}$, but it leaves their *product* $\hat{\kappa}\hat{\theta} = \xi_f\bar{x}_f$ fixed.

Since the **Feller-type nonattainment** inequality sees only that product, it is the same one that governs the original process,

$$2\xi_f\bar{x}_f \geq \sigma_f^2,$$

with no exposure parameter in it at all.

Mean reversion of the twisted X^f, i.e. $\hat{\kappa} > 0$, is a genuinely separate requirement, and it is picked by the same root-selection argument we used for the SDF in (141.18).

Neither condition is sufficient on its own.

Note

Nonattainment is free here, but stability overall is not: changing growth risk can still violate it and invalidate the long-run approximation, so the choice of (γ_f^g, γ_o^g) isn't free.

At this calibration, for instance, a large negative γ_f^g drives the discriminant in (141.18) below zero, so no real exponential-affine eigenfunction exists at all.

To price the cash flow $D_t = D_0 G_t \psi(X_t)$, use the semigroup generated by $M = GS$.

The combined affine parameters are

$$\begin{aligned}
\bar{\beta} &= \bar{\beta}^s + \delta - \frac{(\gamma_o^g)^2}{2}, \\
\beta_f &= \beta_f^s - \frac{(\gamma_f^g)^2}{2}, \\
\beta_o &= \beta_o^s, \\
\gamma_f &= \gamma_f^s + \gamma_f^g, \\
\gamma_o &= \gamma_o^s + \gamma_o^g.
\end{aligned}$$

Let ρ be the principal eigenvalue of this GS semigroup.

Then $-\rho$ is the long-run decay rate in value, and

$$R_\infty = -\rho + \delta$$

is the asymptotic required return net of the cash-flow growth rate.

```
def required_return_for_growth_exposure(γ_g_o, γ_g_f=0.0, δ=0.02):
    """Long-run required return for a cash-flow growth exposure."""
    p = dict(params_sdf)
    p.update({
        "β_bar": params_sdf["β_bar"] + δ - 0.5 * γ_g_o ** 2,
        "β_f": params_sdf["β_f"] - 0.5 * γ_g_f ** 2,
        "β_o": params_sdf["β_o"],
        "γ_f": params_sdf["γ_f"] + γ_g_f,
        "γ_o": params_sdf["γ_o"] + γ_g_o,
    })

    _, _, ρ, _ = solve_affine_eigenfunction(p)
    return -ρ + δ
```

```
γ_g_o_grid = np.linspace(-0.5, 0.5, 101)
required_returns = np.array([
    required_return_for_growth_exposure(g) for g in γ_g_o_grid
])

local_line = (required_return_for_growth_exposure(0.0)
              + local_price_o * γ_g_o_grid)

fig, ax = plt.subplots()
ax.plot(γ_g_o_grid, required_returns, lw=2,
        label="long-run required return")
ax.plot(γ_g_o_grid, local_line, "--", lw=2,
        label="local slope")
ax.set_xlabel("cash-flow exposure $\\gamma_o^g$")
ax.set_ylabel("rate of return")
ax.legend()
plt.show()
```

The slope of the long-run line is the risk price in (141.25), and the dashed line shares the same value at $\gamma_o^g = 0$ but with slope equal to the local Brownian risk price $-\gamma_o^s$.

The gap between the two slopes is the *persistence correction*; the dashed line is a slope comparator, not the actual Breeden local expected-return frontier (which is state dependent in X^f and X^o).

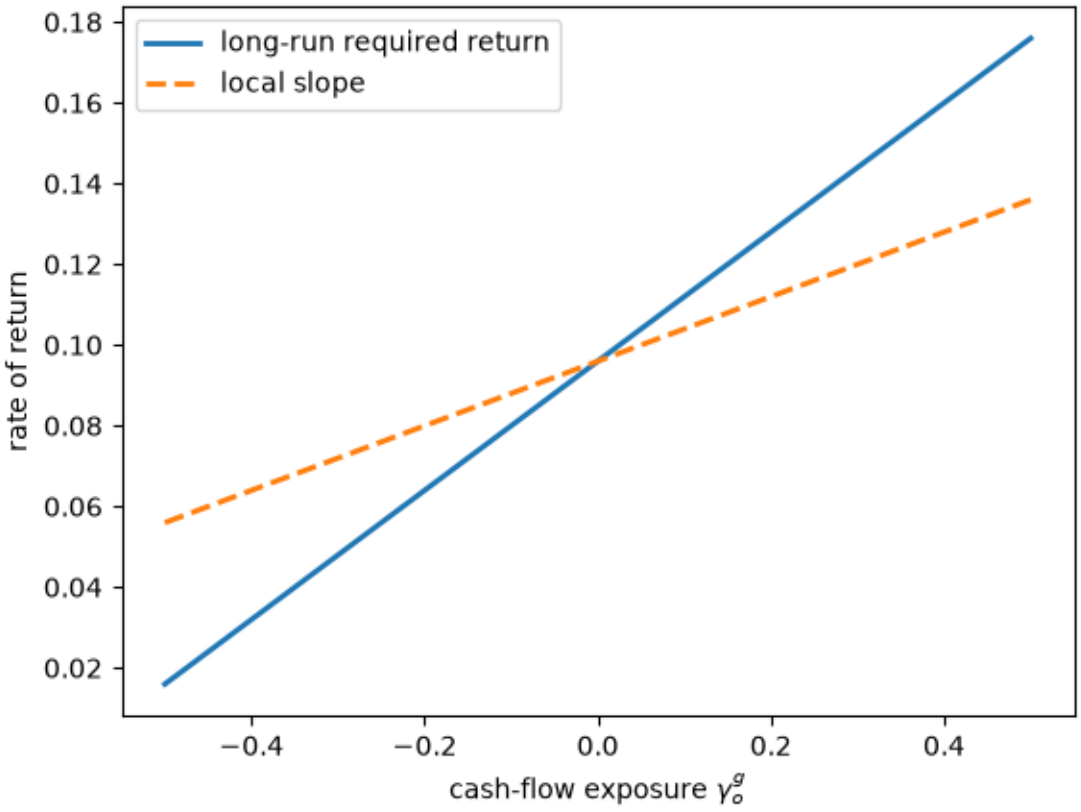


Fig. 141.6: Local and long-run pricing of persistent growth risk

```
finite_difference = (
    required_return_for_growth_exposure(0.001)
    - required_return_for_growth_exposure(-0.001)
) / 0.002

print(f"finite-difference slope = {finite_difference:.6f}")
print(f"formula                 = {long_run_price_o:.6f}")
```

```
finite-difference slope = 0.160000
formula                 = 0.160000
```

141.8.7 Where the two frontiers come apart

We promised above that the valuation-functional and cash-flow frontiers can differ.

In the B^o direction they do not: the two finite differences just computed agree to six decimals, because γ_o enters the eigenvalue (141.19) linearly.

The B^f direction is the interesting one, because there the exposure feeds through c_f, which solves the quadratic (141.16).

```
def central_difference(f, h=1e-5):
    return (f(h) - f(-h)) / (2 * h)

valuation_price_f = central_difference(
    lambda g: valuation_eigenvalue_for_exposure(0.0, g)
)
cashflow_price_f = central_difference(
    lambda g: required_return_for_growth_exposure(0.0, g)
)
local_price_f = -params_sdf["γ_f"] * params_sdf["xbar_f"]

print(f"B^f local price (at x^f = xbar_f) = {local_price_f:.6f}")
print(f"B^f valuation-functional frontier = {valuation_price_f:.6f}")
print(f"B^f cash-flow frontier            = {cashflow_price_f:.6f}")
```

```
B^f local price (at x^f = xbar_f) = 0.009600
B^f valuation-functional frontier = 0.009905
B^f cash-flow frontier            = 0.010716
```

The three numbers differ in the third decimal place.

The local price evaluated at $x^f = \bar{x}_f$ is the smallest, the valuation-functional frontier is slightly larger, and the cash-flow frontier is larger still.

The reason is that a B^f exposure shifts c_f nonlinearly, and the two frontiers load that nonlinearity differently: the valuation frontier also adjusts β_f^v through the local pricing restriction (141.26), whereas the cash-flow frontier adjusts β_f through the Itô compensator in (141.29).

This is the discrepancy anticipated in the discussion of the two frontiers above, and it is absent in the B^o direction precisely because X^o has constant volatility.

141.9 Assumptions behind the scenes

The examples above make the eigenfunction calculation look routine.

For finite-state chains and the affine model, it really is; Perron-Frobenius theory and closed-form algebra handle every requirement.

But in a general state space, three things can go wrong, each corresponding to one of the assumptions we have been carrying.

The rest of this section examines each in turn.

141.9.1 Issue 1: $\hat{M}$ might fail to be a martingale

A positive eigenfunction ϕ gives us a candidate martingale $\hat{M}$ from (141.10), but $\hat{M}$ is only a nonnegative local martingale, hence a supermartingale.

A supermartingale is not enough to define a probability measure: we need $\mathbb{E}[\hat{M}_t] = 1$, i.e. a genuine martingale, the content of Assumption 6.1 in Hansen and Scheinkman [2009].

A standard way to verify this is a two-sided **Girsanov construction**: write the drift and jump distortion induced by $\hat{M}$, check that the distorted Markov process is well-behaved, and verify that the reverse density (the inverse of $\hat{M}$) is locally integrable.

141.9.2 Issue 2: the twisted process might fail to be stable

Even with $\hat{M}$ a martingale, the long-run limit (141.12) requires that the twisted process actually settles into a steady state.

This is where stochastic stability (our trio of stationary distribution, irreducibility of the skeleton, and Harris recurrence) does real work.

The affine example illustrates this concretely, since we *rejected* one of the two algebraically valid eigenfunctions because it implied an explosive twisted square-root process.

141.9.3 Issue 3: a principal eigenfunction might not exist

In a general state space, even *existence* of a strictly positive eigenfunction is not automatic.

A standard sufficient condition starts with a **Lyapunov-type drift bound**: there is a function $V \geq 1$ on the state space and a constant a_0 such that

$$\frac{\mathbb{A}V}{V} \leq a_0.$$

Roughly: V doesn't grow too fast under the semigroup.

With this in hand, for any $\alpha > a_0$ define the **resolvent operator**

$$F_\alpha \psi(x) = \int_0^\infty \exp(-\alpha t)\, \mathbb{E}\left[M_t \, \frac{V(X_t)}{V(x)} \, \psi(X_t) \,\middle|\, X_0 = x \right] dt. \tag{141.30}$$

F_α is the Laplace transform of the semigroup generated by the *rescaled* multiplicative functional $M_t V(X_t)/V(X_0)$.

The existence proof then proceeds in three steps:

1. **Irreducibility for the resolvent.** There exists a reference measure ν such that $F_\alpha \mathbf{1}_\Lambda(x) > 0$ for every x whenever $\nu(\Lambda) > 0$, so the resolvent doesn't "miss" any region of state space.

2. **Nummelin minorization.** Irreducibility yields a *bounded nonnegative function* s on the state space, with $\int s\, d\nu > 0$, such that for every nonnegative ψ,

$$F_\alpha \psi(x) \geq s(x) \int \psi\, d\nu.$$

   The function s (often called the *minorization function*) measures how strongly the resolvent dominates a fixed reference measure ν.

3. **Eigenfunction extraction.** The minorization, combined with additional boundedness or strengthened drift assumptions, identifies a critical spectral value for F_α and an associated positive eigenfunction.
   - Inverting the resolvent transform produces a positive eigenfunction for the original semigroup.

These steps are all nontrivial and are out of scope for this lecture.

The details are in Section 9 of Hansen and Scheinkman [2009].

We can summarize the chain of conditions as:

Want	Need
A factorization (141.1)	A positive eigenfunction ϕ
$\widehat{M}$ to define a probability measure	$\widehat{M}$ is a martingale (Assumption 6.1)
The long-run limit (141.12)	Stochastic stability of the twisted process
A unique principal eigenfunction	Stability selects among positive eigenfunctions

In the finite-state case, all four follow from one Perron-Frobenius calculation; in the affine model, they reduce to picking the right root of a quadratic.

In general, each must be checked separately.

The full theory in Hansen and Scheinkman [2009] also delivers stronger L^p approximation results and Lyapunov criteria for stochastic stability, which we don't reproduce here.

141.10 Summary

The Hansen-Scheinkman approach studies long-run risk by studying positive eigenfunctions of valuation semigroups.

The main steps are:

1. Model discounting, growth, or cumulated returns by a positive multiplicative functional M.
2. Build the semigroup $\mathbb{M}_t \psi(x) = \mathbb{E}[M_t \psi(X_t) \mid X_0 = x]$.
3. When V is a valuation functional, use the local pricing restriction that VS is a martingale to determine the drift of V from its Brownian and jump exposures, and then apply the eigenvalue problem to $M = V$; for cash-flow valuation take $M = Q = GS$, the functional generating the semigroup $\{\mathbb{Q}_t\}$, where the pricing restriction falls on S alone and G enters only as a growth twist.
4. Solve the principal eigenvalue problem $\mathbb{A}\phi = \rho\phi$.
5. Use the factorization $M_t = \exp(\rho t)\widehat{M}_t \phi(X_0)/\phi(X_t)$.
6. Under the twisted probability measure induced by $\widehat{M}$, use stability to obtain long-run approximations of the form (141.12).

In finite-state problems, this is Perron-Frobenius theory.

In affine diffusion problems, exponential-affine eigenfunctions often produce closed-form formulas.

The long-run risk prices that emerge can differ sharply from local risk prices when shocks move persistent state variables.

This persistence effect is the economic channel emphasized in long-run risk asset-pricing models and in the empirical work of Hansen *et al.* [2008].

141.11 Exercises

Exercise 141.11.1

Consider a two-state Markov chain with intensity matrix

$$U = \begin{bmatrix} -\lambda & \lambda \\ \mu & -\mu \end{bmatrix}.$$

Let the multiplicative functional have decay rate $r_1 > 0$ in state 1, decay rate $r_2 = 0$ in state 2, and no jumps.

1. Write down the generator matrix A.
2. Find the principal eigenvalue ρ in terms of λ, μ, and r_1.
3. Verify numerically with $\lambda = 0.4$, $\mu = 0.6$, and $r_1 = 0.05$.
4. Show that $-r_1 < \rho < 0$.

Solution

Here is one solution:

1. The generator is

$$A = \begin{bmatrix} -\lambda - r_1 & \lambda \\ \mu & -\mu \end{bmatrix}.$$

2. The characteristic equation is

$$\rho^2 + (\lambda + \mu + r_1)\rho + \mu r_1 = 0.$$

Hence the principal eigenvalue is the larger root

$$\rho = \frac{-(\lambda + \mu + r_1) + \sqrt{(\lambda + \mu + r_1)^2 - 4\mu r_1}}{2}.$$

3. Numerical verification:

```
λ, μ, r1 = 0.4, 0.6, 0.05

disc = (λ + μ + r1) ** 2 - 4 * μ * r1
ρ_formula = (-(λ + μ + r1) + np.sqrt(disc)) / 2

U_ex = np.array([[-λ, λ],
                 [μ, -μ]])
r_ex = np.array([r1, 0.0])
κ_ex = np.zeros((2, 2))

A_ex = build_generator(U_ex, r_ex, κ_ex)
ρ_numeric, ψ_numeric = principal_eigenpair(A_ex)

print(f"formula   ρ = {ρ_formula:.8f}")
print(f"numeric   ρ = {ρ_numeric:.8f}")
print(f"difference  = {abs(ρ_formula-ρ_numeric):.2e}")
```

```
formula  ρ = -0.02939431
numeric  ρ = -0.02939431
difference   = 0.00e+00
```

4. Let

$$q(x) = x^2 + (\lambda + \mu + r_1)x + \mu r_1.$$

Then $q(0) = \mu r_1 > 0$ and

$$q(-r_1) = -\lambda r_1 < 0.$$

Since the parabola opens upward, one root lies in $(-r_1, 0)$.

The principal eigenvalue is the larger root, so $-r_1 < \rho < 0$.

Exercise 141.11.2

Fig. 141.4 shows the long-run price of B^o exposure falling towards the local price as the mean-reversion speed ξ_o rises.

This exercise asks you to pin down the *rate* at which it falls.

Define the **persistence wedge**

$$w(\xi_o) = \text{long-run price of } B^o - \text{local price of } B^o.$$

1. Using (141.25) together with the Breeden coefficients (141.20), show analytically that

$$w(\xi_o) = \frac{a\sigma_o}{\xi_o},$$

so the wedge is exactly proportional to $1/\xi_o$ and, in particular, does not depend on the consumption loading ϑ_o.

2. Verify this numerically on the grid $\xi_o \in \{0.1, 0.2, 0.5, 1, 2, 5\}$ by checking that the product $\xi_o\, w(\xi_o)$ is constant across the grid and equal to $a\sigma_o$.
3. Plot w against ξ_o on log-log axes and confirm that the fitted slope is -1, the signature of exactly hyperbolic decay.

Solution

Here is one solution:

1. The local price is $-\gamma_o^s$ and, by (141.25), the long-run price is $-\gamma_o^s - (\beta_o^s/\xi_o)\sigma_o$.

Subtracting, the common term $-\gamma_o^s$ cancels and only the persistence correction survives:

$$w(\xi_o) = -\frac{\beta_o^s}{\xi_o}\sigma_o.$$

The Breeden coefficients (141.20) set $\beta_o^s = -a$, which gives $w(\xi_o) = a\sigma_o/\xi_o$.

Because ϑ_o enters only through $\gamma_o^s = -a\vartheta_o$, which appears identically in both prices, it cancels and cannot affect the wedge.

2. and *3.* Numerically:

```
ξ_vals = np.array([0.1, 0.2, 0.5, 1.0, 2.0, 5.0])
wedge = -(params_sdf["β_o"] / ξ_vals) * params_sdf["σ_o"]

print(f"a * σ_o = {a * params_sdf['σ_o']:.6f}\n")
for ξ, w in zip(ξ_vals, wedge):
    print(f"ξ_o = {ξ:3.1f}: wedge = {w:.6f}, ξ_o * wedge = {ξ * w:.6f}")

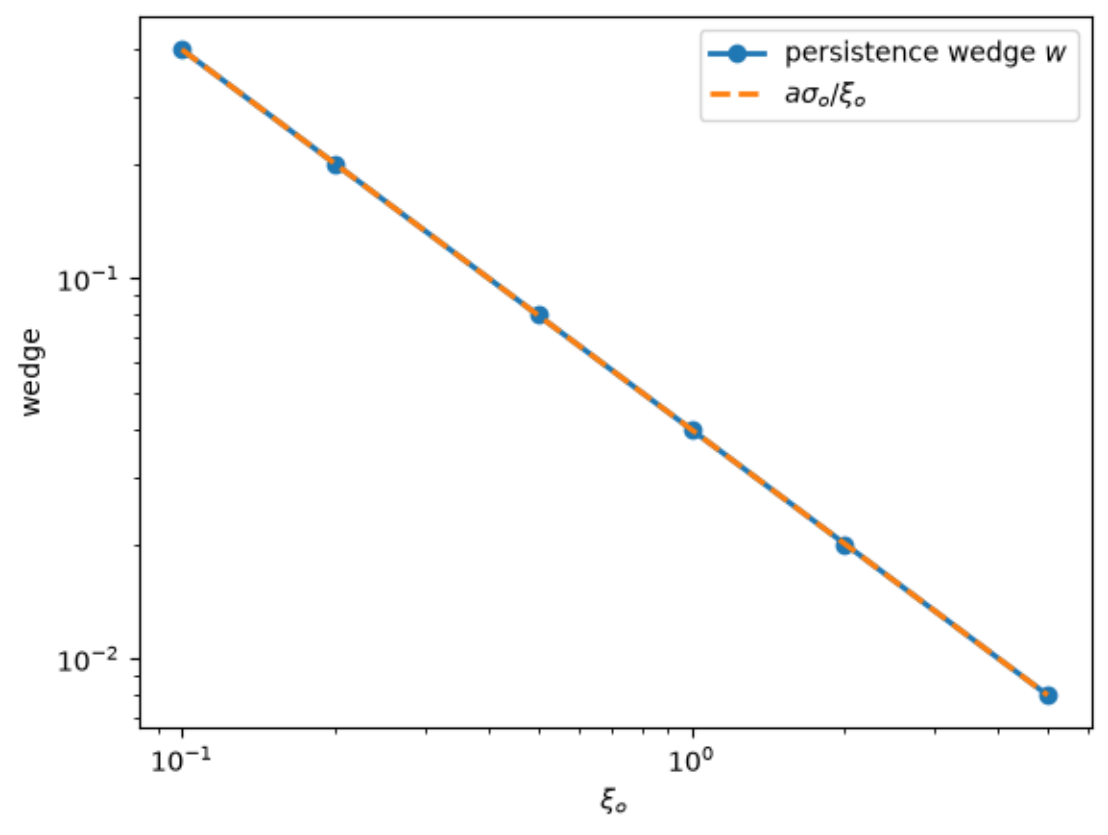
slope, _ = np.polyfit(np.log(ξ_vals), np.log(wedge), 1)
print(f"\nlog-log slope = {slope:.6f}")

fig, ax = plt.subplots()
ax.loglog(ξ_vals, wedge, "o-", lw=2, label="persistence wedge $w$")
ax.loglog(ξ_vals, a * params_sdf["σ_o"] / ξ_vals, "--", lw=2,
          label="$a\\sigma_o/\\xi_o$")
ax.set_xlabel("$\\xi_o$")
ax.set_ylabel("wedge")
ax.legend()
plt.show()
```

```
a * σ_o = 0.040000

ξ_o = 0.1: wedge = 0.400000, ξ_o * wedge = 0.040000
ξ_o = 0.2: wedge = 0.200000, ξ_o * wedge = 0.040000
ξ_o = 0.5: wedge = 0.080000, ξ_o * wedge = 0.040000
ξ_o = 1.0: wedge = 0.040000, ξ_o * wedge = 0.040000
ξ_o = 2.0: wedge = 0.020000, ξ_o * wedge = 0.040000
ξ_o = 5.0: wedge = 0.008000, ξ_o * wedge = 0.040000

log-log slope = -1.000000
```



The product $\xi_o \, w(\xi_o)$ prints as 0.040000 at every grid point, matching $a\sigma_o = 4 \times 0.01$, and the fitted log-log slope is -1.000000.

So the two prices do not merely converge: the wedge decays at the exact hyperbolic rate $a\sigma_o/\xi_o$.

The economics behind the rate is that a B^o shock displaces X^o by σ_o on impact, and that displacement decays at rate ξ_o, so its cumulative effect on expected growth is σ_o/ξ_o.

Each unit of that cumulative effect is priced by the SDF's loading a on X^o.

Exercise 141.11.3

Consider a three-state chain with states {expansion, normal, contraction}, intensity matrix

$$U = \begin{bmatrix} -0.40 & 0.30 & 0.10 \\ 0.20 & -0.50 & 0.30 \\ 0.10 & 0.20 & -0.30 \end{bmatrix},$$

decay-rate vector $r = (0.06, 0.04, 0.01)$, and no jumps in the multiplicative functional.

Let $\psi = (3, 1, 2)$.

1. Compute the principal eigenpair (ρ, ϕ) and twisted stationary distribution $\hat{\varsigma}$, and report the theoretical limit

$$\phi \sum_i \frac{\psi_i}{\phi_i} \hat{\varsigma}_i.$$

2. Plot

$$\max_i \left| \exp(-\rho t)(\mathbb{M}_t \psi)_i - \phi_i \sum_j \frac{\psi_j}{\phi_j} \hat{\varsigma}_j \right|$$

on a logarithmic scale.

3. Compare the convergence rate to the spectral gap between the largest and second-largest real parts of the eigenvalues of A.

Solution

Here is one solution:

```
state_names = ["expansion", "normal", "contraction"]

U3 = np.array([[-0.40,  0.30,  0.10],
               [ 0.20, -0.50,  0.30],
               [ 0.10,  0.20, -0.30]])

r3 = np.array([0.06, 0.04, 0.01])
κ3 = np.zeros((3, 3))

A3 = build_generator(U3, r3, κ3)
ρ3, ψ3 = principal_eigenpair(A3)
A3_hat = twisted_generator(A3, ρ3, ψ3)
ς3 = stationary_distribution(A3_hat)

ψ = np.array([3.0, 1.0, 2.0])
limit = ψ3 * np.sum((ψ / ψ3) * ς3)

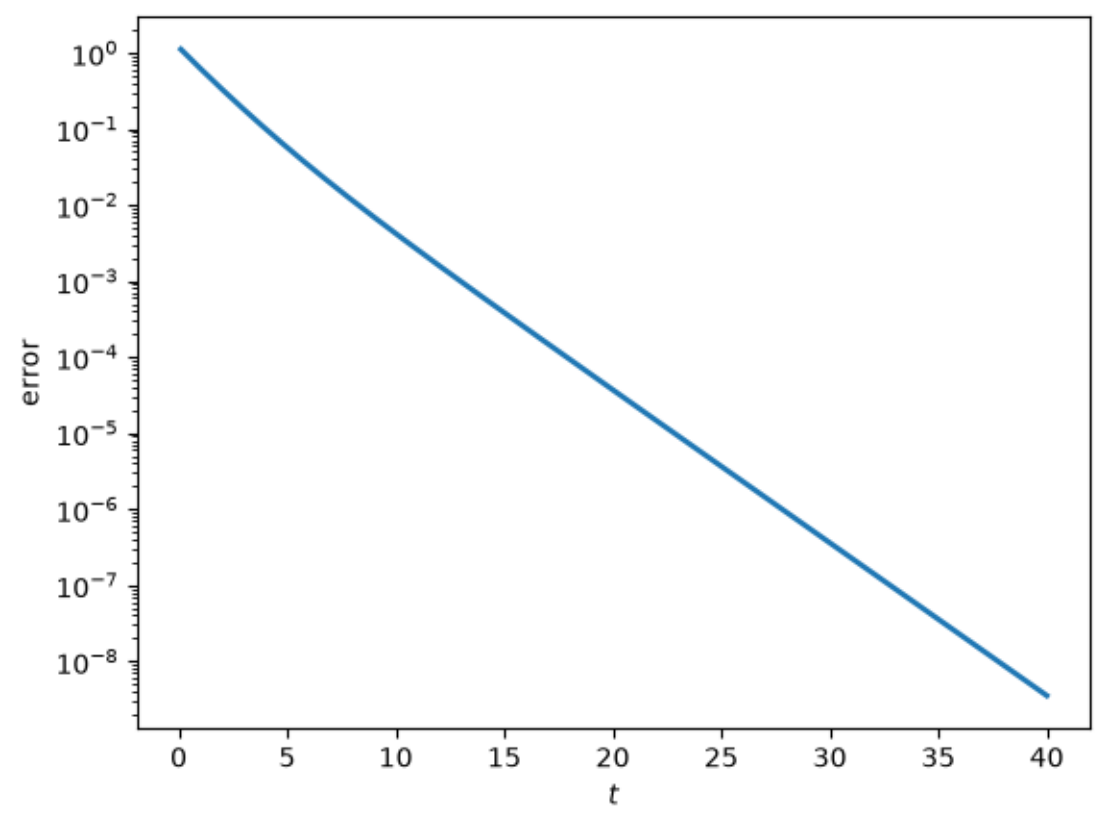
print("limit:")
for name, value in zip(state_names, limit):
    print(f"  {name:11s} {value:.6f}")

t_vals = np.linspace(0.1, 40, 300)
errors = np.array([
    np.max(np.abs(np.exp(-ρ3 * t_val) * expm(t_val * A3) @ ψ - limit))
    for t_val in t_vals
])

eigenvalues = eig(A3, right=False)
real_parts = np.sort(eigenvalues.real)[::-1]
gap = real_parts[0] - real_parts[1]

fig, ax = plt.subplots()
ax.semilogy(t_vals, errors, lw=2)
ax.set_xlabel("$t$")
ax.set_ylabel("error")
plt.show()
```

```
limit:
  expansion   1.808490
  normal      1.906080
  contraction 2.021897
```



For part 3 we estimate the empirical decay rate by fitting a straight line to $\log(\text{error})$ over a window in which the leading correction term dominates, and compare the fitted slope to the gap.

```
mask = (t_vals > 20) & (t_vals < 40)
fitted_slope, _ = np.polyfit(t_vals[mask], np.log(errors[mask]), 1)

print(f"empirical decay rate = {-fitted_slope:.6f}")
print(f"spectral gap         = {gap:.6f}")
```

```
empirical decay rate = 0.462900
spectral gap         = 0.462679
```

The normalized semigroup converges at an exponential rate governed by the separation between the dominant eigenvalue and the remaining eigenvalues.

In this finite-state example, that separation is the spectral gap, and the two numbers printed above agree to three decimal places.

The small remaining discrepancy is the contribution of the third eigenvalue, which has not yet died out completely; fitting over a later window shrinks it further.

Exercise 141.11.4

Derive the local martingale restriction (141.5) from Itô's formula.

Let $M = \exp(A)$ for the additive functional A in (141.3), with parameters (β, γ, κ).

1. Decompose $A_t = A_t^c + A_t^j$ into its continuous and pure-jump parts and write down dA_t^c and the jump mag-

nitudes ΔA_t.

2. Apply Itô's formula for semimartingales to $f(a) = e^a$ to show that

$$dM_t = M_{t-}\, dA_t^c + \tfrac{1}{2}\, M_{t-}\, d\langle A^c, A^c\rangle_t + M_{t-}\big(\exp[\Delta A_t] - 1\big) \quad \text{at jumps.}$$

3. Use $d\langle A^c, A^c\rangle_t = \gamma^\top \gamma\, dt$ and rewrite the jump term as an integral against the random counting measure ζ.
4. Split ζ into its compensator $\eta(dy \mid X_{t-})\, dt$ and the compensated martingale measure $\tilde{\zeta} = \zeta - \eta(dy \mid X_{t-})\, dt$.
5. Collect drift (predictable) and martingale terms and conclude that M is a local martingale iff the drift vanishes at every state, which gives (141.5).

Solution

Here is one solution.

1. From the parameterization (141.3),

$$dA_t^c = \beta(X_t)\, dt + \gamma(X_{t-})^\top\, dB_t, \qquad \Delta A_t = \kappa(X_t, X_{t-}) \text{ at a jump time.}$$

2. For $f(a) = e^a$ we have $f'(a) = f''(a) = e^a$, so $f'(A_{t-}) = f''(A_{t-}) = M_{t-}$.

The (generalized) Itô's formula for a semimartingale gives

$$dM_t = f'(A_{t-})\, dA_t^c + \tfrac{1}{2}\, f''(A_{t-})\, d\langle A^c, A^c\rangle_t + [f(A_t) - f(A_{t-})].$$

Since $A_t = A_{t-} + \Delta A_t$ at a jump, $f(A_t) - f(A_{t-}) = M_{t-}\big(\exp[\Delta A_t] - 1\big)$, which is the stated expression.

3. Substituting $d\langle A^c, A^c\rangle_t = \gamma(X_{t-})^\top \gamma(X_{t-})\, dt$ and rewriting the jump contribution as an integral against the random counting measure ζ of (X, A) gives

$$\begin{aligned}\frac{dM_t}{M_{t-}} &= \beta(X_t)\, dt + \gamma(X_{t-})^\top dB_t + \tfrac{1}{2}\, \gamma(X_{t-})^\top \gamma(X_{t-})\, dt \\ &\quad + \int \big(\exp[\kappa(y, X_{t-})] - 1\big)\, \zeta(dy, dt).\end{aligned}$$

4. Writing $\zeta = \tilde{\zeta} + \eta(dy \mid X_{t-})\, dt$ separates the jump integral into a martingale and a predictable drift contribution:

$$\int \big(\exp[\kappa(y, X_{t-})] - 1\big)\zeta(dy, dt) = \int \big(\exp[\kappa(y, X_{t-})] - 1\big)\tilde{\zeta}(dy, dt) + \int \big(\exp[\kappa(y, X_{t-})] - 1\big)\eta(dy \mid X_{t-})\, dt.$$

5. Collecting drift and martingale terms,

$$\begin{aligned}\frac{dM_t}{M_{t-}} &= \left[\beta(X_t) + \tfrac{1}{2}\gamma^\top\gamma + \int \big(\exp[\kappa(y, X_{t-})] - 1\big)\, \eta(dy \mid X_{t-})\right] dt \\ &\quad + \gamma^\top dB_t + \int \big(\exp[\kappa(y, X_{t-})] - 1\big)\, \tilde{\zeta}(dy, dt).\end{aligned}$$

The Brownian and compensated-jump terms are local martingales, so M is a local martingale iff the bracketed drift vanishes for every state x, giving (141.5).

Exercise 141.11.5

Derive the closed-form expression for the extended generator in (141.8).

Take the jump-diffusion setting in which the Markov state satisfies

$$dX_t^c = \xi(X_t)dt + \Gamma(X_t)dB_t$$

between jumps, with $\Sigma = \Gamma\Gamma^\top$ and jump compensator $\eta(dy \mid x)$.

Let $M = \exp(A)$ for the additive functional A in (141.3) and let ϕ be a smooth, strictly positive function.

Recall from *Definition 141.4.1* that $\mathbb{A}\phi = \chi$ if

$$N_t = M_t\phi(X_t) - \phi(X_0) - \int_0^t M_s\chi(X_s)\,ds$$

is a local martingale.

The task therefore has two pieces: identify the predictable drift of $M_t\phi(X_t)$ to read off a candidate χ, and verify that the residual N_t really is a local martingale.

(1) Apply Itô's formula to $Y_t = \exp(A_t)\phi(X_t)$ between jumps and show that the continuous part of dY_t has drift

$$M_t\left[(\xi + \Gamma\gamma)^\top\frac{\partial\phi}{\partial x} + \frac{1}{2}\operatorname{trace}\left(\Sigma\frac{\partial^2\phi}{\partial x\partial x^\top}\right) + \left(\beta + \frac{\gamma^\top\gamma}{2}\right)\phi\right]dt.$$

(2) Show that at a jump time t with $X_{t-} = x$ and $X_t = y$,

$$\Delta Y_t = M_{t-}\big[\exp[\kappa(y,x)]\phi(y) - \phi(x)\big],$$

and conclude that the predictable jump compensator contributes drift

$$M_{t-}\int\big\{\exp[\kappa(y,x)]\phi(y) - \phi(x)\big\}\eta(dy \mid x)\,dt.$$

(3) Decompose

$$\exp[\kappa(y,x)]\phi(y) - \phi(x) = \exp[\kappa(y,x)]\big[\phi(y) - \phi(x)\big] + \big[\exp[\kappa(y,x)] - 1\big]\phi(x),$$

combine with part (1) to obtain the full predictable drift coefficient $\chi(x)$, and check that it matches the closed form (141.8).

(4) Identify the continuous Itô integrand from part (1) as a local martingale N_t^c and the compensated jump sum from part (2) as a local martingale N_t^j.

Show that the semimartingale decomposition of $Y_t = M_t\phi(X_t)$ implies

$$M_t\phi(X_t) - \phi(X_0) - \int_0^t M_s\,\chi(X_s)\,ds = N_t^c + N_t^j,$$

which is a local martingale.

Conclude via *Definition 141.4.1* that $\mathbb{A}\phi = \chi$.

Solution

Here is one solution:

(1) Set $g(a,x) = e^a \phi(x)$, so that $Y_t = g(A_t, X_t)$.

Between jumps, the continuous parts of A and X are

$$dA_t^c = \beta(X_t)\,dt + \gamma(X_t)^\top dB_t, \qquad dX_t^c = \xi(X_t)\,dt + \Gamma(X_t)\,dB_t,$$

with quadratic covariations

$$d\langle A^c, A^c\rangle_t = \gamma^\top\gamma\,dt, \qquad d\langle X^c, X^c\rangle_t = \Sigma\,dt, \qquad d\langle A^c, X^c\rangle_t = \Gamma\gamma\,dt.$$

The partial derivatives of g satisfy

$$\partial_a g = g, \quad \partial_{aa} g = g, \quad \partial_x g = e^a \frac{\partial\phi}{\partial x}, \quad \partial_{ax} g = e^a \frac{\partial\phi}{\partial x}, \quad \partial_{xx} g = e^a \frac{\partial^2\phi}{\partial x \partial x^\top}.$$

Itô's formula yields a continuous local martingale N_t^c (the dB_t part) plus the drift

$$M_t\Big[\beta\phi + \tfrac{1}{2}\gamma^\top\gamma\,\phi + \xi^\top\tfrac{\partial\phi}{\partial x} + \tfrac{1}{2}\,\mathrm{trace}\big(\Sigma\,\tfrac{\partial^2\phi}{\partial x\partial x^\top}\big) + (\Gamma\gamma)^\top\tfrac{\partial\phi}{\partial x}\Big]\,dt.$$

Grouping the gradient terms gives the expression in the question.

(2) At a jump time t, $\Delta A_t = \kappa(X_t, X_{t-}) = \kappa(y,x)$, so $M_t = M_{t-}\exp[\kappa(y,x)]$ and

$$\Delta Y_t = M_t\phi(X_t) - M_{t-}\phi(X_{t-}) = M_{t-}\big[\exp[\kappa(y,x)]\phi(y) - \phi(x)\big].$$

The sum of jumps minus its predictable compensator is a local martingale N_t^j, and the compensator itself contributes the predictable drift quoted in the question.

(3) Adding the jump drift from (2) to the continuous drift from (1), the predictable drift of $Y_t = M_t\phi(X_t)$ at state x is $M_t\,\chi(x)\,dt$ with

$$\begin{aligned}\chi(x) = (\xi + \Gamma\gamma)^\top\frac{\partial\phi}{\partial x} + \frac{1}{2}\,\mathrm{trace}\left(\Sigma\frac{\partial^2\phi}{\partial x\partial x^\top}\right) + \left(\beta + \frac{\gamma^\top\gamma}{2}\right)\phi(x) \\ + \int \exp[\kappa(y,x)]\big[\phi(y) - \phi(x)\big]\,\eta(dy \mid x) + \phi(x)\int\big[\exp[\kappa(y,x)] - 1\big]\,\eta(dy \mid x).\end{aligned}$$

Collecting the terms multiplying $\phi(x)$ recovers (141.8).

(4) Putting the two martingale pieces together and integrating the drift, steps (1) and (2) give the semimartingale decomposition

$$M_t\phi(X_t) = \phi(X_0) + \int_0^t M_s\,\chi(X_s)\,ds + N_t^c + N_t^j.$$

Rearranging,

$$N_t := M_t\phi(X_t) - \phi(X_0) - \int_0^t M_s\,\chi(X_s)\,ds = N_t^c + N_t^j,$$

which is a local martingale because it is a sum of local martingales.

This is exactly the property required by *Definition 141.4.1*, so $\chi = \mathbb{A}\phi$.

CHAPTER

TWO

MISSPECIFIED RECOVERY

Contents

- *Misspecified Recovery*
  - *Overview*
  - *Three transition matrices*
  - *Martingale component*
  - *From matrices to the general framework*
  - *When recovery fails*
  - *Additional state vector*
  - *Measuring the martingale component*
  - *Lessons*
  - *Exercises*

142.1 Overview

The lecture *The Recovery Theorem* studies the case in which recovery is valid.

There, **transition independence** lets us use Arrow prices to separate investors' beliefs from the pricing kernel.

This lecture asks what the same Perron–Frobenius approach delivers when that restriction is not imposed.

The eigenfunction, martingale-factorization, and twisted-measure machinery that we use here is developed at length in *Long-term Risk: An Operator Approach*.

We will keep three probability measures separate.

The first is the correctly specified probability measure, which governs the Markov state in the model.

The second is the one-period risk-neutral probability measure, which comes from normalizing one-period Arrow prices by bond prices.

The third is the probability measure recovered by Perron–Frobenius Theory, also called the long-term risk-neutral measure.

The central question is whether the recovered probability measure equals the correctly specified probability measure.

Borovička *et al.* [2016] show that, in general, the answer is no.

The paper studies the ratio of the recovered probability measure to the correctly specified probability measure.

The reason is that the stochastic discount factor can contain a martingale component that changes the probability measure.

If that martingale component is identically one, Ross recovery returns the correctly specified transition probabilities.

If it is not identically one, the recovered probability measure absorbs long-term risk adjustments, because martingale increments compound along histories.

In the examples below, the recovered probability measure assigns more probability to adverse long-run-risk states than the correctly specified probability measure.

We will:

- use results from *The Recovery Theorem* without re-proving them,
- study misspecification through the martingale component,
- show why recursive utility and permanent shocks make the recovered probability measure differ from the correctly specified probability measure,
- measure the difference in a long-run risk model.

We will use the following imports.

```
import numpy as np
import matplotlib.pyplot as plt
from scipy import linalg
from scipy.integrate import solve_ivp
from scipy.stats import gaussian_kde
```

The next cell contains code inherited from the previous lecture.

It row-normalizes Arrow prices, finds the Perron–Frobenius eigenvalue and positive right eigenvector, and computes stationary distributions.

142.2 Three transition matrices

Let $\mathbf{P} = [p_{ij}]$ denote the correctly specified transition matrix and $\mathbf{Q} = [q_{ij}]$ the Arrow price matrix.

Here "correctly specified" means that $\mathbf{P}$ is the transition matrix that governs the Markov state in the model.

The one-period stochastic discount factor (SDF) satisfies

$$q_{ij} = s_{ij} p_{ij}. \tag{142.1}$$

We will compare $\mathbf{P}$ with two probability matrices constructed from $\mathbf{Q}$.

The first one is the **one-period risk-neutral matrix**.

It divides each row of $\mathbf{Q}$ by the price of a one-period discount bond in the current state:

$$\bar{p}_{ij} = \frac{q_{ij}}{\sum_k q_{ik}}.$$

This matrix absorbs one-period risk adjustments into transition probabilities.

The second one is the transition matrix associated with the **long-term risk-neutral probability**.

It starts from the Perron–Frobenius eigenvalue and positive right eigenvector of $\mathbf{Q}$.

Let $(\exp(\widehat{\eta}), \widehat{e})$ solve

$$\mathbf{Q}\widehat{e} = \exp(\widehat{\eta})\widehat{e}. \tag{142.2}$$

Then define

$$\widehat{p}_{ij} = \exp(-\widehat{\eta}) q_{ij} \frac{\widehat{e}_j}{\widehat{e}_i}. \tag{142.3}$$

The factor $\widehat{e}_j / \widehat{e}_i$ is chosen to cancel any SDF component of the form $\exp(\widehat{\eta})\widehat{e}_i / \widehat{e}_j$.

The result is a stochastic matrix $\widehat{\mathbf{P}}$.

This construction assumes that $\mathbf{Q}$ has a unique positive right eigenvector up to scale.

For a finite irreducible nonnegative matrix, the Perron root has a strictly positive right eigenvector unique up to scale.

For long-horizon dominance and convergence, one typically imposes a stronger condition such as primitivity or aperiodicity; the paper uses a positivity condition on $\sum_{t=0}^{\infty} \lambda^t \mathbf{Q}^t$.

In general state spaces this guarantee does not carry over: multiple positive eigenfunctions may exist, and an additional selection condition is needed to pin down the long-term risk-neutral measure.

The general framework in the next section makes that selection condition explicit.

Following Borovička *et al.* [2016], $\widehat{\mathbf{P}}$ is called a **long-term risk-neutral** transition matrix.

The name means that the Perron–Frobenius eigenvalue and eigenvector isolate the part of pricing that dominates long-maturity Arrow claims.

It is not the same transition matrix as the one-period risk-neutral matrix $\bar{\mathbf{P}}$.

In *The Recovery Theorem*, transition independence restricts the SDF to

$$s_{ij} = \exp(-\delta) \frac{m_j}{m_i}$$

for a positive vector m and scalar δ, which pins down the split between s_{ij} and p_{ij}.

Here we drop that restriction.

The question is whether the transition matrix associated with the long-term risk-neutral probability, $\widehat{\mathbf{P}}$, still equals the correctly specified matrix $\mathbf{P}$.

142.2.1 Degenerate martingale component

We start with a three-state economy: recession, normal, and expansion.

The correctly specified transition matrix is deliberately simple.

For trend-stationary consumption and power utility, the SDF is

$$s_{ij} = A \left(\frac{c_j}{c_i} \right)^{-\gamma}, \qquad A = \exp(-\delta - \gamma g_c),$$

where δ is the subjective discount rate, γ is risk aversion, and g_c is trend consumption growth.

This is a case where Ross recovery should return the correctly specified transition matrix.

```
P_true = np.array([
    [0.70, 0.25, 0.05],
    [0.15, 0.65, 0.20],
    [0.05, 0.30, 0.65],
])

c_levels = np.array([0.997, 1.000, 1.003])
state_names = ['recession', 'normal', 'expansion']

δ = -np.log(0.99)    # monthly subjective discount rate
γ_power = 5.0        # risk aversion
g_c = 0.002          # monthly trend growth

# Price Arrow claims as actual probabilities times the power-utility SDF
S_power = (
    np.exp(-δ - γ_power * g_c)
    * (c_levels[None, :] / c_levels[:, None])**(-γ_power)
)
Q_power = S_power * P_true
```

We now compute the one-period risk-neutral matrix and the transition matrix associated with the long-term risk-neutral probability from the same Arrow price matrix.

```
P_bar, q_bonds = risk_neutral_probs(Q_power)
η_hat, exp_η, e_hat, P_hat = perron_frobenius(Q_power)
π_true = stationary_dist(P_true)
π_bar = stationary_dist(P_bar)
π_hat = stationary_dist(P_hat)
```

These two matrices should not be expected to agree.

The row-normalized matrix $\bar{\mathbf{P}}$ is a short-horizon risk-neutral change of measure: it folds the one-period SDF into transition probabilities, so it generally differs from the correctly specified matrix $\mathbf{P}$.

The reason $\hat{\mathbf{P}}$ can equal $\mathbf{P}$ comes from the Perron–Frobenius construction in *The Recovery Theorem*.

In the transition-independent case, the pricing kernel has the form $s_{ij} = \exp(\hat{\eta})\hat{e}_i/\hat{e}_j$.

Substituting this into the Perron–Frobenius transition formula gives

$$\hat{p}_{ij} = \exp(-\hat{\eta})q_{ij}\frac{\hat{e}_j}{\hat{e}_i} = \exp(-\hat{\eta})\left(\exp(\hat{\eta})\frac{\hat{e}_i}{\hat{e}_j}p_{ij}\right)\frac{\hat{e}_j}{\hat{e}_i} = p_{ij}.$$

Thus the transition matrix $\hat{\mathbf{P}}$ associated with the long-term risk-neutral probability cancels the transition-independent part of the SDF.

In this power-utility benchmark, the whole SDF has exactly that form, so the remaining martingale increment should be one and $\hat{\mathbf{P}}$ should coincide with $\mathbf{P}$.

The next calculation checks this by comparing the Perron–Frobenius eigenfunction with c_i^γ and then computing the ratio $\hat{\mathbf{P}}/\mathbf{P}$.

Define the one-period martingale increment

$$\hat{h}_{ij} = \frac{\hat{p}_{ij}}{p_{ij}} = \exp(-\hat{\eta})s_{ij}\frac{\hat{e}_j}{\hat{e}_i}.$$

When $\hat{h}_{ij} = 1$ for every transition, $\hat{\mathbf{P}}$ and $\mathbf{P}$ are the same.

The section *Martingale component* defines this ratio formally in (142.4) and explains why it is a one-period martingale increment.

We ask you to verify in Exercise 142.9.2 that $\hat{e}_i = c_i^\gamma$ and $\exp(\hat{\eta}) = A$ in this example, so that $\hat{h}_{ij} \equiv 1$.

```
H_power = np.divide(P_hat, P_true, out=np.ones_like(P_true), where=P_true > 0)
e_theory = c_levels**y_power

print("Perron-Frobenius eigenfunction: numerical vs c^gamma")
for name, e_num, e_th in zip(state_names, e_hat / e_hat[1],
                             e_theory / e_theory[1]):
    print(f"{name:9s}: {e_num:.6f}  {e_th:.6f}")

print("\nmartingale increment h_hat = P_hat / P")
print(np.round(H_power, 6))

print("\nconditional means under P")
print(np.round((P_true * H_power).sum(axis=1), 6))

print(f"\nmax |h_hat - 1| = "
      f"{np.max(np.abs(H_power[P_true > 0] - 1)):.2e}")
```

```
Perron-Frobenius eigenfunction: numerical vs c^gamma
recession: 0.985090  0.985090
normal   : 1.000000  1.000000
expansion: 1.015090  1.015090

martingale increment h_hat = P_hat / P
[[1. 1. 1.]
 [1. 1. 1.]
 [1. 1. 1.]]

conditional means under P
[1. 1. 1.]

max |h_hat - 1| = 1.33e-15
```

The output separates a short-horizon risk adjustment from the Perron–Frobenius approach.

The one-period risk-neutral matrix $\bar{\mathbf{P}}$ is close to, but not the same as, the correctly specified matrix $\mathbf{P}$.

It changes the transition probabilities because one-period Arrow prices include one-period risk adjustments.

By contrast, the long-term risk-neutral matrix $\hat{\mathbf{P}}$ is exactly the same as $\mathbf{P}$ in this example.

The calculation confirms why: the martingale increment $\hat{h}_{ij}$ is one for every transition.

This is the condition under which Ross recovery returns the correctly specified transition matrix.

In this example, that cancellation exhausts the SDF, so the martingale component is degenerate.

142.3 Martingale component

Let $(\hat{\eta}, \hat{e})$ be the Perron–Frobenius eigenvalue exponent and positive right eigenvector of $\mathbf{Q}$:

$$\mathbf{Q}\hat{e} = \exp(\hat{\eta})\hat{e}.$$

The associated long-term risk-neutral transition matrix is

$$\hat{p}_{ij} = \exp(-\hat{\eta}) q_{ij} \frac{\hat{e}_j}{\hat{e}_i}.$$

To see whether recovery has changed the probability measure, compare each recovered transition probability with the corresponding correctly specified transition probability.

For feasible transitions with $p_{ij} > 0$, define the one-period martingale increment

$$\hat{h}_{ij} = \frac{\hat{p}_{ij}}{p_{ij}}. \tag{142.4}$$

If $\hat{h}_{ij} > 1$, the recovered probability measure assigns more probability to transition (i, j) than the correctly specified probability measure.

If $\hat{h}_{ij} < 1$, it assigns less probability to that transition.

For a fixed current state i, the numbers $\hat{h}_{ij}$ average to one under the correctly specified transition probabilities:

$$\sum_j \hat{h}_{ij} p_{ij} = 1.$$

Thus $\hat{h}_{ij}$ is a one-period martingale increment.

Multiplying these increments along a history of states gives the ratio of the recovered probability measure to the correctly specified probability measure for the whole history.

That ratio process is a martingale, which is why the last term in (142.5) is called a martingale component.

Using (142.1), (142.3), and (142.4), the one-period SDF can be written as

$$s_{ij} = \exp(\hat{\eta}) \frac{\hat{e}_i}{\hat{e}_j} \hat{h}_{ij}. \tag{142.5}$$

The Perron–Frobenius approach therefore separates the SDF into:

Part	Role
$\exp(\hat{\eta})$	deterministic long-run discounting
$\hat{e}_i / \hat{e}_j$	state-dependent long-run term
$\hat{h}_{ij}$	martingale increment that changes probabilities

If $\hat{h}_{ij} = 1$ for every feasible transition, then the transition matrix associated with the recovered probability measure and the correctly specified transition matrix are the same.

This is the condition under which Ross recovery returns the correctly specified transition matrix.

Proposition 142.3.1 (Finite-state martingale component)

Under the finite-state assumptions used in this lecture, for a Markov model with correctly specified transition matrix **P** and Arrow matrix **Q**, the probability measure recovered by Perron–Frobenius Theory returns the correctly specified transition matrix if and only if $\hat{h}_{ij} = 1$ for every transition with $p_{ij} > 0$.

Equivalently, recovery returns the correctly specified transition matrix if and only if the SDF in (142.5) has no non-constant martingale component:

$$s_{ij} = \exp(\hat{\eta}) \frac{\hat{e}_i}{\hat{e}_j}.$$

Proof. Using $q_{ij} = s_{ij} p_{ij}$,

$$\hat{h}_{ij} = \frac{\hat{p}_{ij}}{p_{ij}} = \exp(-\hat{\eta}) s_{ij} \frac{\hat{e}_j}{\hat{e}_i}.$$

Thus $\hat{\mathbf{P}} = \mathbf{P}$ if and only if $\hat{h}_{ij} = 1$ on every feasible transition.

This condition is the same as saying that the SDF takes the form displayed in *Proposition 142.3.1*.

This finite-state implication is a special case of the paper's general identification result.

If a pair (S, P) explains asset prices and H is a positive multiplicative martingale, then the same asset prices are also explained by the changed probability measure P^H together with the adjusted stochastic discount factor

$$S_t^H = S_t \frac{H_0}{H_t}.$$

More generally, any strictly positive martingale can change probability measures, but multiplicativity preserves the Markov structure used here.

Thus Arrow prices alone cannot usually distinguish a change in beliefs from a change in the SDF.

Ross recovery becomes an identification result only after imposing a restriction such as

$$S_t = \exp(-\delta t) \frac{m(X_t)}{m(X_0)},$$

which rules out a nontrivial martingale component.

The power-utility example above illustrates the proposition.

In that benchmark, the martingale increment $\hat{h}_{ij}$ is identically one.

142.4 From matrices to the general framework

The finite-state calculation has three objects:

1. the correctly specified transition probabilities p_{ij},
2. the SDF increments s_{ij},
3. the Arrow prices $q_{ij} = s_{ij} p_{ij}$.

It also has one diagnostic object:

$$\hat{h}_{ij} = \frac{\hat{p}_{ij}}{p_{ij}} = \exp(-\hat{\eta}) s_{ij} \frac{\hat{e}_j}{\hat{e}_i}.$$

The numbers $\hat{h}_{ij}$ are one-period martingale increments.

They change the probability of a one-period transition from p_{ij} to $\hat{p}_{ij} = \hat{h}_{ij} p_{ij}$.

The general framework in Borovička *et al.* [2016] does the same thing without assuming that states are finite.

The transition matrix becomes a Markov probability measure, the Arrow price matrix becomes a family of pricing operators, and the one-period ratios $\hat{h}_{ij}$ become increments of a positive multiplicative martingale.

The point of this section is to build that dictionary.

142.4.1 Probability space and state

Start with a probability space $(\Omega, \mathcal{F}, P)$.

Here P is the correctly specified probability measure.

In the rational-expectations interpretation of the paper, this is the actual, or original, probability measure governing the state.

In the finite-state section, P was represented by the transition matrix $\mathbf{P} = [p_{ij}]$.

The index set is either discrete time, $\mathbb{T} = \{0, 1, 2, ...\}$, or continuous time, $\mathbb{T} = \mathbb{R}_+$.

The main state process is $X = \{X_t : t \in \mathbb{T}\}$, which is stationary and Markov under P.

A second process $W = \{W_t : t \in \mathbb{T}\}$ records shocks that drive X and other economic quantities.

In discrete time, the shock increment between dates t and $t+1$ is

$$\Delta W_{t+1} = W_{t+1} - W_t.$$

The known function ϕ_x maps today's state and the next shock increment into tomorrow's state.

The discrete-time state evolution is

$$X_{t+1} = \phi_x(X_t, \Delta W_{t+1}).$$

The conditional law generated by this equation is the general-state replacement for the finite-matrix row indexed by the current state x.

Assumption 142.4.1 (Markov state and shock increments)

The process X is ergodic under P.

The conditional distribution of ΔW_{t+1} given X_t is time invariant and independent of past shock histories conditioned on X_t.

The filtration $\{\mathcal{F}_t\}$ is generated by the initial condition X_0 and by the history of shocks through date t.

142.4.2 Information and Y

The Markov state and the information that reveals shocks need not coincide.

The state X_t is observed at date t.

The next shock ΔW_{t+1} need not be directly observed from the pair (X_t, X_{t+1}).

If (X_t, X_{t+1}) does reveal ΔW_{t+1}, then X alone carries the relevant shock information.

If it does not, introduce an auxiliary process $Y = \{Y_t\}$ with stationary increments.

The known function ϕ_y maps today's state and the next shock increment into the increment of Y.

The discrete-time evolution for the auxiliary increment is

$$Y_{t+1} - Y_t = \phi_y(X_t, \Delta W_{t+1}).$$

The pair $(X_{t+1}, Y_{t+1} - Y_t)$ is then rich enough, together with X_t, to recover the shock increment ΔW_{t+1}.

This device lets the model handle shocks or growth components that affect payoffs and SDFs but are not fully summarized by the next Markov state alone.

Write the enlarged process as $Z = (X, Y)$.

The process Z is Markov with a triangular structure: the conditional distribution of $(X_{t+1}, Y_{t+1} - Y_t)$ depends on the past only through X_t.

Histories of Z, together with X_0, generate the same information as the shock history.

This is why the next Perron–Frobenius problem can first be posed with eigenfunctions of X alone.

The section *Additional state vector* later returns to what changes when the eigenfunction is allowed to depend on Y as well.

142.4.3 Multiplicative functionals

The general framework needs a way to describe objects that compound over time.

This is the role of a positive multiplicative functional $M = \{M_t\}$.

Its log increment is a function of today's state and the next shock increment:

$$\log \frac{M_{t+1}}{M_t} = \kappa(X_t, \Delta W_{t+1}).$$

Equivalently,

$$\frac{M_{t+1}}{M_t} = \exp\{\kappa(X_t, \Delta W_{t+1})\}.$$

Thus M_t / M_0 is a product of positive one-period increments.

This is the Condition-1 version of a multiplicative functional used in the paper.

The formal definition is slightly broader, but this form covers the models studied below.

Under *Assumption 142.4.1*, the logarithm of M has stationary increments.

Products and reciprocals of positive multiplicative functionals are again positive multiplicative functionals.

Exponential functions of linear combinations of the components of Y are examples.

Stochastic discount factors, stochastic growth factors, and positive multiplicative martingales are all modeled this way.

In the finite-state model, the SDF increment s_{ij} is one example of a multiplicative-functional increment.

The ratio h_{ij} that changes probabilities is another.

142.4.4 Stochastic discount factors and pricing operators

A stochastic discount factor $S = \{S_t\}$ is a positive multiplicative functional with $S_0 = 1$ and finite first moments conditional on X_0.

Let Φ_t be a bounded payoff measurable with respect to the date-t information.

The date-τ price of Φ_t is

$$\Pi_{\tau,t}(\Phi_t) = E\left[\frac{S_t}{S_\tau}\Phi_t \mid \mathcal{F}_\tau\right].$$

The ratio S_t/S_τ is the stochastic discount factor from date t back to date τ.

If the payoff is a bounded function $f(X_t)$ of the future Markov state, this pricing formula defines a horizon-t operator Q_t by

$$[Q_t f](x) = E[S_t f(X_t) \mid X_0 = x].$$

This operator is the general-state analogue of multiplying a payoff vector by the Arrow-price matrix $\mathbf{Q}$.

To see the connection, suppose again that the state space is finite and $t = 1$.

Then

$$[Q_1 f]_i = \sum_j s_{ij} p_{ij} f_j = \sum_j q_{ij} f_j.$$

Thus Q_1 is exactly the matrix $\mathbf{Q}$.

In discrete time, the multiplicative property of S implies that Q_t is obtained by applying the one-period operator Q_1 repeatedly.

In continuous time, the family $\{Q_t : t \geq 0\}$ is a semigroup of pricing operators.

142.4.5 Martingales and equivalent probability measures

Different stochastic discount factor / probability pairs can produce the same pricing operators, and that flexibility is the source of the identification problem.

For the Markov setting used here, the probability changes of interest are generated by positive multiplicative martingales.

At the level of probability changes, let $H = \{H_t\}$ be a strictly positive martingale with $E[H_0] = 1$ under P.

For an event A observable by date τ, the changed probability measure P^H is defined by

$$P^H(A) = E[1_A H_\tau].$$

The law of iterated expectations makes this definition consistent across dates.

When H is also a multiplicative functional, making it a multiplicative martingale, the change of probability preserves the Markov structure of Z.

The SDF that represents the same prices under P^H is

$$S_t^H = S_t \frac{H_0}{H_t}.$$

Thus the same pricing operators can be represented by the pair (S, P) or by the pair (S^H, P^H).

In the finite-state model, this is just

$$p_{ij}^H = h_{ij} p_{ij}, \qquad s_{ij}^H = \frac{s_{ij}}{h_{ij}},$$

so that

$$s_{ij}^H p_{ij}^H = s_{ij} p_{ij} = q_{ij}.$$

The same Arrow prices can therefore be explained by changing the probability measure and offsetting that change in the SDF.

This is also the precise sense in which Arrow prices alone do not identify beliefs.

142.4.6 What Perron–Frobenius recovers

Now return to the Perron–Frobenius step.

The finite-state equation was (142.2).

The general-state replacement is an eigenfunction problem for the pricing operators of the kind studied in *Long-term Risk: An Operator Approach*: find a scalar $\hat{\eta}$ and a positive function $\hat{e}$ such that, for every horizon t,

$$[Q_t \hat{e}](x) = \exp(\hat{\eta} t) \hat{e}(x). \tag{142.6}$$

The positive function $\hat{e}$ is the general-state counterpart of the Perron–Frobenius eigenvector.

The scalar $\hat{\eta}$ is the log eigenvalue.

In finite states, $\hat{e}$ is just a positive vector with one entry for each state.

In general state spaces, $\hat{e}(x)$ is a positive function of the current state.

Its job is to record the state-dependent part of long-horizon valuation.

Equation (142.6) says that a future payoff equal to $\hat{e}(X_t)$ has date-0 price $\exp(\hat{\eta} t)\hat{e}(X_0)$.

Thus $\hat{\eta}$ gives the common growth or discount rate, while $\hat{e}$ gives the state-dependent scaling.

Since eigenfunctions are defined only up to scale, uniqueness always means uniqueness up to multiplication by a positive constant.

In general state spaces, existence of a positive eigenfunction is also a substantive condition.

The eigenfunction equation implies the conditional moment restriction

$$E[S_t \hat{e}(X_t) \mid \mathcal{F}_\tau] = \exp((t-\tau)\hat{\eta}) S_\tau \hat{e}(X_\tau), \qquad t \geq \tau.$$

Use this restriction to define

$$\frac{\hat{H}_t}{\hat{H}_0} = \exp(-\hat{\eta} t) S_t \frac{\hat{e}(X_t)}{\hat{e}(X_0)}. \tag{142.7}$$

This process is a martingale because, for $t \geq \tau$,

$$
\begin{aligned}
E\left[\frac{\hat{H}_t}{\hat{H}_0} \mid \mathcal{F}_\tau\right] &= \frac{\exp(-\hat{\eta}t)}{\hat{e}(X_0)} E[S_t \hat{e}(X_t) \mid \mathcal{F}_\tau] \\
&= \frac{\exp(-\hat{\eta}t)}{\hat{e}(X_0)} \exp((t-\tau)\hat{\eta}) S_\tau \hat{e}(X_\tau) \\
&= \exp(-\hat{\eta}\tau) S_\tau \frac{\hat{e}(X_\tau)}{\hat{e}(X_0)} = \frac{\hat{H}_\tau}{\hat{H}_0}.
\end{aligned}
$$

The process is positive because S and $\hat{e}$ are positive.

Its one-period increment is

$$
\frac{\hat{H}_{t+1}}{\hat{H}_t} = \exp(-\hat{\eta}) \frac{S_{t+1}}{S_t} \frac{\hat{e}(X_{t+1})}{\hat{e}(X_t)}. \tag{142.8}
$$

In the finite-state model, when $X_t = i$ and $X_{t+1} = j$, (142.8) becomes

$$
\frac{\hat{H}_{t+1}}{\hat{H}_t} = \exp(-\hat{\eta}) s_{ij} \frac{\hat{e}_j}{\hat{e}_i} = \hat{h}_{ij} = \frac{\hat{p}_{ij}}{p_{ij}}.
$$

This is the same three-component decomposition as (142.5).

In finite states, (142.5) is equivalent to

$$
\frac{S_{t+1}}{S_t} = \exp(\hat{\eta}) \frac{\hat{e}(X_t)}{\hat{e}(X_{t+1})} \frac{\hat{H}_{t+1}}{\hat{H}_t}.
$$

The only change is notation: $\hat{h}_{ij}$ is the one-period density ratio in a finite Markov chain, while $\hat{H}_{t+1}/\hat{H}_t$ is the corresponding one-period density ratio in the general Markov setting.

Conditional on X_0, the likelihood ratio for histories through date t is $\hat{H}_t/\hat{H}_0$.

For the unconditional measure on $\mathcal{F}_t$, the Radon–Nikodym density is $\hat{H}_t$, where $\hat{H}_0$ adjusts the initial distribution.

If $\hat{H}_{t+1}/\hat{H}_t$ is not identically one, the recovered probability measure differs from the correctly specified probability measure.

We show the restriction that rules out this difference in the next section.

Note

The objects $(\hat{e}, \hat{\eta}, \hat{H})$ are the ones that *Long-term Risk: An Operator Approach* writes as $(\phi, \rho, \hat{M})$, working there with a generic multiplicative functional M in place of the stochastic discount factor S used here; what that lecture calls the **twisted measure** is what this lecture calls the long-term risk-neutral measure $\hat{\mathbf{P}} = P^{\hat{H}}$.

One symbol does not travel.

Here Q_t is the semigroup built from the stochastic discount factor alone, which that lecture writes $\mathbb{S}_t$; there, $\mathbb{Q}_t$ is instead the cash-flow valuation semigroup built from $Q = GS$, which also carries a stochastic growth factor.

142.4.7 Selection and recovery

In finite irreducible matrix problems, Perron–Frobenius theory gives a unique positive eigenvector up to scale, so the recovered transition matrix is pinned down by **Q**.

In general state spaces, a positive eigenfunction need not exist, and multiple positive eigenfunctions may solve the same pricing operator problem when one does.

The paper therefore imposes a selection condition on the probability measure induced by the candidate eigenfunction.

This is the same difficulty that *Long-term Risk: An Operator Approach* treats under "Stability of the twisted process"; see *Definition 141.5.6* there.

Assumption 142.4.2 (Ergodicity of the recovered measure)

The process X is stationary and ergodic under $P^{\hat{H}}$, the probability measure induced by the multiplicative martingale $\hat{H}$ defined in the previous section.

Proposition 142.4.1 (Uniqueness of the Perron–Frobenius solution)

There is at most one solution $(\hat{e}, \hat{\eta})$ to the Perron–Frobenius problem such that X is stationary and ergodic under the induced probability measure $P^{\hat{H}}$.

This selected solution, when it exists, identifies the long-term risk-neutral measure.

It does not by itself identify subjective beliefs.

To make recovery identify beliefs, an additional restriction is needed on the SDF.

The restriction used by Ross recovery is the paper's Condition 4:

Assumption 142.4.3

Let

$$S_t = \exp(-\delta t)\frac{m(X_t)}{m(X_0)}$$

for some positive function m and real number δ.

This is the transition-independence restriction from *The Recovery Theorem*, imposed on the SDF representation whose probability measure one wants to recover.

Under this restriction, setting $\hat{e} = 1/m$ and $\hat{\eta} = -\delta$ gives

$$\frac{\hat{H}_t}{\hat{H}_0} = \exp(\delta t)\left[\exp(-\delta t)\frac{m(X_t)}{m(X_0)}\right]\frac{1/m(X_t)}{1/m(X_0)} = 1.$$

Thus the martingale component is identically one after normalization.

This is the general-state version of the finite-state condition $\hat{h}_{ij} = 1$ for every feasible transition.

If this martingale is not identically one, the recovered probability measure absorbs it and generally differs from the correctly specified probability measure.

We will see a few important examples of this in the next section.

The section *Additional state vector* returns to what happens when the eigenfunction is allowed to depend on the auxiliary process Y.

142.4.8 Continuous-time version

Let's briefly restate the framework in continuous time before discussing examples where recovery fails.

We introduce the diffusion notation because the long-run risk example below is written in continuous time.

The objects are the same as before:

- X is the Markov state,
- Y records additional growing or shock-revealing components,
- M is a positive multiplicative functional, such as an SDF, a cash-flow growth process, or a martingale used to change probabilities.

In the continuous-time version, W is a Brownian motion.

The state, auxiliary process, and multiplicative functional satisfy

$$\begin{aligned} dX_t &= \mu_x(X_t)dt + \sigma_x(X_t)dW_t, \\ dY_t &= \mu_y(X_t)dt + \sigma_y(X_t)dW_t, \\ d\log M_t &= \beta(X_t)dt + \alpha(X_t)\cdot dW_t. \end{aligned}$$

Here μ_x and μ_y are drift functions, while σ_x and σ_y are shock-exposure matrices.

The function β is the drift of $\log M$, and α is the Brownian shock exposure of $\log M$.

The invertibility assumption on the stacked shock-exposure matrix is the continuous-time counterpart of the discrete-time condition that $(X_{t+1}, Y_{t+1} - Y_t)$ reveals the shock increment.

It lets the history of $Z = (X, Y)$ reveal the Brownian information.

For M to be a local martingale, its drift must satisfy

$$\beta(x) = -\frac{1}{2}\alpha(x)\cdot\alpha(x).$$

This follows from Ito's formula:

$$\frac{dM_t}{M_t} = \left(\beta(X_t) + \frac{1}{2}\alpha(X_t)\cdot\alpha(X_t)\right)dt + \alpha(X_t)\cdot dW_t.$$

A local martingale has zero drift in dM_t/M_t, which gives the displayed restriction.

Additional integrability conditions then ensure that this local martingale is a true martingale.

Under the probability measure induced by a martingale H with this exposure α, $\widetilde{W}_t = W_t - \int_0^t \alpha(X_s)ds$ is Brownian.

With this sign convention, the drift of X changes from μ_x to $\mu_x + \sigma_x\alpha$, and the drift of Y changes from μ_y to $\mu_y + \sigma_y\alpha$.

This is the continuous-time analogue of replacing p_{ij} by $h_{ij}p_{ij}$ in the finite-state model.

The Markov and triangular structure of Z is preserved, which is why the same Perron–Frobenius decomposition can be applied.

142.5 When recovery fails

Now let's discuss a few examples where the recovered probability measure differs from the correctly specified probability measure.

142.5.1 Recursive utility

We now use the martingale component to see when the recovered probability measure differs from the correctly specified probability measure.

In the previous example, all risk adjustment in the SDF could be written as a ratio of a function of today's state to a function of tomorrow's state.

The Perron–Frobenius transition formula cancels exactly that kind of term.

Recursive utility adds a continuation-value term.

The key point is that this term behaves like the martingale increment defined above.

For the unit-EIS Epstein–Zin case in Borovička *et al.* [2016], with $C_t = \exp(g_c t)c(X_t)$, write the translated continuation value as $V_t = g_c t + v(X_t)$, and define

$$v_i^* = \exp((1-\gamma)v_i).$$

The SDF is

$$s_{ij} = \exp(-\delta - g_c)\frac{c_i}{c_j}\frac{v_j^*}{\sum_k p_{ik} v_k^*}.$$

In this unit-EIS example, the Perron–Frobenius eigenfunction is $\widehat{e}_j = c_j$ and $\widehat{\eta} = -(\delta + g_c)$.

Applying the Perron–Frobenius transition formula therefore leaves

$$\widehat{p}_{ij} = p_{ij}\frac{v_j^*}{\sum_k p_{ik} v_k^*}.$$

The denominator is the conditional expectation of v_j^* given current state i.

Therefore the last fraction has conditional mean one under **P**.

It is therefore a martingale increment.

When v^* is not constant, that ratio varies across next-period states.

That variation is why the probability measure recovered by Perron–Frobenius Theory no longer gives the correctly specified transition matrix.

The next cell solves the finite-state continuation-value equation and builds the SDF.

```
def solve_ez_unit_eis(P, c, δ, γ, g_c, tol=1e-12, max_iter=10_000):
    """Finite-state unit-EIS Epstein-Zin continuation values and SDF."""
    β = np.exp(-δ)
    log_c = np.log(c)
    n = len(c)
    flow = (1 - β) * log_c + β * g_c

    if abs(γ - 1) < 1e-10:
        v = linalg.solve(np.eye(n) - β * P, flow)
        v_star = np.ones(n)
```

(continues on next page)

(continued from previous page)

```
        Pv_star = np.ones(n)
    else:
        v = log_c.copy()
        for _ in range(max_iter):
            v_star = np.exp((1 - γ) * v)
            Pv_star = P @ v_star
            v_new = flow + β / (1 - γ) * np.log(Pv_star)
            if np.max(np.abs(v_new - v)) < tol:
                v = v_new
                break
            v = v_new
        else:
            raise ValueError("Epstein-Zin fixed point did not converge.")

        v_star = np.exp((1 - γ) * v)
        Pv_star = P @ v_star

    S = (
        np.exp(-δ - g_c)
        * (c[:, None] / c[None, :])
        * (v_star[None, :] / Pv_star[:, None])
    )

    return v, v_star, S
```

At log utility, v^* is constant and the martingale increment is one.

As risk aversion rises, continuation values matter more.

The recovered probability measure then differs more from the correctly specified probability measure.

To make the mechanism visible in a small three-state example, the figure below uses the more dispersed consumption vector

$$c = (0.85, 1.00, 1.15).$$

The heatmap reports percentage deviations of the martingale increment from one: $100(\hat{h}_{ij} - 1)$.

Positive entries are transitions that receive more probability under the recovered probability measure than under the correctly specified probability measure.

The right panel reports the increase in the recovered recession probability, measured in percentage points.

```
c_recursive = np.array([0.85, 1.00, 1.15])
γ_demo = 10.0
_, _, S_demo = solve_ez_unit_eis(P_true, c_recursive, δ, γ_demo, g_c)
Q_demo = S_demo * P_true
H_demo, _, _, P_hat_demo = martingale_increment(Q_demo, P_true)
H_dev = 100 * (H_demo - 1)

γ_grid = np.linspace(1, 15, 80)
rec_prob = []
for γ in γ_grid:
    _, _, S_g = solve_ez_unit_eis(P_true, c_recursive, δ, γ, g_c)
    Q_g = S_g * P_true
    _, _, _, P_hat_g = martingale_increment(Q_g, P_true)
    rec_prob.append(stationary_dist(P_hat_g)[0])
rec_prob = np.array(rec_prob)
```

(continues on next page)

(continued from previous page)

```
rec_prob_gain = 100 * (rec_prob - π_true[0])

fig, axes = plt.subplots(1, 2, figsize=(12, 4.5))

bound = np.max(np.abs(H_dev))
im = axes[0].imshow(H_dev, cmap='RdBu_r', vmin=-bound, vmax=bound)
axes[0].set_xticks(range(3))
axes[0].set_yticks(range(3))
axes[0].set_xticklabels(state_names, rotation=20)
axes[0].set_yticklabels(state_names)
axes[0].set_xlabel('next state')
axes[0].set_ylabel(r'current state')
axes[0].set_title(r'martingale increment, $\gamma=10$')

for i in range(3):
    for j in range(3):
        axes[0].text(j, i, f"{H_dev[i, j]:.1f}",
                     ha='center', va='center', fontsize=9)
plt.colorbar(im, ax=axes[0], fraction=0.046, pad=0.04,
             label=r'$100(\hat h_{ij}-1)$')

axes[1].plot(y_grid, rec_prob_gain, lw=2.5)
axes[1].axhline(0, ls='--', lw=1.5, color='0.5')
axes[1].set_xlabel(r"risk aversion $\gamma$")
axes[1].set_ylabel('increase in recession probability\n(percentage points)')
axes[1].set_title('recovered recession probability')
axes[1].set_ylim(0, rec_prob_gain.max() * 1.08)

plt.tight_layout()
plt.show()
```

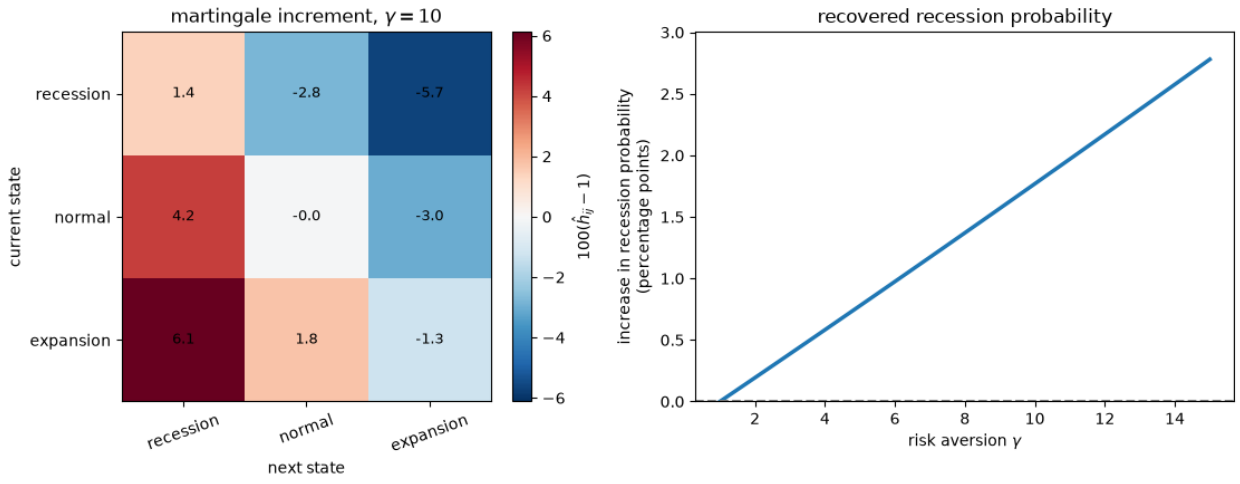


Fig. 142.1: Recursive utility generates a nonconstant martingale increment.

Recursive utility makes the recovered probability measure assign more probability to recession transitions.

At $\gamma = 10$, transitions into recession receive more probability under the recovered probability measure, while transitions into expansion receive less.

As risk aversion rises, the stationary recession probability under the recovered probability measure moves further above its correctly specified value.

Thus, as the continuation-value term creates a nonconstant $\hat{h}_{ij}$, the transition matrix associated with the long-term risk-neutral probability no longer equals the correctly specified transition matrix.

142.5.2 Permanent shocks

Recursive utility gives one nonconstant martingale component.

Permanent shocks provide another, but this example has to be read carefully, because what it delivers depends on the information set on which recovery is run.

Suppose consumption has a permanent shock,

$$\log C_{t+1} - \log C_t = g + x(X_{t+1}) - x(X_t) + \sigma\varepsilon_{t+1},$$

where ε_{t+1} is independent over time and independent of the Markov state X.

With power utility, the SDF contains

$$\exp(-\delta - \gamma g)\exp\{-\gamma[x(X_{t+1}) - x(X_t)]\}\exp(-\gamma\sigma\varepsilon_{t+1}).$$

The middle term depends only on the current and next Markov states.

It is a ratio of state functions, so the Perron–Frobenius transition formula can cancel it.

The last term depends on the new shock ε_{t+1}, which is not summarized by the Markov state.

After dividing by its conditional mean, that term becomes a martingale increment:

$$\frac{\exp(-\gamma\sigma\varepsilon_{t+1})}{E[\exp(-\gamma\sigma\varepsilon_{t+1})]}.$$

Whether that martingale increment breaks recovery depends on the information set on which the eigenfunction problem is posed, so we take the two cases in turn.

Consider first an analyst who prices claims written on the Markov state X alone.

Taking expectations over ε_{t+1}, the Arrow prices over the states of X are

$$q_{ij} = \exp(-\delta - \gamma g)\, E[\exp(-\gamma\sigma\varepsilon_{t+1})]\, \frac{\exp(-\gamma x_j)}{\exp(-\gamma x_i)}\, p_{ij}.$$

Because ε_{t+1} is independent of X, the permanent shock contributes the same constant to every entry of $\mathbf{Q}$ and factors out.

What is left is exactly transition independent in X, with $\hat{e}_i = \exp(\gamma x_i)$.

So Perron–Frobenius recovery of the transition matrix of X is exact here: the martingale increment over the states of X is identically one.

The next cell confirms this in a three-state example.

```
x_perm = np.array([-0.02, 0.0, 0.02])    # state effect on log consumption
σ_perm = 0.01                            # scale of the permanent shock
E_perm = np.exp(0.5 * (γ_power * σ_perm)**2)    # E[exp(-γ σ ε)] for ε ~ N(0, 1)

# Arrow prices over X: the permanent shock enters only through the constant E_perm
S_perm = (np.exp(-δ - γ_power * g_c) * E_perm
```

(continues on next page)

(continued from previous page)

```
          * np.exp(-y_power * x_perm)[None, :]
          / np.exp(-y_power * x_perm)[:, None])
Q_perm = S_perm * P_true

H_perm, _, e_perm, P_hat_perm = martingale_increment(Q_perm, P_true)

print("eigenfunction: numerical vs exp(gamma x)")
print(np.round(e_perm / e_perm[1], 6))
print(np.round(np.exp(y_power * x_perm) / np.exp(y_power * x_perm[1]), 6))
print(f"\nmax |P_hat - P| = {np.max(np.abs(P_hat_perm - P_true)):.1e}")
print(f"max |h_hat - 1| = {np.max(np.abs(H_perm[P_true > 0] - 1)):.1e}")
```

```
eigenfunction: numerical vs exp(gamma x)
[0.904837 1.       1.105171]
[0.904837 1.       1.105171]

max |P_hat - P| = 3.9e-16
max |h_hat - 1| = 1.3e-15
```

The recovered eigenfunction is $\exp(\gamma x_i)$ and both discrepancies are at machine precision.

Now enlarge the filtration so that it reveals consumption, and therefore reveals ε_{t+1} as well.

Relative to investors' information about (X, ε), the martingale increment displayed above is not identically one: it tilts the distribution of the permanent shock toward low realizations.

So the recovered law of consumption differs from investors' beliefs even under ordinary power utility, although the recovered law of the Markov state does not.

The conclusion can now be stated precisely.

Permanent consumption shocks make the recovered probability measure differ from investors' beliefs on the filtration that reveals those shocks, even under ordinary power utility.

Recovery over X alone also breaks down as soon as the permanent shock stops being independent of X, since the shock then no longer factors out of $\mathbf{Q}$ into a constant.

Enlarging the state or information structure to account for the shock does not repair recovery; it leads to the identification problem discussed in *Additional state vector*.

142.5.3 Long-run risk

We now move from small finite-state examples to a standard continuous-time macro-finance model.

The model is the Bansal–Yaron long-run risk model, using the calibration reported by Borovička *et al.* [2016].

The point is to compare the recovered probability measure with the correctly specified probability measure in a standard macro-finance model.

The construction has the same structure as before.

We first write the correctly specified state dynamics, then compute the probability measure implied by the Perron–Frobenius approach.

The state vector $X_t = (X_{1t}, X_{2t})'$ follows

$$
\begin{aligned}
dX_{1t} &= [\mu_{11}(X_{1t} - \iota_1) + \mu_{12}(X_{2t} - \iota_2)]dt + \sqrt{X_{2t}}\sigma_1 dW_t, \\
dX_{2t} &= \mu_{22}(X_{2t} - \iota_2)dt + \sqrt{X_{2t}}\sigma_2 dW_t.
\end{aligned}
$$

Log consumption has an affine drift and the same volatility scaling:

$$d \log C_t = [\beta_{c0} + \beta_{c1}(X_{1t} - \iota_1) + \beta_{c2}(X_{2t} - \iota_2)]dt + \sqrt{X_{2t}}\, \alpha_c \cdot dW_t.$$

Here X_1 is predictable consumption growth and X_2 is a stochastic volatility state.

Two conventions are used throughout this section and both matter for reading the code.

First, every drift is written relative to the long-run means (ι_1, ι_2), so a coefficient triple $(\beta_0, \beta_1, \beta_2)$ always means $\beta_0 + \beta_1(x_1 - \iota_1) + \beta_2(x_2 - \iota_2)$.

Second, each diffusion loading is scaled by $\sqrt{X_{2t}}$, so an Ito correction $\frac{1}{2}|\alpha|^2$ enters the coefficient on x_2 rather than the constant.

The two conventions work together, and it is worth seeing how, because the code below carries a term $-\frac{1}{2}|\alpha_{H^*}|^2$ in the x_2 coefficient of the drift of $\log S$ and a term $-\frac{1}{2}\iota_2|\alpha_{H^*}|^2$ in its constant, where α_{H^*} is the shock exposure of the continuation-value martingale defined below.

That is one correction, not two.

The continuation-value martingale contributes the drift $-\frac{1}{2}x_2|\alpha_{H^*}|^2$ to $\log S$, and writing this single term in centered form splits it into $-\frac{1}{2}|\alpha_{H^*}|^2$ multiplying $(x_2 - \iota_2)$ and $-\frac{1}{2}\iota_2|\alpha_{H^*}|^2$ left over in the constant.

The calibration below sets $\beta_{c1} = 1$, so X_1 *is* predictable consumption growth.

The representative agent has Epstein–Zin utility with unit elasticity of intertemporal substitution.

The continuation value introduces the continuous-time analogue of the martingale component above.

We denote that process by H^*, and the SDF satisfies

$$d \log S_t = -\delta dt - d \log C_t + d \log H_t^*.$$

Here H^* is the continuation-value martingale entering the Epstein–Zin SDF.

The multiplicative martingale $\hat{H}$ associated with the Perron–Frobenius problem is obtained only after also incorporating the Perron–Frobenius eigenfunction.

In models with martingale components in consumption growth, H^* and $\hat{H}$ need not coincide.

The next cell sets the calibration.

```
lrr_params = dict(
    δ=0.002,
    γ=10.0,
    μ11=-0.021,
    μ12=0.0,
    μ22=-0.013,
    ι1=0.0,
    ι2=1.0,
    σ1=np.array([0.0, 0.00034, 0.0]),
    σ2=np.array([0.0, 0.0, -0.038]),
    β_c0=0.0015,
    β_c1=1.0,
    β_c2=0.0,
    α_c=np.array([0.0078, 0.0, 0.0]),
)
```

The next code block computes how the different probability measures change the drift of the state vector.

The first quantity is the continuation value.

In this affine model, the translated continuation value is linear in the state:

$$v(x) = v_0 + v_1 x_1 + v_2 x_2.$$

This is why we call v_1 and v_2 slopes.

They are the derivatives of the continuation value with respect to predictable growth and volatility.

These slopes enter the continuation-value martingale H^*.

In the code, this martingale has shock exposure

$$\alpha_{H^*} = (1-\gamma)(\alpha_c + \sigma_1 v_1 + \sigma_2 v_2).$$

Since the SDF is $d\log S_t = -\delta dt - d\log C_t + d\log H_t^*$, its shock exposure is

$$\alpha_S = -\alpha_c + \alpha_{H^*}.$$

This vector α_S drives the one-period risk-neutral change of measure.

The second quantity is the Perron–Frobenius eigenfunction.

It is exponential-affine:

$$\hat{e}(x) = \exp(e_0 + e_1 x_1 + e_2 x_2).$$

Thus e_1 and e_2 are slopes of the log eigenfunction.

Because X_1 and X_2 have shock loadings σ_1 and σ_2, the Perron–Frobenius eigenfunction contributes the additional shock exposure

$$\sigma_1 e_1 + \sigma_2 e_2.$$

Therefore the one-period risk-neutral dynamics use only α_S, while the dynamics under the long-term risk-neutral measure use

$$\alpha_S + \sigma_1 e_1 + \sigma_2 e_2.$$

The functions below follow this order: compute (v_1, v_2), compute α_S and (e_1, e_2), and then translate these shock exposures into drifts for X.

```
def solve_value_function(p):
    """Slopes of the affine continuation value."""
    δ, γ = p["δ"], p["γ"]
    μ11, μ12, μ22 = p["μ11"], p["μ12"], p["μ22"]
    σ1, σ2 = p["σ1"], p["σ2"]
    β_c1, β_c2 = p["β_c1"], p["β_c2"]
    α_c = p["α_c"]

    # v1 is the coefficient on predictable growth in v(x).
    v1 = β_c1 / (δ - μ11)

    # v2 is the coefficient on volatility.
    # In the affine model it is the stable root of a scalar quadratic.
    A_vec = α_c + σ1 * v1
    B_vec = σ2

    a = 0.5 * (1 - γ) * np.dot(B_vec, B_vec)
    b = (μ22 - δ) + (1 - γ) * np.dot(A_vec, B_vec)
```

(continues on next page)

(continued from previous page)

```
    c = β_c2 + μ12 * v1 + 0.5 * (1 - γ) * np.dot(A_vec, A_vec)

    disc = b**2 - 4 * a * c
    if disc < 0:
        raise ValueError("Value function does not exist for these parameters.")

    v2 = (-b - np.sqrt(disc)) / (2 * a)
    return v1, v2


def solve_pf_lrr(p, v1, v2):
    """Perron-Frobenius eigenfunction slopes and the SDF diffusion loading."""
    δ, γ = p["δ"], p["γ"]
    μ11, μ12, μ22 = p["μ11"], p["μ12"], p["μ22"]
    ι1, ι2 = p["ι1"], p["ι2"]
    σ1, σ2 = p["σ1"], p["σ2"]
    α_c = p["α_c"]
    β_c0, β_c1, β_c2 = p["β_c0"], p["β_c1"], p["β_c2"]

    # Continuation-value martingale exposure and SDF exposure.
    α_h_star = (1 - γ) * (α_c + σ1 * v1 + σ2 * v2)
    α_s = -α_c + α_h_star

    # Drift coefficients of log S before the Perron-Frobenius decomposition.
    β_s11 = -β_c1
    β_s12 = -β_c2 - 0.5 * np.dot(α_h_star, α_h_star)
    β_s0 = -δ - β_c0 - 0.5 * ι2 * np.dot(α_h_star, α_h_star)

    # e1 and e2 are coefficients in log e(x) = e0 + e1 x1 + e2 x2.
    e1 = -β_s11 / μ11

    # e2 solves the remaining quadratic from the Perron-Frobenius eigenvalue equation.
    const = (β_s12 + 0.5 * np.dot(α_s, α_s)
             + e1 * (μ12 + np.dot(σ1, α_s))
             + 0.5 * e1**2 * np.dot(σ1, σ1))
    lin = μ22 + np.dot(σ2, α_s) + e1 * np.dot(σ1, σ2)
    quad = 0.5 * np.dot(σ2, σ2)

    disc = lin**2 - 4 * quad * const
    roots = [(-lin - np.sqrt(disc)) / (2 * quad),
             (-lin + np.sqrt(disc)) / (2 * quad)]

    # Selection rule: keep the root under which X is stationary and ergodic under the
    # induced measure.  In this affine model that requires the twisted volatility
    # process to mean revert (mu_hat_22 < 0) to a positive long-run mean.
    selected = []
    for e2 in roots:
        α_h = α_s + σ1 * e1 + σ2 * e2
        μ_hat_22 = μ22 + np.dot(σ2, α_h)
        ι_hat_2 = (μ22 / μ_hat_22) * ι2
        if μ_hat_22 < 0 and ι_hat_2 > 0:
            eta = (β_s0 - β_s11 * ι1 - β_s12 * ι2
                   - e1 * (μ11 * ι1 + μ12 * ι2) - e2 * μ22 * ι2)
            selected.append((eta, e2))

    if len(selected) != 1:
        raise ValueError("Selection condition does not pin down a unique eigenfunction
```

(continues on next page)

(continued from previous page)

```
↪")

    eta, e2 = selected[0]
    return e1, e2, eta, α_s


def recovered_lrr_dynamics(p, e1, e2, α_s):
    """State dynamics under the long-term risk-neutral measure."""
    μ11, μ12, μ22 = p["μ11"], p["μ12"], p["μ22"]
    ι1, ι2 = p["ι1"], p["ι2"]
    σ1, σ2 = p["σ1"], p["σ2"]

    # The long-term risk-neutral measure uses the SDF exposure plus the
    # eigenfunction exposure.
    α_h = α_s + σ1 * e1 + σ2 * e2

    # A diffusion change of measure shifts each drift by sigma_i dot alpha_h.
    μ_hat_11 = μ11
    μ_hat_12 = μ12 + np.dot(σ1, α_h)
    μ_hat_22 = μ22 + np.dot(σ2, α_h)

    # Rewrite the shifted drift in mean-reversion form.
    ι_hat_2 = (μ22 / μ_hat_22) * ι2
    ι_hat_1 = ι1 + (μ12 * ι2 - μ_hat_12 * ι_hat_2) / μ11

    return dict(
        μ11=μ_hat_11,
        μ12=μ_hat_12,
        μ22=μ_hat_22,
        ι1=ι_hat_1,
        ι2=ι_hat_2,
        σ1=σ1,
        σ2=σ2,
        α_h=α_h,
    )


def risk_neutral_lrr_dynamics(p, α_s):
    """State dynamics under the one-period risk-neutral measure."""
    μ11, μ12, μ22 = p["μ11"], p["μ12"], p["μ22"]
    ι1, ι2 = p["ι1"], p["ι2"]
    σ1, σ2 = p["σ1"], p["σ2"]

    # The one-period risk-neutral measure uses only the SDF exposure.
    μ_bar_11 = μ11
    μ_bar_12 = μ12 + np.dot(σ1, α_s)
    μ_bar_22 = μ22 + np.dot(σ2, α_s)

    # Rewrite the shifted drift in mean-reversion form.
    ι_bar_2 = (μ22 / μ_bar_22) * ι2
    ι_bar_1 = ι1 + (μ12 * ι2 - μ_bar_12 * ι_bar_2) / μ11

    return dict(
        μ11=μ_bar_11,
        μ12=μ_bar_12,
        μ22=μ_bar_22,
        ι1=ι_bar_1,
```

(continues on next page)

(continued from previous page)

```
        ι2=ι_bar_2,
        σ1=σ1,
        σ2=σ2,
    )
```

The quadratic for e_2 has two roots, and `solve_pf_lrr` selects between them using the criterion stated in *Assumption 142.4.2*, which *Proposition 142.4.1* shows leaves at most one solution.

Here the discarded root gives $\hat{\mu}_{22} = +0.0115$ and $\hat{\iota}_2 = -1.13$, an explosive volatility process reverting to a negative long-run mean, so X is certainly not stationary and ergodic under the measure it induces.

For the calibration used here, the recovered probability measure changes the long-run state distribution.

It lowers the mean of expected growth and raises the mean of volatility.

```
v1, v2 = solve_value_function(lrr_params)
e1, e2, η_lrr, α_s = solve_pf_lrr(lrr_params, v1, v2)
dyn_hat = recovered_lrr_dynamics(lrr_params, e1, e2, α_s)
dyn_bar = risk_neutral_lrr_dynamics(lrr_params, α_s)

print(f"value slopes:       v1 = {v1:.4f}, v2 = {v2:.4f}")
print(f"eigenfunction coefficients: e1 = {e1:.4f}, e2 = {e2:.4f}")
print(f"log eigenvalue:     eta = {η_lrr:.6f}  "
      f"(annualized {12 * η_lrr:.4f})")
print()
print("Long-run means under three measures")
print("measure        iota_1     iota_2     mu_12      mu_22")
print("---------   --------   --------   --------   --------")
print(f"actual      {lrr_params['ι1']:8.5f}   {lrr_params['ι2']:8.5f}"
      f"   {lrr_params['μ12']:8.5f}   {lrr_params['μ22']:8.5f}")
print(f"one-period  {dyn_bar['ι1']:8.5f}   {dyn_bar['ι2']:8.5f}"
      f"   {dyn_bar['μ12']:8.5f}   {dyn_bar['μ22']:8.5f}")
print(f"long-term   {dyn_hat['ι1']:8.5f}   {dyn_hat['ι2']:8.5f}"
      f"   {dyn_hat['μ12']:8.5f}   {dyn_hat['μ22']:8.5f}")
```

```
value slopes:       v1 = 43.4783, v2 = -0.0871
eigenfunction coefficients: e1 = -47.6190, e2 = 0.2449
log eigenvalue:     eta = -0.000316  (annualized -0.0038)

Long-run means under three measures
measure        iota_1     iota_2     mu_12      mu_22
---------   --------   --------   --------   --------
actual       0.00000    1.00000    0.00000   -0.01300
one-period  -0.00236    1.09537   -0.00005   -0.01187
long-term   -0.00273    1.12901   -0.00005   -0.01151
```

These numbers show the mechanism clearly.

The positive value slope v_1 says that the continuation value is very sensitive to predictable consumption growth.

The volatility slope v_2 is negative in this calibration, so higher volatility lowers continuation value.

The eigenfunction coefficient e_1 has the opposite sign: the eigenfunction declines in predictable growth.

The change of measure itself loads through $\alpha_S + \sigma_1 e_1 + \sigma_2 e_2$, and the eigenfunction term $\sigma_1 e_1$ reinforces the negative tilt already present in α_S.

Thus the recovered probability measure assigns more probability to histories with lower expected growth.

The positive e_2 works in the other direction for volatility, assigning more probability to higher-volatility states.

The table translates those coefficients into state dynamics.

Relative to the correctly specified probability measure, both risk-neutral measures lower the long-run mean of predictable growth and raise the long-run mean of volatility.

The long-term risk-neutral measure moves further in that direction than the one-period risk-neutral measure: ι_1 falls from 0 to about -0.0027, while ι_2 rises from 1 to about 1.13.

The small negative log eigenvalue means that $\exp(\eta)$ is slightly below one; with the usual yield sign convention, $-\eta$ is the corresponding long-run discount rate.

Stationary densities

The coefficient table gives one summary of the difference between probability measures.

A stationary-density plot gives another.

It shows not only that the means of X_1 and X_2 move, but also which combinations of growth and volatility become more likely.

This matters because treating the recovered probability measure as beliefs changes the whole forecast distribution, not just a pair of long-run averages.

Under the recovered probability measure, probability mass shifts toward adverse long-run-risk states.

These are states with lower predictable growth X_1 and higher volatility X_2.

The dashed contour adds the one-period risk-neutral probability measure.

In this calibration, the one-period risk-neutral and long-term risk-neutral stationary distributions are close to each other, and both are far from the correctly specified distribution.

The whole gap between $\mathbf{P}$ and $\hat{\mathbf{P}}$ is the martingale component $\hat{H}$, whose shock exposure is $\alpha_S + \sigma_1 e_1 + \sigma_2 e_2$.

What the closeness of the two risk-neutral densities shows is how that exposure is split: most of it comes from the one-period risk price α_S, while the Perron–Frobenius eigenfunction supplies the rest, about a tenth of the loading on the growth shock and about a quarter of the loading on the volatility shock.

The paper's Figure 1 reports model-implied stationary densities; the simulation below is a numerical approximation to those densities.

The plot below simulates the state process under each probability measure and estimates the stationary joint density of (X_2, X_1).

The horizontal line marks $X_1 = 0$ and the vertical line marks the correctly specified mean of volatility, $X_2 = \iota_2$.

```
def simulate_lrr(dyn, T=180_000, seed=123):
    """
    Euler simulation of the LRR state process under one probability measure.
    """
    rng = np.random.default_rng(seed)
    X1 = np.zeros(T)
    X2 = np.full(T, dyn["ι2"])

    # Euler step with monthly time increment
    for t in range(1, T):
        X2_prev = max(X2[t-1], 1e-9)
        dW = rng.standard_normal(3)
        sqrt_X2 = np.sqrt(X2_prev)

        X1[t] = (
```

(continues on next page)

(continued from previous page)

```
            X1[t-1]
            + dyn["μ11"] * (X1[t-1] - dyn["ι1"])
            + dyn["μ12"] * (X2_prev - dyn["ι2"])
            + sqrt_X2 * np.dot(dyn["σ1"], dW)
        )
        X2[t] = max(
            X2_prev
            + dyn["μ22"] * (X2_prev - dyn["ι2"])
            + sqrt_X2 * np.dot(dyn["σ2"], dW),
            1e-9,
        )

    burn = T // 5
    return X1[burn:], X2[burn:]


def kde2d_contour(ax, X1, X2, label, levels=7, fill=True,
                  linestyle='solid', outer_only=False):
    """Estimate the stationary density and draw its contours."""
    m = min(25_000, len(X1))
    idx = np.linspace(0, len(X1) - 1, m, dtype=int)
    x1 = X1[idx]
    x2 = X2[idx]

    kde = gaussian_kde(np.vstack([x2, x1]))
    x2_grid = np.linspace(0.6, 1.6, 140)
    x1_grid = np.linspace(-0.006, 0.006, 140)
    X2g, X1g = np.meshgrid(x2_grid, x1_grid)
    Z = kde(np.vstack([X2g.ravel(), X1g.ravel()])).reshape(X2g.shape)

    contour_levels = np.linspace(0.12 * Z.max(), 0.9 * Z.max(), levels)
    if outer_only:
        contour_levels = contour_levels[:1]

    if fill:
        fill_levels = np.r_[contour_levels, Z.max()]
        ax.contourf(X2g, X1g, Z, levels=fill_levels, cmap='Greys',
                    alpha=0.85)
        ax.contour(X2g, X1g, Z, levels=contour_levels, colors='0.55',
                   linewidths=0.4)
        ax.plot([], [], color='0.25', lw=1.5, label=label)
    else:
        ax.contour(X2g, X1g, Z, levels=contour_levels, colors='black',
                   linewidths=1.5, linestyles=linestyle)
        ax.plot([], [], color='black', lw=1.5, ls=linestyle, label=label)


dyn_true = dict(
    μ11=lrr_params["μ11"],
    μ12=lrr_params["μ12"],
    μ22=lrr_params["μ22"],
    ι1=lrr_params["ι1"],
    ι2=lrr_params["ι2"],
    σ1=lrr_params["σ1"],
    σ2=lrr_params["σ2"],
)
```

(continues on next page)

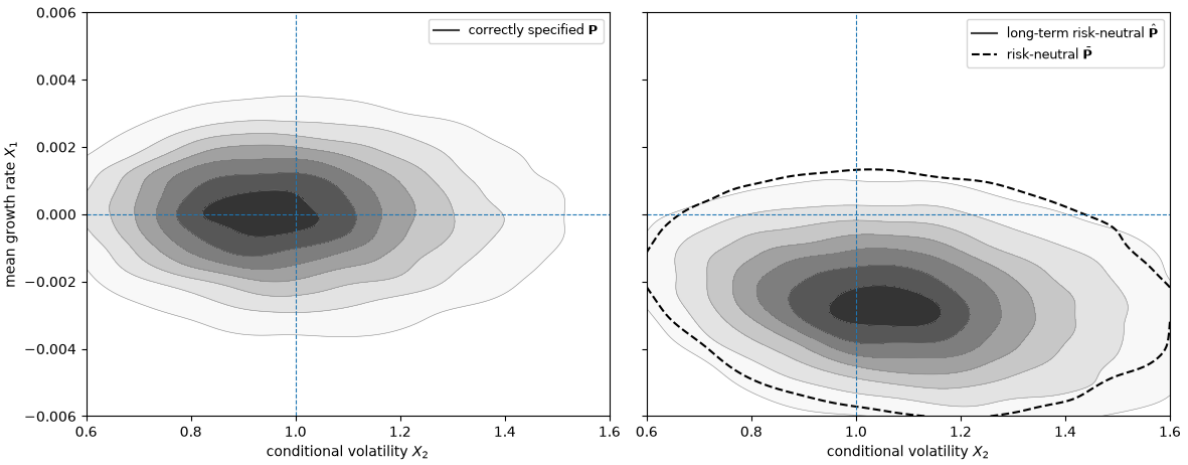
(continued from previous page)

```
X1_P, X2_P = simulate_lrr(dyn_true, seed=1)
X1_H, X2_H = simulate_lrr(dyn_hat, seed=2)
X1_B, X2_B = simulate_lrr(dyn_bar, seed=3)

fig, axes = plt.subplots(1, 2, figsize=(12, 4.8), sharex=True, sharey=True)
kde2d_contour(axes[0], X1_P, X2_P, label=r'correctly specified $\mathbf{P}$')
kde2d_contour(axes[1], X1_H, X2_H,
              label=r'long-term risk-neutral $\hat{\mathbf{P}}$')
kde2d_contour(axes[1], X1_B, X2_B,
              label=r'risk-neutral $\bar{\mathbf{P}}$',
              fill=False, linestyle='--', outer_only=True)

for ax in axes:
    ax.axhline(0, lw=0.8, ls='--')
    ax.axvline(lrr_params["ι2"], lw=0.8, ls='--')
    ax.set_xlim(0.6, 1.6)
    ax.set_ylim(-0.006, 0.006)
    ax.set_xlabel(r"conditional volatility $X_2$")
    ax.legend(fontsize=9)

axes[0].set_ylabel(r"mean growth rate $X_1$")
plt.tight_layout()
plt.show()
```



The movement below the horizontal line means lower expected growth, while movement to the right of the vertical line means higher volatility.

Yield implications

The difference between probability measures matters for asset-pricing interpretation because yields mix two quantities: a payoff forecast and an asset price.

The recovered probability measure is called long-term risk-neutral because it absorbs the martingale component that prices long-horizon risk.

For stochastically growing cash flows, the paper's long-horizon result is that risk premia relative to maturity-matched bonds vanish under the recovered probability measure, subject to the stability and moment conditions used for the limit.

Under the correctly specified probability measure, those same long-term risk premia need not vanish.

For a cash flow G_t, write expectations under the correctly specified probability measure as E_P and expectations under the recovered probability measure as $E_{\hat{P}}$.

The yield computed under the correctly specified probability measure is

$$y_t^P[G](x) = \frac{1}{t}\log E_P[G_t \mid X_0 = x] - \frac{1}{t}\log E_P[S_t G_t \mid X_0 = x].$$

The first term is the payoff forecast.

The second term is the asset price, written using the original SDF representation.

Arrow prices determine the second term.

The question here is what happens to the first term if an analyst treats the recovered probability measure $\hat{\mathbf{P}}$ as investors' beliefs.

In that comparison, prices are held fixed and only the forecast term is recomputed:

$$y_t^{\hat{P}}[G](x) = \frac{1}{t}\log E_{\hat{P}}[G_t \mid X_0 = x] - \frac{1}{t}\log E_P[S_t G_t \mid X_0 = x].$$

For an aggregate-consumption payoff, the answer is substantial.

The recovered probability measure assigns more probability to low-growth, high-volatility states, so it forecasts lower future consumption.

Holding prices fixed, that lower forecast translates into lower consumption yields.

The zero-coupon bond is the comparison case.

Its payoff is one, so the forecast term is always $\log E[1] = 0$.

Changing beliefs therefore does not move the bond-yield panel.

The same solution to the Perron–Frobenius problem also appears in long-bond and forward-measure limits.

The limiting one-period return on a very long bond is

$$R^{\infty}_{t,t+1} = \exp(-\hat{\eta})\frac{\hat{e}(X_{t+1})}{\hat{e}(X_t)}.$$

The martingale increment satisfies

$$\frac{\hat{H}_{t+1}}{\hat{H}_t} = \frac{S_{t+1}}{S_t} R^{\infty}_{t,t+1}.$$

Thus the limiting one-period transition from forward measures coincides with the transition associated with the long-term risk-neutral probability.

The calculation below uses the affine formulas implied by the long-run risk model.

If a multiplicative functional M has log drift affine in X and diffusion proportional to $\sqrt{X_2}$, then

$$E[M_t \mid X_0 = x] = \exp\{\theta_0(t) + \theta_1(t)x_1 + \theta_2(t)x_2\},$$

where the coefficients solve Riccati equations.

The code below computes these affine expectations under the correctly specified measure, recomputes only the consumption forecast under the recovered probability measure, and keeps asset prices fixed.

It then plots median and interquartile yield bands across the same simulated initial states.

```
def affine_expectation_coeffs(dyn, β0, β1, β2, α, horizons):
    """Riccati coefficients for log E[M_t | X_0=x]."""
    μ11, μ12, μ22 = dyn["μ11"], dyn["μ12"], dyn["μ22"]
    ι1, ι2 = dyn["ι1"], dyn["ι2"]
    σ1, σ2 = dyn["σ1"], dyn["σ2"]

    def ode(_, θ):
        θ0, θ1, θ2 = θ
        θ0_dot = (β0 - β1 * ι1 - β2 * ι2
                  - θ1 * (μ11 * ι1 + μ12 * ι2)
                  - θ2 * μ22 * ι2)
        θ1_dot = β1 + μ11 * θ1
        θ2_dot = (β2 + μ12 * θ1 + μ22 * θ2
                  + 0.5 * np.dot(α, α)
                  + θ1 * np.dot(σ1, α)
                  + θ2 * np.dot(σ2, α)
                  + 0.5 * θ1**2 * np.dot(σ1, σ1)
                  + θ1 * θ2 * np.dot(σ1, σ2)
                  + 0.5 * θ2**2 * np.dot(σ2, σ2))
        return [θ0_dot, θ1_dot, θ2_dot]

    sol = solve_ivp(ode, (0, horizons[-1]), np.zeros(3),
                    t_eval=horizons, rtol=1e-8, atol=1e-10)
    if not sol.success:
        raise ValueError("Riccati equation failed to solve")
    return sol.y.T


def log_expectation(θ, X1, X2):
    """Evaluate log E[M_t | X_0=x] on simulated states."""
    return θ[:, 0, None] + θ[:, 1, None] * X1[None, :] + θ[:, 2, None] * X2[None, :]


def yield_quantiles(log_num, log_den, horizons):
    """Quartiles of annualized yields across initial states."""
    yields = 12 * (log_num - log_den) / horizons[:, None]
    return np.quantile(yields, [0.25, 0.5, 0.75], axis=1)


def transform_functional(β0, β1, β2, α, dyn_old, dyn_new, α_h):
    """Rewrite a multiplicative functional after changing probabilities."""
    # Girsanov: under the new measure the drift of log M gains x2 * (alpha . alpha_h),
    # while the shock exposure alpha is unchanged.  The x2 loading therefore moves by
    # alpha . alpha_h, and the constant is re-centered on the new long-run means.
    β_level = β0 - β1 * dyn_old["ι1"] - β2 * dyn_old["ι2"]
    β2_new = β2 + np.dot(α, α_h)
    β0_new = β_level + β1 * dyn_new["ι1"] + β2_new * dyn_new["ι2"]
    return β0_new, β1, β2_new, α


def sdf_coefficients(p, v1, v2):
    """SDF coefficients used in the affine expectation calculation."""
    δ, γ = p["δ"], p["γ"]
    α_c, σ1, σ2 = p["α_c"], p["σ1"], p["σ2"]

    α_h_star = (1 - γ) * (α_c + σ1 * v1 + σ2 * v2)
    α_s = -α_c + α_h_star
```

(continues on next page)

(continued from previous page)

```
    β_s1 = -p["β_c1"]
    β_s2 = -p["β_c2"] - 0.5 * np.dot(α_h_star, α_h_star)
    β_s0 = -δ - p["β_c0"] - 0.5 * p["ι2"] * np.dot(α_h_star, α_h_star)

    return β_s0, β_s1, β_s2, α_s


quarters = np.arange(1, 101)
horizons = 3 * quarters

β_c0, β_c1, β_c2 = (lrr_params["β_c0"],
                    lrr_params["β_c1"],
                    lrr_params["β_c2"])
α_c = lrr_params["α_c"]

β_s0, β_s1, β_s2, α_s = sdf_coefficients(lrr_params, v1, v2)

# Numerators and denominators for yields under the correctly specified measure
θ_C_P = affine_expectation_coeffs(dyn_true, β_c0, β_c1, β_c2, α_c, horizons)
θ_S_P = affine_expectation_coeffs(dyn_true, β_s0, β_s1, β_s2, α_s, horizons)
θ_SC_P = affine_expectation_coeffs(
    dyn_true, β_s0 + β_c0, β_s1 + β_c1, β_s2 + β_c2,
    α_s + α_c, horizons
)

# Numerator for the aggregate-consumption payoff under the recovered probability␣
 ↪measure
β_Ch0, β_Ch1, β_Ch2, α_Ch = transform_functional(
    β_c0, β_c1, β_c2, α_c, dyn_true, dyn_hat, dyn_hat["α_h"]
)
θ_C_H = affine_expectation_coeffs(dyn_hat, β_Ch0, β_Ch1, β_Ch2,
                                  α_Ch, horizons)

log_C_P = log_expectation(θ_C_P, X1_P, X2_P)
log_C_H = log_expectation(θ_C_H, X1_P, X2_P)
log_S_P = log_expectation(θ_S_P, X1_P, X2_P)
log_SC_P = log_expectation(θ_SC_P, X1_P, X2_P)

qC_P = yield_quantiles(log_C_P, log_SC_P, horizons)
qC_H = yield_quantiles(log_C_H, log_SC_P, horizons)
qB_P = yield_quantiles(np.zeros_like(log_S_P), log_S_P, horizons)
# A zero-coupon payoff has the same numerator, log E[1] = 0, under either belief.
qB_H = qB_P.copy()

fig, axes = plt.subplots(1, 2, figsize=(12, 4.8), sharex=True)

def plot_yield_band(ax, x, q, color, label, linestyle='solid',
                    alpha=0.35):
    """Plot quartile band and quartile lines."""
    ax.fill_between(x, q[0], q[2], color=color, alpha=alpha, linewidth=0)
    ax.plot(x, q[1], color=color, lw=2.4, ls=linestyle, label=label)
    ax.plot(x, q[0], color=color, lw=1.3, ls=linestyle)
    ax.plot(x, q[2], color=color, lw=1.3, ls=linestyle)


plot_yield_band(axes[0], quarters, qC_P, color='0.2',
```

(continues on next page)

(continued from previous page)

```
                label='correctly specified measure', alpha=0.45)
plot_yield_band(axes[0], quarters, qC_H, color='0.65',
                label='recovered probability measure', linestyle='--', alpha=0.35)
plot_yield_band(axes[1], quarters, qB_P, color='0.2',
                label='correctly specified measure', alpha=0.45)
plot_yield_band(axes[1], quarters, qB_H, color='0.65',
                label='recovered probability measure', linestyle='--', alpha=0.25)

axes[0].set_xlabel('maturity (quarters)')
axes[0].set_ylabel('consumption yield to maturity')
axes[1].set_xlabel('maturity (quarters)')
axes[1].set_ylabel('bond yield to maturity')

axes[0].legend(fontsize=9)

plt.tight_layout()
plt.show()
```

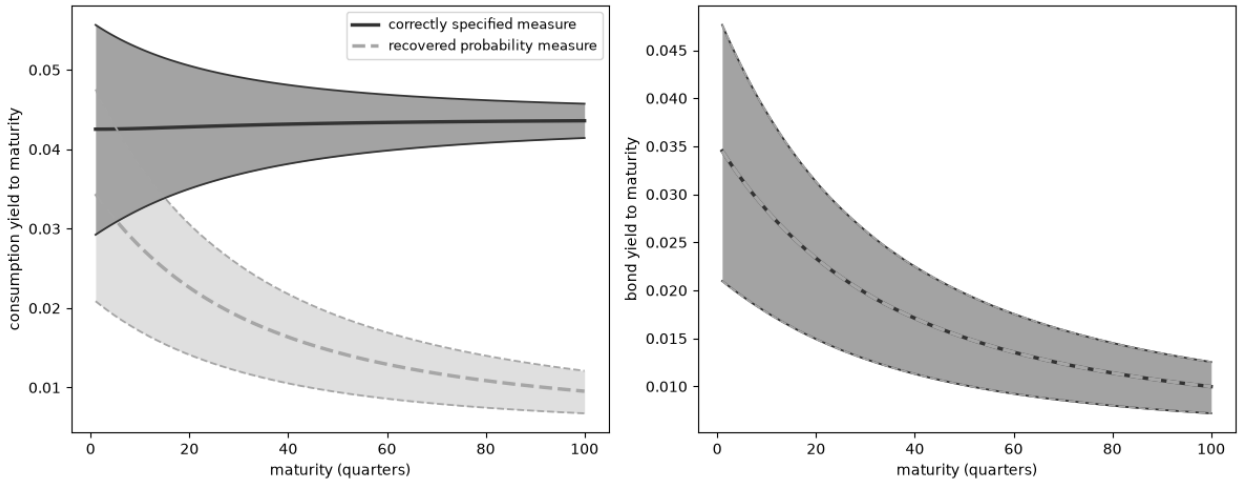


Fig. 142.2: Yield implications of using the recovered probability measure as beliefs. Dashed consumption-yield bands use payoff forecasts under the recovered probability measure with prices fixed; bond yields are unchanged because the zero-coupon payoff has no forecast term.

The left panel is the key one: treating the recovered probability measure as beliefs assigns more probability to low-growth, high-volatility states, so the implied forecast for consumption is lower and consumption yields fall when prices are held fixed.

The bond panel verifies the zero-coupon comparison.

Since $\log E[1] = 0$ under any measure, the solid and dashed bond-yield bands coincide.

Differencing the two panels gives the paper's headline long-horizon result, which the cell above has already computed.

```
prem_P = 1e4 * (qC_P[1] - qB_P[1])
prem_H = 1e4 * (qC_H[1] - qB_P[1])

print("median consumption premium over a maturity-matched bond, basis points")
print("maturity   correctly specified   recovered")
```

(continues on next page)

(continued from previous page)

```
for q in [1, 8, 20, 40, 100]:
    print(f"{q:6d}q {prem_P[q - 1]:18.1f} {prem_H[q - 1]:13.1f}")

print(f"\nlargest absolute premium: {np.max(np.abs(prem_P)):.0f} bp under P, "
      f"{np.max(np.abs(prem_H)):.0f} bp under P-hat")
```

```
median consumption premium over a maturity-matched bond, basis points
maturity   correctly specified   recovered
     1q                 80.2          -1.0
     8q                130.2          -5.3
    20q                194.7          -7.7
    40q                260.6          -7.4
   100q                336.1          -4.5

largest absolute premium: 336 bp under P, 8 bp under P-hat
```

Under the correctly specified probability measure the consumption claim earns 80 to 336 basis points over a maturity-matched bond.

Under the recovered probability measure that premium never exceeds 8 basis points in absolute value, which is the sense in which long-horizon risk premia relative to maturity-matched bonds vanish under the long-term risk-neutral measure.

142.6 Additional state vector

Borovička *et al.* [2016] then ask whether enlarging the state vector changes the recovery problem.

So far, the Perron–Frobenius eigenfunction has depended only on the Markov state X_t.

But many models also contain a growing component Y_t, such as log consumption, with increments driven by the same shock increments.

Here ΔW_{t+1} denotes the shock increment between dates t and $t+1$.

The map ϕ_x sends today's state and the next shock increment into tomorrow's state.

The map ϕ_y sends today's state and the next shock increment into the increment of Y.

$$X_{t+1} = \phi_x(X_t, \Delta W_{t+1}), \qquad Y_{t+1} - Y_t = \phi_y(X_t, \Delta W_{t+1}).$$

Let $\tilde{e}$ denote an eigenfunction candidate that is allowed to depend on both the stationary state X_t and the growing component Y_t.

Let ζ be a vector of loadings on Y, and let $\tilde{e}_\zeta$ be a positive function of X.

Then a natural candidate is

$$\tilde{e}(x, y) = \exp(\zeta \cdot y)\tilde{e}_\zeta(x).$$

This form is natural because Y enters through increments.

Along a path,

$$\exp(\zeta \cdot Y_{t+1}) = \exp(\zeta \cdot Y_t)\exp\{\zeta \cdot (Y_{t+1} - Y_t)\}.$$

Since $Y_{t+1} - Y_t$ is a function of $(X_t, \Delta W_{t+1})$, the ratio $\exp(\zeta \cdot Y_{t+1})/\exp(\zeta \cdot Y_t)$ is a one-period positive multiplicative functional increment.

For a fixed ζ, this factor tilts the one-period pricing operator by $\exp\{\zeta \cdot (Y_{t+1} - Y_t)\}$.

The x-dependent term is therefore not simply the earlier eigenfunction reused.

For each choice of ζ, the remaining x-dependent part solves a different Perron–Frobenius problem:

$$E\left[\frac{S_{t+1}}{S_t}\exp\{\zeta\cdot(Y_{t+1}-Y_t)\}\tilde{e}_\zeta(X_{t+1})\mid X_t=x\right]=\exp(\eta_\zeta)\tilde{e}_\zeta(x).$$

Changing ζ changes how much long-run growth risk is loaded into the eigenfunction.

Thus adding Y_t can make the subjective probability measure one possible solution, but it also creates a family of possible solutions.

The extra state variable therefore does not remove the identification problem; it makes the selection problem more explicit.

The paper also points out a related practical issue.

Highly persistent stationary processes can be hard to distinguish from processes with stationary increments.

A stationary approximation may have a unique solution to the Perron–Frobenius problem for each finite persistence level, but as persistence becomes extreme, the limiting problem can have many approximate solutions.

Numerically, this means the solution to the Perron–Frobenius problem can be sensitive exactly in the cases where a stationary model is being used to approximate stochastic growth.

There is, however, a structured way forward.

If the analyst supplies a reference multiplicative functional Y^r that is known to have the same martingale component as the SDF, then one can restrict the enlarged eigenfunction to the form

$$(Y^r)^{-1}e(x).$$

This restriction chooses which long-run martingale component is allowed into the eigenfunction.

With this extra structure, Arrow prices can again reveal subjective probabilities.

But the key input is external: the long-run martingale component has been supplied by the analyst, not recovered from Arrow prices alone.

142.7 Measuring the martingale component

The paper also asks how large the martingale component is in asset-market data.

Under rational expectations, this measures how important long-term risk adjustments are for valuation.

Under a subjective-beliefs interpretation, it measures the discrepancy between subjective beliefs and the correctly specified probability measure only after imposing that the subjective SDF itself has no martingale component.

With that extra restriction in place, a small martingale component would make the recovered probability measure close to beliefs, while a large one would make long-term risk adjustments more important for the recovered probability measure.

One family of measures applies a convex function to the martingale increment $\hat{H}_{t+1}/\hat{H}_t$.

For example, conditional relative entropy uses

$$E\left[\frac{\hat{H}_{t+1}}{\hat{H}_t}\log\frac{\hat{H}_{t+1}}{\hat{H}_t}\mid X_t=x\right].$$

This expression is zero only when the martingale increment is identically one.

With incomplete asset-market data, the full martingale increment is not observed.

The paper therefore uses pricing restrictions and long-bond return approximations to derive lower bounds on such discrepancy measures.

These bounds are a way to test whether the martingale component is economically small without requiring a full set of Arrow prices.

142.8 Lessons

The Perron–Frobenius approach remains useful under misspecification, but it no longer solves the belief-recovery problem by itself.

It delivers a probability measure that may include long-horizon risk premia.

That measure equals investors' beliefs only when the martingale component is identically one.

Recursive utility, permanent shocks, and long-run risk models give this martingale an economically important role, so transition independence is a substantive economic restriction, not a technical convenience.

142.9 Exercises

Exercise 142.9.1

A two-state martingale component.

Let

$$\mathbf{P} = \begin{pmatrix} 0.8 & 0.2 \\ 0.4 & 0.6 \end{pmatrix}, \qquad \mathbf{Q} = \begin{pmatrix} 0.72 & 0.15 \\ 0.36 & 0.42 \end{pmatrix}.$$

1. Compute the one-period risk-neutral transition matrix $\bar{\mathbf{P}}$.
2. Compute the transition matrix $\hat{\mathbf{P}}$ associated with the recovered probability measure.
3. Compute $\hat{h}_{ij} = \hat{p}_{ij}/p_{ij}$ and decide whether recovery returns the correctly specified transition matrix.

Solution

Here is one solution:

```
P2 = np.array([[0.8, 0.2],
               [0.4, 0.6]])
Q2 = np.array([[0.72, 0.15],
               [0.36, 0.42]])

Pbar2, qb2 = risk_neutral_probs(Q2)
H2, eta2, e2, Phat2 = martingale_increment(Q2, P2)

print("One-period risk-neutral transition matrix P_bar")
print(np.round(Pbar2, 4))
print("\nTransition matrix P_hat associated with the recovered probability measure
 ↪")
print(np.round(Phat2, 4))
print("\nMartingale increment h_hat")
```

```
print(np.round(H2, 4))
print("\nRecovery returns P:", np.allclose(H2[P2 > 0], 1))
```

```
One-period risk-neutral transition matrix P_bar
[[0.8276 0.1724]
 [0.4615 0.5385]]

Transition matrix P_hat associated with the recovered probability measure
[[0.8505 0.1495]
 [0.5039 0.4961]]

Martingale increment h_hat
[[1.0631 0.7476]
 [1.2597 0.8269]]

Recovery returns P: False
```

Exercise 142.9.2

Power utility benchmark.

For trend-stationary consumption and power utility,

$$s_{ij} = A \left(\frac{c_j}{c_i} \right)^{-\gamma}.$$

Show that $\hat{e}_i = c_i^\gamma$ is the Perron–Frobenius eigenvector and that $\hat{\mathbf{P}} = \mathbf{P}$.

Then verify the result numerically using the three-state baseline in the lecture.

Solution

The analytical check is:

$$[\mathbf{Q}\hat{e}]_i = \sum_j A \left(\frac{c_j}{c_i} \right)^{-\gamma} p_{ij} c_j^\gamma = A c_i^\gamma = A \hat{e}_i.$$

Thus $\exp(\hat{\eta}) = A$ and

$$\hat{p}_{ij} = \frac{1}{A} q_{ij} \frac{\hat{e}_j}{\hat{e}_i} = p_{ij}.$$

Below is the numerical check.

```
H_power, _, e_power, P_hat_power = martingale_increment(Q_power, P_true)
e_theory = c_levels**γ_power
e_theory = e_theory / e_theory.sum()

print("Perron-Frobenius eigenvector")
print(np.round(e_power, 6))
print("\nNormalized c^gamma")
print(np.round(e_theory, 6))
print("\nmax |P_hat - P|:",
      np.max(np.abs(P_hat_power - P_true)))
print("max |h_hat - 1|:",
      np.max(np.abs(H_power[P_true > 0] - 1)))
```

```
Perron-Frobenius eigenvector
[0.328344 0.333313 0.338343]

Normalized c^gamma
[0.328344 0.333313 0.338343]

max |P_hat - P|: 3.3306690738754696e-16
max |h_hat - 1|: 1.3322676295501878e-15
```

Exercise 142.9.3

Recursive utility and risk aversion.

Using the finite-state Epstein–Zin example with $c = (0.85, 1.00, 1.15)$, compute the stationary distribution of $\hat{\mathbf{P}}$ for $\gamma \in \{1, 5, 10, 15\}$.

Which state receives the largest increase in stationary probability as γ rises?

Solution

Here is one solution:

```
for γ in [1, 5, 10, 15]:
    _, _, S_g = solve_ez_unit_eis(P_true, c_recursive, δ, γ, g_c)
    Q_g = S_g * P_true
    _, _, _, P_hat_g = martingale_increment(Q_g, P_true)
    π_g = stationary_dist(P_hat_g)
    print(f"gamma={γ:2.0f}: {np.round(π_g, 4)}")

print("\nCorrectly specified:", np.round(π_true, 4))
```

```
gamma= 1: [0.2688 0.4409 0.2903]
gamma= 5: [0.2766 0.4399 0.2835]
gamma=10: [0.2865 0.4385 0.275 ]
gamma=15: [0.2966 0.4368 0.2666]

Correctly specified: [0.2688 0.4409 0.2903]
```

The recession state receives the largest increase.

Part XVII

Data and Empirics

CHAPTER

THREE

PANDAS FOR PANEL DATA

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install --upgrade seaborn
```

We use the following imports.

```
import matplotlib.pyplot as plt
import seaborn as sns
sns.set_theme()
```

143.1 Overview

In an earlier lecture on pandas, we looked at working with simple data sets.

Econometricians often need to work with more complex data sets, such as panels.

Common tasks include

- Importing data, cleaning it and reshaping it across several axes.
- Selecting a time series or cross-section from a panel.
- Grouping and summarizing data.

`pandas` (derived from 'panel' and 'data') contains powerful and easy-to-use tools for solving exactly these kinds of problems.

In what follows, we will use a panel data set of real minimum wages from the OECD to create:

- summary statistics over multiple dimensions of our data
- a time series of the average minimum wage of countries in the dataset
- kernel density estimates of wages by continent

We will begin by reading in our long format panel data from a CSV file and reshaping the resulting `DataFrame` with `pivot_table` to build a `MultiIndex`.

Additional detail will be added to our `DataFrame` using pandas' `merge` function, and data will be summarized with the `groupby` function.

143.2 Slicing and Reshaping Data

We will read in a dataset from the OECD of real minimum wages in 32 countries and assign it to `realwage`.

The dataset can be accessed with the following link:

```
url1 = 'https://raw.githubusercontent.com/QuantEcon/data-lectures/main/lectures/
↪realwage.csv'
```

```
import pandas as pd

# Display 6 columns for viewing purposes
pd.set_option('display.max_columns', 6)

# Reduce decimal points to 2
pd.options.display.float_format = '{:,.2f}'.format

realwage = pd.read_csv(url1)
```

Let's have a look at what we've got to work with

```
realwage.head()  # Show first 5 rows
```

```
   Unnamed: 0        Time  Country                                  Series  \
0           0  2006-01-01  Ireland  In 2015 constant prices at 2015 USD PPPs
1           1  2007-01-01  Ireland  In 2015 constant prices at 2015 USD PPPs
2           2  2008-01-01  Ireland  In 2015 constant prices at 2015 USD PPPs
3           3  2009-01-01  Ireland  In 2015 constant prices at 2015 USD PPPs
4           4  2010-01-01  Ireland  In 2015 constant prices at 2015 USD PPPs

  Pay period     value
0     Annual 17,132.44
1     Annual 18,100.92
2     Annual 17,747.41
3     Annual 18,580.14
4     Annual 18,755.83
```

The data is currently in long format, which is difficult to analyze when there are several dimensions to the data.

We will use `pivot_table` to create a wide format panel, with a `MultiIndex` to handle higher dimensional data.

`pivot_table` arguments should specify the data (values), the index, and the columns we want in our resulting dataframe.

By passing a list in columns, we can create a `MultiIndex` in our column axis

```
realwage = realwage.pivot_table(values='value',
                                index='Time',
                                columns=['Country', 'Series', 'Pay period'])
realwage.head()
```

```
Country                                    Australia         \
Series     In 2015 constant prices at 2015 USD PPPs
Pay period                                    Annual Hourly
Time
2006-01-01                                20,410.65  10.33
2007-01-01                                21,087.57  10.67
```

(continues on next page)

(continued from previous page)

```
2008-01-01                                    20,718.24  10.48
2009-01-01                                    20,984.77  10.62
2010-01-01                                    20,879.33  10.57

Country                                                                ...  \
Series     In 2015 constant prices at 2015 USD exchange rates          ...
Pay period                                                  Annual     ...
Time                                                                   ...
2006-01-01                                                23,826.64    ...
2007-01-01                                                24,616.84    ...
2008-01-01                                                24,185.70    ...
2009-01-01                                                24,496.84    ...
2010-01-01                                                24,373.76    ...

Country                                  United States  \
Series     In 2015 constant prices at 2015 USD PPPs
Pay period                                      Hourly
Time
2006-01-01                                        6.05
2007-01-01                                        6.24
2008-01-01                                        6.78
2009-01-01                                        7.58
2010-01-01                                        7.88

Country
Series     In 2015 constant prices at 2015 USD exchange rates
Pay period                                                  Annual Hourly
Time
2006-01-01                                                12,594.40   6.05
2007-01-01                                                12,974.40   6.24
2008-01-01                                                14,097.56   6.78
2009-01-01                                                15,756.42   7.58
2010-01-01                                                16,391.31   7.88

[5 rows x 128 columns]
```

To more easily filter our time series data, later on, we will convert the index into a `DateTimeIndex`

```
realwage.index = pd.to_datetime(realwage.index)
type(realwage.index)
```

```
pandas.DatetimeIndex
```

The columns contain multiple levels of indexing, known as a `MultiIndex`, with levels being ordered hierarchically (Country > Series > Pay period).

A `MultiIndex` is the simplest and most flexible way to manage panel data in pandas

```
type(realwage.columns)
```

```
pandas.MultiIndex
```

```
realwage.columns.names
```

```
FrozenList(['Country', 'Series', 'Pay period'])
```

Like before, we can select the country (the top level of our `MultiIndex`)

```
realwage['United States'].head()
```

```
Series     In 2015 constant prices at 2015 USD PPPs         \
Pay period                                    Annual Hourly
Time
2006-01-01                                12,594.40   6.05
2007-01-01                                12,974.40   6.24
2008-01-01                                14,097.56   6.78
2009-01-01                                15,756.42   7.58
2010-01-01                                16,391.31   7.88

Series     In 2015 constant prices at 2015 USD exchange rates
Pay period                                              Annual Hourly
Time
2006-01-01                                           12,594.40   6.05
2007-01-01                                           12,974.40   6.24
2008-01-01                                           14,097.56   6.78
2009-01-01                                           15,756.42   7.58
2010-01-01                                           16,391.31   7.88
```

Stacking and unstacking levels of the `MultiIndex` will be used throughout this lecture to reshape our dataframe into a format we need.

`.stack()` rotates the lowest level of the column `MultiIndex` to the row index (`.unstack()` works in the opposite direction - try it out)

```
realwage.stack().head()
```

```
Country                                                Australia  \
Series                 In 2015 constant prices at 2015 USD PPPs
Time       Pay period
2006-01-01 Annual                                      20,410.65
           Hourly                                          10.33
2007-01-01 Annual                                      21,087.57
           Hourly                                          10.67
2008-01-01 Annual                                      20,718.24

Country                                                                    \
Series                 In 2015 constant prices at 2015 USD exchange rates
Time       Pay period
2006-01-01 Annual                                                 23,826.64
           Hourly                                                     12.06
2007-01-01 Annual                                                 24,616.84
           Hourly                                                     12.46
2008-01-01 Annual                                                 24,185.70

Country                                                  Belgium  ...  \
Series                 In 2015 constant prices at 2015 USD PPPs  ...
Time       Pay period                                            ...
2006-01-01 Annual                                      21,042.28  ...
           Hourly                                          10.09  ...
2007-01-01 Annual                                      21,310.05  ...
           Hourly                                          10.22  ...
2008-01-01 Annual                                      21,416.96  ...
```

(continues on next page)

(continued from previous page)

```
Country                                                  United Kingdom  \
Series                In 2015 constant prices at 2015 USD exchange rates
Time       Pay period
2006-01-01 Annual                                              20,376.32
           Hourly                                                   9.81
2007-01-01 Annual                                              20,954.13
           Hourly                                                  10.07
2008-01-01 Annual                                              20,902.87

Country                                          United States  \
Series                In 2015 constant prices at 2015 USD PPPs
Time       Pay period
2006-01-01 Annual                                    12,594.40
           Hourly                                         6.05
2007-01-01 Annual                                    12,974.40
           Hourly                                         6.24
2008-01-01 Annual                                    14,097.56

Country
Series                In 2015 constant prices at 2015 USD exchange rates
Time       Pay period
2006-01-01 Annual                                              12,594.40
           Hourly                                                   6.05
2007-01-01 Annual                                              12,974.40
           Hourly                                                   6.24
2008-01-01 Annual                                              14,097.56

[5 rows x 64 columns]
```

We can also pass in an argument to select the level we would like to stack

```
realwage.stack(level='Country').head()
```

```
Series               In 2015 constant prices at 2015 USD PPPs         \
Pay period                                             Annual Hourly
Time       Country
2006-01-01 Australia                                20,410.65  10.33
           Belgium                                  21,042.28  10.09
           Brazil                                    3,310.51   1.41
           Canada                                   13,649.69   6.56
           Chile                                     5,201.65   2.22

Series               In 2015 constant prices at 2015 USD exchange rates
Pay period                                                     Annual Hourly
Time       Country
2006-01-01 Australia                                          23,826.64  12.06
           Belgium                                            20,228.74   9.70
           Brazil                                              2,032.87   0.87
           Canada                                             14,335.12   6.89
           Chile                                               3,333.76   1.42
```

Using a `DatetimeIndex` makes it easy to select a particular time period.

Selecting one year and stacking the two lower levels of the `MultiIndex` creates a cross-section of our panel data

```
realwage.loc['2015'].stack(level=(1, 2)).transpose().head()
```

```
Time                                        2015-01-01         \
Series     In 2015 constant prices at 2015 USD PPPs
Pay period                                      Annual Hourly
Country
Australia                                   21,715.53  10.99
Belgium                                     21,588.12  10.35
Brazil                                       4,628.63   2.00
Canada                                      16,536.83   7.95
Chile                                        6,633.56   2.80

Time
Series     In 2015 constant prices at 2015 USD exchange rates
Pay period                                                Annual Hourly
Country
Australia                                                25,349.90  12.83
Belgium                                                  20,753.48   9.95
Brazil                                                    2,842.28   1.21
Canada                                                   17,367.24   8.35
Chile                                                     4,251.49   1.81
```

For the rest of lecture, we will work with a dataframe of the hourly real minimum wages across countries and time, measured in 2015 US dollars.

To create our filtered dataframe (`realwage_f`), we can use the `xs` method to select values at lower levels in the multiindex, while keeping the higher levels (countries in this case)

```
realwage_f = realwage.xs(('Hourly', 'In 2015 constant prices at 2015 USD exchange␣
↪rates'),
                         level=('Pay period', 'Series'), axis=1)
realwage_f.head()
```

```
Country     Australia  Belgium  Brazil  ...  Turkey  United Kingdom  \
Time                                    ...
2006-01-01      12.06     9.70    0.87  ...    2.27            9.81
2007-01-01      12.46     9.82    0.92  ...    2.26           10.07
2008-01-01      12.24     9.87    0.96  ...    2.22           10.04
2009-01-01      12.40    10.21    1.03  ...    2.28           10.15
2010-01-01      12.34    10.05    1.08  ...    2.30            9.96

Country     United States
Time
2006-01-01           6.05
2007-01-01           6.24
2008-01-01           6.78
2009-01-01           7.58
2010-01-01           7.88

[5 rows x 32 columns]
```

143.3 Merging Dataframes and Filling NaNs

Similar to relational databases like SQL, pandas has built in methods to merge datasets together.

Using country information from WorldData.info, we'll add the continent of each country to `realwage_f` with the `merge` function.

The dataset can be accessed with the following link:

```
url2 = 'https://raw.githubusercontent.com/QuantEcon/data-lectures/main/lectures/
↪countries.csv'
```

```
worlddata = pd.read_csv(url2, sep=';')
worlddata.head()
```

```
    Country (en) Country (de)        Country (local)  ... Deathrate  \
0    Afghanistan  Afghanistan  Afganistan/Afqanestan  ...     13.70
1          Egypt      Ägypten                  Misr  ...      4.70
2  Åland Islands  Ålandinseln                 Åland  ...      0.00
3        Albania     Albanien              Shqipëria  ...      6.70
4        Algeria     Algerien     Al-Jaza'ir/Algérie  ...      4.30

  Life expectancy                                                Url
0           51.30  https://www.laenderdaten.info/Asien/Afghanista...
1           72.70  https://www.laenderdaten.info/Afrika/Aegypten/...
2            0.00  https://www.laenderdaten.info/Europa/Aland/ind...
3           78.30  https://www.laenderdaten.info/Europa/Albanien/...
4           76.80  https://www.laenderdaten.info/Afrika/Algerien/...

[5 rows x 17 columns]
```

First, we'll select just the country and continent variables from `worlddata` and rename the column to 'Country'

```
worlddata = worlddata[['Country (en)', 'Continent']]
worlddata = worlddata.rename(columns={'Country (en)': 'Country'})
worlddata.head()
```

```
         Country Continent
0    Afghanistan      Asia
1          Egypt    Africa
2  Åland Islands    Europe
3        Albania    Europe
4        Algeria    Africa
```

We want to merge our new dataframe, `worlddata`, with `realwage_f`.

The pandas `merge` function allows dataframes to be joined together by rows.

Our dataframes will be merged using country names, requiring us to use the transpose of `realwage_f` so that rows correspond to country names in both dataframes

```
realwage_f.transpose().head()
```

```
Time       2006-01-01  2007-01-01  2008-01-01  ...  2014-01-01  2015-01-01  \
Country                                        ...
Australia       12.06       12.46       12.24  ...       12.67       12.83
Belgium          9.70        9.82        9.87  ...       10.01        9.95
```

(continues on next page)

(continued from previous page)

```
Brazil            0.87        0.92        0.96  ...        1.21        1.21
Canada            6.89        6.96        7.24  ...        8.22        8.35
Chile             1.42        1.45        1.44  ...        1.76        1.81

Time       2016-01-01
Country
Australia       12.98
Belgium          9.76
Brazil           1.24
Canada           8.48
Chile            1.91

[5 rows x 11 columns]
```

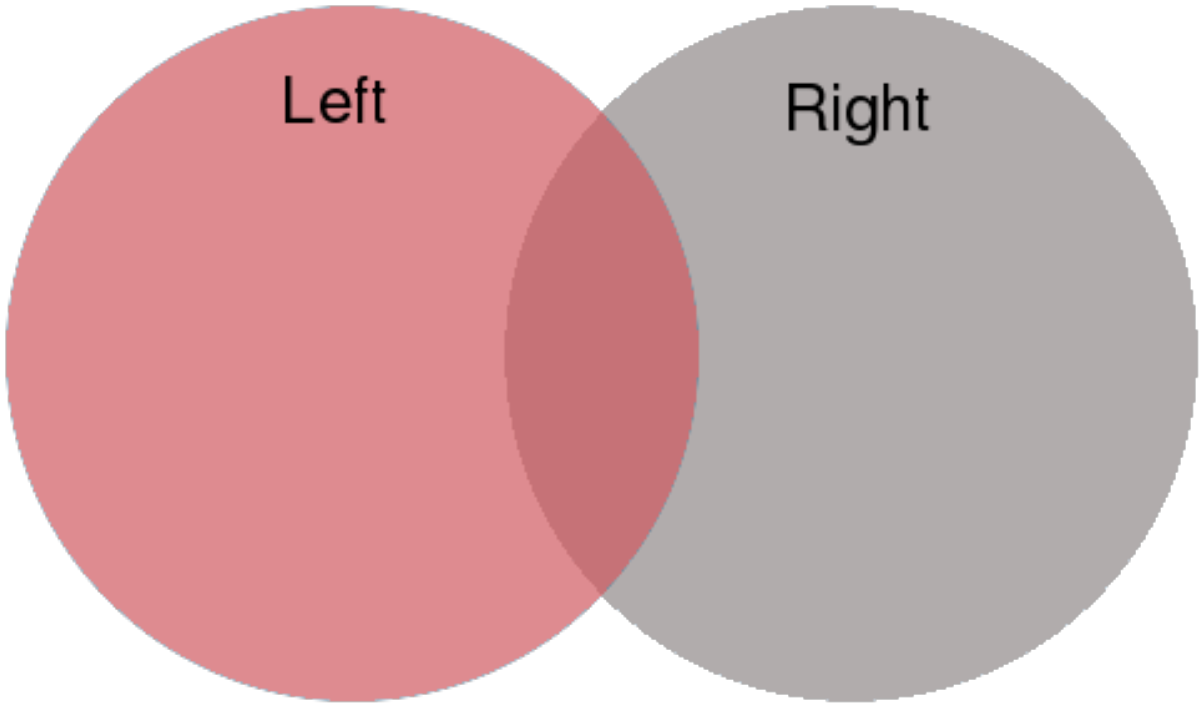
We can use either left, right, inner, or outer join to merge our datasets:

- left join includes only countries from the left dataset
- right join includes only countries from the right dataset
- outer join includes countries that are in either the left and right datasets
- inner join includes only countries common to both the left and right datasets

By default, `merge` will use an inner join.

Here we will pass `how='left'` to keep all countries in `realwage_f`, but discard countries in `worlddata` that do not have a corresponding data entry `realwage_f`.

This is illustrated by the red shading in the following diagram



We will also need to specify where the country name is located in each dataframe, which will be the `key` that is used to merge the dataframes 'on'.

Our 'left' dataframe (`realwage_f.transpose()`) contains countries in the index, so we set `left_index=True`.

Our 'right' dataframe (`worlddata`) contains countries in the 'Country' column, so we set `right_on='Country'`

```
merged = pd.merge(realwage_f.transpose(), worlddata,
                  how='left', left_index=True, right_on='Country')
merged.head()
```

```
        2006-01-01 00:00:00  2007-01-01 00:00:00  2008-01-01 00:00:00  ...  \
17.00                 12.06                12.46                12.24  ...
23.00                  9.70                 9.82                 9.87  ...
32.00                  0.87                 0.92                 0.96  ...
100.00                 6.89                 6.96                 7.24  ...
38.00                  1.42                 1.45                 1.44  ...

        2016-01-01 00:00:00    Country      Continent
17.00                 12.98  Australia      Australia
23.00                  9.76    Belgium         Europe
32.00                  1.24     Brazil  South America
100.00                 8.48     Canada  North America
38.00                  1.91      Chile  South America

[5 rows x 13 columns]
```

Countries that appeared in `realwage_f` but not in `worlddata` will have `NaN` in the Continent column.

To check whether this has occurred, we can use `.isnull()` on the continent column and filter the merged dataframe

```
merged[merged['Continent'].isnull()]
```

```
     2006-01-01 00:00:00  2007-01-01 00:00:00  2008-01-01 00:00:00  ...  \
NaN                 3.42                 3.74                 3.87  ...
NaN                 0.23                 0.45                 0.39  ...
NaN                 1.50                 1.64                 1.71  ...

     2016-01-01 00:00:00             Country  Continent
NaN                 5.28               Korea        NaN
NaN                 0.55  Russian Federation        NaN
NaN                 2.08     Slovak Republic        NaN

[3 rows x 13 columns]
```

We have three missing values!

One option to deal with NaN values is to create a dictionary containing these countries and their respective continents.

`.map()` will match countries in `merged['Country']` with their continent from the dictionary.

Notice how countries not in our dictionary are mapped with `NaN`

```
missing_continents = {'Korea': 'Asia',
                      'Russian Federation': 'Europe',
                      'Slovak Republic': 'Europe'}

merged['Country'].map(missing_continents)
```

```
17.00         NaN
23.00         NaN
32.00         NaN
100.00        NaN
38.00         NaN
108.00        NaN
41.00         NaN
225.00        NaN
53.00         NaN
58.00         NaN
45.00         NaN
68.00         NaN
233.00        NaN
86.00         NaN
88.00         NaN
91.00         NaN
NaN          Asia
117.00        NaN
122.00        NaN
123.00        NaN
138.00        NaN
153.00        NaN
151.00        NaN
174.00        NaN
175.00        NaN
NaN        Europe
NaN        Europe
198.00        NaN
200.00        NaN
227.00        NaN
241.00        NaN
240.00        NaN
Name: Country, dtype: str
```

We don't want to overwrite the entire series with this mapping.

`.fillna()` only fills in `NaN` values in `merged['Continent']` with the mapping, while leaving other values in the column unchanged

```
merged['Continent'] = merged['Continent'].fillna(merged['Country'].map(missing_
↪continents))

# Check for whether continents were correctly mapped

merged[merged['Country'] == 'Korea']
```

```
     2006-01-01 00:00:00  2007-01-01 00:00:00  2008-01-01 00:00:00  ...  \
NaN                 3.42                 3.74                 3.87  ...

     2016-01-01 00:00:00  Country  Continent
NaN                 5.28    Korea       Asia

[1 rows x 13 columns]
```

We will also combine the Americas into a single continent - this will make our visualization nicer later on.

To do this, we will use `.replace()` with a list of the continent values we want to replace.

```
replace = ['Central America', 'North America', 'South America']
merged['Continent'] = merged['Continent'].replace(to_replace=replace, value='America')
```

Now that we have all the data we want in a single `DataFrame`, we will reshape it back into panel form with a `MultiIndex`.

We should also ensure to sort the index using `.sort_index()` so that we can efficiently filter our dataframe later on.

By default, levels will be sorted top-down

```
merged = merged.set_index(['Continent', 'Country']).sort_index()
merged.head()
```

```
                     2006-01-01 00:00:00  2007-01-01 00:00:00  \
Continent Country
America   Brazil                    0.87                 0.92
          Canada                    6.89                 6.96
          Chile                     1.42                 1.45
          Colombia                  1.01                 1.02
          Costa Rica                 NaN                  NaN

                     2008-01-01 00:00:00  ...  2014-01-01 00:00:00  \
Continent Country                         ...
America   Brazil                    0.96  ...                 1.21
          Canada                    7.24  ...                 8.22
          Chile                     1.44  ...                 1.76
          Colombia                  1.01  ...                 1.13
          Costa Rica                 NaN  ...                 2.41

                     2015-01-01 00:00:00  2016-01-01 00:00:00
Continent Country
America   Brazil                    1.21                 1.24
          Canada                    8.35                 8.48
          Chile                     1.81                 1.91
          Colombia                  1.13                 1.12
          Costa Rica                2.56                 2.63

[5 rows x 11 columns]
```

While merging, we lost our `DatetimeIndex`, as we merged columns that were not in datetime format

```
merged.columns
```

```
Index([2006-01-01 00:00:00, 2007-01-01 00:00:00, 2008-01-01 00:00:00,
       2009-01-01 00:00:00, 2010-01-01 00:00:00, 2011-01-01 00:00:00,
       2012-01-01 00:00:00, 2013-01-01 00:00:00, 2014-01-01 00:00:00,
       2015-01-01 00:00:00, 2016-01-01 00:00:00],
      dtype='object')
```

Now that we have set the merged columns as the index, we can recreate a `DatetimeIndex` using `.to_datetime()`

```
merged.columns = pd.to_datetime(merged.columns)
merged.columns = merged.columns.rename('Time')
merged.columns
```

```
DatetimeIndex(['2006-01-01', '2007-01-01', '2008-01-01', '2009-01-01',
               '2010-01-01', '2011-01-01', '2012-01-01', '2013-01-01',
```

(continues on next page)

(continued from previous page)

```
               '2014-01-01', '2015-01-01', '2016-01-01'],
              dtype='datetime64[us]', name='Time', freq=None)
```

The `DatetimeIndex` tends to work more smoothly in the row axis, so we will go ahead and transpose `merged`

```
merged = merged.transpose()
merged.head()
```

```
Continent  America                ...   Europe
Country     Brazil Canada Chile   ... Slovenia Spain United Kingdom
Time                              ...
2006-01-01    0.87   6.89  1.42   ...     3.92  3.99           9.81
2007-01-01    0.92   6.96  1.45   ...     3.88  4.10          10.07
2008-01-01    0.96   7.24  1.44   ...     3.96  4.14          10.04
2009-01-01    1.03   7.67  1.52   ...     4.08  4.32          10.15
2010-01-01    1.08   7.94  1.56   ...     4.81  4.30           9.96

[5 rows x 32 columns]
```

143.4 Grouping and Summarizing Data

Grouping and summarizing data can be particularly useful for understanding large panel datasets.

A simple way to summarize data is to call an aggregation method on the dataframe, such as `.mean()` or `.max()`.

For example, we can calculate the average real minimum wage for each country over the period 2006 to 2016 (the default is to aggregate over rows)

```
merged.mean().head(10)
```

```
Continent  Country
America    Brazil          1.09
           Canada          7.82
           Chile           1.62
           Colombia        1.07
           Costa Rica      2.53
           Mexico          0.53
           United States   7.15
Asia       Israel          5.95
           Japan           6.18
           Korea           4.22
dtype: float64
```

Using this series, we can plot the average real minimum wage over the past decade for each country in our data set

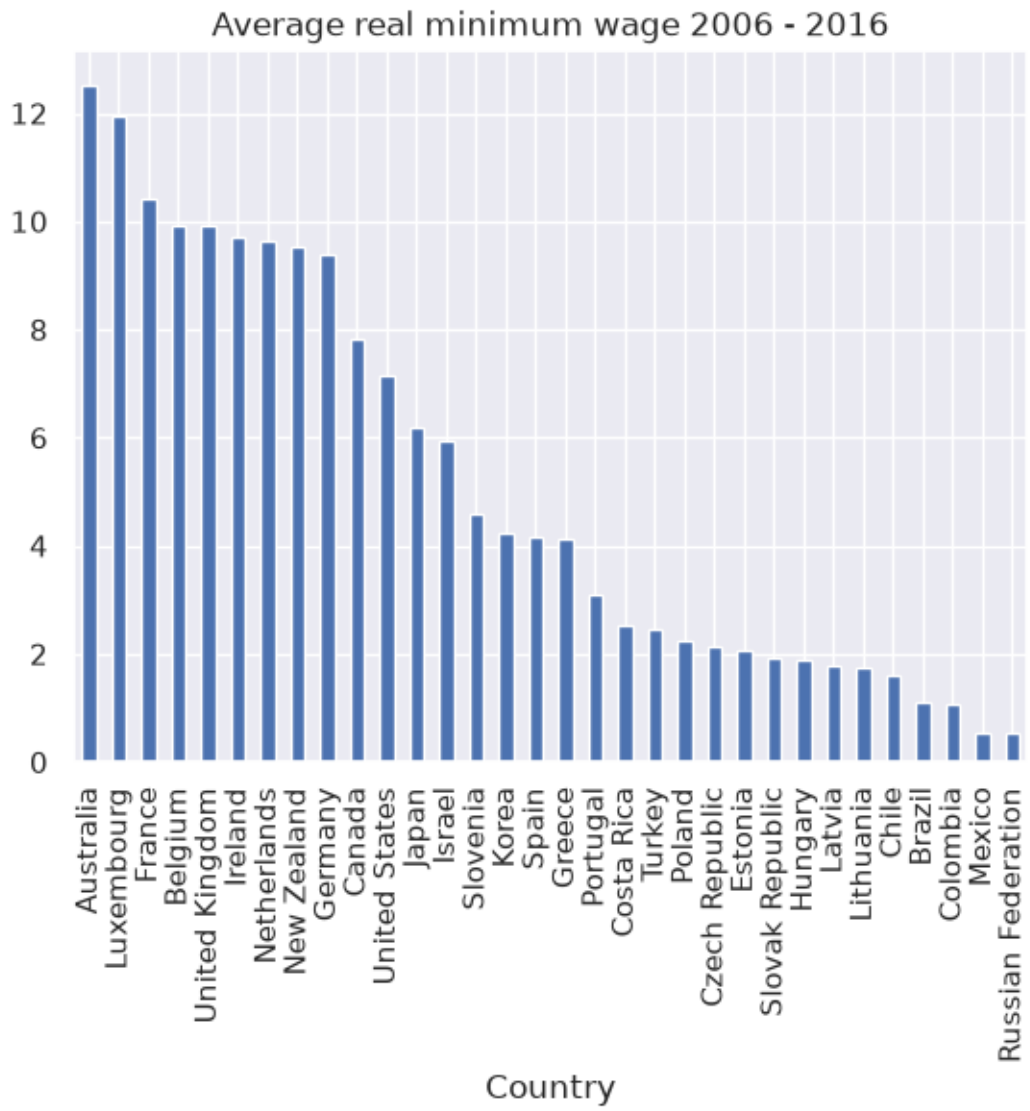
```
merged.mean().sort_values(ascending=False).plot(kind='bar',
                                                title="Average real minimum wage 2006␣
 ↪- 2016")

# Set country labels
country_labels = merged.mean().sort_values(ascending=False).index.get_level_values(
 ↪'Country').tolist()
plt.xticks(range(0, len(country_labels)), country_labels)
plt.xlabel('Country')
```

(continues on next page)

(continued from previous page)

```
plt.show()
```



Passing in `axis=1` to `.mean()` will aggregate over columns (giving the average minimum wage for all countries over time)

```
merged.mean(axis=1).head()
```

```
Time
2006-01-01   4.69
2007-01-01   4.84
2008-01-01   4.90
2009-01-01   5.08
2010-01-01   5.11
dtype: float64
```

We can plot this time series as a line graph

```
merged.mean(axis=1).plot()
plt.title('Average real minimum wage 2006 - 2016')
plt.ylabel('2015 USD')
plt.xlabel('Year')
plt.show()
```



We can also specify a level of the `MultiIndex` (in the column axis) to aggregate over.

In the case of `groupby`, we need to use `.T` to transpose the columns into rows, as `pandas` has removed support for `axis=1` in the `groupby` method.

```
merged.T.groupby(level='Continent').mean().head()
```

```
Time       2006-01-01  2007-01-01  2008-01-01  ...  2014-01-01  2015-01-01  \
Continent                                      ...
America          2.80        2.85        2.99  ...        3.22        3.26
Asia             4.29        4.44        4.45  ...        4.86        5.10
Australia       10.25       10.73       10.76  ...       11.25       11.52
Europe           4.80        4.94        4.99  ...        5.17        5.48

Time       2016-01-01
Continent
America          3.30
Asia             5.44
Australia       11.73
Europe           5.57
```

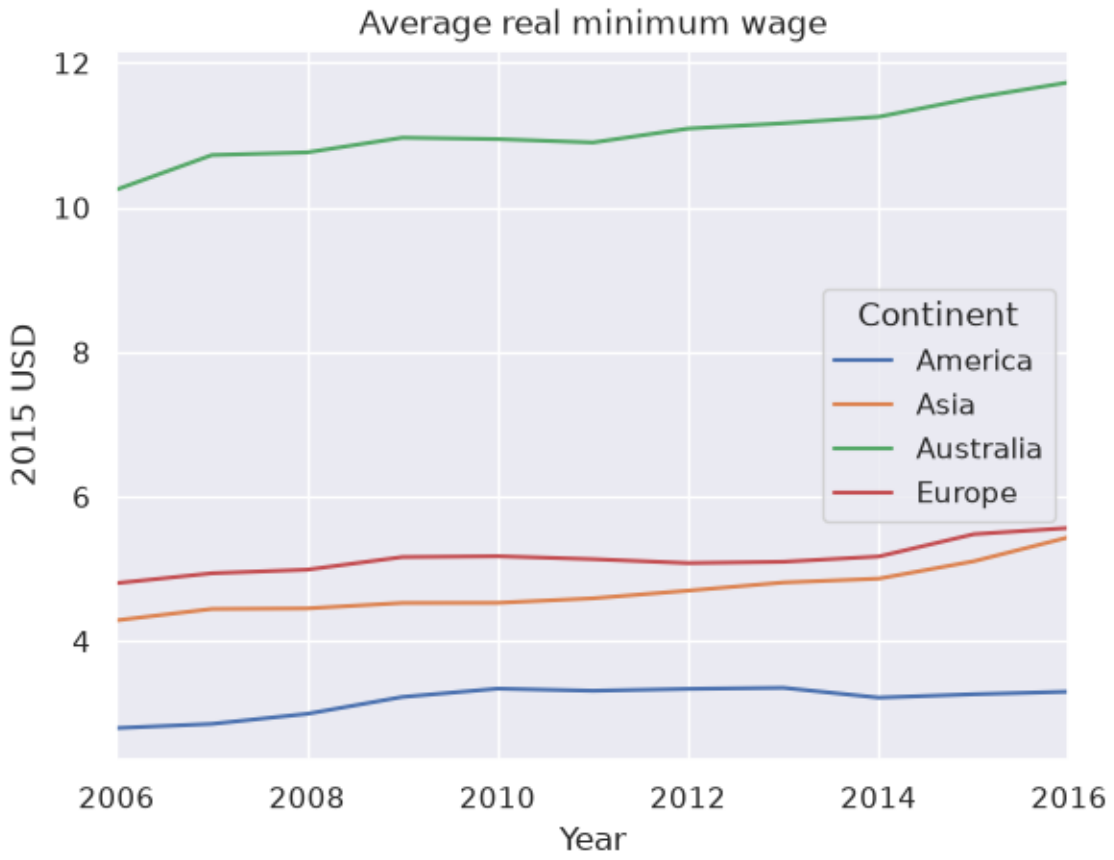
(continues on next page)

(continued from previous page)

```
[4 rows x 11 columns]
```

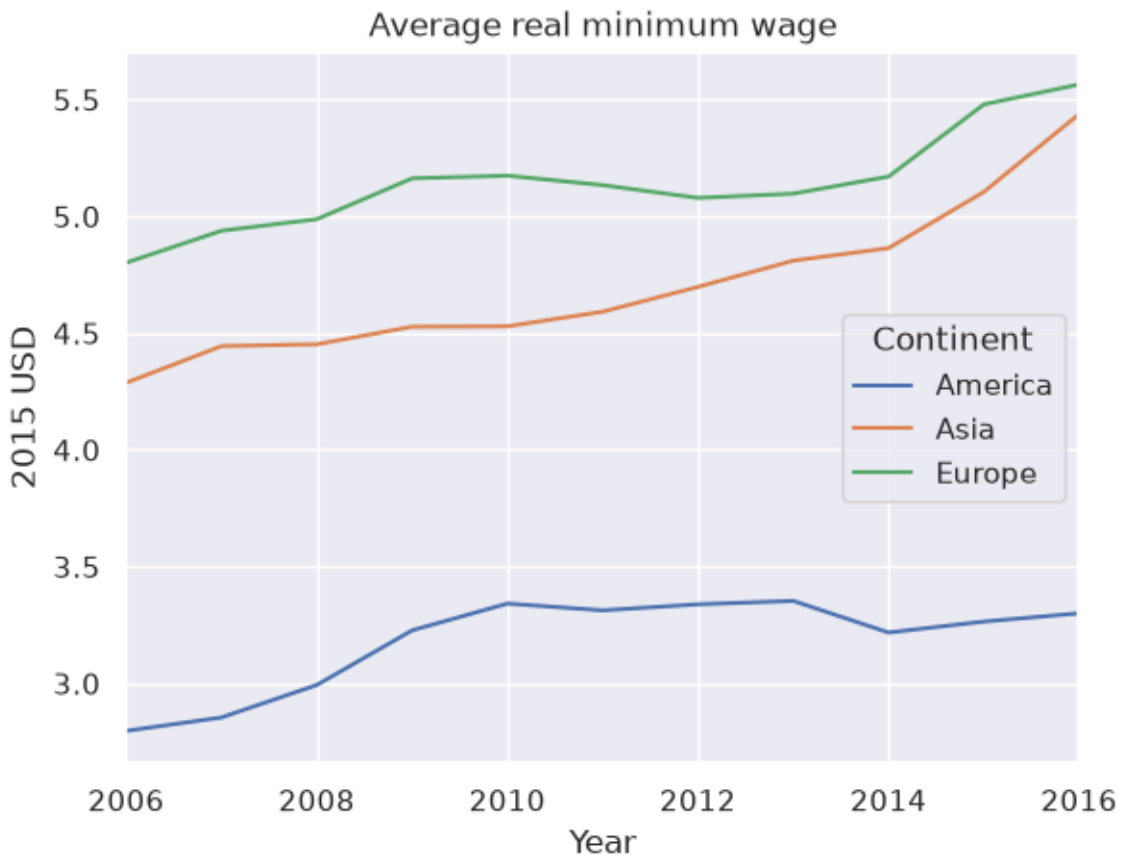
We can plot the average minimum wages in each continent as a time series

```
merged.T.groupby(level='Continent').mean().T.plot()
plt.title('Average real minimum wage')
plt.ylabel('2015 USD')
plt.xlabel('Year')
plt.show()
```



We will drop Australia as a continent for plotting purposes

```
merged = merged.drop('Australia', level='Continent', axis=1)
merged.T.groupby(level='Continent').mean().T.plot()
plt.title('Average real minimum wage')
plt.ylabel('2015 USD')
plt.xlabel('Year')
plt.show()
```



`.describe()` is useful for quickly retrieving a number of common summary statistics

```
merged.stack().describe()
```

```
Continent  America  Asia  Europe
count        69.00 44.00  200.00
mean          3.19  4.70    5.15
std           3.02  1.56    3.82
min           0.52  2.22    0.23
25%           1.03  3.37    2.02
50%           1.44  5.48    3.54
75%           6.96  5.95    9.70
max           8.48  6.65   12.39
```

This is a simplified way to use `groupby`.

Using `groupby` generally follows a 'split-apply-combine' process:

- split: data is grouped based on one or more keys
- apply: a function is called on each group independently
- combine: the results of the function calls are combined into a new data structure

The `groupby` method achieves the first step of this process, creating a new `DataFrameGroupBy` object with data split into groups.

Let's split `merged` by continent again, this time using the `groupby` function, and name the resulting object `grouped`

```
grouped = merged.T.groupby(level='Continent')
grouped
```

```
<pandas.api.typing.DataFrameGroupBy object at 0x7518fa101010>
```

Calling an aggregation method on the object applies the function to each group, the results of which are combined in a new data structure.

For example, we can return the number of countries in our dataset for each continent using `.size()`.

In this case, our new data structure is a `Series`

```
grouped.size()
```

```
Continent
America     7
Asia        4
Europe     19
dtype: int64
```

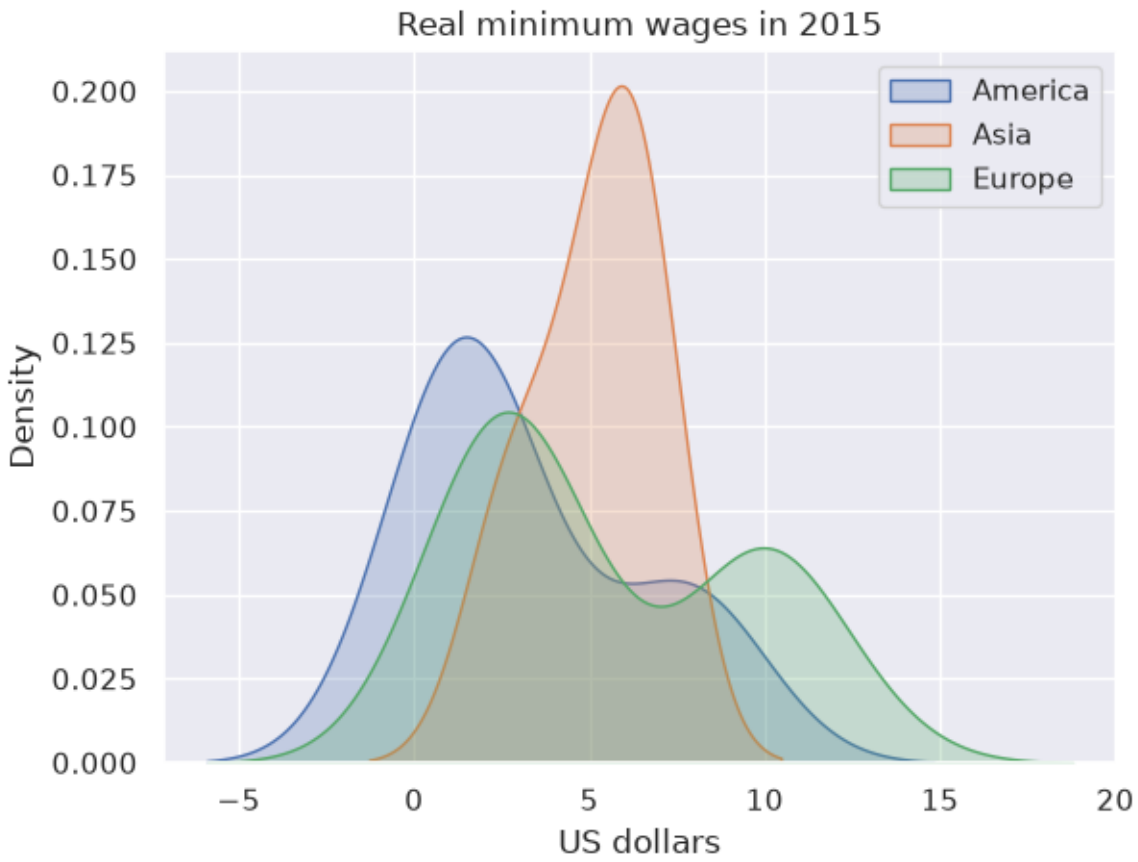
Calling `.get_group()` to return just the countries in a single group, we can create a kernel density estimate of the distribution of real minimum wages in 2016 for each continent.

`grouped.groups.keys()` will return the keys from the `groupby` object

```
continents = grouped.groups.keys()

for continent in continents:
    sns.kdeplot(grouped.get_group(continent).T.loc['2015'].unstack(), label=continent,
 ↪ fill=True)

plt.title('Real minimum wages in 2015')
plt.xlabel('US dollars')
plt.legend()
plt.show()
```



143.5 Final Remarks

This lecture has provided an introduction to some of pandas' more advanced features, including multiindices, merging, grouping and plotting.

Other tools that may be useful in panel data analysis include xarray, a python package that extends pandas to N-dimensional data structures.

143.6 Exercises

Exercise 143.6.1

In these exercises, you'll work with a dataset of employment rates in Europe by age and sex from Eurostat.

The dataset can be accessed with the following link:

```
url3 = 'https://raw.githubusercontent.com/QuantEcon/data-lectures/main/lectures/
↪employ.csv'
```

Reading in the CSV file returns a panel dataset in long format. Use `.pivot_table()` to construct a wide format dataframe with a `MultiIndex` in the columns.

Start off by exploring the dataframe and the variables available in the `MultiIndex` levels.

Write a program that quickly returns all values in the `MultiIndex`.

Solution

```
employ = pd.read_csv(url3)
employ = employ.pivot_table(values='Value',
                            index=['DATE'],
                            columns=['UNIT','AGE', 'SEX', 'INDIC_EM', 'GEO'])
employ.index = pd.to_datetime(employ.index) # ensure that dates are datetime format
employ.head()
```

```
UNIT       Percentage of total population                   ...  \
AGE                    From 15 to 24 years                  ...
SEX                                Females                  ...
INDIC_EM                 Active population                  ...
GEO                              Austria Belgium Bulgaria   ...
DATE                                                        ...
2007-01-01                         56.00   31.60    26.00   ...
2008-01-01                         56.20   30.80    26.10   ...
2009-01-01                         56.20   29.90    24.80   ...
2010-01-01                         54.00   29.80    26.60   ...
2011-01-01                         54.80   29.80    24.80   ...

UNIT                                          Thousand persons           \
AGE                                        From 55 to 64 years
SEX                                                      Total
INDIC_EM   Total employment (resident population concept - LFS)
GEO                                                Switzerland   Turkey
DATE
2007-01-01                                                 NaN   1,282.00
2008-01-01                                                 NaN   1,354.00
2009-01-01                                                 NaN   1,449.00
2010-01-01                                              640.00   1,583.00
2011-01-01                                              661.00   1,760.00

UNIT
AGE
SEX
INDIC_EM
GEO        United Kingdom
DATE
2007-01-01       4,131.00
2008-01-01       4,204.00
2009-01-01       4,193.00
2010-01-01       4,186.00
2011-01-01       4,164.00

[5 rows x 1440 columns]
```

This is a large dataset so it is useful to explore the levels and variables available

```
employ.columns.names
```

```
FrozenList(['UNIT', 'AGE', 'SEX', 'INDIC_EM', 'GEO'])
```

Variables within levels can be quickly retrieved with a loop

```
for name in employ.columns.names:
    print(name, employ.columns.get_level_values(name).unique())
```

```
UNIT Index(['Percentage of total population', 'Thousand persons'], dtype='str',␣
 ↪name='UNIT')
AGE Index(['From 15 to 24 years', 'From 25 to 54 years', 'From 55 to 64 years'],
 ↪ dtype='str', name='AGE')
SEX Index(['Females', 'Males', 'Total'], dtype='str', name='SEX')
INDIC_EM Index(['Active population', 'Total employment (resident population␣
 ↪concept - LFS)'], dtype='str', name='INDIC_EM')
GEO Index(['Austria', 'Belgium', 'Bulgaria', 'Croatia', 'Cyprus', 'Czech␣
 ↪Republic',
       'Denmark', 'Estonia', 'Euro area (17 countries)',
       'Euro area (18 countries)', 'Euro area (19 countries)',
       'European Union (15 countries)', 'European Union (27 countries)',
       'European Union (28 countries)', 'Finland',
       'Former Yugoslav Republic of Macedonia, the', 'France',
       'France (metropolitan)',
       'Germany (until 1990 former territory of the FRG)', 'Greece', 'Hungary',
       'Iceland', 'Ireland', 'Italy', 'Latvia', 'Lithuania', 'Luxembourg',
       'Malta', 'Netherlands', 'Norway', 'Poland', 'Portugal', 'Romania',
       'Slovakia', 'Slovenia', 'Spain', 'Sweden', 'Switzerland', 'Turkey',
       'United Kingdom'],
      dtype='str', name='GEO')
```

Exercise 143.6.2

Filter the above dataframe to only include employment as a percentage of 'active population'.

Create a grouped boxplot using `seaborn` of employment rates in 2015 by age group and sex.

Hint

`GEO` includes both areas and countries.

Solution

To easily filter by country, swap `GEO` to the top level and sort the `MultiIndex`

```
employ.columns = employ.columns.swaplevel(0,-1)
employ = employ.sort_index(axis=1)
```

We need to get rid of a few items in `GEO` which are not countries.

A fast way to get rid of the EU areas is to use a list comprehension to find the level values in `GEO` that begin with 'Euro'

```
geo_list = employ.columns.get_level_values('GEO').unique().tolist()
countries = [x for x in geo_list if not x.startswith('Euro')]
employ = employ[countries]
employ.columns.get_level_values('GEO').unique()
```

```
Index(['Austria', 'Belgium', 'Bulgaria', 'Croatia', 'Cyprus', 'Czech Republic',
       'Denmark', 'Estonia', 'Finland',
       'Former Yugoslav Republic of Macedonia, the', 'France',
       'France (metropolitan)',
       'Germany (until 1990 former territory of the FRG)', 'Greece', 'Hungary',
       'Iceland', 'Ireland', 'Italy', 'Latvia', 'Lithuania', 'Luxembourg',
       'Malta', 'Netherlands', 'Norway', 'Poland', 'Portugal', 'Romania',
       'Slovakia', 'Slovenia', 'Spain', 'Sweden', 'Switzerland', 'Turkey',
       'United Kingdom'],
      dtype='str', name='GEO')
```

Select only percentage employed in the active population from the dataframe

```
employ_f = employ.xs(('Percentage of total population', 'Active population'),
                     level=('UNIT', 'INDIC_EM'),
                     axis=1)
employ_f.head()
```

```
GEO                    Austria                ...      United Kingdom        \
AGE         From 15 to 24 years                ... From 55 to 64 years
SEX                     Females Males Total   ...              Females Males
DATE                                          ...
2007-01-01                56.00 62.90 59.40   ...                49.90 68.90
2008-01-01                56.20 62.90 59.50   ...                50.20 69.80
2009-01-01                56.20 62.90 59.50   ...                50.60 70.30
2010-01-01                54.00 62.60 58.30   ...                51.10 69.20
2011-01-01                54.80 63.60 59.20   ...                51.30 68.40

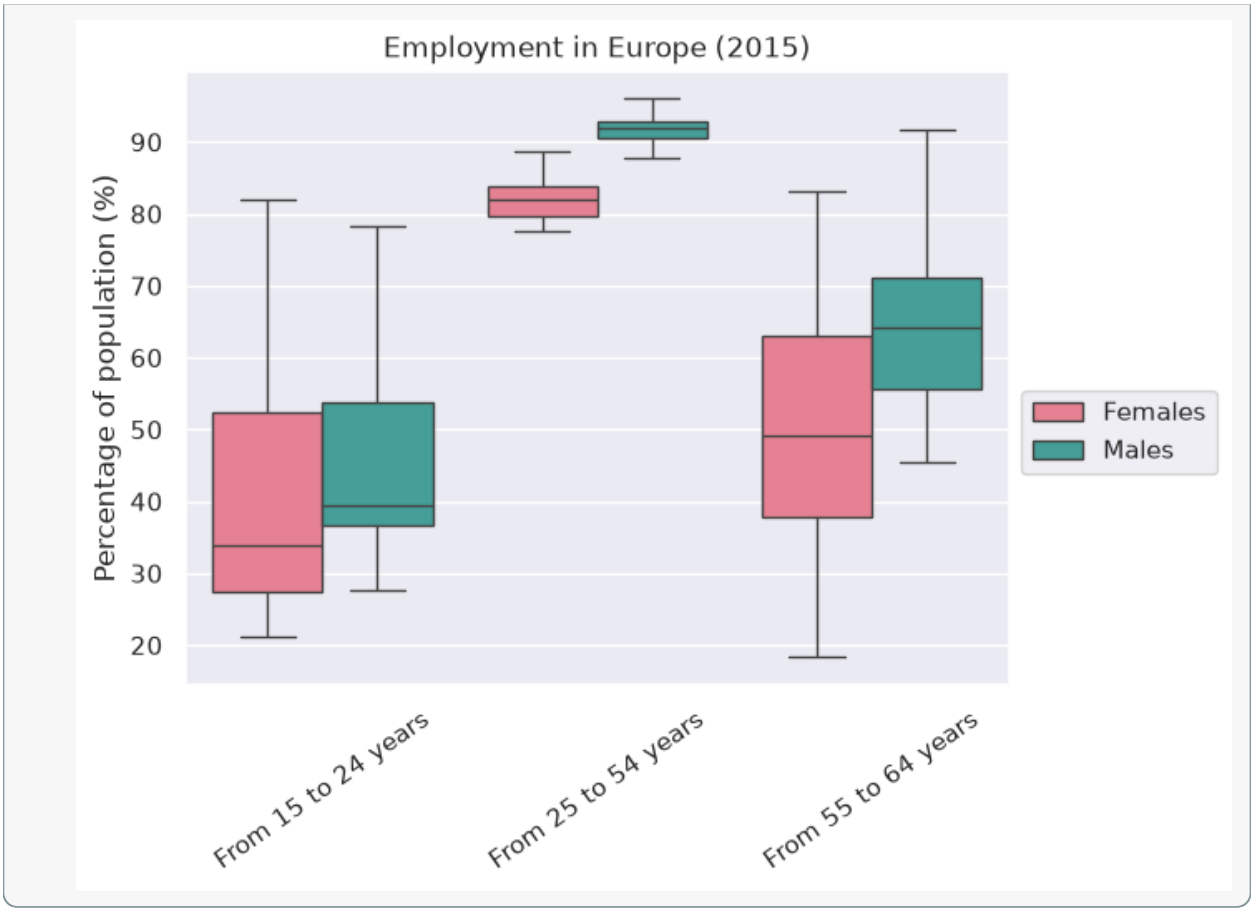
GEO
AGE
SEX         Total
DATE
2007-01-01 59.30
2008-01-01 59.80
2009-01-01 60.30
2010-01-01 60.00
2011-01-01 59.70

[5 rows x 306 columns]
```

Drop the 'Total' value before creating the grouped boxplot

```
employ_f = employ_f.drop('Total', level='SEX', axis=1)
```

```
box = employ_f.loc['2015'].unstack().reset_index()
sns.boxplot(x="AGE", y=0, hue="SEX", data=box, palette=("husl"), showfliers=False)
plt.xlabel('')
plt.xticks(rotation=35)
plt.ylabel('Percentage of population (%)')
plt.title('Employment in Europe (2015)')
plt.legend(bbox_to_anchor=(1,0.5))
plt.show()
```


Employment in Europe (2015)
Percentage of population (%)
90
80
70
60
50
40
30
20
From 15 to 24 years
From 25 to 54 years
From 55 to 64 years
Females
Males

CHAPTER

FOUR

LINEAR REGRESSION IN PYTHON

Contents

- *Linear Regression in Python*
  - *Overview*
  - *Simple Linear Regression*
  - *Extending the Linear Regression Model*
  - *Endogeneity*
  - *Summary*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install linearmodels
```

144.1 Overview

Linear regression is a standard tool for analyzing the relationship between two or more variables.

In this lecture, we'll use the Python package `statsmodels` to estimate, interpret, and visualize linear regression models.

Along the way, we'll discuss a variety of topics, including

- simple and multivariate linear regression
- visualization
- endogeneity and omitted variable bias
- two-stage least squares

As an example, we will replicate results from Acemoglu, Johnson and Robinson's seminal paper [Acemoglu *et al.*, 2001].

- You can download a copy here.

In the paper, the authors emphasize the importance of institutions in economic development.

The main contribution is the use of settler mortality rates as a source of *exogenous* variation in institutional differences.

Such variation is needed to determine whether it is institutions that give rise to greater economic growth, rather than the other way around.

Let's start with some imports:

```
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
import statsmodels.api as sm
from statsmodels.iolib.summary2 import summary_col
from linearmodels.iv import IV2SLS
import seaborn as sns
sns.set_theme()
```

144.1.1 Prerequisites

This lecture assumes you are familiar with basic econometrics.

For an introductory text covering these topics, see, for example, [Wooldridge, 2015].

144.2 Simple Linear Regression

[Acemoglu *et al.*, 2001] wish to determine whether or not differences in institutions can help to explain observed economic outcomes.

How do we measure *institutional differences* and *economic outcomes*?

In this paper,

- economic outcomes are proxied by log GDP per capita in 1995, adjusted for exchange rates.
- institutional differences are proxied by an index of protection against expropriation on average over 1985-95, constructed by the Political Risk Services Group.

These variables and other data used in the paper are available for download on Daron Acemoglu's webpage.

We will use pandas' `.read_stata()` function to read in data contained in the `.dta` files to dataframes

```
df1 = pd.read_stata('https://github.com/QuantEcon/data-lectures/raw/main/lectures/
↪maketable1.dta')
df1.head()
```

```
  shortnam    euro1900  excolony    avexpr  logpgp95  cons1  cons90  democ00a  \
0      AFG    0.000000       1.0       NaN       NaN    1.0     2.0       1.0
1      AGO    8.000000       1.0  5.363636  7.770645    3.0     3.0       0.0
2      ARE    0.000000       1.0  7.181818  9.804219    NaN     NaN       NaN
3      ARG   60.000004       1.0  6.386364  9.133459    1.0     6.0       3.0
4      ARM    0.000000       0.0       NaN  7.682482    NaN     NaN       NaN

   cons00a    extmort4    logem4  loghjypl  baseco
0      1.0   93.699997  4.540098       NaN     NaN
1      1.0  280.000000  5.634789 -3.411248     1.0
2      NaN         NaN       NaN       NaN     NaN
3      3.0   68.900002  4.232656 -0.872274     1.0
4      NaN         NaN       NaN       NaN     NaN
```

Let's use a scatterplot to see whether any obvious relationship exists between GDP per capita and the protection against expropriation index

```
df1.plot(x='avexpr', y='logpgp95', kind='scatter')
plt.show()
```



The plot shows a fairly strong positive relationship between protection against expropriation and log GDP per capita.

Specifically, if higher protection against expropriation is a measure of institutional quality, then better institutions appear to be positively correlated with better economic outcomes (higher GDP per capita).

Given the plot, choosing a linear model to describe this relationship seems like a reasonable assumption.

We can write our model as

$$logpgp95_i = \beta_0 + \beta_1 avexpr_i + u_i$$

where:

- β_0 is the intercept of the linear trend line on the y-axis
- β_1 is the slope of the linear trend line, representing the *marginal effect* of protection against risk on log GDP per capita
- u_i is a random error term (deviations of observations from the linear trend due to factors not included in the model)

Visually, this linear model involves choosing a straight line that best fits the data, as in the following plot (Figure 2 in [Acemoglu *et al.*, 2001])

```
# Dropping NA's is required to use numpy's polyfit
df1_subset = df1.dropna(subset=['logpgp95', 'avexpr'])

# Use only 'base sample' for plotting purposes
df1_subset = df1_subset[df1_subset['baseco'] == 1]

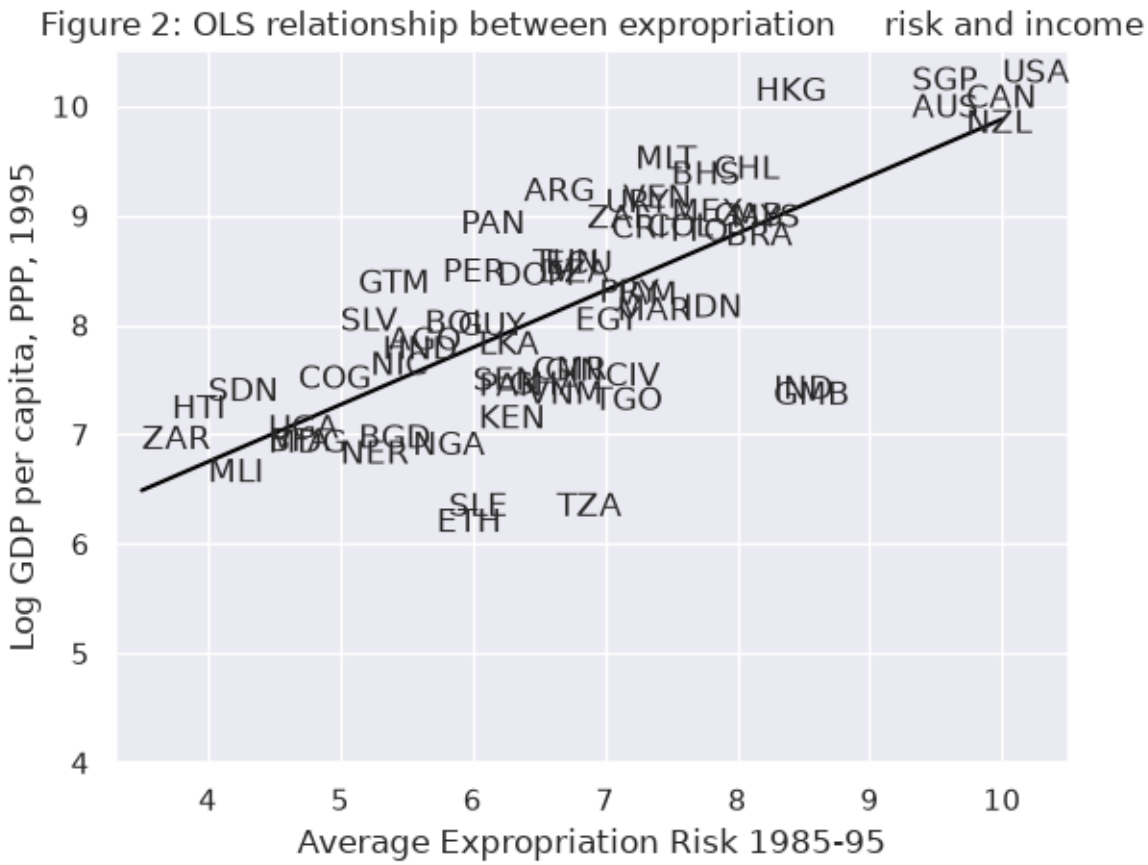
X = df1_subset['avexpr']
y = df1_subset['logpgp95']
labels = df1_subset['shortnam']

# Replace markers with country labels
fig, ax = plt.subplots()
ax.scatter(X, y, marker='')

for i, label in enumerate(labels):
    ax.annotate(label, (X.iloc[i], y.iloc[i]))

# Fit a linear trend line
ax.plot(np.unique(X),
         np.poly1d(np.polyfit(X, y, 1))(np.unique(X)),
         color='black')

ax.set_xlim([3.3,10.5])
ax.set_ylim([4,10.5])
ax.set_xlabel('Average Expropriation Risk 1985-95')
ax.set_ylabel('Log GDP per capita, PPP, 1995')
ax.set_title('Figure 2: OLS relationship between expropriation \
    risk and income')
plt.show()
```



The most common technique to estimate the parameters (β's) of the linear model is Ordinary Least Squares (OLS).

As the name implies, an OLS model is solved by finding the parameters that minimize **the sum of squared residuals**, i.e.

$$\min_{\hat{\beta}} \sum_{i=1}^{N} \hat{u}_i^2$$

where $\hat{u}_i$ is the difference between the observation and the predicted value of the dependent variable.

To estimate the constant term β_0, we need to add a column of 1's to our dataset (consider the equation if β_0 was replaced with $\beta_0 x_i$ and $x_i = 1$)

```
df1['const'] = 1
```

Now we can construct our model in `statsmodels` using the OLS function.

We will use `pandas` dataframes with `statsmodels`, however standard arrays can also be used as arguments

```
reg1 = sm.OLS(endog=df1['logpgp95'], exog=df1[['const', 'avexpr']], \
    missing='drop')
type(reg1)
```

```
statsmodels.regression.linear_model.OLS
```

So far we have simply constructed our model.

We need to use `.fit()` to obtain parameter estimates $\hat{\beta}_0$ and $\hat{\beta}_1$

```
results = reg1.fit()
type(results)
```

```
statsmodels.regression.linear_model.RegressionResultsWrapper
```

We now have the fitted regression model stored in `results`.

To view the OLS regression results, we can call the `.summary()` method.

Note that an observation was mistakenly dropped from the results in the original paper (see the note located in `maketable2.do` from Acemoglu's webpage), and thus the coefficients differ slightly.

```
print(results.summary())
```

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                logpgp95   R-squared:                       0.611
Model:                            OLS   Adj. R-squared:                  0.608
Method:                 Least Squares   F-statistic:                     171.4
Date:                Mon, 02 Nov 2026   Prob (F-statistic):           4.16e-24
Time:                        04:14:55   Log-Likelihood:                -119.71
No. Observations:                 111   AIC:                             243.4
Df Residuals:                     109   BIC:                             248.8
Df Model:                           1
Covariance Type:            nonrobust
==============================================================================
                 coef    std err          t      P>|t|      [0.025      0.975]
------------------------------------------------------------------------------
const          4.6261      0.301     15.391      0.000       4.030       5.222
avexpr         0.5319      0.041     13.093      0.000       0.451       0.612
==============================================================================
Omnibus:                        9.251   Durbin-Watson:                   1.689
Prob(Omnibus):                  0.010   Jarque-Bera (JB):                9.170
Skew:                          -0.680   Prob(JB):                       0.0102
Kurtosis:                       3.362   Cond. No.                         33.2
==============================================================================

Notes:
[1] Standard Errors assume that the covariance matrix of the errors is correctly␣
↪specified.
```

From our results, we see that

- The intercept $\hat{\beta}_0 = 4.63$.
- The slope $\hat{\beta}_1 = 0.53$.
- The positive $\hat{\beta}_1$ parameter estimate implies that. institutional quality has a positive effect on economic outcomes, as we saw in the figure.
- The p-value of 0.000 for $\hat{\beta}_1$ implies that the effect of institutions on GDP is statistically significant (using p < 0.05 as a rejection rule).
- The R-squared value of 0.611 indicates that around 61% of variation in log GDP per capita is explained by protection against expropriation.

Using our parameter estimates, we can now write our estimated relationship as

$$\widehat{logpgp95}_i = 4.63 + 0.53 \, avexpr_i$$

This equation describes the line that best fits our data, as shown in Figure 2.

We can use this equation to predict the level of log GDP per capita for a value of the index of expropriation protection.

For example, for a country with an index value of 7.07 (the average for the dataset), we find that their predicted level of log GDP per capita in 1995 is 8.38.

```
mean_expr = np.mean(df1_subset['avexpr'])
mean_expr
```

```
np.float32(6.515625)
```

```
predicted_logpdp95 = 4.63 + 0.53 * 7.07
predicted_logpdp95
```

```
8.3771
```

An easier (and more accurate) way to obtain this result is to use `.predict()` and set $constant = 1$ and $avexpr_i = mean_expr$

```
results.predict(exog=[1, mean_expr])
```

```
array([8.09156367])
```

We can obtain an array of predicted $logpgp95_i$ for every value of $avexpr_i$ in our dataset by calling `.predict()` on our results.

Plotting the predicted values against $avexpr_i$ shows that the predicted values lie along the linear line that we fitted above.

The observed values of $logpgp95_i$ are also plotted for comparison purposes

```
# Drop missing observations from whole sample

df1_plot = df1.dropna(subset=['logpgp95', 'avexpr'])

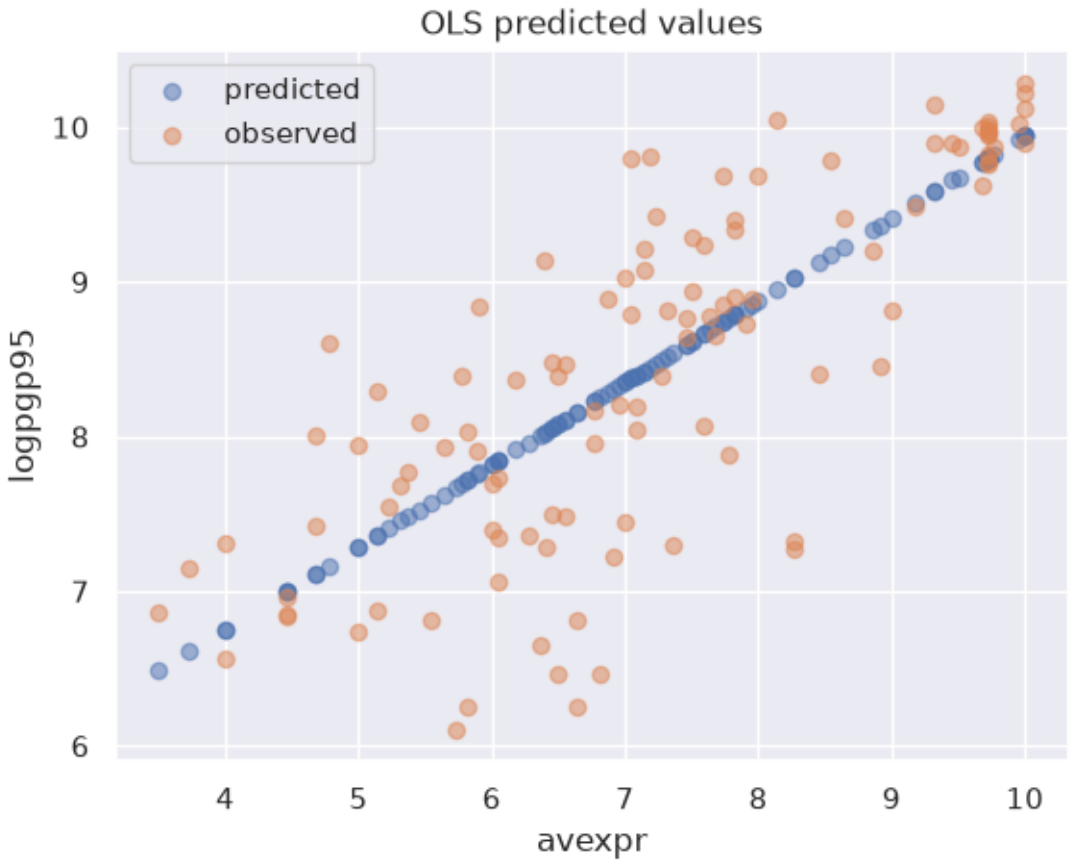
# Plot predicted values

fix, ax = plt.subplots()
ax.scatter(df1_plot['avexpr'], results.predict(), alpha=0.5,
        label='predicted')

# Plot observed values

ax.scatter(df1_plot['avexpr'], df1_plot['logpgp95'], alpha=0.5,
        label='observed')

ax.legend()
ax.set_title('OLS predicted values')
ax.set_xlabel('avexpr')
ax.set_ylabel('logpgp95')
plt.show()
```



144.3 Extending the Linear Regression Model

So far we have only accounted for institutions affecting economic performance - almost certainly there are numerous other factors affecting GDP that are not included in our model.

Leaving out variables that affect $logpgp95_i$ will result in **omitted variable bias**, yielding biased and inconsistent parameter estimates.

We can extend our bivariate regression model to a **multivariate regression model** by adding in other factors that may affect $logpgp95_i$.

[Acemoglu *et al.*, 2001] consider other factors such as:

- the effect of climate on economic outcomes; latitude is used to proxy this
- differences that affect both economic performance and institutions, eg. cultural, historical, etc.; controlled for with the use of continent dummies

Let's estimate some of the extended models considered in the paper (Table 2) using data from `maketable2.dta`

```
df2 = pd.read_stata('https://github.com/QuantEcon/data-lectures/raw/main/lectures/
↪maketable2.dta')

# Add constant term to dataset
df2['const'] = 1
```

(continues on next page)

(continued from previous page)

```
# Create lists of variables to be used in each regression
X1 = ['const', 'avexpr']
X2 = ['const', 'avexpr', 'lat_abst']
X3 = ['const', 'avexpr', 'lat_abst', 'asia', 'africa', 'other']

# Estimate an OLS regression for each set of variables
reg1 = sm.OLS(df2['logpgp95'], df2[X1], missing='drop').fit()
reg2 = sm.OLS(df2['logpgp95'], df2[X2], missing='drop').fit()
reg3 = sm.OLS(df2['logpgp95'], df2[X3], missing='drop').fit()
```

Now that we have fitted our model, we will use `summary_col` to display the results in a single table (model numbers correspond to those in the paper)

```
info_dict={'R-squared' : lambda x: f"{x.rsquared:.2f}",
           'No. observations' : lambda x: f"{int(x.nobs):d}"}

results_table = summary_col(results=[reg1,reg2,reg3],
                            float_format='%0.2f',
                            stars = True,
                            model_names=['Model 1',
                                         'Model 3',
                                         'Model 4'],
                            info_dict=info_dict,
                            regressor_order=['const',
                                             'avexpr',
                                             'lat_abst',
                                             'asia',
                                             'africa'])

results_table.add_title('Table 2 - OLS Regressions')

print(results_table)
```

```
        Table 2 - OLS Regressions
=========================================
                Model 1 Model 3 Model 4
-----------------------------------------
const           4.63*** 4.87*** 5.85***
                (0.30)  (0.33)  (0.34)
avexpr          0.53*** 0.46*** 0.39***
                (0.04)  (0.06)  (0.05)
lat_abst                0.87*   0.33
                        (0.49)  (0.45)
asia                            -0.15
                                (0.15)
africa                          -0.92***
                                (0.17)
other                           0.30
                                (0.37)
R-squared       0.61    0.62    0.72
R-squared Adj.  0.61    0.62    0.70
No. observations 111    111     111
R-squared       0.61    0.62    0.72
=========================================
Standard errors in parentheses.
* p<.1, ** p<.05, ***p<.01
```

144.4 Endogeneity

As [Acemoglu *et al.*, 2001] discuss, the OLS models likely suffer from **endogeneity** issues, resulting in biased and inconsistent model estimates.

Namely, there is likely a two-way relationship between institutions and economic outcomes:

- richer countries may be able to afford or prefer better institutions
- variables that affect income may also be correlated with institutional differences
- the construction of the index may be biased; analysts may be biased towards seeing countries with higher income having better institutions

To deal with endogeneity, we can use **two-stage least squares (2SLS) regression**, which is an extension of OLS regression.

This method requires replacing the endogenous variable $avexpr_i$ with a variable that is:

1. correlated with $avexpr_i$
2. not correlated with the error term (ie. it should not directly affect the dependent variable, otherwise it would be correlated with u_i due to omitted variable bias)

The new set of regressors is called an **instrument**, which aims to remove endogeneity in our proxy of institutional differences.

The main contribution of [Acemoglu *et al.*, 2001] is the use of settler mortality rates to instrument for institutional differences.

They hypothesize that higher mortality rates of colonizers led to the establishment of institutions that were more extractive in nature (less protection against expropriation), and these institutions still persist today.

Using a scatterplot (Figure 3 in [Acemoglu *et al.*, 2001]), we can see protection against expropriation is negatively correlated with settler mortality rates, coinciding with the authors' hypothesis and satisfying the first condition of a valid instrument.

```
# Dropping NA's is required to use numpy's polyfit
df1_subset2 = df1.dropna(subset=['logem4', 'avexpr'])

X = df1_subset2['logem4']
y = df1_subset2['avexpr']
labels = df1_subset2['shortnam']

# Replace markers with country labels
fig, ax = plt.subplots()
ax.scatter(X, y, marker='')

for i, label in enumerate(labels):
    ax.annotate(label, (X.iloc[i], y.iloc[i]))

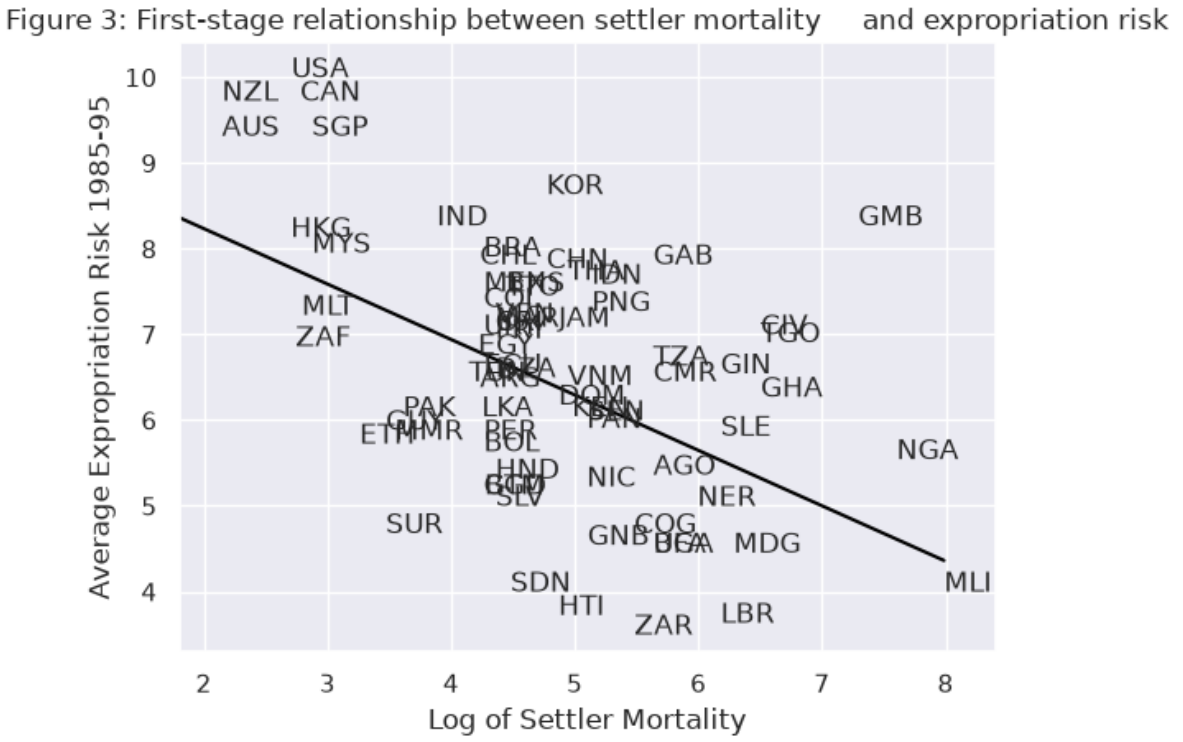
# Fit a linear trend line
ax.plot(np.unique(X),
         np.poly1d(np.polyfit(X, y, 1))(np.unique(X)),
         color='black')

ax.set_xlim([1.8,8.4])
ax.set_ylim([3.3,10.4])
ax.set_xlabel('Log of Settler Mortality')
ax.set_ylabel('Average Expropriation Risk 1985-95')
ax.set_title('Figure 3: First-stage relationship between settler mortality \
```

(continues on next page)

(continued from previous page)

```
    and expropriation risk')
plt.show()
```



The second condition may not be satisfied if settler mortality rates in the 17th to 19th centuries have a direct effect on current GDP (in addition to their indirect effect through institutions).

For example, settler mortality rates may be related to the current disease environment in a country, which could affect current economic performance.

[Acemoglu *et al.*, 2001] argue this is unlikely because:

- The majority of settler deaths were due to malaria and yellow fever and had a limited effect on local people.
- The disease burden on local people in Africa or India, for example, did not appear to be higher than average, supported by relatively high population densities in these areas before colonization.

As we appear to have a valid instrument, we can use 2SLS regression to obtain consistent and unbiased parameter estimates.

First stage

The first stage involves regressing the endogenous variable ($avexpr_i$) on the instrument.

The instrument is the set of all exogenous variables in our model (and not just the variable we have replaced).

Using model 1 as an example, our instrument is simply a constant and settler mortality rates $logem4_i$.

Therefore, we will estimate the first-stage regression as

$$avexpr_i = \delta_0 + \delta_1 logem4_i + v_i$$

The data we need to estimate this equation is located in `maketable4.dta` (only complete data, indicated by `baseco = 1`, is used for estimation)

```
# Import and select the data
df4 = pd.read_stata('https://github.com/QuantEcon/data-lectures/raw/main/lectures/
 ↪maketable4.dta')
df4 = df4[df4['baseco'] == 1]

# Add a constant variable
df4['const'] = 1

# Fit the first stage regression and print summary
results_fs = sm.OLS(df4['avexpr'],
                    df4[['const', 'logem4']],
                    missing='drop').fit()
print(results_fs.summary())
```

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                  avexpr   R-squared:                       0.270
Model:                             OLS   Adj. R-squared:                  0.258
Method:                  Least Squares   F-statistic:                     22.95
Date:                 Mon, 02 Nov 2026   Prob (F-statistic):           1.08e-05
Time:                         04:14:56   Log-Likelihood:                -104.83
No. Observations:                   64   AIC:                             213.7
Df Residuals:                       62   BIC:                             218.0
Df Model:                            1
Covariance Type:             nonrobust
==============================================================================
                 coef    std err          t      P>|t|      [0.025      0.975]
------------------------------------------------------------------------------
const          9.3414      0.611     15.296      0.000       8.121      10.562
logem4        -0.6068      0.127     -4.790      0.000      -0.860      -0.354
==============================================================================
Omnibus:                        0.035   Durbin-Watson:                   2.003
Prob(Omnibus):                  0.983   Jarque-Bera (JB):                0.172
Skew:                           0.045   Prob(JB):                        0.918
Kurtosis:                       2.763   Cond. No.                         19.4
==============================================================================

Notes:
[1] Standard Errors assume that the covariance matrix of the errors is correctly␣
 ↪specified.
```

Second stage

We need to retrieve the predicted values of $avexpr_i$ using `.predict()`.

We then replace the endogenous variable $avexpr_i$ with the predicted values $\widehat{avexpr}_i$ in the original linear model.

Our second stage regression is thus

$$logpgp95_i = \beta_0 + \beta_1 \widehat{avexpr}_i + u_i$$

```
df4['predicted_avexpr'] = results_fs.predict()

results_ss = sm.OLS(df4['logpgp95'],
                    df4[['const', 'predicted_avexpr']]).fit()
print(results_ss.summary())
```

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                logpgp95   R-squared:                       0.477
Model:                            OLS   Adj. R-squared:                  0.469
Method:                 Least Squares   F-statistic:                     56.60
Date:                Mon, 02 Nov 2026   Prob (F-statistic):           2.66e-10
Time:                        04:14:56   Log-Likelihood:                -72.268
No. Observations:                  64   AIC:                             148.5
Df Residuals:                      62   BIC:                             152.9
Df Model:                           1
Covariance Type:            nonrobust
==============================================================================
                 coef    std err          t      P>|t|      [0.025      0.
↪975]
------------------------------------------------------------------------------
↪-
const          1.9097      0.823      2.320      0.024       0.264       3.
↪555
predicted_avexpr     0.9443      0.126      7.523      0.000       0.693       1.
↪195
==============================================================================
Omnibus:                       10.547   Durbin-Watson:                   2.137
Prob(Omnibus):                  0.005   Jarque-Bera (JB):               11.010
Skew:                          -0.790   Prob(JB):                      0.00407
Kurtosis:                       4.277   Cond. No.                         58.1
==============================================================================

Notes:
[1] Standard Errors assume that the covariance matrix of the errors is correctly␣
↪specified.
```

The second-stage regression results give us an unbiased and consistent estimate of the effect of institutions on economic outcomes.

The result suggests a stronger positive relationship than what the OLS results indicated.

Note that while our parameter estimates are correct, our standard errors are not and for this reason, computing 2SLS 'manually' (in stages with OLS) is not recommended.

We can correctly estimate a 2SLS regression in one step using the linearmodels package, an extension of `statsmodels`

Note that when using `IV2SLS`, the exogenous and instrument variables are split up in the function arguments (whereas before the instrument included exogenous variables)

```
iv = IV2SLS(dependent=df4['logpgp95'],
            exog=df4['const'],
            endog=df4['avexpr'],
            instruments=df4['logem4']).fit(cov_type='unadjusted')

print(iv.summary)
```

```
                          IV-2SLS Estimation Summary
==============================================================================
Dep. Variable:                logpgp95   R-squared:                      0.1870
Estimator:                     IV-2SLS   Adj. R-squared:                 0.1739
No. Observations:                  64   F-statistic:                    37.568
Date:                Mon, Nov 02 2026   P-value (F-stat)                0.0000
Time:                        04:14:56   Distribution:                  chi2(1)
```

(continues on next page)

(continued from previous page)

```
Cov. Estimator:            unadjusted

                            Parameter Estimates
==============================================================================
            Parameter  Std. Err.     T-stat    P-value    Lower CI    Upper CI
------------------------------------------------------------------------------
const          1.9097     1.0106     1.8897     0.0588     -0.0710      3.8903
avexpr         0.9443     0.1541     6.1293     0.0000      0.6423      1.2462
==============================================================================

Endogenous: avexpr
Instruments: logem4
Unadjusted Covariance (Homoskedastic)
Debiased: False
```

Given that we now have consistent and unbiased estimates, we can infer from the model we have estimated that institutional differences (stemming from institutions set up during colonization) can help to explain differences in income levels across countries today.

[Acemoglu *et al.*, 2001] use a marginal effect of 0.94 to calculate that the difference in the index between Chile and Nigeria (ie. institutional quality) implies up to a 7-fold difference in income, emphasizing the significance of institutions in economic development.

144.5 Summary

We have demonstrated basic OLS and 2SLS regression in `statsmodels` and `linearmodels`.

If you are familiar with R, you may want to use the formula interface to `statsmodels`, or consider using r2py to call R from within Python.

144.6 Exercises

Exercise 144.6.1

In the lecture, we think the original model suffers from endogeneity bias due to the likely effect income has on institutional development.

Although endogeneity is often best identified by thinking about the data and model, we can formally test for endogeneity using the **Hausman test**.

We want to test for correlation between the endogenous variable, $avexpr_i$, and the errors, u_i

$$H_0 : Cov(avexpr_i, u_i) = 0 \quad (no\ endogeneity)$$
$$H_1 : Cov(avexpr_i, u_i) \neq 0 \quad (endogeneity)$$

This test is running in two stages.

First, we regress $avexpr_i$ on the instrument, $logem4_i$

$$avexpr_i = \pi_0 + \pi_1 logem4_i + \upsilon_i$$

Second, we retrieve the residuals $\hat{v}_i$ and include them in the original equation

$$logpgp95_i = \beta_0 + \beta_1 avexpr_i + \alpha\hat{v}_i + u_i$$

If α is statistically significant (with a p-value < 0.05), then we reject the null hypothesis and conclude that $avexpr_i$ is endogenous.

Using the above information, estimate a Hausman test and interpret your results.

Solution

```
# Load in data
df4 = pd.read_stata('https://github.com/QuantEcon/data-lectures/raw/main/lectures/
 ↪maketable4.dta')

# Add a constant term
df4['const'] = 1

# Estimate the first stage regression
reg1 = sm.OLS(endog=df4['avexpr'],
              exog=df4[['const', 'logem4']],
              missing='drop').fit()

# Retrieve the residuals
df4['resid'] = reg1.resid

# Estimate the second stage residuals
reg2 = sm.OLS(endog=df4['logpgp95'],
              exog=df4[['const', 'avexpr', 'resid']],
              missing='drop').fit()

print(reg2.summary())
```

```
                            OLS Regression Results
==============================================================================
Dep. Variable:                logpgp95   R-squared:                       0.689
Model:                            OLS   Adj. R-squared:                  0.679
Method:                 Least Squares   F-statistic:                     74.05
Date:                Mon, 02 Nov 2026   Prob (F-statistic):           1.07e-17
Time:                        04:14:56   Log-Likelihood:                -62.031
No. Observations:                  70   AIC:                             130.1
Df Residuals:                      67   BIC:                             136.8
Df Model:                           2
Covariance Type:            nonrobust
==============================================================================
                 coef    std err          t      P>|t|      [0.025      0.975]
------------------------------------------------------------------------------
const          2.4782      0.547      4.530      0.000       1.386       3.570
avexpr         0.8564      0.082     10.406      0.000       0.692       1.021
resid         -0.4951      0.099     -5.017      0.000      -0.692      -0.298
==============================================================================
Omnibus:                       17.597   Durbin-Watson:                   2.086
Prob(Omnibus):                  0.000   Jarque-Bera (JB):               23.194
Skew:                          -1.054   Prob(JB):                     9.19e-06
Kurtosis:                       4.873   Cond. No.                         53.8
==============================================================================

Notes:
[1] Standard Errors assume that the covariance matrix of the errors is␣
 ↪correctly specified.
```

The output shows that the coefficient on the residuals is statistically significant, indicating $avexpr_i$ is endogenous.

Exercise 144.6.2

The OLS parameter β can also be estimated using matrix algebra and `numpy` (you may need to review the numpy lecture to complete this exercise).

The linear equation we want to estimate is (written in matrix form)

$$y = X\beta + u$$

To solve for the unknown parameter β, we want to minimize the sum of squared residuals

$$\min_{\hat{\beta}} \hat{u}'\hat{u}$$

Rearranging the first equation and substituting into the second equation, we can write

$$\min_{\hat{\beta}} (Y - X\hat{\beta})'(Y - X\hat{\beta})$$

Solving this optimization problem gives the solution for the $\hat{\beta}$ coefficients

$$\hat{\beta} = (X'X)^{-1}X'y$$

Using the above information, compute $\hat{\beta}$ from model 1 using `numpy` - your results should be the same as those in the `statsmodels` output from earlier in the lecture.

Solution

```
# Load in data
df1 = pd.read_stata('https://github.com/QuantEcon/data-lectures/raw/main/lectures/
↪maketable1.dta')
df1 = df1.dropna(subset=['logpgp95', 'avexpr'])

# Add a constant term
df1['const'] = 1

# Define the X and y variables
y = np.asarray(df1['logpgp95'])
X = np.asarray(df1[['const', 'avexpr']])

# Compute β_hat
β_hat = np.linalg.solve(X.T @ X, X.T @ y)

# Print out the results from the 2 x 1 vector β_hat
print(f'β_0 = {β_hat[0]:.2}')
print(f'β_1 = {β_hat[1]:.2}')
```

```
β_0 = 4.6
β_1 = 0.53
```

It is also possible to use `np.linalg.inv(X.T @ X) @ X.T @ y` to solve for β, however `.solve()` is

preferred as it involves fewer computations.

CHAPTER

FIVE

MAXIMUM LIKELIHOOD ESTIMATION

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

Contents

- *Maximum Likelihood Estimation*
  - *Overview*
  - *Set up and assumptions*
  - *Conditional distributions*
  - *Maximum likelihood estimation*
  - *MLE with numerical methods*
  - *Maximum likelihood estimation with `statsmodels`*
  - *Summary*
  - *Exercises*

In addition to what's in Anaconda, this lecture will need the following libraries:

```
!pip install jax
```

145.1 Overview

In *Linear Regression in Python*, we estimated the relationship between dependent and explanatory variables using linear regression.

But what if a linear relationship is not an appropriate assumption for our model?

One widely used alternative is maximum likelihood estimation, which involves specifying a class of distributions, indexed by unknown parameters, and then using the data to pin down these parameter values.

The benefit relative to linear regression is that it allows more flexibility in the probabilistic relationships between variables.

Here we illustrate maximum likelihood by replicating Daniel Treisman's (2016) paper, Russia's Billionaires, which connects the number of billionaires in a country to its economic characteristics.

The paper concludes that Russia has a higher number of billionaires than economic factors such as market size and tax rate predict.

We'll require the following imports:

```
import numpy as np
import jax.numpy as jnp
import jax
import pandas as pd
from typing import NamedTuple

from jax.scipy.special import factorial, gammaln
from jax.scipy.stats import norm

from statsmodels.api import Poisson
from statsmodels.iolib.summary2 import summary_col

import matplotlib.pyplot as plt
from mpl_toolkits.mplot3d import Axes3D
```

145.1.1 Prerequisites

We assume familiarity with basic probability and multivariate calculus.

145.2 Set up and assumptions

Let's consider the steps we need to go through in maximum likelihood estimation and how they pertain to this study.

145.2.1 Flow of ideas

The first step with maximum likelihood estimation is to choose the probability distribution believed to be generating the data.

More precisely, we need to make an assumption as to which *parametric class* of distributions is generating the data.

- e.g., the class of all normal distributions, or the class of all gamma distributions.

Each such class is a family of distributions indexed by a finite number of parameters.

- e.g., the class of normal distributions is a family of distributions indexed by its mean $\mu \in (-\infty, \infty)$ and standard deviation $\sigma \in (0, \infty)$.

We'll let the data pick out a particular element of the class by pinning down the parameters.

The parameter estimates so produced will be called **maximum likelihood estimates**.

145.2.2 Counting billionaires

Treisman [Treisman, 2016] is interested in estimating the number of billionaires in different countries.

The number of billionaires is integer-valued.

Hence we consider distributions that take values only in the nonnegative integers.

(This is one reason least squares regression is not the best tool for the present problem, since the dependent variable in linear regression is not restricted to integer values.)

One integer distribution is the Poisson distribution, the probability mass function (pmf) of which is

$$f(y) = \frac{\mu^y}{y!} e^{-\mu}, \qquad y = 0, 1, 2, \ldots, \infty$$

We can plot the Poisson distribution over y for different values of μ as follows

```
@jax.jit
def poisson_pmf(y, μ):
    return μ**y / factorial(y) * jnp.exp(-μ)
```

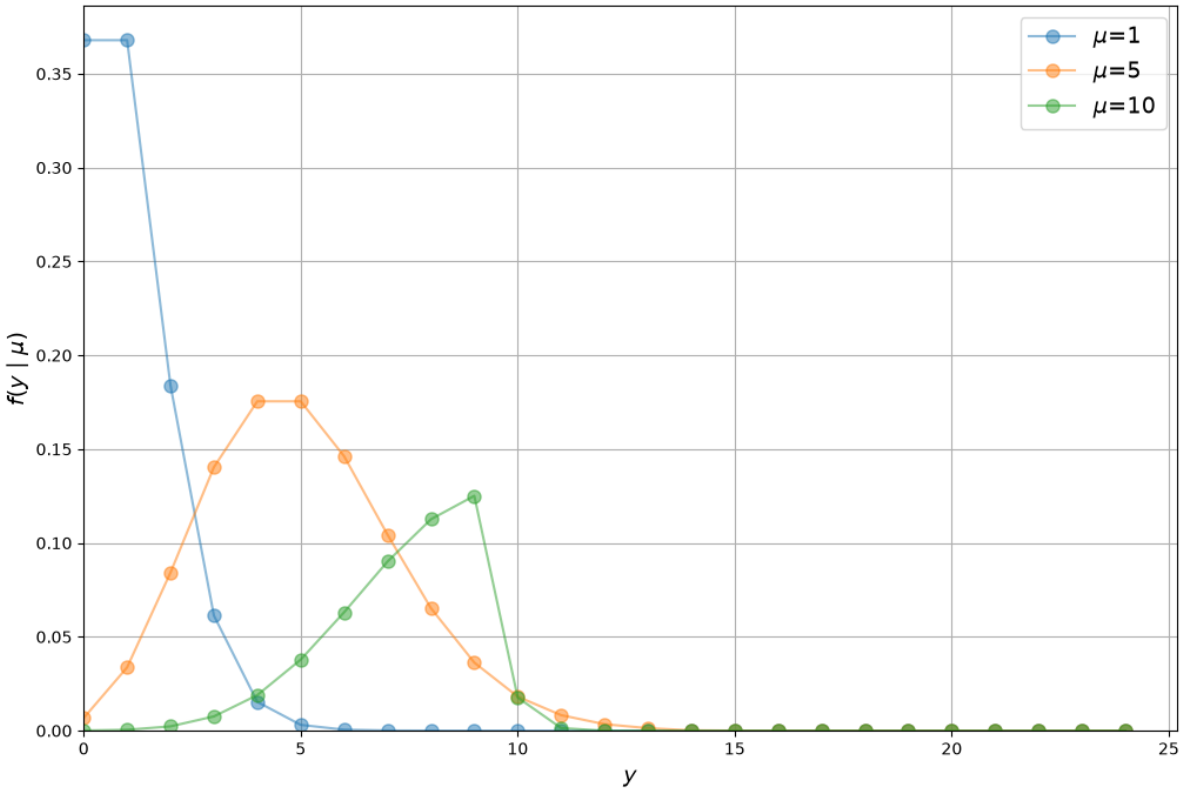
```
y_values = range(0, 25)

fig, ax = plt.subplots(figsize=(12, 8))

for μ in [1, 5, 10]:
    distribution = []
    for y_i in y_values:
        distribution.append(poisson_pmf(y_i, μ))
    ax.plot(
        y_values,
        distribution,
        label=rf"$\mu$={μ}",
        alpha=0.5,
        marker="o",
        markersize=8,
    )

ax.grid()
ax.set_xlabel(r"$y$", fontsize=14)
ax.set_ylabel(r"$f(y \mid \mu)$", fontsize=14)
ax.axis(xmin=0, ymin=0)
ax.legend(fontsize=14)

plt.show()
```



Notice that the Poisson distribution begins to resemble a normal distribution as the mean of y increases.

Let's have a look at the distribution of the data we'll be working with in this lecture.

Treisman's main source of data is *Forbes'* annual rankings of billionaires and their estimated net worth.

The dataset `fp.dta` can be downloaded from here or its AER page.

```
# Load in data and view
df = pd.read_stata(
    "https://github.com/QuantEcon/data-lectures/raw/main/lectures/fp.dta"
)
df.head()
```

```
         country  ccode    year    cyear  numbil  numbil0  numbilall  netw  \
0  United States    2.0  1990.0  21990.0     NaN      NaN        NaN   NaN
1  United States    2.0  1991.0  21991.0     NaN      NaN        NaN   NaN
2  United States    2.0  1992.0  21992.0     NaN      NaN        NaN   NaN
3  United States    2.0  1993.0  21993.0     NaN      NaN        NaN   NaN
4  United States    2.0  1994.0  21994.0     NaN      NaN        NaN   NaN

   netw0  netwall  ...  gattwto08     mcapbdol    mcapbdol08  lnmcap08  \
0    NaN      NaN  ...        0.0  3060.000000  11737.599609  9.370638
1    NaN      NaN  ...        0.0  4090.000000  11737.599609  9.370638
2    NaN      NaN  ...        0.0  4490.000000  11737.599609  9.370638
3    NaN      NaN  ...        0.0  5136.198730  11737.599609  9.370638
4    NaN      NaN  ...        0.0  5067.016113  11737.599609  9.370638

   topintaxnew   topint08     rintr  noyrs  roflaw  nrrents
0    39.799999  39.799999  4.988405   20.0    1.61      NaN
1    39.799999  39.799999  4.988405   20.0    1.61      NaN
```

(continues on next page)

(continued from previous page)

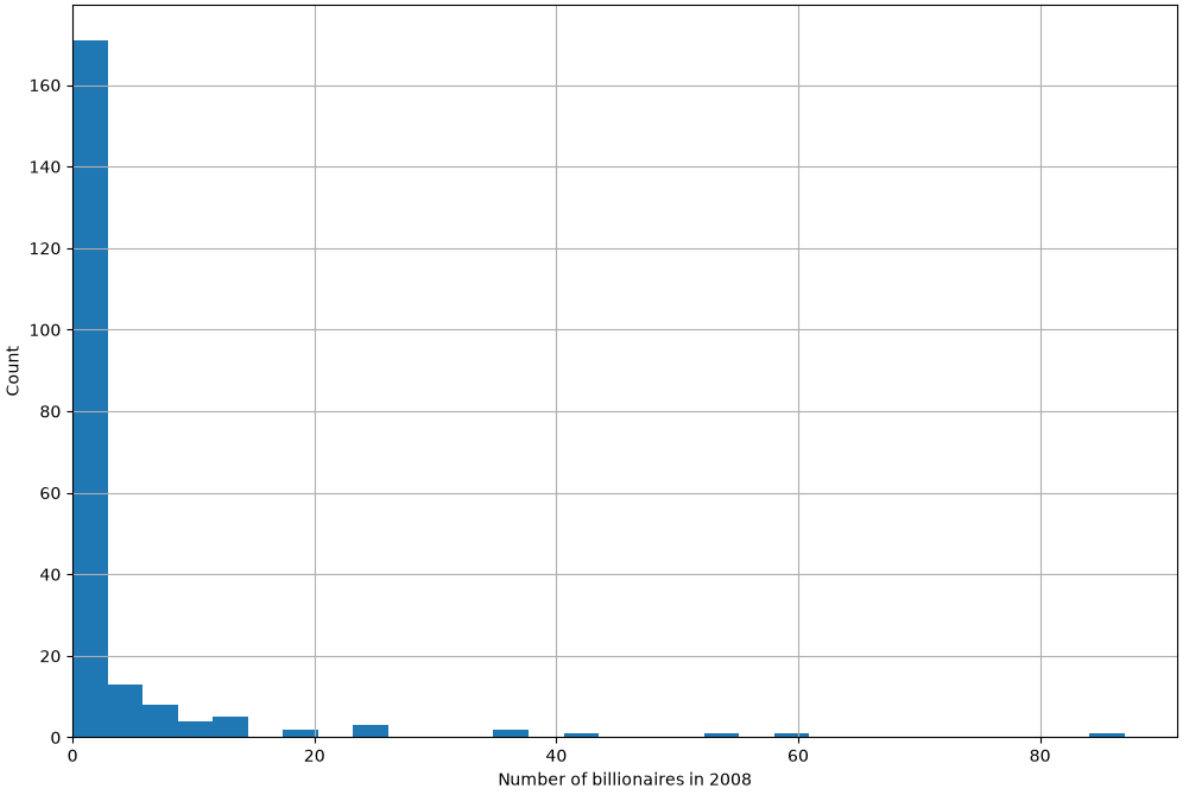
```
2    39.799999  39.799999  4.988405   20.0    1.61      NaN
3    39.799999  39.799999  4.988405   20.0    1.61      NaN
4    39.799999  39.799999  4.988405   20.0    1.61      NaN

[5 rows x 36 columns]
```

Using a histogram, we can view the distribution of the number of billionaires per country, `numbil0`, in 2008 (the United States is dropped for plotting purposes)

```
numbil0_2008 = df[
    (df["year"] == 2008) & (df["country"] != "United States")
].loc[:, "numbil0"]

plt.subplots(figsize=(12, 8))
plt.hist(numbil0_2008, bins=30)
plt.xlim(left=0)
plt.grid()
plt.xlabel("Number of billionaires in 2008")
plt.ylabel("Count")
plt.show()
```



From the histogram, it appears that the Poisson assumption is not unreasonable (albeit with a very low μ and some outliers).

145.3 Conditional distributions

In Treisman's paper, the dependent variable — the number of billionaires y_i in country i — is modeled as a function of GDP per capita, population size, and years membership in GATT and WTO.

Hence, the distribution of y_i needs to be conditioned on the vector of explanatory variables $\mathbf{x}_i$.

The standard formulation — the so-called **Poisson regression** model — is as follows:

$$f(y_i \mid \mathbf{x}_i) = \frac{\mu_i^{y_i}}{y_i!} e^{-\mu_i}; \qquad y_i = 0, 1, 2, \ldots, \infty. \tag{145.1}$$

$$\text{where } \mu_i = \exp(\mathbf{x}_i'\beta) = \exp(\beta_0 + \beta_1 x_{i1} + \ldots + \beta_k x_{ik})$$

To illustrate the idea that the distribution of y_i depends on $\mathbf{x}_i$ let's run a simple simulation.

We use our `poisson_pmf` function from above and arbitrary values for β and $\mathbf{x}_i$

```
y_values = range(0, 20)

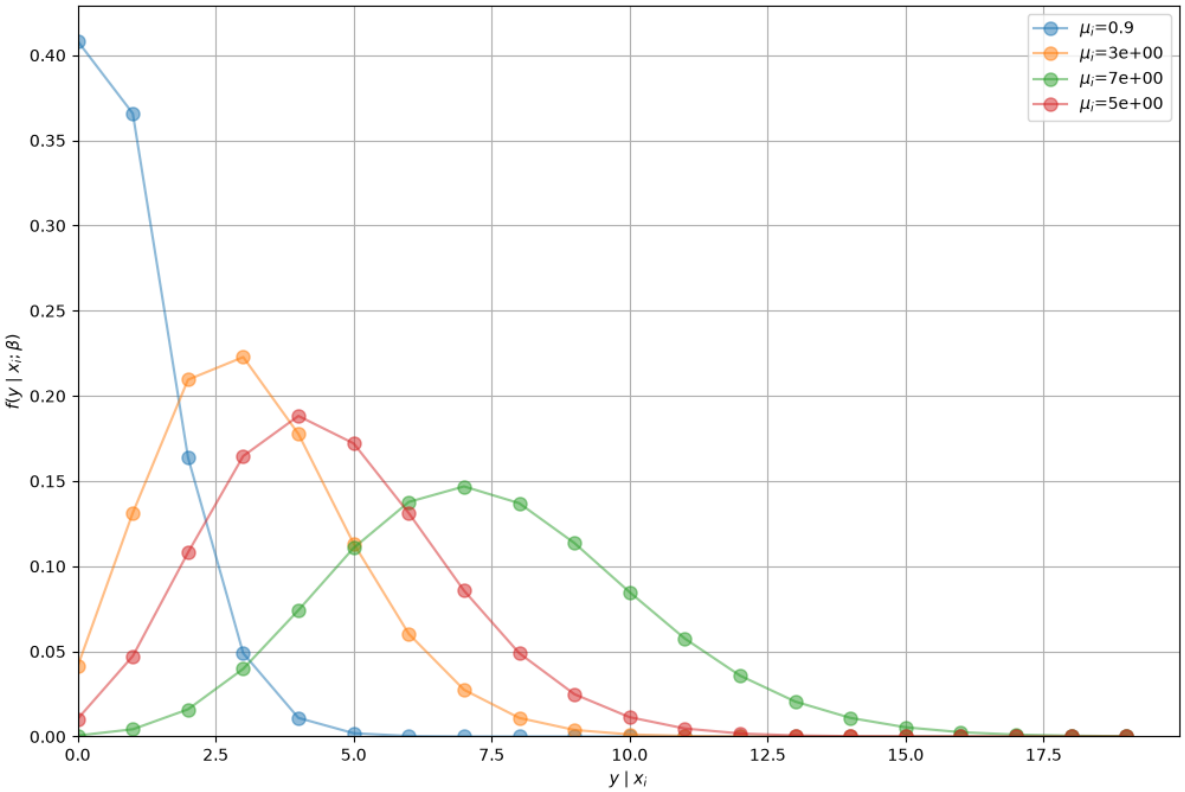
# Define a parameter vector with estimates
β = jnp.array([0.26, 0.18, 0.25, -0.1, -0.22])

# Create some observations X
datasets = [
    jnp.array([0, 1, 1, 1, 2]),
    jnp.array([2, 3, 2, 4, 0]),
    jnp.array([3, 4, 5, 3, 2]),
    jnp.array([6, 5, 4, 4, 7]),
]


fig, ax = plt.subplots(figsize=(12, 8))

for X in datasets:
    μ = jnp.exp(X @ β)
    distribution = []
    for y_i in y_values:
        distribution.append(poisson_pmf(y_i, μ))
    ax.plot(
        y_values,
        distribution,
        label=rf"$\mu_i$={μ:.1}",
        marker="o",
        markersize=8,
        alpha=0.5,
    )

ax.grid()
ax.legend()
ax.set_xlabel(r"$y \mid x_i$")
ax.set_ylabel(r"$f(y \mid x_i; \beta )$")
ax.axis(xmin=0, ymin=0)
plt.show()
```



We can see that the distribution of y_i is conditional on $\mathbf{x}_i$ (μ_i is no longer constant).

145.4 Maximum likelihood estimation

In our model for number of billionaires, the conditional distribution contains 4 ($k = 4$) parameters that we need to estimate.

We will label our entire parameter vector as β where

$$\beta = \begin{bmatrix} \beta_0 \\ \beta_1 \\ \beta_2 \\ \beta_3 \end{bmatrix}$$

To estimate the model using MLE, we want to maximize the likelihood that our estimate $\hat{\beta}$ is the true parameter β.

Intuitively, we want to find the $\hat{\beta}$ that best fits our data.

First, we need to construct the likelihood function $\mathcal{L}(\beta)$, which is similar to a joint probability density function.

Assume we have some data $y_i = \{y_1, y_2\}$ and $y_i \sim f(y_i)$.

If y_1 and y_2 are independent, the joint pmf of these data is $f(y_1, y_2) = f(y_1) \cdot f(y_2)$.

If y_i follows a Poisson distribution with $\lambda = 7$, we can visualize the joint pmf like so

```
def plot_joint_poisson(μ=7, y_n=20):
    yi_values = jnp.arange(0, y_n, 1)
```

(continues on next page)

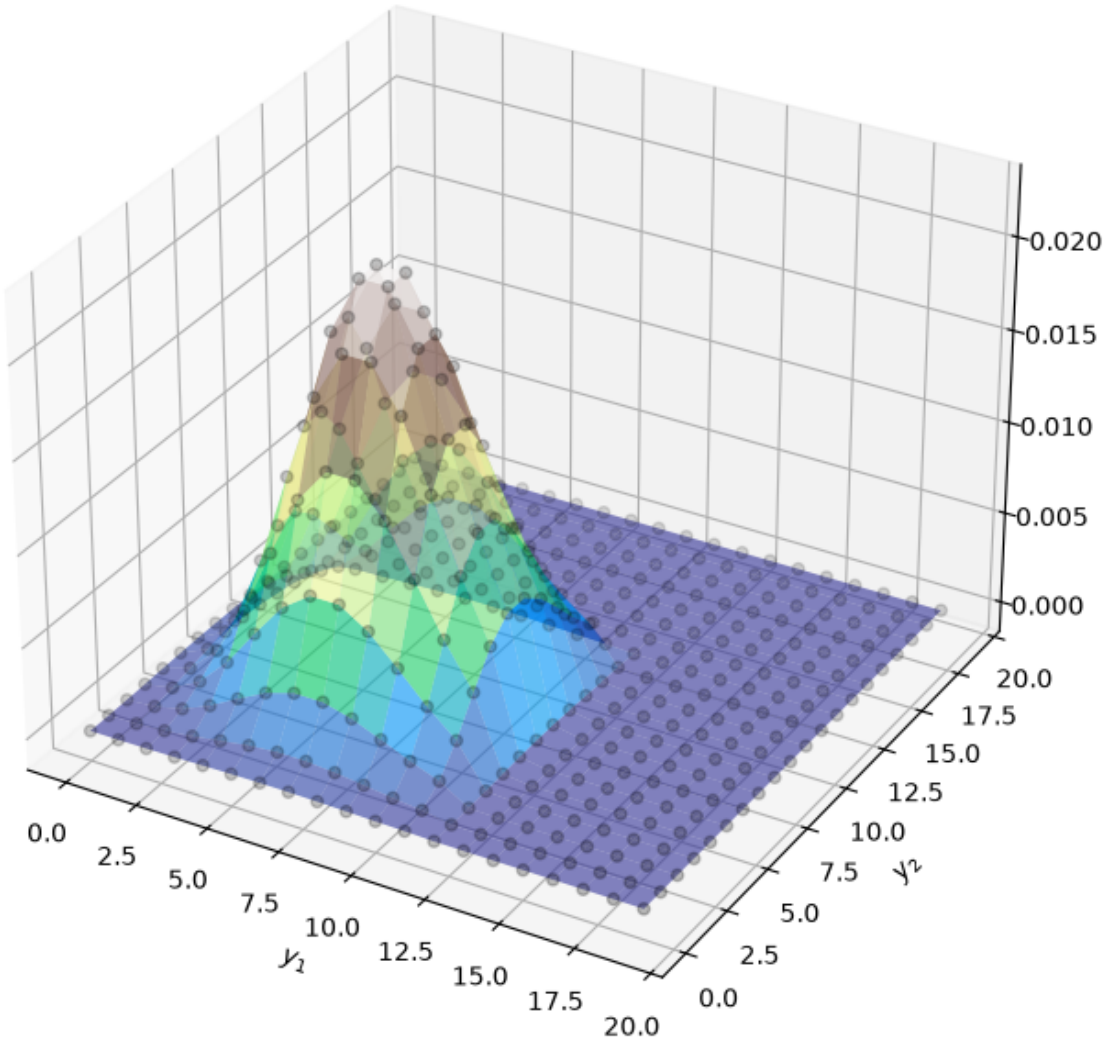
(continued from previous page)

```
    # Create coordinate points of X and Y
    X, Y = jnp.meshgrid(yi_values, yi_values)

    # Multiply distributions together
    Z = poisson_pmf(X, μ) * poisson_pmf(Y, μ)

    fig = plt.figure(figsize=(12, 8))
    ax = fig.add_subplot(111, projection="3d")
    ax.plot_surface(X, Y, Z.T, cmap="terrain", alpha=0.6)
    ax.scatter(X, Y, Z.T, color="black", alpha=0.5, linewidths=1)
    ax.set(xlabel=r"$y_1$", ylabel=r"$y_2$")
    ax.set_zlabel(r"$f(y_1, y_2)$", labelpad=10)
    plt.show()


plot_joint_poisson(μ=7, y_n=20)
```



Similarly, the joint pmf of our data (which is distributed as a conditional Poisson distribution) can be written as

$$f(y_1, y_2, \ldots, y_n \mid \mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n; \beta) = \prod_{i=1}^{n} \frac{\mu_i^{y_i}}{y_i!} e^{-\mu_i}$$

y_i is conditional on both the values of $\mathbf{x}_i$ and the parameters β.

The likelihood function is the same as the joint pmf, but treats the parameter β as a random variable and takes the observations $(y_i, \mathbf{x}_i)$ as given

$$\begin{aligned}\mathcal{L}(\beta \mid y_1, y_2, \ldots, y_n \; ; \; \mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n) &= \prod_{i=1}^{n} \frac{\mu_i^{y_i}}{y_i!} e^{-\mu_i} \\ &= f(y_1, y_2, \ldots, y_n \mid \mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_n; \beta)\end{aligned}$$

Now that we have our likelihood function, we want to find the $\widehat{\beta}$ that yields the maximum likelihood value

$$\max_{\beta} \mathcal{L}(\beta)$$

In doing so it is generally easier to maximize the log-likelihood (consider differentiating $f(x) = x \exp(x)$ vs. $f(x) = \log(x) + x$).

Given that taking a logarithm is a monotone increasing transformation, a maximizer of the likelihood function will also be a maximizer of the log-likelihood function.

In our case the log-likelihood is

$$\begin{aligned}\log \mathcal{L}(\beta) &= \log\Big(f(y_1; \beta) \cdot f(y_2; \beta) \cdot \ldots \cdot f(y_n; \beta)\Big) \\ &= \sum_{i=1}^{n} \log f(y_i; \beta) \\ &= \sum_{i=1}^{n} \log\Big(\frac{\mu_i^{y_i}}{y_i!} e^{-\mu_i}\Big) \\ &= \sum_{i=1}^{n} y_i \log \mu_i - \sum_{i=1}^{n} \mu_i - \sum_{i=1}^{n} \log y_i!\end{aligned}$$

The MLE of the Poisson for $\widehat{\beta}$ can be obtained by solving

$$\max_{\beta}\Big(\sum_{i=1}^{n} y_i \log \mu_i - \sum_{i=1}^{n} \mu_i - \sum_{i=1}^{n} \log y_i!\Big)$$

However, no analytical solution exists to the above problem – to find the MLE we need to use numerical methods.

145.5 MLE with numerical methods

Many distributions do not have nice, analytical solutions and therefore require numerical methods to solve for parameter estimates.

One such numerical method is the Newton-Raphson algorithm.

Our goal is to find the maximum likelihood estimate $\widehat{\beta}$.

At $\widehat{\beta}$, the first derivative of the log-likelihood function will be equal to 0.

Let's illustrate this by supposing

$$\log \mathcal{L}(\beta) = -(\beta - 10)^2 - 10$$

```
@jax.jit
def logL(β):
    return -((β - 10) ** 2) - 10
```

To find the value of the gradient of the above function, we can use jax.grad which auto-differentiates the given function.

We further use jax.vmap which vectorizes the given function i.e. the function acting upon scalar inputs can now be used with vector inputs.

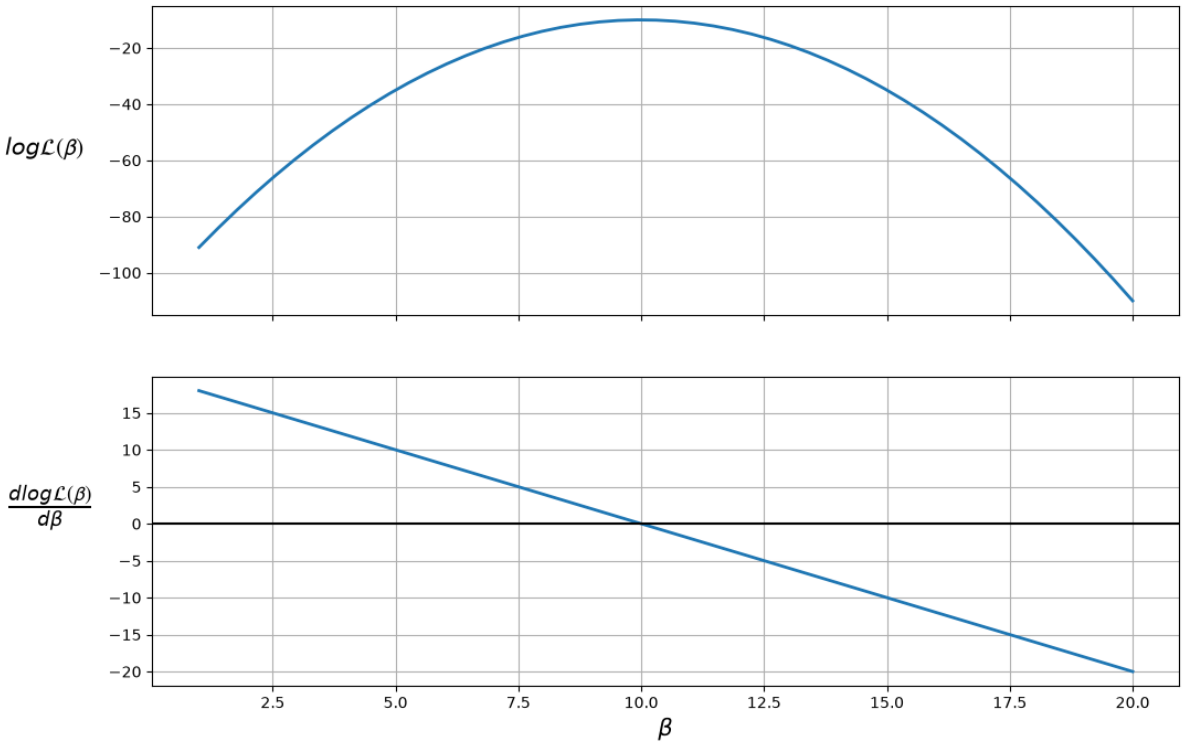
```
dlogL = jax.vmap(jax.grad(logL))
```

```
β = jnp.linspace(1, 20)

fig, (ax1, ax2) = plt.subplots(2, sharex=True, figsize=(12, 8))

ax1.plot(β, logL(β), lw=2)
ax2.plot(β, dlogL(β), lw=2)

ax1.set_ylabel(
    r"$log \mathcal{L(\beta)}$", rotation=0, labelpad=35, fontsize=15
)
ax2.set_ylabel(
    r"$\frac{dlog \mathcal{L(\beta)}}{d \beta}$ ",
    rotation=0,
    labelpad=35,
    fontsize=19,
)
ax2.set_xlabel(r"$\beta$", fontsize=15)
ax1.grid(), ax2.grid()
plt.axhline(c="black")
plt.show()
```



The plot shows that the maximum likelihood value (the top plot) occurs when $\frac{d \log \mathcal{L}(\beta)}{d\beta} = 0$ (the bottom plot).

Therefore, the likelihood is maximized when $\beta = 10$.

We can also ensure that this value is a *maximum* (as opposed to a minimum) by checking that the second derivative (slope of the bottom plot) is negative.

The Newton-Raphson algorithm finds a point where the first derivative is 0.

To use the algorithm, we take an initial guess at the maximum value, β_0 (the OLS parameter estimates might be a reasonable guess), then

1. Use the updating rule to iterate the algorithm

$$\beta_{(k+1)} = \beta_{(k)} - H^{-1}(\beta_{(k)})G(\beta_{(k)})$$

   where:

$$G(\beta_{(k)}) = \frac{d \log \mathcal{L}(\beta_{(k)})}{d\beta_{(k)}}$$
$$H(\beta_{(k)}) = \frac{d^2 \log \mathcal{L}(\beta_{(k)})}{d\beta_{(k)} d\beta'_{(k)}}$$

2. Check whether $\beta_{(k+1)} - \beta_{(k)} < tol$
   - If true, then stop iterating and set $\widehat{\beta} = \beta_{(k+1)}$
   - If false, then update $\beta_{(k+1)}$

As can be seen from the updating equation, $\beta_{(k+1)} = \beta_{(k)}$ only when $G(\beta_{(k)}) = 0$ i.e. where the first derivative is equal to 0.

(In practice, we stop iterating when the difference is below a small tolerance threshold.)

Let's have a go at implementing the Newton-Raphson algorithm.

First, we'll create a class called `PoissonRegression` so we can easily recompute the values of the log likelihood, gradient and Hessian for every iteration

```
class PoissonRegression(NamedTuple):
    X: jnp.ndarray
    y: jnp.ndarray
```

Now we can define the log likelihood function in Python

```
@jax.jit
def logL(β, model):
    y = model.y
    μ = jnp.exp(model.X @ β)
    return jnp.sum(model.y * jnp.log(μ) - μ - jnp.log(factorial(y)))
```

To find the gradient of the `poisson_logL`, we again use jax.grad.

According to the documentation,

- `jax.jacfwd` uses forward-mode automatic differentiation, which is more efficient for "tall" Jacobian matrices, while
- `jax.jacrev` uses reverse-mode, which is more efficient for "wide" Jacobian matrices.

(The documentation also states that when matrices that are near-square, `jax.jacfwd` probably has an edge over `jax.jacrev`.)

Therefore, to find the Hessian, we can directly use `jax.jacfwd`.

```
G_logL = jax.grad(logL)
H_logL = jax.jacfwd(G_logL)
```

Our function `newton_raphson` will take a `PoissonRegression` object that has an initial guess of the parameter vector β_0.

The algorithm will update the parameter vector according to the updating rule, and recalculate the gradient and Hessian matrices at the new parameter estimates.

Iteration will end when either:

- The difference between the parameter and the updated parameter is below a tolerance level.
- The maximum number of iterations has been achieved (meaning convergence is not achieved).

So we can get an idea of what's going on while the algorithm is running, an option `display=True` is added to print out values at each iteration.

```
def newton_raphson(model, β, tol=1e-3, max_iter=100, display=True):

    i = 0
    error = 100  # Initial error value

    # Print header of output
    if display:
        header = f'{"Iteration_k":<13}{"Log-likelihood":<16}{"θ":<60}'
        print(header)
        print("-" * len(header))

    # While loop runs while any value in error is greater
    # than the tolerance until max iterations are reached
    while jnp.any(error > tol) and i < max_iter:
        H, G = jnp.squeeze(H_logL(β, model)), G_logL(β, model)
        β_new = β - (jnp.dot(jnp.linalg.inv(H), G))
        error = jnp.abs(β_new - β)
        β = β_new

        if display:
            β_list = [f"{t:.3}" for t in list(β.flatten())]
            update = f"{i:<13}{logL(β, model):<16.8}{β_list}"
            print(update)

        i += 1

    print(f"Number of iterations: {i}")
    print(f"β_hat = {β.flatten()}")

    return β
```

Let's try out our algorithm with a small dataset of 5 observations and 3 variables in $\mathbf{X}$.

```
X = jnp.array([[1, 2, 5], [1, 1, 3], [1, 4, 2], [1, 5, 2], [1, 3, 1]])

y = jnp.array([1, 0, 1, 1, 0])
```

(continues on next page)

(continued from previous page)

```
# Take a guess at initial βs
init_β = jnp.array([0.1, 0.1, 0.1])

# Create an object with Poisson model values
poi = PoissonRegression(X=X, y=y)

# Use newton_raphson to find the MLE
β_hat = newton_raphson(poi, init_β, display=True)
```

```
Iteration_k  Log-likelihood  θ
--------------------------------------------------------------------------------
 ↪------
0             -4.3447633      ['-1.49', '0.265', '0.244']
1             -3.5742412      ['-3.38', '0.528', '0.474']
2             -3.3999527      ['-5.06', '0.782', '0.702']
3             -3.3788645      ['-5.92', '0.909', '0.82']
4             -3.3783557      ['-6.07', '0.933', '0.843']
5             -3.3783555      ['-6.08', '0.933', '0.843']
6             -3.3783555      ['-6.08', '0.933', '0.843']
Number of iterations: 7
β_hat = [-6.078486   0.9334028  0.8432968]
```

As this was a simple model with few observations, the algorithm achieved convergence in only 7 iterations.

You can see that with each iteration, the log-likelihood value increased.

Remember, our objective was to maximize the log-likelihood function, which the algorithm has worked to achieve.

Also, note that the increase in $\log \mathcal{L}(\beta_{(k)})$ becomes smaller with each iteration.

This is because the gradient is approaching 0 as we reach the maximum, and therefore the numerator in our updating equation is becoming smaller.

The gradient vector should be close to 0 at $\hat{\beta}$

```
G_logL(β_hat, poi)
```

```
Array([ 7.4505806e-09, -2.9802322e-07,  3.7252903e-08], dtype=float32)
```

The iterative process can be visualized in the following diagram, where the maximum is found at $\beta = 10$

```
@jax.jit
def logL(x):
    return -((x - 10) ** 2) - 10


@jax.jit
def find_tangent(β, a=0.01):
    y1 = logL(β)
    y2 = logL(β + a)
    x = jnp.array([[β, 1], [β + a, 1]])
    m, c = jnp.linalg.lstsq(x, jnp.array([y1, y2]), rcond=None)[0]
    return m, c
```

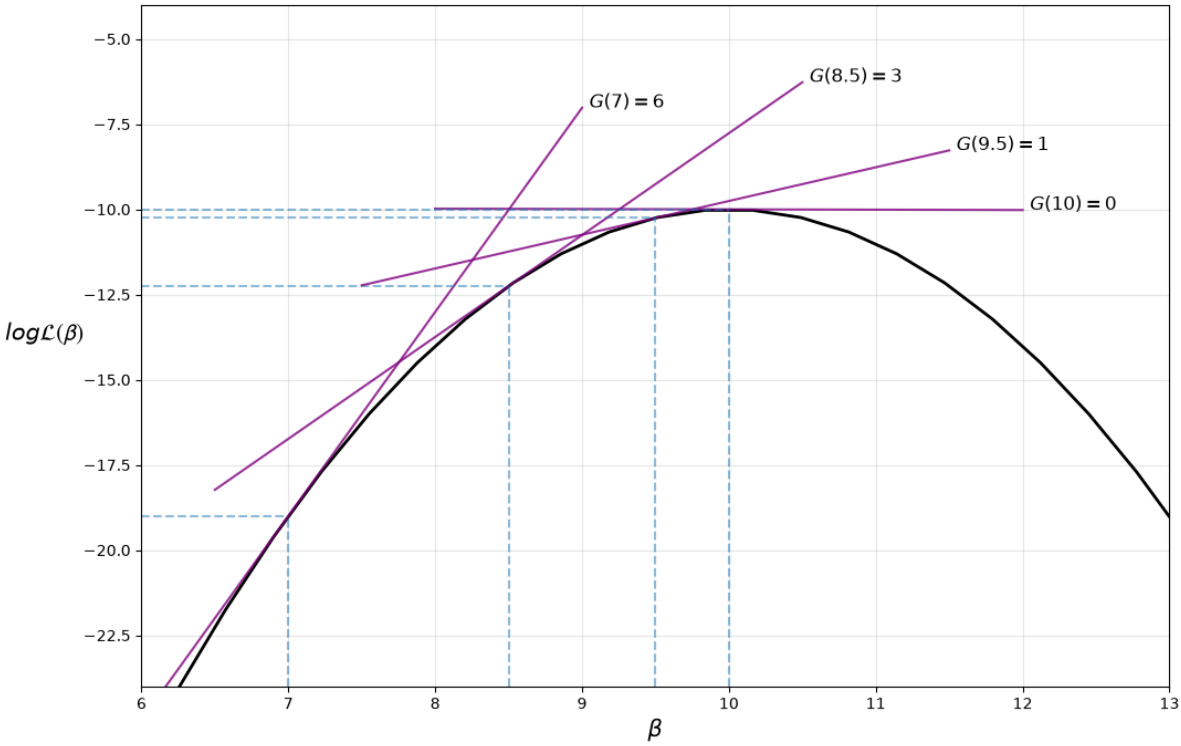
```
β = jnp.linspace(2, 18)
fig, ax = plt.subplots(figsize=(12, 8))
ax.plot(β, logL(β), lw=2, c="black")
```

(continues on next page)

(continued from previous page)

```
for β in [7, 8.5, 9.5, 10]:
    β_line = jnp.linspace(β - 2, β + 2)
    m, c = find_tangent(β)
    y = m * β_line + c
    ax.plot(β_line, y, "-", c="purple", alpha=0.8)
    ax.text(β + 2.05, y[-1], rf"$G({β}) = {abs(m):.0f}$", fontsize=12)
    ax.vlines(β, -24, logL(β), linestyles="--", alpha=0.5)
    ax.hlines(logL(β), 6, β, linestyles="--", alpha=0.5)

ax.set(ylim=(-24, -4), xlim=(6, 13))
ax.set_xlabel(r"$\beta$", fontsize=15)
ax.set_ylabel(
    r"$log \mathcal{L(\beta)}$", rotation=0, labelpad=25, fontsize=15
)
ax.grid(alpha=0.3)
plt.show()
```



Note that our implementation of the Newton-Raphson algorithm is rather basic — for more robust implementations see, for example, scipy.optimize.

145.6 Maximum likelihood estimation with `statsmodels`

Now that we know what's going on under the hood, we can apply MLE to an interesting application.

We'll use the Poisson regression model in `statsmodels` to obtain a richer output with standard errors, test values, and more.

`statsmodels` uses the same algorithm as above to find the maximum likelihood estimates.

Before we begin, let's re-estimate our simple model with `statsmodels` to confirm we obtain the same coefficients and log-likelihood value.

Now, as `statsmodels` accepts only NumPy arrays, we can use `np.array` method to convert them to NumPy arrays.

```
X = jnp.array([[1, 2, 5], [1, 1, 3], [1, 4, 2], [1, 5, 2], [1, 3, 1]])

y = jnp.array([1, 0, 1, 1, 0])

y_numpy = np.array(y)
X_numpy = np.array(X)
stats_poisson = Poisson(y_numpy, X_numpy).fit()
print(stats_poisson.summary())
```

```
Optimization terminated successfully.
         Current function value: 0.675671
         Iterations 7
                          Poisson Regression Results
==============================================================================
Dep. Variable:                      y   No. Observations:                    5
Model:                        Poisson   Df Residuals:                        2
Method:                           MLE   Df Model:                            2
Date:                Mon, 02 Nov 2026   Pseudo R-squ.:                  0.2546
Time:                        04:10:37   Log-Likelihood:                -3.3784
converged:                       True   LL-Null:                       -4.5325
Covariance Type:            nonrobust   LLR p-value:                    0.3153
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
const         -6.0785      5.279     -1.151      0.250     -16.425       4.268
x1             0.9334      0.829      1.126      0.260      -0.691       2.558
x2             0.8433      0.798      1.057      0.291      -0.720       2.407
==============================================================================
```

Now let's replicate results from Daniel Treisman's paper, Russia's Billionaires, mentioned earlier in the lecture.

Treisman starts by estimating equation (145.1), where:

- y_i is $number\ of\ billionaires_i$
- x_{i1} is $\log GDP\ per\ capita_i$
- x_{i2} is $\log population_i$
- x_{i3} is $years\ in\ GATT_i$ – years membership in GATT and WTO (to proxy access to international markets)

The paper only considers the year 2008 for estimation.

We will set up our variables for estimation like so (you should have the data assigned to `df` from earlier in the lecture)

```
# Keep only year 2008
df = df[df["year"] == 2008]

# Add a constant
df["const"] = 1

# Variable sets
reg1 = ["const", "lngdppc", "lnpop", "gattwto08"]
reg2 = [
    "const",
    "lngdppc",
    "lnpop",
    "gattwto08",
    "lnmcap08",
    "rintr",
    "topint08",
]
reg3 = [
    "const",
    "lngdppc",
    "lnpop",
    "gattwto08",
    "lnmcap08",
    "rintr",
    "topint08",
    "nrrents",
    "roflaw",
]
```

Then we can use the `Poisson` function from `statsmodels` to fit the model.

We'll use robust standard errors as in the author's paper

```
# Specify model
poisson_reg = Poisson(df[["numbil0"]], df[reg1], missing="drop").fit(
    cov_type="HC0"
)
print(poisson_reg.summary())
```

```
Optimization terminated successfully.
         Current function value: 2.226090
         Iterations 9
                          Poisson Regression Results
==============================================================================
Dep. Variable:                numbil0   No. Observations:                  197
Model:                        Poisson   Df Residuals:                      193
Method:                           MLE   Df Model:                            3
Date:                Mon, 02 Nov 2026   Pseudo R-squ.:                  0.8574
Time:                        04:10:37   Log-Likelihood:                -438.54
converged:                       True   LL-Null:                       -3074.7
Covariance Type:                 HC0   LLR p-value:                     0.000
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
const        -29.0495      2.578    -11.268      0.000     -34.103     -23.997
lngdppc        1.0839      0.138      7.834      0.000       0.813       1.355
lnpop          1.1714      0.097     12.024      0.000       0.980       1.362
```

(continues on next page)

(continued from previous page)

```
gattwto08      0.0060      0.007      0.868      0.386      -0.008       0.019
==============================================================================
```

Success! The algorithm was able to achieve convergence in 9 iterations.

Our output indicates that GDP per capita, population, and years of membership in the General Agreement on Tariffs and Trade (GATT) are positively related to the number of billionaires a country has, as expected.

Let's also estimate the author's more full-featured models and display them in a single table

```
regs = [reg1, reg2, reg3]
reg_names = ["Model 1", "Model 2", "Model 3"]
info_dict = {
    "Pseudo R-squared": lambda x: f"{x.prsquared:.2f}",
    "No. observations": lambda x: f"{int(x.nobs):d}",
}
regressor_order = [
    "const",
    "lngdppc",
    "lnpop",
    "gattwto08",
    "lnmcap08",
    "rintr",
    "topint08",
    "nrrents",
    "roflaw",
]
results = []

for reg in regs:
    result = Poisson(df[["numbil0"]], df[reg], missing="drop").fit(
        cov_type="HC0", maxiter=100, disp=0
    )
    results.append(result)

results_table = summary_col(
    results=results,
    float_format="%0.3f",
    stars=True,
    model_names=reg_names,
    info_dict=info_dict,
    regressor_order=regressor_order,
)
results_table.add_title(
    "Table 1 - Explaining the Number of Billionaires \
                        in 2008"
)
print(results_table)
```

```
Table 1 - Explaining the Number of Billionaires                        in 2008
================================================
                  Model 1    Model 2    Model 3
------------------------------------------------
const            -29.050*** -19.444*** -20.858***
                 (2.578)    (4.820)    (4.255)
lngdppc          1.084***   0.717***   0.737***
                 (0.138)    (0.244)    (0.233)
```

(continues on next page)

(continued from previous page)

```
lnpop             1.171***   0.806***   0.929***
                  (0.097)    (0.213)    (0.195)
gattwto08         0.006      0.007      0.004
                  (0.007)    (0.006)    (0.006)
lnmcap08                     0.399**    0.286*
                             (0.172)    (0.167)
rintr                        -0.010     -0.009
                             (0.010)    (0.010)
topint08                     -0.051***  -0.058***
                             (0.011)    (0.012)
nrrents                                 -0.005
                                        (0.010)
roflaw                                  0.203
                                        (0.372)
No. observations 197         131        131
Pseudo R-squared 0.86        0.90       0.90
================================================
Standard errors in parentheses.
* p<.1, ** p<.05, ***p<.01
```

The output suggests that the frequency of billionaires is positively correlated with GDP per capita, population size, stock market capitalization, and negatively correlated with top marginal income tax rate.

To analyze our results by country, we can plot the difference between the predicted and actual values, then sort from highest to lowest and plot the first 15

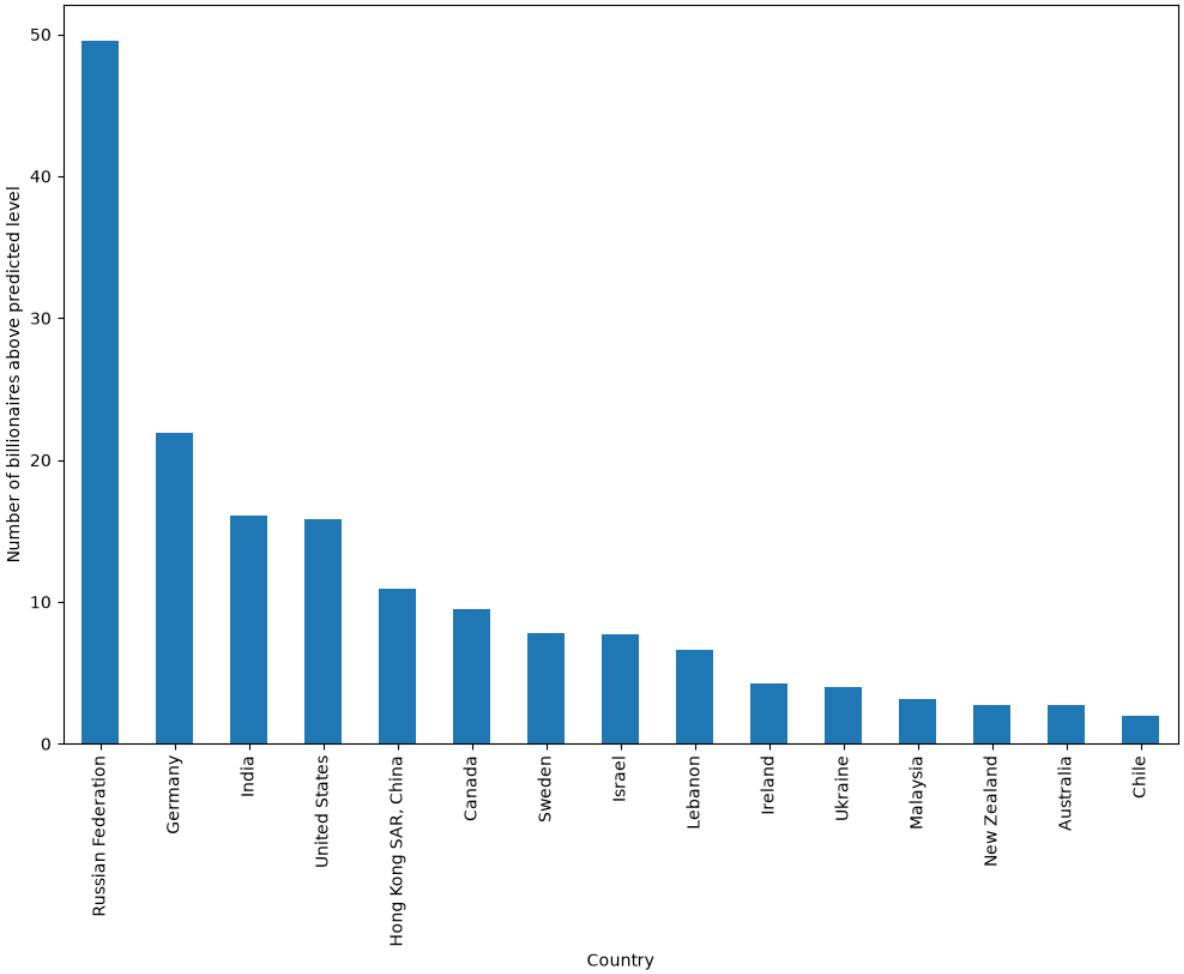
```
data = [
    "const",
    "lngdppc",
    "lnpop",
    "gattwto08",
    "lnmcap08",
    "rintr",
    "topint08",
    "nrrents",
    "roflaw",
    "numbil0",
    "country",
]
results_df = df[data].dropna()

# Use last model (model 3)
results_df["prediction"] = results[-1].predict()

# Calculate difference
results_df["difference"] = results_df["numbil0"] - results_df["prediction"]

# Sort in descending order
results_df.sort_values("difference", ascending=False, inplace=True)

# Plot the first 15 data points
results_df[:15].plot(
    "country", "difference", kind="bar", figsize=(12, 8), legend=False
)
plt.ylabel("Number of billionaires above predicted level")
plt.xlabel("Country")
plt.show()
```



As we can see, Russia has by far the highest number of billionaires in excess of what is predicted by the model (around 50 more than expected).

Treisman uses this empirical result to discuss possible reasons for Russia's excess of billionaires, including the origination of wealth in Russia, the political climate, and the history of privatization in the years after the USSR.

145.7 Summary

In this lecture, we used Maximum Likelihood Estimation to estimate the parameters of a Poisson model.

`statsmodels` contains other built-in likelihood models such as Probit and Logit.

For further flexibility, `statsmodels` provides a way to specify the distribution manually using the `GenericLikelihoodModel` class - an example notebook can be found here.

145.8 Exercises

Exercise 145.8.1

Suppose we wanted to estimate the probability of an event y_i occurring, given some observations.

We could use a probit regression model, where the pmf of y_i is

$$f(y_i; \beta) = \mu_i^{y_i}(1-\mu_i)^{1-y_i}, \quad y_i = 0, 1$$
$$\text{where} \quad \mu_i = \Phi(\mathbf{x}_i'\beta)$$

Φ represents the **cumulative normal distribution** and constrains the predicted y_i to be between 0 and 1 (as required for a probability).

β is a vector of coefficients.

Following the example in the lecture, write a class to represent the Probit model.

To begin, find the log-likelihood function and derive the gradient and Hessian.

The `jax.scipy.stats` module `norm` contains the functions needed to compute the cdf and pdf of the normal distribution.

Solution

The log-likelihood can be written as

$$\log \mathcal{L} = \sum_{i=1}^{n} \left[y_i \log \Phi(\mathbf{x}_i'\beta) + (1-y_i) \log(1 - \Phi(\mathbf{x}_i'\beta)) \right]$$

Using the **fundamental theorem of calculus**, the derivative of a cumulative probability distribution is its marginal distribution

$$\frac{\partial}{\partial s}\Phi(s) = \phi(s)$$

where ϕ is the marginal normal distribution.

The gradient vector of the Probit model is

$$\frac{\partial \log \mathcal{L}}{\partial \beta} = \sum_{i=1}^{n} \left[y_i \frac{\phi(\mathbf{x}_i'\beta)}{\Phi(\mathbf{x}_i'\beta)} - (1-y_i)\frac{\phi(\mathbf{x}_i'\beta)}{1-\Phi(\mathbf{x}_i'\beta)} \right] \mathbf{x}_i$$

The Hessian of the Probit model is

$$\frac{\partial^2 \log \mathcal{L}}{\partial \beta \partial \beta'} = -\sum_{i=1}^{n} \phi(\mathbf{x}_i'\beta) \left[y_i \frac{\phi(\mathbf{x}_i'\beta) + \mathbf{x}_i'\beta\Phi(\mathbf{x}_i'\beta)}{[\Phi(\mathbf{x}_i'\beta)]^2} + (1-y_i)\frac{\phi(\mathbf{x}_i'\beta) - \mathbf{x}_i'\beta(1-\Phi(\mathbf{x}_i'\beta))}{[1-\Phi(\mathbf{x}_i'\beta)]^2} \right] \mathbf{x}_i\mathbf{x}_i'$$

Using these results, we can write a class for the Probit model as follows

```
class ProbitRegression(NamedTuple):
    X: jnp.ndarray
    y: jnp.ndarray
```

```
@jax.jit
def logL(β, model):
    y = model.y
    μ = norm.cdf(model.X @ β.T)
    return y @ jnp.log(μ) + (1 - y) @ jnp.log(1 - μ)
```

```
G_logL = jax.grad(logL)
H_logL = jax.jacfwd(G_logL)
```

Exercise 145.8.2

Use the following dataset and initial values of β to estimate the MLE with the Newton-Raphson algorithm developed earlier in the lecture

$$\mathbf{X} = \begin{bmatrix} 1 & 2 & 4 \\ 1 & 1 & 1 \\ 1 & 4 & 3 \\ 1 & 5 & 6 \\ 1 & 3 & 5 \end{bmatrix} \quad y = \begin{bmatrix} 1 \\ 0 \\ 1 \\ 1 \\ 0 \end{bmatrix} \quad \beta_{(0)} = \begin{bmatrix} 0.1 \\ 0.1 \\ 0.1 \end{bmatrix}$$

Verify your results with `statsmodels` - you can import the Probit function with the following import statement

```
from statsmodels.discrete.discrete_model import Probit
```

Note that the simple Newton-Raphson algorithm developed in this lecture is very sensitive to initial values, and therefore you may fail to achieve convergence with different starting values.

Solution

Here is one solution

```
X = jnp.array([[1, 2, 4], [1, 1, 1], [1, 4, 3], [1, 5, 6], [1, 3, 5]])

y = jnp.array([1, 0, 1, 1, 0])

# Take a guess at initial βs
β = jnp.array([0.1, 0.1, 0.1])

# Create a model of Probit regression
prob = ProbitRegression(y=y, X=X)

# Run Newton-Raphson algorithm
newton_raphson(prob, β)
```

```
Iteration_k  Log-likelihood  θ
--------------------------------------------------------------------------------
 ↪---------
0             -2.3796887      ['-1.34', '0.775', '-0.157']
1             -2.3687525      ['-1.53', '0.775', '-0.0981']
2             -2.3687296      ['-1.55', '0.778', '-0.0971']
3             -2.3687294      ['-1.55', '0.778', '-0.0971']
Number of iterations: 4
β_hat = [-1.5462586   0.7777895  -0.09709753]
```

```
Array([-1.5462586 ,  0.7777895 , -0.09709753], dtype=float32)
```

```
# Use statsmodels to verify results
y_numpy = np.array(y)
X_numpy = np.array(X)
print(Probit(y_numpy, X_numpy).fit().summary())
```

```
Optimization terminated successfully.
         Current function value: 0.473746
         Iterations 6
                          Probit Regression Results
==============================================================================
Dep. Variable:                      y   No. Observations:                    5
Model:                         Probit   Df Residuals:                        2
Method:                           MLE   Df Model:                            2
Date:                Mon, 02 Nov 2026   Pseudo R-squ.:                  0.2961
Time:                        04:10:38   Log-Likelihood:                -2.3687
converged:                       True   LL-Null:                       -3.3651
Covariance Type:            nonrobust   LLR p-value:                    0.3692
==============================================================================
                 coef    std err          z      P>|z|      [0.025      0.975]
------------------------------------------------------------------------------
const         -1.5463      1.866     -0.829      0.407      -5.204       2.111
x1             0.7778      0.788      0.986      0.324      -0.768       2.323
x2            -0.0971      0.590     -0.165      0.869      -1.254       1.060
==============================================================================
```

CHAPTER

SIX

A LINEAR MODEL OF UNEMPLOYMENT

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

In addition to what's in Anaconda, this lecture needs the following libraries.

We first install `numpyro` and `jax`:

```
!pip install numpyro jax
```

We also install `pandas_datareader`, which we use to download data from FRED:

```
!pip install pandas_datareader
```

146.1 Overview

This lecture applies Bayesian estimation to a linear AR(1) model of the US unemployment rate.

We met the machinery in *Posterior Distributions for AR(1) Parameters*, but there the data were simulated and the focus was theoretical.

Here we work carefully through real data, organizing the lecture around a question that divided macroeconomists in the 1980s and 1990s: is unemployment a *random walk*?

In the process, we will observe an asymmetry the model cannot capture.

This will motivate a sequel lecture, titled *Unemployment Dynamics with Asymmetric Shocks*.

As in *Posterior Distributions for AR(1) Parameters* and *Non-Conjugate Priors*, we estimate by sampling posteriors with the NUTS sampler in NumPyro.

(See the *introduction to NUTS* in *Non-Conjugate Priors* for a brief account of how it works.)

Let's start with some imports.

```
import numpy as np
import matplotlib.pyplot as plt
import datetime as dt
from pandas_datareader import data as web

import jax.numpy as jnp
from jax import random
import numpyro
import numpyro.distributions as dist
from numpyro.infer import MCMC, NUTS
```

146.2 The natural rate versus hysteresis

In the early 1980s, Nelson and Plosser [1982] launched the unit-root literature.

They examined 14 macroeconomic time series and found that they could reject the random walk for only one of them.

Roughly speaking, this meant that, for most macroeconomic series, the effects of shocks looked permanent rather than transitory.

For unemployment, that view found expression in the **hysteresis hypothesis** of [Blanchard and Summers, 1986].

This hypothesis states that a shock to unemployment can be more or less permanent, because spells of joblessness erode skills and attachment to the labor force.

In contrast, the **natural rate hypothesis** of [Friedman, 1968] holds that unemployment fluctuates around a stable equilibrium rate — and hence shocks are transitory rather than permanent.

The debate matters because the two views imply different policies.

In particular, if shocks are permanent, a deep recession leaves a lasting scar, encouraging remedial action.

Here we reexamine the question with Bayesian estimation.

Note

There is an irony in the history described above.

While Nelson and Plosser [1982] fueled the unit root debate, and hence the hysteresis hypothesis, the one macroeconomic series they rejected the unit root for was the unemployment rate.

146.3 The data

We use the US civilian unemployment rate, series `UNRATE` from FRED, monthly and seasonally adjusted.

```
start, end = dt.datetime(1948, 1, 1), dt.datetime(2024, 12, 31)
unrate = web.DataReader("UNRATE", "fred", start, end)["UNRATE"]
```

The COVID-19 spike of 2020 is an extreme outlier driven by events our models know nothing about, so we drop it.

```
pre_covid = unrate[unrate.index < "2020-01-01"]
u_monthly = pre_covid.to_numpy()
print(f"{len(u_monthly)} monthly observations")
```

```
864 monthly observations
```

Here is the monthly series.

```
fig, ax = plt.subplots()
ax.plot(pre_covid.index, u_monthly, lw=2)
ax.set_xlabel('year')
ax.set_ylabel('unemployment rate (%)')
plt.show()
```

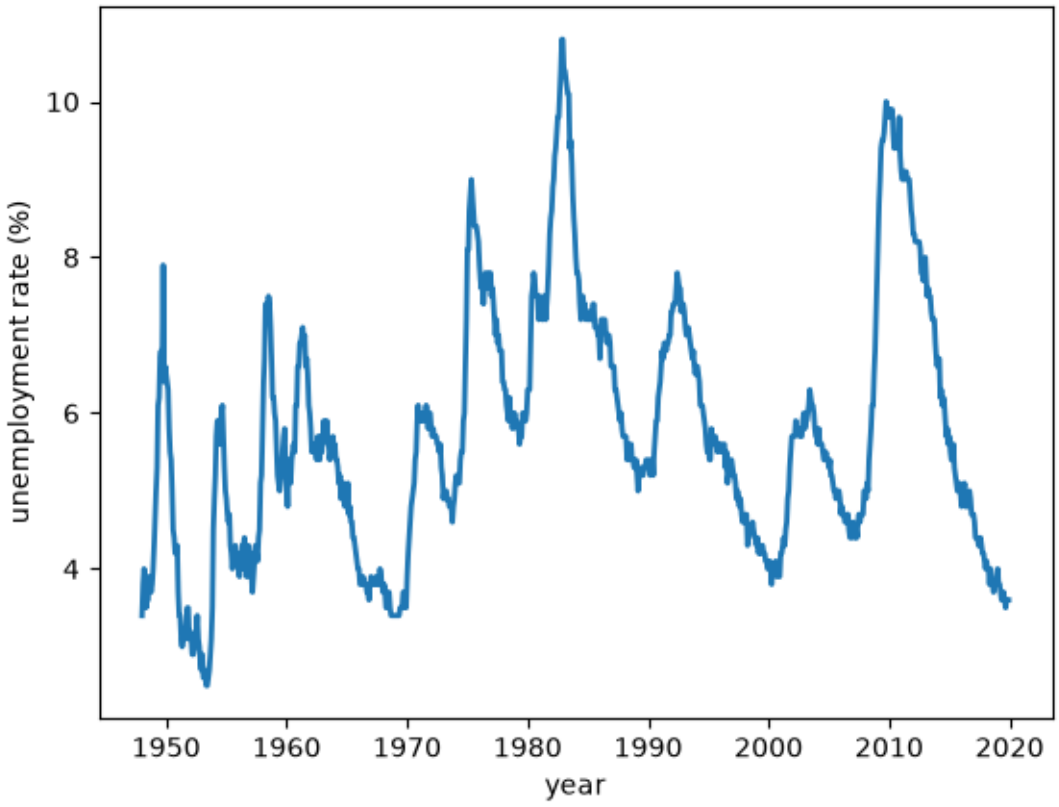


Fig. 146.1: US monthly unemployment rate, 1948–2019

As Fig. 146.1 shows, the rate rises sharply in recessions and drifts down in recoveries, but it always stays inside a band — roughly 3% to 11% over the whole post-war period.

(Keep that band in mind, since it connects back to the unit root debate, as we discuss below.)

146.4 A linear model of unemployment

Let's now set up and estimate our model, using monthly data.

146.4.1 The model

We let unemployment be pulled back toward a normal level $\bar{u}$:

$$u_{t+1} = \bar{u} + \phi\,(u_t - \bar{u}) + \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim N(0, \sigma^2), \tag{146.1}$$

with $0 \leq \phi < 1$.

This is a linear AR(1) model, written so that $\bar{u}$ is the level the series reverts to and ϕ measures persistence.

The closer ϕ is to one, the more slowly the series returns to $\bar{u}$ after a shock; the random walk is the limiting case $\phi = 1$.

146.4.2 Priors

We treat $\bar{u}$, ϕ and σ as unknown and place weakly informative priors on them.

We give ϕ a uniform prior on $[0, 1)$.

The upper endpoint is excluded deliberately, since doing so confines us to the *stationary* region.

This is necessary, as we'll see below.

At the same time, the prior still lets ϕ approach one as closely as the data demand.

We center $\bar{u}$ on a plausible natural rate with a fairly wide normal prior, and give the shock scale σ a half-normal prior.

We write the model as a NumPyro function: each `numpyro.sample` introduces a random variable, and the keyword `obs=` ties the last one to the data, supplying the likelihood.

```
def linear_model(u):
    ubar = numpyro.sample("ubar",  dist.Normal(5.5, 2.0))      # Natural rate prior
    φ    = numpyro.sample("phi",   dist.Uniform(0.0, 1.0))     # Persistence prior
    σ    = numpyro.sample("sigma", dist.HalfNormal(1.0))       # Volatility prior
    μ = ubar + φ * (u[:-1] - ubar)
    numpyro.sample("u_obs", dist.Normal(μ, σ), obs=u[1:])
```

The vector `μ` holds the conditional means $\bar{u} + \phi(u_t - \bar{u})$, and `obs=u[1:]` says that each next value is drawn from $N(\mu_t, \sigma^2)$.

This one statement encodes the whole likelihood. (See *Non-Conjugate Priors* for more on writing NumPyro models.)

146.4.3 Estimation

We sample the posterior with NUTS, running four chains so that we can check convergence.

We use `chain_method="vectorized"`, which evaluates all chains together on a single device, so the same code runs unchanged on a CPU or a GPU.

```
def run_nuts(model, data, seed=0, num_warmup=1000, num_samples=2000, num_chains=4):
    "Sample a NumPyro model with the NUTS sampler."
    mcmc = MCMC(NUTS(model),
                num_warmup=num_warmup, num_samples=num_samples,
```

(continues on next page)

(continued from previous page)

```
                num_chains=num_chains, chain_method="vectorized",
                progress_bar=False)
    mcmc.run(random.key(seed), jnp.asarray(data))
    return mcmc
```

We fit the model to the monthly data.

```
mcmc_monthly = run_nuts(linear_model, u_monthly)
```

Now we inspect the output.

```
mcmc_monthly.print_summary()
```

```
                mean       std    median      5.0%     95.0%     n_eff     r_hat
       phi       0.99      0.00      0.99      0.99      1.00   2867.64      1.00
     sigma       0.21      0.01      0.21      0.20      0.22   4182.48      1.00
      ubar       5.64      1.16      5.66      3.81      7.62   3181.27      1.00

Number of divergences: 0
```

Each row summarizes the posterior for one parameter.

The `mean`, `median`, and `5.0%`/`95.0%` columns give the posterior mean, median, and a 90% credible interval; `std` is the posterior standard deviation.

The last two columns are convergence diagnostics: `n_eff` is the effective number of independent draws, and `r_hat` compares variation within and across chains — a value very close to 1.0 means the chains agree and the sampler has converged.

Here `r_hat` is essentially one and `n_eff` is large, so we can trust the draws.

The number that matters for us is the posterior for ϕ: its mass crowds right up against one.

In particular, the mean and median are very close to one, while the standard deviation is very small.

In other words, at a monthly frequency, US unemployment is *almost* a random walk.

This is the hysteresis boundary.

Thus, we find that the natural-rate and hysteresis views are nearly indistinguishable in the monthly data — the estimate is consistent with both.

This is why the unit-root tests of the era struggled to settle the debate [Røed, 1997].

146.5 A random walk would wander off

While the estimation above seems to give reasonable support to the unit root hypothesis, there is a good reason to view it as false.

To see the argument, suppose that unemployment really is a pure random walk, with $\phi = 1$:

$$u_{t+1} = u_t + \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim N(0, \sigma^2).$$

Then $u_t = u_0 + \sum_{s=1}^{t} \varepsilon_s$, so its variance grows without bound: $\text{Var}(u_t) = t\sigma^2$.

The distribution spreads out forever, and eventually probability mass leaves *every* bounded interval.

We can see this by simulating many random-walk paths, using the observed one-month changes to set the shock size, and watching them fan out.

```
rng = np.random.default_rng(0)
T = len(u_monthly)
σ_rw = np.diff(u_monthly).std()
paths = u_monthly[0] + np.cumsum(rng.normal(0, σ_rw, size=(400, T)), axis=1)

u_min, u_max = u_monthly.min(), u_monthly.max()

fig, ax = plt.subplots()
ax.plot(paths[:60].T, color='C0', lw=0.5, alpha=0.3)
ax.axhspan(u_min, u_max, color='C1', alpha=0.15, label='observed range')
ax.axhline(u_min, color='C1', ls='--', lw=1.5)
ax.axhline(u_max, color='C1', ls='--', lw=1.5)
ax.set_xlabel('months since 1948')
ax.set_ylabel('unemployment rate (%)')
ax.legend()
plt.show()
```

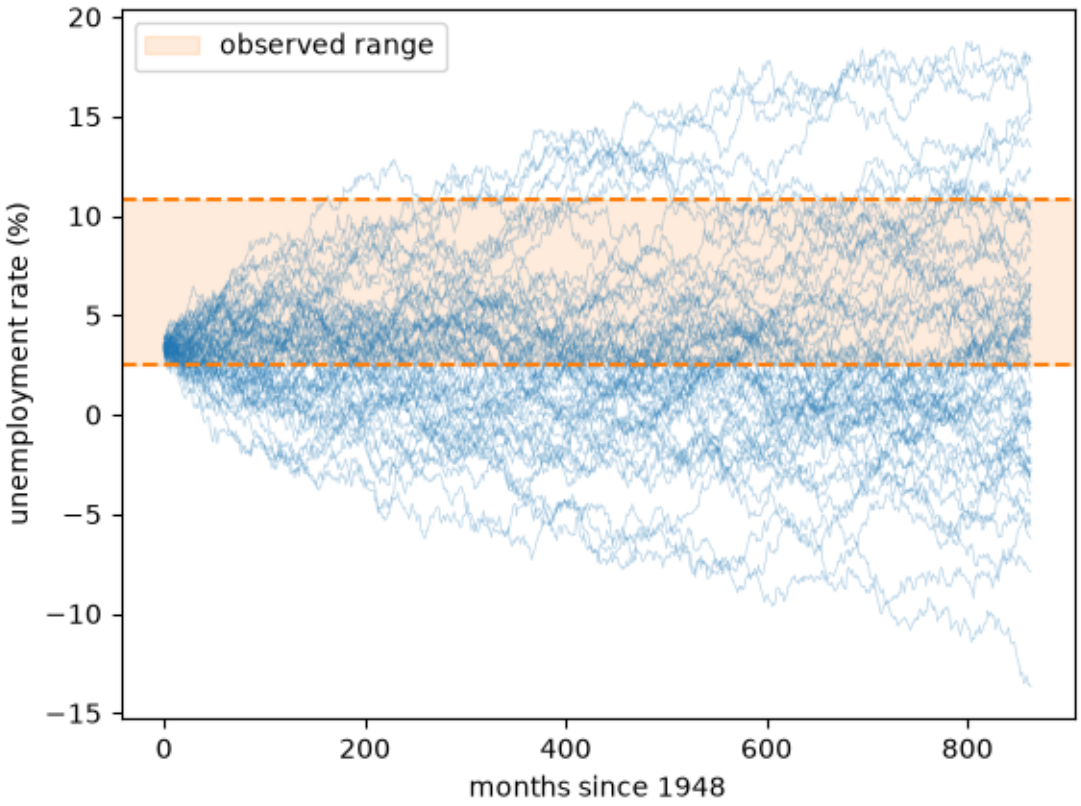


Fig. 146.2: Random walks leave the observed range

In Fig. 146.2 the dashed lines mark the lowest and highest unemployment rates ever seen in the data, and the shaded region is the band between them.

The simulated paths fan out like $\sqrt{t}$ and quickly spread far beyond this band, including into negative rates.

A random walk has no anchor, but unemployment clearly does — it has stayed inside a narrow band for seventy years.

So we can already rule out an exact random walk.

This will become even clearer when we examine annual data.

146.6 Monthly versus annual

The monthly data leave ϕ pinned against one.

Does a different frequency allow us to see something more?

We form an annual series from end-of-year values and fit the same model to it.

```
u_annual = pre_covid.resample("YE").last().to_numpy()
print(f"{len(u_annual)} annual observations")
```

```
72 annual observations
```

We fit the same model to this shorter series.

```
mcmc_annual = run_nuts(linear_model, u_annual)
```

Now we compare the posterior for ϕ across the two frequencies.

```
φ_m = np.asarray(mcmc_monthly.get_samples()["phi"])
φ_a = np.asarray(mcmc_annual.get_samples()["phi"])

fig, ax = plt.subplots()
ax.hist(φ_m, bins=50, density=True, alpha=0.6, label='monthly')
ax.hist(φ_a, bins=50, density=True, alpha=0.6, label='annual')
ax.set_xlabel('$\\phi$')
ax.legend()
plt.show()
```

Fig. 146.3 illustrates our main finding: the annual posterior for ϕ sits well below one, with clear reversion, while the monthly posterior pushes up against the boundary.

This is not a contradiction.

If the monthly persistence is ϕ, then for end-of-year values the persistence is about ϕ^{12}, and raising a number near one to the twelfth power pulls it appreciably below one — in line with our annual estimate.

146.7 What the model misses

Looking again at Fig. 146.1, we notice that unemployment jumps up quickly in recessions and drifts down slowly in recoveries.

This is an *asymmetry* that our current model cannot replicate.

To see why, we look closely at the one part of the model that is random — the shocks.

We proceed in three steps: state what the model assumes about the shocks, recover them from the data, and compare the two.

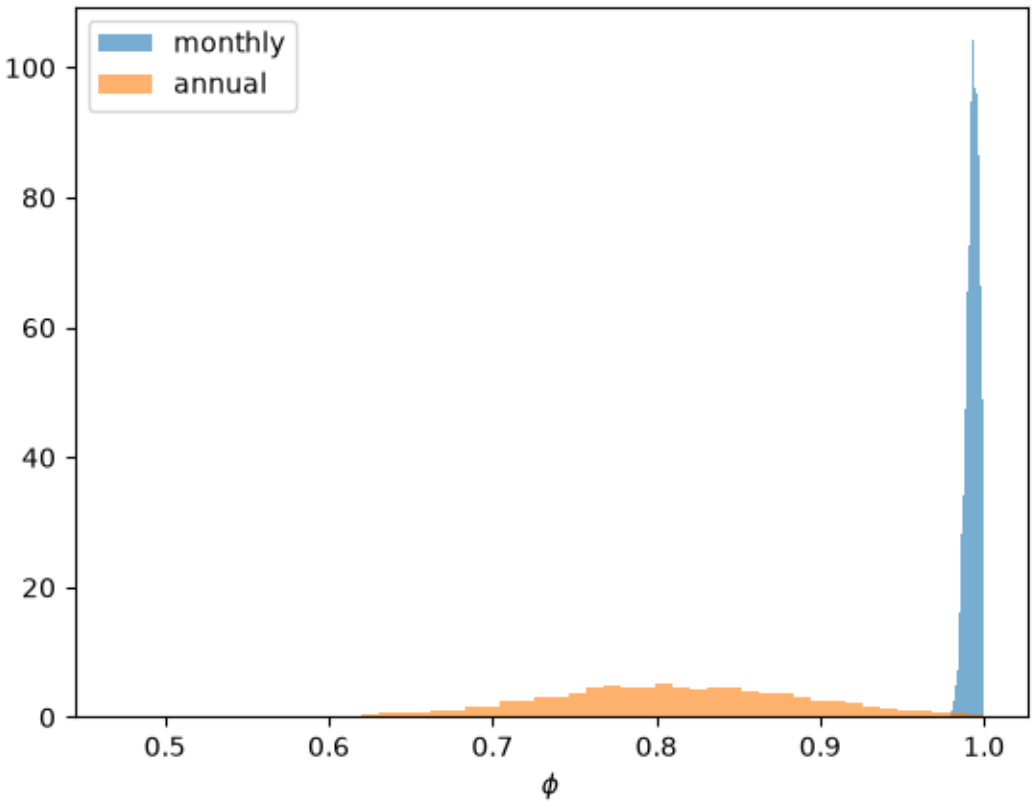


Fig. 146.3: Posterior for the persistence parameter

146.7.1 What the model assumes about the shocks

Given last period's rate u_t and the parameters, the next rate is a deterministic conditional mean plus a shock:

$$u_{t+1} = \underbrace{\bar{u} + \phi\,(u_t - \bar{u})}_{\text{conditional mean}} + \varepsilon_{t+1}, \qquad \varepsilon_{t+1} \sim N(0, \sigma^2).$$

Rearranging, the shock is just the gap between what happened and what the model expected:

$$\varepsilon_{t+1} = u_{t+1} - (\bar{u} + \phi\,(u_t - \bar{u})).$$

The model makes a strong and testable claim about these shocks: they are independent draws from a *symmetric* normal distribution.

If that claim holds, the shocks we recover from the data should look like a bell curve.

If it fails, the way it fails will tell us what the model is missing.

146.7.2 Recovering the shocks

We cannot read the shocks off directly, because we do not know the parameters $\bar{u}$ and ϕ.

So we estimate them, plugging a single representative value in for each.

We use the posterior medians — a reasonable choice for a quick diagnostic.

```
med = {k: np.median(np.asarray(mcmc_monthly.get_samples()[k]))
       for k in ("ubar", "phi", "sigma")}
resid = u_monthly[1:] - (med["ubar"] + med["phi"] * (u_monthly[:-1] - med["ubar"]))
```

The `resid` array holds our estimated shocks, one for each month-to-month transition — the model's **residuals**.

The slicing is what lines up each month with the one before it.

Each entry of `resid` is $u_{t+1} - (\hat{\bar{u}} + \hat{\phi}\,(u_t - \hat{\bar{u}}))$, the model's one-step-ahead forecast error, with hats denoting the median estimates.

(Because our estimate $\hat{\phi}$ is so close to one, this is nearly just the monthly change $u_{t+1} - u_t$.)

146.7.3 Comparing with the Gaussian

Now we ask whether these residuals look Gaussian.

We overlay a normal density whose standard deviation we set equal to that of the residuals themselves.

This is deliberate: the residuals already have mean near zero, so matching the variance makes the mean and the spread agree.

Anything left over is then a difference in *shape* — which is what we want to isolate.

One way to measure the shape is the **skewness**, the third standardized moment:

$$\text{skew} = \frac{\frac{1}{n}\sum_i (\varepsilon_i - \bar{\varepsilon})^3}{\left(\frac{1}{n}\sum_i (\varepsilon_i - \bar{\varepsilon})^2\right)^{3/2}}.$$

This measure is zero for any symmetric distribution and positive when the right tail is the longer one.

We now plot the residuals and compute their skewness:

```
def skewness(x):
    x = x - x.mean()
    return (x**3).mean() / x.std()**3

fig, ax = plt.subplots()
ax.hist(resid, bins=60, density=True, alpha=0.6, label='residuals')
grid = np.linspace(resid.min(), resid.max(), 200)
gauss = np.exp(-grid**2 / (2 * resid.std()**2)) / (resid.std() * np.sqrt(2 * np.pi))
ax.plot(grid, gauss, 'C1', lw=2, label='symmetric Gaussian')
ax.set_xlabel('one-month change not explained by the model')
ax.legend()
plt.show()

print(f"residual skewness = {skewness(resid):.2f}")
```

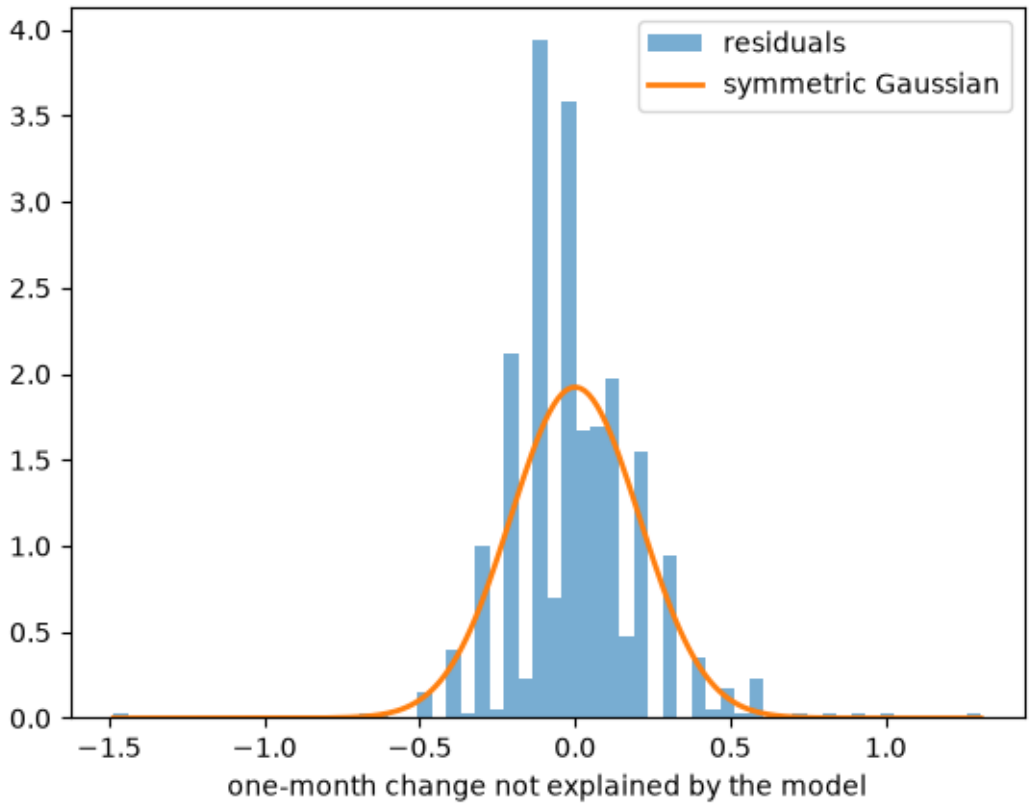


Fig. 146.4: Model residuals are right-skewed

```
residual skewness = 0.39
```

The residuals in Fig. 146.4 depart from the Gaussian in two ways.

They are *heavy-tailed* and *sharply peaked*: more tiny changes, and more large ones, than a bell curve allows.

Moreover, they are *right-skewed*, with the largest surprises on the upside (recessions).

The symmetric Gaussian (orange) can match neither feature: it treats upward and downward shocks as equally likely and has no room for the occasional very large jump.

So our model is bound to misread the data, seeing the rare jump up and the long gentle slide down as the *same* kind of

shock.

146.7.4 A note on the plug-in

A Bayesian purist would object that there is no single residual series here.

Every posterior draw of $(\bar{u}, \phi)$ implies its own, and we simply chose the medians.

That is fair, and the fully Bayesian version of this check — simulating whole datasets from the posterior and comparing a summary statistic — is what we do in *Unemployment Dynamics with Asymmetric Shocks*.

This plug-in check is a quick preview of that fuller test.

Capturing the asymmetry is the task of that next lecture, where we model the shocks themselves rather than the reversion curve.

146.8 Exercises

The lecture *Forecasting an AR(1) Process* forecasts an AR(1) process by simulating future paths, considering both a predictive distribution that **conditions on** fixed parameter values and one that **integrates over** their posterior uncertainty.

The following exercises apply these ideas to our fitted unemployment model.

Exercise 146.8.1

Using the fitted **annual** model, forecast unemployment over the next $H = 15$ years, starting from the last observed value, in two ways:

1. **plug-in**: hold the parameters at their posterior medians and simulate many future paths;
2. **extended**: for each future path, draw a fresh $(\bar{u}, \phi, \sigma)$ from the posterior.

Plot the 90% predictive band for each on the same axes and compare.

Which band is wider, and why?

Solution

We use the annual posterior draws and simulate forward from the last observation.

We work with the annual model because its persistence ϕ is far less certain than at the monthly frequency, so parameter uncertainty has more to say.

```
post = mcmc_annual.get_samples()
ubar_s = np.asarray(post["ubar"])
ψ_s = np.asarray(post["phi"])
σ_s = np.asarray(post["sigma"])

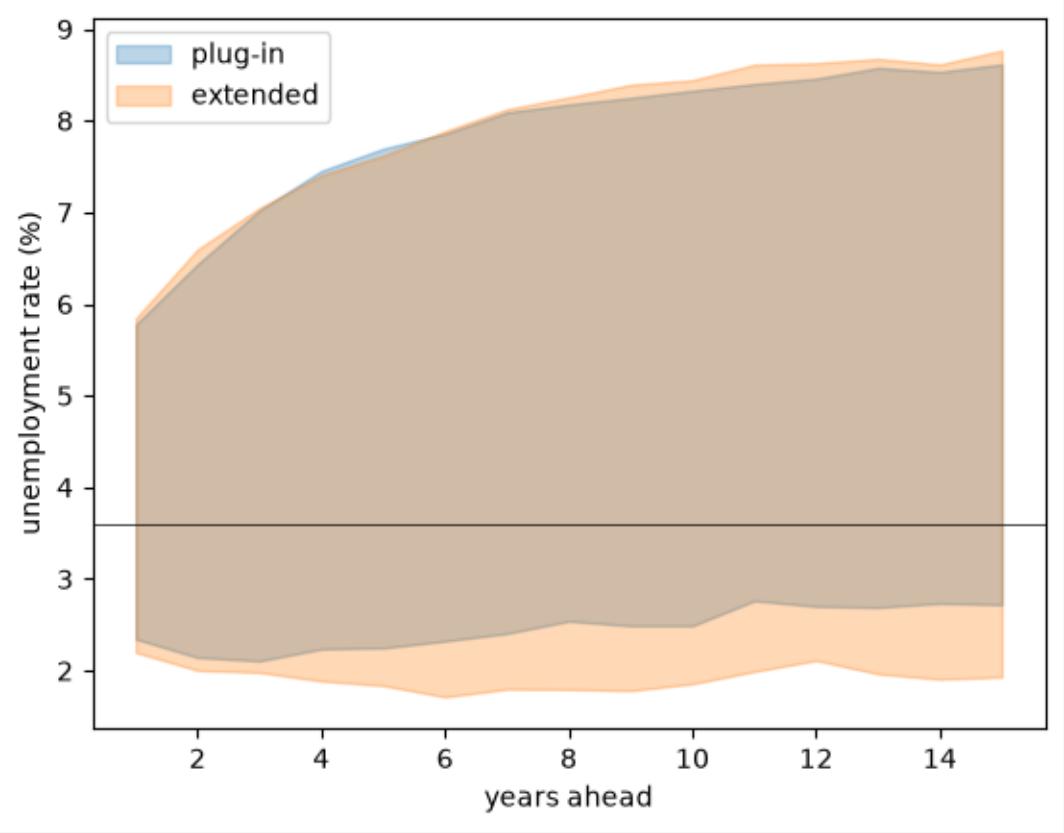
def sim_future(u_last, ubar, ψ, σ, H, rng):
    "Simulate H steps of the linear model forward from u_last."
    u = np.empty(H)
    prev = u_last
    for h in range(H):
        prev = ubar + ψ * (prev - ubar) + rng.normal(0, σ)
        u[h] = prev
    return u
```

```
H, N = 15, 2000
u_last = u_annual[-1]
rng = np.random.default_rng(0)

# plug-in: parameters fixed at posterior medians
ub0, φ0, σ0 = np.median(ubar_s), np.median(φ_s), np.median(σ_s)
plug = np.array([sim_future(u_last, ub0, φ0, σ0, H, rng) for _ in range(N)])

# extended: a fresh posterior draw for each path
idx = rng.integers(0, len(φ_s), N)
ext = np.array([sim_future(u_last, ubar_s[i], φ_s[i], σ_s[i], H, rng)
                for i in idx])

fig, ax = plt.subplots()
horizon = np.arange(1, H + 1)
for data, c, lab in [(plug, 'C0', 'plug-in'), (ext, 'C1', 'extended')]:
    lo, hi = np.percentile(data, [5, 95], axis=0)
    ax.fill_between(horizon, lo, hi, alpha=0.3, color=c, label=lab)
ax.axhline(u_last, color='k', lw=0.5)
ax.set_xlabel('years ahead')
ax.set_ylabel('unemployment rate (%)')
ax.legend()
plt.show()
```



Both forecasts drift up from the low starting value toward the reversion level $\bar{u}$, with bands that widen as we look further ahead.

The extended band is clearly wider, because it adds uncertainty about the parameters on top of the uncertainty from

future shocks — and here the persistence ϕ and the level $\bar{u}$ are both genuinely uncertain.

Exercise 146.8.2

Following Wecker (see *Forecasting an AR(1) Process*), we can also form a predictive distribution for a **path statistic** — a nonlinear function of the whole future path.

Take the **maximum unemployment rate over the next** $H = 8$ **years** as the statistic: a simple measure of how bad the next several years might get.

Using the extended simulation (a posterior draw per path), compute this maximum for each path and plot its predictive distribution.

What is the posterior predictive probability that unemployment reaches at least 7% — recession territory — at some point over the next eight years?

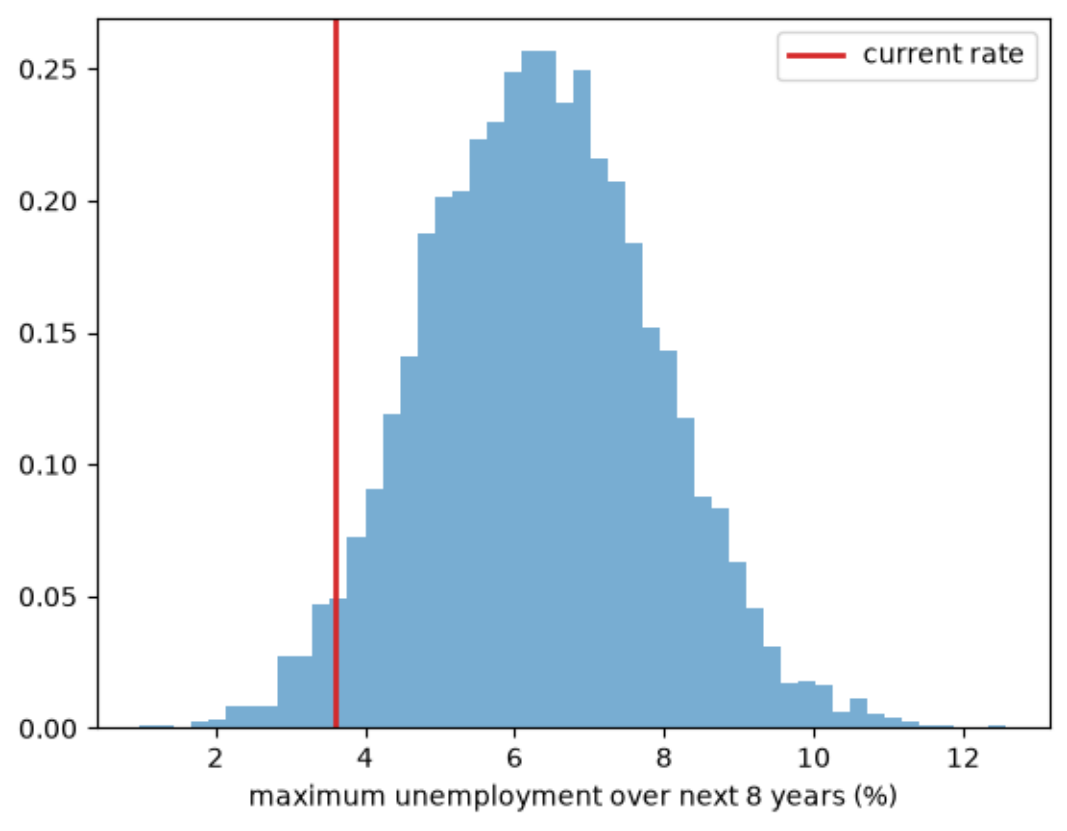
Solution

We reuse `sim_future` and the posterior draws from the previous exercise.

```
H2, M = 8, 5000
idx = rng.integers(0, len(ψ_s), M)
peak = np.array([sim_future(u_last, ubar_s[i], ψ_s[i], σ_s[i], H2, rng).max()
                 for i in idx])

fig, ax = plt.subplots()
ax.hist(peak, bins=50, density=True, alpha=0.6)
ax.axvline(u_last, color='C3', lw=2, label='current rate')
ax.set_xlabel('maximum unemployment over next 8 years (%)')
ax.legend()
plt.show()

prob = (peak > 7.0).mean()
print(f"P(unemployment reaches 7% within 8 years) = {prob:.2f}")
```




```
P(unemployment reaches 7% within 8 years) = 0.33
```

The predictive distribution summarizes the plausible “worst case” over the next several years, integrating over both shocks and parameter uncertainty.

Starting from a cyclical low, the model sees a substantial chance of a return to recession-level unemployment within the horizon.

CHAPTER

SEVEN

UNEMPLOYMENT DYNAMICS WITH ASYMMETRIC SHOCKS

GPU

This lecture was built using a machine with access to a GPU — although it will also run without one.

Google Colab has a free tier with GPUs that you can access as follows:

1. Click on the "play" icon top right
2. Select Colab
3. Set the runtime environment to include a GPU

This lecture is a sequel to *A Linear Model of Unemployment*.

There we fit a linear AR(1) model to the US unemployment rate and examined persistence.

One thing we found but didn't address is that the model's residuals were both *heavy-tailed and right-skewed*.

Unemployment jumps up sharply in recessions and drifts down slowly in recoveries.

A model with symmetric Gaussian shocks cannot reproduce these features.

In this lecture we allow the innovations to be right-skewed and, to some degree, heavy-tailed.

We also use this lecture to develop a very useful Bayesian technique: *model comparison* by leave-one-out cross-validation.

We use this to decide whether one model really predicts better than another.

Our plan is to

1. build a model with a linear mean but asymmetric, occasionally-large shocks,
2. estimate it on the annual data, and
3. compare it to the original Gaussian model with cross-validation.

In addition to what's in Anaconda, this lecture needs the following libraries.

```
!pip install numpyro jax pandas_datareader arviz
```

We'll use the following imports.

```
import numpy as np
import matplotlib.pyplot as plt
import datetime as dt
from pandas_datareader import data as web

import jax.numpy as jnp
```

(continues on next page)

(continued from previous page)

```
from jax import random
import numpyro
import numpyro.distributions as dist
from numpyro.infer import MCMC, NUTS, log_likelihood
```

147.1 The data

We use the same series as in *A Linear Model of Unemployment* — the US unemployment rate (`UNRATE` from FRED), excluding the COVID-19 spike — and work at an annual frequency.

```
start, end = dt.datetime(1948, 1, 1), dt.datetime(2024, 12, 31)
unrate = web.DataReader("UNRATE", "fred", start, end)["UNRATE"]
pre_covid = unrate[unrate.index < "2020-01-01"]
u_annual = pre_covid.resample("YE").last().to_numpy()
years = pre_covid.resample("YE").last().index.year
```

The shape we want to capture is plain in the annual series.

```
fig, ax = plt.subplots()
ax.plot(years, u_annual, lw=2, marker='o', ms=3)
ax.set_xlabel('year')
ax.set_ylabel('unemployment rate (%)')
plt.show()
```

In Fig. 147.1 each recession is a sharp rise followed by a long, gradual decline — fast up, slow down.

147.2 A model with asymmetric shocks

We keep the linear structure from *A Linear Model of Unemployment*:

$$u_{t+1} = \bar{u} + \rho\,(u_t - \bar{u}) + \eta_{t+1}, \qquad 0 \le \rho < 1, \tag{147.1}$$

Here, however, the innovation η_{t+1} is drawn from a *mixture of two normals*:

$$\eta_{t+1} \sim \begin{cases} N(0, \sigma_s^2) & \text{with probability } 1-p \quad \text{(a quiet year)}, \\ N(\mu_J, \sigma_J^2) & \text{with probability } p \quad \text{(a recession jump)}, \end{cases} \qquad \mu_J > 0.$$

The model has six parameters:

symbol	name	role
$\bar{u}$	floor	the low level the series reverts to
ρ	persistence	speed of reversion (as in the linear lecture)
p	jump probability	how often a large shock arrives
μ_J	jump size	the average upward kick of a recession
σ_s, σ_J	shock spreads	the quiet and jumpy volatilities

The mechanism leads to a sawtooth dynamic for the unemployment rate.

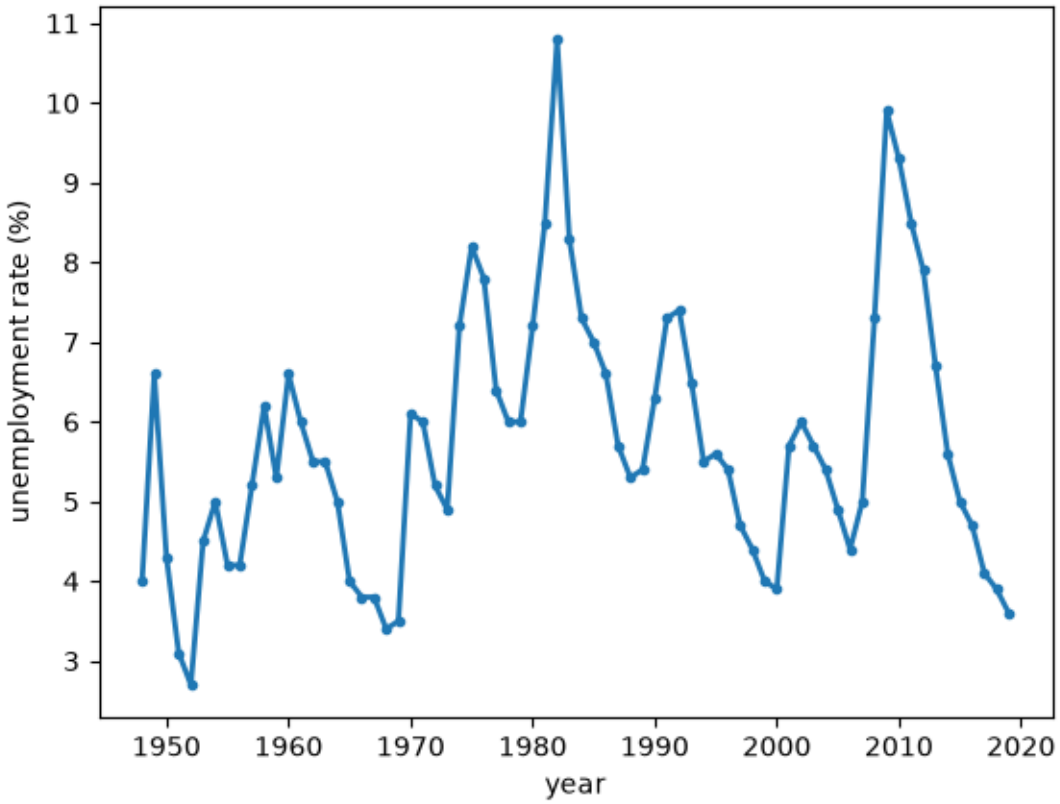


Fig. 147.1: Annual US unemployment

Most years are quiet, with small noise.

Occasionally a positive jump pushes unemployment up by a large amount.

Then, in most cases, the linear reversion pulls it slowly back down toward $\bar{u}$.

Thus, a spike arrives in one step, while the recovery takes time — the asymmetric dynamics we observe in the data.

This is, in spirit, Milton Friedman's "plucking model": a floor near full employment, from which bad shocks pluck the series upward.

The next cell shows the density of the shock and its two components.

```
def normal_pdf(x, m, s):
    return np.exp(-(x - m)**2 / (2 * s**2)) / (s * np.sqrt(2 * np.pi))

η = np.linspace(-3, 6, 400)
p_, μJ_, σs_, σJ_ = 0.25, 2.0, 0.4, 1.2
quiet = (1 - p_) * normal_pdf(η, 0.0, σs_)
jump  = p_ * normal_pdf(η, μJ_, σJ_)

fig, ax = plt.subplots()
ax.plot(η, quiet + jump, lw=2, label='shock density')
ax.fill_between(η, quiet, alpha=0.3, label='quiet years')
ax.fill_between(η, jump, alpha=0.3, label='recession jumps')
ax.set_xlabel('innovation $\\eta$')
ax.legend()
plt.show()
```

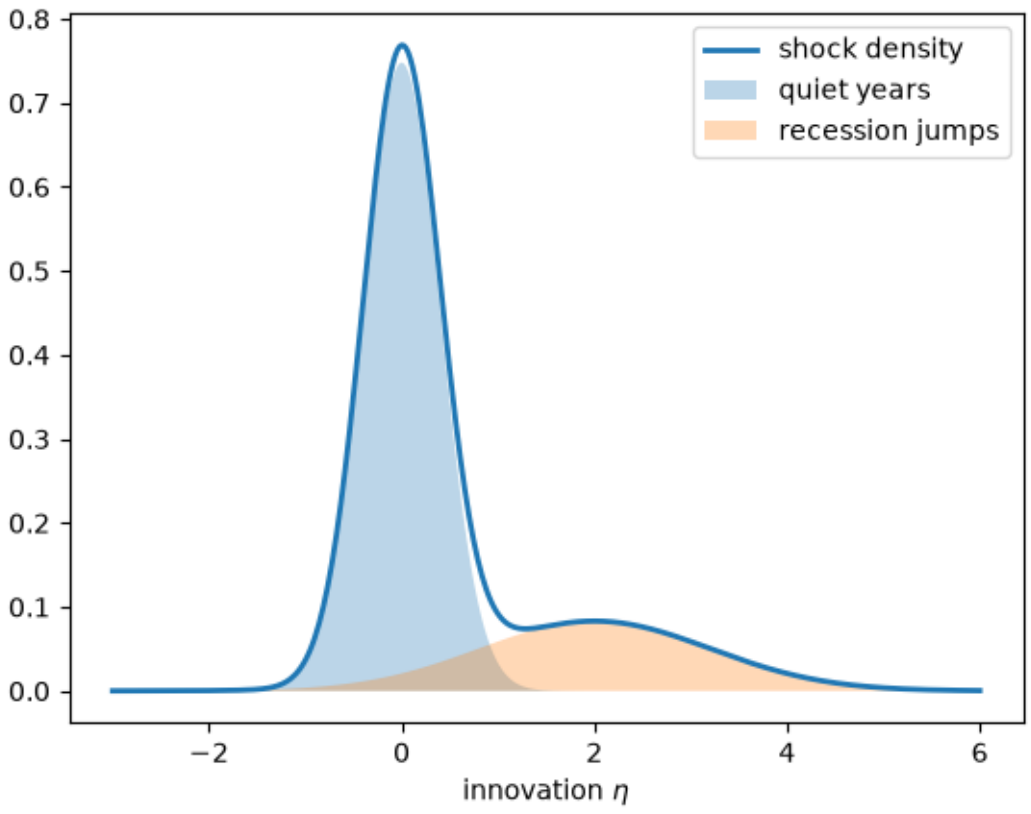


Fig. 147.2: Shock distribution as a mixture

Fig. 147.2 shows a tall, narrow peak at zero for ordinary years, plus a low, broad bump on the positive side for the recession jumps.

147.3 Bayesian estimation

We place weakly informative priors on the six parameters.

The mixture is written with NumPyro's `MixtureSameFamily`, which sums over the two components analytically.

```
def jump_model(u):
    # Set the priors
    ubar = numpyro.sample("ubar",    dist.Normal(4.5, 1.5))
    ρ    = numpyro.sample("rho",     dist.Uniform(0.0, 1.0))
    p    = numpyro.sample("p",       dist.Beta(2.0, 8.0))
    μ_J  = numpyro.sample("mu_J",    dist.HalfNormal(2.0))
    σ_s  = numpyro.sample("sigma_s", dist.HalfNormal(0.5))
    σ_J  = numpyro.sample("sigma_J", dist.HalfNormal(1.5))
    # Build the model
    n = u.shape[0] - 1
    base   = ubar + ρ * (u[:-1] - ubar)
    locs   = jnp.stack([base, base + μ_J], axis=-1)
    scales = jnp.stack([jnp.broadcast_to(σ_s, (n,)),
                        jnp.broadcast_to(σ_J, (n,))], axis=-1)
    probs  = jnp.broadcast_to(jnp.stack([1 - p, p]), (n, 2))
    mix = dist.MixtureSameFamily(dist.Categorical(probs=probs),
                                 dist.Normal(locs, scales))
    numpyro.sample("u_obs", mix, obs=u[1:])
```

We sample with NUTS, four chains, vectorized so the code runs on a CPU or a GPU.

```
def run_nuts(model, data, seed=0, num_warmup=2000, num_samples=4000, num_chains=4):
    "Sample a NumPyro model with the NUTS sampler."
    mcmc = MCMC(NUTS(model),
                num_warmup=num_warmup, num_samples=num_samples,
                num_chains=num_chains, chain_method="vectorized",
                progress_bar=False)
    mcmc.run(random.key(seed), jnp.asarray(data))
    return mcmc
```

We fit the model to the annual data.

```
mcmc_jump = run_nuts(jump_model, u_annual)
```

Let us look at the posterior.

```
mcmc_jump.print_summary()
```

```
                mean       std    median      5.0%     95.0%     n_eff     r_hat
      mu_J      1.25      0.42      1.20      0.58      1.96   3392.90      1.00
         p      0.35      0.09      0.35      0.20      0.50   4640.47      1.00
       rho      0.83      0.04      0.84      0.76      0.90   5896.66      1.00
   sigma_J      1.29      0.26      1.30      0.87      1.73   3866.36      1.00
   sigma_s      0.39      0.09      0.38      0.25      0.54   3978.27      1.00
      ubar      3.02      0.74      3.08      1.85      4.22   5563.53      1.00
```

(continues on next page)

(continued from previous page)

```
Number of divergences: 0
```

The `r_hat` values are essentially one, so the chains have converged.

Notice the distinct jump component: the normal year spread σ_s is small and the jump spread σ_J is large, with a positive mean μ_J.

Let's do a quick check by overlaying the fitted shock density on the model's actual residuals, as we did in *A Linear Model of Unemployment*.

```
m = {k: np.median(np.asarray(mcmc_jump.get_samples()[k]))
     for k in ("ubar", "rho", "p", "mu_J", "sigma_s", "sigma_J")}
resid = u_annual[1:] - (m["ubar"] + m["rho"] * (u_annual[:-1] - m["ubar"]))

η = np.linspace(resid.min() - 0.5, resid.max() + 0.5, 400)
quiet = (1 - m["p"]) * normal_pdf(η, 0.0, m["sigma_s"])
jump  = m["p"] * normal_pdf(η, m["mu_J"], m["sigma_J"])

fig, ax = plt.subplots()
ax.hist(resid, bins=25, density=True, alpha=0.5, label='residuals')
ax.plot(η, quiet + jump, lw=2, label='fitted density')
ax.fill_between(η, jump, alpha=0.3, color='C2', label='jump component')
ax.set_xlabel('innovation $\\eta$')
ax.legend()
plt.show()
```

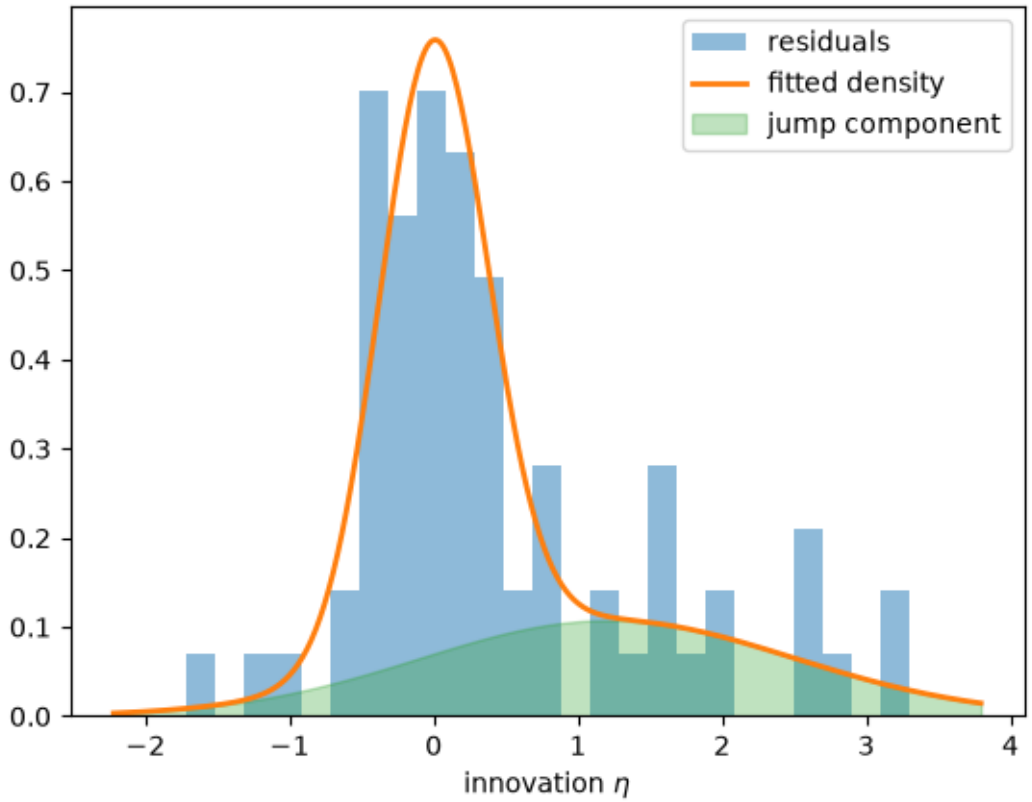


Fig. 147.3: Fitted shock density against the residuals

The fit looks quite good (although we are still struggling to capture some of the large shocks in the right tail).

147.4 Comparing models with cross-validation

The jump model looks reasonable, but is it actually better than the linear one?

To answer this question, we will use a form of *Bayesian model comparison*.

We build it up in stages: the guiding principle, the leave-one-out estimate, and how the computation is actually done.

147.4.1 The guiding principle: out-of-sample predictions

It is tempting to compare two models by asking which one fits the observed data better.

But this isn't the right criterion: more complex models can always be tuned to better match the data they are fitting.

The right question is instead: which model better predicts data it has not seen?

In other words, the criterion we care about is out-of-sample predictive accuracy.

We do not have new data for out-of-sample testing, so we manufacture some using *cross-validation*.

147.4.2 Leave-one-out cross-validation

The idea is to leave out one observation, fit the model to the rest, and then score the held-out point using a predictive distribution constructed from the rest of the data set.

The predictive distribution for data point u_i given the rest of the data u_{-i} is

$$p(u_i \mid u_{-i}) = \int p(u_i \mid \theta)\, p(\theta \mid u_{-i})\, d\theta.$$

Here $p(\theta \mid u_{-i})$ is the posterior obtained from the data in u_{-i}.

We score each held-out point by the *log* of this density.

This rewards a model for putting high probability on what actually happened.

Adding these scores over all data points gives the **expected log predictive density** (elpd), here in its leave-one-out (LOO) form,

$$\text{elpd}_{\text{loo}} = \sum_{i=1}^{n} \log p(u_i \mid u_{-i}).$$

147.4.3 Computational methods

The definition of elpd suggests that we need to refit the model n times, once for each omitted point.

This would be very time-consuming.

Here we discuss a method that removes all but the first fit, using the full posterior — which we have already sampled.

Regarding notation: let $u = (u_1, \ldots, u_n)$, and write u_{-i} for the sample with u_i deleted.

The method depends on the following independence-type assumption:

Assumption 147.4.1

Given θ, the observations are conditionally independent:

$$p(u \mid \theta) = \prod_{j=1}^{n} p(u_j \mid \theta).$$

For our autoregressive data this assumption is false, since each u_i is dependent on earlier observations through the transition dynamics.

We adopt it nonetheless, since

- it's what makes the leave-one-out computation tractable and
- it's a very standard procedure that needs to be understood.

We discuss the consequences in the section on time-series structure below.

Conditional on our assumption, the proposition that follows is exact.

Proposition 147.4.1

Under *Assumption 147.4.1*, we have

$$\frac{1}{p(u_i \mid u_{-i})} = \int \frac{1}{p(u_i \mid \theta)} p(\theta \mid u)\, d\theta \qquad (i = 1, \dots, n).$$

Proof. By Bayes' rule, $p(\theta \mid u) = p(u \mid \theta)\, p(\theta)/p(u)$ and $p(\theta \mid u_{-i}) = p(u_{-i} \mid \theta)\, p(\theta)/p(u_{-i})$.

By *Assumption 147.4.1*, $p(u \mid \theta) = p(u_i \mid \theta)\, p(u_{-i} \mid \theta)$.

Dividing the two posteriors and substituting this factorization,

$$\frac{p(\theta \mid u)}{p(\theta \mid u_{-i})} = \frac{p(u \mid \theta)}{p(u_{-i} \mid \theta)} \cdot \frac{p(u_{-i})}{p(u)} = p(u_i \mid \theta) \cdot \frac{p(u_{-i})}{p(u)},$$

so, rearranging,

$$\frac{p(\theta \mid u)}{p(u_i \mid \theta)} = \frac{p(u_{-i})}{p(u)} p(\theta \mid u_{-i}).$$

Now integrate both sides over θ.

On the right, $\int p(\theta \mid u_{-i})\, d\theta = 1$, leaving the constant factor $p(u_{-i})/p(u)$, so

$$\int \frac{1}{p(u_i \mid \theta)} p(\theta \mid u)\, d\theta = \frac{p(u_{-i})}{p(u)} = \frac{1}{p(u_i \mid u_{-i})},$$

the last equality by $p(u) = p(u_i \mid u_{-i})\, p(u_{-i})$.

The leave-i-out posterior has vanished; only the full posterior remains, and of it we have S draws $\theta^1, \dots, \theta^S$.

Replacing the integral in *Proposition 147.4.1* by the average over these draws gives

$$\frac{1}{p(u_i \mid u_{-i})} \;\approx\; \frac{1}{S}\sum_{s=1}^{S}\frac{1}{p(u_i \mid \theta^s)}, \qquad \text{so} \qquad p(u_i \mid u_{-i}) \;\approx\; \frac{S}{\sum_{s=1}^{S} 1/p(u_i \mid \theta^s)},$$

the harmonic mean of the per-draw likelihoods of u_i.

The integrand has finite mean, equal to $1/p(u_i \mid u_{-i})$ itself, so the law of large numbers applies.

The average therefore converges to that mean as $S \to \infty$, and the approximation becomes exact, with probability one, in the limit.

Its one ingredient is the likelihood of each observation under each draw — the **pointwise log-likelihood** — which is why we have been careful to keep it.

For numerical stability we compute on the log scale, where the harmonic mean becomes

$$\log p(u_i \mid u_{-i}) \;\approx\; \log S - \log \sum_{s=1}^{S} e^{-\ell_i^s}, \qquad \ell_i^s = \log p(u_i \mid \theta^s),$$

which is the `jnp.log(S) - logsumexp(-ll)` we use below.

To compare against the linear model, we first fit it on the same data.

```
def linear_model(u):
    ubar = numpyro.sample("ubar",  dist.Normal(5.5, 2.0))
    ρ    = numpyro.sample("rho",   dist.Uniform(0.0, 1.0))
    σ    = numpyro.sample("sigma", dist.HalfNormal(1.0))
    numpyro.sample("u_obs", dist.Normal(ubar + ρ * (u[:-1] - ubar), σ), obs=u[1:])
```

We sample it on the same annual data.

```
mcmc_lin = run_nuts(linear_model, u_annual)
```

Now we can do the whole leave-one-out calculation by hand in a few lines.

```
from jax.scipy.special import logsumexp

def pointwise_loo(mcmc, model):
    "Leave-one-out log predictive density for each observation."
    ll = log_likelihood(model, mcmc.get_samples(),
                        u=jnp.asarray(u_annual))["u_obs"]  # (draws, obs)
    S = ll.shape[0]
    # log of the harmonic mean of the per-draw likelihoods, per observation
    return jnp.log(S) - logsumexp(-ll, axis=0)

elpd_jump = pointwise_loo(mcmc_jump, jump_model)
elpd_lin  = pointwise_loo(mcmc_lin,  linear_model)

print(f"jump   elpd_loo = {elpd_jump.sum():.1f}")
print(f"linear elpd_loo = {elpd_lin.sum():.1f}")
```

```
jump   elpd_loo = -93.9
linear elpd_loo = -105.7
```

The jump model scores higher (closer to zero), so on out-of-sample prediction it wins.

To help understand whether the gap is real, we compute its standard error:

```
diff = elpd_jump - elpd_lin
n = diff.size
print(f"elpd difference = {diff.sum():.1f}")
print(f"standard error  = {jnp.sqrt(n) * diff.std():.1f}")
```

```
elpd difference = 11.8
standard error  = 5.7
```

The jump model is ahead by about twelve points, against a standard error near six — a bit over two standard errors, so a real improvement.

147.4.4 ArviZ implementation

In practice, most people do this calculation with a library like ArviZ.

We hand it the same pointwise log-likelihoods, packaged in its data format.

```
import arviz as az
import xarray as xr

def to_arviz(mcmc, model, u):
    "Package a NumPyro fit for ArviZ, with pointwise log-likelihoods."
    idata = az.from_numpyro(mcmc)
    ll = np.asarray(log_likelihood(model, mcmc.get_samples(),
                                   u=jnp.asarray(u))["u_obs"])
    ll = ll.reshape(mcmc.num_chains, -1, ll.shape[-1])
    idata["log_likelihood"] = xr.Dataset(
        {"u_obs": (("chain", "draw", "obs"), ll)})
    return idata

az.compare({
    "linear": to_arviz(mcmc_lin,  linear_model, u_annual),
    "jump":   to_arviz(mcmc_jump, jump_model,   u_annual),
})
```

```
        rank  elpd_diff  dse  p_worse diag_diff diag_elpd    p   elpd   se  \
jump       0        0.0  0.0      NaN                     6.5  -90.0  9.3
linear     1      -10.0  5.7     0.98   N < 100           3.5 -110.0  7.4

        weight
jump      0.86
linear    0.14
```

Here `az.compare` confirms the hand calculation and ranks the jump model first (its table rounds the scores for display).

147.4.5 Accommodating the time series structure

Leave-one-out drops one transition at a time, but it does not respect the order of time.

When it scores the step from u_t to u_{t+1}, the model doing the scoring was fit on data lying on *both* sides of that step — including the neighboring values u_t and u_{t+1} themselves.

So the held-out point was never truly unseen.

The formally correct measure is **leave-future-out** cross-validation.

It only ever predicts forward: fit the model on $u_1, \ldots, u_t$, score its one-step-ahead forecast of u_{t+1}, then expand the window by one step and repeat.

Now the model is genuinely blind to the future it is asked to predict, exactly as in real forecasting.

The price is computational.

Leave-one-out reused a single fit through its importance-sampling shortcut, but leave-future-out has no such trick.

For our short annual series leave-future-out cross-validation is still feasible but for a long series it can become prohibitive.

(We ran it without including it here and the conclusion holds: under leave-future-out the jump model still beats the linear one, by an even larger margin.)

147.5 Conclusion

In *A Linear Model of Unemployment* we applied a linear AR(1) model with Gaussian shocks to unemployment and found high levels of persistence in monthly data.

We also argued that the model is overly simplistic.

Here we found that the feature the linear model most conspicuously misses — the asymmetry of recessions — can be addressed by modeling the distribution of the shocks.

A linear reversion with large, one-sided innovations matches the right-skewed, heavy-tailed shocks that the Gaussian model could not.

Cross-validation prefers it.

An even better model than the one considered above would let the economy switch between persistent expansion and recession regimes.

This is the Markov-switching approach, which we leave to further reading.

147.6 Exercises

Exercise 147.6.1

The jump component here is symmetric within itself, $N(\mu_J, \sigma_J^2)$.

Replace the two-component mixture with a single **skew-normal** (or Student-t with a skew) innovation, refit, and compare to the mixture model with `az.compare`.

Does the simpler skewed shock do as well as the mixture?

CHAPTER

EIGHT

TWO ILLUSTRATIONS OF THE QUANTITY THEORY OF MONEY

Contents

- *Two Illustrations of the Quantity Theory of Money*
  - *Overview*
  - *Lucas's moving averages*
  - *Whiteman's reinterpretation*
  - *Data*
  - *Scatter plots*
  - *Regressions on filtered data*
  - *Estimates of $h(0)$ from a VAR*
  - *A model for monetary policy analysis*
  - *Three things worth checking*
  - *Solving the model*
  - *The Sargent-Surico model in canonical form*
  - *Bayesian estimation*
  - *How monetary policy moves the low-frequency slopes*
  - *Estimating with Hamiltonian Monte Carlo*
  - *Concluding remarks*
  - *Exercises*

In addition to what's in Anaconda, this lecture uses `pandas_datareader` to download macroeconomic data and `jax`, `numpyro` and `arviz` for the Hamiltonian Monte Carlo section at the end:

```
!pip install pandas_datareader numpyro jax arviz
```

148.1 Overview

Lucas [1980] plotted long moving averages of U.S. inflation against long moving averages of U.S. money growth, and then long moving averages of a nominal interest rate against the same moving averages of money growth.

Both scatter plots hugged a 45 degree line.

Lucas read those two unit slopes as illustrating "two central implications of the quantity theory of money: that a given change in the rate of change in the quantity of money induces (i) an equal change in the rate of price inflation; and (ii) an equal change in nominal rates of interest."

He also warned that a theory tells us "conditions under which one might expect them to break down."

This lecture studies Sargent and Surico [2011], which takes that warning seriously.

The paper does three things.

First, it extends Lucas's sample and shows that his two slopes are *not* stable across subperiods.

Second, following Whiteman [1984], it interprets a Lucas slope as the sum of coefficients in a two-sided distributed lag regression, an object that a time series model delivers as a ratio of spectral densities at frequency zero.

Third, it estimates a small new Keynesian model on a pre-1984 sample and then perturbs *only* the monetary policy rule, showing that policy alone can move the two slopes across the range that the data display.

We do all of this from scratch in Python.

We write our own solver for linear rational expectations models, our own Kalman filter, and our own Metropolis-Hastings sampler, so that every step is visible.

A final section then uses the estimated model as a test bed for Hamiltonian Monte Carlo, which turns out to require replacing the model solver with a differentiable one.

Along the way we flag several places where the published paper's statement or implementation of its model needs care, and we check each of them numerically.

Most are corrections, and one turns out to sharpen the paper's message rather than weaken it.

Let's start with imports.

```
import datetime
import time
import warnings

import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
import pandas_datareader.data as web
from scipy import optimize, stats
from scipy.linalg import ordqz, solve_discrete_lyapunov, svd
from scipy.stats import invwishart

warnings.filterwarnings('ignore')
plt.rcParams['figure.figsize'] = (10, 6)
plt.rcParams['axes.grid'] = True
plt.rcParams['grid.alpha'] = 0.3
```

148.2 Lucas's moving averages

For a scalar series x_t and $\beta \in [0, 1)$, Lucas [1980] formed the two-sided moving average

$$\bar{x}_t(\beta) = a \sum_{k=-n}^{n} \beta^{|k|} x_{t+k}, \qquad a = \frac{(1-\beta)^2}{1 - \beta^2 - 2\beta^{n+1}(1-\beta)}, \tag{148.1}$$

where a makes the weights sum to one.

Larger β means a smoother series and a longer window.

Setting $\beta = 0$ returns the raw data.

We apply the weights to whatever observations are available, renormalizing so that they always sum to one.

```
def lucas_filter(x, beta=0.95):
    """Two-sided exponentially weighted moving average of Lucas (1980)."""
    x = np.asarray(x, dtype=float)
    T = len(x)
    out = np.empty(T)
    for t in range(T):
        w = beta ** np.abs(np.arange(T) - t)
        out[t] = w @ x / w.sum()
    return out
```

148.3 Whiteman's reinterpretation

Whiteman [1984] observed that fitting a straight line through a scatter of moving averages is an informal way of computing the *sum* of coefficients in a long two-sided distributed lag regression.

Let $\{y_t, z_t\}$ be jointly covariance stationary with zero means, and project y_t on the entire history and future of z,

$$y_t = \sum_{j=-\infty}^{\infty} h_j z_{t-j} + \epsilon_t, \qquad \mathbb{E}\,\epsilon_t z_{t-j} = 0 \;\; \forall j. \tag{148.2}$$

Write $S_y(\omega)$, $S_z(\omega)$ and $S_{yz}(\omega)$ for the spectral and cross-spectral densities and $h(\omega) = \sum_j h_j e^{-i\omega j}$ for the transfer function.

Then

$$h(\omega) = \frac{S_{yz}(\omega)}{S_z(\omega)} \qquad \text{so} \qquad \sum_{j=-\infty}^{\infty} h_j = h(0) = \frac{S_{yz}(0)}{S_z(0)}. \tag{148.3}$$

Whiteman showed that as $\beta \to 1$ the regression coefficient of $\bar{y}_t(\beta)$ on $\bar{z}_t(\beta)$ converges to $h(0)$.

So Lucas's two illustrations are claims that

$$h_{\pi,\Delta m}(0) = 1 \qquad \text{and} \qquad h_{R,\Delta m}(0) = 1,$$

where π is inflation, Δm is money growth, and R is a nominal interest rate.

This is a useful reformulation because $h(0)$ is a *population* object that any time series model delivers.

148.3.1 From a state space model to $h(0)$

Suppose a model implies

$$X_{t+1} = AX_t + BW_{t+1}, \qquad Y_{t+1} = CX_t + DW_{t+1}, \tag{148.4}$$

with W_{t+1} a standard normal vector that is IID over time and A stable.

Iterating gives $X_t = \sum_{j\geq 0} A^j B W_{t-j}$, so the transfer function from W to Y is

$$H(\zeta) = D + \zeta\, C(I - A\zeta)^{-1}B, \qquad \zeta = e^{-i\omega},$$

and the spectral density matrix of Y is $S_Y(\omega) = H(\zeta)H(\zeta)^*/(2\pi)$, that is

$$2\pi S_Y(\omega) = C(I - Ae^{-i\omega})^{-1}BB'(I - A'e^{i\omega})^{-1}C' + DD' + e^{-i\omega}C(I - Ae^{-i\omega})^{-1}BD' + e^{i\omega}DB'(I - A'e^{i\omega})^{-1}C'. \tag{148.5}$$

Note

Equation (7) of Sargent and Surico [2011] reports only the first two terms of (148.5).

The two cross terms drop out exactly when $BD' = 0$, that is when the shocks that move the state are orthogonal to the shocks that hit the observation equation.

A one-line example shows that they matter otherwise: with scalars $A = 1/2$, $B = C = D = 1$, the correct $2\pi S_Y(0)$ is $(1 + 1/(1 - 1/2))^2 = 9$, while the truncated formula gives $1/(1 - 1/2)^2 + 1 = 5$.

Nothing in what follows turns on this, because we write both our VAR and our solved model with $D = 0$, which is the natural way to write each of them.

At $\omega = 0$ with $D = 0$ everything collapses to a long-run multiplier.

Writing $G = C(I - A)^{-1}B$,

$$2\pi S_Y(0) = GG', \qquad h_{y,z}(0) = \frac{[GG']_{yz}}{[GG']_{zz}}. \tag{148.6}$$

```
def h_zero_from_state_space(A, B, C, iy, iz):
    """h(0) for observable iy regressed on observable iz."""
    G = np.linalg.solve(np.eye(A.shape[0]) - A, B)
    S0 = C @ G @ G.T @ C.T
    return S0[iy, iz] / S0[iz, iz]
```

148.4 Data

Sargent and Surico [2011] splice FRED data onto the historical series of Balke and Gordon [1986] to reach back to 1900.

Those historical series are not on FRED, so we work with the post-1959 quarterly data that FRED does supply, and we extend the sample by twenty years beyond the paper's 2005 endpoint.

That extension is the interesting part: it covers the balance sheet expansions after 2008 and after 2020, and the inflation of 2021-2023.

We use M2 for money, the GDP deflator for prices, real GDP for output, and the three-month Treasury bill rate for the nominal interest rate.

The paper's benchmark is a six-month commercial paper rate, which FRED no longer carries; its Appendix Table A reports that the three-month bill gives similar answers.

Money growth, inflation, and output growth are 100 times quarterly log differences, and the interest rate is divided by four so that it too is a quarterly rate in percent.

```
start, end = datetime.datetime(1959, 1, 1), datetime.datetime(2030, 1, 1)
raw = web.DataReader(['M2SL', 'GDPDEF', 'GDPC1', 'TB3MS'], 'fred', start, end)

q = pd.DataFrame({'M2': raw['M2SL'].resample('QS').mean(),
                  'P': raw['GDPDEF'].resample('QS').mean(),
                  'Y': raw['GDPC1'].resample('QS').mean(),
                  'R': raw['TB3MS'].resample('QS').mean() / 4}).dropna()

data = pd.DataFrame({'dm': 100 * np.log(q['M2']).diff(),
                     'pi': 100 * np.log(q['P']).diff(),
                     'dy': 100 * np.log(q['Y']).diff(),
                     'R': q['R']}).dropna()

VARS = ['dm', 'pi', 'R', 'dy']
LABELS = {'dm': 'M2 growth', 'pi': 'inflation',
          'R': 'T-bill rate', 'dy': 'real GDP growth'}
print(f"sample: {data.index[0].date()} to {data.index[-1].date()},  T = {len(data)}")
data.describe().round(3)
```

```
sample: 1959-04-01 to 2026-07-01,  T = 270
```

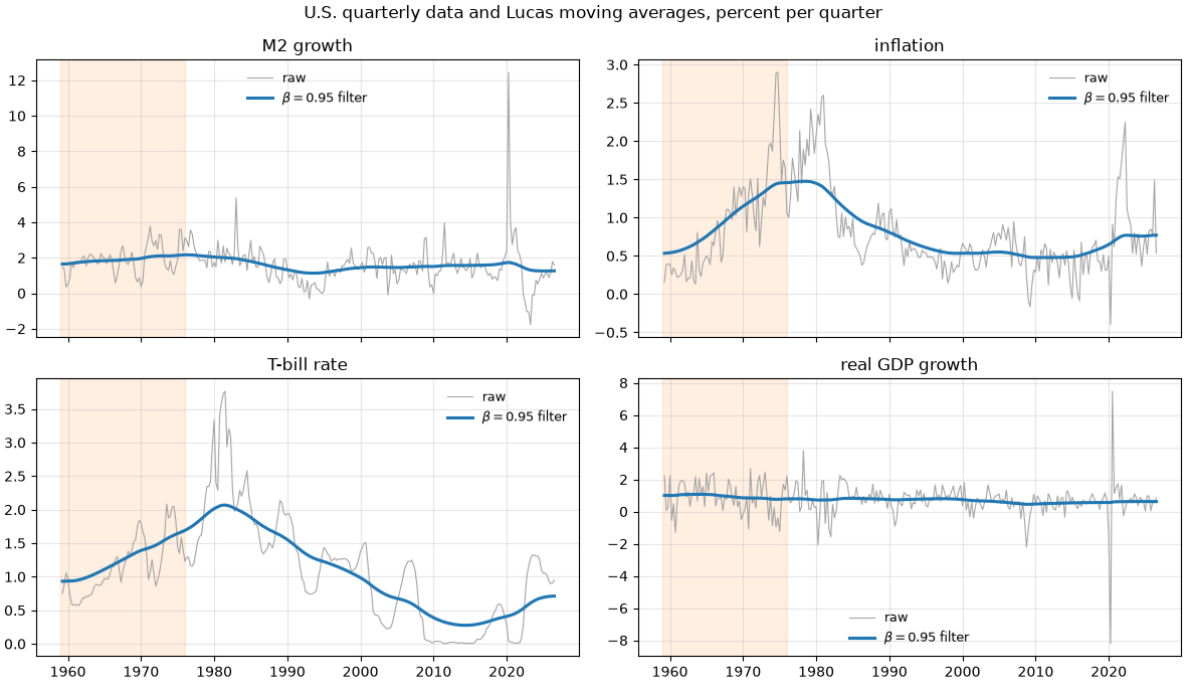
	dm	pi	dy	R
count	270.000	270.000	270.000	270.000
mean	1.627	0.806	0.738	1.089
std	1.114	0.574	1.049	0.769
min	-1.770	-0.399	-8.204	0.003
25%	1.099	0.424	0.349	0.527
50%	1.600	0.627	0.764	1.099
75%	2.087	1.023	1.120	1.433
max	12.448	2.899	7.476	3.763

Here are the raw series and their $\beta = 0.95$ moving averages.

The shaded band marks 1959-1975, the closest we can come to Lucas's window.

```
filtered = pd.DataFrame({c: lucas_filter(data[c].values, 0.95) for c in VARS},
                        index=data.index)

fig, axes = plt.subplots(2, 2, figsize=(12, 7), sharex=True)
for ax, c in zip(axes.flat, VARS):
    ax.plot(data.index, data[c], lw=0.8, color='0.65', label='raw')
    ax.plot(filtered.index, filtered[c], lw=2.2, color='C0',
            label=r'$\beta = 0.95$ filter')
    ax.axvspan(pd.Timestamp('1959-01-01'), pd.Timestamp('1975-12-31'),
               color='C1', alpha=0.12)
    ax.set_title(LABELS[c])
    ax.legend(frameon=False, fontsize=9)
fig.suptitle('U.S. quarterly data and Lucas moving averages, percent per quarter')
plt.tight_layout()
plt.show()
```



Money growth, inflation, and the interest rate rise together into the early 1980s and then fall together.

After 1984 that comovement is much harder to see, and the pandemic surge in M2 growth stands entirely apart from everything else.

148.5 Scatter plots

Following the paper we plot second quarter observations of each year, which reduces the overlap between neighboring filtered values.

```
PERIODS = [('1959-1975', 1959, 1975), ('1960-1983', 1960, 1983),
           ('1984-2007', 1984, 2007), ('2008-present', 2008, 2100)]

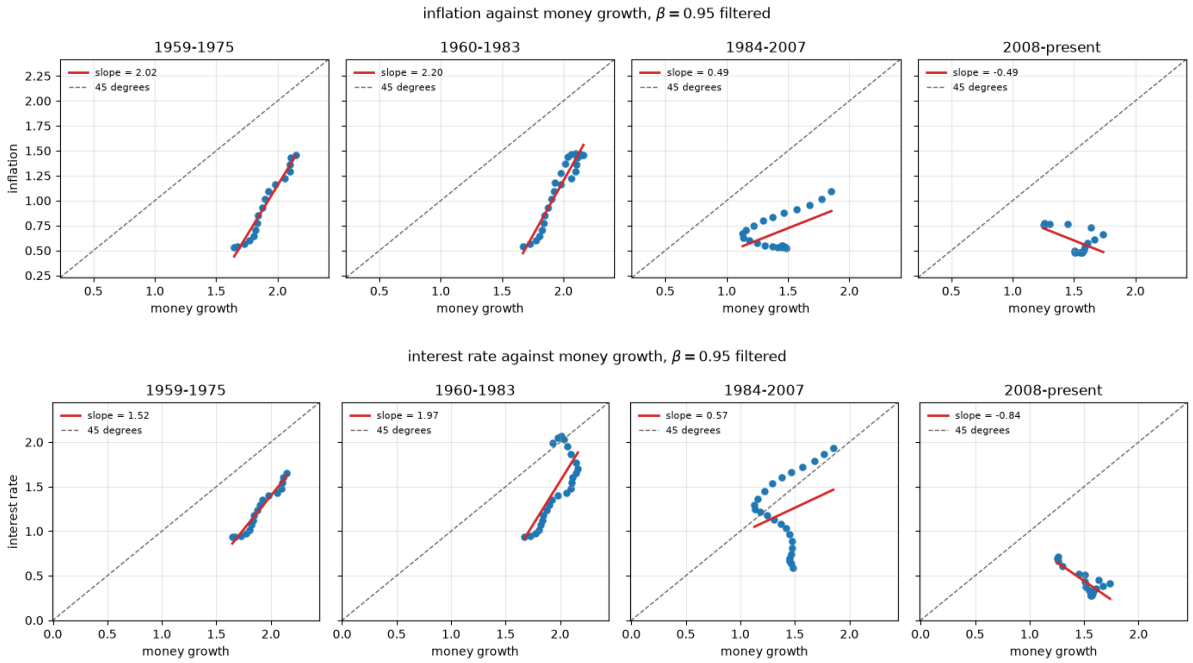

def scatter_panel(yvar, title):
    q2 = filtered[filtered.index.quarter == 2]
    pad = 0.15 * (q2[['dm', yvar]].values.max() - q2[['dm', yvar]].values.min())
    lim = np.array([q2[['dm', yvar]].values.min() - pad,
                    q2[['dm', yvar]].values.max() + pad])
    fig, axes = plt.subplots(1, 4, figsize=(14, 3.8), sharex=True, sharey=True)
    for ax, (name, lo, hi) in zip(axes, PERIODS):
        sub = q2[(q2.index.year >= lo) & (q2.index.year <= hi)]
        x, y = sub['dm'].values, sub[yvar].values
        ax.scatter(x, y, s=26, color='C0')
        b = np.polyfit(x, y, 1)
        xx = np.linspace(x.min(), x.max(), 2)
        ax.plot(xx, np.polyval(b, xx), color='C3', lw=2,
                label=f'slope = {b[0]:.2f}')
        ax.plot(lim, lim, color='0.4', ls='--', lw=1, label='45 degrees')
        ax.set(xlim=lim, ylim=lim, title=name, xlabel='money growth')
        ax.legend(frameon=False, fontsize=8, loc='upper left')
```

(continues on next page)

(continued from previous page)

```
    axes[0].set_ylabel(title)
    fig.suptitle(f'{title} against money growth, $\\beta = 0.95$ filtered')
    plt.tight_layout()
    plt.show()


scatter_panel('pi', 'inflation')
scatter_panel('R', 'interest rate')
```



Lucas's window and the pre-Volcker sample line up with the 45 degree line.

The great moderation and the period since 2008 do not.

148.6 Regressions on filtered data

Table 1 of the paper reports these slopes for a grid of β values.

```
def filtered_slopes(betas=(0.95, 0.8, 0.5, 0.0)):
    rows = {}
    for beta in betas:
        f = pd.DataFrame({c: lucas_filter(data[c].values, beta) for c in VARS},
                         index=data.index)
        for name, lo, hi in PERIODS + [('full sample', 1900, 2100)]:
            m = ((f.index.year >= lo) & (f.index.year <= hi)
                 & (f.index.quarter == 2))
            x = f.loc[m, 'dm'].values
            rows.setdefault(name, {})[('pi on dm', beta)] = np.polyfit(x, f.loc[m, 'pi
↪'], 1)[0]
            rows[name][('R on dm', beta)] = np.polyfit(x, f.loc[m, 'R'], 1)[0]
    tab = pd.DataFrame(rows).T
    order = [(v, b) for v in ['pi on dm', 'R on dm'] for b in betas]
```

(continues on next page)

(continued from previous page)

```
    return tab.reindex(columns=order)


filtered_slopes().round(2)
```

```
              pi on dm                    R on dm
                  0.95  0.80  0.50  0.00    0.95  0.80  0.50  0.00
1959-1975         2.02  1.06  0.24  0.01    1.52  0.54 -0.02 -0.12
1960-1983         2.20  1.14  0.23  0.04    1.97  0.94  0.15  0.05
1984-2007         0.49  0.02 -0.04 -0.03    0.57  0.23  0.21  0.08
2008-present     -0.49 -0.07 -0.05 -0.08   -0.84 -0.33 -0.15 -0.06
full sample       0.85  0.39  0.10 -0.04    0.90  0.30  0.04 -0.03
```

Two patterns stand out.

The closer β is to one, the larger the slopes, exactly as in Lucas's graphs and in the paper's Table 1.

And the slopes fall steadily as we move through the sample, from above one in Lucas's window to negative after 2008.

148.7 Estimates of $h(0)$ from a VAR

The filtered regressions are informal.

Equation (148.6) lets us compute $h(0)$ properly from a fitted VAR.

We fit a VAR(2) in money growth, inflation, the interest rate, and output growth, put it in companion form, and read $h(0)$ off the long-run multiplier matrix.

We use a diffuse prior so that we can report posterior bands, as the paper does in its Figure 5.

Draws whose companion matrix is explosive are discarded, since $h(0)$ is not defined for them.

```
def bvar_h0(Y, p=2, n_draws=500, seed=0):
    """Posterior draws of h(0) from a VAR(p) under a diffuse prior."""
    rng = np.random.default_rng(seed)
    T, n = Y.shape
    X = np.column_stack([np.ones(T - p)] + [Y[p - l - 1:T - l - 1] for l in range(p)])
    Z, k = Y[p:], 1 + n * p
    XXi = np.linalg.inv(X.T @ X)
    B_hat = XXi @ X.T @ Z
    E = Z - X @ B_hat
    S, nu = E.T @ E, T - p - k
    L_x = np.linalg.cholesky(XXi)
    J = np.vstack([np.eye(n), np.zeros((n * (p - 1), n))])

    out = np.full((n_draws, 2), np.nan)
    for d in range(n_draws):
        Sigma = invwishart.rvs(df=nu, scale=S, random_state=rng)
        B = B_hat + L_x @ rng.standard_normal((k, n)) @ np.linalg.cholesky(Sigma).T
        A = np.zeros((n * p, n * p))
        A[:n] = np.hstack([B[1 + l * n:1 + (l + 1) * n].T for l in range(p)])
        if p > 1:
            A[n:, :n * (p - 1)] = np.eye(n * (p - 1))
        if np.max(abs(np.linalg.eigvals(A))) >= 1:
            continue
        G = np.linalg.solve(np.eye(n * p) - A, J)
```

(continues on next page)

(continued from previous page)

```
        S0 = G @ Sigma @ G.T
        out[d] = [S0[1, 0] / S0[0, 0], S0[2, 0] / S0[0, 0]]
    return out
```

```
rows = []
for name, lo, hi in PERIODS + [('full sample', 1900, 2100)]:
    m = (data.index.year >= lo) & (data.index.year <= hi)
    o = bvar_h0(data.loc[m, VARS].values)
    qs = np.nanpercentile(o, [16, 50, 84], axis=0)
    rows.append(dict(period=name, T=int(m.sum()),
                     h_pi=qs[1, 0], h_pi_lo=qs[0, 0], h_pi_hi=qs[2, 0],
                     h_R=qs[1, 1], h_R_lo=qs[0, 1], h_R_hi=qs[2, 1]))
pd.DataFrame(rows).set_index('period').round(2)
```

period	T	h_pi	h_pi_lo	h_pi_hi	h_R	h_R_lo	h_R_hi
1959-1975	67	1.26	0.22	2.36	0.61	-0.04	1.31
1960-1983	96	0.74	-0.31	1.90	0.61	-0.63	1.91
1984-2007	96	-0.07	-0.30	0.14	0.27	-0.29	0.81
2008-present	75	-0.09	-0.50	0.42	-0.50	-1.10	0.28
full sample	270	0.72	0.44	1.08	0.99	0.45	1.61

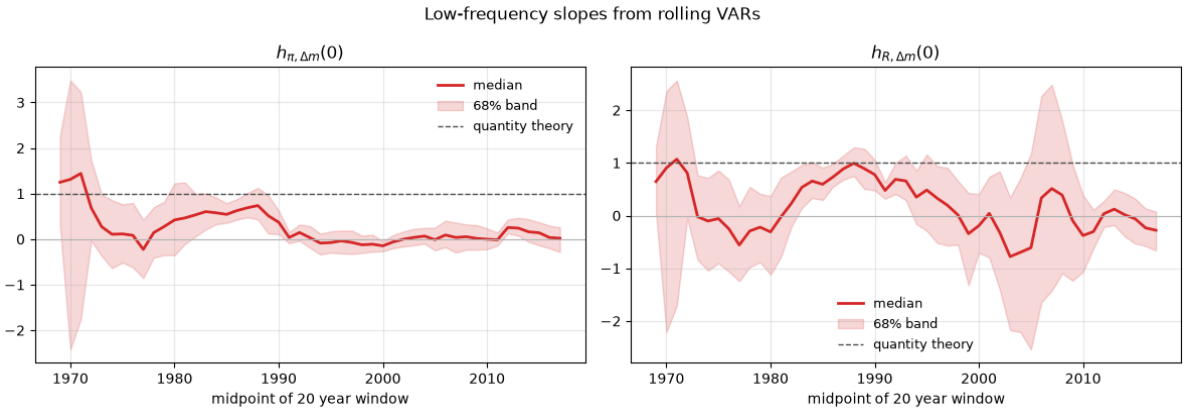
The subperiods are chosen by hand, which is the objection that Boschen and Otrok [1994] raised against this style of evidence.

So we also run the VAR through rolling twenty year windows.

```
years = np.arange(data.index.year.min(), data.index.year.max() - 18)
mid, lo68, hi68 = {}, {}, {}
for y in years:
    m = (data.index.year >= y) & (data.index.year < y + 20)
    if m.sum() < 70:
        continue
    o = bvar_h0(data.loc[m, VARS].values, n_draws=300, seed=int(y))
    qs = np.nanpercentile(o, [16, 50, 84], axis=0)
    mid[y + 10], lo68[y + 10], hi68[y + 10] = qs[1], qs[0], qs[2]

mid = pd.DataFrame(mid).T
lo68, hi68 = pd.DataFrame(lo68).T, pd.DataFrame(hi68).T

fig, axes = plt.subplots(1, 2, figsize=(12, 4.2), sharex=True)
for j, (ax, ttl) in enumerate(zip(axes, [r'$h_{\pi,\Delta m}(0)$',
                                         r'$h_{R,\Delta m}(0)$'])):
    ax.plot(mid.index, mid[j], color='C3', lw=2, label='median')
    ax.fill_between(mid.index, lo68[j], hi68[j], color='C3', alpha=0.18,
                    label='68% band')
    ax.axhline(1, color='0.3', ls='--', lw=1, label='quantity theory')
    ax.axhline(0, color='0.7', lw=0.8)
    ax.set(title=ttl, xlabel='midpoint of 20 year window')
    ax.legend(frameon=False, fontsize=9)
fig.suptitle('Low-frequency slopes from rolling VARs')
plt.tight_layout()
plt.show()
```



The two low-frequency slopes are high early in the sample, when Lucas wrote, and drift toward zero afterwards.

One is rarely inside the 68 percent band after the mid 1980s.

This is the instability that the rest of the lecture tries to explain.

148.8 A model for monetary policy analysis

Section III of Sargent and Surico [2011] uses the log-linearized sticky price model of Ireland [2004], which has price indexation, habit formation, and a unit root technology shock.

The structure of the economy is

$$\pi_t = \theta(1-\alpha_\pi)\mathbb{E}_t\pi_{t+1} + \theta\alpha_\pi\pi_{t-1} + \kappa x_t - \tfrac{1}{\tau}e_t, \tag{148.7}$$

$$x_t = (1-\alpha_x)\mathbb{E}_t x_{t+1} + \alpha_x x_{t-1} - \sigma(R_t - \mathbb{E}_t\pi_{t+1}) + \sigma(1-\xi)(1-\rho_a)a_t, \tag{148.8}$$

$$\Delta m_t = \pi_t + z_t + \tfrac{1}{\sigma\gamma}\Delta x_t - \tfrac{1}{\gamma}\Delta R_t + \tfrac{1}{\gamma}(\Delta\chi_t - \Delta a_t), \tag{148.9}$$

$$\tilde{y}_t = x_t + \xi a_t, \qquad \Delta y_t = \tilde{y}_t - \tilde{y}_{t-1} + z_t. \tag{148.10}$$

Here π_t is inflation, x_t the output gap, Δm_t nominal money growth, R_t the short rate, $\tilde{y}_t$ detrended output, and z_t the growth rate of technology.

Equation (148.7) is a new Keynesian Phillips curve and (148.8) a new Keynesian IS curve.

Equation (148.9) is the money demand relation of McCallum and Nelson [1999] and Ireland [2003].

The discount factor is θ, α_π is indexation to past inflation, α_x is habit formation, κ is the slope of the Phillips curve, σ is the elasticity of intertemporal substitution, τ is the Rotemberg [1982] price adjustment cost, ξ is the inverse Frisch elasticity, and $1/\gamma$ is the interest semi-elasticity of money demand.

Four nonpolicy disturbances drive the economy: a markup shock e_t, a demand shock a_t, a money demand shock χ_t, and technology z_t,

$$e_t = \rho_e e_{t-1} + \varepsilon_{et}, \quad a_t = \rho_a a_{t-1} + \varepsilon_{at}, \quad \chi_t = \rho_\chi\chi_{t-1} + \varepsilon_{\chi t}, \quad z_t = \varepsilon_{zt}. \tag{148.11}$$

Monetary policy is either a money growth rule

$$\Delta m_t = \rho_m\Delta m_{t-1} + (1-\rho_m)(\phi_\pi\pi_t + \phi_x x_t) + \varepsilon_{mt} \tag{148.12}$$

or a Taylor rule

$$R_t = \rho_r R_{t-1} + (1-\rho_r)(\psi_\pi\pi_t + \psi_x x_t) + \varepsilon_{Rt}. \tag{148.13}$$

The observables are $[\Delta m_t, \pi_t, R_t, \Delta y_t]$.

148.9 Three things worth checking

Before estimating anything it pays to look at what this parameterization can and cannot deliver.

148.9.1 τ is not identified

The price adjustment cost τ enters the model in exactly one place, the term e_t/τ in (148.7).

Since e_t is an AR(1) with a freely estimated innovation standard deviation σ_e, the process e_t/τ is an AR(1) with innovation standard deviation σ_e/τ.

So the likelihood depends on (τ, σ_e) only through the ratio σ_e/τ.

We verify this numerically below.

Table 2 of the paper nevertheless reports a posterior for τ with mean 3.51 and a 5-95 interval of $[1.99, 4.99]$, against a prior with mean 4 and interval $[2.51, 5.77]$.

That posterior is not evidence about the Rotemberg adjustment cost.

Its leftward shift is what the prior on σ_e implies once the data pin down the ratio: the data want $\sigma_e/\tau \approx 0.31$, while the inverse gamma prior on σ_e has mean 0.3 and pulls σ_e down, dragging τ with it.

Because τ contributes nothing, we fix it at its prior mean and estimate one fewer parameter.

Nothing in the fit, or in $h(0)$, changes.

148.9.2 The first unit slope is an identity when money is exogenous

Group the terms of the money demand equation (148.9) as

$$\Delta m_t = \pi_t + z_t + \Delta v_t, \qquad v_t = \tfrac{1}{\sigma\gamma} x_t - \tfrac{1}{\gamma} R_t + \tfrac{1}{\gamma}(\chi_t - a_t). \tag{148.14}$$

Everything other than $\pi_t + z_t$ is a *first difference*.

The filter $1 - e^{-i\omega}$ vanishes at $\omega = 0$, so Δv_t contributes nothing to any spectral density or cross spectrum at frequency zero.

Therefore

$$h_{\pi,\Delta m}(0) = \frac{S_{\pi,\pi+z}(0)}{S_{\pi+z}(0)} = 1 - \frac{S_{z,\pi+z}(0)}{S_{\pi+z}(0)}. \tag{148.15}$$

This yields a proposition.

Proposition 148.9.1

Under the money growth rule (148.12) with $\phi_\pi = \phi_x = 0$,

$$h_{\pi,\Delta m}(0) = 1$$

exactly, whatever the values of the other parameters.

Proof. With $\phi_\pi = \phi_x = 0$ the rule reads $\Delta m_t = \rho_m \Delta m_{t-1} + \varepsilon_{mt}$, so money growth is driven by ε_m alone.

From (148.14), $\pi_t + z_t = \Delta m_t - \Delta v_t$, and Δv contributes nothing at frequency zero, so $S_{z,\pi+z}(0) = S_{z,\Delta m}(0)$.

Technology is exogenous and orthogonal to ε_m, so $S_{z,\Delta m}(\omega) \equiv 0$.

Now apply (148.15).

This is worth thinking about.

Lucas's *first* illustration prevails whenever money growth is econometrically exogenous, which is the case Whiteman [1984] assumed.

It is a property of the money demand specification, not of a deep monetary neutrality.

Departures from one arise only through $S_{z,\pi+z}(0)$, and that term is nonzero only when the policy rule makes money growth respond to the endogenous variables, so that technology becomes correlated with money growth at frequency zero.

The *second* illustration has no such backing.

The same argument gives $h_{R,\Delta m}(0) = S_{R,\pi+z}(0)/S_{\pi+z}(0)$, which equals one only if the low-frequency behavior of the ex ante real rate happens to cooperate.

So within this structural model, Lucas's two illustrations are not on the same footing.

148.9.3 The reported posterior for α_x stops at one half

Table 2 reports a posterior for the habit parameter α_x with mean 0.4775 and a 95th percentile of exactly 0.5000, under a beta prior on $[0, 1]$.

An interval that ends on a round number is usually a constraint rather than a finding.

It is not a determinacy constraint: we check below that the model has a unique stable solution on both sides of $\alpha_x = 0.5$.

We therefore let α_x range over the whole unit interval.

148.10 Solving the model

The model is a linear rational expectations system, which we solve with the method of Sims [2002].

Write it in the canonical form

$$\Gamma_0 y_t = \Gamma_1 y_{t-1} + \Psi \varepsilon_t + \Pi \eta_t, \tag{148.16}$$

where ε_t collects the structural shocks and η_t collects expectational errors, one for each variable whose expectation appears in the system.

The trick is to add $\mathbb{E}_t \pi_{t+1}$ and $\mathbb{E}_t x_{t+1}$ to the vector of variables, together with the identities

$$\pi_t = \mathbb{E}_{t-1}\pi_t + \eta^\pi_t, \qquad x_t = \mathbb{E}_{t-1} x_t + \eta^x_t.$$

The solution algorithm takes a generalized Schur decomposition of the matrix pencil (Γ_0, Γ_1), orders the stable generalized eigenvalues first, and then asks whether the expectational errors can be chosen to kill the explosive block.

If they can, a stable solution exists; if they can be chosen in only one way, it is unique.

```
SMALL = 1e-6


def gensys(g0, g1, psi, pi, div=1.01):
    """
    Solve  g0 @ y[t] = g1 @ y[t-1] + psi @ eps[t] + pi @ eta[t].

    Returns (G1, impact, eu) so that, when eu == (1, 1),

        y[t] = G1 @ y[t-1] + impact @ eps[t]

    is the unique stable solution.  eu[0] flags existence, eu[1] uniqueness.
    """
    n = g0.shape[0]
    S, T, alpha, beta, Q, Z = ordqz(g0, g1, output='complex',
                                    sort=lambda a, b: abs(b) < div * abs(a))
    nunstab = int(np.sum(abs(beta) >= div * abs(alpha)))
    ns = n - nunstab
    if np.any((abs(alpha) < SMALL) & (abs(beta) < SMALL)):
        return None, None, (-2, -2)

    q = Q.conj().T
    q1, q2 = q[:ns], q[ns:]

    def trimmed_svd(M):
        u, d, vh = svd(M)
        keep = d > SMALL
        r = len(d)
        return u[:, :r][:, keep], d[keep], vh.conj().T[:, :r][:, keep]

    u2, d2, v2 = trimmed_svd(q2 @ pi)        # can eta kill the explosive block?
    u1, d1, v1 = trimmed_svd(q1 @ pi)        # is the choice unique?

    exist = len(d2) >= nunstab
    if v1.shape[1] == 0:
        unique = True
    else:
        loose = v1 - v2 @ v2.conj().T @ v1
        unique = np.sum(svd(loose, compute_uv=False) > SMALL * n) == 0
    eu = (int(exist), int(unique))
    if not exist:
        return None, None, eu

    W = u2 @ np.diag(1 / d2) @ v2.conj().T @ v1 @ np.diag(d1) @ u1.conj().T
    tmat = np.hstack([np.eye(ns), -W.conj().T])
    G0 = np.vstack([tmat @ S,
                    np.hstack([np.zeros((nunstab, ns)), np.eye(nunstab)])])
    G0i = np.linalg.inv(G0)
    G1 = np.real(Z @ (G0i @ np.vstack([tmat @ T, np.zeros((nunstab, n))]))
                 @ Z.conj().T)
    impact = np.real(Z @ (G0i @ np.vstack([tmat @ q @ psi,
                                           np.zeros((nunstab, psi.shape[1]))])))
    return G1, impact, eu
```

Before trusting it, we check it against a model whose solution we can write down.

In the textbook three-equation new Keynesian model

$$\pi_t = \theta \mathbb{E}_t \pi_{t+1} + \kappa x_t, \quad x_t = \mathbb{E}_t x_{t+1} - \sigma(R_t - \mathbb{E}_t \pi_{t+1}) + g_t, \quad R_t = \phi_\pi \pi_t,$$

with g_t an AR(1), the equilibrium is $\pi_t = \pi_g g_t$ and $x_t = x_g g_t$ where (π_g, x_g) solves a two by two linear system, and the equilibrium is unique if and only if $\phi_\pi > 1$.

```
def toy_nk(phi_pi, theta=0.99, kappa=0.1, sigma=1.0, rho=0.8):
    """y = [pi, x, R, g, E pi(+1), E x(+1)]."""
    g0, g1 = np.zeros((6, 6)), np.zeros((6, 6))
    psi, pie = np.zeros((6, 1)), np.zeros((6, 2))
    g0[0, 0], g0[0, 4], g0[0, 1] = 1, -theta, -kappa
    g0[1, 1], g0[1, 5], g0[1, 2], g0[1, 4], g0[1, 3] = 1, -1, sigma, -sigma, -1
    g0[2, 2], g0[2, 0] = 1, -phi_pi
    g0[3, 3], g1[3, 3], psi[3, 0] = 1, rho, 1
    g0[4, 0], g1[4, 4], pie[4, 0] = 1, 1, 1
    g0[5, 1], g1[5, 5], pie[5, 1] = 1, 1, 1
    return g0, g1, psi, pie


for phi in [0.5, 1.5, 3.0]:
    G1, impact, eu = gensys(*toy_nk(phi))
    msg = 'unique stable solution' if eu == (1, 1) else 'indeterminate'
    print(f'phi_pi = {phi}:  eu = {eu}  ({msg})')

theta, kappa, sigma, rho, phi = 0.99, 0.1, 1.0, 0.8, 1.5
M = np.array([[1 - theta * rho, -kappa], [sigma * (phi - rho), 1 - rho]])
exact = np.linalg.solve(M, np.array([0.0, 1.0]))
G1, impact, _ = gensys(*toy_nk(phi))
print(f'\nanalytic (pi_g, x_g) = {np.round(exact, 8)}')
print(f'gensys   (pi_g, x_g) = {np.round(impact[:2, 0], 8)}')
```

```
phi_pi = 0.5:  eu = (1, 0)  (indeterminate)
phi_pi = 1.5:  eu = (1, 1)  (unique stable solution)
phi_pi = 3.0:  eu = (1, 1)  (unique stable solution)

analytic (pi_g, x_g) = [0.89605735 1.86379928]
gensys   (pi_g, x_g) = [0.89605735 1.86379928]
```

The solver reproduces the analytic solution to machine precision and correctly refuses to deliver one when $\phi_\pi < 1$.

148.11 The Sargent-Surico model in canonical form

Now we write (148.7) through (148.13) in the form (148.16).

The vector of variables is

$$y_t = [\pi_t,\ x_t,\ \Delta m_t,\ R_t,\ e_t,\ a_t,\ \chi_t,\ z_t,\ \mathbb{E}_t\pi_{t+1},\ \mathbb{E}_t x_{t+1}]',$$

and the shocks are $\varepsilon_t = [\varepsilon_{et}, \varepsilon_{at}, \varepsilon_{\chi t}, \varepsilon_{zt}, \varepsilon_{mt}]'$.

```
PI, X, DM, R, E, A_, CH, Z, EPI, EX = range(10)
NY = 10


def canonical(p, rule='money'):
    """Matrices g0, g1, psi, pi of equations (SS-NKPC) through (SS-Taylor)."""
    th, api, kap, tau = p['theta'], p['alpha_pi'], p['kappa'], p['tau']
    ax, sig, xi, gam = p['alpha_x'], p['sigma'], p['xi'], p['gamma']
    g0, g1 = np.zeros((NY, NY)), np.zeros((NY, NY))
```

(continues on next page)

(continued from previous page)

```
    psi, pie = np.zeros((NY, 5)), np.zeros((NY, 2))

    # Phillips curve
    g0[0, PI], g0[0, EPI], g0[0, X], g0[0, E] = 1, -th * (1 - api), -kap, 1 / tau
    g1[0, PI] = th * api
    # IS curve
    g0[1, X], g0[1, EX], g0[1, R], g0[1, EPI] = 1, -(1 - ax), sig, -sig
    g0[1, A_] = -sig * (1 - xi) * (1 - p['rho_a'])
    g1[1, X] = ax
    # money demand
    g0[2, DM], g0[2, PI], g0[2, Z] = 1, -1, -1
    for col, coef in [(X, -1 / (sig * gam)), (R, 1 / gam),
                      (CH, -1 / gam), (A_, 1 / gam)]:
        g0[2, col], g1[2, col] = coef, coef
    # policy rule
    if rule == 'money':
        rm = p['rho_m']
        g0[3, DM] = 1
        g0[3, PI], g0[3, X] = -(1 - rm) * p['phi_pi'], -(1 - rm) * p['phi_x']
        g1[3, DM] = rm
    else:
        rr = p['rho_r']
        g0[3, R] = 1
        g0[3, PI], g0[3, X] = -(1 - rr) * p['psi_pi'], -(1 - rr) * p['psi_x']
        g1[3, R] = rr
    psi[3, 4] = 1
    # exogenous shocks
    for row, (v, rho, k) in enumerate([(E, p['rho_e'], 0), (A_, p['rho_a'], 1),
                                       (CH, p['rho_chi'], 2), (Z, 0.0, 3)], 4):
        g0[row, v], g1[row, v], psi[row, k] = 1, rho, 1
    # expectational identities
    g0[8, PI], g1[8, EPI], pie[8, 0] = 1, 1, 1
    g0[9, X], g1[9, EX], pie[9, 1] = 1, 1, 1
    return g0, g1, psi, pie
```

Output growth in (148.10) needs x_{t-1} and a_{t-1}, so the state carries those two lags,

$$S_t = [y_t', \, x_{t-1}, \, a_{t-1}]', \qquad S_t = AS_{t-1} + B\varepsilon_t, \qquad Y_t = CS_t.$$

There is no measurement error, so $D = 0$ and (148.6) applies directly.

The following issue arises here.

The `div` argument of `gensys` splits stable from explosive generalized eigenvalues, and it has to sit a little above one for the split to be numerically reliable.

So a solution that `gensys` reports as unique and stable can still have a root of the transition matrix at or just above one.

Both $h(0)$ and the Kalman filter need covariance stationarity, so we test for it explicitly rather than trusting the flags.

```
def state_space(p, rule='money'):
    """
    Return A, B, C for S[t] = A S[t-1] + B eps[t] and Y[t] = C S[t],
    together with a status string that is 'ok' when the equilibrium is
    unique and covariance stationary.
    """
    G1, impact, eu = gensys(*canonical(p, rule))
    if eu[0] != 1:
```

(continues on next page)

(continued from previous page)

```
        return None, None, None, 'no stable solution'
    if eu[1] != 1:
        return None, None, None, 'indeterminate'
    sd = np.diag([p['sig_e'], p['sig_a'], p['sig_chi'], p['sig_z'], p['sig_m']])
    n = NY + 2
    A, B = np.zeros((n, n)), np.zeros((n, 5))
    A[:NY, :NY] = G1
    A[NY, X], A[NY + 1, A_] = 1, 1
    B[:NY] = impact @ sd
    C = np.zeros((4, n))                                # [dm, pi, R, dy]
    C[0, DM], C[1, PI], C[2, R] = 1, 1, 1
    C[3, X], C[3, NY] = 1, -1
    C[3, A_], C[3, NY + 1] = p['xi'], -p['xi']
    C[3, Z] = 1
    if np.max(np.abs(np.linalg.eigvals(A))) > 1 - 1e-9:
        return None, None, None, 'unit or explosive root'
    return A, B, C, 'ok'


def h_zero(p, rule='money'):
    """(h_pi,dm(0), h_R,dm(0)) implied by the model at parameters p."""
    A, B, C, status = state_space(p, rule)
    if status != 'ok':
        return np.nan, np.nan
    return (h_zero_from_state_space(A, B, C, 1, 0),
            h_zero_from_state_space(A, B, C, 2, 0))
```

This test matters.

Without it, a point like $\phi_\pi = 1$, $\phi_x = 0$ returns $h_{\pi,\Delta m}(0) = h_{R,\Delta m}(0) = 1.000$, an apparently perfect confirmation of both of Lucas's illustrations.

The transition matrix there has a root of exactly one, the spectral density at frequency zero does not exist, and those two numbers come from inverting a matrix whose condition number is about 10^{16}.

That point sits in the upper right corner of the range that Figure 6 of the paper plots.

We show below that at the paper's own Table 2 posterior means a whole strip of that range has no covariance stationary equilibrium.

We can now check our implementation against the paper.

At the posterior means of Table 2, Sargent and Surico [2011] report implied values $h_{\pi,\Delta m}(0) = 1.0068$ and $h_{R,\Delta m}(0) = 0.8163$.

```
PAPER = dict(theta=0.9901, alpha_pi=0.8815, kappa=0.0324, tau=3.5126,
             alpha_x=0.4775, sigma=0.0997, xi=3.0319, gamma=3.8128,
             phi_pi=0.2312, phi_x=-0.1971, rho_m=0.7428, rho_e=0.5645,
             rho_a=0.9241, rho_chi=0.5024, sig_e=1.0922, sig_a=0.7226,
             sig_chi=0.2388, sig_z=1.5845, sig_m=1.1457)

print('h at the posterior means of Table 2: %.4f, %.4f' % h_zero(PAPER))
print('reported in Table 2:                 1.0068, 0.8163')
```

```
h at the posterior means of Table 2: 1.0093, 0.8060
reported in Table 2:                 1.0068, 0.8163
```

Our independent implementation lands on the paper's numbers.

That gives us confidence that we have read equations (148.7) through (148.12) the way their authors intended.

148.11.1 Checking the three claims

Now we verify the three observations of the previous section.

```
print('scaling tau and sigma_e together leaves h(0) untouched:')
for f in [0.5, 1.0, 2.0, 8.0]:
    q = dict(PAPER, tau=PAPER['tau'] * f, sig_e=PAPER['sig_e'] * f)
    print(f'   tau, sigma_e scaled by {f:4.1f}:  h_pi = {h_zero(q)[0]:.10f}')

print('\nwith phi_pi = phi_x = 0, h_pi is exactly one for any other parameters:')
rng = np.random.default_rng(0)
for i in range(5):
    q = dict(PAPER, phi_pi=0.0, phi_x=0.0,
             kappa=rng.uniform(0.01, 0.3), alpha_pi=rng.uniform(0.1, 0.9),
             alpha_x=rng.uniform(0.1, 0.9), sigma=rng.uniform(0.05, 0.5),
             gamma=rng.uniform(1, 8), rho_m=rng.uniform(0, 0.95),
             sig_z=rng.uniform(0.2, 3.0), xi=rng.uniform(0.5, 5.0))
    print(f'   draw {i}:  h_pi = {h_zero(q)[0]:.12f}   h_R = {h_zero(q)[1]:.4f}')

print('\ndeterminacy on both sides of alpha_x = 0.5:')
for ax in [0.30, 0.45, 0.499, 0.501, 0.60, 0.80]:
    print(f'   alpha_x = {ax:5.3f}:  {state_space(dict(PAPER, alpha_x=ax))[3]}')
```

```
scaling tau and sigma_e together leaves h(0) untouched:
   tau, sigma_e scaled by  0.5:  h_pi = 1.0092936992
   tau, sigma_e scaled by  1.0:  h_pi = 1.0092936992
   tau, sigma_e scaled by  2.0:  h_pi = 1.0092936992
   tau, sigma_e scaled by  8.0:  h_pi = 1.0092936992

with phi_pi = phi_x = 0, h_pi is exactly one for any other parameters:
   draw 0:  h_pi = 1.000000000000   h_R = 0.4849
   draw 1:  h_pi = 1.000000000000   h_R = -0.0141
   draw 2:  h_pi = 1.000000000000   h_R = 0.2936
   draw 3:  h_pi = 1.000000000000   h_R = 0.7060
   draw 4:  h_pi = 1.000000000000   h_R = 0.7182

determinacy on both sides of alpha_x = 0.5:
   alpha_x = 0.300:  ok
   alpha_x = 0.450:  ok
   alpha_x = 0.499:  ok
   alpha_x = 0.501:  ok
   alpha_x = 0.600:  ok
   alpha_x = 0.800:  ok
```

All three claims hold.

The likelihood is invariant to scaling (τ, σ_e), the first unit slope is an identity under exogenous money growth, and the model is determinate on both sides of $\alpha_x = 1/2$.

Notice also that $h_{R,\Delta m}(0)$ moves all over the place across those same draws, which is the asymmetry between the two illustrations that *Proposition 148.9.1* predicts.

Finally, here is the strip of the Figure 6 policy grid on which, at the paper's own posterior means, no covariance stationary equilibrium exists.

```
for phi_x in [0.0, -0.25, -0.5]:
    bad = [round(g, 2) for g in np.linspace(-3, 1, 41)
           if state_space(dict(PAPER, phi_pi=g, phi_x=phi_x))[3] != 'ok']
    print(f'phi_x = {phi_x:5.2f}: no stationary equilibrium at phi_pi in '
          f'{bad if bad else "none of the grid"}')
```

```
phi_x =  0.00: no stationary equilibrium at phi_pi in [np.float64(-3.0), np.
↪float64(-2.9), np.float64(-2.8), np.float64(-2.7), np.float64(-2.6), np.float64(-
↪2.5), np.float64(-2.4), np.float64(-2.3), np.float64(-2.2), np.float64(-2.1), np.
↪float64(-2.0), np.float64(-1.9), np.float64(-1.8), np.float64(-1.7), np.float64(-
↪1.6), np.float64(1.0)]
phi_x = -0.25: no stationary equilibrium at phi_pi in none of the grid
phi_x = -0.50: no stationary equilibrium at phi_pi in none of the grid
```

The strip is far from the estimated policy rule, so the paper's conclusions are not at stake.

But an $h(0)$ reported over it would be a number computed from a spectral density that does not exist.

148.12 Bayesian estimation

We estimate the money growth rule version on 1960:I-1983:IV, the paper's sample.

The four observables are demeaned, since the model is written in deviations from a steady state.

```
mask = (data.index.year >= 1960) & (data.index.year <= 1983)
Y_est = data.loc[mask, VARS].values
Y_est = Y_est - Y_est.mean(0)
print(f'estimation sample: {mask.sum()} quarters, '
      f'{data.index[mask][0].date()} to {data.index[mask][-1].date()}')
```

```
estimation sample: 96 quarters, 1960-01-01 to 1983-10-01
```

148.12.1 The likelihood

Given (A, B, C) the likelihood follows from the Kalman filter.

We initialize the state at its unconditional mean and covariance, the latter solving the discrete Lyapunov equation $P = APA' + BB'$.

```
def loglik(p, Y, rule='money'):
    """Kalman filter log likelihood of Y (T x 4) at parameters p."""
    A, B, C, status = state_space(p, rule)
    if status != 'ok':
        return -np.inf
    Q = B @ B.T
    try:
        P = solve_discrete_lyapunov(A, Q)
    except Exception:
        return -np.inf
    s = np.zeros(A.shape[0])
    ll, const = 0.0, Y.shape[1] * np.log(2 * np.pi)
    for t in range(Y.shape[0]):
        s = A @ s
```

(continues on next page)

(continued from previous page)

```
        P = A @ P @ A.T + Q
        v = Y[t] - C @ s                                # forecast error
        PCt = P @ C.T
        F = C @ PCt                                     # its covariance
        try:
            L = np.linalg.cholesky(F)
        except np.linalg.LinAlgError:
            return -np.inf
        u = np.linalg.solve(L, v)
        ll -= 0.5 * (const + 2 * np.sum(np.log(np.diag(L))) + u @ u)
        K = np.linalg.solve(F, PCt.T).T                 # Kalman gain
        s, P = s + K @ v, P - K @ PCt.T
    return ll if np.isfinite(ll) else -np.inf
```

148.12.2 Priors

We use the prior means and standard deviations of Table 2, matching moments to pick the parameters of each family.

```
def beta_prior(m, s):
    nu = m * (1 - m) / s ** 2 - 1
    return stats.beta(m * nu, (1 - m) * nu)


def gamma_prior(m, s):
    return stats.gamma(m ** 2 / s ** 2, scale=s ** 2 / m)


def invgamma_prior(m, s):
    a = m ** 2 / s ** 2 + 2
    return stats.invgamma(a, scale=m * (a - 1))


PRIOR, SUPPORT = {}, {}
for n, (m, s) in [('theta', (0.99, 0.005)), ('alpha_pi', (0.5, 0.2)),
                  ('alpha_x', (0.5, 0.2)), ('rho_m', (0.5, 0.05)),
                  ('rho_e', (0.5, 0.1)), ('rho_a', (0.5, 0.1)),
                  ('rho_chi', (0.5, 0.1))]:
    PRIOR[n], SUPPORT[n] = beta_prior(m, s), (1e-6, 1 - 1e-6)
for n, (m, s) in [('kappa', (0.3, 0.1)), ('sigma', (0.1, 0.05)),
                  ('xi', (2.0, 1.0)), ('gamma', (4.0, 1.0))]:
    PRIOR[n], SUPPORT[n] = gamma_prior(m, s), (1e-8, np.inf)
for n in ['phi_pi', 'phi_x']:
    PRIOR[n], SUPPORT[n] = stats.norm(0, 0.5), (-np.inf, np.inf)
for n in ['sig_e', 'sig_a', 'sig_chi', 'sig_z', 'sig_m']:
    PRIOR[n], SUPPORT[n] = invgamma_prior(0.3, 1.0), (1e-8, np.inf)

FREE = ['theta', 'alpha_pi', 'kappa', 'alpha_x', 'sigma', 'xi', 'gamma',
        'phi_pi', 'phi_x', 'rho_m', 'rho_e', 'rho_a', 'rho_chi',
        'sig_e', 'sig_a', 'sig_chi', 'sig_z', 'sig_m']
TAU = 4.0                       # not identified; fixed at its prior mean


def unpack(v):
    return dict({n: float(x) for n, x in zip(FREE, v)}, tau=TAU)
```

(continues on next page)

(continued from previous page)

```
def log_post(v, Y):
    lp = 0.0
    for n, x in zip(FREE, v):
        lo, hi = SUPPORT[n]
        if not lo < x < hi:
            return -np.inf
        lp += PRIOR[n].logpdf(x)
    return lp + loglik(unpack(v), Y)
```

148.12.3 The posterior mode

We start the search at the paper's posterior means, use Powell's derivative free method, and polish with a gradient method.

```
v_start = np.array([PAPER[n] for n in FREE])
neg_log_post = lambda v: -log_post(v, Y_est)

res = optimize.minimize(neg_log_post, v_start, method='Powell',
                        options=dict(maxiter=20000, maxfev=20000))
res = optimize.minimize(neg_log_post, res.x, method='L-BFGS-B',
                        bounds=[SUPPORT[n] for n in FREE])
v_mode = res.x
print(f'log posterior at the paper\'s means: {log_post(v_start, Y_est):10.3f}')
print(f'log posterior at the mode:          {-res.fun:10.3f}')
print('h(0) at the mode: %.4f, %.4f' % h_zero(unpack(v_mode)))
```

```
log posterior at the paper's means:   -531.052
log posterior at the mode:            -320.204
h(0) at the mode: 0.9765, 0.7694
```

Note

Starting the search at the paper's estimates finds the mode nearest to them.

The Hamiltonian Monte Carlo section at the end of this lecture discovers that this is a *local* mode, and that the posterior has a second one with higher density.

148.12.4 Random walk Metropolis-Hastings

The standard proposal covariance is the inverse Hessian of the negative log posterior at the mode, which we compute by finite differences.

```
def numerical_hessian(f, v, rel=1e-4):
    n = len(v)
    h = rel * np.maximum(np.abs(v), 1e-2)
    H = np.zeros((n, n))
    for i in range(n):
        for j in range(i, n):
            ei, ej = np.zeros(n), np.zeros(n)
            ei[i], ej[j] = h[i], h[j]
            H[i, j] = H[j, i] = (f(v + ei + ej) - f(v + ei - ej)
```

(continues on next page)

(continued from previous page)

```
                                - f(v - ei + ej) + f(v - ei - ej)) / (4 * h[i] *␣
↪h[j])
    return H


H = numerical_hessian(neg_log_post, v_mode)
Sigma_prop = np.linalg.inv(H)
print('proposal covariance is positive definite:',
      np.all(np.linalg.eigvalsh(H) > 0))
```

```
proposal covariance is positive definite: True
```

```
def rwmh(Y, v0, Sigma, n_draws, c=0.45, seed=42):
    """Random walk Metropolis-Hastings in the natural parameter space."""
    rng = np.random.default_rng(seed)
    L = np.linalg.cholesky(Sigma)
    v, lp = v0.copy(), log_post(v0, Y)
    draws, n_acc = np.empty((n_draws, len(v))), 0
    for i in range(n_draws):
        cand = v + c * (L @ rng.standard_normal(len(v)))
        lp_cand = log_post(cand, Y)
        if np.log(rng.random()) < lp_cand - lp:
            v, lp, n_acc = cand, lp_cand, n_acc + 1
        draws[i] = v
    return draws, n_acc / n_draws
```

Draws that fall outside the support of a prior get $-\infty$ and are rejected, so we can work in the natural parameter space and skip transformations.

The chain below is short enough to run while you read; a serious application would use many more draws.

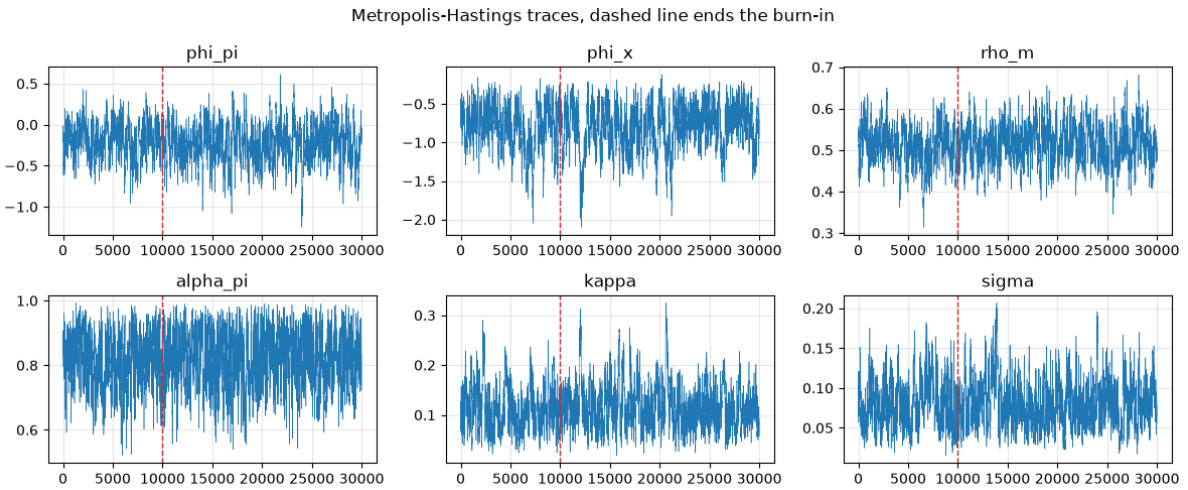
```
N_DRAWS, BURN = 30_000, 10_000

t0 = time.time()
draws, acc_rate = rwmh(Y_est, v_mode, Sigma_prop, N_DRAWS)
rwmh_seconds = time.time() - t0

kept = draws[BURN::5]
print(f'acceptance rate {acc_rate:.3f},  {len(kept)} retained draws, '
      f'{rwmh_seconds:.0f} seconds')
```

```
acceptance rate 0.283,  4000 retained draws, 278 seconds
```

```
fig, axes = plt.subplots(2, 3, figsize=(12, 5))
for ax, n in zip(axes.flat, ['phi_pi', 'phi_x', 'rho_m',
                             'alpha_pi', 'kappa', 'sigma']):
    j = FREE.index(n)
    ax.plot(draws[:, j], lw=0.4, color='C0')
    ax.axvline(BURN, color='C3', ls='--', lw=1)
    ax.set_title(n)
fig.suptitle('Metropolis-Hastings traces, dashed line ends the burn-in')
plt.tight_layout()
plt.show()
```




```
summary = pd.DataFrame({
    'prior mean': [PRIOR[n].mean() for n in FREE],
    'post. mean': kept.mean(0),
    '5th': np.percentile(kept, 5, axis=0),
    '95th': np.percentile(kept, 95, axis=0),
    'paper': [PAPER[n] for n in FREE]}, index=FREE)
summary.round(4)
```

	prior mean	post. mean	5th	95th	paper
theta	0.99	0.9887	0.9787	0.9960	0.9901
alpha_pi	0.50	0.8162	0.6641	0.9455	0.8815
kappa	0.30	0.1174	0.0581	0.1957	0.0324
alpha_x	0.50	0.3799	0.3026	0.4313	0.4775
sigma	0.10	0.0820	0.0396	0.1345	0.0997
xi	2.00	3.5793	2.1664	5.3518	3.0319
gamma	4.00	2.0430	1.5863	2.6046	3.8128
phi_pi	0.00	-0.2361	-0.6368	0.1261	0.2312
phi_x	0.00	-0.7989	-1.3512	-0.3665	-0.1971
rho_m	0.50	0.5162	0.4381	0.5949	0.7428
rho_e	0.50	0.2580	0.1561	0.3750	0.5645
rho_a	0.50	0.8499	0.7584	0.9133	0.9241
rho_chi	0.50	0.5169	0.3609	0.6777	0.5024
sig_e	0.30	1.1648	0.8815	1.4654	1.0922
sig_a	0.30	0.2468	0.1511	0.3717	0.7226
sig_chi	0.30	0.1408	0.0589	0.2533	0.2388
sig_z	0.30	0.7131	0.5513	0.9239	1.5845
sig_m	0.30	0.7143	0.6161	0.8311	1.1457

Our estimates are recognizably those of Table 2, with the differences one expects from a different interest rate series and different data vintages.

The Phillips curve is strongly backward looking and flat, the IS curve much less so.

Most important for what follows, the money growth rule responds weakly to inflation, with a posterior for ϕ_π that straddles zero, and it carries substantial smoothing.

That is the paper's central reading of the pre-1984 regime: the Federal Reserve put persistent, nearly exogenous, movement into money growth.

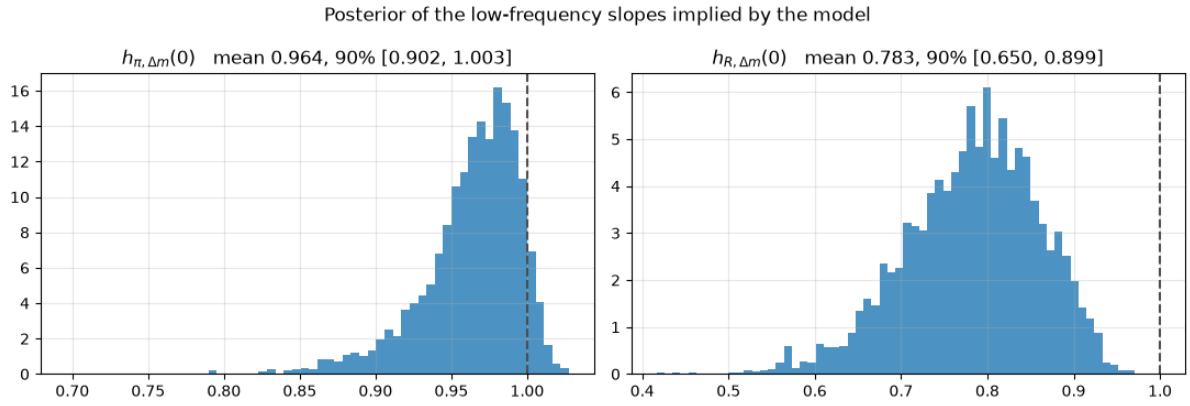
By *Proposition 148.9.1* that is precisely the configuration in which Lucas's first illustration holds, which is also why the paper's own posterior for $h_{\pi,\Delta m}(0)$ in Table 2 is so remarkably tight around one.

148.12.5 The posterior distribution of $h(0)$

Every draw of the structural parameters implies a pair of low-frequency slopes.

```
h_draws = np.array([h_zero(unpack(v)) for v in kept])

fig, axes = plt.subplots(1, 2, figsize=(11, 3.8))
for ax, j, ttl in zip(axes, [0, 1],
                      [r'$h_{\pi,\Delta m}(0)$', r'$h_{R,\Delta m}(0)$']):
    ax.hist(h_draws[:, j], bins=60, color='C0', alpha=0.8, density=True)
    ax.axvline(1, color='0.3', ls='--', lw=1.5)
    ax.set_title(f'{ttl}   mean {h_draws[:, j].mean():.3f}, '
                 f'90% [{np.percentile(h_draws[:, j], 5):.3f}, '
                 f'{np.percentile(h_draws[:, j], 95):.3f}]')
fig.suptitle('Posterior of the low-frequency slopes implied by the model')
plt.tight_layout()
plt.show()
```



The estimated model reproduces Lucas's two illustrations over a sample much like his.

Compare these with the VAR estimates for 1960-1983 that we computed earlier, and with the paper's reported posterior means of 1.0068 and 0.8163.

148.13 How monetary policy moves the low-frequency slopes

Now for the paper's main experiment.

We lock every structural parameter at its posterior mean, vary only the two policy coefficients, and recompute $h(0)$ at each point.

```
p_bar = unpack(kept.mean(0))


def h_grid(p, rule, k1, k2, g1_vals, g2_vals):
    H1 = np.full((len(g2_vals), len(g1_vals)), np.nan)
    H2 = np.full_like(H1, np.nan)
    for i, b in enumerate(g2_vals):
        for j, a in enumerate(g1_vals):
            H1[i, j], H2[i, j] = h_zero(dict(p, **{k1: a, k2: b}), rule)
    return H1, H2
```

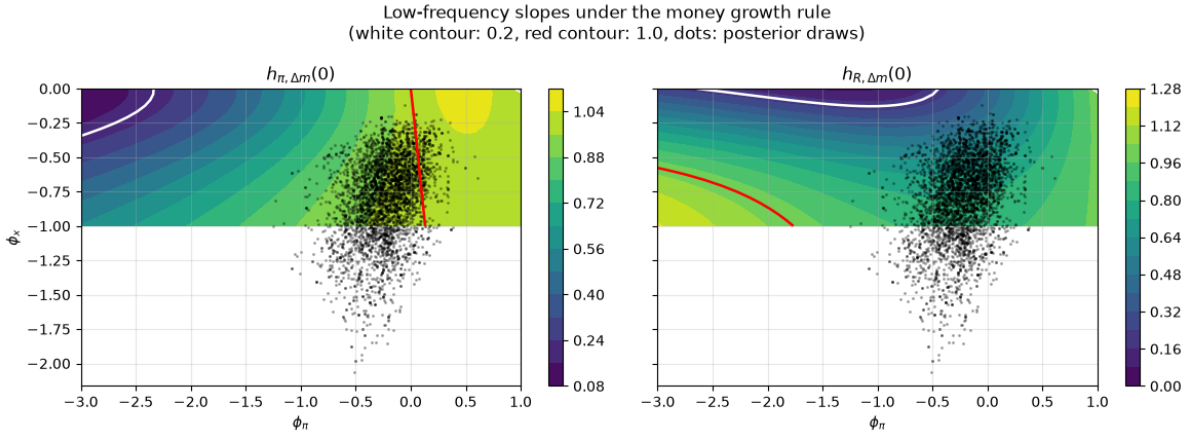
(continues on next page)

(continued from previous page)

```
phi_pi_grid = np.linspace(-3, 1, 49)
phi_x_grid = np.linspace(-1, 0, 25)
Hpi, HR = h_grid(p_bar, 'money', 'phi_pi', 'phi_x', phi_pi_grid, phi_x_grid)
```

```
def contour_panel(g1_vals, g2_vals, H1, H2, xlab, ylab, suptitle, scatter=None):
    fig, axes = plt.subplots(1, 2, figsize=(12, 4.4), sharey=True)
    for ax, Hm, ttl in zip(axes, [H1, H2],
                           [r'$h_{\pi,\Delta m}(0)$', r'$h_{R,\Delta m}(0)$']):
        cs = ax.contourf(g1_vals, g2_vals, Hm, levels=14, cmap='viridis')
        ax.contour(g1_vals, g2_vals, Hm, levels=[0.2, 1.0],
                   colors=['white', 'red'], linewidths=1.8)
        ax.contourf(g1_vals, g2_vals, np.isnan(Hm).astype(float),
                    levels=[0.5, 1.5], colors=['0.85'])
        if scatter is not None:
            ax.scatter(*scatter, s=1.5, color='k', alpha=0.25)
        fig.colorbar(cs, ax=ax)
        ax.set(title=ttl, xlabel=xlab)
    axes[0].set_ylabel(ylab)
    fig.suptitle(suptitle)
    plt.tight_layout()
    plt.show()


contour_panel(phi_pi_grid, phi_x_grid, Hpi, HR, r'$\phi_\pi$', r'$\phi_x$',
              'Low-frequency slopes under the money growth rule\n'
              '(white contour: 0.2, red contour: 1.0, dots: posterior draws)',
              scatter=(kept[:, FREE.index('phi_pi')],
                       kept[:, FREE.index('phi_x')]))
```



Three things come through, and they are the paper's three findings.

The low-frequency slopes are *not* policy invariant: moving (ϕ_π, ϕ_x) alone sweeps $h_{\pi,\Delta m}(0)$ across most of the range that the rolling VARs displayed.

The cloud of posterior draws sits in the region where both slopes are near one, which is where the 1960-1983 data put us.

And a much more aggressive anti-inflation stance, meaning a large *negative* ϕ_π so that money growth is cut hard when inflation rises, moves both slopes far below one, with the smallest values in the corner where ϕ_x is also near zero.

The red contour marks $h(0) = 1$ and, as *Proposition 148.9.1* requires, it passes exactly through $\phi_\pi = \phi_x = 0$.

Any gray region would be one with no unique covariance stationary equilibrium, where $h(0)$ does not exist.

Here is where the slopes bottom out, a slice through the grid at the posterior mean of ϕ_x, and the fraction of the grid the stationarity test removes.

```
print(f'no stationary equilibrium at {np.isnan(Hpi).mean():.0%} of grid points')
for name, Hm in [('h_pi', Hpi), ('h_R', HR)]:
    i, j = np.unravel_index(np.nanargmin(Hm), Hm.shape)
    print(f'   smallest {name:5s} = {Hm[i, j]:6.3f} at '
          f'phi_pi = {phi_pi_grid[j]:5.2f}, phi_x = {phi_x_grid[i]:5.2f}')

row = np.argmin(np.abs(phi_x_grid - p_bar['phi_x']))
print(f'\nslice at phi_x = {phi_x_grid[row]:.3f}')
for j in range(0, len(phi_pi_grid), 6):
    print(f'   phi_pi = {phi_pi_grid[j]:6.2f}:  '
          f'h_pi = {Hpi[row, j]:7.3f},  h_R = {HR[row, j]:7.3f}')
```

```
no stationary equilibrium at 0% of grid points
   smallest h_pi  =  0.097 at phi_pi = -3.00, phi_x =  0.00
   smallest h_R   =  0.059 at phi_pi = -1.25, phi_x =  0.00

slice at phi_x = -0.792
   phi_pi =  -3.00:  h_pi =   0.463,  h_R =   1.137
   phi_pi =  -2.50:  h_pi =   0.550,  h_R =   1.056
   phi_pi =  -2.00:  h_pi =   0.646,  h_R =   0.972
   phi_pi =  -1.50:  h_pi =   0.748,  h_R =   0.892
   phi_pi =  -1.00:  h_pi =   0.847,  h_R =   0.828
   phi_pi =  -0.50:  h_pi =   0.933,  h_R =   0.796
   phi_pi =   0.00:  h_pi =   0.992,  h_R =   0.808
   phi_pi =   0.50:  h_pi =   1.017,  h_R =   0.870
   phi_pi =   1.00:  h_pi =   1.005,  h_R =   0.980
```

148.13.1 The same exercise with a Taylor rule

The paper repeats the experiment with the interest rate rule (148.13).

We keep the smoothing coefficient and the policy shock standard deviation at their money rule estimates and vary (ψ_π, ψ_x).

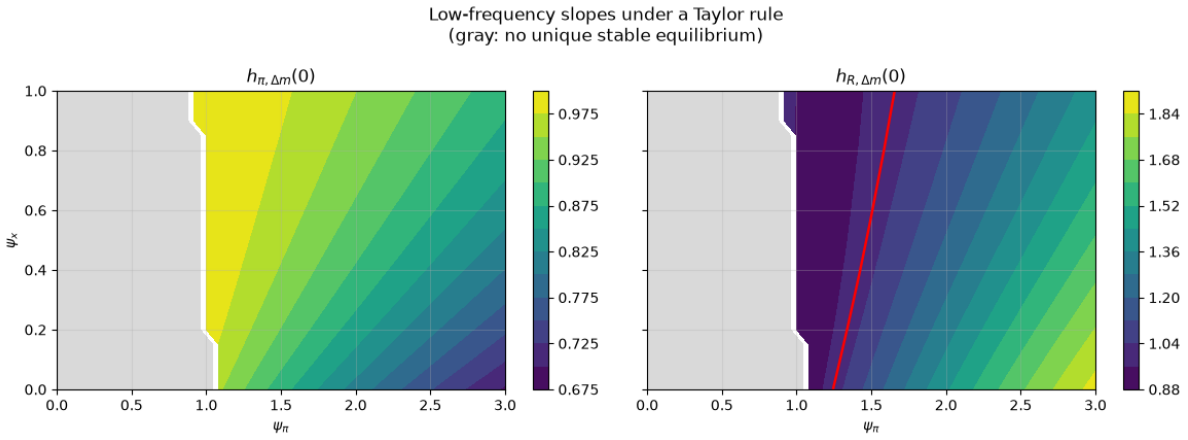
Where the Taylor rule leaves the equilibrium indeterminate we simply record it, shading those regions in gray; the paper instead selects the orthogonality solution of Lubik and Schorfheide [2004].

```
p_taylor = dict(p_bar, rho_r=p_bar['rho_m'], psi_pi=1.5, psi_x=0.5)
psi_pi_grid = np.linspace(0, 3, 37)
psi_x_grid = np.linspace(0, 1, 21)
Tpi, TR = h_grid(p_taylor, 'taylor', 'psi_pi', 'psi_x',
                 psi_pi_grid, psi_x_grid)

print(f'indeterminate or nonexistent at {np.isnan(Tpi).mean():.0%} of grid points')
print(f'h_pi ranges over [{np.nanmin(Tpi):.2f}, {np.nanmax(Tpi):.2f}]')
print(f'h_R  ranges over [{np.nanmin(TR):.2f}, {np.nanmax(TR):.2f}]')

contour_panel(psi_pi_grid, psi_x_grid, Tpi, TR, r'$\psi_\pi$', r'$\psi_x$',
              'Low-frequency slopes under a Taylor rule\n'
              '(gray: no unique stable equilibrium)')
```

```
indeterminate or nonexistent at 33% of grid points
h_pi ranges over [0.70, 1.00]
h_R  ranges over [0.89, 1.87]
```



The dependence on policy survives, but it is much weaker.

Over the determinate region the Taylor rule cannot drive either slope anywhere near zero.

This is the paper's finding that the interest rate rule is "less able to replicate the estimates."

The decomposition (148.15) supplies a reason.

Driving $h_{\pi,\Delta m}(0)$ toward zero requires making the ratio $S_{z,\pi+z}(0)/S_{\pi+z}(0)$ close to one, which in practice means squeezing the zero-frequency power of inflation until $\pi_t + z_t$ is dominated by technology.

A money growth rule with a large negative ϕ_π does exactly that, because the central bank is choosing the left side of (148.14) directly.

A Taylor rule instead sets the interest rate and lets money growth be whatever money demand implies, so it never gets much purchase on low-frequency inflation relative to money growth, and $h_{\pi,\Delta m}(0)$ stays near one.

148.14 Estimating with Hamiltonian Monte Carlo

The random walk sampler that produced our posterior is the traditional workhorse of DSGE estimation.

It is also wasteful, and we can measure how wasteful.

This section takes the Sargent-Surico model as a guinea pig for a modern alternative.

148.14.1 What the random walk delivers

A correlated chain of length N is worth fewer than N independent draws, and the **effective sample size** says how many.

```
import arviz as az

rwmh_idata = az.from_dict(
    {'posterior': {n: kept[:, j][None, :] for j, n in enumerate(FREE)}})
rwmh_summary = az.summary(rwmh_idata, var_names=FREE)
print(rwmh_summary[['mean', 'sd', 'ess_bulk']].to_string())
print(f'\n{len(kept)} retained draws out of {N_DRAWS}, in {rwmh_seconds:.0f} seconds')
print(f'smallest effective sample size = {rwmh_summary["ess_bulk"].min():.0f}')
```

```
            mean      sd ess_bulk
theta     0.9887  0.0053      197
alpha_pi   0.816   0.085      415
kappa      0.117   0.045      195
alpha_x     0.38   0.039      112
sigma      0.082   0.029      165
xi           3.6       1      111
gamma       2.04     0.3      125
phi_pi     -0.24    0.24      220
phi_x       -0.8    0.31       91
rho_m      0.516   0.047      221
rho_e      0.258   0.065      328
rho_a       0.85   0.047       53
rho_chi     0.52   0.096      307
sig_e       1.16    0.17      279
sig_a       0.25   0.068       93
sig_chi    0.141   0.061      200
sig_z       0.71    0.12       96
sig_m       0.71   0.065      149

4000 retained draws out of 30000, in 278 seconds
smallest effective sample size = 53
```

Tens of thousands of iterations buy us on the order of a hundred independent draws for the worst-mixing parameter.

The reason is visible in the proposal.

A random walk perturbs all eighteen parameters in a direction that knows nothing about the local shape of the posterior, so the step size has to respect the most tightly constrained direction while the loosest direction is explored at that same crawl.

Our posterior is badly scaled.

```
w = np.linalg.eigvalsh(H)
print(f'condition number of the posterior Hessian     {w.max() / w.min():10.3g}')
print(f'ratio of longest to shortest posterior scale  {np.sqrt(w.max() / w.min()):10.
↪0f}')
```

```
condition number of the posterior Hessian        4.1e+04
ratio of longest to shortest posterior scale         202
```

148.14.2 Why Hamiltonian methods are awkward for DSGE models

Hamiltonian Monte Carlo [Duane *et al.*, 1987, Neal, 2011] treats the parameter vector as the position of a particle, gives it a random momentum, and simulates Hamiltonian dynamics in which the negative log posterior plays the role of potential energy.

Trajectories follow the contours of the distribution instead of blundering across them, so a proposal can travel a long way and still be accepted.

The **No-U-Turn sampler** [Hoffman and Gelman, 2014] removes the need to tune the trajectory length by extending each trajectory until it starts to double back on itself.

Betancourt [2017] gives a conceptual introduction.

The price of admission is the gradient $\nabla_\theta \log p(\theta \mid Y)$.

That is where DSGE models become awkward.

Our likelihood runs through `gensys`, whose first act is an *ordered* generalized Schur decomposition: it computes generalized eigenvalues, **sorts** them by modulus, and reorders the decomposition to match.

Sorting is a discrete operation, and no automatic differentiation library propagates derivatives through it.

This, rather than any statistical objection, is the reason Hamiltonian methods remain uncommon in DSGE estimation.

148.14.3 A differentiable solver

The obstacle is the *algorithm*, not the model, so we change algorithms.

Write the equilibrium conditions as

$$F_A \, \mathbb{E}_t y_{t+1} + F_B \, y_t + F_C \, y_{t-1} + F_E \, \varepsilon_t = 0, \tag{148.17}$$

where

$$y_t = [\pi_t, \; x_t, \; \Delta m_t, \; R_t, \; e_t, \; a_t, \; \chi_t, \; z_t, \; u_t]'$$

and $u_t = \varepsilon_{mt}$ carries the money rule shock.

The auxiliary expectation variables that `gensys` needs are not required here.

Conjecture a solution

$$y_t = G \, y_{t-1} + \Theta \, \varepsilon_t. \tag{148.18}$$

Then $\mathbb{E}_t y_{t+1} = G y_t$, and (148.17) becomes

$$(F_A G + F_B) \, y_t + F_C \, y_{t-1} + F_E \, \varepsilon_t = 0.$$

Solving for y_t and matching coefficients with (148.18) gives a fixed point in G alone,

$$G = -(F_A G + F_B)^{-1} F_C, \qquad \Theta = -(F_A G + F_B)^{-1} F_E. \tag{148.19}$$

Iterating (148.19) from $G = 0$ involves nothing but matrix products and linear solves, every one of them differentiable.

```
import jax
import jax.numpy as jnp
import numpyro
import numpyro.distributions as dist
from jax import lax
from numpyro.infer import MCMC, NUTS

jax.config.update('jax_enable_x64', True)
jax.config.update('jax_platform_name', 'cpu')

U, NJ = 8, 9          # y = [pi, x, dm, R, e, a, chi, z, u]


def structural(p):
    """Return F_A, F_B, F_C, F_E of equation (SS-structural)."""
    th, api, kap = p['theta'], p['alpha_pi'], p['kappa']
    ax, sig, xi, gam = p['alpha_x'], p['sigma'], p['xi'], p['gamma']
    rm, ra = p['rho_m'], p['rho_a']
    FA = jnp.zeros((NJ, NJ)); FB = jnp.zeros((NJ, NJ))
    FC = jnp.zeros((NJ, NJ)); FE = jnp.zeros((NJ, 5))
```

(continues on next page)

(continued from previous page)

```
    FA = FA.at[0, PI].set(-th * (1 - api))                          # Phillips curve
    FB = FB.at[0, PI].set(1).at[0, X].set(-kap).at[0, E].set(1 / TAU)
    FC = FC.at[0, PI].set(-th * api)

    FA = FA.at[1, X].set(-(1 - ax)).at[1, PI].set(-sig)             # IS curve
    FB = (FB.at[1, X].set(1).at[1, R].set(sig)
            .at[1, A_].set(-sig * (1 - xi) * (1 - ra)))
    FC = FC.at[1, X].set(-ax)

    FB = FB.at[2, DM].set(1).at[2, PI].set(-1).at[2, Z].set(-1)     # money demand
    FB = (FB.at[2, X].set(-1 / (sig * gam)).at[2, R].set(1 / gam)
            .at[2, CH].set(-1 / gam).at[2, A_].set(1 / gam))
    FC = (FC.at[2, X].set(1 / (sig * gam)).at[2, R].set(-1 / gam)
            .at[2, CH].set(1 / gam).at[2, A_].set(-1 / gam))

    FB = (FB.at[3, DM].set(1).at[3, U].set(-1)                      # policy rule
            .at[3, PI].set(-(1 - rm) * p['phi_pi'])
            .at[3, X].set(-(1 - rm) * p['phi_x']))
    FC = FC.at[3, DM].set(-rm)

    for row, (v, rho, k) in enumerate([(E, p['rho_e'], 0), (A_, ra, 1),
                                       (CH, p['rho_chi'], 2), (Z, 0.0, 3)], 4):
        FB = FB.at[row, v].set(1)
        FC = FC.at[row, v].set(-rho)
        FE = FE.at[row, k].set(-1)
    FB = FB.at[8, U].set(1)
    FE = FE.at[8, 4].set(-1)
    return FA, FB, FC, FE


def solve_fixed_point(p, n_iter=60):
    """Iterate (SS-fixedpoint) to convergence; differentiable throughout."""
    FA, FB, FC, FE = structural(p)

    def step(G, _):
        return -jnp.linalg.solve(FA @ G + FB, FC), None

    G, _ = lax.scan(step, jnp.zeros((NJ, NJ)), None, length=n_iter)
    return G, -jnp.linalg.solve(FA @ G + FB, FE)
```

Before trusting it we check it against `gensys` at the paper's posterior means.

The two algorithms share no code, so agreement to machine precision is a real test.

```
p_check = {n: float(PAPER[n]) for n in FREE}
G_fp, Theta_fp = solve_fixed_point(p_check)

A_gen, B_gen, C_gen, _ = state_space(dict(p_check, tau=TAU))
idx = [PI, X, DM, R, E, A_, CH, Z]
sd_check = np.array([p_check[n] for n in
                     ['sig_e', 'sig_a', 'sig_chi', 'sig_z', 'sig_m']])

print('fixed point versus gensys')
print(f'   transition matrix  max abs difference  '
      f'{np.abs(np.asarray(G_fp)[np.ix_(idx, idx)] - A_gen[:10, :10][np.ix_(idx,␣
 ↪idx)]).max():.2e}')
print(f'   shock loadings     max abs difference  '
```

(continues on next page)

(continued from previous page)

```
    f'{np.abs(np.asarray(Theta_fp * sd_check)[idx] - B_gen[:10][idx]).max():.2e}')
```

```
fixed point versus gensys
   transition matrix  max abs difference  1.20e-14
   shock loadings     max abs difference  3.50e-15
```

Note

The fixed point converges to the stable solution when there is one, but unlike `gensys` it returns no existence and uniqueness flags.

We therefore keep `gensys` for the determinacy map of the previous section and use the fixed point only where we need derivatives.

148.14.4 The likelihood in JAX

The Kalman filter carries over unchanged, written with `lax.scan` so that it too can be differentiated.

The only substantive change is the initial covariance: instead of calling a Lyapunov solver we use $\mathrm{vec}(P) = (I - A \otimes A)^{-1}\,\mathrm{vec}(Q)$, which is one linear solve.

```
def state_space_jax(p):
    G, Theta = solve_fixed_point(p)
    sd = jnp.array([p['sig_e'], p['sig_a'], p['sig_chi'], p['sig_z'], p['sig_m']])
    n = NJ + 2                                  # append x(-1) and a(-1)
    A = jnp.zeros((n, n)).at[:NJ, :NJ].set(G)
    A = A.at[NJ, X].set(1).at[NJ + 1, A_].set(1)
    B = jnp.zeros((n, 5)).at[:NJ].set(Theta * sd)
    C = jnp.zeros((4, n))
    C = C.at[0, DM].set(1).at[1, PI].set(1).at[2, R].set(1)
    C = C.at[3, X].set(1).at[3, NJ].set(-1)
    C = C.at[3, A_].set(p['xi']).at[3, NJ + 1].set(-p['xi'])
    C = C.at[3, Z].set(1)
    return A, B, C


def loglik_jax(p, Y):
    A, B, C = state_space_jax(p)
    n = A.shape[0]
    Q = B @ B.T
    P0 = jnp.linalg.solve(jnp.eye(n * n) - jnp.kron(A, A),
                          Q.reshape(-1)).reshape(n, n)
    const = Y.shape[1] * jnp.log(2 * jnp.pi)

    def step(carry, y):
        s, P = carry
        s, P = A @ s, A @ P @ A.T + Q
        v = y - C @ s
        PCt = P @ C.T
        F = C @ PCt
        L = jnp.linalg.cholesky(F)
        u = jax.scipy.linalg.solve_triangular(L, v, lower=True)
        ll = -0.5 * (const + 2 * jnp.sum(jnp.log(jnp.diag(L))) + u @ u)
        K = jnp.linalg.solve(F, PCt.T).T
```

(continues on next page)

(continued from previous page)

```
        return (s + K @ v, P - K @ PCt.T), ll

    _, lls = lax.scan(step, (jnp.zeros(n), P0), Y)
    return jnp.sum(lls)
```

Two checks: the new likelihood must agree with the one we used for the random walk, and its gradient must agree with finite differences.

```
Y_jax = jnp.asarray(Y_est)
print(f'log likelihood, JAX    {float(loglik_jax(p_check, Y_jax)):.8f}')
print(f'log likelihood, NumPy  {loglik(dict(p_check, tau=TAU), Y_est):.8f}')

grad_ll = jax.jit(jax.grad(lambda q: loglik_jax(q, Y_jax)))
g = grad_ll(p_check)

print('\n                  autodiff    finite difference')
for n in ['phi_pi', 'kappa', 'rho_m', 'sig_z']:
    step_n = 1e-5
    up = dict(p_check); up[n] = p_check[n] + step_n
    dn = dict(p_check); dn[n] = p_check[n] - step_n
    fd = (loglik_jax(up, Y_jax) - loglik_jax(dn, Y_jax)) / (2 * step_n)
    print(f'   {n:8s} {float(g[n]):12.5f} {float(fd):17.5f}')
```

```
log likelihood, JAX    -493.21564637
log likelihood, NumPy  -493.21564637

                  autodiff    finite difference
   phi_pi       -1.00005          -1.00005
   kappa        -0.68555          -0.68555
   rho_m         1.92667           1.92667
   sig_z       -56.71510         -56.71510
```

```
t0 = time.time()
for _ in range(20):
    gg = grad_ll(p_check)
jax.block_until_ready(gg)
grad_ms = 1000 * (time.time() - t0) / 20

t0 = time.time()
for _ in range(20):
    loglik(dict(p_check, tau=TAU), Y_est)
loglik_ms = 1000 * (time.time() - t0) / 20

print(f'one NumPy log likelihood         {loglik_ms:6.2f} ms')
print(f'one JAX gradient (18 partials)   {grad_ms:6.2f} ms')
```

```
one NumPy log likelihood          7.91 ms
one JAX gradient (18 partials)    4.92 ms
```

A full gradient costs about what a single likelihood evaluation costs, which is the whole point of reverse-mode differentiation.

A finite-difference gradient would need at least nineteen likelihood evaluations.

148.14.5 Sampling with NUTS

We hand the same priors to NumPyro [Phan *et al.*, 2019], which supplies NUTS and handles the transformations to unconstrained space that Hamiltonian dynamics require.

```
def beta_np(m, s):
    nu = m * (1 - m) / s ** 2 - 1
    return dist.Beta(m * nu, (1 - m) * nu)


def gamma_np(m, s):
    return dist.Gamma(m ** 2 / s ** 2, m / s ** 2)


def invgamma_np(m, s):
    a = m ** 2 / s ** 2 + 2
    return dist.InverseGamma(a, m * (a - 1))


PRIORS_NP = {
    'theta': beta_np(0.99, 0.005), 'alpha_pi': beta_np(0.5, 0.2),
    'alpha_x': beta_np(0.5, 0.2), 'rho_m': beta_np(0.5, 0.05),
    'rho_e': beta_np(0.5, 0.1), 'rho_a': beta_np(0.5, 0.1),
    'rho_chi': beta_np(0.5, 0.1),
    'kappa': gamma_np(0.3, 0.1), 'sigma': gamma_np(0.1, 0.05),
    'xi': gamma_np(2.0, 1.0), 'gamma': gamma_np(4.0, 1.0),
    'phi_pi': dist.Normal(0, 0.5), 'phi_x': dist.Normal(0, 0.5),
    'sig_e': invgamma_np(0.3, 1.0), 'sig_a': invgamma_np(0.3, 1.0),
    'sig_chi': invgamma_np(0.3, 1.0), 'sig_z': invgamma_np(0.3, 1.0),
    'sig_m': invgamma_np(0.3, 1.0)}


def ss_model(Y):
    p = {n: numpyro.sample(n, PRIORS_NP[n]) for n in FREE}
    numpyro.factor('loglik', loglik_jax(p, Y))
```

Three settings matter.

We ask for a **dense mass matrix**, because the condition number reported above says the posterior has correlations that a diagonal preconditioner cannot absorb.

We cap the trajectory length, since without a cap NUTS spends most of its time on very long trajectories in the flattest directions.

And we run **four chains** rather than one, all started from the posterior mode.

That last choice is the one that earns its keep.

```
kernel = NUTS(ss_model, target_accept_prob=0.8, dense_mass=True,
              max_tree_depth=8,
              init_strategy=numpyro.infer.init_to_value(
                  values={n: float(v) for n, v in zip(FREE, v_mode)}))
mcmc = MCMC(kernel, num_warmup=400, num_samples=400, num_chains=4,
            chain_method='sequential', progress_bar=False)

t0 = time.time()
mcmc.run(jax.random.PRNGKey(1), Y_jax, extra_fields=('num_steps', 'diverging'))
jax.block_until_ready(mcmc.get_samples())
nuts_seconds = time.time() - t0
```

(continues on next page)

(continued from previous page)

```
extra = mcmc.get_extra_fields()
print(f'{nuts_seconds:.0f} seconds for 4 chains of 400 draws')
print(f'mean leapfrog steps per iteration  '
      f'{np.asarray(extra["num_steps"]).mean():.0f}')
print(f'divergences                        '
      f'{int(np.asarray(extra["diverging"]).sum())}')
```

```
410 seconds for 4 chains of 400 draws
mean leapfrog steps per iteration  22
divergences                        0
```

```
nuts_idata = az.from_numpyro(mcmc)
nuts_summary = az.summary(nuts_idata, var_names=FREE,
                          ci_kind='hdi', ci_prob=0.94)
nuts_kept = np.column_stack([np.asarray(mcmc.get_samples()[n]) for n in FREE])
nuts_summary[['mean', 'sd', 'hdi94_lb', 'hdi94_ub', 'ess_bulk', 'r_hat']]
```

	mean	sd	hdi94_lb	hdi94_ub	ess_bulk	r_hat
theta	0.9889	0.0052	0.98	1	1811	1.00
alpha_pi	0.6	0.3	0.074	0.96	12	1.26
kappa	0.111	0.045	0.035	0.2	1369	1.00
alpha_x	0.387	0.039	0.31	0.45	926	1.01
sigma	0.085	0.031	0.032	0.14	1389	1.00
xi	3.46	1.04	1.7	5.3	1950	1.00
gamma	1.99	0.29	1.5	2.5	45	1.06
phi_pi	-0.3	0.24	-0.71	0.18	43	1.06
phi_x	-0.82	0.3	-1.3	-0.27	193	1.03
rho_m	0.518	0.046	0.43	0.6	1522	1.01
rho_e	0.4	0.3	0.16	0.87	12	1.25
rho_a	0.846	0.047	0.76	0.93	1004	1.01
rho_chi	0.514	0.101	0.33	0.7	1836	1.00
sig_e	0.9	0.4	0.29	1.4	12	1.25
sig_a	0.255	0.077	0.12	0.4	1875	1.00
sig_chi	0.149	0.08	0.032	0.28	1940	1.00
sig_z	0.714	0.108	0.54	0.94	591	1.01
sig_m	0.716	0.066	0.6	0.85	1156	1.00

148.14.6 A warning sign in the diagnostics

Most parameters look excellent, with $\hat{R}$ at one and effective sample sizes in the hundreds or thousands.

A few do not, and it is worth asking which.

```
worst = nuts_summary['ess_bulk'].nsmallest(3).index.tolist()
print('weakest mixing:')
print(nuts_summary.loc[worst, ['mean', 'sd', 'ess_bulk', 'r_hat']])

by_chain = mcmc.get_samples(group_by_chain=True)
print('\nper-chain posterior means')
print(f'{"":10s}' + ''.join(f'{"chain " + str(c):>10s}' for c in range(4)))
for n in worst:
    v = np.asarray(by_chain[n])
    print(f'{n:10s}' + ''.join(f'{v[c].mean():10.4f}' for c in range(4)))
```

```
weakest mixing:
          mean   sd ess_bulk r_hat
sig_e      0.9  0.4       12  1.25
rho_e      0.4  0.3       12  1.25
alpha_pi   0.6  0.3       12  1.26

per-chain posterior means
             chain 0   chain 1   chain 2   chain 3
sig_e         0.5609    0.9451    1.1382    1.1475
rho_e         0.6876    0.4083    0.2615    0.2633
alpha_pi      0.3025    0.6328    0.8050    0.8063
```

The chains started from a common point, so any disagreement between them means some of them have wandered somewhere the others have not.

The identity of the badly behaved parameters is a clue.

Inflation in the data is persistent, and equation (148.7) offers two ways to deliver that persistence.

Intrinsic persistence comes from indexation to past inflation, a large α_π, with the markup shock left transitory.

Inherited persistence comes from a persistent markup shock, a large ρ_e, with little indexation.

The parameters that mix worst are exactly the ones that distinguish these two stories.

That is a reason to suspect the posterior has more than one mode, and it tells us where to look for the second one.

148.14.7 Two modes

Sampler diagnostics are a noisy instrument, so we settle the question with the optimizer instead.

Our mode search in the estimation section started from the paper's posterior means and climbed to a high-indexation mode.

We now start the same optimizer from the opposite corner.

```
v_start_a = v_mode.copy()
for n, val in [('alpha_pi', 0.12), ('rho_e', 0.84), ('sig_e', 0.35)]:
    v_start_a[FREE.index(n)] = val

res_a = optimize.minimize(neg_log_post, v_start_a, method='Powell',
                          options=dict(maxiter=20000, maxfev=20000))
res_a = optimize.minimize(neg_log_post, res_a.x, method='L-BFGS-B',
                          bounds=[SUPPORT[n] for n in FREE])
v_mode_a = res_a.x

print(f'log posterior, low-indexation mode   {-res_a.fun:10.3f}')
print(f'log posterior, high-indexation mode  {log_post(v_mode, Y_est):10.3f}')
print('\n                low-index   high-index      paper')
for n in ['alpha_pi', 'rho_e', 'sig_e', 'kappa', 'phi_pi', 'phi_x']:
    i = FREE.index(n)
    print(f'   {n:9s}{v_mode_a[i]:10.4f}{v_mode[i]:13.4f}{PAPER[n]:11.4f}')
```

```
log posterior, low-indexation mode     -318.746
log posterior, high-indexation mode    -320.204

                low-index   high-index      paper
   alpha_pi        0.1096       0.8296     0.8815
```

(continues on next page)

(continued from previous page)

```
rho_e        0.8376       0.2359      0.5645
sig_e        0.3288       1.1516      1.0922
kappa        0.0817       0.0967      0.0324
phi_pi      -0.4359      -0.1957      0.2312
phi_x       -0.7975      -0.5739     -0.1971
```

There are two distinct modes, and the low-indexation one has the **higher** posterior density.

The mode we found earlier, the one nearest the paper's published estimates and the one our random walk explored for thirty thousand draws, is a *local* mode.

The pooled NUTS draws show the same thing from the sampling side.

```
fig, axes = plt.subplots(1, 3, figsize=(13, 3.8))
for ax, n, ttl in zip(axes, ['alpha_pi', 'rho_e', 'sig_e'],
                      [r'$\alpha_\pi$: indexation',
                       r'$\rho_e$: markup persistence',
                       r'$\sigma_e$: markup volatility']):
    i = FREE.index(n)
    ax.hist(nuts_kept[:, i], bins=60, density=True, color='C0', alpha=0.8)
    ax.axvline(v_mode_a[i], color='C2', lw=2, label='low-index mode')
    ax.axvline(v_mode[i], color='C3', lw=2, label='high-index mode')
    ax.axvline(PAPER[n], color='k', ls='--', lw=1.2, label='paper')
    ax.set(title=ttl, xlabel=n)
axes[0].legend(fontsize=8)
fig.suptitle('The marginals spread across both modes rather than concentrating on one
↪')
fig.tight_layout()
plt.show()
```

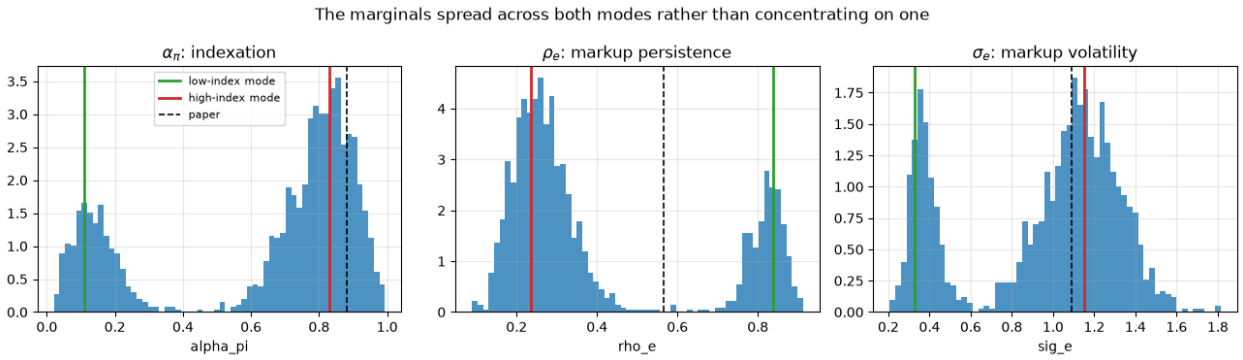


Fig. 148.1: Marginal posteriors for the parameters that separate the two modes

The marginals are wide and lumpy rather than the tidy bells that every other parameter produces.

The chains do move between the two regions, but slowly, and that slow movement is precisely what the poor effective sample sizes were reporting.

Thirty thousand random walk draws never made the trip at all.

The economics behind the two modes is a familiar identification problem.

The data over 1960-1983 cannot tell us whether inflation is persistent because prices are indexed to past inflation or because the shocks pushing inflation around are themselves persistent.

Both stories fit, and the published estimates describe the one the data like slightly less.

148.14.8 Does the bimodality change the economics?

This is the question that matters for the rest of the lecture.

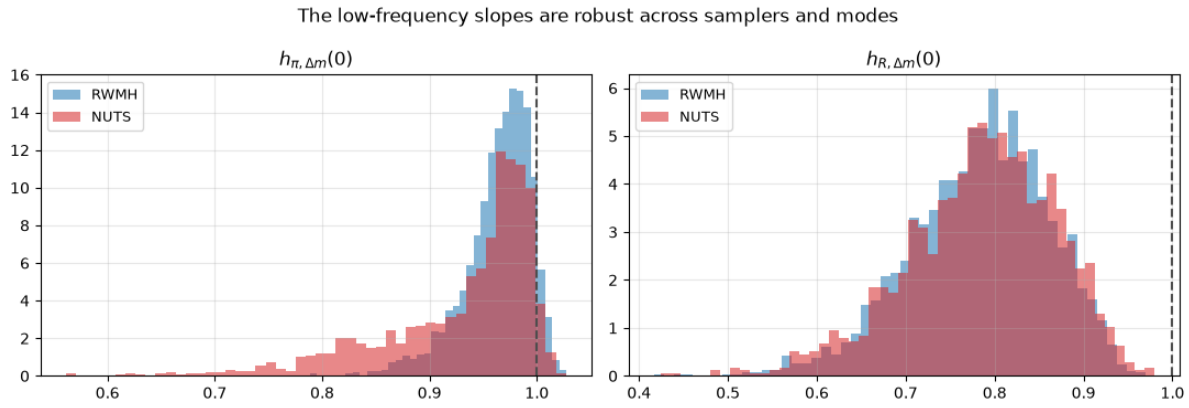
```
h_mode_a = h_zero(unpack(v_mode_a))
h_mode_b = h_zero(unpack(v_mode))
print(f'h(0) at the low-indexation mode   {h_mode_a[0]:.4f}, {h_mode_a[1]:.4f}')
print(f'h(0) at the high-indexation mode  {h_mode_b[0]:.4f}, {h_mode_b[1]:.4f}')

h_nuts = np.array([h_zero(unpack(v)) for v in nuts_kept])
low = nuts_kept[:, FREE.index('alpha_pi')] < 0.5
print(f'\n{low.sum()} of {len(low)} NUTS draws sit on the low-indexation side')
for tag, m in [('low-indexation draws ', low), ('high-indexation draws', ~low)]:
    if m.sum() < 20:
        continue
    print(f'{tag}  h_pi {h_nuts[m, 0].mean():.4f} '
          f'[{np.percentile(h_nuts[m, 0], 5):.4f}, {np.percentile(h_nuts[m, 0], 95):.
↪4f}]'
          f'   h_R {h_nuts[m, 1].mean():.4f} '
          f'[{np.percentile(h_nuts[m, 1], 5):.4f}, {np.percentile(h_nuts[m, 1], 95):.
↪4f}]')
```

```
h(0) at the low-indexation mode   0.8632, 0.8292
h(0) at the high-indexation mode  0.9765, 0.7694

415 of 1600 NUTS draws sit on the low-indexation side
low-indexation draws   h_pi 0.8323 [0.7082, 0.9163]   h_R 0.8018 [0.6645, 0.9159]
high-indexation draws  h_pi 0.9635 [0.9042, 1.0014]   h_R 0.7793 [0.6273, 0.9063]
```

```
fig, axes = plt.subplots(1, 2, figsize=(11, 3.8))
for ax, j, ttl in zip(axes, [0, 1],
                      [r'$h_{\pi,\Delta m}(0)$', r'$h_{R,\Delta m}(0)$']):
    ax.hist(h_draws[:, j], bins=50, density=True, alpha=0.55,
            color='C0', label='RWMH')
    ax.hist(h_nuts[:, j], bins=50, density=True, alpha=0.55,
            color='C3', label='NUTS')
    ax.axvline(1, color='0.3', ls='--', lw=1.5)
    ax.set_title(ttl)
    ax.legend(fontsize=9)
fig.suptitle('The low-frequency slopes are robust across samplers and modes')
plt.tight_layout()
plt.show()
```



The reassuring answer is that it does not.

The two modes disagree sharply about *why* inflation is persistent and far less about the low-frequency slopes that this lecture is about.

Whether inflation persistence is intrinsic or inherited, $h_{\pi,\Delta m}(0)$ stays between about four fifths and one, and $h_{R,\Delta m}(0)$ sits near four fifths in both cases.

The paper's substantive conclusions survive the discovery that its parameter estimates describe a local mode.

148.14.9 Comparing the samplers

With that settled we can return to efficiency.

Effective sample sizes for the three parameters that separate the modes are not comparable across the two samplers, because for NUTS they partly measure movement *between* modes while for the random walk they measure movement within one.

We therefore compare medians rather than minima.

```
compare = pd.DataFrame({
    'RWMH mean': rwmh_summary['mean'], 'NUTS mean': nuts_summary['mean'],
    'RWMH ESS': rwmh_summary['ess_bulk'], 'NUTS ESS': nuts_summary['ess_bulk'],
    'NUTS r_hat': nuts_summary['r_hat']})
compare['ESS ratio'] = compare['NUTS ESS'] / compare['RWMH ESS']
compare.round(3)
```

	RWMH mean	NUTS mean	RWMH ESS	NUTS ESS	NUTS r_hat	ESS ratio
theta	0.989	0.989	197.674	1811.677	1.000	9.165
alpha_pi	0.816	0.637	415.958	12.378	1.259	0.030
kappa	0.117	0.111	195.344	1369.515	1.004	7.011
alpha_x	0.380	0.387	112.547	926.736	1.007	8.234
sigma	0.082	0.085	165.631	1389.912	1.003	8.392
xi	3.579	3.464	111.759	1950.431	1.004	17.452
gamma	2.043	1.994	125.072	45.640	1.060	0.365
phi_pi	-0.236	-0.299	220.898	43.144	1.062	0.195
phi_x	-0.799	-0.825	91.547	193.422	1.025	2.113
rho_m	0.516	0.518	221.684	1522.476	1.006	6.868
rho_e	0.258	0.405	328.990	12.208	1.253	0.037
rho_a	0.850	0.846	53.416	1004.125	1.006	18.798
rho_chi	0.517	0.514	307.121	1836.621	1.003	5.980
sig_e	1.165	0.948	279.498	12.005	1.246	0.043
sig_a	0.247	0.255	93.286	1875.823	1.001	20.108

(continues on next page)

(continued from previous page)

```
sig_chi       0.141      0.149   200.244  1940.388       1.003       9.690
sig_z         0.713      0.714    96.781   591.203       1.012       6.109
sig_m         0.714      0.716   149.238  1156.163       1.005       7.747
```

```
unimodal = ~nuts_summary.index.isin(worst)
r_med = rwmh_summary['ess_bulk'][unimodal].median()
n_med = nuts_summary['ess_bulk'][unimodal].median()

print(f'{"":26s}{"RWMH":>12s}{"NUTS":>12s}')
print(f'{"draws kept":26s}{len(kept):12d}{len(nuts_kept):12d}')
print(f'{"seconds":26s}{rwmh_seconds:12.0f}{nuts_seconds:12.0f}')
print(f'{"median ESS":26s}{r_med:12.0f}{n_med:12.0f}')
print(f'{"median ESS per 1000 draws":26s}{1000 * r_med / len(kept):12.1f}'
      f'{1000 * n_med / len(nuts_kept):12.1f}')
print(f'{"median ESS per second":26s}{r_med / rwmh_seconds:12.2f}'
      f'{n_med / nuts_seconds:12.2f}')
print(f'\n{"speed-up in ESS per second":30s}'
      f'{(n_med / nuts_seconds) / (r_med / rwmh_seconds):6.1f}x')
```

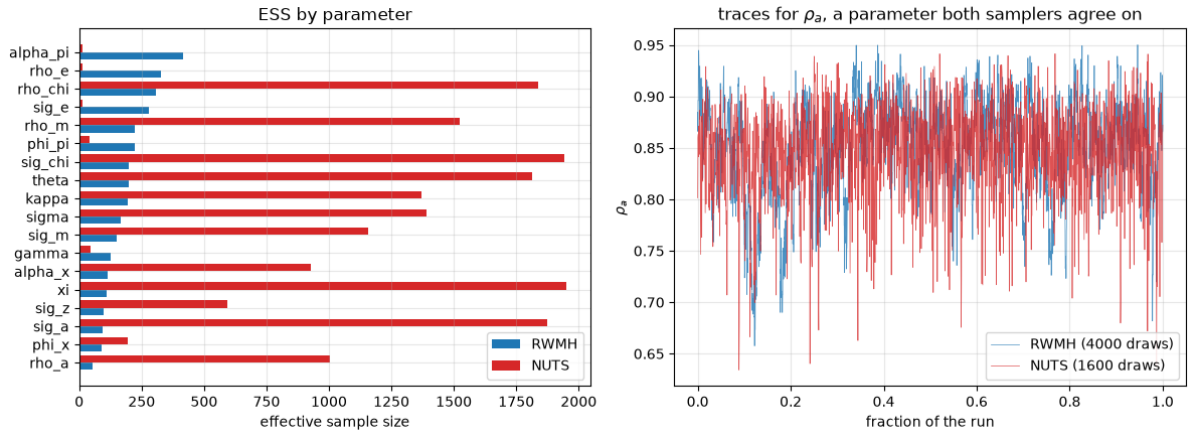
```
                                RWMH        NUTS
draws kept                      4000        1600
seconds                          278         410
median ESS                       149        1370
median ESS per 1000 draws       37.3       855.9
median ESS per second           0.54        3.34

speed-up in ESS per second      6.2x
```

Per draw the gap is very large, because each NUTS draw ends a trajectory that has travelled across the posterior rather than taking one small random step.

Cost claws some of that back, since a NUTS iteration needs many gradient evaluations where a random walk iteration needs one likelihood evaluation.

Effective draws per second is the ratio that matters, and it still favours NUTS by a wide margin.



148.14.10 What to take away

The obstacle to Hamiltonian methods in DSGE estimation is the model solver, not the statistics.

Replacing an eigenvalue-sorting solver by a fixed point that uses only linear algebra buys exact gradients for about the cost of one extra likelihood evaluation, and that is enough to put NUTS within reach.

The payoff was not only speed.

Cheap chains made it cheap to run several of them and inspect their diagnostics, and the parameters that mixed worst pointed straight at a second and better mode that thirty thousand random walk draws had never visited.

Two caveats are worth stating.

First, the fixed point returns something meaningless where no stable solution exists, so a sampler that wanders outside the determinacy region will produce nonsense rather than a rejection; the posterior here sits well inside that region, but a flatter one might not.

Second, the efficiency gain reported above is specific to this posterior, and a well-scaled posterior in fewer dimensions would narrow the gap considerably.

148.15 Concluding remarks

Lucas [1980] was careful to say that his two unit slopes should be expected to hold under some monetary policies and to break down under others.

Extending his sample by half a century confirms the breakdowns.

Estimating a small structural model on a sample like his, and then moving only the monetary policy rule, produces slopes that range over most of what the data display.

The lesson is the one Lucas, Sargent [1971], and Whiteman [1984] all drew, namely that low-frequency regression coefficients are not structural.

Working through the model from scratch adds something to the paper's account.

Lucas's first illustration, the unit slope of inflation on money growth, is an *identity* in this model whenever money growth is econometrically exogenous, because money demand is a quantity equation whose remaining terms are all first differences.

His second illustration, the unit slope of the interest rate on money growth, has no such support and moves around freely.

So the two illustrations, which look symmetric in a scatter plot, are quite different objects from the perspective of a structural model.

That asymmetry also marks the limit of the exercise, as Exercise 148.16.2 brings out.

Because $h_{\pi,\Delta m}(0)$ is anchored at one by money demand, driving it to the near-zero values that the post-1984 data display takes a money growth rule far more anti-inflationary than the one those same data select.

On the computational side, the model turned out to be a useful test bed for Hamiltonian Monte Carlo.

The barrier to using it on DSGE models is not statistical but algorithmic: the standard solvers sort eigenvalues, and sorting has no derivative.

Swapping in a fixed-point solver that uses only linear algebra restores exact gradients, and the sampler that becomes available is far more efficient per unit of computing time on a posterior as badly scaled as this one.

The larger dividend was a diagnostic one.

Cheap chains and their convergence diagnostics led us to a second mode, and the posterior turned out to be bimodal: the data cannot choose between intrinsic and inherited inflation persistence, and the mode nearest the published estimates is the *lower* of the two.

That the two modes nonetheless agree about $h_{\pi,\Delta m}(0)$ and $h_{R,\Delta m}(0)$ is a piece of good news for the paper that only the more thorough sampler could deliver.

148.16 Exercises

Exercise 148.16.1

The paper takes technology growth to be serially uncorrelated.

Replace $z_t = \varepsilon_{zt}$ by $z_t = \rho_z z_{t-1} + \varepsilon_{zt}$ and ask two questions.

Does *Proposition 148.9.1* survive, and does persistent technology growth offer a route to breaking the unit slope that has nothing to do with monetary policy?

Solution

The technology block is row 7 of the canonical form, where `g1[7, Z]` is currently zero.

```
def canonical_rho_z(p, rule='money'):
    g0, g1, psi, pie = canonical(p, rule)
    g1[7, Z] = p['rho_z']
    return g0, g1, psi, pie


def h_zero_rho_z(p):
    G1, impact, eu = gensys(*canonical_rho_z(p))
    if eu != (1, 1):
        return np.nan, np.nan
    sd = np.diag([p['sig_e'], p['sig_a'], p['sig_chi'], p['sig_z'], p['sig_m']])
    n = NY + 2
    A, B = np.zeros((n, n)), np.zeros((n, 5))
    A[:NY, :NY], A[NY, X], A[NY + 1, A_] = G1, 1, 1
    B[:NY] = impact @ sd
    C = np.zeros((4, n))
    C[0, DM], C[1, PI], C[2, R], C[3, Z] = 1, 1, 1, 1
    C[3, X], C[3, NY] = 1, -1
    C[3, A_], C[3, NY + 1] = p['xi'], -p['xi']
    if np.max(np.abs(np.linalg.eigvals(A))) > 1 - 1e-9:
        return np.nan, np.nan
    return (h_zero_from_state_space(A, B, C, 1, 0),
            h_zero_from_state_space(A, B, C, 2, 0))


print('exogenous money growth, phi_pi = phi_x = 0')
for rho_z in [0.0, 0.3, 0.6, 0.9]:
    q = dict(p_bar, phi_pi=0.0, phi_x=0.0, rho_z=rho_z)
    print(f'   rho_z = {rho_z:.1f}:  h_pi = {h_zero_rho_z(q)[0]:.10f}, '
          f'h_R = {h_zero_rho_z(q)[1]:.4f}')

print('\nwith policy feedback, phi_pi = -2, phi_x = -0.5')
for rho_z in [0.0, 0.3, 0.6, 0.9]:
    q = dict(p_bar, phi_pi=-2.0, phi_x=-0.5, rho_z=rho_z)
    print(f'   rho_z = {rho_z:.1f}:  h_pi = {h_zero_rho_z(q)[0]:.4f}, '
          f'h_R = {h_zero_rho_z(q)[1]:.4f}')
```

```
exogenous money growth, phi_pi = phi_x = 0
   rho_z = 0.0:  h_pi = 1.0000000000, h_R = 0.4111
   rho_z = 0.3:  h_pi = 1.0000000000, h_R = 0.4111
   rho_z = 0.6:  h_pi = 1.0000000000, h_R = 0.4111
   rho_z = 0.9:  h_pi = 1.0000000000, h_R = 0.4111

with policy feedback, phi_pi = -2, phi_x = -0.5
   rho_z = 0.0:  h_pi = 0.4897, h_R = 0.7767
   rho_z = 0.3:  h_pi = 0.2340, h_R = 0.6392
   rho_z = 0.6:  h_pi = -0.1354, h_R = 0.3046
   rho_z = 0.9:  h_pi = -0.4584, h_R = -0.2820
```

Proposition 148.9.1 survives untouched, because its proof used only that technology is orthogonal to the policy shock, never how persistent it is.

Under exogenous money growth *both* slopes are invariant to ρ_z, and for a reason worth noticing: z_t enters the model only contemporaneously, through money demand, so ρ_z changes the dynamics of technology without changing anyone's response to a money shock.

Once policy feeds back the picture is different, and persistent technology growth pulls $h_{\pi,\Delta m}(0)$ well away from one.

So the answer to the second question is yes, and it points at a limitation of the paper's experiment: the technology process is held fixed while policy varies, but the two interact.

Exercise 148.16.2

The lecture estimates the model on 1960:I-1983:IV.

Find the posterior mode on 1984:I-2007:IV and compare the policy coefficients and the implied $h(0)$ with what we found for the earlier sample.

The VAR put $h_{\pi,\Delta m}(0)$ near zero over 1984-2007, so does the structural model, re-estimated, agree?

Solution

```
mask2 = (data.index.year >= 1984) & (data.index.year <= 2007)
Y2 = data.loc[mask2, VARS].values
Y2 = Y2 - Y2.mean(0)

res2 = optimize.minimize(lambda v: -log_post(v, Y2), v_mode, method='Powell',
                         options=dict(maxiter=20000, maxfev=20000))
res2 = optimize.minimize(lambda v: -log_post(v, Y2), res2.x, method='L-BFGS-B',
                         bounds=[SUPPORT[n] for n in FREE])
v2 = res2.x
print('                        1960-1983    1984-2007')
for n in ['phi_pi', 'phi_x', 'rho_m']:
    j = FREE.index(n)
    print(f'   {n:9s} {v_mode[j]:11.4f} {v2[j]:12.4f}')
print('   %-9s %5.3f, %.3f  %5.3f, %.3f'
      % (('model h(0)',) + h_zero(unpack(v_mode)) + h_zero(unpack(v2))))
for name, lo, hi in [('1960-1983', 1960, 1983), ('1984-2007', 1984, 2007)]:
    m = (data.index.year >= lo) & (data.index.year <= hi)
    med = np.nanmedian(bvar_h0(data.loc[m, VARS].values), axis=0)
    print(f'   VAR   h(0) on {name}: {med[0]:.3f}, {med[1]:.3f}')
```

```
                      1960-1983    1984-2007
phi_pi          -0.1957      -0.0399
phi_x           -0.5739      -1.5633
rho_m            0.5301       0.3080
model h(0) 0.976, 0.769  0.989, 0.858
VAR   h(0) on 1960-1983: 0.739, 0.607
VAR   h(0) on 1984-2007: -0.073, 0.268
```

The re-estimated model does *not* agree with the VAR.

The policy coefficients move, but the implied $h_{\pi,\Delta m}(0)$ stays close to one on both samples, while the VAR puts it near zero after 1984.

Proposition 148.9.1 explains why this had to be hard.

Within this model $h_{\pi,\Delta m}(0)$ is anchored at one and is dragged away from it only by policy feedback, so reaching values near zero needs a large negative ϕ_π, and the post-1984 data do not put the money growth rule anywhere near there.

That is a genuine tension between the paper's Section II and its Section III: the contour plot shows a policy rule that *could* generate the low slopes, but the rule the later data actually select is not that rule.

A full comparison would run the sampler on both samples rather than compare modes.

Exercise 148.16.3

Figure 8 of the paper asks whether a fall in the variance of supply shocks, rather than a change in policy, could account for the decline in $h(0)$.

Redo the money rule contour plot with σ_e cut to one quarter of its posterior mean and report how much the picture moves.

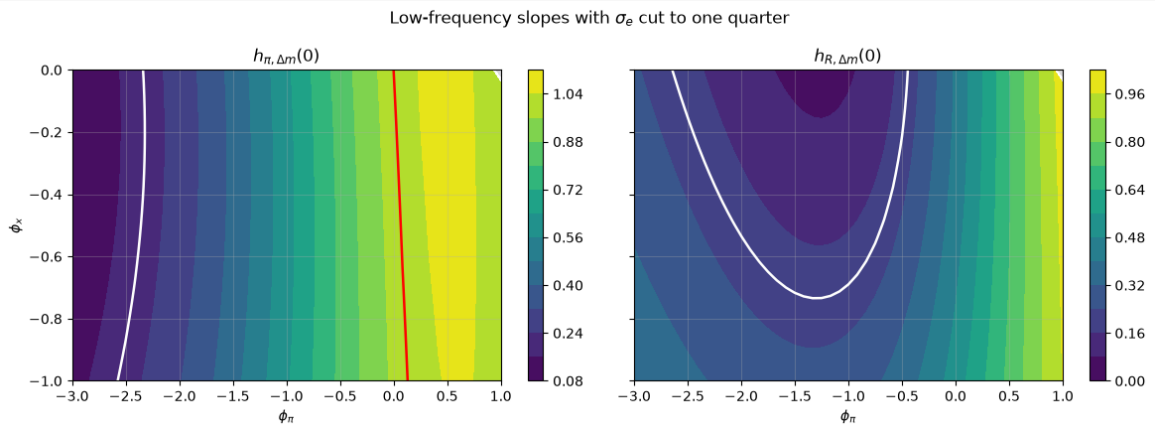
Solution

```
p_small_e = dict(p_bar, sig_e=p_bar['sig_e'] / 4)
Hpi2, HR2 = h_grid(p_small_e, 'money', 'phi_pi', 'phi_x',
                   phi_pi_grid, phi_x_grid)

print('baseline    h_pi in [%.2f, %.2f],  h_R in [%.2f, %.2f]'
      % (np.nanmin(Hpi), np.nanmax(Hpi), np.nanmin(HR), np.nanmax(HR)))
print('small sig_e h_pi in [%.2f, %.2f],  h_R in [%.2f, %.2f]'
      % (np.nanmin(Hpi2), np.nanmax(Hpi2), np.nanmin(HR2), np.nanmax(HR2)))
print('largest change in h_pi across the grid: %.3f'
      % np.nanmax(np.abs(Hpi2 - Hpi)))

contour_panel(phi_pi_grid, phi_x_grid, Hpi2, HR2, r'$\phi_\pi$', r'$\phi_x$',
              r'Low-frequency slopes with $\sigma_e$ cut to one quarter')
```

```
baseline    h_pi in [0.10, 1.06],  h_R in [0.06, 1.21]
small sig_e h_pi in [0.10, 1.06],  h_R in [0.06, 1.00]
largest change in h_pi across the grid: 0.432
```



The range that $h_{\pi,\Delta m}(0)$ can reach is unchanged, and the mapping from policy to $h(0)$ keeps the same shape.

The contours do move within that range, by up to a few tenths at some grid points, so the claim is not that shock variances are irrelevant.

It is the weaker claim the paper makes, that a change in shock variances of this magnitude is not enough to substitute for a change in policy.

Part XVIII

Auctions

CHAPTER

NINE

FIRST-PRICE AND SECOND-PRICE AUCTIONS

This lecture is designed to set the stage for a subsequent lecture about *Multiple Good Allocation Mechanisms*

In that lecture, a planner or auctioneer simultaneously allocates several goods to set of people.

In the present lecture, a single good is allocated to one person within a set of people.

Here we'll learn about and simulate two classic auctions :

- a First-Price Sealed-Bid Auction (FPSB)
- a Second-Price Sealed-Bid Auction (SPSB) created by William Vickrey [Vickrey, 1961]

We'll also learn about and apply a

- Revenue Equivalent Theorem

We recommend watching this video about second price auctions by Anders Munk-Nielsen:

https://youtu.be/qwWk_Bqtue8

and

https://youtu.be/eYTGQCGpmXI

Anders Munk-Nielsen put his code on GitHub.

Much of our Python code below is based on his.

149.1 First-price sealed-bid auction (FPSB)

Protocols:

- A single good is auctioned.
- Prospective buyers simultaneously submit sealed bids.
- Each bidder knows only his/her own bid.
- The good is allocated to the person who submits the highest bid.
- The winning bidder pays price she has bid.

Detailed Setting:

There are $n \geq 2$ prospective buyers named $i = 1, 2, \ldots, n$.

Buyer i attaches value v_i to the good being sold.

Buyer i wants to maximize the expected value of her **surplus** defined as $v_i - p$, where p is the price that she pays, conditional on her winning the auction.

Evidently,

- If i bids exactly v_i, she pays what she thinks it is worth and gathers no surplus value.
- Buyer i will never want to bid more than v_i.
- If buyer i bids $b < v_i$ and wins the auction, then she gathers surplus value $v_i - b > 0$.
- If buyer i bids $b < v_i$ and someone else bids more than b, buyer i loses the auction and gets no surplus value.
- To proceed, buyer i wants to know the probability that she wins the auction as a function of her bid v_i
  - this requires that she know a probability distribution of bids v_j made by prospective buyers $j \neq i$
- Given her idea about that probability distribution, buyer i wants to set a bid that maximizes the mathematical expectation of her surplus value.

Bids are sealed, so no bidder knows bids submitted by other prospective buyers.

This means that bidders are in effect participating in a game in which players do not know **payoffs** of other players.

This is a **Bayesian game**, a Nash equilibrium of which is called a **Bayesian Nash equilibrium**.

To complete the specification of the situation, we'll assume that prospective buyers' valuations are independently and identically distributed according to a probability distribution that is known by all bidders.

Bidder optimally chooses to bid less than v_i.

149.1.1 Characterization of FPSB auction

We assume throughout that

- valuations are **private** and **independent** across bidders
- bidders are **symmetric**: their valuations are drawn from a common distribution F that is continuous and strictly increasing on its support
- bidders are **risk neutral**

Under these assumptions a FPSB auction has a unique Bayesian Nash equilibrium in symmetric, strictly increasing bidding strategies.

Because the equilibrium bidding strategy is strictly increasing, the bidder with the highest valuation submits the highest bid, and so wins.

The optimal bid of buyer i is

$$\mathbb{E}[y_i | y_i < v_i] \tag{149.1}$$

where v_i is the valuation of bidder i and y_i is the maximum valuation of all other bidders:

$$y_i = \max_{j \neq i} v_j \tag{149.2}$$

For a derivation, see the Wikipedia page about first-price sealed-bid auctions, or Krishna [2009], chapter 2.

We'll verify this formula by simulation below, and Exercise 149.13.2 asks you to derive an equivalent expression that is easy to evaluate for any distribution F.

149.2 Second-price sealed-bid auction (SPSB)

Protocols: In a second-price sealed-bid (SPSB) auction, the winner pays the second-highest bid.

149.3 Characterization of SPSB auction

In a SPSB auction bidders optimally choose to bid their values.

Formally, in a SPSB auction with a single, indivisible item, bidding one's own value is a **weakly dominant** strategy.

It is *weakly* dominant because a bidder who bids something other than her value never does better, and sometimes does worse, whatever the other bidders do.

Notice how much stronger this is than the FPSB result: it requires no assumption at all about the distribution of other bidders' valuations, nor about how they bid.

A proof is provided at the Wikipedia page about Vickrey auctions

149.4 Uniform distribution of private values

We assume valuation v_i of bidder i is distributed $v_i \overset{\text{IID}}{\sim} U(0,1)$.

Under this assumption, we can analytically compute probability distributions of prices bid in both FPSB and SPSB.

We'll simulate outcomes and, by using a law of large numbers, verify that the simulated outcomes agree with analytical ones.

We can use our simulation to illustrate a **revenue equivalence theorem** that asserts that on average first-price and second-price sealed bid auctions provide a seller the same revenue.

The theorem requires hypotheses that both of our auctions satisfy:

- valuations are independent and private, and bidders are symmetric and risk neutral
- the two mechanisms award the good to the bidder with the highest valuation
- a bidder with the lowest possible valuation expects zero surplus

Under these hypotheses any two such mechanisms yield the same expected payment for each bidder, and therefore the same expected revenue for the seller.

Exercise 149.13.4 shows what happens when one of these hypotheses – risk neutrality – fails.

To read about the revenue equivalence theorem, see this Wikipedia page

149.5 Setup

There are n bidders.

Each bidder knows that there are $n-1$ other bidders.

149.6 First price sealed bid auction

An optimal bid for bidder i in a **FPSB** is described by equations (149.1) and (149.2).

When bids are IID draws from a uniform distribution, the CDF of y_i is

$$\begin{aligned}\tilde{F}_{n-1}(y) = \mathbb{P}\{y_i \leq y\} &= \mathbb{P}\{\max_{j \neq i} v_j \leq y\} \\ &= \prod_{j \neq i} \mathbb{P}\{v_j \leq y\} \\ &= y^{n-1}\end{aligned}$$

and the PDF of y_i is $\tilde{f}_{n-1}(y) = (n-1)y^{n-2}$.

Then bidder i's optimal bid in a **FPSB** auction is:

$$\begin{aligned}\mathbb{E}[y_i | y_i < v_i] &= \frac{\int_0^{v_i} y_i \tilde{f}_{n-1}(y_i) dy_i}{\int_0^{v_i} \tilde{f}_{n-1}(y_i) dy_i} \\ &= \frac{\int_0^{v_i} (n-1) y_i^{n-1} dy_i}{\int_0^{v_i} (n-1) y_i^{n-2} dy_i} \\ &= \frac{n-1}{n} y_i \Big|_0^{v_i} \\ &= \frac{n-1}{n} v_i\end{aligned}$$

149.7 Second price sealed bid auction

In a **SPSB**, it is optimal for bidder i to bid v_i.

149.8 Python code

```
import numpy as np
import matplotlib.pyplot as plt
import seaborn as sns
import scipy.stats as stats
import scipy.interpolate as interp

# for plots
plt.rcParams.update({'font.size': 14})
colors = plt.rcParams['axes.prop_cycle'].by_key()['color']

# ensure the notebook generates the same randomness
rng = np.random.default_rng(1337)
```

We repeat an auction with 5 bidders for 100,000 times.

The valuations of each bidder is distributed $U(0, 1)$.

```
N = 5
R = 100_000

v = rng.uniform(0, 1, (N, R))

# BNE in first-price sealed bid

b_star = lambda vi, N: ((N-1)/N) * vi
b = b_star(v,N)
```

We compute and sort bid price distributions that emerge under both FPSB and SPSB.

```
# Bidders' values are sorted in ascending order in each auction.
# We record the order because we want to apply it to bid price and their id.
idx = np.argsort(v, axis=0)

# same as np.sort(v, axis=0), except now we retain the idx
v = np.take_along_axis(v, idx, axis=0)
b = np.take_along_axis(b, idx, axis=0)

# In FPSB and SPSB the winner is the bidder with the highest valuation,
# which after sorting is the last row.

# highest bid
winner_pays_fpsb = b[-1, :]
# 2nd-highest valuation
winner_pays_spsb = v[-2, :]
```

Let's now plot the *winning* bids $b_{(n)}$ (i.e. the payment) against valuations, $v_{(n)}$ for both FPSB and SPSB.

Note that

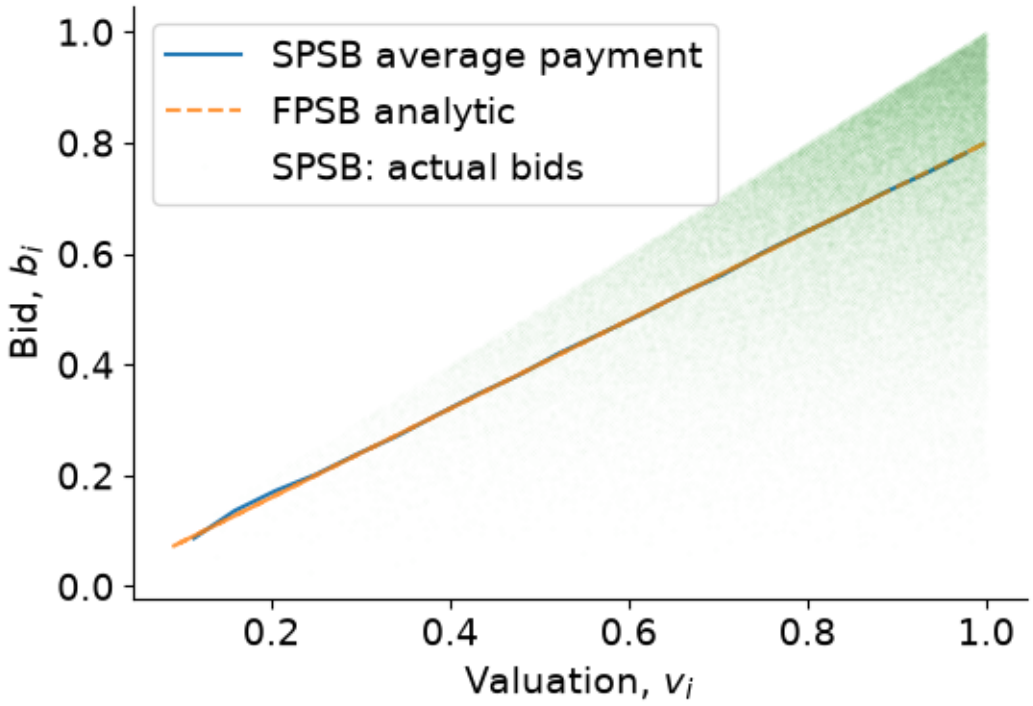
- FPSB: There is a unique bid corresponding to each valuation
- SPSB: Because it equals the valuation of a second-highest bidder, what a winner pays varies even holding fixed the winner's valuation. So here there is a frequency distribution of payments for each valuation.

```
# We intend to compute average payments of different groups of bidders
binned = stats.binned_statistic(v[-1, :], v[-2, :], statistic='mean', bins=20)
xx = binned.bin_edges
xx = [(xx[ii]+xx[ii+1])/2 for ii in range(len(xx)-1)]
yy = binned.statistic

fig, ax = plt.subplots(figsize=(6, 4))

ax.plot(xx, yy, label='SPSB average payment')
ax.plot(v[-1, :], b[-1, :], '--', alpha=0.8, label='FPSB analytic')
ax.plot(v[-1, :], v[-2, :], 'o', alpha=0.05,
                markersize=0.1, label='SPSB: actual bids')

ax.legend(loc='best')
ax.set_xlabel('Valuation, $v_i$')
ax.set_ylabel('Bid, $b_i$')
sns.despine()
```



149.9 Revenue equivalence theorem

We now compare FPSB and a SPSB auctions from the point of view of the revenues that a seller can expect to acquire.

Expected Revenue FPSB:

The winner with valuation y pays $\frac{n-1}{n}y$, where n is the number of bidders.

Above we computed that the CDF is $F_n(y) = y^n$ and the PDF is $f_n = ny^{n-1}$.

Consequently, expected revenue is

$$R = \int_0^1 \frac{n-1}{n} v_i \times n v_i^{n-1} dv_i = \frac{n-1}{n+1}$$

Expected Revenue SPSB:

The expected revenue equals n $\times$ expected payment of a bidder.

Computing this we get

$$\begin{aligned}
\text{TR} &= n\mathbb{E}_{v_i}\left[\mathbb{E}_{y_i}[y_i|y_i < v_i]\mathbb{P}\{y_i < v_i\} + 0 \times \mathbb{P}\{y_i > v_i\}\right] \\
&= n\mathbb{E}_{v_i}\left[\mathbb{E}_{y_i}[y_i|y_i < v_i]\tilde{F}_{n-1}(v_i)\right] \\
&= n\mathbb{E}_{v_i}[\frac{n-1}{n} \times v_i \times v_i^{n-1}] \\
&= (n-1)\mathbb{E}_{v_i}[v_i^n] \\
&= \frac{n-1}{n+1}
\end{aligned}$$

Thus, while probability distributions of winning bids typically differ across the two types of auction, we deduce that expected payments are identical in FPSB and SPSB.

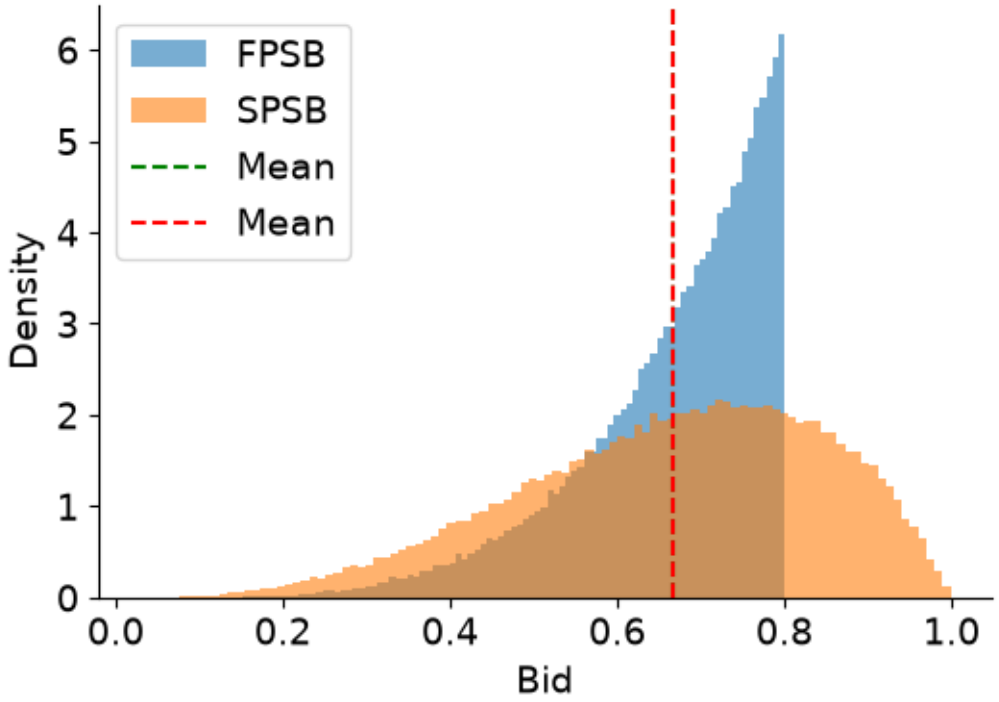
```
fig, ax = plt.subplots(figsize=(6, 4))

for payment, label in zip([winner_pays_fpsb, winner_pays_spsb], ['FPSB', 'SPSB']):
    print('The average payment of %s: %.4f. Std.: %.4f. Median: %.4f' % (
        label, payment.mean(), payment.std(), np.median(payment)))
    ax.hist(payment, density=True, alpha=0.6, label=label, bins=100)

ax.axvline(winner_pays_fpsb.mean(), ls='--', c='g', label='Mean')
ax.axvline(winner_pays_spsb.mean(), ls='--', c='r', label='Mean')

ax.legend(loc='best')
ax.set_xlabel('Bid')
ax.set_ylabel('Density')
sns.despine()
```

```
The average payment of FPSB: 0.6670. Std.: 0.1122. Median: 0.6959
The average payment of SPSB: 0.6660. Std.: 0.1782. Median: 0.6850
```



Summary of FPSB and SPSB results with uniform distribution on $[0, 1]$

Auction: Sealed-Bid	First-Price	Second-Price
Winner	Agent with highest bid	Agent with highest bid
Winner pays	Winner's bid	Second-highest bid
Loser pays	0	0
Dominant strategy	No dominant strategy	Bidding truthfully is dominant strategy
Bayesian Nash equilibrium	Bidder i bids $\frac{n-1}{n} v_i$	Bidder i truthfully bids v_i
Auctioneer's revenue	$\frac{n-1}{n+1}$	$\frac{n-1}{n+1}$

Detour: Computing a Bayesian Nash Equibrium for FPSB

The revenue equivalence theorem lets us find an optimal bidding strategy for a FPSB auction from outcomes of a SPSB auction.

Let $b(v_i)$ be the optimal bid in a FPSB auction.

The revenue equivalence theorem tells us that a bidder agent with value v_i on average receives the same **payment** in the two types of auction.

Consequently,

$$b(v_i)\mathbb{P}\{y_i < v_i\} + 0 \cdot \mathbb{P}\{y_i \geq v_i\} = \mathbb{E}_{y_i}[y_i | y_i < v_i]\mathbb{P}\{y_i < v_i\} + 0 \cdot \mathbb{P}\{y_i \geq v_i\}$$

It follows that an optimal bidding strategy in a FPSB auction is $b(v_i) = \mathbb{E}_{y_i}[y_i | y_i < v_i]$.

149.10 Calculation of bid price in FPSB

In equations (149.1) and (149.2), we displayed formulas for optimal bids in a symmetric Bayesian Nash equilibrium of a FPSB auction.

$$\mathbb{E}[y_i | y_i < v_i]$$

where

- $v_i =$ value of bidder i
- $y_i =:$ maximum value of all bidders except i, i.e., $y_i = \max_{j \neq i} v_j$

Above, we computed an optimal bid price in a FPSB auction analytically for a case in which private values are uniformly distributed.

For most probability distributions of private values, analytical solutions aren't easy to compute.

Instead, we can compute bid prices in FPSB auctions numerically as functions of the distribution of private values.

```
def evaluate_largest(v_hat, array, order=1):
    """
    A method to estimate the largest (or certain-order largest) value of the other␣
↪biders,
    conditional on player 1 wins the auction.

    We estimate E[y | y < v_hat], where y is the highest valuation among the
    other bidders.  We do this by taking bidder 1 as the reference bidder
    (valuations are IID, so the choice does not matter), discarding her row,
    and averaging the highest remaining valuation over those auctions in which
    every other bidder's valuation falls below v_hat.

    Parameters:
    ----------
    v_hat : float, the valuation of the reference bidder.

    array: 2 dimensional array of bidders' values in shape of (N,R),
           where N: number of players, R: number of auctions

    order: int. Which order statistic of the losing bidders to average.
                order=1 gives the highest losing valuation,
                order=2 the second highest, and so on.

    """
```

(continues on next page)

(continued from previous page)

```
    N, R = array.shape

    # discard the reference bidder's row; condition on the rest losing
    array_residual = array[1:, :].copy()

    winning_auctions_mask = (array_residual < v_hat).all(axis=0)

    num_winning_auctions = np.sum(winning_auctions_mask)

    if num_winning_auctions == 0:
        return np.nan

    array_conditional = array_residual[:, winning_auctions_mask]

    array_conditional_sorted = np.sort(array_conditional, axis=0)

    order_largest_bids = array_conditional_sorted[-order, :]

    return np.mean(order_largest_bids)
```

We can check the accuracy of our `evaluate_largest` method by comparing it with an analytical solution.

We find that the `evaluate_largest` method functions well

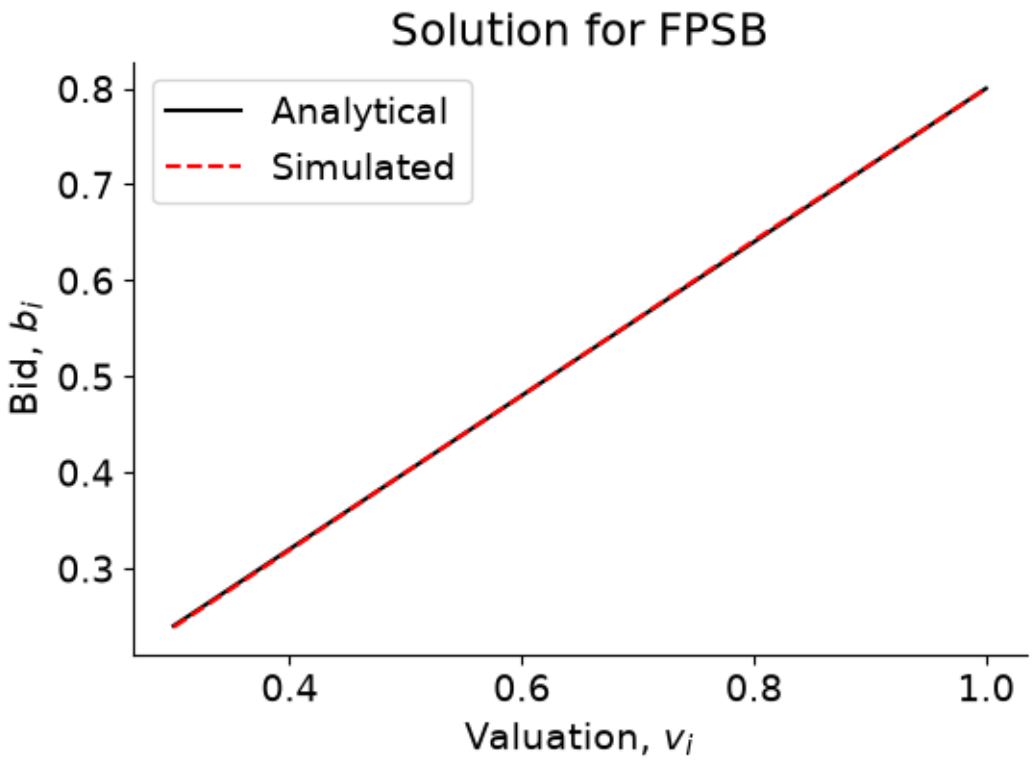
```
v_grid = np.linspace(0.3, 1, 8)
bid_analytical = b_star(v_grid, N)

# Redraw valuations
v = rng.uniform(0, 1, (N, R))
bid_simulated = [evaluate_largest(ii, v) for ii in v_grid]

fig, ax = plt.subplots(figsize=(6, 4))

ax.plot(v_grid, bid_analytical, '-', color='k', label='Analytical')
ax.plot(v_grid, bid_simulated, '--', color='r', label='Simulated')

ax.legend(loc='best')
ax.set_xlabel('Valuation, $v_i$')
ax.set_ylabel('Bid, $b_i$')
ax.set_title('Solution for FPSB')
sns.despine()
```



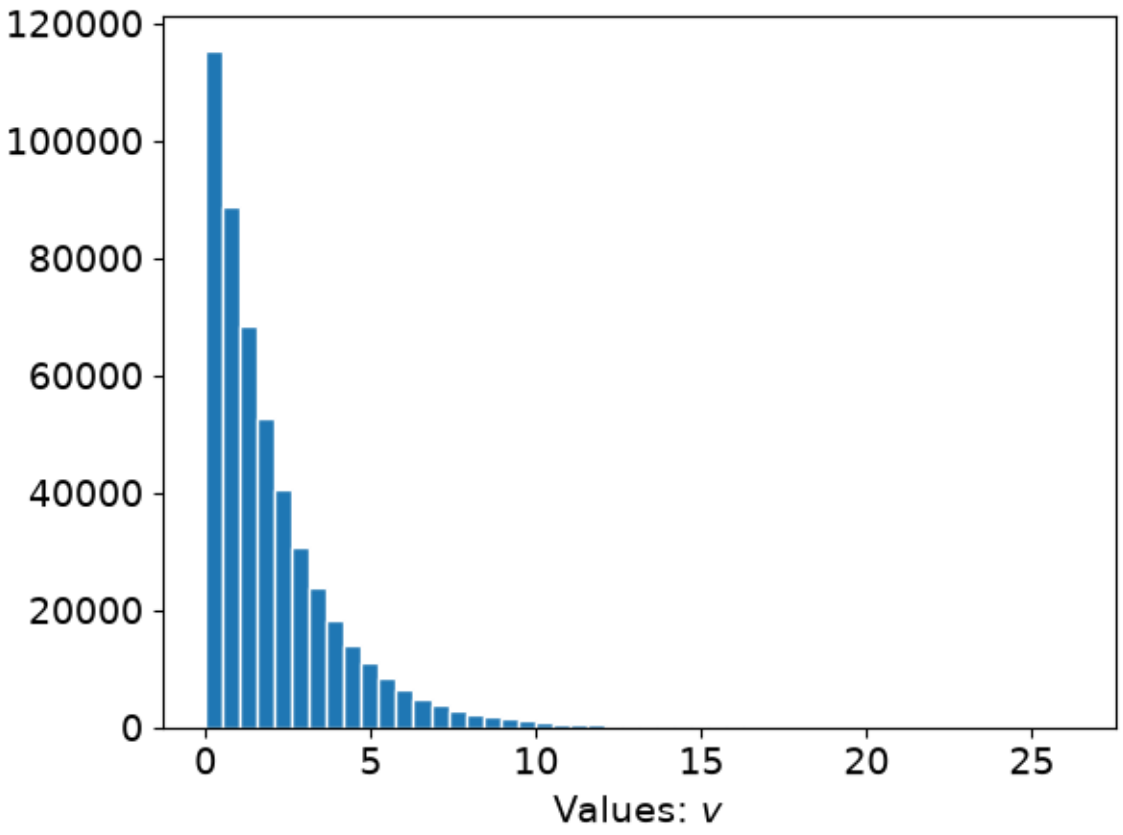
149.11 χ^2 Distribution

Let's try an example in which the distribution of private values is a χ^2 distribution.

We'll start by taking a look at a χ^2 distribution with the help of the following Python code:

```
rng = np.random.default_rng(1337)
v = rng.chisquare(df=2, size=(N * R,))

plt.hist(v, bins=50, edgecolor='w')
plt.xlabel('Values: $v$')
plt.show()
```



Now we'll get Python to construct a bid price function

```
rng = np.random.default_rng(1337)
v = rng.chisquare(df=2, size=(N, R))

# we compute the quantile of v as our grid
pct_quantile = np.linspace(0, 100, 101)[1:-1]
v_grid = np.percentile(v.flatten(), q=pct_quantile)

# nan values are returned for some low quantiles due to lack of observations
EV = [evaluate_largest(ii, v) for ii in v_grid]
```

```
# we insert 0 into our grid and bid price function as a complement
EV = np.insert(EV, 0, 0)
v_grid = np.insert(v_grid, 0, 0)

b_star_num = interp.interp1d(v_grid, EV, fill_value="extrapolate")
```

We check our bid price function by computing and visualizing the result.

```
pct_quantile_fine = np.linspace(0, 100, 1001)[1:-1]
v_grid_fine = np.percentile(v.flatten(), q=pct_quantile_fine)

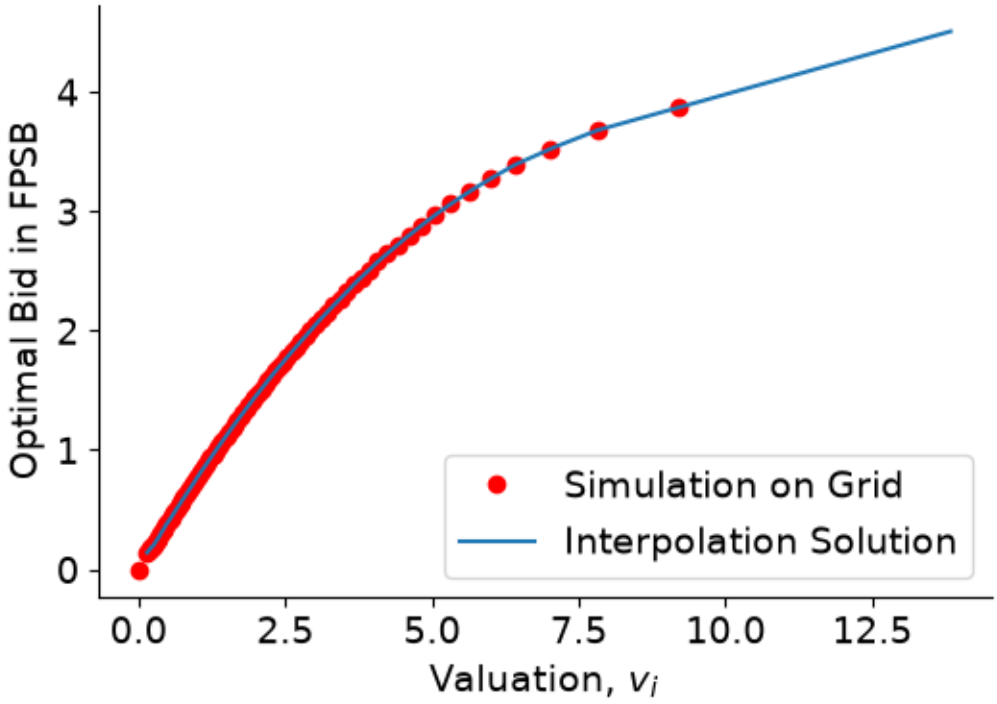
fig, ax = plt.subplots(figsize=(6, 4))

ax.plot(v_grid, EV, 'or', label='Simulation on Grid')
ax.plot(v_grid_fine, b_star_num(v_grid_fine),
                '-', label='Interpolation Solution')
```

(continues on next page)

(continued from previous page)

```
ax.legend(loc='best')
ax.set_xlabel('Valuation, $v_i$')
ax.set_ylabel('Optimal Bid in FPSB')
sns.despine()
```



Now we can use Python to compute the probability distribution of the price paid by the winning bidder

```
b = b_star_num(v)

idx = np.argsort(v, axis=0)
# same as np.sort(v, axis=0), except now we retain the idx
v = np.take_along_axis(v, idx, axis=0)
b = np.take_along_axis(b, idx, axis=0)

# highest bid
winner_pays_fpsb = b[-1, :]
# 2nd-highest valuation
winner_pays_spsb = v[-2, :]
```

```
fig, ax = plt.subplots(figsize=(6, 4))

for payment, label in zip([winner_pays_fpsb, winner_pays_spsb],
                          ['FPSB', 'SPSB']):
    print('The average payment of %s: %.4f. Std.: %.4f. Median: %.4f' % (
        label, payment.mean(), payment.std(), np.median(payment)))
    ax.hist(payment, density=True, alpha=0.6, label=label, bins=100)

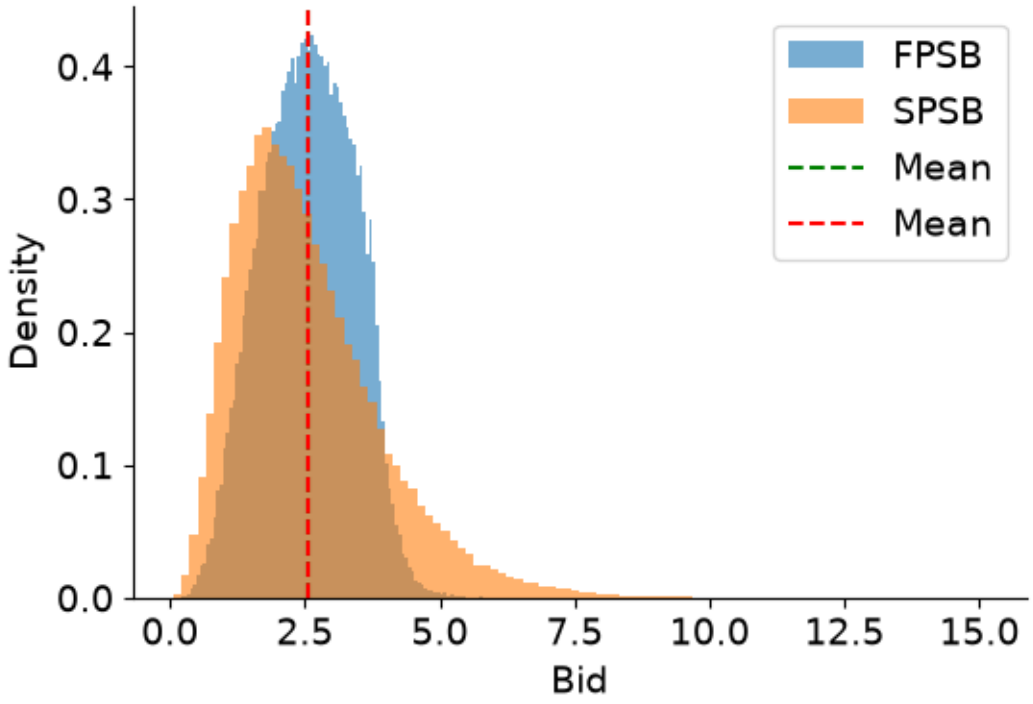
ax.axvline(winner_pays_fpsb.mean(), ls='--', c='g', label='Mean')
ax.axvline(winner_pays_spsb.mean(), ls='--', c='r', label='Mean')
```

(continues on next page)

(continued from previous page)

```
ax.legend(loc='best')
ax.set_xlabel('Bid')
ax.set_ylabel('Density')
sns.despine()
```

```
The average payment of FPSB: 2.5672. Std.: 0.8373. Median: 2.5794
The average payment of SPSB: 2.5634. Std.: 1.3637. Median: 2.3075
```



149.12 Code summary

We assemble the functions that we have used into a Python class

```
class bid_price_solution:

    def __init__(self, array):
        """
        A class that can plot the value distribution of bidders,
        compute the optimal bid price for bidders in FPSB
        and plot the distribution of winner's payment in both FPSB and SPSB

        Parameters:
        ----------

        array: 2 dimensional array of bidders' values in shape of (N, R),
               where N: number of players, R: number of auctions

        """
        self.value_mat = array.copy()
```

(continues on next page)

(continued from previous page)

```
        return None

    def plot_value_distribution(self):
        plt.hist(self.value_mat.flatten(), bins=50, edgecolor='w')
        plt.xlabel('Values: $v$')
        plt.show()

        return None

    def evaluate_largest(self, v_hat, order=1):
        N, R = self.value_mat.shape

        # drop the first row because we assume first row is the winner's bid
        array_residual = self.value_mat[1:, :].copy()

        winning_auctions_mask = (array_residual < v_hat).all(axis=0)

        num_winning_auctions = np.sum(winning_auctions_mask)

        if num_winning_auctions == 0:
            return np.nan

        array_conditional = array_residual[:, winning_auctions_mask]
        array_conditional_sorted = np.sort(array_conditional, axis=0)
        order_largest_bids = array_conditional_sorted[-order, :]

        return np.mean(order_largest_bids)

    def compute_optimal_bid_FPSB(self, plot=True):
        # we compute the quantile of v as our grid
        pct_quantile = np.linspace(0, 100, 101)[1:-1]
        v_grid = np.percentile(self.value_mat.flatten(), q=pct_quantile)

        # nan values are returned for some low quantiles due to lack of observations
        EV = [self.evaluate_largest(ii) for ii in v_grid]

        # we insert 0 into our grid and bid price function as a complement
        EV = np.insert(EV, 0, 0)
        v_grid = np.insert(v_grid, 0, 0)

        self.b_star_num = interp.interp1d(v_grid, EV,
                                          fill_value="extrapolate")

        if not plot:
            return None

        pct_quantile_fine = np.linspace(0, 100, 1001)[1:-1]
        v_grid_fine = np.percentile(self.value_mat.flatten(),
                                    q=pct_quantile_fine)

        fig, ax = plt.subplots(figsize=(6, 4))

        ax.plot(v_grid, EV, 'or', label='Simulation on Grid')
        ax.plot(v_grid_fine, self.b_star_num(v_grid_fine),
                             '-', label='Interpolation Solution')
```

(continues on next page)

(continued from previous page)

```
        ax.legend(loc='best')
        ax.set_xlabel('Valuation, $v_i$')
        ax.set_ylabel('Optimal Bid in FPSB')
        sns.despine()

        return None

    def plot_winner_payment_distribution(self):
        if not hasattr(self, 'b_star_num'):     # bids have not been computed yet
            self.compute_optimal_bid_FPSB(plot=False)

        self.b = self.b_star_num(self.value_mat)

        idx = np.argsort(self.value_mat, axis=0)
        # same as np.sort(v, axis=0), except now we retain the idx
        self.v = np.take_along_axis(self.value_mat, idx, axis=0)
        self.b = np.take_along_axis(self.b, idx, axis=0)

        # highest bid
        winner_pays_fpsb = self.b[-1, :]
        # 2nd-highest valuation
        winner_pays_spsb = self.v[-2, :]

        fig, ax = plt.subplots(figsize=(6, 4))

        for payment, label in zip([winner_pays_fpsb, winner_pays_spsb],
                                  ['FPSB', 'SPSB']):
            print('The average payment of %s: %.4f. Std.: %.4f. Median: %.4f' %
                  (label, payment.mean(), payment.std(), np.median(payment)))
            ax.hist(payment, density=True, alpha=0.6, label=label, bins=100)

        ax.axvline(winner_pays_fpsb.mean(), ls='--', c='g', label='Mean')
        ax.axvline(winner_pays_spsb.mean(), ls='--', c='r', label='Mean')

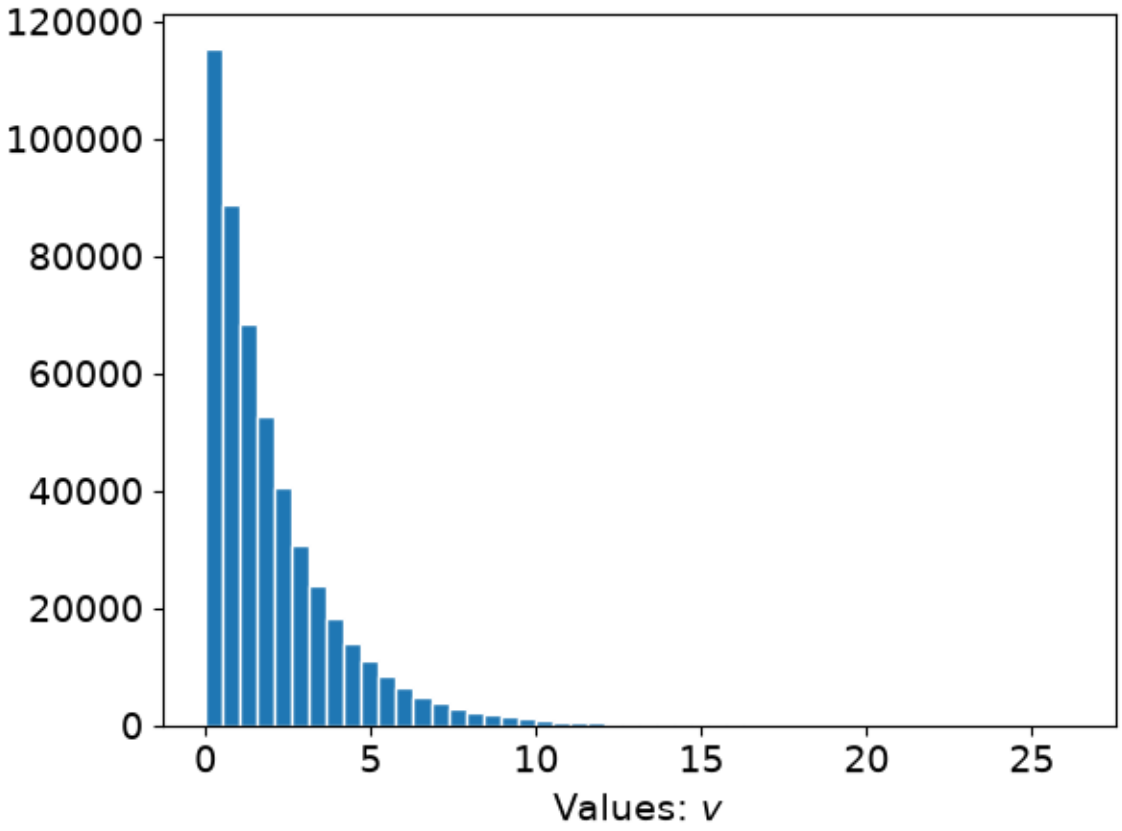
        ax.legend(loc='best')
        ax.set_xlabel('Bid')
        ax.set_ylabel('Density')
        sns.despine()

        return None
```

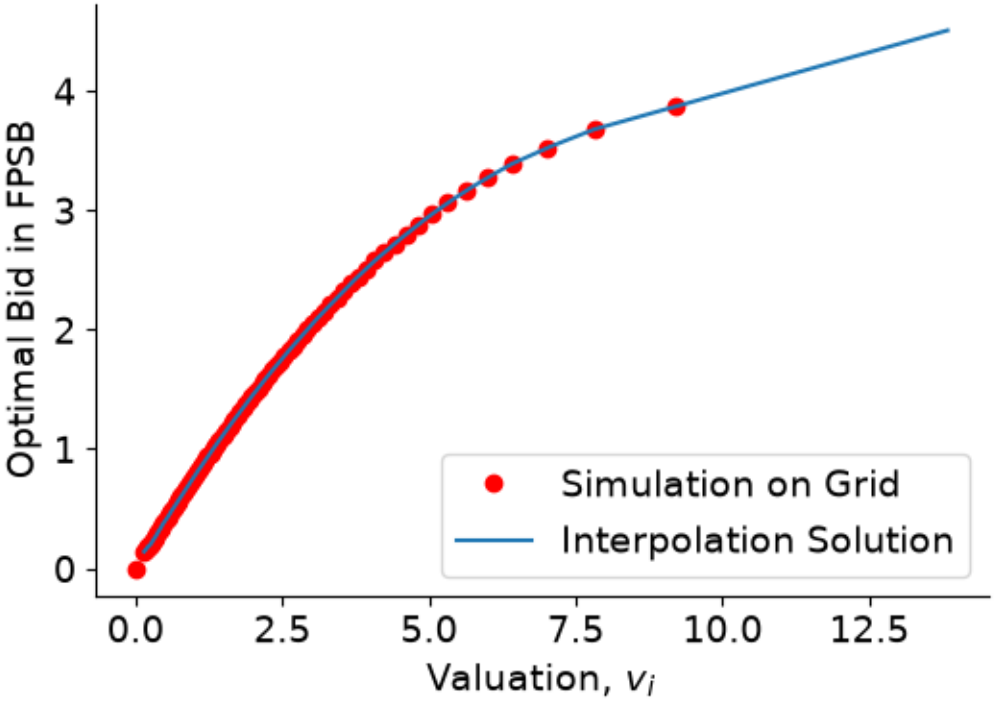
```
rng = np.random.default_rng(1337)
v = rng.chisquare(df=2, size=(N, R))

chi_squ_case = bid_price_solution(v)
```

```
chi_squ_case.plot_value_distribution()
```

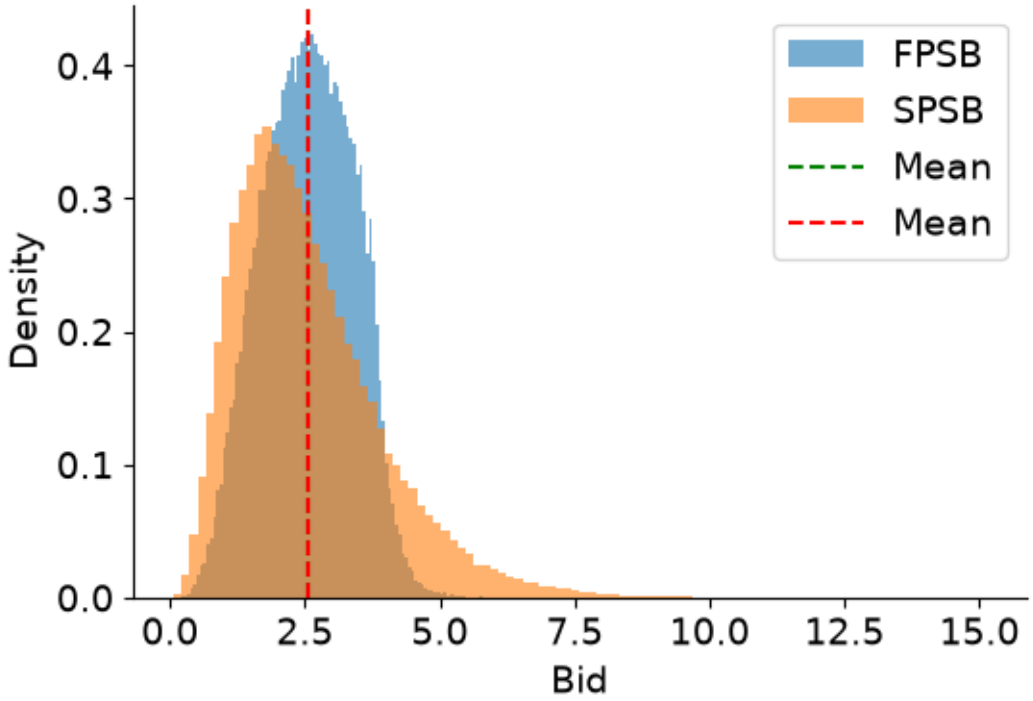



```
chi_squ_case.compute_optimal_bid_FPSB()
```



```
chi_squ_case.plot_winner_payment_distribution()
```

```
The average payment of FPSB: 2.5672. Std.: 0.8373. Median: 2.5794
The average payment of SPSB: 2.5634. Std.: 1.3637. Median: 2.3075
```



149.13 Exercises

Exercise 149.13.1

Verify the revenue equivalence theorem by simulation.

For $n = 2, 3, 5, 10$ bidders with valuations drawn independently from $U(0, 1)$, simulate many auctions and compute

1. the average payment of the winner in a FPSB auction, in which each bidder bids $\frac{n-1}{n} v_i$
2. the average payment of the winner in a SPSB auction, in which each bidder bids v_i

Compare both with the theoretical expected revenue $\frac{n-1}{n+1}$, and comment on how the seller's revenue changes with the number of bidders.

Solution

```
R_ex = 200_000
rng_ex = np.random.default_rng(1234)

print(f"{'n':>4}{'FPSB':>12}{'SPSB':>12}{'(n-1)/(n+1)':>14}")
for n in (2, 3, 5, 10):
    v_ex = np.sort(rng_ex.uniform(0, 1, (n, R_ex)), axis=0)
```

```
    fpsb = (n - 1)/n * v_ex[-1, :]       # winner's own bid
    spsb = v_ex[-2, :]                   # second highest valuation
    print(f"{n:>4}{fpsb.mean():>12.4f}{spsb.mean():>12.4f}{(n-1)/(n+1):>14.4f}")
```

```
   n        FPSB        SPSB   (n-1)/(n+1)
   2      0.3335      0.3340        0.3333
   3      0.5000      0.5000        0.5000
   5      0.6671      0.6676        0.6667
  10      0.8185      0.8185        0.8182
```

The two auctions raise the same expected revenue, and both converge to the highest possible valuation as n grows.

With more bidders, competition pushes the winning payment toward the top of the support of valuations.

Notice that the two auctions raise the same revenue on average even though the *distributions* of the winner's payment differ: in a FPSB auction the payment is a deterministic function of the winner's valuation, while in a SPSB auction it is the second-highest valuation, which is random given the winner's valuation.

Exercise 149.13.2

Equation (149.1) says that an optimal bid in a FPSB auction is $\mathbb{E}[y_i \mid y_i < v_i]$.

1. Show that this can be written

$$b(v) = v - \frac{\int_0^v F(x)^{n-1} dx}{F(v)^{n-1}}$$

   where F is the distribution function of a valuation.

2. Verify that this reduces to $\frac{n-1}{n} v$ when F is uniform on $[0, 1]$.

3. Evaluate the formula for a $\chi^2(2)$ distribution of valuations and compare it with the simulation-based bid function computed in the lecture.

Solution

The distribution function of $y_i = \max_{j \neq i} v_j$ is $\tilde{F}_{n-1}(y) = F(y)^{n-1}$.

Hence

$$\mathbb{E}[y \mid y < v] = \frac{1}{F(v)^{n-1}} \int_0^v y \, d\left[F(y)^{n-1}\right].$$

Integrating by parts,

$$\int_0^v y \, d\left[F(y)^{n-1}\right] = vF(v)^{n-1} - \int_0^v F(y)^{n-1} dy,$$

which gives the formula.

For $F(x) = x$ on $[0, 1]$ we get $b(v) = v - \frac{v^n/n}{v^{n-1}} = \frac{n-1}{n} v$.

The formula has a nice reading: a bidder shades her bid below her valuation by an amount that shrinks as the number of competitors grows.

```
from scipy.integrate import quad

def b_closed_form(v, F, n):
    "Optimal FPSB bid for a bidder with valuation v when rivals' values ~ F."
    shading = quad(lambda x: F(x)**(n - 1), 0, v)[0] / F(v)**(n - 1)
    return v - shading

# check against the analytical solution for the uniform case
print("uniform check")
for v0 in (0.3, 0.6, 0.9):
    print(f"  v = {v0}:  closed form {b_closed_form(v0, lambda x: x, N):.4f}, "
          f"analytical {b_star(v0, N):.4f}")
```

```
uniform check
  v = 0.3:  closed form 0.2400, analytical 0.2400
  v = 0.6:  closed form 0.4800, analytical 0.4800
  v = 0.9:  closed form 0.7200, analytical 0.7200
```

```
# now the chi-squared case studied in the lecture
F_chi2 = stats.chi2(df=2).cdf
v_test = np.percentile(v.flatten(), [10, 30, 50, 70, 90])

print(f"{'v':>8}{'closed form':>14}{'simulated':>12}")
for v0 in v_test:
    print(f"{v0:>8.3f}{b_closed_form(v0, F_chi2, N):>14.4f}"
          f"{float(b_star_num(v0)):>12.4f}")
```

```
       v   closed form   simulated
   0.210        0.1663      0.1825
   0.711        0.5513      0.5544
   1.383        1.0370      1.0433
   2.397        1.6965      1.6984
   4.607        2.8055      2.7974
```

The closed form and the simulation agree closely, which is a useful check on both.

Exercise 149.13.3

This exercise asks you to see *why* truthful bidding is a weakly dominant strategy in a SPSB auction but not in a FPSB auction.

Fix $n = 5$ and consider a bidder whose valuation is $v = 0.75$, facing rivals whose valuations are $U(0, 1)$.

1. In a SPSB auction the rivals bid truthfully. Compute this bidder's expected surplus as a function of her own bid b and plot it.
2. In a FPSB auction the rivals bid $\frac{n-1}{n} v_j$. Compute and plot her expected surplus as a function of b.
3. Where does each curve peak? What surplus does she earn in the FPSB auction if she bids her valuation?

Solution

In the SPSB auction she wins when $y < b$ and then pays y, so her expected surplus is

$$\int_0^b (v - y)\,(n-1) y^{n-2} dy.$$

Differentiating with respect to b gives $(v-b)(n-1)b^{n-2}$, which is positive for $b < v$ and negative for $b > v$, so $b = v$ is optimal.

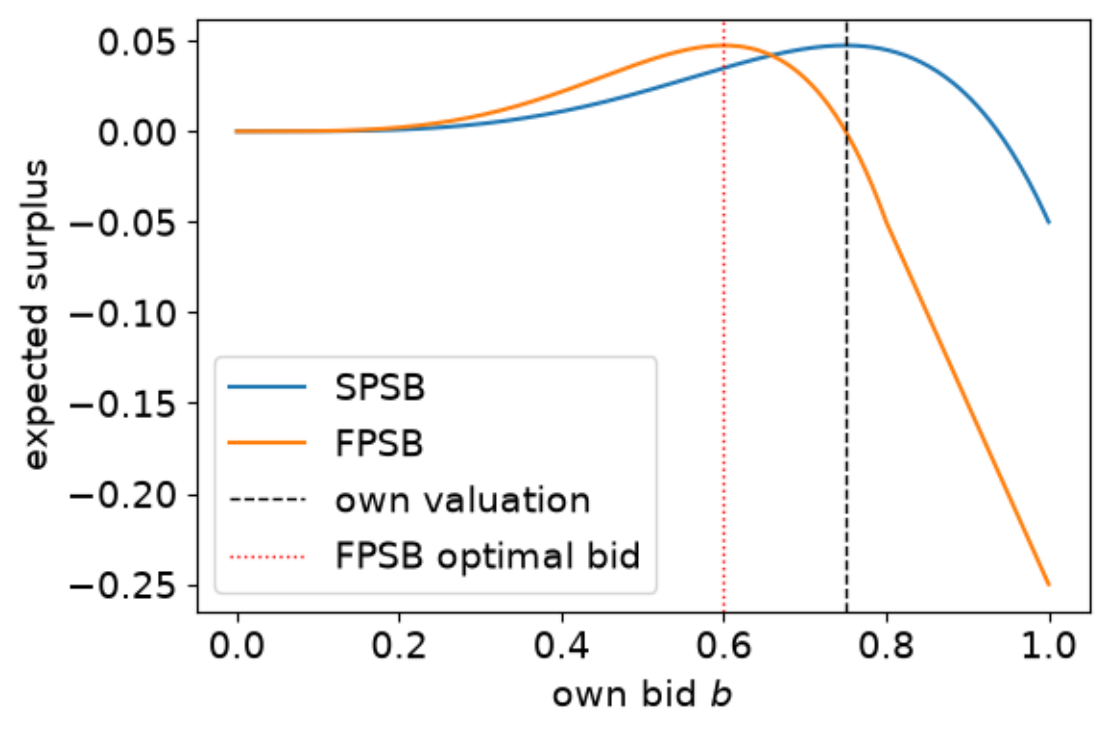
In the FPSB auction she wins when every rival's bid falls below b, which happens with probability $\left(\frac{nb}{n-1}\right)^{n-1}$, and then she pays b.

```
n_ex, v_own = 5, 0.75
bids = np.linspace(0, 1, 401)

spsb_surplus = [quad(lambda y: (v_own - y)*(n_ex - 1)*y**(n_ex - 2),
                     0, min(bb, 1))[0] for bb in bids]
fpsb_surplus = [(v_own - bb)*min(1, n_ex*bb/(n_ex - 1))**(n_ex - 1)
                for bb in bids]

fig, ax = plt.subplots(figsize=(6, 4))
ax.plot(bids, spsb_surplus, label='SPSB')
ax.plot(bids, fpsb_surplus, label='FPSB')
ax.axvline(v_own, ls='--', c='k', lw=1, label='own valuation')
ax.axvline((n_ex - 1)/n_ex*v_own, ls=':', c='r', lw=1, label='FPSB optimal bid')
ax.set_xlabel('own bid $b$')
ax.set_ylabel('expected surplus')
ax.legend()
plt.show()

print(f"SPSB surplus is maximized at b = {bids[int(np.argmax(spsb_surplus))]:.3f}")
print(f"FPSB surplus is maximized at b = {bids[int(np.argmax(fpsb_surplus))]:.3f}"
      f"  (theory: {(n_ex - 1)/n_ex*v_own:.3f})")
print(f"FPSB surplus from bidding one's valuation: {(v_own - v_own):.3f}")
```



```
SPSB surplus is maximized at b = 0.750
FPSB surplus is maximized at b = 0.600  (theory: 0.600)
FPSB surplus from bidding one's valuation: 0.000
```

The SPSB curve peaks exactly at the bidder's valuation.

The FPSB curve peaks strictly below it, and bidding one's valuation in a FPSB auction earns a surplus of exactly zero: the bidder wins more often, but pays her full valuation whenever she does.

Exercise 149.13.4

The revenue equivalence theorem requires bidders to be *risk neutral*.

Suppose instead that each bidder has utility $u(x) = x^\rho$ with $0 < \rho \leq 1$, so that $\rho < 1$ means risk aversion, and that valuations are $U(0, 1)$.

One can show that the symmetric equilibrium bid in a FPSB auction becomes

$$b(v) = \frac{n-1}{n-1+\rho} v.$$

1. Verify this numerically: for $n = 5$ and a bidder with $v = 0.8$, compute expected utility as a function of her own bid when rivals use this rule, and check where it is maximized.
2. Compute the seller's expected revenue in FPSB and SPSB for $\rho = 1, 0.6, 0.3$.
3. Explain the intuition.

Solution

```
def expected_utility(b, v_own, n, ρ):
    "Expected utility of bidding b when rivals bid (n-1)v/(n-1+ρ)."
    win_prob = np.minimum(1, b*(n - 1 + ρ)/(n - 1))**(n - 1)
    return win_prob * np.maximum(v_own - b, 0)**ρ

n_ex, v_own = 5, 0.8
grid = np.linspace(0.001, v_own, 2001)

print(f"{'ρ':>6}{'theory b*':>12}{'numerical':>12}")
for ρ in (1.0, 0.5, 0.2):
    theory = (n_ex - 1)*v_own/(n_ex - 1 + ρ)
    numerical = grid[int(np.argmax(expected_utility(grid, v_own, n_ex, ρ)))]
    print(f"{ρ:>6}{theory:>12.4f}{numerical:>12.4f}")
```

```
     ρ   theory b*   numerical
   1.0      0.6400      0.6398
   0.5      0.7111      0.7109
   0.2      0.7619      0.7620
```

```
rng_ra = np.random.default_rng(42)
v_ra = np.sort(rng_ra.uniform(0, 1, (n_ex, 200_000)), axis=0)

print(f"{'ρ':>6}{'FPSB revenue':>15}{'SPSB revenue':>15}")
for ρ in (1.0, 0.6, 0.3):
    fpsb = ((n_ex - 1)/(n_ex - 1 + ρ)) * v_ra[-1, :]
    print(f"{ρ:>6}{fpsb.mean():>15.4f}{v_ra[-2, :].mean():>15.4f}")
```

```
     ρ   FPSB revenue   SPSB revenue
1.0          0.6666         0.6668
0.6          0.7245         0.6668
0.3          0.7751         0.6668
```

With risk neutrality ($\rho = 1$) the two auctions raise the same revenue, as the theorem says.

With risk aversion ($\rho < 1$) the FPSB auction raises *more*.

The intuition is that in a FPSB auction, shading one's bid is a gamble: it raises the surplus conditional on winning but lowers the probability of winning.

A risk-averse bidder dislikes that gamble and so shades less, which transfers revenue to the seller.

In a SPSB auction the winner's payment does not depend on her own bid, so risk aversion changes nothing: bidding one's valuation remains weakly dominant, and the seller's revenue is unaffected.

149.14 Further reading

The second-price sealed-bid auction was proposed by Vickrey [1961].

For textbook treatments of the material in this lecture, see Krishna [2009] and Milgrom [2004].

Klemperer [1999] surveys the literature.

The revenue equivalence theorem in the general form sketched above is due to Myerson [1981] and Riley and Samuelson [1981].

Both auctions studied here are naturally described in terms of order statistics of bidders' valuations, a subject treated at length by David and Nagaraja [2003].

MULTIPLE GOOD ALLOCATION MECHANISMS

```
!pip install "prettytable<3.18"
```

150.1 Overview

This lecture describes two mechanisms for allocating n private goods ("houses") to m people ("buyers").

We assume that $m > n$ so that there are more potential buyers than there are houses.

Prospective buyers regard the houses as **substitutes**.

Buyer j attaches value v_{ij} to house i.

These values are **private**

- v_{ij} is known only to person j unless person j chooses to tell someone.

We require that a mechanism allocate **at most** one house to one prospective buyer.

We describe two distinct mechanisms

- A multiple rounds, ascending bid auction
- A sequential, "greedy" second-price mechanism administered by a benevolent social planner, inspired by the ideas of Groves [1973] and Clarke [1971]

Note

In 1994, the multiple rounds, ascending bid auction was actually used by Stanford University to sell leases to 9 lots on the Stanford campus to eligible faculty members.

We begin with overviews of the two mechanisms.

150.2 Ascending bids auction for multiple goods

An auction is administered by an **auctioneer**

The auctioneer has an $n \times 1$ vector r of reservation prices on the n houses.

The auctioneer sells house i only if the final price bid for it exceeds r_i

The auctioneer allocates all n houses **simultaneously**

The auctioneer does not know bidders' private values v_{ij}

There are multiple **rounds**

- during each round, active participants can submit bids on any of the n houses
- each bidder can bid on only one house during one round
- a person who was high bidder on a particular house in one round is understood to submit that same bid for the same house in the next round
- between rounds, a bidder who was not a high bidder can change the house on which he/she chooses to bid
- the auction ends when the price of no house changes from one round to the next
- all n houses are allocated after the final round
- house i is retained by the auctioneer if no prospective buyer offers more than r_i for the house

In this auction, person j never tells anyone else his/her private values v_{ij}

150.3 A benevolent planner

This mechanism is designed so that all prospective buyers voluntarily choose to reveal their private values to a **social planner** who uses them to construct a socially optimal allocation.

Among all feasible allocations, a **socially optimal allocation** maximizes the sum of private values across all prospective buyers.

The planner tells everyone in advance how he/she will allocate houses based on the matrix of values that prospective buyers report.

The mechanism is intended to give every prospective buyer an incentive to reveal his vector of private values to the planner.

After the planner receives everyone's vector of private values, the planner deploys a **sequential** algorithm to determine an **allocation** of houses and a set of **fees** that he charges awardees for the negative **externality** that their presence impose on other prospective buyers.

150.4 Comparing the two mechanisms

Remarkably, these two mechanisms can produce virtually identical allocations.

But "virtually identical" is not "identical", and the differences turn out to be instructive.

We'll discover that

- the ascending bid auction maximizes the total value of the allocation, and, as the bid increment ϵ becomes small, the prices that it sets converge to the payments prescribed by the **Vickrey-Clarke-Groves** (VCG) mechanism
- the greedy mechanism of the benevolent planner does *not* always maximize total value, and the prices that it charges are *not* VCG payments, so truthful reporting is not always in a buyer's interest

We construct Python code for both mechanisms.

We also work out some examples by hand or almost by hand.

Next, let's dive down into the details.

150.5 Ascending bid auction

150.5.1 Basic setting

We start with a more detailed description of the setting.

- A seller owns n houses that he wants to sell for the maximum possible amounts to a set of m prospective eligible buyers.
- The seller wants to sell at most one house to each potential buyer.
- There are m potential eligible buyers, identified by $j = [1, 2, \dots, m]$
  - Each potential buyer is permitted to buy at most one house.
  - Buyer j would be willing to pay at most v_{ij} for house i.
  - Buyer j knows $v_{ij}, i = 1, \dots, n$, but no one else does.
  - If buyer j pays p_i for house i, he enjoys surplus value $v_{ij} - p_i$.
  - Each buyer j wants to choose the i that maximizes his/her surplus value $v_{ij} - p_i$.
  - The seller wants to maximize $\sum_i p_i$.

The seller conducts a **simultaneous, multiple goods, ascending bid auction**.

Outcomes of the auction are

- An $n \times 1$ vector p of sales prices $p = [p_1, \dots, p_n]$ for the n houses.
- An $n \times m$ matrix Q of 0's and 1's, where $Q_{ij} = 1$ if and only if person j bought house i.
- An $n \times m$ matrix S of surplus values consisting of all zeros unless person j bought house i, in which case $S_{ij} = v_{ij} - p_i$

We describe rules for the auction it terms of **pseudo code**.

The pseudo code will provide a road map for writing Python code to implement the auction.

150.6 Pseudocode

Here is a quick sketch of a possible simple structure for our Python code

Inputs:

- n, m.
- an $n \times m$ non-negative matrix v of private values
- an $n \times 1$ vector r of seller-specified reservation prices
- the seller will not accept a price less than r_i for house i
- we are free to think of these reservation prices as private values of a fictitious $m + 1$ th buyer who does not actually participate in the auction
- initial bids can be thought of starting at r
- a scalar ϵ of seller-specified minimum price-bid increments

For each round of the auction, new bids on a house must be at least the prevailing highest bid so far **plus** ϵ

Auction Protocols

- the auction consists of a finite number of **rounds**
- in each round, a prospective buyer can bid on one and only one house
- after each round, a house is temporarily awarded to the person who made the highest bid for that house
  - temporarily winning bids on each house are announced
  - this sets the stage to move on to the next round
- a new round is held
  - bids for temporary winners from the previous round are again attached to the houses on which they bid; temporary winners of the last round leave their bids from the previous round unchanged
  - all other active prospective buyers must submit a new bid on some house
  - new bids on a house must be at least equal to the prevailing temporary price that won the last round **plus** ϵ
  - if a person does not submit a new bid and was also not a temporary winner from the previous round, that person must drop out of the auction permanently
  - for each house, the highest bid, whether it is a new bid or was the temporary winner from the previous round, is announced, with the person who made that new (temporarily) winning bid being (temporarily) awarded the house to start the next round
- rounds continue until no price on **any** house changes from the previous round
- houses are sold to the winning bidders from the final round at the prices that they bid

Outputs:

- an $n \times 1$ vector p of sales prices
- an $n \times m$ matrix S of surplus values consisting of all zeros unless person j bought house i, in which case $S_{ij} = v_{ij} - p_i$
- an $n \times (m+1)$ matrix Q of 0's and 1's that tells which buyer bought which house. (The last column accounts for unsold houses.)

Proposed buyer strategy:

In this pseudo code and the actual Python code below, we'll assume that all buyers choose to use the following strategy, often called **straightforward bidding**: in each round, bid on the house that currently offers the buyer the largest surplus.

Each buyer $j = 1, \ldots, m$ uses the same strategy.

Straightforward bidding is *myopic*: a buyer looks only at prevailing prices and ignores how his bid today might affect prices tomorrow.

That myopia is less restrictive than it might sound.

Our setting is an **assignment market** of the kind studied by Shapley and Shubik [1971]: each buyer wants at most one house, and values are private.

For such markets, Crawford and Knoer [1981] and Demange *et al.* [1986] showed that an ascending auction like ours, in which all buyers bid straightforwardly, converges to a competitive equilibrium of the assignment market.

It converges in particular to the *minimum* competitive equilibrium price vector, the one most favorable to buyers.

We'll return below to a striking consequence of this fact.

The strategy has the form:

- Let $\breve{p}^t$ be the $n \times 1$ vector of prevailing highest-bid prices at the beginning of round t
- Let $\epsilon > 0$ be the minimum bid increment specified by the seller

- For each prospective buyer j, compute the index of the best house to bid on during round t, namely $\hat{i}_t = \text{argmax}_i\{[v_{ij} - \check{p}_i^t - \epsilon]\}$
- If $\max_i\{[v_{ij} - \check{p}_i^t - \epsilon]\} \leq 0$, person j permanently drops out of the auction at round t
- If $v_{\hat{i}_t,j} - \check{p}_{\hat{i}_t}^t - \epsilon > 0$, person j bids $\check{p}_{\hat{i}_t}^t + \epsilon$ on house $\hat{i}_t$

Resolving ambiguities: The protocols we have described so far leave open two possible sources of ambiguity.

(1) **The optimal bid choice for buyers in each round.** It is possible that a buyer has the same surplus value for multiple houses. The argmax function in Python always returns the first argmax element. We instead prefer to randomize among such houses. For that reason, we write our own argmax function below.

(2) **Seller's choice of winner if same price bid cast by several buyers.** To resolve this ambiguity, we use the np.random.choice function below.

Given the randomness in outcomes, it is possible that different allocations of houses could emerge from the same inputs.

However, this will happen only when the bid price increment ϵ is nonnegligible.

```
import numpy as np
import prettytable as pt

np.random.seed(100)
```

```
np.set_printoptions(precision=3, suppress=True)
```

150.7 An example

Before building a Python class, let's step by step solve things almost "by hand" to grasp how the auction proceeds.

A step-by-step procedure also helps reduce bugs, especially when the value matrix is peculiar (e.g. the differences between values are negligible, a column containing identical values or multiple buyers have the same valuation etc.).

Fortunately, our auction behaves well and robustly with various peculiar matrices.

We provide some examples later in this lecture.

```
v = np.array([[8, 5, 9, 4],
              [4, 11, 7, 4],
              [9, 7, 6, 4]])
n, m = v.shape
r = np.array([2, 1, 0])
ϵ = 1
p = r.copy()
buyer_list = np.arange(m)
house_list = np.arange(n)
```

```
v
```

```
array([[ 8,  5,  9,  4],
       [ 4, 11,  7,  4],
       [ 9,  7,  6,  4]])
```

Remember that column indexes j indicate buyers and row indexes i indicate houses.

The above value matrix v is peculiar in the sense that Buyer 3 (indexed from 0) puts the same value 4 on every house being sold.

Maybe buyer 3 is a bureaucrat who purchases these house simply by following instructions from his superior.

```
r
```

```
array([2, 1, 0])
```

```
def find_argmax_with_randomness(v):
    """
    We build our own version of argmax function such that the argmax index will be␣
 ↪returned randomly
    when there are multiple maximum values. This function is similar to np.argmax(v,
 ↪axis=0)

    Parameters:
    ----------
    v: 2 dimensional np.array

    """

    n, m = v.shape
    index_array = np.arange(n)
    result=[]

    for ii in range(m):
        max_value = v[:,ii].max()
        result.append(np.random.choice(index_array[v[:,ii] == max_value]))

    return np.array(result)
```

```
def present_dict(dt):
    """
    A function that present the information in table.

    Parameters:
    ----------
    dt: dictionary.

    """

    ymtb = pt.PrettyTable()
    ymtb.field_names = ['House Number', *dt.keys()]
    ymtb.add_row(['Buyer', *dt.values()])
    print(ymtb)
```

Check Kick Off Condition

```
def check_kick_off_condition(v, r, ϵ):
    """
    A function that checks whether the auction could be initiated given the␣
 ↪reservation price and value matrix.
    To avoid the situation that the reservation prices are so high that no one would␣
 ↪even bid in the first round.

    Parameters:
    ----------
    v : value matrix of the shape (n,m).
```

(continues on next page)

(continued from previous page)

```
    r: the reservation price

    ϵ: the minimum price increment in each round

    """

    n, m = v.shape

    # we convert the price vector to a matrix in the same shape as value matrix to␣
↪facilitate subtraction
    p_start = (ϵ+r)[:,None] @ np.ones(m)[None,:]

    surplus_value = v - p_start
    buyer_decision = (surplus_value > 0).any(axis = 0)
    return buyer_decision.any()
```

```
check_kick_off_condition(v, r, ϵ)
```

```
np.True_
```

150.7.1 Round 1

submit bid

```
def submit_initial_bid(p_initial, ϵ, v):
    """
    A function that describes the bid information in the first round.

    Parameters:
    ----------
    p_initial: the price (or the reservation prices) at the beginning of auction.

    v: the value matrix

    ϵ: the minimum price increment in each round

    Returns:
    ----------
    p: price array after this round of bidding

    bid_info: a dictionary that contains bidding information (house number as keys␣
↪and buyer as values).

    """

    n, m = v.shape
    buyer_list = np.arange(m)

    p = p_initial.copy()
    p_start_mat = (ϵ + p)[:,None] @ np.ones(m)[None,:]
    surplus_value = v - p_start_mat

    # we only care about active buyers who have positive surplus values
    active_buyer_diagnosis = (surplus_value > 0).any(axis = 0)
```

(continues on next page)

(continued from previous page)

```
    active_buyer_list = buyer_list[active_buyer_diagnosis]
    active_buyer_surplus_value = surplus_value[:,active_buyer_diagnosis]
    active_buyer_choice = find_argmax_with_randomness(active_buyer_surplus_value)
    # choice means the favourite houses given the current price and ϵ

    # we only retain the unique house index because prices increase once at one round
    house_bid =  list(set(active_buyer_choice))
    p[house_bid] += ϵ

    bid_info = {}
    for house_num in house_bid:
        bid_info[house_num] = active_buyer_list[active_buyer_choice == house_num]

    return p, bid_info
```

```
p, bid_info = submit_initial_bid(p, ϵ, v)
```

```
p
```

```
array([3, 2, 1])
```

```
present_dict(bid_info)
```

```
+--------------+-----+-----+-------+
| House Number |  0  |  1  |   2   |
+--------------+-----+-----+-------+
|    Buyer     | [2] | [1] | [0 3] |
+--------------+-----+-----+-------+
```

check terminal condition

Notice that two buyers bid for house 2 (indexed from 0).

Because the auction protocol does not specify a selection rule in this case, we simply select a winner **randomly**.

This is reasonable because the seller can't distinguish these buyers and doesn't know the valuation of each buyer.

It is both convenient and practical for him to just pick a winner randomly.

There is a 50% probability that Buyer 3 is chosen as the winner for house 2, although he values it less than buyer 0.

In this case, buyer 0 has to bid one more time with a higher price, which crowds out Buyer 3.

Therefore, final price could be 3 or 4, depending on the winner in the last round.

```
def check_terminal_condition(bid_info, p, v, ϵ):
    """
    A function that checks whether the auction ends.

    Recall that the auction ends when either losers have non-positive surplus values␣
↪for each house
    or there is no loser (every buyer gets a house).

    Parameters:
    ----------
    bid_info: a dictionary that contains bidding information of house numbers (as␣
↪keys) and buyers (as values).
```

(continues on next page)

(continued from previous page)

```
    p: np.array. price array of houses

    v: value matrix

    ϵ: the minimum price increment in each round

    Returns:
    ----------
    allocation: a dictionary that describes how the houses bid are assigned.

    winner_list: a list of winners

    loser_list: a list of losers

    """

    n, m = v.shape
    buyer_list = np.arange(m)

    # there may be several buyers bidding one house, we choose a winner randomly
    winner_list=[np.random.choice(bid_info[ii]) for ii in bid_info.keys()]

    allocation = {house_num:winner for house_num,winner in zip(bid_info.keys(),winner_
↪list)}

    loser_set = set(buyer_list).difference(set(winner_list))
    loser_list = list(loser_set)
    loser_num = len(loser_list)

    if loser_num == 0:
        print('The auction ends because every buyer gets one house.')
        return allocation,winner_list,loser_list

    p_mat = (ϵ + p)[:,None] @ np.ones(loser_num)[None,:]
    loser_surplus_value = v[:,loser_list] - p_mat
    loser_decision = (loser_surplus_value > 0).any(axis = 0)

    print('auction ends:', not loser_decision.any())

    return allocation,winner_list,loser_list
```

```
allocation,winner_list,loser_list = check_terminal_condition(bid_info, p, v, ϵ)
```

```
auction ends: False
```

```
present_dict(allocation)
```

```
+--------------+---+---+---+
| House Number | 0 | 1 | 2 |
+--------------+---+---+---+
|    Buyer     | 2 | 1 | 0 |
+--------------+---+---+---+
```

```
winner_list
```

```
[np.int64(2), np.int64(1), np.int64(0)]
```

```
loser_list
```

```
[np.int64(3)]
```

150.7.2 Round 2

From the second round on, the auction proceeds differently from the first round.

Now only active losers (those who have positive surplus values) have an incentive to submit bids to displace temporary winners from the previous round.

```
def submit_bid(loser_list, p, ϵ, v, bid_info):
    """
    A function that executes the bid operation after the first round.
    After the first round, only active losers would cast a new bid with price as old␣
↪price + increment.
    By such bid, winners of last round are replaced by the active losers.

    Parameters:
    ----------
    loser_list: a list that includes the indexes of losers

    p: np.array. price array of houses

    ϵ: minimum increment of bid price

    v: value matrix

    bid_info: a dictionary that contains bidding information of house numbers (as␣
↪keys) and buyers (as values).

    Returns:
    ----------
    p_end: a price array after this round of bidding

    bid_info: a dictionary that contains updated bidding information.

    """

    p_end=p.copy()

    loser_num = len(loser_list)
    p_mat = (ϵ + p_end)[:,None] @ np.ones(loser_num)[None,:]
    loser_surplus_value = v[:,loser_list] - p_mat
    loser_decision = (loser_surplus_value > 0).any(axis = 0)

    active_loser_list = np.array(loser_list)[loser_decision]
    active_loser_surplus_value = loser_surplus_value[:,loser_decision]
    active_loser_choice = find_argmax_with_randomness(active_loser_surplus_value)

    # we retain the unique house index and increasing the corresponding bid price
```

(continues on next page)

(continued from previous page)

```
    house_bid = list(set(active_loser_choice))
    p_end[house_bid] += ϵ

    # we record the bidding information from active losers
    bid_info_active_loser = {}
    for house_num in house_bid:
        bid_info_active_loser[house_num] = active_loser_list[active_loser_choice ==␣
↪house_num]

    # we update the bidding information according to the bidding from actice losers
    for house_num in bid_info_active_loser.keys():
        bid_info[house_num] = bid_info_active_loser[house_num]

    return p_end,bid_info
```

```
p,bid_info = submit_bid(loser_list, p, ϵ, v, bid_info)
```

```
p
```

```
array([3, 2, 2])
```

```
present_dict(bid_info)
```

```
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Buyer     | [2] | [1] | [3] |
+--------------+-----+-----+-----+
```

```
allocation,winner_list,loser_list = check_terminal_condition(bid_info, p, v, ϵ)
```

```
auction ends: False
```

```
present_dict(allocation)
```

```
+--------------+---+---+---+
| House Number | 0 | 1 | 2 |
+--------------+---+---+---+
|    Buyer     | 2 | 1 | 3 |
+--------------+---+---+---+
```

150.7.3 Later rounds

Every remaining round repeats exactly the same two steps, so rather than writing them out one at a time we iterate until no price changes.

Watch for two things as the rounds unfold: a buyer who is outbid switches to whichever house now offers him the largest surplus, and the auction stops as soon as a round passes in which no losing buyer is willing to raise any price.

```
round_number = 2
```

(continues on next page)

(continued from previous page)

```
while True:
    round_number += 1
    p_previous = p.copy()

    p, bid_info = submit_bid(loser_list, p, ϵ, v, bid_info)
    allocation, winner_list, loser_list = check_terminal_condition(bid_info, p, v, ϵ)

    print(f'round {round_number}: prices {p}')
    present_dict(allocation)

    if np.array_equal(p, p_previous):
        print('no price changed during this round, so the auction ends')
        break
```

```
auction ends: False
round 3: prices [3 2 3]
+--------------+---+---+---+
| House Number | 0 | 1 | 2 |
+--------------+---+---+---+
|    Buyer     | 2 | 1 | 0 |
+--------------+---+---+---+
auction ends: False
round 4: prices [3 3 3]
+--------------+---+---+---+
| House Number | 0 | 1 | 2 |
+--------------+---+---+---+
|    Buyer     | 2 | 3 | 0 |
+--------------+---+---+---+
auction ends: True
round 5: prices [3 4 3]
+--------------+---+---+---+
| House Number | 0 | 1 | 2 |
+--------------+---+---+---+
|    Buyer     | 2 | 1 | 0 |
+--------------+---+---+---+
auction ends: True
round 6: prices [3 4 3]
+--------------+---+---+---+
| House Number | 0 | 1 | 2 |
+--------------+---+---+---+
|    Buyer     | 2 | 1 | 0 |
+--------------+---+---+---+
no price changed during this round, so the auction ends
```

```
# as for the houses unsold

house_unsold_list = list(set(house_list).difference(set(allocation.keys())))
house_unsold_list
```

```
[]
```

```
total_revenue = p[list(allocation.keys())].sum()
total_revenue
```

```
np.int64(10)
```

150.8 A Python class

Above we simulated an ascending bid auction step by step.

When defining functions, we repeatedly computed some intermediate objects because our Python function loses track of variables once the function is executed.

That of course led to redundancy in our code

It is much more efficient to collect all of the aforementioned code into a class that records information about all rounds.

```
class ascending_bid_auction:

    def __init__(self, v, r, ϵ):
        """
        A class that simulates an ascending bid auction for houses.

        Given buyers' value matrix, sellers' reservation prices and minimum increment␣
↪of bid prices,
        this class can execute an ascending bid auction and present information round␣
↪by round until the end.

        Parameters:
        ----------
        v: 2 dimensional value matrix

        r: np.array of reservation prices

        ϵ: minimum increment of bid price

        """

        self.v = v.copy()
        self.n,self.m = self.v.shape
        # prices are stored as floats so that a bid increment ϵ < 1 is allowed
        self.r = np.asarray(r, dtype=float)
        self.ϵ = ϵ
        self.p = self.r.copy()
        self.buyer_list = np.arange(self.m)
        self.house_list = np.arange(self.n)
        self.log = print          # replaced by a no-op when verbose=False
        self.bid_info_history = []
        self.allocation_history = []
        self.winner_history = []
        self.loser_history = []


    def find_argmax_with_randomness(self, v):
        n,m = v.shape
        index_array = np.arange(n)
        result=[]

        for ii in range(m):
            max_value = v[:,ii].max()
```

(continues on next page)

(continued from previous page)

```
            result.append(np.random.choice(index_array[v[:,ii] == max_value]))

        return np.array(result)


    def check_kick_off_condition(self):
        # we convert the price vector to a matrix in the same shape as value matrix␣
↪to facilitate subtraction
        p_start = (self.ϵ + self.r)[:,None] @ np.ones(self.m)[None,:]
        self.surplus_value = self.v - p_start
        buyer_decision = (self.surplus_value > 0).any(axis = 0)
        return buyer_decision.any()


    def submit_initial_bid(self):
        # we intend to find the optimal choice of each buyer
        p_start_mat = (self.ϵ + self.p)[:,None] @ np.ones(self.m)[None,:]
        self.surplus_value = self.v - p_start_mat

        # we only care about active buyers who have positive surplus values
        active_buyer_diagnosis = (self.surplus_value > 0).any(axis = 0)
        active_buyer_list = self.buyer_list[active_buyer_diagnosis]
        active_buyer_surplus_value = self.surplus_value[:,active_buyer_diagnosis]
        active_buyer_choice = self.find_argmax_with_randomness(active_buyer_surplus_
↪value)

        # we only retain the unique house index because prices increase once at one␣
↪round
        house_bid =  list(set(active_buyer_choice))
        self.p[house_bid] += self.ϵ

        bid_info = {}
        for house_num in house_bid:
            bid_info[house_num] = active_buyer_list[active_buyer_choice == house_num]
        self.bid_info_history.append(bid_info)

        self.log('The bid information is')
        ymtb = pt.PrettyTable()
        ymtb.field_names = ['House Number', *bid_info.keys()]
        ymtb.add_row(['Buyer', *bid_info.values()])
        self.log(ymtb)

        self.log('The bid prices for houses are')
        ymtb = pt.PrettyTable()
        ymtb.field_names = ['House Number', *self.house_list]
        ymtb.add_row(['Price', *self.p])
        self.log(ymtb)

        self.winner_list=[np.random.choice(bid_info[ii]) for ii in bid_info.keys()]
        self.winner_history.append(self.winner_list)

        self.allocation = {house_num:[winner] for house_num,winner in zip(bid_info.
↪keys(),self.winner_list)}
        self.allocation_history.append(self.allocation)

        loser_set = set(self.buyer_list).difference(set(self.winner_list))
        self.loser_list = list(loser_set)
```

(continues on next page)

(continued from previous page)

```
        self.loser_history.append(self.loser_list)

        self.log('The winners are')
        self.log(self.winner_list)

        self.log('The losers are')
        self.log(self.loser_list)
        self.log('\n')


    def check_terminal_condition(self):
        loser_num = len(self.loser_list)

        if loser_num == 0:
            self.log('The auction ends because every buyer gets one house.')
            self.log('\n')
            return True

        p_mat = (self.ϵ + self.p)[:,None] @ np.ones(loser_num)[None,:]
        self.loser_surplus_value = self.v[:,self.loser_list] - p_mat
        self.loser_decision = (self.loser_surplus_value > 0).any(axis = 0)

        return ~(self.loser_decision.any())


    def submit_bid(self):
        bid_info = self.allocation_history[-1].copy()  # we only record the bid info␣
↪of winner

        loser_num = len(self.loser_list)
        p_mat = (self.ϵ + self.p)[:,None] @ np.ones(loser_num)[None,:]
        self.loser_surplus_value = self.v[:,self.loser_list] - p_mat
        self.loser_decision = (self.loser_surplus_value > 0).any(axis = 0)

        active_loser_list = np.array(self.loser_list)[self.loser_decision]
        active_loser_surplus_value = self.loser_surplus_value[:,self.loser_decision]
        active_loser_choice = self.find_argmax_with_randomness(active_loser_surplus_
↪value)

        # we retain the unique house index and increasing the corresponding bid price
        house_bid = list(set(active_loser_choice))
        self.p[house_bid] += self.ϵ

        # we record the bidding information from active losers
        bid_info_active_loser = {}
        for house_num in house_bid:
            bid_info_active_loser[house_num] = active_loser_list[active_loser_choice␣
↪== house_num]

        # we update the bidding information according to the bidding from actice␣
↪losers
        for house_num in bid_info_active_loser.keys():
            bid_info[house_num] = bid_info_active_loser[house_num]
        self.bid_info_history.append(bid_info)

        self.log('The bid information is')
        ymtb = pt.PrettyTable()
```

(continues on next page)

(continued from previous page)

```
        ymtb.field_names = ['House Number', *bid_info.keys()]
        ymtb.add_row(['Buyer', *bid_info.values()])
        self.log(ymtb)

        self.log('The bid prices for houses are')
        ymtb = pt.PrettyTable()
        ymtb.field_names = ['House Number', *self.house_list]
        ymtb.add_row(['Price', *self.p])
        self.log(ymtb)

        self.winner_list=[np.random.choice(bid_info[ii]) for ii in bid_info.keys()]
        self.winner_history.append(self.winner_list)

        self.allocation = {house_num:[winner] for house_num,winner in zip(bid_info.
↪keys(),self.winner_list)}
        self.allocation_history.append(self.allocation)

        loser_set = set(self.buyer_list).difference(set(self.winner_list))
        self.loser_list = list(loser_set)
        self.loser_history.append(self.loser_list)

        self.log('The winners are')
        self.log(self.winner_list)

        self.log('The losers are')
        self.log(self.loser_list)
        self.log('\n')


    def start_auction(self, verbose=True):
        # set verbose=False to run the auction without printing a report
        self.log = print if verbose else (lambda *args, **kwargs: None)
        self.log('The Ascending Bid Auction for Houses')
        self.log('\n')

        self.log('Basic Information: %d houses, %d buyers'%(self.n, self.m))

        self.log('The valuation matrix is as follows')
        ymtb = pt.PrettyTable()
        ymtb.field_names = ['Buyer Number', *(np.arange(self.m))]
        for ii in range(self.n):
            ymtb.add_row(['House %d'%(ii), *self.v[ii,:]])
        self.log(ymtb)

        self.log('The reservation prices for houses are')
        ymtb = pt.PrettyTable()
        ymtb.field_names = ['House Number', *self.house_list]
        ymtb.add_row(['Price', *self.r])
        self.log(ymtb)
        self.log('The minimum increment of bid price is %.2f' % self.ϵ)
        self.log('\n')

        ctr = 1
        if self.check_kick_off_condition():
            self.log('Auction starts successfully')
            self.log('\n')
            self.log('Round %d'% ctr)
```

(continues on next page)

(continued from previous page)

```
            self.submit_initial_bid()

            while True:
                if self.check_terminal_condition():
                    self.log('Auction ends')
                    self.log('\n')

                    self.log('The final result is as follows')
                    self.log('\n')
                    self.log('The allocation plan is')
                    ymtb = pt.PrettyTable()
                    ymtb.field_names = ['House Number', *self.allocation.keys()]
                    ymtb.add_row(['Buyer', *self.allocation.values()])
                    self.log(ymtb)

                    self.log('The bid prices for houses are')
                    ymtb = pt.PrettyTable()
                    ymtb.field_names = ['House Number', *self.house_list]
                    ymtb.add_row(['Price', *self.p])
                    self.log(ymtb)

                    self.log('The winners are')
                    self.log(self.winner_list)

                    self.log('The losers are')
                    self.log(self.loser_list)

                    self.house_unsold_list = list(set(self.house_list).
↪difference(set(self.allocation.keys())))
                    self.log('The houses unsold are')
                    self.log(self.house_unsold_list)

                    self.total_revenue = self.p[list(self.allocation.keys())].sum()
                    self.log('The total revenue is %.2f' % self.total_revenue)

                    break

                ctr += 1
                self.log('Round %d'% ctr)
                self.submit_bid()

            # we compute the surplus matrix S and the quantity matrix X as required␣
↪in 1.1
            self.S = np.zeros((self.n, self.m))
            for ii,jj in zip(self.allocation.keys(),self.allocation.values()):
                self.S[ii,jj] = self.v[ii,jj] - self.p[ii]

            self.Q = np.zeros((self.n, self.m + 1))  # the last column records the␣
↪houses unsold
            for ii,jj in zip(self.allocation.keys(),self.allocation.values()):
                self.Q[ii,jj] = 1
            for ii in self.house_unsold_list:
                self.Q[ii,-1] = 1

            # we sort the allocation result by the house number
            house_sold_list = list(self.allocation.keys())
```

(continues on next page)

(continued from previous page)

```
            house_sold_list.sort()

            dict_temp = {}
            for ii in house_sold_list:
                dict_temp[ii] = self.allocation[ii]
            self.allocation = dict_temp

        else:
            self.log('The auction can not start because of high reservation prices')
```

Let's use our class to conduct the auction described in one of the above examples.

```
v = np.array([[8,5,9,4],[4,11,7,4],[9,7,6,4]])
r = np.array([2,1,0])
ϵ = 1

auction_1 = ascending_bid_auction(v, r, ϵ)

auction_1.start_auction()
```

```
The Ascending Bid Auction for Houses


Basic Information: 3 houses, 4 buyers
The valuation matrix is as follows
+--------------+---+----+---+---+
| Buyer Number | 0 | 1  | 2 | 3 |
+--------------+---+----+---+---+
|   House 0    | 8 | 5  | 9 | 4 |
|   House 1    | 4 | 11 | 7 | 4 |
|   House 2    | 9 | 7  | 6 | 4 |
+--------------+---+----+---+---+
The reservation prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 2.0 | 1.0 | 0.0 |
+--------------+-----+-----+-----+
The minimum increment of bid price is 1.00


Auction starts successfully


Round 1
The bid information is
+--------------+-----+-----+-------+
| House Number |  0  |  1  |   2   |
+--------------+-----+-----+-------+
|    Buyer     | [2] | [1] | [0 3] |
+--------------+-----+-----+-------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 2.0 | 1.0 |
```

(continues on next page)

(continued from previous page)

```
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[np.int64(3)]


Round 2
The bid information is
+--------------+---------------+---------------+-----+
| House Number |       0       |       1       |  2  |
+--------------+---------------+---------------+-----+
|    Buyer     | [np.int64(2)] | [np.int64(1)] | [3] |
+--------------+---------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 2.0 | 2.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(3)]
The losers are
[np.int64(0)]


Round 3
The bid information is
+--------------+---------------+---------------+-----+
| House Number |       0       |       1       |  2  |
+--------------+---------------+---------------+-----+
|    Buyer     | [np.int64(2)] | [np.int64(1)] | [0] |
+--------------+---------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 2.0 | 3.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[np.int64(3)]


Round 4
The bid information is
+--------------+---------------+-----+---------------+
| House Number |       0       |  1  |       2       |
+--------------+---------------+-----+---------------+
|    Buyer     | [np.int64(2)] | [3] | [np.int64(0)] |
+--------------+---------------+-----+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 3.0 | 3.0 |
```

(continues on next page)

(continued from previous page)

```
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(3), np.int64(0)]
The losers are
[np.int64(1)]


Round 5
The bid information is
+--------------+---------------+-----+---------------+
| House Number |       0       |  1  |       2       |
+--------------+---------------+-----+---------------+
|    Buyer     | [np.int64(2)] | [1] | [np.int64(0)] |
+--------------+---------------+-----+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 4.0 | 3.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[np.int64(3)]


Auction ends


The final result is as follows


The allocation plan is
+--------------+---------------+---------------+---------------+
| House Number |       0       |       1       |       2       |
+--------------+---------------+---------------+---------------+
|    Buyer     | [np.int64(2)] | [np.int64(1)] | [np.int64(0)] |
+--------------+---------------+---------------+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 4.0 | 3.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[np.int64(3)]
The houses unsold are
[]
The total revenue is 10.00
```

```
# the surplus matrix S

auction_1.S
```

```
array([[0., 0., 6., 0.],
       [0., 7., 0., 0.],
       [6., 0., 0., 0.]])
```

```
# the quantity matrix X

auction_1.Q
```

```
array([[0., 0., 1., 0., 0.],
       [0., 1., 0., 0., 0.],
       [1., 0., 0., 0., 0.]])
```

150.9 Robustness checks

Let's do some stress testing of our code by applying it to auctions characterized by different matrices of private values.

1. number of houses = number of buyers

```
v2 = np.array([[8,5,9],[4,11,7],[9,7,6]])

auction_2 = ascending_bid_auction(v2, r, ϵ)

auction_2.start_auction()
```

```
The Ascending Bid Auction for Houses


Basic Information: 3 houses, 3 buyers
The valuation matrix is as follows
+--------------+---+----+---+
| Buyer Number | 0 | 1  | 2 |
+--------------+---+----+---+
|   House 0    | 8 | 5  | 9 |
|   House 1    | 4 | 11 | 7 |
|   House 2    | 9 | 7  | 6 |
+--------------+---+----+---+
The reservation prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 2.0 | 1.0 | 0.0 |
+--------------+-----+-----+-----+
The minimum increment of bid price is 1.00


Auction starts successfully


Round 1
The bid information is
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Buyer     | [2] | [1] | [0] |
```

(continues on next page)

(continued from previous page)

```
+--------------+-----+-----+-----+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 2.0 | 1.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[]


The auction ends because every buyer gets one house.


Auction ends


The final result is as follows


The allocation plan is
+--------------+---------------+---------------+---------------+
| House Number |       0       |       1       |       2       |
+--------------+---------------+---------------+---------------+
|    Buyer     | [np.int64(2)] | [np.int64(1)] | [np.int64(0)] |
+--------------+---------------+---------------+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 2.0 | 1.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[]
The houses unsold are
[]
The total revenue is 6.00
```

2. multiple excess buyers

```
v3 = np.array([[8,5,9,4,3],[4,11,7,4,6],[9,7,6,4,2]])

auction_3 = ascending_bid_auction(v3, r, ϵ)

auction_3.start_auction()
```

```
The Ascending Bid Auction for Houses


Basic Information: 3 houses, 5 buyers
The valuation matrix is as follows
+--------------+---+----+---+---+---+
```

(continues on next page)

(continued from previous page)

```
| Buyer Number | 0 | 1  | 2 | 3 | 4 |
+--------------+---+----+---+---+---+
|   House 0    | 8 | 5  | 9 | 4 | 3 |
|   House 1    | 4 | 11 | 7 | 4 | 6 |
|   House 2    | 9 | 7  | 6 | 4 | 2 |
+--------------+---+----+---+---+---+
The reservation prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 2.0 | 1.0 | 0.0 |
+--------------+-----+-----+-----+
The minimum increment of bid price is 1.00


Auction starts successfully


Round 1
The bid information is
+--------------+-----+-------+-------+
| House Number |  0  |   1   |   2   |
+--------------+-----+-------+-------+
|    Buyer     | [2] | [1 4] | [0 3] |
+--------------+-----+-------+-------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 2.0 | 1.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(4), np.int64(3)]
The losers are
[np.int64(0), np.int64(1)]


Round 2
The bid information is
+--------------+---------------+-----+-----+
| House Number |       0       |  1  |  2  |
+--------------+---------------+-----+-----+
|    Buyer     | [np.int64(2)] | [1] | [0] |
+--------------+---------------+-----+-----+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 3.0 | 2.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[np.int64(3), np.int64(4)]


Round 3
```

(continues on next page)

(continued from previous page)

```
The bid information is
+--------------+---------------+-----+-----+
| House Number |       0       |  1  |  2  |
+--------------+---------------+-----+-----+
|    Buyer     | [np.int64(2)] | [4] | [3] |
+--------------+---------------+-----+-----+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 4.0 | 3.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(4), np.int64(3)]
The losers are
[np.int64(0), np.int64(1)]


Round 4
The bid information is
+--------------+---------------+-----+-----+
| House Number |       0       |  1  |  2  |
+--------------+---------------+-----+-----+
|    Buyer     | [np.int64(2)] | [1] | [0] |
+--------------+---------------+-----+-----+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 5.0 | 4.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
The losers are
[np.int64(3), np.int64(4)]


Auction ends


The final result is as follows


The allocation plan is
+--------------+---------------+---------------+---------------+
| House Number |       0       |       1       |       2       |
+--------------+---------------+---------------+---------------+
|    Buyer     | [np.int64(2)] | [np.int64(1)] | [np.int64(0)] |
+--------------+---------------+---------------+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+
| House Number |  0  |  1  |  2  |
+--------------+-----+-----+-----+
|    Price     | 3.0 | 5.0 | 4.0 |
+--------------+-----+-----+-----+
The winners are
[np.int64(2), np.int64(1), np.int64(0)]
```

(continues on next page)

(continued from previous page)

```
The losers are
[np.int64(3), np.int64(4)]
The houses unsold are
[]
The total revenue is 12.00
```

3. more houses than buyers

```
v4 = np.array([[8,5,4],[4,11,7],[9,7,9],[6,4,5],[2,2,2]])
r2 = np.array([2,1,0,1,1])

auction_4 = ascending_bid_auction(v4, r2, ϵ)

auction_4.start_auction()
```

```
The Ascending Bid Auction for Houses


Basic Information: 5 houses, 3 buyers
The valuation matrix is as follows
+--------------+---+----+---+
| Buyer Number | 0 | 1  | 2 |
+--------------+---+----+---+
|   House 0    | 8 | 5  | 4 |
|   House 1    | 4 | 11 | 7 |
|   House 2    | 9 | 7  | 9 |
|   House 3    | 6 | 4  | 5 |
|   House 4    | 2 | 2  | 2 |
+--------------+---+----+---+
The reservation prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 1.0 | 0.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The minimum increment of bid price is 1.00


Auction starts successfully


Round 1
The bid information is
+--------------+-----+-------+
| House Number |  1  |   2   |
+--------------+-----+-------+
|    Buyer     | [1] | [0 2] |
+--------------+-----+-------+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 2.0 | 1.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2)]
```

(continues on next page)

(continued from previous page)

```
The losers are
[np.int64(0)]


Round 2
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [0] |
+--------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 2.0 | 2.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
The losers are
[np.int64(2)]


Round 3
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [2] |
+--------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 2.0 | 3.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2)]
The losers are
[np.int64(0)]


Round 4
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [0] |
+--------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 2.0 | 4.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
```

(continues on next page)

(continued from previous page)

```
The losers are
[np.int64(2)]


Round 5
The bid information is
+--------------+-----+---------------+
| House Number |  1  |       2       |
+--------------+-----+---------------+
|    Buyer     | [2] | [np.int64(0)] |
+--------------+-----+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 3.0 | 4.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(2), np.int64(0)]
The losers are
[np.int64(1)]


Round 6
The bid information is
+--------------+-----+---------------+
| House Number |  1  |       2       |
+--------------+-----+---------------+
|    Buyer     | [1] | [np.int64(0)] |
+--------------+-----+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 4.0 | 4.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
The losers are
[np.int64(2)]


Round 7
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [2] |
+--------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 2.0 | 4.0 | 5.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2)]
```

(continues on next page)

(continued from previous page)

```
The losers are
[np.int64(0)]


Round 8
The bid information is
+--------------+---------------+---------------+-----+
| House Number |       1       |       2       |  0  |
+--------------+---------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [np.int64(2)] | [0] |
+--------------+---------------+---------------+-----+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 3.0 | 4.0 | 5.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2), np.int64(0)]
The losers are
[]


The auction ends because every buyer gets one house.


Auction ends


The final result is as follows


The allocation plan is
+--------------+---------------+---------------+---------------+
| House Number |       1       |       2       |       0       |
+--------------+---------------+---------------+---------------+
|    Buyer     | [np.int64(1)] | [np.int64(2)] | [np.int64(0)] |
+--------------+---------------+---------------+---------------+
The bid prices for houses are
+--------------+-----+-----+-----+-----+-----+
| House Number |  0  |  1  |  2  |  3  |  4  |
+--------------+-----+-----+-----+-----+-----+
|    Price     | 3.0 | 4.0 | 5.0 | 1.0 | 1.0 |
+--------------+-----+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2), np.int64(0)]
The losers are
[]
The houses unsold are
[np.int64(3), np.int64(4)]
The total revenue is 12.00
```

4. some houses have extremely high reservation prices

```
v5 = np.array([[8,5,4],[4,11,7],[9,7,9],[6,4,5],[2,2,2]])
r3 = np.array([10,1,0,1,1])
```

(continues on next page)

(continued from previous page)

```
auction_5 = ascending_bid_auction(v5, r3, ϵ)

auction_5.start_auction()
```

```
The Ascending Bid Auction for Houses


Basic Information: 5 houses, 3 buyers
The valuation matrix is as follows
+--------------+---+----+---+
| Buyer Number | 0 | 1  | 2 |
+--------------+---+----+---+
|   House 0    | 8 | 5  | 4 |
|   House 1    | 4 | 11 | 7 |
|   House 2    | 9 | 7  | 9 |
|   House 3    | 6 | 4  | 5 |
|   House 4    | 2 | 2  | 2 |
+--------------+---+----+---+
The reservation prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 1.0 | 0.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The minimum increment of bid price is 1.00


Auction starts successfully


Round 1
The bid information is
+--------------+-----+-------+
| House Number |  1  |   2   |
+--------------+-----+-------+
|    Buyer     | [1] | [0 2] |
+--------------+-----+-------+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 2.0 | 1.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
The losers are
[np.int64(2)]


Round 2
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [2] |
+--------------+---------------+-----+
```

(continues on next page)

(continued from previous page)

```
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 2.0 | 2.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2)]
The losers are
[np.int64(0)]


Round 3
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [0] |
+--------------+---------------+-----+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 2.0 | 3.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
The losers are
[np.int64(2)]


Round 4
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [2] |
+--------------+---------------+-----+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 2.0 | 4.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2)]
The losers are
[np.int64(0)]


Round 5
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [0] |
+--------------+---------------+-----+
```

(continues on next page)

(continued from previous page)

```
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 2.0 | 5.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
The losers are
[np.int64(2)]


Round 6
The bid information is
+--------------+-----+---------------+
| House Number |  1  |       2       |
+--------------+-----+---------------+
|    Buyer     | [2] | [np.int64(0)] |
+--------------+-----+---------------+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 3.0 | 5.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(2), np.int64(0)]
The losers are
[np.int64(1)]


Round 7
The bid information is
+--------------+-----+---------------+
| House Number |  1  |       2       |
+--------------+-----+---------------+
|    Buyer     | [1] | [np.int64(0)] |
+--------------+-----+---------------+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 4.0 | 5.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(0)]
The losers are
[np.int64(2)]


Round 8
The bid information is
+--------------+---------------+-----+
| House Number |       1       |  2  |
+--------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [2] |
+--------------+---------------+-----+
```

(continues on next page)

(continued from previous page)

```
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 4.0 | 6.0 | 1.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2)]
The losers are
[np.int64(0)]


Round 9
The bid information is
+--------------+---------------+---------------+-----+
| House Number |       1       |       2       |  3  |
+--------------+---------------+---------------+-----+
|    Buyer     | [np.int64(1)] | [np.int64(2)] | [0] |
+--------------+---------------+---------------+-----+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 4.0 | 6.0 | 2.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2), np.int64(0)]
The losers are
[]


The auction ends because every buyer gets one house.


Auction ends


The final result is as follows


The allocation plan is
+--------------+---------------+---------------+---------------+
| House Number |       1       |       2       |       3       |
+--------------+---------------+---------------+---------------+
|    Buyer     | [np.int64(1)] | [np.int64(2)] | [np.int64(0)] |
+--------------+---------------+---------------+---------------+
The bid prices for houses are
+--------------+------+-----+-----+-----+-----+
| House Number |  0   |  1  |  2  |  3  |  4  |
+--------------+------+-----+-----+-----+-----+
|    Price     | 10.0 | 4.0 | 6.0 | 2.0 | 1.0 |
+--------------+------+-----+-----+-----+-----+
The winners are
[np.int64(1), np.int64(2), np.int64(0)]
The losers are
[]
The houses unsold are
```

(continues on next page)

(continued from previous page)

```
[np.int64(0), np.int64(4)]
The total revenue is 12.00
```

5. reservation prices are so high that the auction can't start

```
r4 = np.array([15,15,15])

auction_6 = ascending_bid_auction(v, r4, ϵ)

auction_6.start_auction()
```

```
The Ascending Bid Auction for Houses


Basic Information: 3 houses, 4 buyers
The valuation matrix is as follows
+--------------+---+----+---+---+
| Buyer Number | 0 | 1  | 2 | 3 |
+--------------+---+----+---+---+
|   House 0    | 8 | 5  | 9 | 4 |
|   House 1    | 4 | 11 | 7 | 4 |
|   House 2    | 9 | 7  | 6 | 4 |
+--------------+---+----+---+---+
The reservation prices for houses are
+--------------+------+------+------+
| House Number |  0   |  1   |  2   |
+--------------+------+------+------+
|    Price     | 15.0 | 15.0 | 15.0 |
+--------------+------+------+------+
The minimum increment of bid price is 1.00


The auction can not start because of high reservation prices
```

150.10 A Groves-Clarke mechanism

We now describe an alternative way for society to allocate n houses to m possible buyers in a way that maximizes total value across all potential buyers.

We continue to assume that each buyer can purchase at most one house.

The mechanism is a very special case of a Groves-Clarke mechanism [Groves, 1973], [Clarke, 1971].

Its special structure substantially simplifies writing Python code to find an optimal allocation.

Our mechanism works like this.

- The values V_{ij} are private information to person j
- The mechanism makes each person j willing to tell a social planner his private values $V_{i,j}$ for all $i = 1, \ldots, n$.
- The social planner asks all potential bidders to tell the planner their private values V_{ij}
- The social planner tells no one these, but uses them to allocate houses and set prices
- The mechanism is *intended* to make all prospective buyers want to tell the planner their private values
  - we shall put this intention to a test below, and find that it is not achieved

- The planner finds a house, bidder pair with highest private value by computing $(\tilde{i}, \tilde{j}) = \text{argmax}(V_{ij})$
- The planner assigns house $\tilde{i}$ to buyer $\tilde{j}$
- The planner charges buyer $\tilde{j}$ the price $\max_{-\tilde{j}} V_{\tilde{i},j}$, where $-\tilde{j}$ means all j's except $\tilde{j}$.
- The planner creates a matrix of private values for the remaining houses $-\tilde{i}$ by deleting row (i.e., house) $\tilde{i}$ and column (i.e., buyer) $\tilde{j}$ from V.
  - (But in doing this, the planner keeps track of the real names of the bidders and the houses).
- The planner returns to the original step and repeat it.
- The planner iterates until all n houses are allocated and the charges for all n houses are set.

150.11 An example solved by hand

Let's see how our Groves-Clarke algorithm would work for the following simple matrix V matrix of private values

$$V = \begin{bmatrix} 10 & 9 & 8 & 7 & 6 \\ 9 & 9 & 7 & 6 & 6 \\ 8 & 6 & 6 & 9 & 4 \\ 7 & 5 & 6 & 4 & 9 \end{bmatrix}$$

Remark: In the first step, when the highest private value corresponds to more than one house, bidder pairs, we choose the pair with the highest sale price. If a highest sale price corresponds to two or more pairs with highest private values, we randomly choose one.

```
np.random.seed(666)

V_orig = np.array([[10, 9, 8, 7, 6],  # record the original values
                   [9, 9, 7, 6, 6],
                   [8, 6, 6, 9, 4],
                   [7, 5, 6, 4, 9]])
V = np.copy(V_orig)  # used iteratively
n, m = V.shape
p = np.zeros(n) # prices of houses
Q = np.zeros((n, m)) # keep record the status of houses and buyers
```

First assignment

First, we find house, bidder pair with highest private value.

```
i, j = np.where(V==np.max(V))
i, j
```

```
(array([0]), array([0]))
```

So, house 0 will be sold to buyer 0 at a price of 9. We then update the sale price of house 0 and the status matrix Q.

```
p[i] = np.max(np.delete(V[i, :], j))
Q[i, j] = 1
p, Q
```

```
(array([9., 0., 0., 0.]),
 array([[1., 0., 0., 0., 0.],
        [0., 0., 0., 0., 0.],
```

(continues on next page)

(continued from previous page)

```
        [0., 0., 0., 0., 0.],
        [0., 0., 0., 0., 0.]]))
```

Then we remove row 0 and column 0 from V. To keep the real number of houses and buyers, we set this row and this column to -1, which will have the same result as removing them since $V \geq 0$.

```
V[i, :] = -1
V[:, j] = -1
V
```

```
array([[-1, -1, -1, -1, -1],
       [-1,  9,  7,  6,  6],
       [-1,  6,  6,  9,  4],
       [-1,  5,  6,  4,  9]])
```

Second assignment

We find house, bidder pair with the highest private value again.

```
i, j = np.where(V==np.max(V))
i, j
```

```
(array([1, 2, 3]), array([1, 3, 4]))
```

In this special example, there are three pairs (1, 1), (2, 3) and (3, 4) with the highest private value. To solve this problem, we choose the one with highest sale price.

```
p_candidate = np.zeros(len(i))
for k in range(len(i)):
    p_candidate[k] = np.max(np.delete(V[i[k], :], j[k]))
k, = np.where(p_candidate==np.max(p_candidate))
i, j = i[k], j[k]
i, j
```

```
(array([1]), array([1]))
```

So, house 1 will be sold to buyer 1 at a price of 7. We update matrices.

```
p[i] = np.max(np.delete(V[i, :], j))
Q[i, j] = 1
V[i, :] = -1
V[:, j] = -1
p, Q, V
```

```
(array([9., 7., 0., 0.]),
 array([[1., 0., 0., 0., 0.],
        [0., 1., 0., 0., 0.],
        [0., 0., 0., 0., 0.],
        [0., 0., 0., 0., 0.]]),
 array([[-1, -1, -1, -1, -1],
        [-1, -1, -1, -1, -1],
        [-1, -1,  6,  9,  4],
        [-1, -1,  6,  4,  9]]))
```

Third assignment

```
i, j = np.where(V==np.max(V))
i, j
```

```
(array([2, 3]), array([3, 4]))
```

In this special example, there are two pairs (2, 3) and (3, 4) with the highest private value.

To resolve the assignment, we choose the one with highest sale price.

```
p_candidate = np.zeros(len(i))
for k in range(len(i)):
    p_candidate[k] = np.max(np.delete(V[i[k], :], j[k]))
k, = np.where(p_candidate==np.max(p_candidate))
i, j = i[k], j[k]
i, j
```

```
(array([2, 3]), array([3, 4]))
```

The two pairs even have the same sale price.

We randomly choose one pair.

```
k = np.random.choice(len(i))
i, j = i[k], j[k]
i, j
```

```
(np.int64(2), np.int64(3))
```

Finally, house 2 will be sold to buyer 3.

We update matrices accordingly.

```
p[i] = np.max(np.delete(V[i, :], j))
Q[i, j] = 1
V[i, :] = -1
V[:, j] = -1
p, Q, V
```

```
(array([9., 7., 6., 0.]),
 array([[1., 0., 0., 0., 0.],
        [0., 1., 0., 0., 0.],
        [0., 0., 0., 1., 0.],
        [0., 0., 0., 0., 0.]]),
 array([[-1, -1, -1, -1, -1],
        [-1, -1, -1, -1, -1],
        [-1, -1, -1, -1, -1],
        [-1, -1,  6, -1,  9]]))
```

Fourth assignment

```
i, j = np.where(V==np.max(V))
i, j
```

```
(array([3]), array([4]))
```

House 3 will be sold to buyer 4.

The final outcome follows.

```
p[i] = np.max(np.delete(V[i, :], j))
Q[i, j] = 1
V[i, :] = -1
V[:, j] = -1
S = V_orig*Q - np.diag(p)@Q
p, Q, V, S
```

```
(array([9., 7., 6., 6.]),
 array([[1., 0., 0., 0., 0.],
        [0., 1., 0., 0., 0.],
        [0., 0., 0., 1., 0.],
        [0., 0., 0., 0., 1.]]),
 array([[-1, -1, -1, -1, -1],
        [-1, -1, -1, -1, -1],
        [-1, -1, -1, -1, -1],
        [-1, -1, -1, -1, -1]]),
 array([[1., 0., 0., 0., 0.],
        [0., 2., 0., 0., 0.],
        [0., 0., 0., 3., 0.],
        [0., 0., 0., 0., 3.]]))
```

150.12 Another Python class

It is efficient to assemble our calculations in a single Python class.

```
class GC_Mechanism:

    def __init__(self, V):
        """
        Implementation of the special Groves Clarke Mechanism for house auction.

        Parameters:
        ----------
        V: 2 dimensional private value matrix

        """

        self.V_orig = V.copy()
        self.V = V.copy()
        self.n, self.m = self.V.shape
        self.p = np.zeros(self.n)
        self.Q = np.zeros((self.n, self.m))
        self.S = np.copy(self.Q)

    def find_argmax(self):
        """
        Find the house-buyer pair with the highest value.
        When the highest private value corresponds to more than one house, bidder␣
↪pairs,
        we choose the pair with the highest sale price.
        Moreover, if the highest sale price corresponds to two or more pairs with␣
↪highest private value,
        We randomly choose one.
```

(continues on next page)

(continued from previous page)

```
        Parameters:
        ----------
        V: 2 dimensional private value matrix with -1 indicating removed rows and␣
↪columns

        Returns:
        ----------
        i: the index of the sold house

        j: the index of the buyer

        """
        i, j = np.where(self.V==np.max(self.V))

        if (len(i)>1):
            p_candidate = np.zeros(len(i))
            for k in range(len(i)):
                p_candidate[k] = np.max(np.delete(self.V[i[k], :], j[k]))
            k, = np.where(p_candidate==np.max(p_candidate))
            i, j = i[k], j[k]

            if (len(i)>1):
                k = np.random.choice(len(i))
                k = np.array([k])
                i, j = i[k], j[k]
        return i, j

    def update_status(self, i, j):
        # the price is the best forgone value of a rival buyer for the same house;
        # entries equal to -1 mark buyers and houses already dealt with, and if no
        # rival remains there is no competing bid, so the price is zero
        rivals = np.delete(self.V[i, :], j)
        rivals = rivals[rivals >= 0]
        self.p[i] = rivals.max() if rivals.size > 0 else 0.0
        self.Q[i, j] = 1
        self.V[i, :] = -1
        self.V[:, j] = -1

    def calculate_surplus(self):
        self.S = self.V_orig*self.Q - np.diag(self.p)@self.Q

    def start(self, verbose=True):
        log = print if verbose else (lambda *args, **kwargs: None)
        while (np.max(self.V)>=0):
            i, j = self.find_argmax()
            self.update_status(i, j)
            log("House %i is sold to buyer %i at price %i"%(i[0], j[0], self.p[i[0]]))
            log("\n")
        self.calculate_surplus()
        log("Prices of house:\n", self.p)
        log("\n")
        log("The status matrix:\n", self.Q)
        log("\n")
        log("The surplus matrix:\n", self.S)
```

```
np.random.seed(666)

V_orig = np.array([[10, 9, 8, 7, 6],
                   [9, 9, 7, 6, 6],
                   [8, 6, 6, 9, 4],
                   [7, 5, 6, 4, 9]])
gc_mechanism = GC_Mechanism(V_orig)
gc_mechanism.start()
```

```
House 0 is sold to buyer 0 at price 9


House 1 is sold to buyer 1 at price 7


House 2 is sold to buyer 3 at price 6


House 3 is sold to buyer 4 at price 6


Prices of house:
 [9. 7. 6. 6.]


The status matrix:
 [[1. 0. 0. 0. 0.]
 [0. 1. 0. 0. 0.]
 [0. 0. 0. 1. 0.]
 [0. 0. 0. 0. 1.]]


The surplus matrix:
 [[1. 0. 0. 0. 0.]
 [0. 2. 0. 0. 0.]
 [0. 0. 0. 3. 0.]
 [0. 0. 0. 0. 3.]]
```

150.12.1 Elaborations

Here we use some additional notation designed to conform with standard notation in parts of the VCG literature.

We want to verify that our pseudo code is indeed a **pivot mechanism**, also called a **VCG** (Vickrey-Clarke-Groves) mechanism.

- The mechanism is named after Groves [1973], Clarke [1971], and Vickrey [1961].

To prepare for verifying this, we add some notation.

Let X be the set of feasible allocations of houses under the protocols above (i.e., at most one house to each person).

Let $X(v)$ be the allocation that the mechanism chooses for matrix v of private values.

The mechanism maps a matrix v of private values into an $x \in X$.

Let $v_j(x)$ be the value that person j attaches to allocation $x \in X$.

Let $\check{t}_j(v)$ the payment that the mechanism charges person j.

The VCG mechanism chooses the allocation

$$X(v) = \text{argmax}_{x \in X} \sum_{j=1}^{m} v_j(x) \tag{150.1}$$

and charges person j a "social cost"

$$\check{t}_j(v) = \max_{x \in X} \sum_{k \neq j} v_k(x) - \sum_{k \neq j} v_k(X(v)) \tag{150.2}$$

In our setting, equation (150.1) says that the VCG allocation allocates houses to maximize the total value of the successful prospective buyers.

In our setting, equation (150.2) says that the mechanism charges people for the externality that their presence in society imposes on other prospective buyers.

Thus, notice that according to equation (150.2):

- unsuccessful prospective buyers pay 0 because removing them from "society" would not affect the allocation chosen by the mechanism
- successful prospective buyers pay the difference between the total value society could achieve without them present and the total value that others present in society do achieve under the mechanism.

It is natural to ask whether the greedy mechanism that we coded above satisfies (150.1) and (150.2).

We shall compute $\check{t}_j$ for $j = 1, \ldots, m$ and compare it with the price p_j that the greedy mechanism actually charges.

The answer, which we develop in the next two sections, is that it satisfies neither requirement in general.

150.12.2 Social cost

Using the GC_Mechanism class, we can calculate the social cost of each buyer.

Let's see a simpler example with private value matrix

$$V = \begin{bmatrix} 10 & 9 & 8 & 7 & 6 \\ 9 & 8 & 7 & 6 & 6 \\ 8 & 7 & 6 & 5 & 4 \end{bmatrix}$$

To begin with, we implement the GC mechanism and see the outcome.

```
np.random.seed(666)

V_orig = np.array([[10, 9, 8, 7, 6],
                   [9, 8, 7, 6, 6],
                   [8, 7, 6, 5, 4]])
gc_mechanism = GC_Mechanism(V_orig)
gc_mechanism.start()
```

```
House 0 is sold to buyer 0 at price 9


House 1 is sold to buyer 1 at price 7


House 2 is sold to buyer 2 at price 5
```

(continues on next page)

(continued from previous page)

```
Prices of house:
 [9. 7. 5.]


The status matrix:
 [[1. 0. 0. 0. 0.]
 [0. 1. 0. 0. 0.]
 [0. 0. 1. 0. 0.]]


The surplus matrix:
 [[1. 0. 0. 0. 0.]
 [0. 1. 0. 0. 0.]
 [0. 0. 1. 0. 0.]]
```

We exclude buyer 0 and calculate the allocation.

```
V_exc_0 = np.copy(V_orig)
V_exc_0[:, 0] = -1
V_exc_0
gc_mechanism_exc_0 = GC_Mechanism(V_exc_0)
gc_mechanism_exc_0.start()
```

```
House 0 is sold to buyer 1 at price 8


House 1 is sold to buyer 2 at price 6


House 2 is sold to buyer 3 at price 4


Prices of house:
 [8. 6. 4.]


The status matrix:
 [[0. 1. 0. 0. 0.]
 [0. 0. 1. 0. 0.]
 [0. 0. 0. 1. 0.]]


The surplus matrix:
 [[-0.  1.  0.  0.  0.]
 [-0.  0.  1.  0.  0.]
 [-0.  0.  0.  1.  0.]]
```

Calculate the social cost of buyer 0.

```
print("The social cost of buyer 0:",
     np.sum(gc_mechanism_exc_0.Q*gc_mechanism_exc_0.V_orig)-np.sum(np.delete(gc_
 ↪mechanism.Q*gc_mechanism.V_orig, 0, axis=1)))
```

```
The social cost of buyer 0: 7.0
```

Repeat this process for buyer 1 and buyer 2

```
V_exc_1 = np.copy(V_orig)
V_exc_1[:, 1] = -1
V_exc_1
gc_mechanism_exc_1 = GC_Mechanism(V_exc_1)
gc_mechanism_exc_1.start()

print("\nThe social cost of buyer 1:",
     np.sum(gc_mechanism_exc_1.Q*gc_mechanism_exc_1.V_orig)-np.sum(np.delete(gc_
 ↪mechanism.Q*gc_mechanism.V_orig, 1, axis=1)))
```

```
House 0 is sold to buyer 0 at price 8


House 1 is sold to buyer 2 at price 6


House 2 is sold to buyer 3 at price 4


Prices of house:
 [8. 6. 4.]


The status matrix:
 [[1. 0. 0. 0. 0.]
 [0. 0. 1. 0. 0.]
 [0. 0. 0. 1. 0.]]


The surplus matrix:
 [[ 2. -0.  0.  0.  0.]
 [ 0. -0.  1.  0.  0.]
 [ 0. -0.  0.  1.  0.]]

The social cost of buyer 1: 6.0
```

```
V_exc_2 = np.copy(V_orig)
V_exc_2[:, 2] = -1
V_exc_2
gc_mechanism_exc_2 = GC_Mechanism(V_exc_2)
gc_mechanism_exc_2.start()

print("\nThe social cost of buyer 2:",
     np.sum(gc_mechanism_exc_2.Q*gc_mechanism_exc_2.V_orig)-np.sum(np.delete(gc_
 ↪mechanism.Q*gc_mechanism.V_orig, 2, axis=1)))
```

```
House 0 is sold to buyer 0 at price 9


House 1 is sold to buyer 1 at price 6


House 2 is sold to buyer 3 at price 4


Prices of house:
```

(continues on next page)

(continued from previous page)

```
 [9. 6. 4.]


The status matrix:
 [[1. 0. 0. 0. 0.]
 [0. 1. 0. 0. 0.]
 [0. 0. 0. 1. 0.]]


The surplus matrix:
 [[ 1.  0. -0.  0.  0.]
 [ 0.  2. -0.  0.  0.]
 [ 0.  0. -0.  1.  0.]]

The social cost of buyer 2: 5.0
```

150.13 Is the greedy mechanism a VCG mechanism?

The mechanism that our benevolent planner runs is **greedy**: it repeatedly grabs the largest remaining entry of the value matrix.

A genuine VCG mechanism must do two things, namely, satisfy (150.1) and (150.2).

Let's check them one at a time.

We'll need a way to compute the value-maximizing allocation, which is a classic **assignment problem** that can be solved by linear programming.

```
from scipy.optimize import linear_sum_assignment

def optimal_assignment(V):
    """
    Solve the assignment problem: allocate at most one house to each buyer
    so as to maximize the total value.
    Returns the total value and a dictionary mapping buyer -> house.
    """
    rows, cols = linear_sum_assignment(-V)
    return V[rows, cols].sum(), {int(c): int(r) for r, c in zip(rows, cols)}

def greedy_outcome(V):
    "Run the planner's greedy mechanism quietly and report what it achieves."
    gc = GC_Mechanism(V.copy())
    gc.start(verbose=False)
    return (gc.Q * gc.V_orig).sum(), gc
```

150.13.1 Does the greedy allocation maximize total value?

Consider this simple matrix of private values for two houses and two buyers.

```
V_trap = np.array([[10, 9],
                   [9, 1]])

greedy_value, _ = greedy_outcome(V_trap)
best_value, best_alloc = optimal_assignment(V_trap)

print(f"greedy total value:  {greedy_value:.0f}")
print(f"optimal total value: {best_value:.0f}   (buyer -> house: {best_alloc})")
```

```
greedy total value:  11
optimal total value: 18   (buyer -> house: {1: 0, 0: 1})
```

The greedy rule starts from the largest entry, $V_{00} = 10$, and gives house 0 to buyer 0.

That leaves only $V_{11} = 1$, for a total of 11.

But assigning house 0 to buyer 1 and house 1 to buyer 0 yields $9 + 9 = 18$.

The largest entry was a trap: grabbing it destroyed a better pair of assignments.

So the greedy mechanism does *not* satisfy (150.1).

Exercise 150.14.1 asks you to investigate how often this happens.

The value matrices used earlier in this lecture happen to be ones for which greedy does find the best allocation, which is why the difficulty did not show up before.

150.13.2 Are the prices VCG payments?

Now let's compare the prices that the greedy mechanism charges with the social costs (150.2).

```
def vcg_payments(V):
    """
    The social cost of each buyer:
    the total value others could achieve without this buyer,
    minus the total value others do achieve when this buyer is present.
    """
    total, alloc = optimal_assignment(V)
    payments = np.zeros(V.shape[1])

    for j in range(V.shape[1]):
        total_without_j, _ = optimal_assignment(np.delete(V, j, axis=1))
        value_of_others = total - (V[alloc[j], j] if j in alloc else 0.0)
        payments[j] = total_without_j - value_of_others

    return payments
```

```
t_vcg = vcg_payments(V_orig)

print(f"{'buyer':>6}{'house':>7}{'greedy price':>15}{'VCG payment':>14}")
for j in range(V_orig.shape[1]):
    house = [i for i in range(V_orig.shape[0]) if gc_mechanism.Q[i, j] == 1]
    if house:
```

(continues on next page)

(continued from previous page)

```
        print(f"{j:>6}{house[0]:>7}{gc_mechanism.p[house[0]]:>15.1f}{t_vcg[j]:>14.1f}
↪")
```

```
 buyer  house   greedy price   VCG payment
     0      0            9.0           7.0
     1      1            7.0           6.0
     2      2            5.0           5.0
```

The prices differ from the VCG payments.

Notice that the VCG payments are exactly the social costs that we computed in the previous section by excluding one buyer at a time.

So the planner's mechanism collects *more* than the externality that a winning buyer imposes on others.

A mechanism that assigns objects greedily and charges each winner the best forgone offer for the same object is known as a **generalized second-price** mechanism.

Generalized second-price mechanisms are widely used – they are how search engines sell advertising slots – and it is well known that they are *not* VCG mechanisms [Edelman *et al.*, 2007].

150.13.3 Is truthful reporting a dominant strategy?

Because the mechanism is not VCG, the argument that makes truth telling a dominant strategy does not apply.

In fact truth telling is *not* a dominant strategy here, and Exercise 150.14.2 asks you to demonstrate that by finding a profitable lie.

150.13.4 What about the ascending bid auction?

Here is the striking fact promised earlier.

Our setting is an assignment market in which each buyer wants at most one house.

For such markets, Leonard [1983] and Demange *et al.* [1986] showed that the **minimum competitive equilibrium prices** coincide with the payments prescribed by the VCG mechanism.

Since an ascending auction with straightforward bidding converges to those minimum competitive equilibrium prices, it follows that our ascending bid auction implements the VCG outcome as the bid increment ϵ becomes small.

It is the auctioneer, not the benevolent planner, who runs a VCG mechanism here.

Exercise 150.14.3 asks you to verify this numerically.

150.14 Exercises

Exercise 150.14.1

We saw that the greedy mechanism can fail to maximize total value.

How often does it fail?

Draw many random value matrices with integer entries, run both the greedy mechanism and `optimal_assignment`, and report

1. the fraction of matrices for which greedy is suboptimal
2. the largest percentage shortfall that you find

Then explain in words why a greedy rule goes wrong.

Solution

```
rng_ha = np.random.default_rng(0)
trials = 200
failures, worst = 0, 0.0

for _ in range(trials):
    n_h, m_b = rng_ha.integers(2, 5), rng_ha.integers(2, 6)
    if m_b < n_h:
        n_h, m_b = m_b, n_h
    V_rand = rng_ha.integers(0, 12, size=(n_h, m_b))

    greedy_value, _ = greedy_outcome(V_rand)
    best_value, _ = optimal_assignment(V_rand)

    if greedy_value < best_value - 1e-9:
        failures += 1
        worst = max(worst, (best_value - greedy_value)/best_value)

print(f"greedy was suboptimal in {failures} of {trials} matrices")
print(f"largest shortfall: {100*worst:.1f} per cent of the attainable total value")
```

```
greedy was suboptimal in 66 of 200 matrices
largest shortfall: 43.8 per cent of the attainable total value
```

Greedy fails often, and sometimes badly.

The reason is that the largest entry of the value matrix need not belong to the best assignment.

Taking it removes a house *and* a buyer from further consideration, and that can destroy two good matches in order to secure one very good one.

Choosing the best assignment requires weighing all pairings simultaneously, which is what the linear programming routine does.

Exercise 150.14.2

Show that truthful reporting is not a dominant strategy under the planner's greedy mechanism.

Use the value matrix `V_orig` of the previous section, and consider buyer 0, whose true values for the three houses are $(10, 9, 8)$.

1. What surplus does buyer 0 receive when everyone reports truthfully?
2. Search over reports (a, b, c) with integer entries between 0 and 10 that buyer 0 might submit instead, holding fixed the reports of all other buyers, and find the report that maximizes buyer 0's *true* surplus.
3. Explain what the manipulation accomplishes.

Solution

```
def surplus_of_buyer_0(V_true, V_reported):
    "True surplus of buyer 0 when the mechanism is run on the reported values."
    gc = GC_Mechanism(V_reported.copy())
    gc.start(verbose=False)
    for i in range(V_true.shape[0]):
        if gc.Q[i, 0] == 1:
            return V_true[i, 0] - gc.p[i], i, gc.p[i]
    return 0.0, None, None      # buyer 0 wins nothing

np.random.seed(0)
truth_surplus, house, price = surplus_of_buyer_0(V_orig, V_orig)
print(f"truthful report: buyer 0 wins house {house} at price {price:.0f}, "
      f"surplus {truth_surplus:.0f}")

best = (truth_surplus, None)
for a in range(11):
    for b in range(11):
        for c in range(11):
            V_lie = V_orig.copy()
            V_lie[:, 0] = [a, b, c]
            surplus, h, pr = surplus_of_buyer_0(V_orig, V_lie)
            if surplus > best[0] + 1e-9:
                best = (surplus, (a, b, c, h, pr))

report = best[1]
print(f"best misreport found: {report[:3]} wins house {report[3]} "
      f"at price {report[4]:.0f}, surplus {best[0]:.0f}")
```

```
truthful report: buyer 0 wins house 0 at price 9, surplus 1
best misreport found: (0, 0, 5) wins house 2 at price 5, surplus 3
```

By reporting that he places no value on the two expensive houses, buyer 0 steps out of the competition for them.

The mechanism then awards him a cheaper house, and charges him the best forgone offer for *that* house, which is much less than what he paid when he told the truth.

His true value for the house he ends up with is high enough that this is a profitable deviation.

Under a genuine VCG mechanism this could not happen: a buyer's payment never depends on his own report, only on the harm he does to others.

Exercise 150.14.3

Verify that the ascending bid auction implements the VCG outcome.

Use the value matrix

$$v = \begin{bmatrix} 8 & 5 & 9 & 4 \\ 4 & 11 & 7 & 4 \\ 9 & 7 & 6 & 4 \end{bmatrix}$$

with reservation prices of zero, so that prices are free to settle at their competitive levels.

1. Compute the VCG payments and the value-maximizing allocation.
2. Run the ascending bid auction for $\epsilon = 1, 0.5, 0.25, 0.1, 0.05$ and record what each winning buyer pays.

3. Does the auction maximize total value? Do its payments approach the VCG payments?

Solution

```
v_ex = np.array([[8, 5, 9, 4],
                 [4, 11, 7, 4],
                 [9, 7, 6, 4]])

t_vcg_ex = vcg_payments(v_ex)
best_value_ex, best_alloc_ex = optimal_assignment(v_ex)

print(f"VCG payments by buyer: {t_vcg_ex}")
print(f"value-maximizing allocation (buyer -> house): {best_alloc_ex}")
print(f"maximum total value: {best_value_ex:.0f}\n")

print(f"{'ε':>7}{'total value':>14}   payments by buyer")
for ϵ_try in (1.0, 0.5, 0.25, 0.1, 0.05):
    np.random.seed(0)
    auction = ascending_bid_auction(v_ex, np.zeros(3), ϵ_try)
    auction.start_auction(verbose=False)

    payments = {j: round(float(auction.p[i]), 3)
                for i in range(v_ex.shape[0])
                for j in range(v_ex.shape[1]) if auction.Q[i, j] == 1}
    total_value = (auction.Q[:, :v_ex.shape[1]] * v_ex).sum()
    print(f"{ϵ_try:>7}{total_value:>14.0f}   {payments}")
```

```
VCG payments by buyer: [4. 4. 4. 0.]
value-maximizing allocation (buyer -> house): {2: 0, 1: 1, 0: 2}
maximum total value: 29

      ε   total value   payments by buyer
    1.0            29   {2: 4.0, 1: 3.0, 0: 3.0}
    0.5            29   {2: 4.0, 1: 3.5, 0: 3.5}
   0.25            29   {2: 4.0, 1: 3.75, 0: 3.75}
    0.1            29   {2: 4.0, 1: 3.9, 0: 3.9}
   0.05            29   {2: 4.0, 1: 4.05, 0: 4.05}
```

The auction always attains the maximum total value.

As ϵ shrinks, each winner's payment approaches that buyer's VCG payment.

This illustrates the theorem of Leonard [1983] and Demange *et al.* [1986]: in an assignment market, the minimum competitive equilibrium prices are exactly the VCG payments.

The auctioneer never asks anyone to reveal a private value, yet the ascending auction arrives at the outcome that a correctly designed planner would compute from full information.

Part XIX

Other

CHAPTER

ONE

TROUBLESHOOTING

Contents

- *Troubleshooting*
  - *Fixing Your Local Environment*
  - *Reporting an Issue*

This page is for readers experiencing errors when running the code from the lectures.

151.1 Fixing Your Local Environment

The basic assumption of the lectures is that code in a lecture should execute whenever

1. it is executed in a Jupyter notebook and
2. the notebook is running on a machine with the latest version of Anaconda Python.

You have installed Anaconda, haven't you, following the instructions in this lecture?

Assuming that you have, the most common source of problems for our readers is that their Anaconda distribution is not up to date.

Here's a useful article on how to update Anaconda.

Another option is to simply remove Anaconda and reinstall.

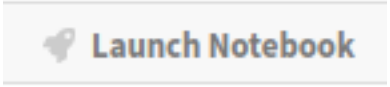
You also need to keep the external code libraries, such as QuantEcon.py up to date.

For this task you can either

- use `pip install --upgrade quantecon` on the command line, or
- execute `!pip install --upgrade quantecon` within a Jupyter notebook.

If your local environment is still not working you can do two things.

First, you can use a remote machine instead, by clicking on the Launch Notebook icon available for each lecture



Second, you can report an issue, so we can try to fix your local set up.

We like getting feedback on the lectures so please don't hesitate to get in touch.

151.2 Reporting an Issue

One way to give feedback is to raise an issue through our issue tracker.

Please be as specific as possible. Tell us where the problem is and as much detail about your local set up as you can provide.

Finally, you can provide direct feedback to contact@quantecon.org

CHAPTER

TWO

REFERENCES

CHAPTER

THREE

EXECUTION STATISTICS

This table contains the latest execution statistics.

Document	Modified	Method	Run Time (s)	Status
affine_risk_prices	2026-11-02 03:19	cache	14.38	✓
aiyagari	2026-11-02 03:19	cache	16.68	✓
aiyagari_egm	2026-11-02 03:19	cache	16.24	✓
ak2	2026-11-02 03:19	cache	12.4	✓
ak_aiyagari	2026-11-02 03:20	cache	33.61	✓
ar1_bayes	2026-11-02 03:24	cache	223.07	✓
ar1_turningpts	2026-11-02 03:27	cache	197.93	✓
back_prop	2026-11-02 03:28	cache	91.89	✓
bayes_nonconj	2026-11-02 03:29	cache	51.35	✓
blackwell_kihlstrom	2026-11-02 03:29	cache	7.17	✓
bounded_rationality	2026-11-08 13:12	cache	10.4	✓
career	2026-11-02 03:30	cache	9.21	✓
cass_fiscal	2026-11-02 03:30	cache	18.06	✓
cass_fiscal_2	2026-11-02 03:31	cache	5.62	✓
cass_koopmans_1	2026-11-02 03:31	cache	13.76	✓
cass_koopmans_2	2026-11-02 03:31	cache	6.7	✓
chow_business_cycles	2026-11-02 03:31	cache	9.44	✓
cross_product_trick	2026-11-02 03:31	cache	0.93	✓
divergence_measures	2026-11-02 03:31	cache	7.12	✓
eig_circulant	2026-11-02 03:31	cache	3.05	✓
endogenous_lake	2026-11-02 03:40	cache	504.94	✓
exchange_rate_learning	2026-11-08 13:13	cache	92.48	✓
exchangeable	2026-11-02 03:41	cache	10.74	✓
finite_markov	2026-11-02 03:41	cache	7.96	✓
gallant_tauchen	2026-11-02 03:42	cache	34.23	✓
ge_arrow	2026-11-02 03:42	cache	3.01	✓
genetic_classifier	2026-11-08 13:15	cache	79.52	✓
hansen_singleton_1982	2026-11-02 03:43	cache	13.51	✓
hansen_singleton_1983	2026-11-02 03:43	cache	25.63	✓
harrison_kreps	2026-11-02 03:43	cache	3.59	✓
hoist_failure	2026-11-02 03:45	cache	93.04	✓
house_auction	2026-11-02 03:45	cache	3.41	✓
ifp_advanced	2026-11-02 03:46	cache	25.0	✓
ifp_discrete	2026-11-02 03:46	cache	12.6	✓
ifp_egm	2026-11-02 03:46	cache	9.78	✓
ifp_egm_transient_shocks	2026-11-02 03:46	cache	16.86	✓

continues on next page

Table 153.1 – continued from previous page

Document	Modified	Method	Run Time (s)	Status
ifp_opi	2026-11-02 03:46	cache	8.32	✓
imp_sample	2026-11-02 03:48	cache	121.13	✓
information_market_equilibrium	2026-11-02 03:49	cache	14.75	✓
intro	2026-11-02 03:49	cache	1.13	✓
inventory_dynamics	2026-11-02 03:49	cache	15.85	✓
inventory_q	2026-11-02 03:49	cache	7.46	✓
jv	2026-11-02 03:49	cache	11.88	✓
kalman	2026-11-02 03:49	cache	7.87	✓
kalman_2	2026-11-02 03:50	cache	34.65	✓
kalman_filter_var	2026-11-02 03:50	cache	7.85	✓
kesten_processes	2026-11-02 03:50	cache	10.93	✓
lagrangian_lqdp	2026-11-02 03:51	cache	25.75	✓
lake_model	2026-11-02 03:51	cache	10.12	✓
learning_approximation	2026-11-08 13:16	cache	98.57	✓
likelihood_bayes	2026-11-02 03:53	cache	53.67	✓
likelihood_ratio_process	2026-11-02 03:53	cache	17.9	✓
likelihood_ratio_process_2	2026-11-02 03:53	cache	23.86	✓
likelihood_var	2026-11-02 03:54	cache	32.46	✓
linear_algebra	2026-11-02 03:54	cache	2.42	✓
linear_models	2026-11-02 03:54	cache	7.79	✓
lln_clt	2026-11-02 03:54	cache	10.47	✓
long_run_risk_operator	2026-11-02 03:55	cache	18.51	✓
lq_bewley_complete_markets	2026-11-02 03:55	cache	4.66	✓
lq_inventories	2026-11-02 03:55	cache	13.01	✓
lq_mean_field_games	2026-11-02 03:55	cache	5.6	✓
lq_mean_field_games_extensions	2026-11-02 03:55	cache	2.77	✓
lq_permanent_income	2026-11-02 03:55	cache	3.57	✓
lq_robust_bewley	2026-11-02 03:55	cache	11.0	✓
lq_robust_smoothing	2026-11-02 03:56	cache	10.92	✓
lqcontrol	2026-11-02 03:56	cache	5.05	✓
ls_learning	2026-11-02 03:56	cache	9.76	✓
lucas_prescott_investment	2026-11-02 03:56	cache	14.04	✓
lucas_prescott_search	2026-11-08 13:18	cache	86.34	✓
marimon_mcgrattan_sargent	2026-11-08 13:20	cache	136.07	✓
market_diffusion	2026-11-02 04:04	cache	9.25	✓
markov_asset	2026-11-02 04:04	cache	5.54	✓
markov_perf	2026-11-02 04:04	cache	4.39	✓
mccall_fitted_vfi	2026-11-02 04:05	cache	45.14	✓
mccall_model	2026-11-02 04:05	cache	12.95	✓
mccall_model_with_sep_markov	2026-11-02 04:05	cache	26.28	✓
mccall_model_with_separation	2026-11-02 04:05	cache	15.07	✓
mccall_persist_trans	2026-11-02 04:06	cache	10.63	✓
mccall_q	2026-11-02 04:06	cache	36.56	✓
mccall_risk	2026-11-02 04:06	cache	13.89	✓
measurement_models	2026-11-02 04:07	cache	22.62	✓
merging_of_opinions	2026-11-02 04:07	cache	7.31	✓
misspecified_recovery	2026-11-02 04:07	cache	28.26	✓
mix_model	2026-11-02 04:10	cache	147.44	✓
mle	2026-11-02 04:10	cache	13.76	✓
morris_learn	2026-11-02 04:10	cache	15.92	✓

continues on next page

Table 153.1 – continued from previous page

Document	Modified	Method	Run Time (s)	Status
multi_hyper	2026-11-02 04:11	cache	30.16	✓
multivariate_normal	2026-11-02 04:11	cache	5.66	✓
navy_captain	2026-11-02 04:12	cache	46.91	✓
newton_method	2026-11-02 04:13	cache	52.74	✓
odu	2026-11-02 04:14	cache	48.39	✓
olg_adaptive_money	2026-11-08 13:21	cache	20.39	✓
ols	2026-11-02 04:14	cache	7.89	✓
opt_transport	2026-11-02 04:15	cache	14.77	✓
optimal_growth_uncertainty	2026-11-02 04:15	cache	4.12	✓
organization_capital	2026-11-02 04:15	cache	5.72	✓
os	2026-11-02 04:15	cache	1.88	✓
os_egm	2026-11-02 04:15	cache	3.76	✓
os_egm_jax	2026-11-02 04:15	cache	6.63	✓
os_numerical	2026-11-02 04:15	cache	14.2	✓
os_stochastic	2026-11-02 04:16	cache	58.23	✓
os_time_iter	2026-11-02 04:16	cache	5.06	✓
pandas_panel	2026-11-02 04:17	cache	6.39	✓
perm_income	2026-11-02 04:17	cache	3.82	✓
perm_income_cons	2026-11-02 04:17	cache	7.14	✓
phillips_adaptive	2026-11-08 13:21	cache	5.49	✓
phillips_credibility	2026-11-08 13:21	cache	7.39	✓
phillips_credible_policies	2026-11-08 13:21	cache	4.32	✓
phillips_drifts_volatilities	2026-11-08 13:37	cache	984.59	✓
phillips_drifts_volatilities_2	2026-11-08 13:51	cache	800.43	✓
phillips_escaping_nash	2026-11-08 13:51	cache	29.6	✓
phillips_learning	2026-11-08 13:55	cache	205.17	✓
phillips_lost_conquest	2026-11-08 13:55	cache	34.64	✓
phillips_misspecified	2026-11-08 13:55	cache	5.58	✓
phillips_policy_evaluation	2026-11-08 13:56	cache	25.36	✓
phillips_priors	2026-11-08 13:56	cache	33.19	✓
phillips_self_confirming	2026-11-08 13:57	cache	23.43	✓
phillips_two_stories	2026-11-08 13:57	cache	5.36	✓
pricing_information	2026-11-02 04:43	cache	19.44	✓
prob_matrix	2026-11-02 04:43	cache	10.81	✓
prob_meaning	2026-11-02 04:44	cache	39.07	✓
prospects_bounded_rationality	2026-11-08 13:57	cache	31.81	✓
qr_decomp	2026-11-02 04:44	cache	1.45	✓
rand_resp	2026-11-02 04:44	cache	2.28	✓
rational_expectations	2026-11-02 04:44	cache	3.51	✓
rational_learning_re	2026-11-02 04:44	cache	5.41	✓
re_with_feedback	2026-11-02 04:44	cache	10.26	✓
robust_permanent_income	2026-11-02 04:44	cache	6.33	✓
ross_recovery	2026-11-02 04:44	cache	6.43	✓
rs_inventory_q	2026-11-02 04:45	cache	8.53	✓
samuelson	2026-11-02 04:45	cache	10.51	✓
sargent_surico	2026-11-02 04:59	cache	879.17	✓
sir_model	2026-11-02 04:59	cache	3.2	✓
social_security_depletion	2026-11-02 05:10	cache	640.8	✓
stats_examples	2026-11-02 05:10	cache	4.08	✓
status	2026-11-02 05:10	cache	6.84	✓

continues on next page

Table 153.1 – continued from previous page

Document	Modified	Method	Run Time (s)	Status
survival_recursive_preferences	2026-11-02 05:10	cache	6.32	✓
svd_intro	2026-11-02 05:10	cache	2.62	✓
theil_1	2026-11-02 05:11	cache	4.42	✓
theil_2	2026-11-02 05:11	cache	5.7	✓
troubleshooting	2026-11-02 03:49	cache	1.13	✓
two_auctions	2026-11-02 05:11	cache	6.78	✓
two_computation	2026-11-02 05:12	cache	45.6	✓
uncertainty_traps	2026-11-02 05:12	cache	2.61	✓
unemployment_linear	2026-11-02 05:12	cache	29.11	✓
unemployment_shocks	2026-11-02 05:13	cache	53.78	✓
util_rand_resp	2026-11-02 05:13	cache	2.4	✓
var_dmd	2026-11-02 05:13	cache	2.62	✓
var_subsets	2026-11-02 05:13	cache	7.87	✓
von_neumann_model	2026-11-02 05:13	cache	3.02	✓
wald_friedman	2026-11-02 05:14	cache	18.75	✓
wald_friedman_2	2026-11-02 05:14	cache	35.87	✓
wealth_dynamics	2026-11-02 05:15	cache	23.52	✓
zreferences	2026-11-02 03:49	cache	1.13	✓

These lectures are built on `linux` instances through `github actions`.

These lectures are using the following python version

```
!python --version
```

```
Python 3.13.14
```

and the following package versions

```
!conda list
```

This lecture series has access to the following GPU

```
!nvidia-smi
```

```
Mon Nov  2 05:10:45 2026
+-----------------------------------------------------------------------------
 ↪-------+
| NVIDIA-SMI 580.105.08             Driver Version: 580.105.08     CUDA Version:␣
 ↪13.0     |
+-----------------------------------------+------------------------+----------
 ↪-------+
| GPU  Name                 Persistence-M | Bus-Id          Disp.A | Volatile␣
 ↪Uncorr. ECC |
| Fan  Temp   Perf          Pwr:Usage/Cap |           Memory-Usage | GPU-Util ␣
 ↪Compute M. |
|                                         |                        |          ␣
 ↪MIG M. |
|=========================================+========================+==========
|   0  Tesla T4                       On  |   00000000:00:1E.0 Off |          ␣
 ↪     0 |
| N/A   24C    P8             13W /   70W |       0MiB /  15360MiB |      0%  ␣
```

(continues on next page)

(continued from previous page)

```
 ↪Default |
|                                         |                         |              ␣
 ↪   N/A |
+-----------------------------------------+-------------------------+--------------
 ↪-------+

+-----------------------------------------------------------------------------------
 ↪-------+
| Processes:                                                                       ␣
 ↪       |
|  GPU   GI   CI              PID   Type   Process name                          GPU␣
 ↪Memory |
|        ID   ID                                                                  ␣
 ↪Usage      |
|=================================================================================|
|  No running processes found                                                      ␣
 ↪       |
+-----------------------------------------------------------------------------------
 ↪-------+
```

You can check the backend used by JAX using:

```
import jax
# Check if JAX is using GPU
print(f"JAX backend: {jax.devices()[0].platform}")
```

```
JAX backend: gpu
```

BIBLIOGRAPHY

[Imrohorouglu93] Selahattin İmrohoroğlu. Testing for sunspot equilibria in the German hyperinflation. *Journal of Economic Dynamics and Control*, 17(1-2):289–317, 1993. doi:10.1016/S0165-1889(06)80013-0.

[Abr86] Dilip Abreu. Extremal equilibria of oligopolistic supergames. *Journal of Economic Theory*, 39(1):191–225, 1986.

[Abr88] Dilip Abreu. On the theory of infinitely repeated games with discounting. *Econometrica*, 56:383–396, 1988.

[APS90] Dilip Abreu, David Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, September 1990.

[ACY16] Daron Acemoglu, Victor Chernozhukov, and Muhamet Yildiz. Fragility of asymptotic agreement under Bayesian learning. *Theoretical Economics*, 11(1):187–225, 2016. doi:10.3982/TE436.

[AJR01] Daron Acemoglu, Simon Johnson, and James A Robinson. The colonial origins of comparative development: an empirical investigation. *The American Economic Review*, 91(5):1369–1401, 2001.

[AHL+22] Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: a continuous-time approach. *Review of Economic Studies*, 89(1):45–86, 2022.

[AP86] Anat R. Admati and Paul Pfleiderer. A monopolistic market for information. *Journal of Economic Theory*, 39(2):400–438, 1986.

[Aiy94] S Rao Aiyagari. Uninsured Idiosyncratic Risk and Aggregate Saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994.

[Alc50] Armen A Alchian. Uncertainty, evolution, and economic theory. *Journal of political economy*, 58(3):211–221, 1950.

[AS66] S. M. Ali and S. D. Silvey. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society, Series B*, 28(1):131–142, 1966.

[AA26] Fernando Alvarez and David Argente. Strategic complementarity and persistence in linear-quadratic mean field games. Working Paper, University of Chicago and Yale University, October 2026.

[AJ05] Fernando Alvarez and Urban J. Jermann. Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016, 2005. doi:10.1111/j.1468-0262.2005.00643.x.

[AL14] Fernando Alvarez and Francesco Lippi. Price setting with menu cost for multiproduct firms. *Econometrica*, 82(1):89–135, 2014.

[AV99] Fernando Alvarez and Marcelo Veracierto. Labor-market policies in an equilibrium search model. In Ben S. Bernanke and Julio J. Rotemberg, editors, *NBER Macroeconomics Annual 1999*, volume 14, pages 265–304. MIT Press, Cambridge, MA, 1999.

[AV01] Fernando Alvarez and Marcelo Veracierto. Severance payments in an economy with frictions. *Journal of Monetary Economics*, 47(3):477–498, 2001. doi:10.1016/S0304-3932(01)00057-2.

[AN00] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of Information Geometry*. Volume 191 of Translations of Mathematical Monographs. American Mathematical Society and Oxford University Press, Providence, RI, 2000.

[ABL02] Torben G. Andersen, Luca Benzoni, and Jesper Lund. An empirical investigation of continuous-time equity return models. *Journal of Finance*, 57(3):1239–1284, 2002. doi:10.1111/1540-6261.00460.

[AM05] D. B. O. Anderson and J. B. Moore. *Optimal Filtering*. Dover Publications, 2005.

[AHMS96] E. W. Anderson, L. P. Hansen, E. R. McGrattan, and T. J. Sargent. Mechanics of Forming and Estimating Dynamic Linear Economies. In *Handbook of Computational Economics*. Elsevier, vol 1 edition, 1996.

[AHS03a] Evan W. Anderson, Lars Peter Hansen, and Thomas J. Sargent. A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association*, 1(1):68–123, 2003.

[AHS03b] Evan W. Anderson, Lars Peter Hansen, and Thomas J. Sargent. A Quartet of Semigroups for Model Specification, Robustness, Prices of Risk, and Model Detection. *Journal of the European Economic Association*, 1(1):68–123, March 2003. URL: https://ideas.repec.org/a/tpr/jeurec/v1y2003i1p68-123.html, doi:.

[And76] Harald Anderson. Estimation of a proportion through randomized response. *International Statistical Review/Revue Internationale de Statistique*, pages 213–217, 1976.

[And51] T. W. Anderson. Estimating linear restrictions on regression coefficients for multivariate normal distributions. *Annals of Mathematical Statistics*, 22(3):327–351, 1951.

[And91] Donald W. K. Andrews. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59(3):817–858, 1991. doi:10.2307/2938229.

[And93] Donald W. K. Andrews. Tests for parameter instability and structural change with unknown change point. *Econometrica*, 61(4):821–856, 1993.

[AP03] Andrew Ang and Monika Piazzesi. A no-arbitrage vector autoregression of term structure dynamics with macroeconomic and latent variables. *Journal of Monetary Economics*, 50(4):745–787, 2003.

[Apo90] George Apostolakis. The concept of probability in safety assessments of technological systems. *Science*, 250(4986):1359–1364, 1990.

[Ari94] Jasmina Arifovic. Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, 18(1):3–28, 1994.

[Ari96] Jasmina Arifovic. The behavior of the exchange rate in the genetic algorithm and experimental economies. *Journal of Political Economy*, 104(3):510–541, 1996. doi:10.1086/262032.

[AB17] Martin Arjovsky and Léon Bottou. Towards principled methods for training generative adversarial networks. In *International Conference on Learning Representations*. 2017.

[Arr64] Kenneth J. Arrow. The role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies*, 31(2):91–96, 1964.

[ABG49] Kenneth J. Arrow, David Blackwell, and M. A. Girshick. Bayes and minimax solutions of sequential decision problems. *Econometrica*, 17(3/4):213–244, 1949.

[Art91] W. Brian Arthur. Designing economic agents that act like human agents: a behavioral approach to bounded rationality. *American Economic Review, Papers and Proceedings*, 81(2):353–359, 1991.

[AK87] Alan J Auerbach and Laurence J Kotlikoff. *Dynamic fiscal policy*. Cambridge University Press, Cambridge, 1987.

[Aum76] Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

[Axe87] Robert Axelrod. The evolution of strategies in the iterated prisoner's dilemma. In Lawrence Davis, editor, *Genetic Algorithms and Simulated Annealing*, pages 32–41. Morgan Kaufmann, Los Altos, CA, 1987.

[Axt01] Robert L Axtell. Zipf distribution of us firm sizes. *science*, 293(5536):1818–1820, 2001.

[BD85] David Backus and John Driffill. Inflation and reputation. *American Economic Review*, 75(3):530–538, 1985.

[BLS23] Isaac Baley, Lars Ljungqvist, and Thomas J. Sargent. Cross-phenomenon restrictions: unemployment effects of layoff costs and quit turbulence. *Review of Economic Dynamics*, 50:43–60, 2023. doi:10.1016/j.red.2023.07.005.

[BLS25] Isaac Baley, Lars Ljungqvist, and Thomas J. Sargent. Returns to labour mobility. *The Economic Journal*, 135(666):430–454, 2025. doi:10.1093/ej/ueae054.

[BG86] Nathan S. Balke and Robert J. Gordon. Appendix b: historical data. In Robert J. Gordon, editor, *The American Business Cycle: Continuity and Change*, pages 781–850. University of Chicago Press, Chicago, 1986.

[Bal95] Laurence Ball. Time-consistent policy and persistent changes in inflation. *Journal of Monetary Economics*, 36(2):329–350, 1995.

[BR90] Laurence Ball and David Romer. Real rigidities and the non-neutrality of money. *Review of Economic Studies*, 57(2):183–203, 1990. doi:10.2307/2297377.

[BGHT95] Ravi Bansal, A. Ronald Gallant, Robert Hussey, and George Tauchen. Nonparametric estimation of structural models for high-frequency currency market data. *Journal of Econometrics*, 66(1-2):251–287, 1995. doi:10.1016/0304-4076(94)01618-A.

[BY04] Ravi Bansal and Amir Yaron. Risks for the long run: a potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004. doi:10.1111/j.1540-6261.2004.00670.x.

[BHS09] Francisco Barillas, Lars Peter Hansen, and Thomas J. Sargent. Doubts or variability? *Journal of Economic Theory*, 144(6):2388–2418, November 2009. URL: https://ideas.repec.org/a/eee/jetheo/v144y2009i6p2388-2418.html, doi:.

[Bar79] Robert J Barro. On the Determination of the Public Debt. *Journal of Political Economy*, 87(5):940–971, 1979.

[Bar06] Robert J Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.

[Bar86] Robert J. Barro. Reputation in a model of monetary policy with incomplete information. *Journal of Monetary Economics*, 17(1):3–20, 1986.

[BG83] Robert J. Barro and David B. Gordon. Rules, discretion and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12(1):101–121, 1983.

[BK99] Marianne Baxter and Robert G. King. Measuring business cycles: approximate band-pass filters for economic time series. *Review of Economics and Statistics*, 81(4):575–593, 1999.

[BacsarB08] Tamer Başar and Pierre Bernhard. *H-infinity optimal control and related minimax design problems: a dynamic game approach*. Springer Science & Business Media, 2008.

[Bec75] Gary S. Becker. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. National Bureau of Economic Research, New York, 2nd edition, 1975.

[BS08] Luca Benati and Paolo Surico. Evolving U.S. monetary policy and the decline of inflation predictability. *Journal of the European Economic Association*, 6(2–3):634–646, 2008.

[BB18] Jess Benhabib and Alberto Bisin. Skewed wealth distributions: theory and empirics. *Journal of Economic Literature*, 56(4):1261–91, 2018.

[BBZ15] Jess Benhabib, Alberto Bisin, and Shenghao Zhu. The wealth distribution in bewley economies with capital income risk. *Journal of Economic Theory*, 159:489–515, 2015.

[BF94] Jess Benhabib and Roger E. A. Farmer. Indeterminacy and increasing returns. *Journal of Economic Theory*, 63(1):19–41, 1994. doi:10.1006/jeth.1994.1031.

[BMetivierP90] Albert Benveniste, Michel Métivier, and Pierre Priouret. *Adaptive Algorithms and Stochastic Approximations*. Springer-Verlag, Berlin, 1990.

[BS79] L M Benveniste and J A Scheinkman. On the Differentiability of the Value Function in Dynamic Models of Economics. *Econometrica*, 47(3):727–732, 1979.

[BB15] Dirk Bergemann and Alessandro Bonatti. Selling cookies. *American Economic Journal: Microeconomics*, 7(3):259–294, 2015.

[BBS18] Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018. doi:10.1257/aer.20161079.

[BValimaki97] Dirk Bergemann and Juuso Välimäki. Market diffusion with two-sided learning. *RAND Journal of Economics*, 28(4):773–795, 1997.

[BValimaki00] Dirk Bergemann and Juuso Välimäki. Experimentation in markets. *Review of Economic Studies*, 67(2):213–234, 2000. doi:10.1111/1467-937X.00128.

[Ber66] Robert H. Berk. Limiting behavior of posterior distributions when the model is incorrect. *Annals of Mathematical Statistics*, 37(1):51–58, 1966.

[Ber22] Ben S. Bernanke. *21st Century Monetary Policy: The Federal Reserve from the Great Inflation to COVID-19*. W. W. Norton and Company, New York, 2022.

[BM98] Ben S. Bernanke and Ilian Mihov. Measuring monetary policy. *Quarterly Journal of Economics*, 113(3):869–902, 1998.

[Ber76] Dimitri P. Bertsekas. *Dynamic Programming and Stochastic Control*. Academic Press, New York, 1976.

[Bet17] Michael Betancourt. A conceptual introduction to Hamiltonian Monte Carlo. *arXiv preprint arXiv:1701.02434*, 2017.

[Bew77] Truman Bewley. The permanent income hypothesis: a theoretical formulation. *Journal of Economic Theory*, 16(2):252–292, 1977.

[Bew86] Truman F Bewley. Stationary monetary equilibrium with a continuum of independently fluctuating consumers. In Werner Hildenbran and Andreu Mas-Colell, editors, *Contributions to Mathematical Economics in Honor of Gerard Debreu*, pages 27–102. North-Holland, Amsterdam, 1986.

[Bis06] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[Bla95] Fischer Black. Interest rates as options. *Journal of Finance*, 50(5):1371–1376, 1995.

[Bla51] David Blackwell. Comparison of experiments. In Jerzy Neyman, editor, *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 93–102. Berkeley, CA, 1951. University of California Press.

[Bla53] David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272, 1953. doi:10.1214/aoms/1177729032.

[Bla65] David Blackwell. Discounted dynamic programming. *Annals of Mathematical Statistics*, 36(1):226–235, 1965.

[BD62] David Blackwell and Lester E. Dubins. Merging of opinions with increasing information. *Annals of Mathematical Statistics*, 33(3):882–886, 1962.

[BS86a] Olivier J. Blanchard and Lawrence H. Summers. Hysteresis and the European unemployment problem. In Stanley Fischer, editor, *NBER Macroeconomics Annual 1986, Volume 1*, pages 15–78. MIT Press, 1986.

[BK80] Olivier Jean Blanchard and Charles M Kahn. The Solution of Linear Difference Models under Rational Expectations. *Econometrica*, 48(5):1305–1311, July 1980.

[BE92] Lawrence Blume and David Easley. Evolution and market behavior. *Journal of Economic Theory*, 58(1):9–40, 1992.

[BE06] Lawrence Blume and David Easley. If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets. *Econometrica*, 74(4):929–966, 2006.

[BCE+18] Lawrence E Blume, Timothy Cogley, David A Easley, Thomas J Sargent, and Viktor Tsyrennikov. A case for incomplete markets. *Journal of Economic Theory*, 178:191–221, 2018.

[BE82] Lawrence E. Blume and David Easley. Learning to be rational. *Journal of Economic Theory*, 26(2):340–351, 1982. doi:10.1016/0022-0531(82)90008-4.

[BSS49] H. F. Bohnenblust, Lloyd S. Shapley, and Seymour Sherman. Reconnaissance in game theory. Technical Report RM-208, The RAND Corporation, Santa Monica, CA, 1949. Cited for the economic criterion for comparing experiments.

[BW92] Tim Bollerslev and Jeffrey M. Wooldridge. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. *Econometric Reviews*, 11(2):143–172, 1992. doi:10.1080/07474939208800229.

[BH99] Patrick Bolton and Christopher Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.

[Bor20] Jaroslav Borovička. Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. *Journal of Political Economy*, 128(1):206–251, 2020. doi:10.1086/704072.

[BHS16] Jaroslav Borovička, Lars Peter Hansen, and José A. Scheinkman. Misspecified recovery. *Journal of Finance*, 71(6):2493–2544, 2016. doi:10.1111/jofi.12404.

[BO94] John F. Boschen and Christopher M. Otrok. Long-run neutrality and superneutrality in an ARIMA framework: comment. *American Economic Review*, 84(5):1470–1473, 1994.

[Bra82] Margaret M. Bray. Learning, estimation, and the stability of rational expectations. *Journal of Economic Theory*, 26(2):318–339, 1982. doi:10.1016/0022-0531(82)90007-2.

[BK87] Margaret M. Bray and David M. Kreps. Rational learning and rational expectations. In George R. Feiwel, editor, *Arrow and the Ascent of Modern Economic Theory*, pages 597–625. Palgrave Macmillan, London, 1987. doi:10.1007/978-1-349-07239-2_19.

[BS86b] Margaret M. Bray and N. E. Savin. Rational expectations equilibria, learning and model specification. *Econometrica*, 54(5):1129–1160, 1986. doi:10.2307/1912325.

[Bre79] Douglas T. Breeden. An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7(3):265–296, 1979. doi:10.1016/0304-405X(79)90016-3.

[BL78] Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978. doi:10.1086/296025.

[Bro82] William A Brock. Asset prices in a production economy. In John J. McCall, editor, *The Economics of Information and Uncertainty*, pages 1–46. University of Chicago Press, Chicago, 1982.

[Bro72] William A. Brock. On models of expectations that arise from maximizing behavior of economic agents over time. *Journal of Economic Theory*, 5(3):348–376, 1972.

[Bro74] William A. Brock. Money and growth: the case of long run perfect foresight. *International Economic Review*, 15(3):750–777, 1974. doi:10.2307/2525739.

[BH97] William A. Brock and Cars H. Hommes. A rational route to randomness. *Econometrica*, 65(5):1059–1095, 1997.

[BM72] William A. Brock and Leonard J. Mirman. Optimal economic growth and uncertainty: the discounted case. *Journal of Economic Theory*, 4(3):479–513, 1972.

[BF90] Michael Bruno and Stanley Fischer. Seigniorage, operating rules, and the high inflation trap. *Quarterly Journal of Economics*, 105(2):353–374, 1990. doi:10.2307/2937791.

[BK19] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering: Machine Learning, Dynamical Systems, and Control*. Cambridge University Press, 2019.

[BK22] Steven L. Brunton and J. Nathan Kutz. *Data-Driven Science and Engineering, Second Edition*. Cambridge University Press, New York, 2022.

[Buc04] James A. Bucklew. *An Introduction to Rare Event Simulation*. Springer Verlag, New York, 2004.

[Bul94] James Bullard. Learning equilibria. *Journal of Economic Theory*, 64(2):468–485, 1994.

[Bul18] James Bullard. The case of the disappearing Phillips curve. Presentation at the ECB Forum on Central Banking, Sintra, Portugal, June 2018.

[BD93] James Bullard and John Duffy. Learning in a large square economy. Mimeo, Federal Reserve Bank of St. Louis, March 1993.

[BDM+16] Dariusz Buraczewski, Ewa Damek, Thomas Mikosch, and others. *Stochastic models with power-law tails*. Springer, 2016.

[Bur23] Jennifer Burns. *Milton Friedman: The Last Conservative*. Farrar, Straus, and Giroux, New York, 2023.

[Cag56] Philip Cagan. The monetary dynamics of hyperinflation. In Milton Friedman, editor, *Studies in the Quantity Theory of Money*, pages 25–117. University of Chicago Press, Chicago, 1956.

[Cam87] John Y. Campbell. Does Saving Anticipate Declining Labor Income? An Alternative Test of the Permanent Income Hypothesis. *Econometrica*, 55(6):1249–1273, 1987.

[CS88] John Y. Campbell and Robert J. Shiller. The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors. *Review of Financial Studies*, 1(3):195–228, 1988.

[CS91] John Y. Campbell and Robert J. Shiller. Yield spreads and interest rate movements: A bird's eye view. *Review of Economic Studies*, 58(3):495–514, 1991.

[Cap85] Andrew S Caplin. The variability of aggregate demand with (s, s) inventory policies. *Econometrica*, pages 1395–1409, 1985.

[CE09] Giacomo Carboni and Martin Ellison. The Great Inflation and the Greenbook. *Journal of Monetary Economics*, 56(6):831–841, 2009.

[CFP09] Charles T. Carlstrom, Timothy S. Fuerst, and Matthias Paustian. Inflation persistence, monetary policy, and the Great Moderation. *Journal of Money, Credit and Banking*, 41(4):767–786, 2009.

[CD18] René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications*. Springer, Cham, 2018.

[CY12] Peter Carr and Jiming Yu. Risk, return, and Ross recovery. *Journal of Derivatives*, 20(1):38–59, 2012. doi:10.3905/jod.2012.20.1.038.

[Car01] Christopher D Carroll. A Theory of the Consumption Function, with and without Liquidity Constraints. *Journal of Economic Perspectives*, 15(3):23–45, 2001.

[Car06] Christopher D Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006.

[CK94] C. K. Carter and R. Kohn. On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553, 1994.

[Cas65] David Cass. Optimum growth in an aggregative model of capital accumulation. *Review of Economic Studies*, 32(3):233–240, 1965.

[Cha98] Roberto Chang. Credible monetary policy in an infinite horizon model: recursive approaches. *Journal of Economic Theory*, 81(2):431–461, 1998.

[CM88] A. Chaudhuri and R. Mukerjee. *Randomized Response: Theory and Techniques*. Marcel Dekker, New York, 1988.

[CW98] Xiaohong Chen and Halbert White. Nonparametric adaptive learning with feedback. *Journal of Economic Theory*, 82(1):190–222, 1998. doi:10.1006/jeth.1998.2426.

[Che52] Herman Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23(4):493–507, 1952.

[CM95] In-Koo Cho and Akihiko Matsui. Induction and the Ramsey policy. *Journal of Economic Dynamics and Control*, 19(5–7):1113–1140, 1995.

[CWS02] In-Koo Cho, Noah Williams, and Thomas J. Sargent. Escaping Nash inflation. *Review of Economic Studies*, 69(1):1–40, 2002.

[Cho68] Gregory C. Chow. The acceleration principle and the nature of business cycles. *The Quarterly Journal of Economics*, 82(3):403–418, aug 1968.

[CL69] Gregory C. Chow and Richard E. Levitan. Nature of business cycles implicit in a linear economic model. *The Quarterly Journal of Economics*, 83(3):504–517, aug 1969.

[Chu90] Heetaik Chung. *Did Policy Makers Really Believe in the Phillips Curve? An Econometric Test*. PhD thesis, University of Minnesota, 1990.

[CGaliG00] Richard Clarida, Jordi Galí, and Mark Gertler. Monetary policy rules and macroeconomic stability: evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–180, 2000.

[Cla71] Edward H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[Coa37] Ronald H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.

[CS01] Timothy Cogley and Thomas J. Sargent. Evolving post-World War II U.S. inflation dynamics. *NBER Macroeconomics Annual*, 16:331–373, 2001.

[CS05a] Timothy Cogley and Thomas J. Sargent. Drifts and volatilities: monetary policies and outcomes in the post WWII US. *Review of Economic Dynamics*, 8(2):262–302, 2005.

[CS05b] Timothy Cogley and Thomas J. Sargent. The conquest of US inflation: learning and robustness to model uncertainty. *Review of Economic Dynamics*, 8(2):528–563, 2005.

[CDGJ02] Courtney Coile, Peter Diamond, Jonathan Gruber, and Alain Jousten. Delays in claiming Social Security benefits. *Journal of Public Economics*, 84(3):357–385, 2002. doi:10.1016/S0047-2727(01)00129-3.

[Col90] Wilbur John Coleman. Solving the Stochastic Growth Model by Policy-Function Iteration. *Journal of Business & Economic Statistics*, 8(1):27–29, 1990.

[CP73] Thomas F. Cooley and Edward C. Prescott. An adaptive regression model. *International Economic Review*, 14(2):364–371, 1973.

[CH06] Russell W. Cooper and John C. Haltiwanger. On the nature of capital adjustment costs. *Review of Economic Studies*, 73(3):611–633, 2006. doi:10.1111/j.1467-937X.2006.00389.x.

[CIR85] John C. Cox, Jonathan E. Ingersoll, and Stephen A. Ross. A theory of the term structure of interest rates. *Econometrica*, 53(2):385–407, 1985.

[CK81] Vincent P. Crawford and Elsie Marie Knoer. Job matching with heterogeneous firms and workers. *Econometrica*, 49(2):437–450, 1981.

[Csiszar63] Imre Csiszár. Eine informationstheoretische Ungleichung und ihre Anwendung auf den Beweis der Ergodizität von Markoffschen Ketten. *Magyar Tud. Akad. Mat. Kutató Int. Közl.*, 8:85–108, 1963.

[DN03] Herbert A. David and Haikady N. Nagaraja. *Order Statistics*. Wiley, Hoboken, NJ, 3rd edition, 2003.

[DFH06] Steven J. Davis, R. Jason Faberman, and John Haltiwanger. The flow approach to labor markets: new data sources and micro-macro links. *Journal of Economic Perspectives*, 20(3):3–26, 2006.

[dF37] Bruno de Finetti. La prevision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1 – 68, 1937. English translation in Kyburg and Smokler (eds.), \it Studies in Subjective Probability, Wiley, New York, 1964.

[DLSSW91] J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. The survival of noise traders in financial markets. *Journal of Business*, 64(1):1–19, 1991.

[Dea91] Angus Deaton. Saving and Liquidity Constraints. *Econometrica*, 59(5):1221–1248, 1991.

[DP94] Angus Deaton and Christina Paxson. Intertemporal Choice and Inequality. *Journal of Political Economy*, 102(3):437–467, 1994.

[DeC79] Stephen J. DeCanio. Rational expectations and learning from experience. *Quarterly Journal of Economics*, 93(1):47–57, 1979.

[DNP15] Marco Del Negro and Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy: a corrigendum. *Review of Economic Studies*, 82(4):1342–1345, 2015.

[DeL97] J. Bradford DeLong. America's peacetime inflation: the 1970s. In Christina D. Romer and David H. Romer, editors, *Reducing Inflation: Motivation and Strategy*, pages 247–280. University of Chicago Press, 1997.

[DGS86] Gabrielle Demange, David Gale, and Marilda Sotomayor. Multi-item auctions. *Journal of Political Economy*, 94(4):863–872, 1986.

[DZ98] Amir Dembo and Ofer Zeitouni. *Large Deviations Techniques and Applications*. Springer-Verlag, New York, 2nd edition, 1998.

[DH10] Wouter J Den Haan. Comparison of solutions to the incomplete markets model with aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):4–27, 2010.

[dHM90] Wouter J. den Haan and Albert Marcet. Solving the stochastic growth model by parameterizing expectations. *Journal of Business & Economic Statistics*, 8(1):31–34, 1990.

[DF86] Persi Diaconis and David Freedman. On the consistency of Bayes estimates. *The Annals of Statistics*, 14(1):1–26, 1986. doi:10.1214/aos/1176349830.

[Dia65] Peter A Diamond. National debt in a neoclassical growth model. *The American Economic Review*, 55(5):1126–1150, 1965.

[Dia67] Peter A. Diamond. The role of a stock market in a general equilibrium model with technological uncertainty. *American Economic Review*, 57(4):759–776, 1967.

[Doo49] Joseph L. Doob. Application of the theory of martingales. In *Le Calcul des Probabilités et ses Applications*, number 13 in Colloques Internationaux du Centre National de la Recherche Scientifique, pages 23–27. Centre National de la Recherche Scientifique, Paris, 1949.

[Doo53] Joseph L. Doob. *Stochastic Processes*. Wiley, New York, 1953.

[DS10] Ulrich Doraszelski and Mark Satterthwaite. Computable markov-perfect industry dynamics. *The RAND Journal of Economics*, 41(2):215–243, 2010.

[DSS58] Robert Dorfman, Paul A. Samuelson, and Robert M. Solow. *Linear Programming and Economic Analysis: Revised Edition*. McGraw Hill, New York, 1958.

[DLP13] Y E Du, Ehud Lehrer, and A D Y Pauzner. Competitive economy as a ranking device over networks. submitted, 2013.

[DKPR87] Simon Duane, A. D. Kennedy, Brian J. Pendleton, and Duncan Roweth. Hybrid Monte Carlo. *Physics Letters B*, 195(2):216–222, 1987. doi:10.1016/0370-2693(87)91197-X.

[Dud02] R M Dudley. *Real Analysis and Probability*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2002.

[DS08] Gregory R. Duffee and Richard H. Stanton. Evidence on simulation inference for near unit-root processes with implications for term structure estimation. *Journal of Financial Econometrics*, 6(1):108–142, 2008. doi:10.1093/jjfinec/nbm014.

[DE92] Darrell Duffie and Larry G. Epstein. Stochastic differential utility. *Econometrica*, 60(2):353–394, 1992.

[DS93] Darrell Duffie and Kenneth J. Singleton. Simulated moments estimation of Markov models of asset prices. *Econometrica*, 61(4):929–952, 1993. doi:10.2307/2951768.

[DUW00] Bernard Dumas, Raman Uppal, and Tan Wang. Efficient intertemporal allocations with recursive utility. *Journal of Economic Theory*, 93(2):154–183, 2000.

[DRS89] Timothy Dunne, Mark J Roberts, and Larry Samuelson. The growth and failure of us manufacturing plants. *The Quarterly Journal of Economics*, 104(4):671–698, 1989.

[DE97] Paul Dupuis and Richard S. Ellis. *A Weak Convergence Approach to the Theory of Large Deviations*. John Wiley and Sons, New York, 1997.

[DK87] Paul Dupuis and Harold J. Kushner. Asymptotic behavior of constrained stochastic approximations via the theory of large deviations. *Probability Theory and Related Fields*, 75(2):223–244, 1987.

[DK89] Paul Dupuis and Harold J. Kushner. Stochastic approximation and large deviations: upper bounds and w.p.1 convergence. *SIAM Journal on Control and Optimization*, 27(5):1108–1135, 1989.

[Dut91] Prajit K. Dutta. What do discounted optima converge to? a theory of discount rate asymptotics in economic models. *Journal of Economic Theory*, 55(1):64–94, 1991.

[EOS07] Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.

[ESAW18] Ashraf Ben El-Shanawany, Keith H Ardron, and Simon P Walker. Lognormal approximations of fault tree uncertainty distributions. *Risk Analysis*, 38(8):1576–1584, 2018.

[EY07] Martin Ellison and Tony Yates. Escaping volatile inflation. *Journal of Money, Credit and Banking*, 39(4):981–993, 2007.

[EG87] Robert F Engle and Clive W J Granger. Co-integration and Error Correction: Representation, Estimation, and Testing. *Econometrica*, 55(2):251–276, 1987.

[EW94] Larry G. Epstein and Tan Wang. Intertemporal Asset Pricing under Knightian Uncertainty. *Econometrica*, 62(2):283–322, 1994.

[EZ89] Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica*, 57(4):937–969, 1989. doi:10.2307/1913778.

[EP95] Richard Ericson and Ariel Pakes. Markov-perfect industry dynamics: a framework for empirical work. *The Review of Economic Studies*, 62(1):53–82, 1995.

[EP16] Ignacio Esponda and Demian Pouzo. Berk–Nash equilibrium: a framework for modeling agents with misspecified models. *Econometrica*, 84(3):1093–1130, 2016.

[Eva87] David S Evans. The relationship between firm growth, size, and age: estimates for 100 manufacturing industries. *The Journal of Industrial Economics*, pages 567–581, 1987.

[EH01] G W Evans and S Honkapohja. *Learning and Expectations in Macroeconomics*. Frontiers of Economic Research. Princeton University Press, 2001.

[EH05] George W Evans and Seppo Honkapohja. An interview with thomas j. sargent. *Macroeconomic Dynamics*, 9(4):561–583, 2005.

[Eva85] George W. Evans. Expectational stability and the multiple equilibria problem in linear rational expectations models. *Quarterly Journal of Economics*, 100(4):1217–1233, 1985. doi:10.2307/1885681.

[Eva89] George W. Evans. The fragility of sunspots and bubbles. *Journal of Monetary Economics*, 23(2):297–317, 1989.

[Fab82] Joseph F Faber. Life tables for the United States: 1900–2050. Actuarial Study Actuarial Study No. 87, Social Security Administration, Office of the Actuary, 1982.

[FSTD15] Pablo Fajgelbaum, Edouard Schaal, and Mathieu Taschereau-Dumouchel. Uncertainty traps. Technical Report, National Bureau of Economic Research, 2015.

[FB87] Eugene F. Fama and Robert R. Bliss. The information in long-maturity forward rates. *American Economic Review*, 77(4):680–692, 1987.

[FR24] Lasse Fischer and Aaditya Ramdas. Improving the (approximate) sequential probability ratio test by avoiding overshoot. *arXiv preprint arXiv:2410.16076*, 2024.

[FPS77] Michael A Fligner, George E Policello, and Jagbir Singh. A comparison of two randomized response survey methods with consideration for the level of respondent protection. *Communications in Statistics-Theory and Methods*, 6(15):1511–1524, 1977.

[FW98] Mark I. Freidlin and Alexander D. Wentzell. *Random Perturbations of Dynamical Systems*. Springer-Verlag, New York, 2nd edition, 1998.

[Fri56] M. Friedman. *A Theory of the Consumption Function*. Princeton University Press, 1956.

[Fri53] Milton Friedman. *Essays in positive economics*. University of Chicago press, 1953.

[Fri57] Milton Friedman. *A Theory of the Consumption Function*. Princeton University Press, Princeton, New Jersey, 1957.

[Fri68] Milton Friedman. The role of monetary policy. *American Economic Review*, 58(1):1–17, 1968.

[FF98] Milton Friedman and Rose D Friedman. *Two Lucky People*. University of Chicago Press, 1998.

[Fri33] Ragnar Frisch. Propagation problems and impulse problems in dynamic economics. In *Economic Essays in Honour of Gustav Cassel*, pages 171–203. Allen and Unwin, 1933.

[FL93] Drew Fudenberg and David K. Levine. Self-confirming equilibrium. *Econometrica*, 61(3):523–545, 1993.

[Gab16] Xavier Gabaix. Power laws in economics: an introduction. *Journal of Economic Perspectives*, 30(1):185–206, 2016.

[Gal89] David Gale. *The theory of linear economic models*. University of Chicago press, 1989.

[Gal16] Alfred Galichon. *Optimal Transport Methods in Economics*. Princeton University Press, Princeton, New Jersey, 2016.

[Gal87] A. Ronald Gallant. *Nonlinear Statistical Models*. Wiley, 1987. doi:10.1002/9780470316719.

[GHT97] A. Ronald Gallant, David Hsieh, and George Tauchen. Estimation of stochastic volatility models with diagnostics. *Journal of Econometrics*, 81(1):159–192, 1997. doi:10.1016/S0304-4076(97)00039-0.

[GN87] A. Ronald Gallant and Douglas W. Nychka. Semi-nonparametric maximum likelihood estimation. *Econometrica*, 55(2):363–390, 1987. doi:10.2307/1913241.

[GT89] A. Ronald Gallant and George Tauchen. Seminonparametric estimation of conditionally constrained heterogeneous processes: asset pricing applications. *Econometrica*, 57(5):1091–1120, 1989. doi:10.2307/1913624.

[GT96] A. Ronald Gallant and George Tauchen. Which moments to match? *Econometric Theory*, 12(4):657–681, 1996. doi:10.1017/S0266466600006976.

[GT10] A. Ronald Gallant and George Tauchen. Simulated score methods and indirect inference for continuous-time models. In Yacine Aït-Sahalia and Lars Peter Hansen, editors, *Handbook of Financial Econometrics: Tools and Techniques*, volume 1, pages 427–477. Elsevier, 2010. doi:10.1016/B978-0-444-50897-3.50011-0.

[Gew82] John Geweke. Measurement of linear dependence and feedback between multiple time series. *Journal of the American Statistical Association*, 77(378):304–313, 1982.

[Gia05] Chryssi Giannitsarou. E-stability does not imply learnability. *Macroeconomic Dynamics*, 9(2):276–287, 2005.

[Gib31] Robert Gibrat. *Les inégalités économiques: Applications d'une loi nouvelle, la loi de l'effet proportionnel*. PhD thesis, Recueil Sirey, 1931.

[GS89] Itzhak Gilboa and David Schmeidler. Maxmin Expected Utility with Non-Unique Prior. *Journal of Mathematical Economics*, 18(2):141–153, apr 1989.

[Gir60] Igor V. Girsanov. On transforming a certain class of stochastic processes by absolutely continuous substitution of measures. *Theory of Probability and Its Applications*, 5(3):285–301, 1960.

[Gol89] David E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley, Reading, MA, 1989.

[GPAM+14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2672–2680. 2014.

[Gor95] Geoffrey J Gordon. Stable function approximation in dynamic programming. In *Machine Learning Proceedings 1995*, pages 261–268. Elsevier, 1995.

[GMR93] Christian Gourieroux, Alain Monfort, and Eric Renault. Indirect inference. *Journal of Applied Econometrics*, 8(S1):S85–S118, 1993. doi:10.1002/jae.3950080507.

[Gra66] C. W. J. Granger. The typical spectral shape of an economic variable. *Econometrica*, 34(1):150–161, 1966.

[GAESH69] Bernard G Greenberg, Abdel-Latif A Abul-Ela, Walt R Simmons, and Daniel G Horvitz. The unrelated question randomized response model: theoretical framework. *Journal of the American Statistical Association*, 64(326):520–539, 1969.

[GKAH77] Bernard G Greenberg, Roy R Kuebler, James R Abernathy, and Daniel G Horvitz. Respondent hazards in the unrelated question randomized response model. *Journal of Statistical Planning and Inference*, 1(1):53–60, 1977.

[GS93] Moses A. Greenfield and Thomas J. Sargent. A probabilistic analysis of a catastrophic transuranic waste hoist accident at the wipp. Environmental Evaluation Group, Albuquerque, New Mexico, June 1993. URL: http://www.tomsargent.com/research/EEG-53.pdf.

[Gro76] Sanford J. Grossman. On the efficiency of competitive stock markets where trades have diverse information. *Journal of Finance*, 31(2):573–585, 1976.

[GS80] Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.

[Gro73] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[Hal87] Bronwyn H Hall. The relationship between firm size and firm growth in the us manufacturing sector. *The Journal of Industrial Economics*, pages 583–606, 1987.

[Hal71] Robert E Hall. The dynamic effects of fiscal policy in an economy with foresight. *The Review of Economic Studies*, 38(2):229–244, 1971.

[Hal78] Robert E Hall. Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence. *Journal of Political Economy*, 86(6):971–987, 1978.

[HM82] Robert E. Hall and Frederic S. Mishkin. The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households. *Econometrica*, 50(2):461–481, 1982.

[HTW67] Michael J Hamburger, Gerald L Thompson, and Roman L Weil. Computation of expansion rates for the generalized von neumann model of an expanding economy. *Econometrica, Journal of the Econometric Society*, pages 542–547, 1967.

[Ham05] James D Hamilton. What's real about the business cycle? *Federal Reserve Bank of St. Louis Review*, pages 435–452, 2005.

[Han92] Bruce E. Hansen. Testing for parameter instability in linear models. *Journal of Policy Modeling*, 14(4):517–533, 1992.

[Han93] Gary D Hansen. The cyclical and secular behaviour of the labour input: comparing efficiency units and hours worked. *Journal of Applied Econometrics*, 8(1):71–80, 1993.

[HS08] L P Hansen and T J Sargent. *Robustness*. Princeton University Press, 2008.

[HS13] L P Hansen and T J Sargent. *Recursive Models of Dynamic Linear Economies*. The Gorman Lectures in Economics. Princeton University Press, 2013.

[Han82] Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982.

[Han14] Lars Peter Hansen. Nobel lecture: uncertainty outside and inside economic models. *Journal of Political Economy*, 122(5):945–987, 2014.

[HHL08] Lars Peter Hansen, John C. Heaton, and Nan Li. Consumption strikes back? measuring long-run risk. *Journal of Political Economy*, 116(2):260–302, 2008. doi:10.1086/588200.

[HJ91] Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991. doi:10.1086/261750.

[HR87] Lars Peter Hansen and Scott F Richard. The role of conditioning information in deducing testable restrictions implied by dynamc asset pricing models. *Econometrica*, 55(3):587–613, May 1987.

[HS80] Lars Peter Hansen and Thomas J Sargent. Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and control*, 2:7–46, 1980.

[HS95] Lars Peter Hansen and Thomas J Sargent. Discounted linear exponential quadratic gaussian control. *IEEE Transactions on Automatic control*, 40(5):968–971, 1995.

[HS04] Lars Peter Hansen and Thomas J Sargent. Certainty equivalence and model uncertainty. In *Conference on Models and Monetary Policy: Research in the Tradition of Dale Henderson, Richard Porter, and Peter Tinsley (http://www. federalreserve. gov/events/conferences/mmp2004/pdf/hansensargent. pdf)*. 2004.

[HS93] Lars Peter Hansen and Thomas J. Sargent. Seasonality and approximation errors in rational expectations models. *Journal of Econometrics*, 55(1–2):21–55, 1993.

[HS01] Lars Peter Hansen and Thomas J. Sargent. Robust control and model uncertainty. *American Economic Review*, 91(2):60–66, 2001.

[HST99] Lars Peter Hansen, Thomas J. Sargent, and Thomas D. Tallarini. Robust Permanent Income and Pricing. *Review of Economic Studies*, 66(4):873–907, 1999. URL: https://ideas.repec.org/a/oup/restud/v66y1999i4p873-907..html, doi:.

[HS09] Lars Peter Hansen and José A. Scheinkman. Long-term risk: an operator approach. *Econometrica*, 77(1):177–234, 2009. doi:10.3982/ECTA6761.

[HS82] Lars Peter Hansen and Kenneth J Singleton. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society*, pages 1269–1286, 1982.

[HS83] Lars Peter Hansen and Kenneth J Singleton. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of political economy*, 91(2):249–265, 1983.

[HK78] J. Michael Harrison and David M. Kreps. Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 92(2):323–336, 1978.

[HK79] J. Michael Harrison and David M. Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3):381–408, June 1979.

[Har67] John C. Harsanyi. Games with incomplete information played by "Bayesian" players, I–III part I. The basic model. *Management Science*, 14(3):159–182, 1967.

[Har68a] John C. Harsanyi. Games with incomplete information played by "Bayesian" players, I–III part III. The basic probability distribution of the game. *Management Science*, 14(7):486–502, 1968.

[Har68b] John C. Harsanyi. Games with incomplete information played by "Bayesian" players, I–III part II. Bayesian equilibrium points. *Management Science*, 14(5):320–334, 1968.

[Har75] Oliver D. Hart. On the optimality of equilibrium when the market structure is incomplete. *Journal of Economic Theory*, 11(3):418–443, 1975.

[Hay82] Fumio Hayashi. Tobin's marginal q and average q: a neoclassical interpretation. *Econometrica*, 50(1):213–224, 1982. doi:10.2307/1912538.

[HL96] John Heaton and Deborah J Lucas. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy*, pages 443–487, 1996.

[HLL96] O Hernandez-Lerma and J B Lasserre. *Discrete-Time Markov Control Processes: Basic Optimality Criteria*. Number Vol 1 in Applications of Mathematics Stochastic Modelling and Applied Probability. Springer, 1996.

[HS55] Edwin Hewitt and Leonard J. Savage. Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society*, 80(2):470–501, 1955.

[Hil75] Bruce M. Hill. A simple general approach to inference about the tail of a distribution. *Annals of Statistics*, 3(5):1163–1174, 1975.

[HG14] Matthew D. Hoffman and Andrew Gelman. The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15(47):1593–1623, 2014.

[Hol75] John H. Holland. *Adaptation in Natural and Artificial Systems*. University of Michigan Press, Ann Arbor, 1975.

[HHNT86] John H. Holland, Keith J. Holyoak, Richard E. Nisbett, and Paul R. Thagard. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA, 1986.

[HMMS60] Charles Holt, Franco Modigliani, John F. Muth, and Herbert Simon. *Planning Production, Inventories, and Work Force*. Prentice-Hall International Series in Management, New Jersey, 1960.

[Hop92] Hugo A Hopenhayn. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica: Journal of the Econometric Society*, pages 1127–1150, 1992.

[Hop82] John J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79(8):2554–2558, 1982.

[HIS97] He Huang, Selahattin Imrohoroglu, and Thomas J Sargent. Two computations to fund social security. *Macroeconomic Dynamics*, 1(1):7–44, 1997.

[HMalhameC06] Minyi Huang, Roland P. Malhamé, and Peter E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the nash certainty equivalence principle. *Communications in Information and Systems*, 6(3):221–252, 2006.

[Hug93] Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969, 1993.

[Hur50] Leonid Hurwicz. Least squares bias in time series. *Statistical inference in dynamic economic models*, 10:365–383, 1950.

[Hur62] Leonid Hurwicz. On the structural form of interdependent systems. In *Logic, Methodology and Philosophy of Science*, pages 232–239. Stanford University Press, Stanford, CA, 1962.

[HP62] Stephen Hymer and Peter Pashigian. Firm size and rate of growth. *Journal of Political Economy*, 70(6):556–569, 1962.

[Haggstrom02] Olle Häggström. *Finite Markov chains and algorithmic applications*. Volume 52. Cambridge University Press, 2002.

[Int02] Michael D Intriligator. *Mathematical optimization and economic theory*. SIAM, 2002.

[Ire03] Peter N. Ireland. Endogenous money or sticky prices? *Journal of Monetary Economics*, 50(8):1623–1648, 2003.

[Ire04] Peter N. Ireland. Technology shocks in the new keynesian model. *Review of Economics and Statistics*, 86(4):923–936, 2004.

[Janich94] K Jänich. *Linear Algebra*. Springer Undergraduate Texts in Mathematics and Technology. Springer, 1994.

[JK99] Matthew O. Jackson and Ehud Kalai. Reputation versus social learning. *Journal of Economic Theory*, 88(1):40–59, 1999. doi:10.1006/jeth.1999.2538.

[JKS99] Matthew O. Jackson, Ehud Kalai, and Rann Smorodinsky. Bayesian representation of stochastic processes under learning: de Finetti revisited. *Econometrica*, 67(4):875–893, 1999. doi:10.1111/1468-0262.00055.

[Jac73] D. H. Jacobson. Optimal stochastic linear systems with exponential performance criteria and their relation to differential games. *IEEE Transactions on Automatic Control*, 18(2):124–131, 1973.

[JPR94] Eric Jacquier, Nicholas G. Polson, and Peter E. Rossi. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics*, 12(4):371–389, 1994.

[Jef46] Harold Jeffreys. An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 186(1007):453–461, 1946.

[Jor82] James S. Jordan. The generic existence of rational expectations equilibrium in the higher dimensional case. *Journal of Economic Theory*, 26(2):224–243, 1982. doi:10.1016/0022-0531(82)90002-3.

[Jov79] Boyan Jovanovic. Firm-specific capital and turnover. *Journal of Political Economy*, 87(6):1246–1260, 1979.

[Jud90] K L Judd. Cournot versus bertrand: a dynamic resolution. Technical Report, Hoover Institution, Stanford University, 1990.

[Kak48] Shizuo Kakutani. On equivalence of infinite product measures. *Annals of Mathematics*, 49(1):214–224, 1948.

[KL93a] Ehud Kalai and Ehud Lehrer. Rational learning leads to Nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993. doi:10.2307/2951492.

[KL93b] Ehud Kalai and Ehud Lehrer. Subjective equilibrium in repeated games. *Econometrica*, 61(5):1231–1240, 1993. doi:10.2307/2951500.

[KL94] Ehud Kalai and Ehud Lehrer. Weak and strong merging of opinions. *Journal of Mathematical Economics*, 23(1):73–86, 1994. doi:10.1016/0304-4068(94)90037-X.

[KLS99] Ehud Kalai, Ehud Lehrer, and Rann Smorodinsky. Calibrated forecasting and merging. *Games and Economic Behavior*, 29(1–2):151–169, 1999. doi:10.1006/game.1998.0608.

[KG11] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[KMR93] Michihiro Kandori, George J. Mailath, and Rafael Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–56, 1993.

[KW81] John Kareken and Neil Wallace. On the indeterminacy of equilibrium exchange rates. *Quarterly Journal of Economics*, 96(2):207–222, 1981. doi:10.2307/1882388.

[KT81] Samuel Karlin and Howard M. Taylor. *A Second Course in Stochastic Processes*. Academic Press, 1981.

[Kas04] Kenneth Kasa. Learning, large deviations, and recurrent currency crises. *International Economic Review*, 45(1):141–173, 2004.

[KMT56] John G Kemeny, Oskar Morgenstern, and Gerald L Thompson. A generalization of the von neumann model of an expanding economy. *Econometrica, Journal of the Econometric Society*, pages 115–135, 1956.

[KW52] Jack Kiefer and Jacob Wolfowitz. Stochastic estimation of the maximum of a regression function. *Annals of Mathematical Statistics*, 23(3):462–466, 1952.

[Kih74] Richard E. Kihlstrom. A general theory of demand for information about product quality. *Journal of Economic Theory*, 8(4):413–439, 1974. doi:10.1016/0022-0531(74)90019-2.

[Kih84] Richard E. Kihlstrom. A Bayesian exposition of Blackwell's theorem on the comparison of experiments. In Marcel Boyer and Richard E. Kihlstrom, editors, *Bayesian Models in Economic Theory*, volume 5 of Studies in Bayesian Econometrics, pages 13–31. North-Holland, Amsterdam, 1984.

[Kih19] Richard E. Kihlstrom. The Role of Securities in the Allocation of Risk: Arrow's 1964 REStud Paper. Class notes, Wharton School, University of Pennsylvania, 2019.

[KM75] Richard E. Kihlstrom and Leonard J. Mirman. Information and market equilibrium. *The Bell Journal of Economics*, 6(1):357–376, 1975.

[KN99] Chang-Jin Kim and Charles R. Nelson. Has the U.S. economy become more stable? a Bayesian approach based on a Markov-switching model of the business cycle. *Review of Economics and Statistics*, 81(4):608–616, 1999.

[KSC98] Sangjoon Kim, Neil Shephard, and Siddhartha Chib. Stochastic volatility: likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3):361–393, 1998.

[Kim95] Miles S. Kimball. The quantitative analytics of the basic neomonetarist model. *Journal of Money, Credit and Banking*, 27(4):1241–1277, 1995.

[KW94] Robert G. King and Mark W. Watson. The post-war U.S. Phillips curve: a revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy*, 41:157–219, 1994.

[Kit14] Sagiri Kitao. Sustainable Social Security: four options. *Review of Economic Dynamics*, 17(4):756–779, 2014. doi:10.1016/j.red.2013.11.004.

[KW89] Nobuhiro Kiyotaki and Randall Wright. On money as a medium of exchange. *Journal of Political Economy*, 97(4):927–954, 1989. doi:10.1086/261634.

[Kle99] Paul Klemperer. Auction theory: a guide to the literature. *Journal of Economic Surveys*, 13(3):227–286, 1999.

[KW16] Peter J. Klenow and Jonathan L. Willis. Real rigidities and nominal price changes. *Economica*, 83(331):443–472, 2016.

[Koo47] Tjalling C. Koopmans. Measurement without theory. *Review of Economics and Statistics*, 29(3):161–172, 1947.

[Koo65] Tjalling C. Koopmans. On the concept of optimal economic growth. In Tjalling C. Koopmans, editor, *The Economic Approach to Development Planning*, pages 225–287. Rand McNally, Chicago, 1965.

[Kre88] David M. Kreps. *Notes on the Theory of Choice*. Westview Press, Boulder, Colorado, 1988.

[Kre98] David M. Kreps. Anticipated utility and dynamic choice. In Donald P. Jacobs, Ehud Kalai, and Morton I. Kamien, editors, *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983–1997*, pages 242–274. Cambridge University Press, Cambridge, 1998.

[KP78] David M. Kreps and Evan L. Porteus. Temporal resolution of uncertainty and dynamic choice theory. *Econometrica*, 46(1):185–200, 1978. doi:10.2307/1913656.

[Kri09] Vijay Krishna. *Auction Theory*. Academic Press, San Diego, 2nd edition, 2009.

[KL95] Chung-Ming Kuan and Tung Liu. Forecasting exchange rates using feedforward and recurrent neural networks. *Journal of Applied Econometrics*, 10(4):347–364, 1995.

[Kuh13] Moritz Kuhn. Recursive Equilibria In An Aiyagari-Style Economy With Permanent Income Shocks. *International Economic Review*, 54:807–835, 2013.

[KL51] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[KC78] Harold J. Kushner and Dean S. Clark. *Stochastic Approximation Methods for Constrained and Unconstrained Systems*. Springer-Verlag, New York, 1978.

[KY03] Harold J. Kushner and G. George Yin. *Stochastic Approximation and Recursive Algorithms and Applications*. Springer-Verlag, New York, 2nd edition, 2003.

[KBBP16] J. Nathan Kutz, Steven L. Brunton, Bingni W. Brunton, and Joshua L. Proctor. *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM, 2016.

[KP77] Finn E. Kydland and Edward C. Prescott. Rules rather than discretion: the inconsistency of optimal plans. *Journal of Political Economy*, 85(3):473–492, 1977.

[LR95] Peter Lancaster and Leiba Rodman. *Algebraic Riccati Equations*. Oxford University Press, Oxford, 1995.

[Lan75] Jan Lanke. On the choice of the unrelated question in simmons' version of randomized response. *Journal of the American Statistical Association*, 70(349):80–83, 1975.

[Lan76] Jan Lanke. On the degree of protection in randomized interviews. *International Statistical Review/Revue Internationale de Statistique*, pages 197–203, 1976.

[LL07] Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.

[LS96a] Ehud Lehrer and Rann Smorodinsky. Compatible measures and merging. *Mathematics of Operations Research*, 21(3):697–706, 1996. doi:10.1287/moor.21.3.697.

[LS96b] Ehud Lehrer and Rann Smorodinsky. Merging and learning. In Thomas S. Ferguson, Lloyd S. Shapley, and James B. MacQueen, editors, *Statistics, Probability and Game Theory: Papers in Honor of David Blackwell*, volume 30 of IMS Lecture Notes – Monograph Series, pages 147–168. Institute of Mathematical Statistics, Hayward, CA, 1996. doi:10.1214/lnms/1215453571.

[Leo83] Herman B. Leonard. Elicitation of honest preferences for the assignment of individuals to positions. *Journal of Political Economy*, 91(3):461–479, 1983.

[LL01] Martin Lettau and Sydney Ludvigson. Consumption, Aggregate Wealth, and Expected Stock Returns. *Journal of Finance*, 56(3):815–849, 06 2001.

[LL04] Martin Lettau and Sydney C. Ludvigson. Understanding Trend and Cycle in Asset Values: Reevaluating the Wealth Effect on Consumption. *American Economic Review*, 94(1):276–299, March 2004.

[LM80] David Levhari and Leonard J Mirman. The great fish war: an example using a dynamic cournot-nash solution. *The Bell Journal of Economics*, pages 322–334, 1980.

[LW76] Frederick W Leysieffer and Stanley L Warner. Respondent jeopardy and optimal designs in randomized response models. *Journal of the American Statistical Association*, 71(355):649–656, 1976.

[Lie12] Friedrich Liese. phi-divergences, sufficiency, Bayes sufficiency, and deficiency. *Kybernetika*, 48(4):690–713, 2012.

[LS77] Robert S. Liptser and Albert N. Shiryaev. *Statistics of Random Processes I: General Theory*. Springer-Verlag, New York, 1977.

[Lju77] Lennart Ljung. Analysis of recursive stochastic algorithms. *IEEE Transactions on Automatic Control*, 22(4):551–575, 1977. doi:10.1109/TAC.1977.1101561.

[LS18] L Ljungqvist and T J Sargent. *Recursive Macroeconomic Theory*. MIT Press, 4 edition, 2018.

[Lju93] Lars Ljungqvist. A unified approach to measures of privacy in randomized response models: a utilitarian perspective. *Journal of the American Statistical Association*, 88(421):97–103, 1993.

[LS04] Thomas A. Lubik and Frank Schorfheide. Testing for indeterminacy: an application to U.S. monetary policy. *American Economic Review*, 94(1):190–217, 2004.

[Luc72] Jr. Lucas, Robert E. Econometric testing of the natural rate hypothesis. In Otto Eckstein, editor, *The Econometrics of Price Determination*, pages 50–59. Board of Governors of the Federal Reserve System, Washington, D.C., 1972.

[Luc80] Jr. Lucas, Robert E. Two illustrations of the quantity theory of money. *American Economic Review*, 70(5):1005–1014, 1980.

[LP71a] Jr. Lucas, Robert E. and Edward C. Prescott. Investment under uncertainty. *Econometrica*, 39(5):659–681, 1971.

[LP74] Jr. Lucas, Robert E. and Edward C. Prescott. Equilibrium search and unemployment. *Journal of Economic Theory*, 7(2):188–209, 1974. doi:10.1016/0022-0531(74)90106-9.

[LS81] Robert E Lucas and Thomas J Sargent. *Rational expectations and econometric practice*. U of Minnesota Press, Minneapolis, Minnesota, 1981.

[Luc78] Robert E Lucas, Jr. Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445, 1978.

[Luc76] Robert E Jr Lucas. Econometric policy evaluation: a critique. In *Carnegie-Rochester conference series on public policy*, volume 1, 19–46. North-Holland, 1976.

[LP71b] Robert E. Lucas and Edward C. Prescott. Investment under uncertainty. *Econometrica*, 39(5):659–681, 1971. doi:10.2307/1909571.

[LJ67] Robert E Lucas Jr. Adjustment costs and the theory of supply. *Journal of political economy*, 75(4, Part 1):321–334, 1967.

[MST20] Qingyin Ma, John Stachurski, and Alexis Akira Toda. The income fluctuation problem and the evolution of wealth. *Journal of Economic Theory*, 187:105003, 2020.

[MaC82] Thomas E. MaCurdy. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics*, 18(1):83–114, 1982.

[MN16] Ulrike Malmendier and Stefan Nagel. Learning from inflation experiences. *Quarterly Journal of Economics*, 131(1):53–87, 2016.

[Man62] Edwin Mansfield. Entry, Gibrat's law, innovation, and the growth of firms. *American Economic Review*, 52(5):1023–1051, 1962.

[MM92] Albert Marcet and David A. Marshall. Convergence of approximate model solutions to rational expectations equilibria using the method of parameterized expectations. In *Universitat Pompeu Fabra Economics Working Paper*, number 13. 1992.

[MS89a] Albert Marcet and Thomas J Sargent. Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information. *Journal of Political Economy*, 97(6):1306–1322, 1989.

[MS89b] Albert Marcet and Thomas J. Sargent. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48(2):337–368, 1989. doi:10.1016/0022-0531(89)90032-X.

[MS89c] Albert Marcet and Thomas J. Sargent. Least squares learning and the dynamics of hyperinflation. In William A. Barnett, John Geweke, and Karl Shell, editors, *Economic Complexity: Chaos, Sunspots, Bubbles, and Nonlinearity*, pages 119–137. Cambridge University Press, Cambridge, 1989.

[MMS90] Ramon Marimon, Ellen McGrattan, and Thomas J. Sargent. Money as a medium of exchange in an economy with artificially intelligent agents. *Journal of Economic Dynamics and Control*, 14(2):329–373, 1990. doi:10.1016/0165-1889(90)90025-C.

[MS93] Ramon Marimon and Shyam Sunder. Indeterminacy of equilibria in a hyperinflationary world: experimental evidence. *Econometrica*, 61(5):1073–1107, 1993. doi:10.2307/2951494.

[MCWG95] A Mas-Colell, M D Whinston, and J R Green. *Microeconomic Theory*. Volume 1. Oxford University Press, 1995.

[McC70] J J McCall. Economics of Information and Job Search. *The Quarterly Journal of Economics*, 84(1):113–126, 1970.

[MN99] Bennett T. McCallum and Edward Nelson. Performance of operational policy rules in an estimated semiclassical structural model. In John B. Taylor, editor, *Monetary Policy Rules*, pages 15–45. University of Chicago Press, Chicago, 1999.

[MPQ00] Margaret M. McConnell and Gabriel Perez-Quiros. Output fluctuations in the United States: what has changed since the early 1980's? *American Economic Review*, 90(5):1464–1476, 2000.

[McF89] Daniel McFadden. A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica*, 57(5):995–1026, 1989. doi:10.2307/1913621.

[MT19] Michael McLeay and Silvana Tenreyro. Optimal inflation and the identification of the Phillips curve. In Martin Eichenbaum, Erik Hurst, and Jonathan A. Parker, editors, *NBER Macroeconomics Annual 2018, Volume 33*, pages 199–255. University of Chicago Press, Chicago, 2019.

[MP04] Costas Meghir and Luigi Pistaferri. Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32, 2004.

[MP85] Rajnish Mehra and Edward C Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985.

[MT98] Enrique G Mendoza and Linda L Tesar. The international ramifications of tax reforms: supply-side economics in a global economy. *American Economic Review*, pages 226–245, 1998.

[Mid11] Virgiliu Midrigan. Menu costs, multiproduct firms, and aggregate fluctuations. *Econometrica*, 79(4):1139–1180, 2011.

[Mil04] Paul Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, Cambridge, 2004.

[Mil77] Edward M Miller. Risk, uncertainty, and divergence of opinion. *The Journal of finance*, 32(4):1151–1168, 1977.

[MS99] Ronald I. Miller and Chris William Sanchirico. The role of absolute continuity in "Merging of Opinions" and "Rational Learning". *Games and Economic Behavior*, 29(1–2):170–190, 1999. doi:10.1006/game.1999.0752.

[MP69] Marvin Minsky and Seymour Papert. *Perceptrons: An Introduction to Computational Geometry*. MIT Press, Cambridge, MA, 1969.

[MB54] F. Modigliani and R. Brumberg. Utility analysis and the consumption function: An interpretation of cross-section data. In K.K Kurihara, editor, *Post-Keynesian Economics*. 1954.

[Mor63] Tetsuzo Morimoto. Markov Processes and the H-Theorem. *Journal of the Physical Society of Japan*, 18(3):328–331, 1963. doi:10.1143/JPSJ.18.328.

[Mor96] Stephen Morris. Speculative investor behavior and learning. *The Quarterly Journal of Economics*, 111(4):1111–1133, 1996.

[MAM10] Iain Murray, Ryan P. Adams, and David J. C. MacKay. Elliptical slice sampling. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of Proceedings of Machine Learning Research, 541–548. 2010.

[Mut60] John F Muth. Optimal properties of exponentially weighted forecasts. *Journal of the american statistical association*, 55(290):299–306, 1960.

[Mut61] John F. Muth. Rational expectations and the theory of price movements. *Econometrica*, 29(3):315–335, 1961.

[Mye81] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[Nea99] Derek Neal. The Complexity of Job Mobility among Young Men. *Journal of Labor Economics*, 17(2):237–261, 1999.

[Nea11] Radford M. Neal. MCMC using Hamiltonian dynamics. In Steve Brooks, Andrew Gelman, Galin L. Jones, and Xiao-Li Meng, editors, *Handbook of Markov Chain Monte Carlo*, pages 113–162. Chapman and Hall/CRC, 2011.

[Neg60] Takashi Negishi. Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica*, 12(2–3):92–97, 1960.

[NP82] Charles R. Nelson and Charles I. Plosser. Trends and random walks in macroeconomic time series: some evidence and implications. *Journal of Monetary Economics*, 10(2):139–162, 1982.

[NP33] J. Neyman and E. S. Pearson. On the problem of the most efficient tests of statistical hypotheses. *Philosophical Transactions of the Royal Society of London*, 231(694–706):289–337, 1933.

[Nov72] Alexander A. Novikov. On an identity for stochastic integrals. *Theory of Probability and Its Applications*, 17(4):717–720, 1972.

[Nya94] Yaw Nyarko. Bayesian learning leads to correlated equilibria in normal form games. *Economic Theory*, 4(6):821–841, 1994. doi:10.1007/BF01213814.

[Nyb89] Jukka Nyblom. Testing for the constancy of parameters over time. *Journal of the American Statistical Association*, 84(405):223–230, 1989.

[OW69] Guy H. Orcutt and Herbert S. Winokur. First order autoregression: inference, estimation, and prediction. *Econometrica*, 37(1):1–14, 1969.

[Orp01] Athanasios Orphanides. Monetary policy rules based on real-time data. *American Economic Review*, 91(4):964–985, 2001.

[PP89] Ariel Pakes and David Pollard. Simulation and the asymptotics of optimization estimators. *Econometrica*, 57(5):1027–1057, 1989. doi:10.2307/1913622.

[Par99] Jonathan A Parker. The Reaction of Household Consumption to Predictable Changes in Social Security Taxes. *American Economic Review*, 89(4):959–973, 1999.

[Per96] Lawrence Perko. *Differential Equations and Dynamical Systems*. Springer-Verlag, New York, 2nd edition, 1996.

[PPJ19] Du Phan, Neeraj Pradhan, and Martin Jankowiak. Composable effects for flexible and accelerated probabilistic programming in NumPyro. *arXiv preprint arXiv:1912.11554*, 2019.

[Phe67] Edmund S. Phelps. Phillips curves, expectations of inflation and optimal unemployment over time. *Economica*, 34(135):254–281, 1967.

[PT77] Edmund S. Phelps and John B. Taylor. Stabilizing powers of monetary policy under rational expectations. *Journal of Political Economy*, 85(1):163–190, 1977.

[PSS15] Monika Piazzesi, Juliana Salomao, and Martin Schneider. Trend and Cycle in Bond Premia. Working Paper, Stanford University, 2015.

[PR07] Frederic Pivetta and Ricardo Reis. The persistence of inflation in the United States. *Journal of Economic Dynamics and Control*, 31(4):1326–1358, 2007.

[PANS14] Luciano Pomatto, Nabil I. Al-Najjar, and Alvaro Sandroni. Merging and testing opinions. *The Annals of Statistics*, 42(3):1003–1028, 2014. doi:10.1214/14-AOS1212.

[PM80] Edward C Prescott and Rajnish Mehra. Recursive competitive equilibrium: the case of homogeneous households. *Econometrica*, 48(6):1365–1379, 1980.

[PV80] Edward C. Prescott and Michael Visscher. Organization capital. *Journal of Political Economy*, 88(3):446–461, 1980.

[Pri05] Giorgio E. Primiceri. Time varying structural vector autoregressions and monetary policy. *Review of Economic Studies*, 72(3):821–852, 2005.

[Pri06] Giorgio E. Primiceri. Why inflation rose and fell: policymakers' beliefs and U.S. postwar stabilization policy. *Quarterly Journal of Economics*, 121(3):867–901, 2006.

[PalS13] Jenő Pál and John Stachurski. Fitted value function iteration with probability one contractions. *Journal of Economic Dynamics and Control*, 37(1):251–264, 2013.

[Rab02] Guillaume Rabault. When do borrowing constraints bind? Some new results on the income fluctuation problem. *Journal of Economic Dynamics and Control*, 26(2):217–245, 2002.

[Rad72] Roy Radner. Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica*, 40(2):289–303, 1972.

[Rad79] Roy Radner. Rational expectations equilibrium: generic existence and the information revealed by prices. *Econometrica*, 47(3):655–678, 1979. doi:10.2307/1910413.

[RS84] Roy Radner and Joseph E. Stiglitz. A nonconcavity in the value of information. In Marcel Boyer and Richard E. Kihlstrom, editors, *Bayesian Models in Economic Theory*, pages 33–52. North-Holland, Amsterdam, 1984.

[Ref96] Kevin L Reffett. Production-based asset pricing in monetary economies with transactions costs. *Economica*, pages 427–443, 1996.

[Rei09] Michael Reiter. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665, 2009.

[RZ83] John Riley and Richard Zeckhauser. Optimal selling strategies: when to haggle, when to hold firm. *Quarterly Journal of Economics*, 98(2):267–289, 1983.

[RS81] John G. Riley and William F. Samuelson. Optimal auctions. *American Economic Review*, 71(3):381–392, 1981.

[RM51] Herbert Robbins and Sutton Monro. A stochastic approximation method. *Annals of Mathematical Statistics*, 22(3):400–407, 1951.

[Rog85] Kenneth Rogoff. The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100(4):1169–1189, 1985.

[Ros15] Stephen A. Ross. The recovery theorem. *Journal of Finance*, 70(2):615–648, 2015. doi:10.1111/jofi.12092.

[Rot82] Julio J. Rotemberg. Sticky prices in the united states. *Journal of Political Economy*, 90(6):1187–1211, 1982.

[Rya12] Stephen P Ryan. The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):1019–1061, 2012.

[Red97] Knut Røed. Hysteresis in unemployment. *Journal of Economic Surveys*, 11(4):389–418, 1997.

[Sam39] Paul A. Samuelson. Interactions between the multiplier analysis and the principle of acceleration. *The Review of Economics and Statistics*, 21(2):75–78, 1939.

[Sam58] Paul A. Samuelson. An exact consumption-loan model of interest with or without the social contrivance of money. *Journal of Political Economy*, 66(6):467–482, 1958. doi:10.1086/258100.

[SS60] Paul A. Samuelson and Robert M. Solow. Analytical aspects of anti-inflation policy. *American Economic Review*, 50(2):177–194, 1960.

[San98] Alvaro Sandroni. Necessary and sufficient conditions for convergence to Nash equilibrium: the almost absolute continuity hypothesis. *Games and Economic Behavior*, 22(1):121–147, 1998. doi:10.1006/game.1997.0572.

[San00] Alvaro Sandroni. Do markets favor agents able to make accurate predictions? *Econometrica*, 68(6):1303–1341, 2000.

[Sar77] Thomas J Sargent. The Demand for Money During Hyperinflations under Rational Expectations: I. *International Economic Review*, 18(1):59–82, February 1977.

[Sar81] Thomas J Sargent. Interpreting economic time series. *Journal of political Economy*, 89(2):213–248, 1981.

[Sar87] Thomas J Sargent. *Macroeconomic Theory*. Academic Press, New York, 2nd edition, 1987.

[Sar89] Thomas J Sargent. Two models of measurements and the investment accelerator. *Journal of Political Economy*, 97(2):251–287, 1989.

[Sar71] Thomas J. Sargent. A note on the `accelerationist' controversy. *Journal of Money, Credit and Banking*, 3(3):721–725, 1971.

[Sar80] Thomas J. Sargent. "Tobin's q" and the rate of investment in general equilibrium. *Carnegie-Rochester Conference Series on Public Policy*, 12:107–154, 1980.

[Sar93] Thomas J. Sargent. *Bounded Rationality in Macroeconomics*. The Arne Ryde Memorial Lectures. Oxford University Press, Oxford, 1993.

[Sar99] Thomas J. Sargent. *The Conquest of American Inflation*. Princeton University Press, Princeton, New Jersey, 1999.

[Sar08] Thomas J. Sargent. Evolution and intelligent design. *American Economic Review*, 98(1):5–37, 2008.

[SSY26a] Thomas J. Sargent, Yatheesan J. Selvakumar, and Ziyue Yang. Aggregate shocks and cross-section dynamics: quantifying redistribution and insurance in US household data. Technical Report, New York University, 2026.

[SSY26b] Thomas J. Sargent, Yatheesan J. Selvakumar, and Ziyue Yang. Dynamic mode decompositions and vector autoregressions. *International Economic Review*, 2026.

[SS26] Thomas J. Sargent and Zejin Shi. Social Security depletion: worker responses, bailout financing, and policy timing. Technical Report, University of Pennsylvania and Wuhan University, 2026.

[SS24] Thomas J. Sargent and John Stachurski. *Economic Networks: Theory and Computation*. Structural Analysis in the Social Sciences. Cambridge University Press, 2024.

[SS25] Thomas J. Sargent and John Stachurski. *Dynamic Programming: Finite States*. Cambridge University Press, 2025.

[SS11] Thomas J. Sargent and Paolo Surico. Two illustrations of the quantity theory of money: breakdowns and revivals. *American Economic Review*, 101(1):109–128, 2011. doi:10.1257/aer.101.1.109.

[SV95] Thomas J. Sargent and François R. Velde. Macroeconomic features of the French Revolution. *Journal of Political Economy*, 103(3):474–518, 1995.

[SW05] Thomas J. Sargent and Noah Williams. Impacts of priors on convergence and escapes from Nash inflation. *Review of Economic Dynamics*, 8(2):360–391, 2005.

[SW25a] Thomas J. Sargent and Noah Williams. Conquest lost and regained: American inflation in the 2020s. Extended abstract, 2025.

[SW25b] Thomas J. Sargent and Noah Williams. Rationalizing Fed interest rate decisions in the 2020s. *Journal of Political Economy: Macroeconomics*, 2025. Forthcoming.

[SWZ06] Thomas J. Sargent, Noah Williams, and Tao Zha. Shocks and government beliefs: the rise and fall of American inflation. *American Economic Review*, 96(4):1193–1224, 2006.

[SE77] Jack Schechtman and Vera L S Escudero. Some results on an income fluctuation problem. *Journal of Economic Theory*, 16(2):151–166, 1977.

[Sch14] Jose A. Scheinkman. *Speculation, Trading, and Bubbles*. Columbia University Press, New York, 2014.

[Sch10] Peter J Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of fluid mechanics*, 656:5–28, 2010.

[SS82] Avner Shaked and John Sutton. Relaxing price competition through product differentiation. *Review of Economic Studies*, 49(1):3–13, 1982.

[Sha48] Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.

[SS71] Lloyd S. Shapley and Martin Shubik. The assignment game I: the core. *International Journal of Game Theory*, 1(1):111–130, 1971.

[SZT17] Ravid Shwartz-Ziv and Naftali Tishby. Opening the Black Box of Deep Neural Networks via Information. *arXiv preprint arXiv:1703.00810*, 2017.

[Sim56] Herbert A Simon. Dynamic programming under uncertainty with a quadratic criterion function. *Econometrica, Journal of the Econometric Society*, pages 74–81, 1956.

[Sim80] Christopher A. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.

[Sim82] Christopher A. Sims. Policy analysis with econometric models. *Brookings Papers on Economic Activity*, 1982(1):107–164, 1982.

[Sim88] Christopher A. Sims. Projecting policy effects with statistical models. *Revista de Análisis Económico*, 3(1):3–20, 1988.

[Sim99] Christopher A. Sims. Drifts and breaks in monetary policy. Unpublished manuscript, Department of Economics, Princeton University, 1999.

[Sim01] Christopher A. Sims. Comment on Sargent and Cogley's `Evolving Post World War II US Inflation Dynamics'. *NBER Macroeconomics Annual*, 16:373–379, 2001.

[Sim02] Christopher A. Sims. Solving linear rational expectations models. *Computational Economics*, 20(1–2):1–20, 2002.

[Sim07] Christopher A. Sims. Evidence for passive monetary policy pre-1979. Comments at the Federal Reserve Bank of Cleveland Conference, May 2007.

[SZ06] Christopher A. Sims and Tao Zha. Were there regime switches in U.S. monetary policy? *American Economic Review*, 96(1):54–81, 2006.

[Slu37] Eugen Slutzky. The summation of random causes as the source of cyclic processes. *Econometrica*, 5(2):105–146, 1937. doi:10.2307/1907241.

[Smi93] Anthony A. Smith. Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics*, 8(S1):S63–S84, 1993. doi:10.1002/jae.3950080506.

[Sol68] Robert M. Solow. Recent controversy on the theory of inflation: an eclectic view. In Stephen W. Rousseaus, editor, *Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control*. New York University, New York, 1968.

[Sta08] John Stachurski. Continuous state dynamic programming via nonexpansive approximation. *Computational Economics*, 31(2):141–160, 2008.

[ST19] John Stachurski and Alexis Akira Toda. An impossibility theorem for wealth in heterogeneous-agent models with limited heterogeneity. *Journal of Economic Theory*, 182:1–24, 2019.

[Sti58] George J. Stigler. The economies of scale. *Journal of Law and Economics*, 1:54–71, 1958.

[Sto01] James H. Stock. Discussion of Cogley and Sargent `Evolving Post World War II US Inflation Dynamics'. *NBER Macroeconomics Annual*, 16:379–387, 2001.

[SW98] James H. Stock and Mark W. Watson. Median unbiased estimation of coefficient variance in a time-varying parameter model. *Journal of the American Statistical Association*, 93(441):349–358, 1998.

[SW02] James H. Stock and Mark W. Watson. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, 97(460):1167–1179, 2002.

[SW07] James H. Stock and Mark W. Watson. Why has U.S. inflation become harder to forecast? *Journal of Money, Credit and Banking*, 39(S1):3–33, 2007.

[SLP89] N L Stokey, R E Lucas, and E C Prescott. *Recursive Methods in Economic Dynamics*. Harvard University Press, 1989.

[Sto89] Nancy L Stokey. Reputation and time consistency. *The American Economic Review*, pages 134–139, 1989.

[Sto91] Nancy L. Stokey. Credible public policy. *Journal of Economic Dynamics and Control*, 15(4):627–656, October 1991.

[STY04] Kjetil Storesletten, Christopher I Telmer, and Amir Yaron. Consumption and risk sharing over the life cycle. *Journal of Monetary Economics*, 51(3):609–633, 2004.

[Str65] Volker Strassen. The existence of probability measures with given marginals. *Annals of Mathematical Statistics*, 36(2):423–439, 1965.

[Sun96] R K Sundaram. *A First Course in Optimization Theory*. Cambridge University Press, 1996.

[SB18] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

[Tau86] George Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986.

[Tay97] John B. Taylor. Comment on `America's Peacetime Inflation: The 1970s'. In Christina D. Romer and David H. Romer, editors, *Reducing Inflation: Motivation and Strategy*, pages 276–280. University of Chicago Press, 1997.

[The57] Henri Theil. A note on certainty equivalence in dynamic planning. *Econometrica: Journal of the Econometric Society*, pages 346–349, 1957.

[TPB99] Naftali Tishby, Fernando C. Pereira, and William Bialek. The Information Bottleneck Method. In *Proceedings of the 37th Annual Allerton Conference on Communication, Control, and Computing*, 368–377. 1999.

[Tob68] James Tobin. Discussion. In Stephen W. Rousseaus, editor, *Proceedings of a Symposium on Inflation: Its Causes, Consequences, and Control*. New York University, New York, 1968.

[Toi11] Juuso Toikka. Ironing without control. *Journal of Economic Theory*, 146(6):2510–2526, 2011.

[Tow83] Robert M. Townsend. Forecasting the forecasts of others. *Journal of Political Economy*, 91(4):546–588, 1983. doi:10.1086/261166.

[Tre16] Daniel Treisman. Russia's billionaires. *The American Economic Review*, 106(5):236–241, 2016.

[TRL+14] J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. On dynamic mode decomposition: theory and applications. *Journal of Computational Dynamics*, 1(2):391–421, 2014.

[VL11] Ngo Van Long. Dynamic games in the economics of natural resources: a survey. *Dynamic Games and Applications*, 1(1):115–148, 2011.

[Vic61] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[vN28] John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1):295–320, 1928.

[vN37] John von Neumann. Uber ein okonomsiches gleichungssystem und eine verallgemeinering des browerschen fixpunktsatzes. In *Erge. Math. Kolloq.*, volume 8, 73–83. 1937.

[Wal47] Abraham Wald. *Sequential Analysis*. John Wiley and Sons, New York, 1947.

[WW48] Abraham Wald and Jacob Wolfowitz. Optimum character of the sequential probability ratio test. *Annals of Mathematical Statistics*, 19(3):326–339, 1948.

[Wal80] W Allen Wallis. The statistical research group, 1942–1945. *Journal of the American Statistical Association*, 75(370):320–330, 1980.

[War65] Stanley L Warner. Randomized response: a survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.

[Wec79] William E Wecker. Predicting the turning points of a time series. *Journal of business*, pages 35–50, 1979.

[Wei89] Philippe Weil. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24(3):401–421, 1989. doi:10.1016/0304-3932(89)90028-7.

[Whi82] Halbert White. Maximum likelihood estimation of misspecified models. *Econometrica*, 50(1):1–25, 1982. doi:10.2307/1912526.

[Whi83] Charles Whiteman. *Linear Rational Expectations Models: A User's Guide*. University of Minnesota Press, Minneapolis, Minnesota, 1983.

[Whi84] Charles H. Whiteman. Lucas on the quantity theory: hypothesis testing without theory. *American Economic Review*, 74(4):742–749, 1984.

[Whi53] Peter Whittle. The analysis of multiple stationary time series. *Journal of the Royal Statistical Society, Series B*, 15(1):125–139, 1953.

[Whi81] Peter Whittle. Risk-sensitive linear/quadratic/gaussian control. *Advances in Applied Probability*, 13(4):764–777, 1981.

[Whi90] Peter Whittle. *Risk-Sensitive Optimal Control*. Wiley, New York, 1990.

[Wil01] Noah Williams. *Escape Dynamics in Learning Models*. PhD thesis, University of Chicago, 2001.

[Wil19] Noah Williams. Escape dynamics in learning models. *Review of Economic Studies*, 86(2):882–912, 2019.

[Wil75] Oliver E. Williamson. *Markets and Hierarchies: Analysis and Antitrust Implications*. Free Press, New York, 1975.

[Wol23] Christian K. Wolf. The missing intercept: a demand equivalence approach. *American Economic Review*, 113(8):2232–2269, 2023. doi:10.1257/aer.20211751.

[Woo90] Michael Woodford. Learning to believe in sunspots. *Econometrica*, 58(2):277–307, 1990.

[Woo15] Jeffrey M Wooldridge. *Introductory econometrics: A modern approach*. Nelson Education, 2015.

[YS05] G Alastair Young and Richard L Smith. *Essentials of statistical inference*. Cambridge University Press, 2005.

[vCencov81] Nikolai N. Čencov. *Statistical Decision Rules and Optimal Inference*. Volume 53 of Translations of Mathematical Monographs. American Mathematical Society, Providence, RI, 1981.

[BoardoTrusteesFOAaSIaFDITFunds26] Board of Trustees, Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. The 2026 annual report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds. Technical Report, U.S. Government Publishing Office, 2026.

[DeGroot62] Morris H. DeGroot. Uncertainty, information, and sequential experiments. *Annals of Mathematical Statistics*, 33(2):404–419, 1962. doi:10.1214/aoms/1177704567.

PROOF INDEX

absolute_continuity
absolute_continuity (*merging_of_opinions*), 819

ak-aiyagari-algorithm
ak-aiyagari-algorithm (*ak_aiyagari*), 2817

algorithm-0
algorithm-0 (*two_computation*), 2850

as:util-rand-discomfort
as:util-rand-discomfort (*util_rand_resp*), 341

as:util-rand-truth
as:util-rand-truth (*util_rand_resp*), 342

ass-ms-a1
ass-ms-a1 (*ls_learning*), 3178

ass-ms-a2
ass-ms-a2 (*ls_learning*), 3178

ass-ms-a3
ass-ms-a3 (*ls_learning*), 3179

ass-ms-a4
ass-ms-a4 (*ls_learning*), 3179

ass-ms-a5
ass-ms-a5 (*ls_learning*), 3179

ass-ms-a6
ass-ms-a6 (*ls_learning*), 3179

ass-ms-a7
ass-ms-a7 (*ls_learning*), 3179

assm-cond-ind
assm-cond-ind (*unemployment_shocks*), 3433

assum-bk-borel
assum-bk-borel (*rational_learning_re*), 3165

assum-bk-continuity
assum-bk-continuity (*rational_learning_re*), 3166

assum-bk-identification
assum-bk-identification (*rational_learning_re*), 3166

assumption-0
assumption-0 (*sir_model*), 10

assumption-mr-condition-4
assumption-mr-condition-4 (*misspecified_recovery*), 3325

assumption-mr-ergodicity
assumption-mr-ergodicity (*misspecified_recovery*), 3325

assumption-mr-markov-shocks
assumption-mr-markov-shocks (*misspecified_recovery*), 3320

assumption1
assumption1 (*von_neumann_model*), 991

assumption2
assumption2 (*von_neumann_model*), 991

blackwell_dubins
blackwell_dubins (*merging_of_opinions*), 822

bm_policies
bm_policies (*optimal_growth_uncertainty*), 2668

bm_theorem
bm_theorem (*optimal_growth_uncertainty*), 2669

boundary_drift_gap
boundary_drift_gap (*survival_recursive_preferences*), 865

bs_theorem
bs_theorem (*optimal_growth_uncertainty*), 2671

com_eq_tax
com_eq_tax (*cass_fiscal*), 2507

cor-ms-exogenous
cor-ms-exogenous (*ls_learning*), 3182

cor-ms-projection-dichotomy
cor-ms-projection-dichotomy (*ls_learning*), 3182

cor-risk-neutral-recovery
cor-risk-neutral-recovery (*ross_recovery*), **??**

cp_admissible
cp_admissible (*phillips_credible_policies*), 2167

cp_prop_best
cp_prop_best (*phillips_credible_policies*), 2171

cp_prop_worst
cp_prop_worst (*phillips_credible_policies*), 2170

cp_selfenf
cp_selfenf (*phillips_credible_policies*), 2176

cp_selfgen
cp_selfgen (*phillips_credible_policies*), 2168

cp_spe
cp_spe (*phillips_credible_policies*), 2167

cp_strategy
cp_strategy (*phillips_credible_policies*), 2166

def-circulant-matrix
def-circulant-matrix (*eig_circulant*), 56

def-degroot-uncertainty
def-degroot-uncertainty (*blackwell_kihlstrom*), 686

def-economic-criterion
def-economic-criterion (*blackwell_kihlstrom*), 686

def-markov-perfect-equilibrium
def-markov-perfect-equilibrium (*markov_perf*), 2705

def-sufficiency
def-sufficiency (*blackwell_kihlstrom*), 684

def-transition-independence
def-transition-independence (*ross_recovery*), **??**

definition-0
definition-0 (*rational_expectations*), **??**

definition-2
definition-2 (*von_neumann_model*), 991

definition-3
definition-3 (*von_neumann_model*), 994

definition-4
definition-4 (*von_neumann_model*), 994

definition-6
definition-6 (*von_neumann_model*), 995

definition-8
definition-8 (*von_neumann_model*), 1001

dichotomy_at_infinity
dichotomy_at_infinity (*merging_of_opinions*), 840

equilibrium_asset_price
equilibrium_asset_price (*morris_learn*), 3143

hellinger_affinity
hellinger_affinity (*merging_of_opinions*), 836

ime_def_sufficiency
ime_def_sufficiency (*information_market_equilibrium*), 718

ime_lemma_posterior_sufficiency
ime_lemma_posterior_sufficiency (*information_market_equilibrium*), 718

ime_lemma_same_price_same_allocation
ime_lemma_same_price_same_allocation (*information_market_equilibrium*), 720

ime_lemma_unique_posterior
ime_lemma_unique_posterior (*information_market_equilibrium*), 721

ime_theorem_bayesian_convergence

ime_theorem_bayesian_convergence (*information_market_equilibrium*), 728

ime_theorem_invertibility_conditions

ime_theorem_invertibility_conditions (*information_market_equilibrium*), 721

ime_theorem_price_revelation

ime_theorem_price_revelation (*information_market_equilibrium*), 719

kakutani_dichotomy

kakutani_dichotomy (*merging_of_opinions*), 836

lc_prop1

lc_prop1 (*phillips_lost_conquest*), 2446

lc_prop2

lc_prop2 (*phillips_lost_conquest*), 2446

lp_equilibrium_def

lp_equilibrium_def (*lucas_prescott_investment*), 2618

lp_theorem1

lp_theorem1 (*lucas_prescott_investment*), 2619

lp_theorem2

lp_theorem2 (*lucas_prescott_investment*), 2624

lp_theorem3

lp_theorem3 (*lucas_prescott_investment*), 2627

lp_theorem4

lp_theorem4 (*lucas_prescott_investment*), 2627

lps_prop_v

lps_prop_v (*lucas_prescott_search*), 2640

lrr-def-additive-functional

lrr-def-additive-functional (*long_run_risk_operator*), 3265

lrr-def-cash-flow-valuation-semigroup

lrr-def-cash-flow-valuation-semigroup (*long_run_risk_operator*), 3266

lrr-def-exponential-affine-eigenfunction

lrr-def-exponential-affine-eigenfunction (*long_run_risk_operator*), 3283

lrr-def-extended-generator

lrr-def-extended-generator (*long_run_risk_operator*), 3270

lrr-def-functional

lrr-def-functional (*long_run_risk_operator*), 3263

lrr-def-generator-eigenfunction

lrr-def-generator-eigenfunction (*long_run_risk_operator*), 3272

lrr-def-harris-recurrence

lrr-def-harris-recurrence (*long_run_risk_operator*), 3274

lrr-def-intensity-matrix

lrr-def-intensity-matrix (*long_run_risk_operator*), 3276

lrr-def-irreducible-skeleton

lrr-def-irreducible-skeleton (*long_run_risk_operator*), 3274

lrr-def-local-brownian-risk-price

lrr-def-local-brownian-risk-price (*long_run_risk_operator*), 3292

lrr-def-long-run-risk-price

lrr-def-long-run-risk-price (*long_run_risk_operator*), 3292

lrr-def-martingale-component

lrr-def-martingale-component (*long_run_risk_operator*), 3272

lrr-def-multiplicative-functional

lrr-def-multiplicative-functional (*long_run_risk_operator*), 3264

lrr-def-multiplicative-semigroup

lrr-def-multiplicative-semigroup (*long_run_risk_operator*), 3266

lrr-def-one-parameter-semigroup

lrr-def-one-parameter-semigroup (*long_run_risk_operator*), 3265

lrr-def-stationary-distribution

`lrr-def-stationary-distribution` (*long_run_risk_operator*), 3273

lrr-def-stochastic-discount-factor

`lrr-def-stochastic-discount-factor` (*long_run_risk_operator*), 3266

lrr-def-stochastic-growth-functional

`lrr-def-stochastic-growth-functional` (*long_run_risk_operator*), 3266

lrr-def-stochastic-stability

`lrr-def-stochastic-stability` (*long_run_risk_operator*), 3274

lrr-def-valuation-functional

`lrr-def-valuation-functional` (*long_run_risk_operator*), 3266

martingale_convergence

`martingale_convergence` (*merging_of_opinions*), 820

md_prop_belief

`md_prop_belief` (*market_diffusion*), 915

md_prop_efficient

`md_prop_efficient` (*market_diffusion*), 918

md_prop_equilibrium

`md_prop_equilibrium` (*market_diffusion*), 919

md_prop_sshape

`md_prop_sshape` (*market_diffusion*), 924

mfg_prop_common_noise

`mfg_prop_common_noise` (*lq_mean_field_games*), 2963

mfg_prop_eigen_aI

`mfg_prop_eigen_aI` (*lq_mean_field_games*), 2968

mfg_prop_eigen_symmetric

`mfg_prop_eigen_symmetric` (*lq_mean_field_games*), 2968

mfg_prop_equivalence

`mfg_prop_equivalence` (*lq_mean_field_games*), 2948

mfg_prop_existence

`mfg_prop_existence` (*lq_mean_field_games*), 2949

mfg_prop_identification

`mfg_prop_identification` (*lq_mean_field_games*), 2963

mfg_prop_identification_more

`mfg_prop_identification_more` (*lq_mean_field_games*), 2966

mfg_prop_planner

`mfg_prop_planner` (*lq_mean_field_games*), 2973

mfg_prop_riccati

`mfg_prop_riccati` (*lq_mean_field_games*), 2948

mfg_prop_roots

`mfg_prop_roots` (*lq_mean_field_games*), 2947

mfg_prop_trace

`mfg_prop_trace` (*lq_mean_field_games*), 2966

mfgx_prop_cross

`mfgx_prop_cross` (*lq_mean_field_games_extensions*), 2994

mfgx_prop_forcing

`mfgx_prop_forcing` (*lq_mean_field_games_extensions*), 2998

mfgx_prop_moments

`mfgx_prop_moments` (*lq_mean_field_games_extensions*), 3001

most_optimistic_valuation

`most_optimistic_valuation` (*morris_learn*), 3143

normalized_price

`normalized_price` (*morris_learn*), 3143

one_sided_vs_mutual

`one_sided_vs_mutual` (*merging_of_opinions*), 822

pa_prop

`pa_prop` (*phillips_adaptive*), 2195

pi_prop_structure

`pi_prop_structure` (*pricing_information*), 891

pm_bmap
pm_bmap (*phillips_misspecified*), 2210

pm_equilibrium
pm_equilibrium (*phillips_misspecified*), 2211

prop-bk-event-convergence
prop-bk-event-convergence (*rational_learning_re*), 3165

prop-bk-measure-convergence
prop-bk-measure-convergence (*rational_learning_re*), 3166

prop-bk-sharpening
prop-bk-sharpening (*rational_learning_re*), 3167

prop-dmd-eigenvectors
prop-dmd-eigenvectors (*var_dmd*), 103

prop-dmd-left
prop-dmd-left (*var_dmd*), 104

prop-loo-identity
prop-loo-identity (*unemployment_shocks*), 3434

prop-mean-preservation
prop-mean-preservation (*blackwell_kihlstrom*), 688

prop-misspecified-recovery-martingale-component
prop-misspecified-recovery-martingale-component (*misspecified_recovery*), 3318

prop-mr-uniqueness
prop-mr-uniqueness (*misspecified_recovery*), 3325

prop-ms-convergence
prop-ms-convergence (*ls_learning*), 3180

prop-ms-jacobian-eigenvalues
prop-ms-jacobian-eigenvalues (*ls_learning*), 3180

prop-ms-necessity
prop-ms-necessity (*ls_learning*), 3181

prop-rbew-types
prop-rbew-types (*lq_robust_bewley*), 1850

prop:hst_oe
prop:hst_oe (*robust_permanent_income*), 1754

q_differentiability
q_differentiability (*optimal_growth_uncertainty*), 2672

rate_dominance_beta
rate_dominance_beta (*morris_learn*), 3144

remark-0
remark-0 (*qr_decomp*), 50

sc_def
sc_def (*phillips_self_confirming*), **??**

sharpness
sharpness (*merging_of_opinions*), 823

speculative_premium
speculative_premium (*morris_learn*), 3144

ss_prop
ss_prop (*sargent_surico*), 3449

survival_conditions
survival_conditions (*survival_recursive_preferences*), 859

theorem-0
theorem-0 (*os_stochastic*), 1543

theorem1
theorem1 (*von_neumann_model*), 994

theorem2
theorem2 (*von_neumann_model*), 1000

thm-blackwell
thm-blackwell (*blackwell_kihlstrom*), 687

thm-data-processing
thm-data-processing (*blackwell_kihlstrom*), 705

thm-kihlstrom
thm-kihlstrom (*blackwell_kihlstrom*), 689

thm-perron-frobenius
thm-perron-frobenius (*ross_recovery*), **??**

thm-rcs-oe1
thm-rcs-oe1 (*lq_robust_smoothing*), 1827

thm-rcs-oe2

thm-rcs-oe2 (*lq_robust_smoothing*), 1830

thm-ross-recovery

thm-ross-recovery (*ross_recovery*), **??**

total_variation_distance

total_variation_distance (*merging_of_opinions*), 819

two_comp_algo_buyout

two_comp_algo_buyout (*two_computation*), 2872

two_comp_algo_endo

two_comp_algo_endo (*two_computation*), 2879

two_comp_algo_govfund

two_comp_algo_govfund (*two_computation*), 2886

two_trader_optimist

two_trader_optimist (*morris_learn*), 3144

vs_prop_blockexog

vs_prop_blockexog (*var_subsets*), 1224

vs_prop_full

vs_prop_full (*var_subsets*), 1224

vs_prop_innov

vs_prop_innov (*var_subsets*), 1222

INDEX

A

A Bayesian Formulation of Friedman and Wald's Problem, 599
A Problem that Stumped Milton Friedman, 565
An Introduction to Job Search, 1319, 1335
Arrow Securities
    competitive equilibrium, 3085
Asset Pricing
    GMM Estimation, 3065
    MLE Estimation, 3035
Asset Pricing: Finite State Models, 3013

B

Bewley Model, 1807, 1811
Blackwell
    and screening, 885
Bounded Rationality, 1887
    Exchange Rate Indeterminacy, 1977
    Genetic Algorithms, 2005
    Money as a Medium of Exchange, 2037
    Overlapping Generations, 1913
    Parameterized Expectations, 1951
    Prospects, 2083

C

Central Limit Theorem, 195, 200
    Intuition, 200
    Multivariate Case, 205
Certainty Equivalence
    Robustness, 1769, 1777
CLT, 195
Complete Markets, 1807
    Arrow securities, 1813
Consumption Smoothing, 1793
Cross-Section Distributions
    consumption, 1809

D

Detection Error Probabilities, 1838
Distorted Expectations
    Stackelberg multiplier game, 1829
Dynamic Programming
    Computation, 1544
    Theory, 1543
    Unbounded Utility, 1544

E

Eigenvalues, 21, 36
Eigenvectors, 21, 36
Ergodicity, 1005, 1019

F

Finite Markov Asset Pricing
    Lucas Tree, 3021
Finite Markov Chains, 1005, 1006
    Stochastic Matrices, 1006
Frequency Domain
    permanent income model, 1836

G

Gram-Schmidt process, 1203

H

Hansen Sargent Tallarini
    model, 1824
hidden Markov model, 1203
History Dependence, 1797

I

Information
    pricing, 885
    strategic experimentation, 913
Irreducibility and Aperiodicity, 1005, 1012

J

Job Search VIII: On-the-Job Search, 1415
Job Search VIII: On-the-Job Search, 1415

K

Kalman Filter, 1163, 1199
    and vector autoregressions, 1205, 1219

Programming Implementation, 1172
Recursive Procedure, 1171
Kalman Filter 2, 1181
Kesten processes
heavy tails, 1134

L

Lake Model, 2565
Lake Model, Endogenous, 2585
Law of Large Numbers, 195, 196
Illustration, 198
Multivariate Case, 205
Proof, 197
Learning
two-sided, 913
Linear Algebra, 21
Differentiating Linear and Quadratic Forms, 40
Eigenvalues, 36
Eigenvectors, 36
Matrices, 29
Matrix Norms, 39
Neumann's Theorem, 40
Positive Definite Matrices, 40
SciPy, 36
Series Expansions, 39
Spectral Radius, 40
Vectors, 22
Linear Markov Perfect Equilibria, 2705
Linear State Space Models, 1041, 1133
Distributions, 1048
Ergodicity, 1055
Martingale Difference Shocks, 1043
Moments, 1048
Moving Average Representations, 1047
Prediction, 1061
Seasonals, 1046
Stationarity, 1055
Time Trends, 1046
Univariate Autoregressive Processes, 1044
Vector Autoregressions, 1046
LLN, 195
LQ Control, 1667
Infinite Horizon, 1679
Optimality (*Finite Horizon*), 1671
Permanent Income, 1769, 1777
LQ Permanent Income Model, 1793
cross section, 1807
standard, 1794

M

Marginal Distributions, 1005, 1010
Markov Asset Pricing
Overview, 3013
Markov Chains, 1006
Calculating Stationary Distributions, 1017
Convergence to Stationarity, 1018
Cross-Sectional Distributions, 1012
Ergodicity, 1019
Forecasting Future Values, 1020
Future Probabilities, 1012
Irreducibility, Aperiodicity, 1012
Marginal Distributions, 1010
Simulation, 1008
Stationary Distributions, 1015
Markov Perfect Equilibrium, 2703
Applications, 2708
Background, 2704
Overview, 2703
Markov process, inventory, 1029
Matrix
Determinants, 34
Inverse, 34
Maps, 31
NumPy, 30
Operations, 29
Solving Systems of Equations, 31
Modeling
Career Choice, 1401
Modeling COVID-19, 9
Modeling COVID-19, 9
Models
Harrison Kreps, 3127
Linear State Space, 1042
Markov Asset Pricing, 3013
McCall, 1296
Morris Learning, 3139
On-the-Job Search, 1415
Permanent Income, 1715, 1731
Pricing, 3014
Sequential analysis, 565, 599

N

Neumann's Theorem, 40

O

Observational Equivalence
heterogeneous agents, 1847
robustness and discounting, 1824
Theorem 1, 1827
Theorem 2, 1830
On-the-Job Search
Model, 1416
Model Features, 1416
Parameterization, 1417
Programming Implementation, 1418

Solving for Policies, 1420
Optimal Savings
Computation, 1612, 1652
Model, 1540
Policy Function, 1551
Policy Function Approach, 1541
Problem, 1610, 1628
Programming Implementation, 1616, 1633
Optimal Savings III: Stochastic Returns, 1539
Optimal Savings IV: Time Iteration, 1555
Optimal Savings V: The Endogenous Grid Method, 1567
Optimal Savings VI: EGM with JAX, 1575
Organization Capital, 1241
Bayesian Screening, 1251
Human Capital, 1255
Industry Equilibrium, 1249
Personnel Information, 1242
Team Information, 1253
Transfer Costs, 1256

P

Pandas for Panel Data, 3351
Permanent Income Hypothesis
Friedman, 1794
Permanent Income II: LQ Techniques, 1731
Permanent Income Model
examples, 1797
Hall's Representation, 1723
Savings Problem, 1716
Phillips Curve
Adaptive Expectations, 2191
Adaptive Learning and Escape Dynamics, 2247
Credibility Problem, 2141
Credible Policies, 2161
Drifts and Volatilities, 2365
Drifts and Volatilities after 2000, 2409
Econometric Policy Evaluation, 2341
Escaping Nash Inflation, 2279
Fed Policy in the 2020s, 2429
Optimal Misspecified Beliefs, 2207
Priors and Learning Cycles, 2309
Rise and Fall of U.S. Inflation, 2111
Self-Confirming Equilibria, 2223
Positive Definite Matrices, 40
Precautionary Savings
robustness, 1821
robustness vs convex marginal utility, 1828
Pricing Models, 3013, 3014
Risk Aversion, 3014
Risk-Neutral, 3014
Python
Pandas, 3351
python, 114, 984

Q

Quantity Theory of Money, 3439

R

Rational Expectations Equilibrium, 2599
Competitive Equilbrium (*w. Adjustment Costs*), 2602
Computation, 2605
Definition, 2602
Planning Problem Approach, 2606
Riccati equation
matrix difference, 1202
Robust Bewley Model, 1847
Robust Control
breakdown point, 1832
permanent income, 1821
Robustness
payoff evaluation, 1841

S

spectral factorization identity, 1205
Spectral Radius, 40
Stability in Linear Rational Expectations Models, 2683
Stationary Distributions, 1005, 1015
Stochastic Matrices, 1006
Survival
Recursive Preferences, 853

T

The Income Fluctuation Problem III: The Endogenous Grid Method, 1609
The Income Fluctuation Problem IV: Transient Income Shocks, 1627
The Income Fluctuation Problem V: Stochastic Returns on Assets, 1649
The Permanent Income Model, 1715

U

Unbounded Utility, 1544

V

Vector Autoregression
and Kalman filter, 1199
subsystems, 1219
vector autoregression, 1205
Vectors, 21, 22
Inner Product, 25

Linear Independence, 28
Norm, 25
Operations, 23
Span, 26